Energy Efficient Communication in Next Generation Rural-Area Wireless Networks

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

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

ABSTRACT

White space frequencies are highly attractive for long-distance communication due to greater signal propagation. The lack of standards and licensing issues with increased flexibility provided by the cognitive radio allow for sophisticated customized solutions for white spaces. Rural-area networks are seen as the main beneficiaries and white spaces communication is expected to outperform current wireless solutions in this domain. However, rural networks often have to rely on a constrained energy budget and highly benefit from energy-efficient operation.

We investigate the efficiency of flexible wireless transmission over long-distance white space links. We theoretically and experimentally examine the impact of channel width, modulation and coding and transmission amplitude on energy consumption. From our findings we derive the physical layer parameter settings that achieve energy optimality and develop PowerRate, a protocol that dynamically adjusts transmission parameters according to channel state. We implement PowerRate in GNUradio and evaluate its energy-saving potential in various fading environments.

Categories and Subject Descriptors

C.5 [**Computer-Communication Networks**]

General Terms

Algorithms, Design, Experimentation, Performance, Theory

1. INTRODUCTION

It is estimated that about two thirds of the world's population live in rural areas of, mostly developing, world. In addition to the low standard of living, poor education and infrastructure, these people often lack even basic communication amenities. Access to online services can directly lead to more opportunities for economic development, while enabling global connectivity prevents the digital divide and facilitates equal economical growth. In the mid 2000s providing rural connectivity by means of wireless networks based on commodity WiFi hardware emerged as a viable but suboptimal alternative to expensive wired solutions [14, 11].

A new paradigm, "white space WRANs", promises to revolutionize the way rural connectivity is achieved. "White spaces" encompass frequencies around 700MHz that used to be reserved for television broadcast, but are now becoming free, thanks to the roll-out of digital television (DTV). Due to the propagation characteristics, for the same equivalent isotropically radiated power (EIRP), the communication range in the white spaces is an order of magnitude larger than in the frequencies occupied by WiFi or WiMax [4]. This renders white spaces extremely attractive for rural area wireless networks. Moreover, the lack of licensing issues permits community networking in the environments where commercial providers have no incentive to deploy infrastructure, such as the developing world or low-income areas. Finally, the lack of protocols operating in white spaces allows us to reconsider the existing spectrum access schemes and fully utilize the PHY parameters' flexibility to increase communication efficiency [19, 13, 18]. Low-level parameters are no longer fixed, but through software defined radio (SDR) allowed to adjust to the channel state and/or application needs.

Despite their superior communication performance, next generation white-space-based networks will still face the same obstacles that current WiFi networks experience when deployed in the rural areas of the developing world. The lack of reliable energy supply is the most commonly reported problem in rural area wireless networks [16, 1]. If it exists, the grid infrastructure is often poor, and networks have to rely on alternative sources such as wind and solar; this is why energy efficient operation is of key importance.

Compared to WiFi hardware, SDR platforms enable flexibility that uncovers much wider space for energy efficiency optimization. Parameters such as signal modulation, channel width and transmission amplitude influence not only communication quality but also power consumption. To the best of our knowledge, the problem of identifying the energy optimal PHY settings for flexible wireless transmissions remains unaddressed. Some theoretical work in the field of energy efficiency is applicable in this domain [9]; however, the experimental analysis of power consumption and insights that can directly translate to network protocols are not yet available.

In this paper we examine energy efficiency in the flexible wireless transmission space. We start with a theoretical analysis and extend the existing postulates about energy efficiency with considerations about the solution practicality. We experimentally determine the power consumption and the delivery performance of emulated long distance links and highlight the differences between theory and practice. From our findings we develop a protocol that controls the transmission in an energy efficient way. Our protocol monitors the link state and decides to transmit with the optimal PHY parameters or to defer until more favorable conditions are observed, while keeping the application needs satisfied. Through GNUradio-based experimentation we show that PowerRate provides substantial energy savings in fading channels expected in rural areas.

2. FLEXIBLE RADIO COMMUNICATION

Orthogonal frequency-division multiplexing (OFDM) has emerged as the technique of choice for many wireless standards such as

CoRoNet'10, September 20, 2010, Chicago, Illinois, USA.

WiFi, WiMax and IEEE 802.22 (WRAN in white spaces) due to its robustness in harsh channel conditions. We analyze energy efficiency in a wireless standard agnostic way, but we restrict it to an OFDM scheme. An OFDM channel consists of multiple narrowband subcarriers. A high level of flexibility can be achieved if we manipulate the OFDM subcarriers. For example, varying the number and distribution of active subcarriers can help avoid harmful in-channel interference, while adapting signal modulation to the subcarrier state improves communication robustness. Recent research efforts have demonstrated the practicality of such highly adaptable OFDM on SDR platforms [19, 13].

In this paper we analyze the effect of changing subcarrier allocation[1], signal modulation and the transmission amplitude on energy efficiency. These parameters impact both communication quality and power consumption substantially and in a non-trivial way.

3. ENERGY CONSIDERATIONS

Communication equipment is a major energy consumer in rural-area wireless networks [1]. In digital communications, as a general rule, energy consumption is lowered by either shortening transmission time or lowering transmission power. Higher bitrates lower the transmission time, but are sustainable only when the power is high enough to result in sufficient SNR. Thus, unless we allow for data to be dropped, a tradeoff between the time and the power exists.

The theoretical relationship between bitrate and transmission power is given by Shannon's formula, which defines the boundary of the channel capacity. Since the formula does not provide a means to achieve the boundary bitrates, a theoretical solution can be practically infeasible. Moreover, in theory, the transmitter power is usually analyzed in isolation, while in reality the transmitter needs supporting hardware, which has non-zero power consumption.

Practical solutions, that determine the energy optimal communication settings by exploring the domain of parameter's values, work well when the domain is small [2]. However, once the domain defined by the number of channel width values, modulation and coding schemes (MCSs) and the resolution of transmission power knob becomes large, such as in flexible white space communication, we need to rely on the theory to restrict the empirical search range.

4. THEORETICAL INVESTIGATION

With some approximation (discarding the guard intervals), we can consider the subcarriers individually, and for each of them Shannon's formula defines the maximum achievable bitrate:

$$(1) \quad R_i = Wlog_2\left(1 + SNR_i\right) = Wlog_2\left(1 + \frac{P_{Tx,i}g_i}{N_0 W}\right),$$

where W represents the bandwidth occupied by a single subcarrier. SNR_i represents signal-to-noise ratio, g_i channel gain and $P_{Tx,i}$ transmission power at the i^{th} subcarrier. N_0 represents power spectral density of white Gaussian noise.

A bit of information is transmitted with energy:

$$(2) \quad E_{Tx} = \frac{P_{Tx}}{R}$$

where $P_{Tx} = \sum_{i=0}^{k} P_{Tx,i}$ and $R = \sum_{i=0}^{k} R_i$ are the cumulative transmission power and bitrate of all k active subcarriers, respectively. If we assume that all subcarriers operate at the same transmission power level ($P_{Tx,i} = p_{Tx}, \forall i$) over a flat fading channel ($g_i = g, \forall i$) the energy consumption becomes:

$$(3) \quad E_{Tx} = \frac{kp_{Tx}}{kWlog_2\left(1 + \dfrac{p_{Tx}g}{N_0 W}\right)}$$

[1] We use *subcarrier allocation* and *channel width* exchangeably.

From the above equation we observe: *(i-a)* the number of active subcarriers (k) does not influence the energy consumption, and *(ii-a)* the most energy efficient communication is the one that uses the lowest possible transmission power and bitrate per subcarrier.

However, the conclusions hold only if the transmission power is the sole factor that consumes the energy. It has been shown [9] that a significant part of the energy goes to the *transceiver circuit power* (P_{TC}), which takes into account the consumption of device electronics, such as mixers, filters and DACs, and is bitrate independent. With a non-zero P_{TC} the energy consumption is:

$$(4) \quad E_{Tx} = \frac{kp_{Tx} + P_{TC}}{kWlog_2\left(1 + \dfrac{p_{Tx}g}{N_0 W}\right)}$$

Miao et al. [9] proved that in this case: *(i-b)* energy efficiency increases with an increasing number of active subcarriers. *(ii-b)* the energy optimal transmission power and bitrate are not the lowest ones but depend on the P_{TC} and k values.

The bitrate used in the calculations represents an upper bound. In physical systems the choice of MCS determines the actual bitrate. This bitrate is below the optimal for the given SNR, but is equal to the optimal for a channel with a factor Γ lower SNR. This factor is called the "SNR gap" and depends on the MCS used, as well as the desired bit error rate (BER) [9]. The energy per bit becomes:

$$(5) \quad E_{Tx} = \frac{kp_{Tx} + P_{TC}}{kWlog_2\left(1 + \dfrac{p_{Tx}g}{N_0 W\Gamma}\right)}$$

Finally, this equation considers the bitrate as a continuous function, while in reality we have a very limited number of modulation and coding schemes to pick from. The energy per bit under a given modulation level (M) is:

$$(6) \quad E_{Tx} = \frac{kp_{Tx} + P_{TC}}{kWlog_2 M}$$

M has to be bounded by the capacity formula, i.e. $M \leq 1 + \frac{p_{Tx}g}{N_0 W\Gamma}$, and the equality yields the energy optimal solution. However, in digital communications $M = 2^n$ where n is a positive integer. Thus, $M = \lfloor 1 + \frac{SNR}{\Gamma} \rfloor_{4,8,16,\ldots} = \lfloor 1 + \frac{p_{Tx}g}{N_0 W\Gamma} \rfloor_{4,8,16,\ldots}$, where $\lfloor x \rfloor_{4,8,16,\ldots}$ levels x to the nearest floor value in a set of achievable modulation levels (for QAM these are 4, 8, 16, etc.).

In Figure 1 we plot the energy-per-bit function in three flavors[2]: when neither the SNR gap nor the feasibility of the modulation levels is considered ($\Gamma = 0dB$; labeled "Shannon limit based"); when the SNR gap corresponds to the case of QAM modulation and coding that yields 6dB gain ($\Gamma = 3.8dB$[3]; labeled "Theoretical QAM coded"); and when both SNR gap and the feasibility of the modulation levels is considered ($\Gamma = 3.8dB$; labeled "Feasible QAM coded"). As a consequence of the modulation level discreteness, the energy is a piecewise linear function of the transmission power. A label next to each slope marks the highest feasible modulation level that yields a BER below the threshold. Points where each of the modulation levels touches the "Theoretical QAM coded" curve are the points where the modulation levels are at the edge of sustaining the desired BER.

From the plot we see that the optimal solution (${P_{Tx}}^*, {E_{Tx}}^*$) is not on the "Feasible QAM coded" line unless ${P_{Tx}}^*$ results in just enough SNR that the next modulation level becomes the optimal one at that transmission power. When this is not the case, as in our

[2] In the example we use a flat fading link with 70dB path loss, 50 subcarriers, each 1kHz wide. $P_{TC} = 0.2W$, $BER < 10^{-7}$

[3] Uncoded QAM introduces $9.8dB$ SNR gap; with $6dB$ coding gain the resulting gap is $3.8dB$.

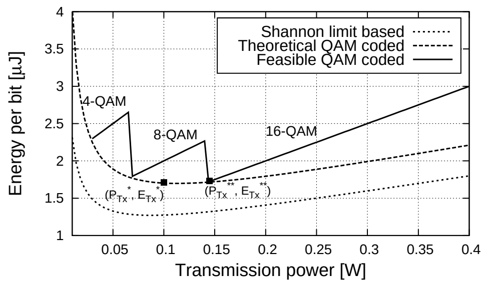

Figure 1: Energy per bit vs. transmission power (P_{Tx}).

example, the minimum feasible energy might not even be achieved with the modulation level which is optimal for ${P_{Tx}}^*$. In Figure 1, the minimum achievable energy consumption is located at ${P_{Tx}}^{**}$ not ${P_{Tx}}^*$, thus is provided with 16-QAM, not 8-QAM.

We augment *(i-b)* and *(ii-b)* with our analysis results:

- ***observation 1***: for a given transmission power level the optimal coding scheme is the one that minimizes the SNR gap Γ, while the optimal modulation level is the highest one for which the BER remains below a desired threshold.[4]
- ***observation 2***: the optimal transmission power is located at one of the points where the available modulation levels perform at the edge of the acceptable BER, i.e. at one of the points where the *feasible* characteristic touches the *theoretical* one.
- ***observation 3***: energy efficiency increases with the increasing number of active subcarriers.

The observations point out that the energy optimal transmission depends on previously unconsidered parameters, such as the feasibility of modulation levels, coding scheme and BER threshold.

5. EXPERIMENTAL INVESTIGATION

Software-defined radio platforms gained wide popularity only recently and little work that investigates their performance has been published [10]. To the best of our knowledge, none of the previous work studies energy efficiency. In order to address this gap we conduct an energy consumption study of the USRP/GNUradio platform due to its high prominence in academic research, support for communication in white spaces and comparatively low cost.

Experimental Testbed. We use two different testbeds for power consumption and communication performance measurements. We measure the power consumption in a local lab where we have physical access to the nodes, while we use CMU's channel emulator [6] to model a long-distance rural wireless link and profile its performance in a controlled environment. In both testbeds we create links between nodes composed of USRP devices and PCs running GNUradio software. In the local testbed we instrument one node with a multimeter to measure the current drain of a USRP powered by a constant DC voltage adapter. The main difference between the two setups is that in our lab we have a newer model of USRP (USRP2), while the CMU emulator uses the older version of the hardware.

For the experiments we split a 320kHz wide band into 512 FFT bins. Each of the bins can host one active OFDM subcarrier. The GNUradio code base allows for BPSK, QPSK, 8-QAM, 16-QAM, 64-QAM and 256-QAM modulation.

5.1 Power Consumption

In a flexible OFDM system the active subcarriers can be distributed inside the channel in an arbitrary manner. We implement

[4]For simplicity we discarded the coding overhead; in practice, it can be considered individually for each of the supported codes.

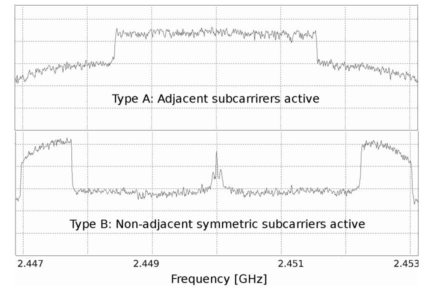

Figure 2: Spectrum analyzer output for two distinct subcarrier allocations.

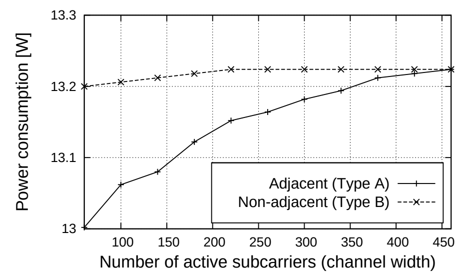

Figure 3: Power consumption - varying number and position of active subcarriers.

the two spectrum sculpting cases that are the most interesting for real-world deployments (Figure 2). In the first case (Type A) the active subcarriers are symmetrically distributed around the center frequency. This corresponds to a channel bonding case in which adjacent TV channels are sensed to be free, bonded and used by our device. In the second case (Type B) active subcarriers are as far away from the central frequency as possible. Such an allocation can be observed if mutually distant TV channels are bonded. We expected the transmission power (P_{Tx}) to grow linearly with the number of active subcarriers, but remain independent of the subcarrier distribution.

In Figure 3 we show the total USRP2 power consumption for the two types of OFDM transmissions depicted in Figure 2. We see that, unless the number of active subcarriers is high, the adjacent active subcarriers are more energy efficient. We speculate that either energy leakage due to improper filtering or the frequency-dependable power efficiency of the transmitter's power amplifier is the reason for this discrepancy [7]. We are aware that with high probability this phenomenon is specific to the hardware we use. However, to discard possible biases, in the rest of the paper we focus on communication over more efficient adjacent subcarriers.

The total power consumption (P_{total}) consists of the base USRP power (P_{base}) needed for keeping a board powered on; the transceiver circuit power (P_{TC}) consumed by the supporting hardware when the card is in the transmission mode; and the actual transmitted signal power (P_{Tx}). A breakdown of the consumption allows us to account for P_{Tx} and P_{TC} when calculating the energy optimal transmission. We modify the GNUradio exposed parameter *tx_amplitude* and channel width and record the total consumption.

We summarize the findings in Table 1. We notice that the power consumption is indeed constant when the device is in the transmission mode with *tx_amplitude* set to zero, no matter what the channel width is. However, the consumption is 1.01W higher than in the case when the device is idle, although no data is being sent in either case[5]. The 1.01W difference goes towards the transceiver circuit power. The transmission power varies from 0W to 0.49W, thus in the transceiver circuit power is more than comparable to

[5]We confirmed this with a spectrum analyzer.

Transceiver circuit power	
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.	
Min power when transmitting $P_{total}(tx\_amp = 0, any\_width)$	12.77 W
Max power when transmitting $P_{total}(tx\_amp = 1, max\_width)$	13.26 W
Transmission power range P_{Tx}	**[0W, 0.49W]**

Table 1: Power consumption breakdown.

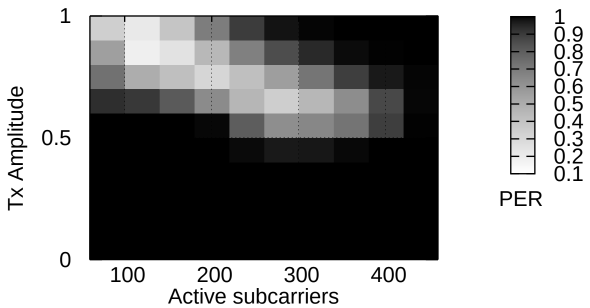

Figure 4: Packet error rate (PER) with varying *tx_amplitude* and channel width.

the transmission power[6]. As a consequence, P_{TC} should not be overlooked from the energy optimal transmission rate calculation.

We also measure the transmission power against *tx_amplitude* and channel width. Both parameters, as expected, stand in a nearly linear relationship with P_{Tx}, thus we omit the results for brevity.

We would like to emphasize the difference in the orders of magnitude between the transmission power ($\sim 1W$) and the total power consumption ($\sim 10W$). It suggests that any transmission energy optimization can be overshadowed by the high consumption of an idle device. A power-saving "sleep" mode, currently not present in SDRs, would reduce the idle power consumption, but it would also change the way in which we think about energy consumption optimization. Instead of optimizing the transmission energy, we would have to consider the total power consumption and sacrifice the transmission power savings for the shorter sending times that would get us into the "sleep" mode as quickly as possible. In [12] Radunovic et al. consider an energy-aware SDR design based on Lyrtech hardware, while in [8] Liu et al. implement a rapid sleep mode in the SDR. We believe that low-power mode will become available as the SDR becomes a part of commodity devices.

5.2 Communication Performance

Communication performance over long-distance links operating at a fixed channel width has been studied in [15, 3]. The studies find that in rural areas the link abstraction holds and that external interference is the main source of dropped packets. The impact of channel width change has been addressed in [2]. The authors find that wider channels result in lower communication range due to lower power-per-Hz and higher susceptibility to delay spread in a multipath environment. However, in their implementation they modify the reference clock of a commodity Atheros WiFi chipset, and keep the same number and position of occupied OFDM subcarriers; therefore the results cannot be generalized to our case.

We measure the packet error rate (PER) on an emulated long-distance link which we keep interference and multipath free. We send 20B packets back-to-back and change the channel width and transmission amplitude every 2000 packets. The packets are not encoded and, unlike the commodity solutions, the GNUradio is not engineered to perform well in high path loss settings; therefore we observe high PER. In Figure 4 we plot the PER at different transmission amplitudes and channel widths. Unexpectedly, we see that at a fixed *tx_amplitude* the transmission is successful only at a subset of the available channel widths. Moreover, the range of well-performing widths changes as we modify the amplitude. To ensure this is not an artifact of a faulty USRP we confirmed the results with two different devices in an interference free lab environment.

Although we were not in a position to experimentally verify our assumption, we believe that the GNUradio's OFDM implementation is limited with respect to the amount of spectral energy it needs for successful decoding. In Figure 4 we observe that wider channels need lower transmission amplitude to get the same PER as the narrower channels operating at higher amplitudes. Since the GNUradio OFDM implementation is often used in the academic research [19, 18] it is important to be aware of this shortcoming.

Hardware deficiencies aside, channel width change can influence the communication quality if the frequency-selective fading is present. In the rural area long-distance setting we expect slow flat fading[7] caused by low mobility and large physical obstacles, also known as *shadowing*, to be the predominant fading effect.

6. ENERGY OPTIMAL PHY SETTINGS

We concluded Section 4 with three observations that represent guidelines for the energy efficient PHY parameters settings. We proceed with a practical solution to parameter adjustment in a slow fading flat channel called *PowerRate*.

Optimal parameters - time invariant case. The energy optimal modulation level (M^*) and transmission power (P^*_{Tx}) can be determined from eq. (5). However, according to *observation 2* the optimal transmission power has to be set to one of the points where the highest modulation level that keeps the BER under the threshold changes. Fortunately, only two points have to be checked - where $\lfloor M^* \rfloor_{4,8,16,...}$ crosses the energy efficiency curve and where the next highest modulation level crosses it. Thus, if the channel gain is known, the search complexity is constant.

We concluded that the most efficient communication takes place over the widest possible channel (*observation 3*). In white spaces the channel availability is subject to the presence of the primary users - TV stations and wireless microphones. Primary detection is an active research area [5] and we assume that our energy efficient parameter adjustment solution relies on a module that provides correct identification of vacant frequencies. Once the availability is known, we allocate the largest free contiguous chunk of spectrum.

Optimal parameters - time varying case. In this case, determining the optimal settings for each of the periods where the channel gain is approximately constant might not result in the overall energy optimal solution. To see why, consider the case of a channel whose state changes from very good to poor and back. It can be tempting to calculate the optimal PHY settings in each of the three periods individually. Nonetheless, that can lead to efficient communication during the high SNR periods that gets undermined with the inefficiency of forcing the communication during the low SNR period. The efficiency could be increased if the sender defers from the transmission. The choice of whether to transmit or not is not a simple one as the sender needs to know the future channel state before it makes the decision. Moreover, there is a non-zero energy cost associated with delaying the transmission until the channel improves. This cost depends on the system implementation and the hardware

[6] We speculate that the measured P_{Tx} includes the power amplifier inefficiencies, thus is higher than the data sheet reported value.

[7] The amplitude and phase change imposed by the channel can be considered roughly constant over the packet duration.

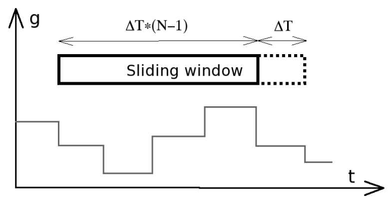

Figure 5: PowerRate protocol - sliding window over a time varying channel.

platform: idle power consumption, existence of low-power sleep modes, and state switching energy, among others. Finally, different applications exhibit different levels of vulnerability to high packet delivery delay and jitter. Note that our problem is more challenging than the power allocation by water-filling in time [17] since the optimal transmission power level varies with the channel state.

6.1 Algorithm - PowerRate

We derive a suboptimal but practical solution to the problem by limiting the time window over which we consider the energy efficiency. As summarized in Figure 6.1, first we divide time into slots where the slot length (ΔT) corresponds to the fading speed, i.e. the channel gain within a slot is roughly constant. Then, we consider a moving window of slots over time. The window size is selected according to the application, such that the application does not suffer if no data is transmitted during the length of the window. Finally, we move the window one slot at the time, and in each iteration we calculate the optimal transmission parameters for that time slot and evaluate the utility function to determine whether it is more efficient to transmit in the current slot or to defer from sending. To guarantee that the maximum packet delay remains bounded, if the decision not to transmit was made in each of the previous time slots of the window, the current slot is used for transmission irrespective of the utility function value.

The utility function. At the beginning of each time slot, based on the observed channel gain, the transmitter makes a decision to use this slot or not. The slot will be used if by transmitting in it the transmitter increases its energy efficiency. Let n be the number of time slots within the time window, E_{n-1} the energy consumed in the first $n-1$ slots and S_{n-1} the amount of data transmitted in the first $n-1$ slots. The utility represents the amount of data sent per unit of the total device energy in the first $n-1$ slots:

$$(9) \quad U_{n-1} = \frac{S_{n-1}}{E_{n-1}}$$

If the n^{th} time slot is used for transmission the utility function is:

$$(10) \quad U_n = \frac{S_{n-1} + kW log_2\left(1 + \frac{{p_{Tx}}^n g^n}{N_0 W \Gamma}\right)\Delta T}{E_{n-1} + (k{p_{Tx}}^n + P_{TC})\,\Delta T + P_{base}\Delta T}$$

where ${p_{Tx}}^n$ represents the optimal power level per subcarrier in the n^{th} time slot, characterized by channel gain g^n. It follows that the utility can increase or decrease depending on the function value in the previous $n-1$ slot and the power/bitrate in the current slot.

If the transmission does not take place in the n^{th} time slot the utility function still changes because of the energy needed to keep the transmitter running. The actual amount of that energy depends on the power of the idle device. In the ideal case the device can go to a very low power sleep mode essentially making U_{n-1} equal to U_n. In the most inefficient case the device consumes the same amount of base power as when transmitting. We restrict ourselves to the latter case as it is more general. The utility function, in case the transmission does not take place in the n^{th} time slot is:

$$(11) \quad U_n = \frac{S_n}{E_n + P_{base}\Delta T}$$

PowerRate compares the utility value in case the transmission takes place (10) and in case it does not take place (11) in the current time slot, and decides on whether to transmit or not.

6.2 Implementation

We prototype PowerRate in GNUradio, relying on the GNUradio and Jello code base [19]. To obtain the channel gain in a timely manner we send a short packet with a known bit sequence at the beginning of each time slot. The transmitter uses the channel gain feedback to identify the most energy efficient transmission power and modulation level as per Section 4. The channel width is adjusted to avoid the bad performing *tx_amplitude*, channel width combinations mapped in Section 5.2.

USRP2 is not optimized for energy efficiency and two important consequences of the hardware properties prevent us from successfully evaluating PowerRate. First, for the given USRP2 characteristics the optimal power often lies above the device limits. Second, the base power of having the device on dominates and as a result it is always better to transmit data then to keep the device idle.

We evaluate the protocol as follows. We use the channel gain feedback to obtain measurements in various fading environments with GNUradio, and then rerun the traces offline under different power/rate allocation schemes. This allows us to have repeatable tests while experimenting with the hardware parameters.

6.3 Evaluation

We gather traces with USRP2s with white-space enabled WBX daughterboards. In an indoor setting we establish three fading environments: **1) static**, where a line-of-sight (LOS) link is maintained at all times; **2) slow dynamic**, where the link gets obstructed by human subjects and **3) fast dynamic**, where a substantial human mobility along with the rapid device movement disturbs the LOS. We feel that these cases represent a good starting approximation of the target area as they can correspond to a strong LOS link between villages, a local link between houses and an indoor access point.

PowerRate maintains the BER above the given threshold at all times, thus we compare it to a solution (labeled *Fixed*) that operates with fixed modulation and power levels that guarantee the BER threshold. We analyze two different flavors of PowerRate: one that takes the feasibility of modulation levels into account (*PowerRate-discrete*) and another that does not consider modulation level feasibility called *PowerRate-continuous*[8].

In Figures 6 and 7 we plot the transmission power and the bitrate in each of the fading scenarios averaged over five minutes. PowerRate adjusts the PHY parameters according to the observed conditions and lowers the transmission power substantially while not sacrificing the bitrate. The benefits are more pronounced for highly varying channels that provide more opportunities for rate/power adaptation. Interestingly, Fixed performs better than PowerRate-continuous in the static case. PowerRate-continuous identifies the theoretically optimal solution; however, that by itself is not enough - the feasibility of the modulation levels has to be considered as well. Before the previous experiment we tuned the transceiver circuit power to 100mW to get the optimal transmission power level within the USRP2 supported transmission power values. Next, we lower the base power to make it attractive for the PowerRate to defer from transmission in the case of a bad channel. For this case,

[8]Although the calculation is performed under the assumption that any modulation level ($M | M \in \mathbb{R}$) is feasible, the actual transmission is modulated with the feasible modulation level $\lfloor M \rfloor_{4,8,16,...}$.

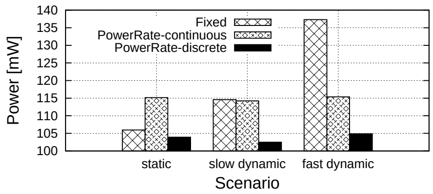

Figure 6: Average transmission power consumption.

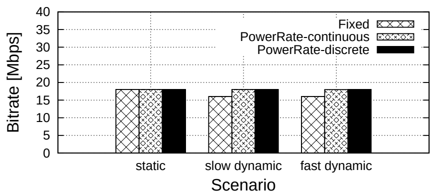

Figure 7: Average bitrate.

we plot the total utility, defined as the number of bits transferred divided by the total energy consumption, for each of the solutions after five minute runs (Figure 8). PowerRate approaches lead to better energy utilization in the two dynamic channel cases.

7. RELATED WORK

Communication performance in real world rural area networks has been analyzed in [3, 15]. These studies provide a detailed profile of the wireless medium in rural area and as such serve as an invaluable resource for our analysis. However, the networks that are subject to this analysis are WiFi-based and do not expose the same level of flexibility as the SDR. The concept of flexible PHY communication is well represented in [19] where the authors propose Jello, a MAC overlay that assigns the OFDM subcarriers to users according to the dynamic application needs, while FARA [13] takes into account frequency-selective fading and adapts bitrate on a per-subcarrier basis. Both approaches are energy agnostic and do not explore all of the PHY parameters concurrently.

From a large body of work on the energy efficient communication, [9] is the most relevant to our work, since it investigates the energy efficiency of a flexible OFDM transmission. We build upon this work, and we go a step further to the real-world modeling as we introduce the notion of physically possible modulation levels. Moreover, the authors of [9] approach the problem exclusively from a theoretical standpoint, while we derive our conclusions on energy-efficiency after a thorough experimental analysis and implement a practical solution to the parameter adjustment problem.

8. CONCLUSION

In this paper we presented a study of energy efficiency of emerging rural-area networks based on flexible wireless communication. We start from the state of the art approaches to energy efficient PHY parameter adjustment and add into consideration the notion of physically achievable modulation and coding schemes. We perform an experimental power consumption and communication performance analysis of the USRP2/GNUradio platform. From our analysis we derive guidelines for energy optimal communication settings and uncover the platform limitations. The generality of our conclusions can provide pointers for future work in this field.

We proceed with PowerRate, a protocol that dynamically adjusts the PHY parameters to achieve energy efficiency. PowerRate identifies the optimal modulation and coding scheme, channel width and transmission amplitude for the currently observed

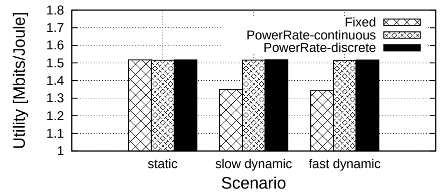

Figure 8: Utility - Energy efficiency (low base power case).

channel state. The protocol decides between the transmission with the optimal parameters or postponement based on the balance between the packet delivery and energy saving benefits while making sure that the application requirements are met. Finally, we implemented PowerRate in GNUradio and through initial trace-based experiments showed that it provides substantial energy savings in the case of a slow fading wireless channel.

9. ACKNOWLEDGMENTS

This work was funded in part through a gift from Google.

10. REFERENCES

[1] G. Bernardi, P. Buneman, and M. K. Marina. Tegola Tiered Mesh Network Testbed in Rural Scotland. In *WiNS-DR '08*, San Francisco, CA, Sept. 2008.

[2] 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, August 2008.

[3] K. Chebrolu, B. Raman, and S. Sen. Long-distance 802.11b Links: Performance Measurements and Experience. In *MobiCom'06*, Los Angeles, CA, Sept. 2006.

[4] 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.

[5] S. Deb, V. Srinivasan, and R. Maheshwari. Dynamic Spectrum Access in DTV Whitespaces: Design Rules, Architecture and Algorithms. In *MobiCom'09*, Beijing, China, September 2009.

[6] G. Judd and P. Steenkiste. Using Emulation to Understand and Improve Wireless Networks and Applications. In *NSDI'05*, Boston, MA, May 2005.

[7] J. Liu, H. Arslan, L. Dunleavy, A. Webster, and J. Paviol. Impact of Carrier Frequency Dependent Power Amplifier Behavior on 802.11a WLAN System. In *7th European Conference on Microwave Technology (EuMW)*, Amsterdam, Netherlands, October 2004.

[8] J. Liu and L. Zhong. Micro power management of active 802.11 interfaces. In *MobiSys'08*, Breckenridge, CO, June 2008.

[9] G. W. Miao, N. Himayat, and G. Y. Li. Energy-Efficient Link Adaptation in Frequency-Selective Channels. *IEEE Transactions on Communications*, to appear (accepted Apr. 2009).

[10] 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.

[11] R. Patra, S. Nedevschi, S. Surana, A. Sheth, L. Subramanian, and E. Brewer. WiLDNet: Design and Implementation of High Performance WiFi Based Long Distance Networks. In *NSDI'07*, Cambridge, MA, USA, April 2007.

[12] B. Radunovic, D. Gunawardena, P. Key, A. Proutiere, N. Singh, V. Balan, and G. Dejean. Rethinking Indoor Wireless Mesh Design: Low Power, Low Frequency, Full-duplex. In *WiMesh'10*, Boston, MA, June 2010.

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

[14] B. Raman and K. Chebrolu. Design and Evaluation of a new MAC Protocol for Long-Distance 802.11 Mesh Networks. In *MobiCom'05*, Cologne, Germany, August/September 2005.

[15] A. Sheth, S. Nedevschi, R. Patra, S. Surana, E. Brewer, and L. Subramanian. Packet Loss Characterization in WiFi-Based Long Distance Networks. In *INFOCOM'07*, Anchorage, AK, May 2007.

[16] S. Surana, R. Patra, S. Nedevschi, 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.

[17] D. Tse and P. Viswanath. *Fundamentals of Wireless Communication*. Cambridge Press, 2005.

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

[19] 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.