Adaptive Transmission for Opportunistic Use of Spectrum

Michael B. Pursley, Thomas C. Royster, IV and Jason S. Skinner

Clemson University, Clemson, SC, USA

ABSTRACT

An adaptive transmission protocol is described and evaluated for use in wireless communication networks that employ dynamic spectrum assignment. The protocol interacts with a spectrum assignment system to select the modulation format and symbol transmission rate for each session. The power is adjusted during the initial part of each session, and code-rate adaptation is employed to compensate for changes in propagation loss that occur during the session.

I. INTRODUCTION

New approaches to spectrum utilization depart from fixed frequency assignments and rely on opportunistic use of spectrum (e.g., [3]). These new approaches require more sophisticated radios that can change frequency bands, modulation formats, and error-control codes from one session to the next and even within a session. One component of such a radio is an adaptive transmission protocol that controls the changes in transmission parameters that are made from packet to packet within a session. In this paper, we examine adaptive transmission for opportunistic use of spectrum in a fully connected wireless communication network.

When one terminal, the *source*, wishes to send a set of packets to another terminal, the *destination*, a session is established between the two terminals and a request is made for the assignment of a frequency band. The adaptive transmission protocol is employed to select the modulation format and symbol rate that are used for the duration of the session. At the beginning of the session, the adaptive transmission protocol adjusts the transmitter power to match the initial channel conditions. As the session progresses, the code rate is adapted to compensate for changes in propagation loss. The protocol must interact with a *spectrum assignment system*, which may be a central controller or a network of devices with distributed protocols. If the terminals have cognitive radios (e.g., see [4] or [5]), then the source and destination terminals may be part of the spectrum assignment system. The spectrum assignment system gathers data and makes assignments of frequency bands, time slots, and power levels for the communication devices within its geographic area. Even though a source is assigned a frequency band for the duration of a session, it may not have exclusive use of the band. We anticipate that a variety of frequency reuse methods will be employed, including spread-spectrum multiple-access communication techniques.

A *packet* is an information packet unless specified otherwise (e.g., acknowledgment packet, control packet, or reservation packet). Our primary mode of operation requires the destination to send an acknowledgment packet to the source in response to each packet it receives from the source. If an acknowledgment packet is not received, the source retransmits the packet. In one possible alternative mode, a single acknowledgment packet is sent in response to a specified number of consecutive packets from the source.

The information bits are included in the *payload* of the packet (i.e., the data portion). The *overhead* is the time between the end of the transmission of the payload of one packet and the beginning of the transmission of the payload of the next. The overhead includes the time required to transmit the preamble, header, control fields, etc.; the time required for the destination to receive, demodulate, decode, and process the packet; and the time required for the source to receive and process the destination's acknowledgment packet. Of these contributions to the overhead, only the destination receiver's demodulation and decoding times depend in any significant way on the code and modulation used for the packet and the demodulation and decoding techniques employed in the receiver.

II. THE ADAPTIVE TRANSMISSION PROTOCOL

The adaptive transmission protocol operates in four phases: session initiation, request and assignment, power adjustment, and code-rate adaptation. In the initiation phase, the source and destination exchange session initiation messages, which may be similar to the request-to-send (RTS) and clear-to-send (CTS) packets that are employed in a variety of media-access protocols. The destination's reply message includes the set of frequency bands in which the destination's receiver is capable of operating, the set of time slots that the destination has available to receive packets during the session, and the set of modulation formats that can be used by the destination's demodulator. The *available modulation formats* are the modulation formats that are compatible with both the transmitting system at the source and the receiving system at the destination. In the initiation phase, the source and destination specify the *preferred completion time* for

This research was supported by the DoD MURI Program and the Office of Naval Research under Grants N00014-00-1-0565 and N00014-04-1-0563.

the session, which is chosen according to the session's quality-of-service priorities. The *completion time* for a session is the time required for the source to deliver all of the session's packets to the destination.

The power level for the first packet transmission in the session is also determined during the initiation phase. The selection of the initial power level is necessarily based on rough estimates, because the exchange of session initiation messages is unlikely to take place in the frequency band that will be used for the session. Moreover, the initiation messages are typically too short for accurate estimation of the propagation loss. The source's choice of the initial power level is subject to the approval of the spectrum assignment system, which must account for the potential disruption of other sessions by the interference from the source's transmissions during the proposed session.

The session-initiation phase provides information that enables the source to request time and frequency assignments and recommend an initial power level. In making assignments for the session, the spectrum allocation system must account for all previously approved sessions that will be active during the proposed session. It must also be aware of fixed allocations (e.g., for emergency services) used by receiving devices located within the estimated range of the source when it transmits at the proposed power level.

After the request has been submitted and the assignment is made, the source selects a modulation format and symbol transmission rate that are compatible with the bandwidth allocation and the preferred completion time. An *admissible modulation format* is an available modulation format that permits transmission at the preferred information rate within the allocated bandwidth. The full set $\{\mathcal{M}_j : 1 \leq j \leq n_m\}$ of modulation formats is indexed in order of increasing spectral efficiency (i.e., in order of decreasing bandwidth for a fixed information rate). Thus, the set of admissible modulation formats is a subset of $\{\mathcal{M}_j : J \leq j \leq n_m\}$ for some integer J. Each modulation format could be used in conjunction with OFDM or another form of multiple-carrier modulation, but such combinations are not evaluated in this paper.

Power adjustment takes place during the transmission of the first several packets of a session. The source obtains feedback from the destination, which permits it to adjust its power level to approximately the minimum that is possible for the desired probability of packet error. Especially for a long session, the spectrum assignment system may wait until the source adjusts its power before making other assignments in the same or adjacent frequency bands. By accounting for any reductions in the initial power that occur during the power-adjustment phase, the spectrum assignment system may be able to provide more efficient frequency reuse in the assignments for other sessions

For the rate-adaptation phase, the adaptive system has a set $\{C_i : 1 \leq i \leq n_c\}$ of codes that are indexed in order of increasing rate. The corresponding set of available code rates is $\{r_i : 1 \leq i \leq n_c\}$. In the current version of our adaptive transmission protocol, the code for the first packet of each session is C_{n_c}, the code of highest rate, and the selection of the modulation format from the set of admissible modulation formats is based on performance results for that code. This enables the source to compensate for moderate increases in propagation loss by reducing the code rate rather than increasing the power level, since an increase in power may increase the interference to other sessions that are in progress.

III. Modulation Formats and Codes

We illustrate the performance of the adaptive transmission protocol for bit-interleaved coded modulation [2] with a set of five binary turbo product codes that are available commercially on a single chip [1]. Several of the modulation formats that we evaluated are listed in Table I along with their null-to-null bandwidths for transmission of uncoded data at an information rate of 1 b/s. The M-ary biorthogonal signal set is the union of a set of $M/2$ orthogonal sequences of $M/2$ binary pulses (e.g., from the rows of an $M/2 \times M/2$ Hadamard matrix) and the set of complements of these sequences. If a source is granted exclusive use of a relatively narrow frequency band, QAM is clearly the best choice from Table I in terms of spectral efficiency. For a given bandwidth, 64-QAM provides six times the information rate of BPSK and 12 times the information rate of 16-ary biorthogonal modulation, which translates to shorter session times for QAM when operating in an interference-free frequency band. If a source is assigned a relatively wide frequency band, then biorthogonal modulation permits transmission at lower power than is possible with QAM, BPSK, or QPSK. Other advantages are that biorthogonal modulation does not have the large amplitude fluctuations of QAM and it can be used with a signature sequence to provide protection against multipath and multiple-access interference [7]. For direct-sequence spread spectrum with one sequence chip per modulation chip, there is no additional bandwidth expansion, yet the signal design permits communication in the presence of interference and provides the flexibility to schedule simultaneous sessions in the same frequency band.

Let E_b/N_0 denote the ratio of the energy per information bit to the one-sided density of the noise, and define $\mathrm{ENR} = 10\log_{10}(E_b/N_0)$. The values of ENR that are required to achieve a packet error probability of 10^{-2} are given in Table II for five turbo product codes employed with binary antipodal modulation (e.g., BPSK), coherent demodulation, and soft-decision decoding. The log-likelihood-ratio (LLR) metric is used for all modulation formats except QAM. A simpler metric [6] with approximately the

Notation	Modulation	Bandwidth
$\mathcal{M}_1$	64-ary Biorthogonal	10.67 Hz
$\mathcal{M}_2$	32-ary Biorthogonal	6.40 Hz
$\mathcal{M}_3$	16-ary Biorthogonal	4.00 Hz
$\mathcal{M}_4$	8-ary Biorthogonal	2.67 Hz
$\mathcal{M}_5$	BPSK	2.0 Hz
$\mathcal{M}_6$	QPSK/4-QAM	1.0 Hz
$\mathcal{M}_7$	16-QAM	0.50 Hz
$\mathcal{M}_8$	64-QAM	0.33 Hz

TABLE I

Modulation formats

Code	Rate	Block Length	Capacity Limit	ENR	Difference
C_1	0.236	2048	-0.8 dB	1.5 dB	2.3 dB
C_2	0.325	4096	-0.5 dB	1.1 dB	1.6 dB
C_3	0.495	4096	0.2 dB	1.7 dB	1.5 dB
C_4	0.660	1024	1.0 dB	2.9 dB	1.9 dB
C_5	0.793	4096	2.0 dB	2.9 dB	0.9 dB

TABLE II

Code performance for BPSK

same performance as the LLR metric is employed for QAM. Comparisons of performance results for alternative metrics are given in [6]–[10]. Each packet represents 4096 binary code symbols, so there are two code words per packet for code C_1 and four code words per packet for C_4. The rate and block length are listed for each code in Table II, and the capacity limit for each rate is included for comparison.

For the version of our protocol described and evaluated herein, the modulation format and symbol rate are selected at the beginning of each session and they are not changed during the session. After the frequency assignment is made, the source determines the set of admissible modulation formats, from which it selects the one that requires the least power. Once the preferred completion time, bandwidth, and modulation format have been selected, the source determines the symbol rate that will be used throughout the session.

Although the modulation format and symbol rate are not changed during the session, the protocol adjusts the transmitter power and code rate in response to information provided by the destination. We assume the radios employ half-duplex transmission, which implies that feedback information from the destination cannot be received during the transmission of a packet. To minimize the time period in which the channel is not available for packet transmission, the amount of feedback information is limited to a few bits in each acknowledgment packet. A high-rate code gives a short session completion time if the packet error rate is not too large; however, if the code rate is too high for the channel conditions, then the need to retransmit several packets will result in a long completion time. Although the packet error rate is not known to the adaptive transmission system, statistics from the receiver permit our protocol to adjust the code rate to achieve high throughput when channel conditions are good but avoid excessive retransmissions when channel conditions deteriorate.

IV. Bandwidth and Session Duration

Let η_j be the number of signature sequence chips per modulation chip for modulation format $\mathcal{M}_j$. For all results in this paper, $\eta_j = 1$, but signal designs with $\eta_j > 1$ are of interest for channels with interference (e.g., [7] and [8]). The number of binary symbols represented by one modulation symbol is denoted by m_j. For example, $m_j = 1$ if $\mathcal{M}_j$ is BPSK and $m_j = \log_2(M)$ if $\mathcal{M}_j$ is an M-ary modulation format such as M-QAM, M-ary orthogonal modulation, or M-ary biorthogonal modulation. The number of modulation chips per symbol is denoted by L_j. For some modulation formats, there is a relationship between m_j and L_j; for example, $m_j = \log_2(L_j)$ for M-ary orthogonal modulation and $m_j = 1 + \log_2(L_j)$ for M-ary biorthogonal modulation.

Let n_b denote the number of binary code symbols in each packet, and let k_i denote the number of information bits per packet when code C_i is used. The rate of code C_i is $r_i = k_i/n_b$. For our numerical results, $n_b = 4096$. If the code is C_i and the modulation format is $\mathcal{M}_j$, then the number of chips per information bit is

$$\lambda_{i,j} = \frac{\eta_j L_j}{m_j r_i} = \frac{\eta_j L_j n_b}{m_j k_i}, \tag{1}$$

which is referred to as the *bandwidth expansion factor*. For BPSK with no error-control coding, $\eta_j = L_j = m_j = r_i = 1$; therefore, $\lambda_{i,j} = 1$ for uncoded BPSK. For modulation format $\mathcal{M}_j$ and code C_i, the parameter $\lambda_{i,j}$ represents the bandwidth expansion relative to uncoded BPSK.

The energy and bandwidth requirements are illustrated in Figure 1 for bit-interleaved coded modulation (BICM) with the modulation formats listed in Table I and the binary turbo product code of rate 0.793. The measure of bandwidth in Figure 1 is the bandwidth expansion factor $\lambda_{i,j}$. The null-to-null bandwidth for modulation format $\mathcal{M}_j$, code C_i, and information transmission rate R_b b/s is $2\lambda_{i,j}R_b$ Hz. For example, the bandwidth expansion factor for BPSK with the code of rate 0.793 is approximately 1.26, and the corresponding bandwidth is approximately $2.52R_b$ Hz. The points connected by the dashed line represent the capacities for the modulation formats with BICM using a code of rate 0.793.

If the number of information bits to be delivered in the session is N_b, then the number of packets that must be delivered is $N_i = \lceil N_b/k_i \rceil \approx N_b/k_i$ when code C_i is used. For example, for the codes in Table II, the number of packets per session ranges from approximately 2500 to approximately 8500 for transmission of a 1 MB file. If the packet error probability is $P_{i,j}$ when the modulation format is $\mathcal{M}_j$ and the code is C_i, then the expected number of transmissions (including retransmissions) required to complete the session is $L_{i,j} = N_i/Q_{i,j}$, where $Q_{i,j} = 1 - P_{i,j}$ is the packet success probability. The

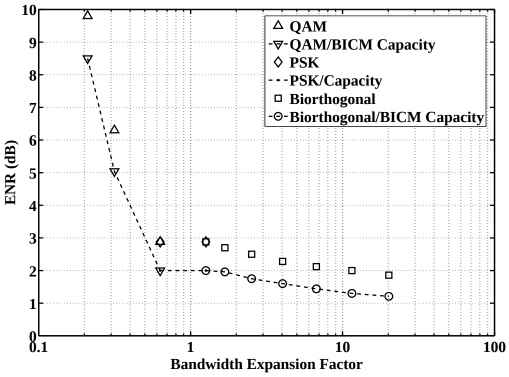

Figure. 1. Relative performance of available modulation formats

number of chips per packet payload is $\eta_j L_j n_b / m_j$, and the average number of chips per session is

$$\Lambda_{i,j} = \frac{\eta_j L_j n_b N_i}{m_j(1 - P_{i,j})} \approx \frac{\eta_j L_j n_b N_b}{m_j k_i Q_{i,j}} = \frac{\lambda_{i,j} N_b}{Q_{i,j}}. \quad (2)$$

If the chip waveform is rectangular and the null-to-null bandwidth is $B_{i,j}$ when modulation format $\mathcal{M}_j$ and code C_i are employed, then the chip rate is $B_{i,j}/2$. The average time required to transmit all payloads for a session is therefore $U_{i,j} = 2\Lambda_{i,j}/B_{i,j}$.

If the overhead time for each packet is $\tau_{i,j}$, then the average overhead time for the session is $V_{i,j} = N_i \tau_{i,j}/Q_{i,j}$ and the average time required for a session is

$$\begin{aligned} T_{i,j} &= U_{i,j} + V_{i,j} = \frac{N_i}{Q_{i,j}} \left[\frac{2\eta_j L_j n_b}{m_j B_{i,j}} + \tau_{i,j} \right] \\ &\approx \frac{N_b}{k_i Q_{i,j}} \left[\frac{2\eta_j L_j n_b}{m_j B_{i,j}} + \tau_{i,j} \right]. \end{aligned} \quad (3)$$

If we solve (3) for the bandwidth, we obtain

$$B_{i,j} \approx \frac{2\eta_j L_j n_b N_b}{m_j(k_i Q_{i,j} T_{i,j} - N_b \tau_{i,j})}, \quad (4)$$

which is used to estimate the bandwidth that will be needed for the session, and this forms the basis for the bandwidth request to the spectrum allocation system.

The *expected throughput* per packet transmission for modulation format $\mathcal{M}_j$ and code C_i is $S_{i,j} = k_i Q_{i,j}$. From (3) we see that once the bandwidth and modulation format have been selected, the average transmission time is minimized by using the code that gives the largest expected throughput. If the feedback information from the previous transmission leads to the prediction that $S_{n,j} = \max\{S_{i,j} : 1 \le i \le n_c\}$ for the next transmission, then our protocol selects code C_n for the next packet.

V. Power Adjustment

Because the source typically does not know the propagation loss in the frequency band assigned for the session, we anticipate that a generous margin will be included in the choice of the power level for the first packet of a new session. It is highly desirable for the first packet to be received and acknowledged so that the source obtains reliable feedback information.

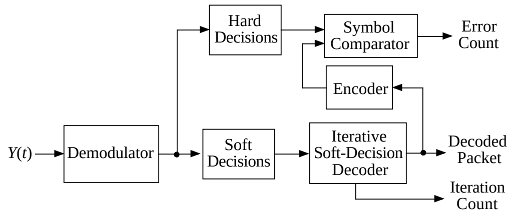

Figure. 2. Statistics for adaptive transmission

If the first packet is transmitted at a power level that is 5–10 dB too low, for example, the destination's receiver may not be able to detect the presence of the packet, synchronize during the preamble, or demodulate the header, let alone demodulate and decode the payload. Although there is a clear advantage in erring on the side of too much rather than too little power for the first few packets, there are significant disadvantages in continued use of more power than necessary, since it causes excessive energy consumption and unnecessary interference for nearby radio receivers.

During the first several packets of the session, our protocol adjusts the power according to feedback information from the destination that consists of two adaptation statistics that are extracted from the demodulator and decoder, as illustrated in Figure 2. The *iteration count for a received word* is the number of iterations performed by the decoder for the word, and the *iteration count for a packet* is the average of the iteration counts for the words in the packet. The *error count* for a packet is an estimate of the number of errors among the binary symbols derived from the hard decisions that are made at the demodulator output. As shown in Figure 2, the error count is obtained by encoding the information symbols at the output of the soft-decision decoder and comparing the resulting code symbols with the demodulated symbols. If there are multiple code words per packet, then the error count for the packet is the sum of the error counts for the received words.

The power-adjustment protocol that uses only the error count q for the previous packet operates as follows: If $q \le q_1$, then the power is decreased by 1 dB; if $q_1 < q \le q_2$, then the power is decreased by 0.5 dB; otherwise, the power is not changed for the next packet. We tested the protocol that uses only the error count ($q_1 = 6$ and $q_2 = 53$) for a target packet error probability of 10^{-2} by simulating 1000 sessions for 64-ary biorthogonal modulation and the code of rate 0.793. Each session begins with a random initial power level, and initial power levels for different sessions are independent. Each initial power level is uniformly distributed on the interval from $P + 10$ dB to $P + 11$ dB, where P is the power (in dB) that gives a packet error probability of 10^{-2}. We found that in 99% of the sessions the power had decreased to less than $P + 1$ dB after 18 packets and in all sessions it had decreased to less than $P + 1$ dB after 25 packets. Similar results were obtained for BPSK.

These results suggest that the power-adjustment phase should terminate after 25 to 30 packets have been transmitted.

VI. Adaptation of the Code Rate

After power adjustment is completed, the adaptation statistics are employed to adapt the code rate to changes in propagation loss. For the version of the protocol described here, only the code rate can be changed in response to a perceived increase in the propagation loss. One reason for decreasing the code rate rather than increasing the power is the desire to avoid interfering with other sessions that might be operating in the same or adjacent frequency bands. As discussed in Section IV, the goal is to choose the code that will give the maximum throughput. For our performance evaluations, we define the *average throughput* as the total number of information bits in packets that are decoded correctly at the destination divided by the total number of packet transmissions that are made. For an information bit to be counted in the numerator of the throughput expression, the entire packet must be decoded correctly. However, each attempt to send a packet from the source to the destination adds to the denominator, whether the packet is decoded correctly or not. Consequently, the adaptive coding system is penalized for unsuccessful transmissions. If the propagation loss is constant for the session, then the average throughput converges to the expected throughput as the number of packets per session increases.

For each transmission, the source desires to use the highest-rate code that will provide a sufficiently small probability of packet error. However, the source does not know the propagation loss. If the code rate is too high, then there will be too many packet errors. If the code rate is too low, then too few information bits are delivered in each packet. In either extreme, the average throughput is low for the session. The adaptive protocol selects the code for the next packet transmission according to adaptation statistics derived from the previous packet transmission. If the previous transmission was not decoded correctly, the destination notifies the source to switch to a lower-rate code if possible (i.e., if the code of lowest rate was not used). A high-rate CRC code is employed to verify correctness.

If the most recent packet was decoded correctly, then the adaptation statistics are compared with adaptation parameters to determine whether to change the code rate. The error count uses adaptation parameters $\nu_0, \nu_1, \ldots, \nu_{n_c}$, in decreasing order. The extreme values are $\nu_0 = n_b$ and $\nu_{n_c} = 0$. If the error count is q and $\nu_{i-1} \geq q > \nu_i$, then rate r_i is chosen for the next transmission. Rate r_{n_c} is chosen if $q = \nu_{n_c}$. For the iteration count, each code rate r_i is associated with adaptation parameters μ_i and λ_i. We define $\mu_1 = \infty$ and $\lambda_{n_c} = 0$. If the iteration count is j and the current

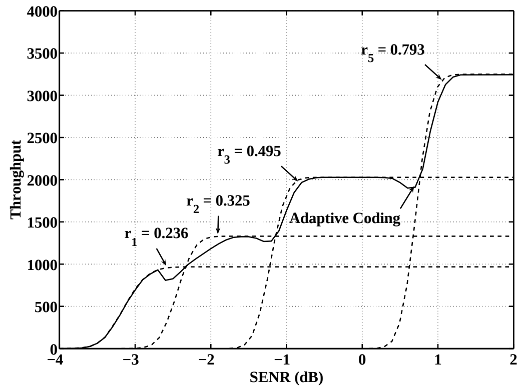

Figure. 3. Performance of the adaptive protocol with the error count for 64-ary biorthogonal modulation

code rate is r_i, then rate r_{i-1} is selected for the next transmission if $j > \mu_i$, rate r_{i+1} is selected if $j < \lambda_i$, and rate r_i is selected if $\lambda_i \leq j \leq \mu_i$.

Define $\mathrm{SENR} = 10\log_{10}(E_s/N_0)$, where E_s is the energy per binary code symbol. When code C_i is used, $E_s = r_i E_b$, which is equivalent to $\mathrm{SENR} = \mathrm{ENR} + 10\log_{10}(r_i)$. We first demonstrate the steady-state performance of the adaptive protocol on a channel with a fixed but unknown propagation loss; that is, SENR is constant but unknown to both the source and destination. Results for 64-ary biorthogonal modulation are given in Figure 3. The dashed lines are the results for fixed-rate coding with each of the four codes C_1, C_2, C_3, and C_5. For biorthogonal modulation, very little is gained by including code C_4, so we set $\nu_3 = \nu_4$ in the code selection algorithm. The solid line represents the adaptive protocol that is based on the error count, with adaptation parameters $\nu_1 = 696$, $\nu_2 = 500$, and $\nu_3 = \nu_4 = 180$. The performance of an ideal protocol that has perfect knowledge of the value of SENR is the upper envelope of the four dashed curves in Figure 3. Our adaptive protocol, which is given no information about the value of SENR, performs almost as well as an ideal protocol with perfect information.

Results for 64-QAM are given in Figure 4. The dashed lines are the results for fixed-rate coding. The solid line represents the performance of the adaptive protocol that uses the error count with adaptation parameters $\nu_1 = 737$, $\nu_2 = 508$, $\nu_3 = 284$, and $\nu_4 = 184$. The upper envelope of the four dashed curves in Figure 4 represents the performance of an ideal protocol that is told the exact value of SENR. As with biorthogonal modulation, the performance of the adaptive protocol that uses only the error count is almost as good as the performance of the ideal protocol that has perfect information. For the performance curves in Figure 5, the iteration count is employed in the adaptive protocol for BPSK modulation. The adaptation parameters are $\lambda_1 = 5$, $\lambda_2 = \lambda_3 = \lambda_4 = 4$, $\mu_2 = 21$, $\mu_3 = 20$, $\mu_4 = 14$, and $\mu_5 = 13$. Again, the performance of the adaptive protocol is nearly as good as the performance of an ideal protocol.

Code-rate adaptation is initiated after the power-

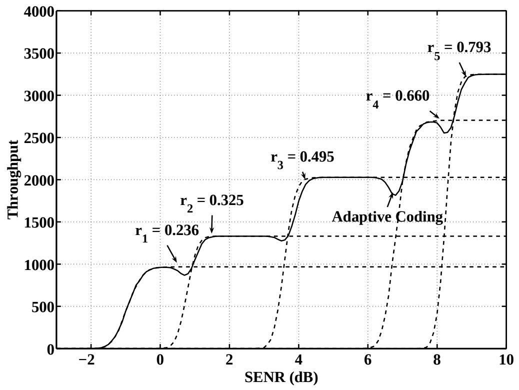

Figure. 4. Performance of the adaptive protocol with the error count for 64-QAM

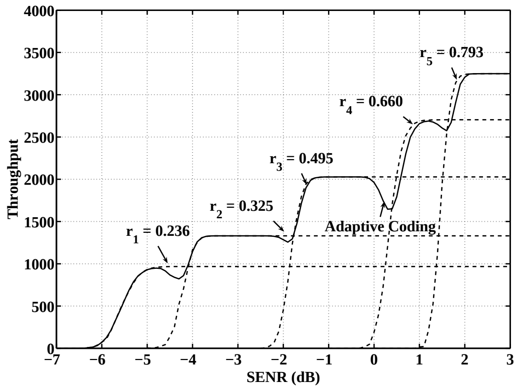

Figure. 5. Performance of the adaptive protocol with the iteration count for BPSK

adjustment phase is completed, so the transmitter power is fixed throughout the code-rate adaptation phase. As a result, there is a one-to-one relationship between the propagation loss and SENR. In the Markov models that we use for the channel, it is convenient to identify the states with values of SENR rather than values of the propagation loss.

We tested our protocol on dynamic channels for which the propagation loss varies according to a four-state Markov chain, and the results are compared in Table III with those obtained from the analytical methods described in [10] for an ideal protocol with perfect previous-state information. All our channel models use the same Markov chain structure. The four states represent four values of SENR that are equally spaced. Transitions are allowed from any state to each state that represents the next lower or next higher value of SENR and these transition probabilities are 0.03. Thus, when the chain is in a state corresponding to the highest or lowest value of SENR, the probability of staying in the same state for the next packet is 0.97, whereas this probability is 0.94 for the states that represent intermediate values of SENR. The results for three modulation formats and four sets of values for SENR are given in Table III, as are the results based on the analytical methods of [10] in which the protocol is given perfect previous-state information and it selects the code that maximizes

Mod.	SENR	Protocol	Ideal	Capacity
$\mathcal{M}_1$	$-3, -2, -1, 0$	1378.4	1442	1664.3
$\mathcal{M}_1$	$-2.5, -1.5, 0.5, 1.5$	2074	2122.4	2614.1
$\mathcal{M}_5$	$-5, -2, 1, 4$	1946	2036.3	2328
$\mathcal{M}_8$	$0, 3, 6, 9$	1827.5	1890.9	2328

TABLE III

Throughput results for four-state Markov channel models.

the conditional expected throughput for the given previous state. The last column in Table III gives the results that are obtained if the codes are replaced by capacity-achieving codes of the same rates and the ideal protocol is used. This column represents the capacity bound on the achievable throughput for any protocol and any codes of the same rates.

VII. Conclusions

The results of our initial investigations indicate that adaptive transmission is highly beneficial for opportunistic use of spectrum. The adaptive transmission protocol makes the power adjustments that are necessitated by the uncertainty in the propagation loss for the assigned frequency band. The protocol also adapts the code rate as the propagation loss changes throughout the session. For the same codes, the performance of our protocol is nearly as good as the performance of an ideal protocol with perfect channel-state information. Furthermore, the performance of our protocol with the codes listed in Table II is often very close to the capacity bound for any codes of the same rates.

References

[1] Advanced Hardware Architectures, Inc., Product Specification for AHA4501 Astro 36 Mbits/sec Turbo Product Code Encoder/Decoder. Available: http://www.aha.com

[2] G. Caire, G. Taricco, and E. Biglieri, "Bit-interleaved coded modulation," *IEEE Transactions on Information Theory*, vol. 44, no. 3, pp. 927–946, May 1998.

[3] FCC Spectrum Policy Task Force Report, ET Docket No. 02-135, U.S. Federal Communications Commission, Nov. 2002.

[4] S. Haykin, "Cognitive radio: Brain-empowered wireless communications," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 2, pp. 201–220, Feb. 2005.

[5] J. Mitola III, "Cognitive radio: Making software radios more personal." *IEEE Personal Communications*, vol. 6, no. 4, pp. 13–18, Aug. 1999.

[6] W. G. Phoel, J. A. Pursley, M. B. Pursley, and J. S. Skinner, "Frequency-hop spread spectrum with quadrature amplitude modulation and error-control coding," *Proceedings of the 2004 IEEE Military Communications Conference*, Nov. 2004.

[7] M. B. Pursley and T. C. Royster, IV, "Coding alternatives for high-rate direct-sequence spread spectrum," *Proceedings of the 2004 IEEE Military Communications Conference*, Nov. 2004.

[8] M. B. Pursley and T. C. Royster, IV, "Adaptive coding for high-rate direct-sequence spread spectrum," *Proceedings of the 2005 IEEE International Conference on Wireless Networks, Communications, and Mobile Computing*, June 2005.

[9] M. B. Pursley and J. S. Skinner, "Decoding strategies for turbo product codes in frequency-hop wireless communications," *Proceedings of the 2003 IEEE International Conference on Communications* (Anchorage, AK), vol. 4, pp. 2963–2968, May 2003.

[10] M. B. Pursley and J. S. Skinner, "Adaptive coding for mobile ad hoc networks with frequency-hop transmission and soft-decision decoding," *Proceedings of the 2005 IEEE International Conference on Wireless Networks, Communications, and Mobile Computing*, June 2005.