A Context-aware Approach to Wireless Transmission Adaptation

Veljko Pejovic
University of California, Santa Barbara
veljko@cs.ucsb.edu

Elizabeth M. Belding
University of California, Santa Barbara
ebelding@cs.ucsb.edu

***Abstract*—Recent advancements in wireless transmission have enabled networks with a high level of physical layer flexibility. Unfortunately, these new opportunities are not harnessed by modern wireless systems. Due to inefficient resource allocation, systems typically encounter problems such as spectrum scarcity, energy depletion or low quality of service.**

In this paper we consider the problem of physical layer parameter adaptation in a flexible wireless system. We observe that for many practical purposes the acceptable quality of communication depends on the interplay among the packet loss ratio, energy savings and spectrum utilization. We harness this fact and propose a context-aware physical layer parameter adaptation solution, WhiteRate. Our solution adjusts the modulation level, coding scheme and channel width to achieve the 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.

I. INTRODUCTION

The demand for wireless connectivity pushes service provisioning from standard home and office environments to situations that were not envisioned 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 expectations. However, higher proliferation of wireless communication is hampered by both typical and unconventional obstacles. For example, spectrum scarcity is becoming an increasingly important problem in highly populated areas where the sheer density of mobile wireless devices causes deterioration in quality of service. In another example, when brought to rural areas that are characterized by erratic power supply, traditional WiFi equipment fails to deliver reliable connectivity due to its low energy efficiency, which leaves battery resources depleted. While these problems are usually attributed to low infrastructure scalability, we find that inefficient wireless resource allocation significantly contributes to poor performance.

Recent advances in software defined radio (SDR) and the provision of newly unlicensed bands (white spaces) have facilitated research on flexible wireless systems [6]. As a result, contemporary devices have a multitude of parameters that can be adapted. However, current devices react to changes in the wireless environment without putting the reaction in the greater context of the situation. For example, a device might lower the communication rate in order to boost the packet delivery rate without considering the actual application needs for reliable delivery. Similarly, devices contend for the same amount of wireless spectrum irrespective of their offered load.

In this work we tackle the problem of context-aware wireless resource allocation. We leverage recent advances in physical layer flexibility realized through adjustable orthogonal frequency division multiplexing (OFDM). 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 [26] and 802.22 [12]: we change the width by OFDM subcarrier (de)activation and bind the available spectrum into a single OFDM channel, even if the available spectrum resides in non-contiguous bands.[1]

In spite of a substantial body of related academic work that analyzes each of the individual parameters separately [11], [15], [20], it is unclear how concurrent adjustments of all the PHY knobs should be approached so that the solution is identified with minimum communication and processing overhead. In our initial analysis we show that tuning these parameters changes energy efficiency, packet delivery and spectrum utilization of a system. We leverage the fact that applications have different requirements for each of the three aspects and develop WhiteRate, a solution that adjusts the PHY layer parameters according to these requirements. WhiteRate achieves the above in a practical manner, with little computational and communication overhead.

In summary, we make the following contributions:

- We theoretically and experimentally examine the impact of different PHY layer parameters on three important aspects of wireless communication systems: packet delivery, energy consumption and spectrum utilization.
- We justify context awareness in wireless resource allocation. We show that wireless systems can harness PHY layer flexibility to satisfy application requirements.
- We develop WhiteRate, a protocol that adapts wireless channel width and modulation and coding scheme (MCS) according to application requirements.

---

[1] In the paper we use "channel width" and "number of active subcarriers" interchangeably unless otherwise stated.

While WhiteRate targets white spaces due to the vast potential combinations of channel widths and data rates, its applicability is by no means limited to white space networks. We implement WhiteRate in GNUradio, a software-defined radio platform, and USRP2[2] hardware, and experimentally evaluate the solution 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 [4]. Therefore, in our lab testbed we investigate the communication quality of WhiteRate as the available spectrum and the number of active clients vary. On the other hand, we envision WhiteRate as a solution for 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 testbed in Pretoria, South Africa. We show that WhiteRate operates along the tradeoff line at which the energy savings are maximized for the given application loss tolerance. We compare WhiteRate with a standard rate adaptation solution and intuitive alternatives. Through the experiments we demonstrate that by using WhiteRate we can support more clients without sacrificing the application quality. We also show that WhiteRate saves up to three quarters of the transmission energy per bit as compared to context agnostic solutions.

This paper is structured as follows. In section II we investigate the impact of PHY parameters on the communication quality, energy consumption and spectrum utilization in the flexible OFDM system. In section III we discuss the opportunities for context-aware operation. We design the WhiteRate channel width and MCS adaptation protocol in section IV. In section V we describe our GNUradio-based experimental methodology and present the results of the WhiteRate performance analysis. We complete the paper with an overview of the related work (sec. VI) and our conclusions (sec. VII).

II. PHY PARAMETERS AND SYSTEM PERFORMANCE

A. Communication Performance

Packet delivery depends on core PHY parameters such as the 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:

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

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:

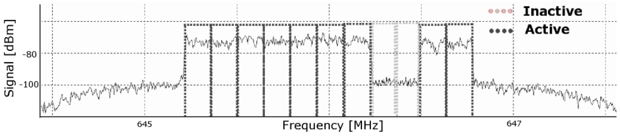

Fig. 1: OFDM subcarrier activation in groups.

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

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 coding scheme introduces redundancy that improves the bit error resiliency, thus allowing higher modulation levels, but at the same time reduces the fraction of "pure" information bits transferred by the factor r. The robustness of MCSs is a well studied topic [20], 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[3].

Wider channels always result in a higher capacity; however, the relationship between channel width and transmission errors is not as clear. In [11] Chandra et al. report lower packet loss over narrower channels and explain this phenomenon with higher power-per-Hz and better resilience to delay spread of narrow channels. The same reasoning does not hold in our case as we keep the same power per subcarrier irrespective of the number of active subcarriers and, unlike in [11], we do not change the subcarrier width. In our problem setting, the most probable cause of varying errors is frequency selective fading. In the case of frequency-selective fading the 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 of the bits are more likely to be corrupted than others.

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 frequent sweeping through the whole frequency range so that the well performing frequencies can be isolated. This increases the protocol delay

---

[2] http://gnuradio.org and http://www.ettus.com

[3] Viterbi decoder occupies a significant number of logical gates on a wireless NIC [13]. The operation is also the most processing intensive part of the demodulation in case of the SDR.

Total power $P_{total}(tx\_amplitude = 0)$	12.77 W
Total power when $idle$, $P_{total}(idle) = P_{base}$	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.	
Changing MCS, $tx\_amplitude = 0.5$	
$P_{Tx} = P_{total}(max\_width, \textbf{any\_mcs}) - P_{base} - P_{TC}$	0.45 W
Changing width, $tx\_amplitude = 0.5$	
$P_{Tx} = P_{total}(\textbf{all\_widths}, any\_mcs) - P_{base} - P_{TC}$	[0.22 - 0.45W]

TABLE I: Power consumption breakdown.

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. If the channel is flat fading, the groups are activated according to the *contiguous* scheme.

B. Energy Efficiency

The energy consumption of a network device, whether it is a smartphone, laptop or a self-powered wireless router in a rural area network, is highly influenced by its wireless network interfaces. Not only is a network interface card (NIC) one of the most power-demanding parts of the system, but its activity often prevents the rest of the system from switching to a low power mode [16]. A good understanding of the energy consumption of wireless NICs operating in a flexible communication environment helps us to identify the space for improvement of the energy efficiency of the whole system.

A bit of information over a single subcarrier i is transmitted with energy:

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

where $P_{Tx,i}$ represents the transmission power and P_{TC} the transceiver circuit power of a wireless NIC [15]. 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. With k active subcarriers the energy per bit becomes:

$$(4) \quad E_{Tx} = \frac{\sum_{i=1}^{k} P_{Tx,i} + P_{TC}}{\sum_{i=1}^{k} R_i}$$

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 [5], [23]. In our previous work [18] we performed a thorough analysis of the USRP2 SDR power consumption (table I) and confirmed that the choice of MCS does not impact the power consumption on this platform. It follows that higher modulations are more energy efficient as they reduce the transmission time without any power consumption penalty.

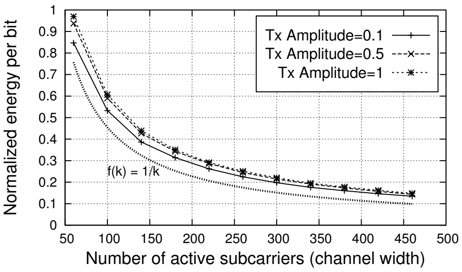

Fig. 2: Normalized energy per bit (E_{Tx}) against $tx\_amplitude$ and channel width (number of active subcarriers).

Previous investigation of the impact of channel width on the power consumption of a wireless device shows that narrow channels result in power efficient operation [11]. Compared to this study, our flexible OFDM system changes the channel width in a fundamentally different way: the number of FFT bins[4] that determines the total number of subcarriers remains constant with channel width, and so does the system clock; only the number of active subcarriers changes. In [18] we investigated the power consumption change as we modified the number of active subcarriers and observed that the transmission power increases approximately linearly with the number of active subcarriers.

We substitute the experimental results for P_{Tx} and P_{TC} into (4) and in figure 2 show the normalized (over the bitrate-per-subcarrier) transmission energy-per-bit under varying channel width for three different $tx\_amplitude$[5] values. At any amplitude E_{Tx} declines almost inversely with the rising number of active subcarriers. Intuitively, it is clear that, if the power consumption is dominated by a fixed factor ($P_{TC} = 1.01W$) while the varying factor ($P_{Tx} = [0.22 - 0.45W]$) remains linearly bounded, reducing the transmission time by extending the channel leads to the energy-optimal transmission.

C. Spectrum Utilization

In practical terms efficient spectrum utilization can be described as the wireless operation in which the medium capacity is used as close to its maximum as possible. Physical layer parameters directly determine the utilization: higher modulations and the appropriate coding push the transmission rate (eq. *(2)* in section II-A) closer to the Shannon capacity (eq. *(1)* in section II-A); aggressive subcarrier allocation that leaves few unutilized frequencies also improves spectrum utilization.

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 [22]. 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.

---

[4] Resolution of Fast Fourier Transform.

[5] A GNUradio exposed parameter that determines the signal amplitude.

III. 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.

A. 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 the communication success is not a binary value, but can be described as more or less satisfactory depending on the individual user's tolerance. In some cases extra effort put into ensuring flawless packet delivery can even hurt the 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.

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%

TABLE II: Commonly used voice and video codecs and their packet loss tolerance.

In table II 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. Table II indicates that the threshold is often significantly above zero. Therefore, if relaxing the packet delivery constraints increases protocol efficiency, we can allow PER to be non-zero but below the application packet loss tolerance.

B. 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 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 the communication protocol's energy inefficiencies. Fortunately, these two instances can highly benefit from context aware PHY layer adaptation.

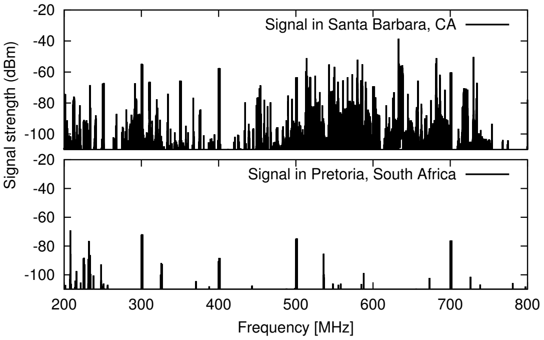

Fig. 3: Spectrum scan in Santa Barbara and Pretoria.

In section II-B we analyzed the impact of PHY parameters on energy consumption. The need for energy efficiency depends on the battery charge, which is usually programmatically available on mobile devices. 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 [19] we developed an energy flow model that allows fine grain tracking and prediction of energy trends. Our protocol designed in section IV 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.

C. Spectrum Demanding Operation

Spectrum is a finite resource and the problem of its scarcity has become widespread since the mobile and pervasive computing revolution of the last decade. The core issue stems from the fact that a large number of devices running bandwidth-hungry applications contend for the same frequencies. The problem is especially pronounced in high population density urban areas. High spatial and temporal correlation of the bandwidth demand among users often leads to intermittent connectivity with virtually no quality of service guarantees. Unfortunately, current protocols usually operate on a fixed central frequency and with a fixed bandwidth, which prevents them from using all the available resources for communication.

In figure 3 we show a measurement of TV spectrum in Santa Barbara, CA and Pretoria, South Africa. Different strategies are needed for efficient spectrum utilization in the two locations. In Pretoria, operation in a frequency division fashion, where each link resides on a separate band, is an attractive option due to the large number of vacant channels. If the application requires a change in the bitrate, modifying the channel width enables the change, often without the bit error rate penalty. On the other hand, when only a small part of the spectrum is free (Santa Barbara in fig. 3), a time ordered CSMA or TDMA scheme is favorable. In this case the bitrate improvement is achieved by the MCS modification, i.e. possibly impacting the bit error rate. Note that the available spectrum is often fragmented; our channel width and MCS adaptation protocol allows for non-adjacent subcarrier activation (fig. 1).

IV. WHITERATE

Physical layer parameters impact communication performance (section II-A), energy efficiency (section II-B) and spectrum utilization (section II-C). Often the impact is beneficial from one angle and detrimental from another. The usage context defines the importance of each aspect of the system. 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.

A. 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 packet error rate. As seen in section III, 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.

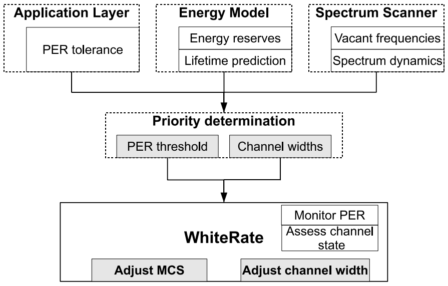

Fig. 4: Relation among the system components.

In figure 4 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 III-A, III-B, and III-C, respectively. *Priority determination* enables WhiteRate to answer to complex and perhaps contradicting demands through a practical algorithm. The key point of its operation is the fact that restriction on the PER and the guidelines on the available spectrum are sufficient to control rate adaptation: high rate positively impacts both energy savings and spectrum efficiency, yet can have adverse effect on packet delivery. Thus, we design WhiteRate to lower the transmission time, but maintain the packet delivery according to the imposed threshold. The threshold is set by the priority determination module and its exact value depends on the context as well as external properties, such as the system hardware and user preferences. In our evaluation we use a stub implementation of the modules in the dotted boxes and concentrate on WhiteRate.

B. 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. At the same time, WhiteRate 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. 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.

C. The Algorithm - Details

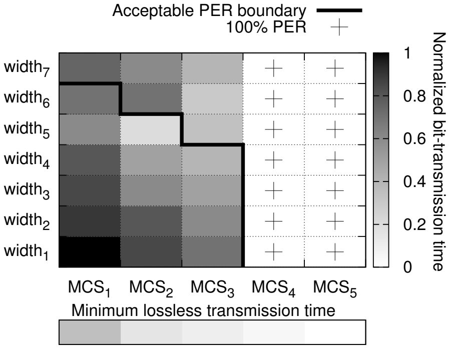

Fig. 5: Average bit-transmission times and PER for all [MCS, channel width] combinations.

We give a fictional example of the algorithm flow for a single source-destination pair in figure 5. The columns represent available MCSs and rows represent channel widths, which can be composed of non-adjacent subcarriers if needed. Wider channels/higher MCSs are labelled with a higher subscript value (widths from 1 to 7, MCSs from 1 to 5). Normalized transmission time per bit of an *[MCS, channel width]* combination is represented by different shades of grey, where lighter means better. Per-bit transmission times are initially unknown, and are computed as the algorithm visits each of the cells. Shades in the bar below the MCSs represent the minimum theoretically possible transmission times for the respective

MCS (with the widest channel), and are known in advance. Depicted is also an initially unknown region within which *[MCS, channel width]* combinations yield PER acceptable for the application. Those combinations that result in 100% packet loss are crossed out.

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 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. That is exactly what happens in our example. The algorithm starts transmission over the highest possible MCS and the widest available channel, $[MCS_5, width_7]$. 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 $[MCS_3, width_7]$.

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 5 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.

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 and spectrum usage. For that, 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. 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 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).

The WhiteRate algorithm bears some resemblance to SampleRate [7] in its PER guided rate adaptation. Similar to practical implementations of SampleRate[6], 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.

V. EVALUATION

A. Methodology

The GNUradio code base contains many low level functionalities necessary for the WhiteRate implementation, such as OFDM modulation/demodulation and packetized transmissions. Additionally, Jello [26] allows non-adjacent subcarrier activation. We extend this to support on-the-fly change of the modulation and coding along with the number of active OFDM subcarriers; we also build a simple MAC layer that handles packet flow control.

The processing latency between the USRP and the host machine prevents the ACK-based stop-and-wait transmission that we need in order to keep track of the packet error rate [17]. Additionally, each channel width change takes a few milliseconds as the GNUradio flow graph has to be stopped, reconfigured and then recompiled. Therefore, we use our USRP testbeds to obtain real-world channel behavior under different PHY parameters and in different environments, while we perform the WhiteRate evaluation by replaying the traces offline. This method allows artefact-free comparison of WhiteRate with solutions such as SampleRate, that do not require time-expensive GNUradio operations, as well as with "oracle" solutions that know the channel behavior over the whole parameter domain.

B. Experiment Setup

As described in the introduction, we utilize two testbeds to represent our two target operating environments. We install two USRP2 nodes with white space enabled WBX cards in suburban Pretoria, South Africa. We establish a direct 500m

---

[6] A version of SampleRate is used in the MadWiFi driver.

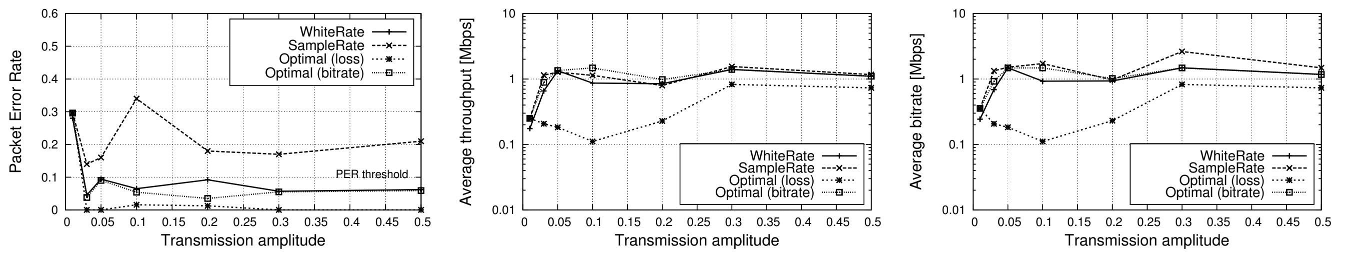

(a) Packet error rate.

(b) Throughput.

(c) Bitrate.

Fig. 6: Rate adaptation under varying transmission amplitude.

line of sight link between a tall building at the CSIR[7] and a nearby hill. We deploy a second testbed in a UCSB campus building in Santa Barbara, California. In the indoor testbed we have two USRP2 nodes with WBX cards comprising a 4m link. Besides their physical configuration, the two testbeds differ in the white space spectrum available at each of the locations, as presented in figure 3.

Our outdoor link can be maintained only with maximum signal power, thus we use the traces from our indoor testbed to analyze WhiteRate's performance under varying transmission amplitudes. We use the outdoor link to demonstrate how WhiteRate enables a trade-off between energy efficiency and communication quality. We revert back to the indoor testbed to analyze WhiteRate's ability to adapt in spectrum-scarce environments. We are especially interested in VoIP traffic, which is crucial in our target environments and has an attractive property of non-zero loss tolerance. Moreover, VoIP packet loss non-subjectively translates to call quality through the e-model [1] allowing for confident evaluation. To obtain traces, we send a train of 1000 VoIP packets, each 80B long, for each of the channel width-MCS combinations. We transmit over white spaces with the central frequency set at 694MHz. We (de)activate groups of 48 subcarriers; the minimum width is 174, while the maximum is 462 subcarriers, resulting in seven different channel widths. We use seven MCSs: BPSK, QPSK, 16-QAM with 1/2 and 3/4 rate convolution coding and 64-QAM with 3/4 convolution coding.

C. Experimental Results

Packet loss, Throughput and Bitrate

In this experiment we measure the average packet loss, throughput and bitrate as we deteriorate the link quality by lowering the transmission amplitude. The results are calculated over the whole experiment, which consists of 2000 packets sent at channel width-MCS combinations picked by the protocol under study.

Figure 6 compares WhiteRate and SampleRate fixed at the highest channel width[8], and the two empirically optimal *[MCS, channel width]* combinations. The first oracle combination, *Optimal(loss)*, is a fixed combination that provides the lowest PER. If multiple combinations result in the same PER, the one that yields the higher bitrate is selected. Similarly, *Optimal(bitrate)* is a fixed combination that provides the highest bitrate, while keeping the PER below the threshold if possible.

In Figure 6(a) we see that WhiteRate keeps the PER below the 10% threshold we imposed. It manages to do so until the link quality falls dramatically (transmission amplitude=0.01). SampleRate, on the other hand, does not explicitly take application requirements into account. The resulting PER is variable and above the threshold value. The two optimal strategies result in the PER below the threshold as long as the link is of good quality.

Despite being constrained by the PER threshold, WhiteRate manages to keep the throughput on 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.

We observe that a near zero PER has its price. *Optimal (loss)* uses robust MCSs and channels, thus provides only very limited bitrate and throughput. Ideally, *Optimal(bitrate)* draws the upper limit on the energy and spectrum efficiency achievable through rate adaptation. Yet, for any practical purposes an approach that simply pushes the bitrate while keeping the PER below the threshold is not acceptable. The MAC layer often implements automatic repeat request (ARQ) for error correction. The transmission time per bit on which WhiteRate relies takes into account any MAC overheads, while blindly sticking to the highest bitrate can be counter productive as energy and spectrum are wasted on retransmissions.

Energy Savings

To evaluate energy savings we fix the transmission amplitude and compare energy per bit performance as we vary the level of PER tolerance. We calculate the energy per bit by multiplying the time needed to transmit a bit of information with the platform power consumption. The first factor, time to transmit a bit of information, is directly influenced by the MCS and channel width used. The second one, power consumption, depends on the transmission device we use and differs drastically between an experimental SDR platform such as USRP2 and a commodity device such as a WiFi NIC. For the system evaluation purposes we use power consumption

---

[7] Council for Scientific and Industrial Research, Pretoria

[8] We also tested SampleRate with the lowest channel width, but leave the results out as the algorithm performance does not improve.

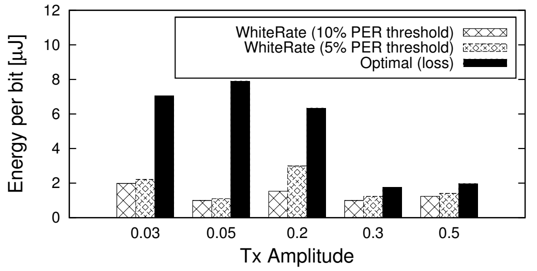

Fig. 7: Energy profile of various application PER requirements.

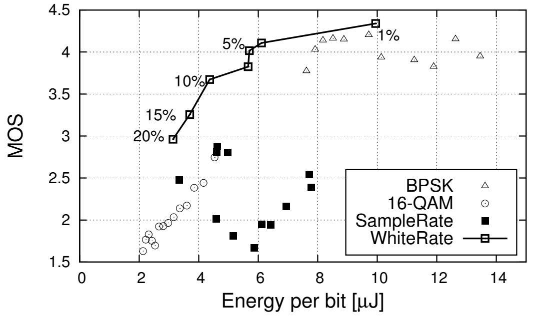

Fig. 8: Tradeoff between voice quality and energy savings.

figures of a USRP2 device from [18]. We discard the base power, since it is MCS and channel width independent and an order of magnitude higher than in the commodity devices and would incur bias towards any protocol that lowers the transmission time, including WhiteRate.

The results in figure 7 show that the energy per bit can be cut to a quarter of its initial value if the delivery requirements are relaxed to the PER threshold of 10%. Even a relatively tight error rate requirement (5%) yields multi-fold energy savings.

Energy Efficiency and Application Performance

We examine the tradeoff between energy efficiency and transmitted voice quality. As the application tightens the PER threshold and improves the voice quality we expect WhiteRate to have less space for energy optimization.

Perceived quality of received audio is often measured in *Mean Opinion Score (MOS)* units. The MOS is a metric derived from the results of subjective tests where listeners evaluate the quality of heard audio. The MOS scale runs from 1 (poor) to 5 (excellent) audio quality. The voice flow ranked above 3.5 is generally considered satisfactory. We use the *e-model* [2] to map network level performance (packet delay, jitter and loss) to the MOS metric. We adjust the e-model according to recommendations [3] for the G.711 codec.

We run a stream of packets to model G.711 encoded speech and vary the application PER requirement from 1% to 20%. At every point we measure the actual PER and substitute it into the e-model to obtain the MOS value. We also inspect the *[MCS, channel width]* combinations used and calculate the energy per bit at each point.

The results (figure 8) show the energy saving benefits of loosening the packet delivery requirements. Relaxing the threshold from 1% to 5% PER yields 43% energy per bit savings with minimal impact on MOS. In the same figure we plot the performance of SampleRate that we run independently under different channel widths (from 174 to 462 active subcarriers). SampleRate is application-agnostic and fails to provide the needed MOS. 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 (we calculate the statistics and plot a point every time we change the channel width). Fixed MCS gives us very 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 imposed (5%), WhiteRate halves the energy needs of BPSK, while significantly improving on 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.

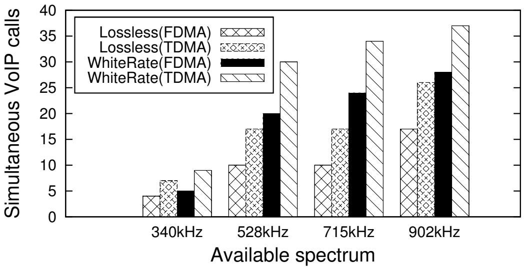

Fig. 9: Spectrum efficiency: Number of supported VoIP calls with WhiteRate (10% loss tolerance) and the lossless approach.

Spectrum Utilization

Real-time communication has well defined packet delay and jitter requirements. As a result, the system has to guarantee a certain throughput to real-time flows.

We analyze two access schemes, FDMA and TDMA, that guarantee quality of service to VoIP flows. We take a single TV white space channel and divide the spectrum among a number of emulated links. Each of the links has the same properties as a physical link we established in the Santa Barbara testbed; all links are assumed to belong to the same same collision domain. We change the amount of spectrum on the channel that we mark as available and calculate the number of flows that can be supported simultaneously. In the FDMA case each link is assigned the same amount of bandwidth, while in the TDMA case, we use the full band for all links and try to pack as many flows as possible in time.

In figure 9 we show the change in the number of supported flows with the available bandwidth. Due to relaxed delivery constraints WhiteRate uses higher MCSs, thus needs less width for the same throughput than *Lossless*, which is application-agnostic and strives for zero packet loss. As a result, in the FDMA case more calls can be packed in the same band. When TDMA is used, by maximizing the bitrate according to the PER threshold WhiteRate minimizes the transmission time. Consequently, more links can send their packets in a 60ms gap between two successive VoIP packets. We note that figure 9 depicts relationship between WhiteRate and Lossless, yet the difference between TDMA and FDMA depends on a multitude of factors such as frequency guard bands and MAC overheads.

VI. RELATED WORK

High flexibility of PHY parameters opens tremendous possibilities for context-aware operation. In [26] Yang et al. propose a system for agile spectrum allocation. In their work application throughput demand drives spectrum distribution. We extend this idea further with application PER requirements and the device's energy profile. In SoftCast [14] the authors provide real-time traffic delivery improvement through cross-layer video encoding. However, unlike WhiteRate, their approach is restricted to one specific application.

The goal of most wireless rate adaptation protocols is to maximize throughput [7], [25]. The introduction of SDR platforms provided rate adaptation protocols with sophisticated PHY-level information such as noise and signal levels [21], decoder confidence levels [24] and channel coherence time [10]. To best of our knowledge, energy and spectrum efficiency have not been addressed in SDR-based rate adaptation. Energy-efficient rate selection for real-time traffic has been proposed in [27]. The authors take a strictly theoretical approach, assume that all consumed energy goes towards useful transmission, and come to the conclusion that the slowest rate that satisfies the application delivery delay constraints is the most energy efficient one. In the preliminary analysis we showed that energy consumption has to be observed in its entirety, and then the highest bitrate becomes the most energy efficient.

VII. CONCLUSION

Wireless network proliferation has resulted in a high demand for wireless resources and current context-agnostic paradigms have proven to be under-performing when these resources are limited. Recent progress in the field of software-defined radio and the release of white space spectrum allows us to consider clean slate solutions that fully utilize the physical layer flexibility.

In this paper we propose WhiteRate, a solution that harnesses the flexibility to answer to the specific operation conditions. Our protocol adapts channel width and modulation and coding scheme according to the requirements of the context: packet error tolerance, energy reserves and available spectrum. WhiteRate is a cross-layer protocol that listens to the system for hints, yet it retains packet-based operation that is easily and incrementally implementable. Using our GNUradio testbed we have shown that WhiteRate delivers both energy efficiency and spectrum utilization improvements while meeting the application delivery requirements. Moreover, WhiteRate allows us to fine tune the balance between the system efficiency and the application performance.

In the future, we see further interleaving of resource utilization and communication adaptation as a large variety of PHY parameters can be manipulated. 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.

ACKNOWLEDGMENTS

We would like to thank David L. Johnson for the spectrum scans. This work was funded in part by a grant from Google.

REFERENCES

[1] http://www.voiptroubleshooter.com/diagnosis/emodel.html.

[2] *ITU-T Recommendation G.107, "'The E-Model, a computational model for use in transmission planning"*, December 1998.

[3] *ITU-T G.113 "Appendix I: Provisional planning values for the equipment impairment factor Ie and packet-loss robustness factor Bpl"*, May 2002.

[4] Mobile Broadband: The Benefits of Additional Spectrum. Tech. report, FCC, October 2010.

[5] Atheros. Power Consumption and Energy Efficiency Comparisons of WLAN Products, 2003.

[6] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh. White Space Networking with Wi-Fi Like Connectivity. In *SIGCOMM'09*, Barcelona, Spain, Aug. 2009.

[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] J. Camp and E. Knightly. Modulation Rate Adaptation in Urban and Vehicular Environments: Cross-layer Implementation and Experimental Evaluation. In *MobiCom'08*, San Francisco, CA, September 2008.

[11] R. Chandra, R. Mahajan, T. Moscibroda, R. Raghavendra, and P. Bahl. A Case for Adapting Channel Width in Wireless Networks. In *SIGCOMM'08*, Seattle, WA, Oct. 2008.

[12] 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.

[13] J. Heiskala and J. Terry. *OFDM Wireless LANs: A Theoretical and Practical Guide*. SAMS, Indianapolis, IN, 2001.

[14] S. Jakubczak, H. Rahul, and D. Katabi. One-Size-Fits-All Wireless Video. In *HotNets'09*, pages 1–6, New York City, NY, October 2009.

[15] 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.

[16] N. Mishra, K. Chebrolu, B. Raman, and A. Pathak. Wake-on-WLAN. In *WWW'06*, Edinburgh, Scotland, May 2006.

[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, 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.

[20] J. G. Proakis. *Digital Communications*. McGraw-Hill, 2000.

[21] H. Rahul, F. Edalat, D. Katabi, and C. Sodini. Frequency-aware Rate Adaptation and MAC Protocols. In *MobiCom'09*, Beijing, China, Sept. 2009.

[22] 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.

[23] Ubiquiti Networks. *XR5 Power Consumption Study*. www.ubnt.net.

[24] M. Vutukuru, H. Balakrishnan, and K. Jamieson. Cross-layer Wireless Bit Rate Adaptation. In *SIGCOMM'09*, Barcelona, Spain, August 2009.

[25] 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.

[26] 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.

[27] 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.