Optimal Transmission Scheduling over a Fading Channel with Energy and Deadline Constraints

Alvin Fu, Eytan Modiano, and John N. Tsitsiklis

***Abstract*— We seek to maximize the average data throughput of a single transmitter sending data over a fading channel to a single user class. The transmitter has a fixed amount of energy and a limited amount of time to send data. Given that the channel state determines the throughput obtained per unit of energy expended, the goal is to obtain a policy for scheduling transmissions that maximizes the expected data throughput. We develop a dynamic programming formulation that leads to an optimal transmission schedule, first where the present channel state is known just before transmission, and then to the case where the current channel state is unknown before transmission, but observed after transmission and evolves according to a Markov process. We then extend our approach to the problem of minimizing the expected energy required to send a fixed amount of data over a fading channel given deadline constraints.**

***Index Terms*— Resource allocation, fading channel, wireless networks, dynamic programming, scheduling.**

I. Introduction

FOR MANY wireless transmitters, increased efficiency in data transmission provides significant benefits. Most such devices are battery powered, and often the energy required for data transmission is a significant drain on the battery. Higher energy efficiency may result in the use of a smaller battery or in a longer battery lifetime. Alternatively, increasing data throughput leads to more efficient bandwidth utilization and higher revenue.

The requirements for optimizing performance are frequently contradictory and must be balanced. For example, increasing transmission rates often result in decreased energy efficiency. A well-designed mobile transmitter must not only maximize data throughput, but also optimize the use of resources, effectively cope with a fading channel, and meet operational constraints. These constraints may include a limit on available energy, and a deadline by which transmission must be completed. In this paper, we attempt to address these issues for a lone transmitter sending data to a single user class.

A successful solution to this class of problems would be useful in a wide variety of applications. For instance, a battery-powered laptop computer or cellular phone might want to upload a file to the internet using the minimum amount of energy possible while still meeting network timeout constraints. A communications satellite, remote sensor, or deep space probe might want to maximize data transfer in the face of a specific time window to transmit data and a limited amount of available energy. Many other wireless devices have similar or more stringent energy, power, and time limitations.

The tradeoff between expended energy and throughput is of prime importance in increasing transmitter efficiency. This relationship will depend on the fade state of the channel being used by the transmitter. For a given fade state, the data throughput is usually concave in expended energy (and the expended energy is convex in throughput). This concavity property results from a number of factors. First, the channel capacity is a concave function of power. It is a well-known result of information theory that the capacity C of a Gaussian channel, with noise spectral density $N/2$ watts, power P watts, and bandwidth W hertz, is given by the Shannon capacity equation

$$C = W \log_2 \left(1 + \frac{P}{NW}\right) \tag{1}$$

where C is given in bits per second [1]. The maximum possible information throughput in a given amount of time is hence a logarithmic, concave function of the energy expended. Moreover, channel capacity is an approximately linear (and concave) function of energy in a low signal-to-noise ratio or high bandwidth environment. Second, under a fixed modulation scheme, throughput usually has a locally linear relationship to expended energy. If, in addition, a power limit is imposed - a maximum on the amount of energy that can be consumed at any time - then this linear relation becomes piecewise linear and concave.

Resource allocation for fading channels is a popular topic in information theory. However, the analysis is most often performed for multiple user classes, the resource being allocated is usually average power or bandwidth, and the quantity to be maximized is most often Shannon capacity. Goldsmith and Li [2] [3] and Tse and Hanly [4] found capacity limits and optimal resource allocation policies for such channels. Biglieri et al. [5] examined power allocation schemes for the block-fading Gaussian channel. Tse and Hanly [6] also studied channel allocations in multi-access fading channels that minimize power consumption. Of particular relevance is the paper of Goldsmith and Varaiya [7], which computed the expected Shannon capacity for fading channels, under the condition that both the receiver and transmitter know the current channel state. This work was extended by Negi and Cioffi [15], who calculated capacity and provided power allocation strategies under an additional delay constraint and assuming that a Gaussian codebook is used. None of these papers,

Manuscript received January 16, 2004; revised December 17, 2004; accepted February 18, 2005. The associate editor coordinating the review of this paper and approving it for publication was A. Conti.

The authors are with the Laboratory for Information and Decision Systems at MIT. This research was supported by DARPA under the Next Generation Internet Initiative and by the ARO under grant DAAD10-00-1-0466 (e-mail: {modiano, jnt}@mit.edu).

Digital Object Identifier 10.1109/TWC.2006.03020.

1536-1276/06$20.00 © 2006 IEEE

however, explore the effects of scheduling transmissions with a finite amount of available energy.

The problem of energy efficient scheduling over a fading channel has received much attention recently. Ferracioli et al. [8] propose a scheduling scheme for the third generation cellular air interface standard that takes channel state into account and seeks to balance service priority and energy efficiency. Wong et al. [9] study channel allocation algorithms for cellular base stations. Given known channel characteristics, the authors seek to assign channels in such a way as to minimize total power consumed by all the mobile users communicating with a base station. A recent paper by Tsybakov examines the problem of transmission time minimization between cellular base stations and mobile users [10], but does not consider situations with limited energy. Berry and Gallager [11] analyze schemes that trade off transmission power and buffer delay, while Zhang and Wasserman [12] use dynamic programming and partially observable Markov models to study the tradeoff between throughput and energy efficiency under incomplete channel state information. Many other publications on similar topics can be found in the literature [13] [14] [15] [16] [17].

Perhaps the work closest to this paper is that of El Gamal et al. [18] and Collins and Cruz [19]. Both papers study the problem of minimizing expended energy for a transmitter with a buffer accepting packets arriving according to a random process. El Gamal et al. postulate a hard deadline constraint for all data packets, and increased energy efficiency with slower transmission rates. The problem is to choose transmission rates for each data packet that would allow transmission after arrival and before the deadline, while minimizing expended energy. The paper does not include the effects of a fading channel. Collins and Cruz use dynamic programming and a duality argument to develop a near-optimal transmission policy for minimizing energy in a fading channel with an average delay constraint and a power limit. They assume energy expenditure that is linear with transmitted data and only two possible channel fade states.

In this paper, we show that dynamic programming can be used to generate optimal solutions to the dual problems of maximizing expected throughput given limited energy, and of minimizing expected energy given minimum throughput constraints. We solve both problems for a single energy-limited transmitter serving a single user class in the presence of a fading channel and hard deadline constraints.

For each problem, we first describe a method of finding an optimal solution for the problem under the unrealistic assumption that channel states are known for all time. The procedure is then generalized to the realistic case where only the current channel state and probability distributions for future channel states are known, and further extended to the case where the current channel state is unknown, but the channel state evolves according to a Markov process. We provide tractable numerical methods for the general case where data throughput is concave in expended energy, and closed form optimal policies for special cases. Finally, a number of numerical examples and simulations are provided.

II. Throughput Maximization

A. System Model

We consider a single transmitter operating over a fading channel, sending information to a single user or user type. Time is assumed to be discrete, and in each time slot the channel state changes according to a known probabilistic model. The channel state determines the throughput that can be obtained per unit energy expended by the transmitter, and is modeled as a random process. The transmitter is also assumed to have a battery with a fixed amount of energy units available for use. The objective is to find a transmission schedule that maximizes expected throughput, subject to a constraint on the total energy that can be expended, and a deadline by which the energy must be consumed (or otherwise wasted).

Let a_k be the available energy in the battery at time slot k. The battery starts with a_1 units of energy and must complete transmission by time slot n. The energy consumed at time slot k is denoted by c_k. Thus, the available energy a_k evolves according to $a_{k+1} = a_k - c_k$, with $c_k \leq a_k$ for all k.

The throughput obtained by consuming energy depends on the channel fade state. Let q_k be the channel quality at time k, and let $f(c,q)$ be the throughput obtained by consuming c units of energy in the presence of channel quality q. The function $f(c,q)$ is assumed concave and non-decreasing in c (for example, it may be linear in c).

The objective is to maximize the expected data throughput achieved by the transmitter given n time slots to transmit data and a_1 units of initial energy. Thus, the problem is to maximize

$$E\left[\sum_{k=1}^{n} f(c_k, q_k)\right] \tag{2}$$

subject to the constraints that $c_k \geq 0$ for all k and

$$\sum_{k=1}^{n} c_k \leq a_1 \tag{3}$$

In the following subsections, we first study throughput maximization under the conditions that the channel fade state q_k is known ahead of time and the throughput/energy tradeoff $f(c,q)$ is a concave function. Next, we assume that q_k is random with known distribution function $p_{q_k}(q)$ (independent across time), and that q_k is not revealed until just before transmission at time k. We develop a dynamic programming algorithm that provides an optimal policy for the case where $f(c,q)$ is concave, and obtain a closed-form optimal policy for the special case where $f(c,q)$ is linear, but subject to a power limit. Finally, we extend our results to the the case where q_k follows a Markov process and is only observed at the end of time slot k.

B. Known Channel Quality

Let us start by examining the throughput maximization problem in the simple case where channel quality q_k is known at time k for all $k \geq 0$. Although knowing the channel fade state for all time is an unrealistic assumption, the solution to this problem provides insight, and is used to solve the problem when the channel fade state is unknown. Since the tradeoff between throughput and energy is precisely known for each time slot, we may define $f_k(c) = f(c, q_k)$. Then expression (2) can be restated as maximizing

$$\sum_{k=1}^{n} f_k(c_k) \tag{4}$$

subject to the same constraints as before, and where each function $f_k(c)$ is concave and known. Furthermore, since it cannot hurt to use up all available energy, note that the constraint given in (3) is active and met with equality.

Assuming that each $f_k(c)$ is differentiable, we may apply the Kuhn-Tucker optimality conditions. It is well known that when the objective function is concave and the constraints linear, any solution satisfying the Kuhn-Tucker conditions is optimal. The optimality conditions are the following: for all k,

$$f_k'(c_k) - \lambda - \mu_k = 0$$

$$\mu_k c_k = 0, \;\; \mu_k \geq 0, \;\; c_k \geq 0, \;\; \sum_{k=1}^{n} c_k = a_1$$

where $f_k'(c)$ is the derivative of $f_k(c)$, and where λ and the μ_k are Lagrange multipliers. The last two conditions are simply the constraints of the maximization. In addition, complementary slackness holds; that is, either $\mu_k = 0$ or $c_k = 0$. From this we conclude that for every k, an optimal solution has either $f_k'(c) = \lambda$ or $c_k = 0$.

Given that each $f_k(c)$ is concave, this solution has an interpretation similar to that of waterfilling in the parallel Gaussian channel. In the waterfilling process, one allocates energy to the least noisy Gaussian channel until the marginal return is lower than that of the next best channel, at which point energy is allocated evenly. Here, we allocate energy to the best time slot until the marginal throughput (determined by $f_k'(c)$) is reduced to that of the next best time slot, at which point energy is allocated in such a fashion as to keep marginal throughput identical for both time slots, and so forth.

Note that a similar waterfilling-type solution applies more generally, as long as $f_k(c)$ is concave, so that the differentiability assumption is in fact not necessary. One only needs to work with one-sided derivatives, which are guaranteed to exist under the concavity assumption.

C. Unknown Channel Quality

Now, let us examine the problem of throughput maximization under the assumptions that the channel quality q_k is not known until just before transmission at time k, and that q_k is random and independent across time, with a known distribution function $p_{q_k}(q)$, which may be different at different times k.

This is a much more realistic scenario than the one of subsection II-B. This model can be applied in situations where the channel coherence time is on the order of the duration of the time slot. For instance, studies have shown that satellite channels are stable enough over relatively short periods of time (less than an hour) such that attenuation is well-described by a stationary log-normal distribution [20]. Furthermore, the spectral power of signal level variation is strongly biased toward the region of 0.01 to 0.1 Hz in clear weather, and the variation is much slower in rain [21]. Given a time slot on the order of several seconds in length, it is hence realistic to assume that the channel throughput q_k per unit energy can be measured before a significant change in its value, and before transmission takes place.

Under these assumptions, the dynamic programming algorithm can be used to find an optimal policy. As usual in dynamic programming, we introduce the value function $J_k(a, q)$, which provides a measure of the desirability of the transmitter having energy level a at time k, given that the current channel quality is q. The functions $J_k(a, q)$ for each stage k are related by the dynamic programming recursion:

$$J_n(a, q) = f(a, q)$$

$$J_k(a, q) = \max_{0 \leq c \leq a} \left[f(c, q) + \overline{J}_{k+1}(a - c) \right] \tag{5}$$

where $\overline{J}_k(a) = E[J_k(a, q_k)]$. The first term in the right hand side of equation (5), $f(c, q)$, represents the data throughput that can be obtained in the current stage by consuming c units of energy. The available energy in the next stage is then $a - c$, and the term $\overline{J}_{k+1}(a - c)$ represents the expected throughput that can be obtained in the future given $a - c$ units of energy.

We claim that $J_k(a, q)$ and $\overline{J}_k(a)$ are concave functions of a for all k and q. Indeed, $J_n(a, q) = f(a, q)$ is concave by assumption. If $J_{k+1}(a, q)$ is concave, it is clear that $\overline{J}_{k+1}(a) = E[J_{k+1}(a, q_{k+1})]$ is also concave, since it is a weighted sum of concave functions. Finally, $J_k(a)$, as given by equation (5), is a supremal convolution of two concave functions. Using the fact that the infimal convolution of two convex functions is convex [22], it follows that the supremal convolution of two concave functions is concave.

We now observe that the maximization in equation (5) is of the same form as the problem of allocating energy between two channels of known quality. To obtain an optimal policy for the unknown channel problem of this subsection, we solve a two-stage known channel problem for each possible value of a_k, at each stage of the dynamic programming recursion. This is a computationally tractable problem and can be readily solved numerically. A specific example is given in section IV.

D. Special Case: Piecewise Linear $f(c, q)$

We now assume that throughput is a piecewise linear function of expended energy, of the form $f(c, q) = q \min(c, P)$, where P represents a limit on energy expended per time slot, which we denote as the power limit. Such a model is not unreasonable; we have already seen that a linear energy-throughput relationship arises in a low signal-to-noise ratio or high bandwidth environment, while the power limit might naturally arise as a result of physical limitations in the transmitter hardware or government regulation.

Substituting into (5), the dynamic programming recursion becomes

$$J_k(a, q) = \max_{0 \leq c \leq a} \left[q \min(c, P) + \overline{J}_{k+1}(a - c) \right] \tag{6}$$

$$J_n(a, q) = q \min(a, P)$$

Here it is possible to precisely identify an optimal policy and obtain a closed-form formula for the value function.

Theorem 1: The expected value function $\overline{J}_k(a)$, for $1 \leq k \leq n$, is piecewise linear, with the form

$$\begin{aligned}\overline{J}_k(a) =& \gamma_k^k \min(a,P) \\ &+ \gamma_k^{k+1}[\min(a,2P) - \min(a,P)] \\ &+ \gamma_k^{k+2}[\min(a,3P) - \min(a,2P)] \\ &\vdots \\ &+ \gamma_k^{n}[\min(a,(n-k+1)P) - \min(a,(n-k)P)]\end{aligned} \tag{7}$$

where the number of linear segments is equal to $(n-k+1)$ and where $\gamma_k^k, \ldots, \gamma_k^n$ are constants that give the slopes of each segment, and are determined recursively. The base case is

$$\gamma_n^n = E[q_n]$$

and in the recursion $\gamma_k^k, \gamma_k^{k+1}, \ldots, \gamma_k^n$ are calculated from $\gamma_{k+1}^{k+1}, \ldots, \gamma_{k+1}^n$ for $k < n$. The constants γ_k^k and γ_k^n are given by

$$\gamma_k^k = E[\max(q_k, \gamma_{k+1}^{k+1})] \text{ and } \gamma_k^n = E[\min(q_k, \gamma_{k+1}^{n})]$$

and $\gamma_k^{k+1}, \ldots, \gamma_k^{n-1}$ are given by

$$\gamma_k^i = E[\min(q_k, \gamma_{k+1}^i) - \min(q_k, \gamma_{k+1}^{i+1})] + \gamma_{k+1}^{i+1}$$

Corollary: An optimal policy for $1 \leq k < n$ is to set the consumption c_k to:

$$\begin{array}{ll}\min(P, a_k) & \text{for } \gamma_{k+1}^{k+1} < q_k \\ \min(P, \max(a_k - P, 0)) & \text{for } \gamma_{k+1}^{k+2} < q_k \leq \gamma_{k+1}^{k+1} \\ \vdots & \\ \min(P, \max(a_k - (n-k)P, 0)) & \text{for } q_k \leq \gamma_{k+1}^{n}\end{array} \tag{8}$$

and to set $c_k = \min(a_k, P)$ when $k = n$.

Proof:

We first show that $\overline{J}_k(a_k)$ satisfies equation (7). From the base case of the dynamic programming recursion, we have

$$\begin{aligned}\overline{J}_n(a) =& E_{q_n}[J_n(a,q_n)] \\ =& E_{q_n}[q_n]\min(a,P) \\ =& \gamma_n^n \min(a,P),\end{aligned}$$

which establishes equation (7) for the case $k = n$.

We now assume that $\overline{J}_{k+1}(a)$ satisfies equation (7), and show that $\overline{J}_k(a)$ has the same property. Repeating equation (6), we have

$$J_k(a,q) = \max_{0 \leq c \leq a} \{q \min(c,P) + \overline{J}_{k+1}(a-c)\}.$$

Substituting equation (7) into equation (6), we obtain (9), shown at the top of the next page.

Using the expression in (9), one may employ an algebraic approach to prove the theorem. It can be shown that the choice of

$$c_k = \begin{cases} 0 & \text{if } \ a_k \leq \phi_k(q_k) \\ \min(a_k - \phi_k(q_k), P) & \text{if } \ \phi_k(q_k) < a_k \end{cases} \tag{10}$$

attains the maximum in the right hand side of equation (6). Here, $\phi_k(q_k)$ is a value of u that maximizes the expression

$$q_k(a_k - u) + \overline{J}_{k+1}(u)$$

over all $u \geq 0$, i.e.,

$$\phi_k(q_k) = \arg\max_{u \geq 0} \left[\overline{J}_{k+1}(u) - q_k u\right].$$

It is possible to show that $\phi_k(q_k)$ can be taken to be an integer multiple of P, substitute the value of c_k specified in equation (10), take the expectation over q_k, and then demonstrate that $\overline{J}_k(a)$ has the proper form. However, because this approach is somewhat tedious, we discuss an alternative method. The results from section II-C indicate that the maximizing value of consumption c in equation (6) can be obtained by solving a two-stage known channel problem. One "channel" represents the throughput that can be obtained by consuming immediately, $q\min(c,P)$, while the other channel represents the expected throughput obtained by saving, $\overline{J}_{k+1}(a-c)$.

In this special case, the two channels have a special structure: they are both piecewise linear. We may take advantage of this property when applying the waterfilling solution (as outlined in Section II-B). The derivatives of both $\overline{J}_{k+1}(a-c)$ and $q\min(c,P)$ are decreasing piecewise constant functions whose values change every P units. Allocating energy to the function with the highest marginal throughput simply consists of picking the function with the highest slope. The resulting $J_k(a,q)$ is again piecewise linear and can be determined precisely since its slopes are known.

△

The optimal policy can be explained as follows: Assume a_k units of energy are available at time k. At each time slot at most P units of energy may be consumed. If q_k were known for all k, maximizing throughput would consist of selecting the $\lceil \frac{a_k}{P} \rceil$ time slots with the best channel quality and allocating energy to the best time slots. Assuming there are enough time slots available, this would entail consuming P units of energy in $\lfloor \frac{a_k}{P} \rfloor$ time slots and $a_k - P\lfloor \frac{a_k}{P} \rfloor$, which is the remaining energy, in another time slot.

Of course, channel quality is in fact unknown. However, the constants γ_k^i are representative of expected channel qualities during future time slots as seen just before time k. The γ_k^i values are ordered: γ_k^k is the expected value of the best channel and γ_k^n is the expected value of the worst, in the sense that

$$\gamma_k^k = \max_{\tau} E\left[q_\tau\right], \quad \text{and} \quad \gamma_k^n = \min_{\tau} E\left[q_\tau\right]$$

where the optimization is over all non-anticipative stopping times, i.e. stopping policies, that satisfy $k \leq \tau \leq n$. If we assume that the ordered list $\gamma_{k+1}^{k+1}, \ldots, \gamma_{k+1}^n$ comprises the actual future channel fade states, sorted in order of quality, we may derive an optimal policy from the earlier case with known channel quality. The policy would be as follows: Take

$$\begin{aligned} J_k(a, q) = \max_{0 \le c \le a} \{ & q \min(c, P) + \gamma_{k+1}^{k+1} \min(a - c, P) \\ & + \gamma_{k+1}^{k+2} [\min(a - c, 2P) - \min(a - c, P)] \\ & \vdots \\ & + \gamma_{k+1}^{n} [\min(a - c, (n - k)P) - \min(a - c, (n - k - 1)P)] \} \end{aligned} \quad (9)$$

the current channel state q_k, insert it into the ordered list. If q_k is among the best $\lfloor \frac{a_k}{P} \rfloor$ channel qualities, consume P units of energy. If this is not the case and q_k is the $\lceil \frac{a_k}{P} \rceil^{th}$ best channel quality, consume $a_k - P\lfloor \frac{a_k}{P} \rfloor$ units. Otherwise, do not consume any energy. Theorem 1 and its corollary state that this policy is in fact optimal; the assumption that the constants $\gamma_k^{k+1}, \ldots, \gamma_k^n$ are the actual future channel qualities is unnecessary.

E. Additional Problem Variations

The approach we have developed for the throughput maximization problem can be used to solve several other variants of the main problem. One important situation is the case where the channel state for the current time slot cannot be analyzed quickly enough for it to be known at the transmitter before transmission.

Under this circumstance, if the channel state for each time slot is independent of the channel state of previous time slots, the problem of throughput maximization becomes trivial. However, if channel states are dependent and information about the channel state in previous time slots is available, a Markov model becomes attractive.

The simplest of these is a Markov chain where the probability distribution for the state of the channel in the current time step is dependent only on the state of the channel in the previous time step. The model is more general than the one examined in section II-C, because time slots can be chosen to be shorter in length than the coherence time of the channel, and the independence assumption used in the previous model is a special case of the Markov chain model. Such models can be found in the literature; Wu and Negi have used a Rayleigh-Rician Markov chain to model a flat-fading cellular communications system [24].

The earlier results can be extended to this more general case where channel dependency is modeled as a Markov chain. The objective is again to maximize the quantity

$$E \left[\sum_{k=1}^{n} f(c_k, q_k) \right]$$

subject to the constraints that $c_k \geq 0$ for all k and

$$\sum_{k=1}^{n} c_k \leq a_1$$

The value function satisfies

$$\begin{aligned} J_k(a, q) = \max_{0 \le c \le a} \{ & E_{q_k} [f(c, q_k)) | q_{k-1} = q] \\ & + E_{q_k} [J_{k+1}(a - c, q_k) | q_{k-1} = q] \} \end{aligned} \quad (11)$$

and at the last stage, stage n, the value function is

$$J_n(a, q) = E_{q_n} [f(a, q_n) | q_{n-1} = q].$$

The value function $J_k(a, q_{k-1})$ is concave in a for any fixed q_{k-1}. This can be shown by induction. By assumption, $f(c, q)$ is concave. Its expectation is again concave since the weighted sum of concave functions is concave. Then both $J_n(a, q_n)$ and its expectation, $E_{q_n}[J_n(a - c, q_n) | q_{n-1}]$ are concave. Now assume that $J_{k+1}(a, q_k)$ is concave. Then its expectation $E_{q_k}[J_{k+1}(a - c, q_k) | q_{k-1}]$ is concave as well. Furthermore, $E_{q_k}[f(c, q_k)) | q_{k-1}]$ is another expectation of a concave function, so this term is concave as well. Finally, $J_k(a, q_{k-1})$ is a supremal convolution of two concave functions and so must be concave.

We have seen from the above discussion that both terms on the right hand side of equation (11) are concave. An optimal policy can thus be obtained using our earlier techniques. More precisely, in the Markovian model, the expectation $E_{q_k}[q_k | q_{k-1}]$ and probability distribution function $p_{q_k}(q_k | q_{k-1})$ take the place of q_k and $p_{q_k}(q_k)$ in the case of independent channels. Once this substitution is made, we obtain a problem of the type analyzed in Section II-C.

In the special case where the energy/throughput tradeoff $f(c, q)$ is piecewise linear and of the form $f(c, q) = q \min(c, P)$, it is also possible to obtain an optimal policy in closed form. The expected value function $\overline{J}_k(a, q)$ has the same form as equation (7) for q fixed, and an optimal policy in the form of expression (8) can be found.

More specifically, the expected value function $\overline{J}_k(a)$ and each γ_k^i in the value function can be computed in the following fashion: Let us suppose that the channel quality takes discrete states $q_k \in \{s_1, \ldots, s_m\}$, where each possible state s_i is a distinct discrete value of q_k. We also assume that the transition probability matrix A of size $m \times m$ governs state transitions, where row i of the matrix gives the probability of transitioning from state $q_k = s_i$ to state $q_{k+1} = \{s_1, \ldots, s_m\}$.

Next, we recursively define matrix B_k as a matrix of size $m \times (n - k + 1)$ at each stage $k \in \{1, \ldots, n\}$. At stage n, B_n is an $m \times 1$ vector and is given by

$$B_n = As$$

where the column vector s has elements $s_1, \ldots, s_m$. The matrix B_{k-1} is computed from B_k by

$$B_{k-1} = \mathcal{H} \{ A [B_k | s] \}$$

where $\mathcal{H}\{\cdot\}$ is an operator that sorts the elements of each row of its operand matrix in decreasing order from left to right, and $[B_k | s]$ is a matrix formed by appending the column vector s to B_k.

Given $q_{k-1} = s_i$, the values of $\gamma_k^k, \ldots, \gamma_k^n$ that determine $J_k(a, q_{k-1})$ are simply the $n-k+1$ elements of the ith row of B_k. That is, at time k, for j in the range $k \leq j \leq n$,

$$\gamma_k^j = b_{i,j-k+1}$$

where $b_{l,m}$ is the element of B in the lth row and mth column.

Once the value function is calculated, obtaining the optimal policy is straightfoward; it is the same as expression (8), except that q_k is replaced by its expected value $E_{q_k}[q_k|q_{k-1}]$. The proof that the value function takes the form stated and that the policy is optimal is similar to that of Theorem 1, and is omitted for brevity.

In addition to the above extension to a Markov model, there are also a wide variety of other problem variations. For instance, one might allow the transmitter to receive additional energy input at each stage, and to have a battery of finite size. These formulations are a straightforward extension of the previous result and can be found in [23].

III. Energy Minimization

A. System Model

We have thus far analyzed a situation where we have a given amount of energy, and wish to maximize the expected throughput within a fixed time period. These results can be extended to the case where the transmitter has a given amount of data that must be sent within a fixed time period of length n, and wishes to minimize the expected amount of energy required to do so. Let the variable d_k be the number of data units remaining to be sent at time k, and let s_k be the amount of data that is actually sent at time k. Thus d_k evolves according to $d_{k+1} = d_k - s_k$. The channel quality at time k is given by a variable q_k, which is random. Transmitting s_k units of data requires $g(s_k, q_k)$ units of energy, and the function $g(s_k, q_k)$ is assumed to be convex and differentiable in s_k. Since the transmission must be completed by time n, the objective is to find a transmission policy that minimizes the expected energy

$$E\left[\sum_{k=1}^{n} g(s_k, q_k)\right] \tag{12}$$

subject to the constraints that $s_k \geq 0$ for all k and $\sum_{k=1}^{n} s_k \geq d_1$.

We show that the energy minimization problem, in the presence of a convex energy/throughput function $g(s,q)$, can be solved using methods similar to those used for the throughput maximization problem. We first examine energy minimization for the case where the channel quality q_k is known at time $k = 0$ for all k. We assume throughout that the random variables q_k are independent, with known probability distribution $p_{q_k}(q)$. Next, we study the case where q_k is revealed to the transmitter just before transmission at time k. We present a dynamic programming algorithm that can be used to obtain an optimal policy. Furthermore, when $g(s,q)$ is linear and subject to a power limit, and q_k only takes values which are integer multiples of a minimum channel quality q_{min}, we are able to describe an optimal policy in closed form.

B. Known Channel Quality

We first examine the energy minimization problem in the simple case where the channel quality q_k is completely known ahead of time. This problem is analogous to the known channel throughput maximization problem, and the solution is similar. Since the channel quality is known, the tradeoff between throughput and energy is known for all time. Then we may define $g_k(s) = g(s, q_k)$. The objective is then to solve the problem

$$\min \sum_{k=1}^{n} g_k(s_k)$$

subject to the constraints that $s_k \geq 0$ for all k and $\sum_{k=1}^{n} s_k \geq d_1$. Applying the Kuhn-Tucker optimality conditions, we see that for every k, the optimal solution has either $g_k'(s_k) = \lambda$ or $s_k = 0$, where λ is a constant and $g_k'(s_k)$ is the derivative of $g_k(s_k)$. This solution has a waterfilling interpretation: it is optimal to send data during the best time slot until the marginal energy cost (determined by $g_k'(\cdot)$) is increased to that of the next best time slot, at which point data is allocated in such a fashion as to keep marginal energy costs identical for both time slots, and so forth.

C. Unknown Channel Quality

We now assume that the channel quality q_k is not known until just before transmission at time k. This problem is similar to that of section II-C, and as before, we may use dynamic programming to solve it. The value functions $J_k(d,q)$ for each stage k are related by the following recursion:

$$J_k(d,q) = \min_{0 \leq s \leq d} \left[g(s,q) + \overline{J}_{k+1}(d-s)\right] \tag{13}$$

where the base case is given by $J_n(d,q) = g(s,q)$ and the expected value function $\overline{J}_k(d)$ is defined by $\overline{J}_k(d) = E[J_k(d, q_k)]$. It can be shown that since $g_k(s,q)$ is convex in s, $J_k(d,q)$ and $\overline{J}_k(d)$ are also convex in d. This property implies that the problem reduces to a series of two-stage known channel problems. These problems are computationally tractable and can be solved to obtain an optimal policy.

D. Special Case: Linear $g(s,q)$

We now examine the special case where $g(s,q)$ is linear in s/q, so that q is proportional to the amount of data transmitted per unit energy consumed. A linear function $g(s,q)$ implies that there is no limit on the amount of data that can be sent or on the energy that can be consumed in a single time step. In such a situation, the problem reduces to an optimal stopping problem. However, if we impose a power limit, the problem becomes more difficult.

The power limit effectively imposes a limit of Pq_k on the throughput, where P is the power limit and q_k is the channel quality. If d is the amount of data remaining to be sent, the dynamic programming recursion becomes

$$J_k(d,q) = \min_{0 \leq s \leq \min(d, Pq)} \{\frac{s}{q} + \overline{J}_{k+1}(d-s)\} \tag{14}$$

where $\overline{J}_k(d) = E[J_k(d, q_k)]$. We impose an infinite cost for not sending all the data by the last stage; the terminal cost function is

$$J_{n+1}(d, q) = \begin{cases} 0 & \text{for } d \leq 0 \\ \infty & \text{for } d > 0 \end{cases} \tag{15}$$

For any possible channel quality q, let $\phi_k(q)$ be a value of u that minimizes the expression

$$\frac{d-u}{q} + \overline{J}_{k+1}(u) \tag{16}$$

over all $u \geq 0$. Thus,

$$\phi_k(q) = \arg\min_{u \geq 0} \left[\overline{J}_{k+1}(u) - \frac{u}{q} \right]$$

A value of s that attains the minimum in the right-hand side of equation (14) can be expressed in terms of $\phi_k(q)$, leading to an optimal policy of the following form:

Theorem 2: There exists an optimal policy of the form

$$s_k = \begin{cases} 0 & \text{if} \quad d_k \leq \phi_k(q_k) \\ \min(d_k - \phi_k(q_k), Pq_k) & \text{if} \quad \phi_k(q_k) < d_k \end{cases}$$

Proof:

Equation (14) can be rewritten as

$$J_k(d, q) = \min_{\max(0, d-Pq) \leq u \leq d} \{ \frac{d-u}{q} + \overline{J}_{k+1}(u) \}$$

We know that $\overline{J}_k(u)$ is convex in u. Expression (16) is also convex in u since it is a sum of convex functions. We further notice that the range $u \geq 0$ contains the range $\max(0, d - Pq) \leq u \leq d$.

As a result, an optimizing value of u is simply $\phi_k(q)$ projected on the interval $[\max(0, d - Pq), d]$. The theorem follows. $\triangle$

In effect, $\phi_k(q_k)$ is a threshold beyond which the energy cost of sending data immediately exceeds the cost of saving data for later transmission. It does not depend on the amount of remaining data d_k, and is hence easy to compute. This property allows the development of numerical methods that considerably speed the process of calculating the value function, and which are detailed in [23].

When q_k is discrete and is restricted in value to integer multiples of a constant $q_{\min}$, it is possible to obtain closed form expressions for the optimal policy and value function. It turns out that the expected value function $\overline{J}_k(a)$ is a piecewise linear function with $n-k+1$ segments, each with slope $1/\eta_k^i$, where $1 \leq i \leq n-k+1$, and where η_k^i is defined by the following:

Definition: Given an m-dimensional list $(\alpha_1, \ldots, \alpha_m)$ sorted in ascending order, and an i-dimensional list consisting of i repetitions of the same number x, let $\theta(i, x, \alpha_1, \ldots, \alpha_m)$ be the $(m+1)$ dimensional sorted list obtained by (i) merging and sorting the two lists, and (ii) keeping the largest $m+1$ elements.

Definition: Define the constants η_k^i for $1 \leq k \leq n$ and $1 \leq i \leq n-k+1$ recursively in the following fashion: The base case for $k = n$ (and $i = 1$) is given by

$$\frac{1}{\eta_n^1} = E\left[\frac{1}{q_n}\right]$$

and the recursion to obtain $\eta_{k-1}^1, \ldots, \eta_{k-1}^{n-k+2}$ from $\eta_k^1, \ldots, \eta_k^{n-k+1}$ is given by

$$\left(\frac{1}{\eta_{k-1}^1}, \ldots, \frac{1}{\eta_{k-1}^{n-k+2}}\right) = E\left[\theta\left(\frac{q_k}{q_{\min}}, \frac{1}{q_k}, \frac{1}{\eta_k^1}, \ldots, \frac{1}{\eta_k^{n-k+1}}\right)\right] \tag{17}$$

The slopes $1/\eta_k^1, \ldots, 1/\eta_k^{n-k+1}$ reflect the expected marginal energy cost of sending a data packet. At each time slot, data may be sent immediately for a cost of $1/q_k$ energy units per unit of data. Since there is a power limit P, a maximum of Pq_k units may be sent during each time slot. Alternatively, data may be sent in future stages for an expected cost determined by $\overline{J}_{k+1}(d)$. This function has slope $1/\eta_{k+1}^1$ for the first $Pq_{\min}$ units of data, and $1/\eta_{k+1}^i$ for each i^{th} additional $Pq_{\min}$ units of data. By following the approach of section III-C, the minimum energy cost may be obtained. The resulting value function $J_k(d, q)$ is a piecewise linear function with slopes

$$\theta\left(\frac{q_{k+1}}{q_{\min}}, \frac{1}{q_{k+1}}, \frac{1}{\eta_k^1}, \ldots, \frac{1}{\eta_k^{n-k+1}}\right) \tag{18}$$

for $0 \leq d \leq (n-k+1)Pq_{\min}$. Furthermore, the slopes for the expected value function $\overline{J}_{k-1}(a)$ at time $k-1$ are given by equation (17). The theorem below formalizes these notions.

Theorem 3: Suppose the channel quality q_k is restricted to integer multiples of $q_{\min}$. Then the expected value function is given by

$$\begin{aligned} \overline{J}_k(d) =& \frac{1}{\eta_k^1} \min(d, Pq_{\min}) \\ &+ \frac{1}{\eta_k^2}[\min(d, 2Pq_{\min}) - \min(d, Pq_{\min})] \\ &\vdots \\ &+ \frac{1}{\eta_k^{n-k+1}}[\min(d, (n-k+1)Pq_{\min}) \\ &\quad - \min(d, (n-k)Pq_{\min})] \end{aligned}$$

Corollary: An optimal policy at time k (for $1 \leq k \leq n-1$) is to set s_k as follows:

$$s_k = \begin{cases} \min(d_k, Pq_k), & \text{if } q_k > \eta_{k+1}^1, \\ \min(\max(d_k - Pq_{\min}, 0), Pq_k), & \text{if } \eta_{k+1}^2 < q_k \leq \eta_{k+1}^1, \\ \min(\max(d_k - 2Pq_{\min}, 0), Pq_k), & \text{if } \eta_{k+1}^3 < q_k \leq \eta_{k+1}^2, \end{cases}$$

and so forth until $q_k < \eta_{k+1}^{n-k}$, where $s_k = \min(\max(d_k - (n-k)Pq_{\min}, 0), Pq_k)$.

The proof of the theorem is similar to that of Theorem 1. The major difference arises because of the power limit. In the throughput maximization problem, the limiting resource is energy and the maximum amount of energy that can be consumed during each time step is P. In this energy minimization problem, the constraining resource is data and the maximum amount of data that can be sent at each time step is Pq_k. There is hence a dependence on q_k that is not present in the earlier problem. However, by imposing an integer constraint on the possible values of q_k, we can obtain a closed form expression for the expected value function. Once this is done, the problem is reduced to a two-stage known channel quality problem, and the waterfilling property discussed in section II-B dictates the optimal policy.

As in the case of throughput maximization, there are a number of variations of the energy minimization problem which can be solved using the approach outlined above. For example, our methods can accommodate Markov channel fade states, and also additional incoming data that arrive after time $k = 0$ and have different deadlines [23].

IV. Numerical Examples

A. Throughput Maximization, Nonlinear $f(c, q)$

We first consider a specific instance of the throughput maximization problem where the energy/throughput tradeoff function $f(c, q)$ is nonlinear and derived from the Shannon capacity of a bandlimited Gaussian channel. We first examine a scenario using realistic parameters for a satellite transmitter, and then explore the effects of varying the parameters and using suboptimal heuristics.

Our use of the Shannon capacity function (1) is motivated by the fact that the channel capacity can be nearly achieved in a Gaussian channel using variable rate coding. Accordingly, for this example we define the energy/throughput tradeoff function $f(c, q)$ as

$$f(c, q) = W \log\left(1 + \frac{c\, q}{W}\right)$$

Here, W is again bandwidth, while cq represents the signal to noise ratio. The term c represents power from the transmitter and is controlled by the satellite operator. In contrast, q is an attenuation factor that varies randomly and includes antenna gains, free space losses, atmospheric effects, and noise. It represents the channel quality and is the average signal-to-noise ratio (SNR) when the transmitter is operating at maximum power.

As mentioned earlier, satellite channels can be modeled using a stationary log-normal distribution with time slots of appropriate length. Accordingly, we take each time slot to be ten seconds long. As a result, we can assume that q changes every time slot and follows a log-normal distribution with a fixed mean and variance. Furthermore, given a time slot ten seconds long, it is also realistic to assume that the channel throughput q per unit energy can be measured before a significant change in the channel quality, and before transmission takes place.

A realistic set of numbers for an actual satellite operating in the 5 GHz frequency band might be for it to have a transmitter with 100 watts of maximum power, a bandwidth of 30 MHz, and an average received SNR on the ground of 30 dB at maximum power. 100 watts of maximum power, applied over a ten-second time slot, implies that c can vary between 0 and 1000, and that the limit on expended energy is 1000 watt-seconds per time slot. We assume that the variable q, which represents the SNR, follows a discrete 50-point log-normal distribution with a standard deviation of 1.3 on the underlying Gaussian. We further assume that the satellite has 10500 watt-seconds of energy to expend over 60 time slots (ten minutes).

Fig. 1 shows the energy consumption pattern generated by the optimal policy for a single trajectory of q. The y-axis plots the SNR in dB, and also gives the percentage of the

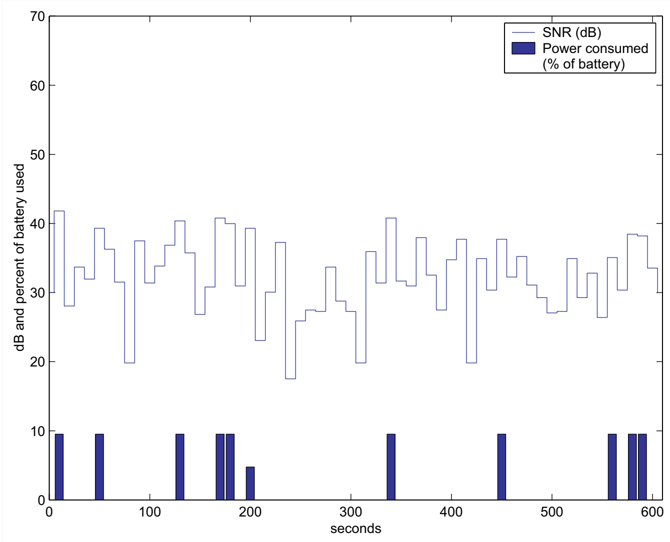

Fig. 1. Channel quality and energy consumption: SNR = 30 dB, W = 30 MHz.

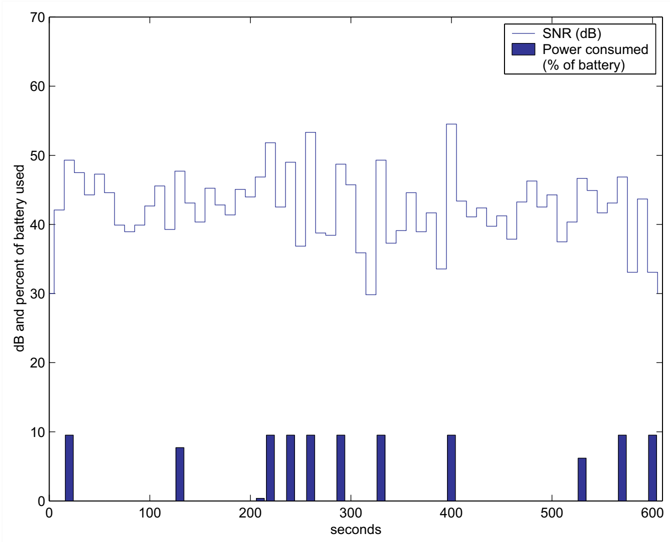

Fig. 2. Channel quality and energy consumption: SNR = 40 dB, W = 1 MHz.

available energy that is consumed during each time step (hence a figure of 10% represents 1050 watt-seconds). Notice that the policy radiates at maximum power over the minimum number of time slots. This is due to the fact that with the parameters given above, the function $f(c, q)$ is very nearly linear. Under such conditions, the throughput is directly proportional to the power, and if it makes sense to radiate one watt of power, it makes sense to radiate at maximum power. There is one time slot where the transmitter does not radiate at maximum power; this is due to the fact that the initial energy in the battery was purposely chosen so as not to be precisely divisible by the power limit. (We further explore the special case of a linear $f(c, q)$ in the next section.)

However, if radiated power increases and bandwidth decreases, $f(c, q)$ becomes increasingly concave, and increasing power rapidly meets with diminishing throughput returns. Under such conditions, the optimal policy is to radiate in more

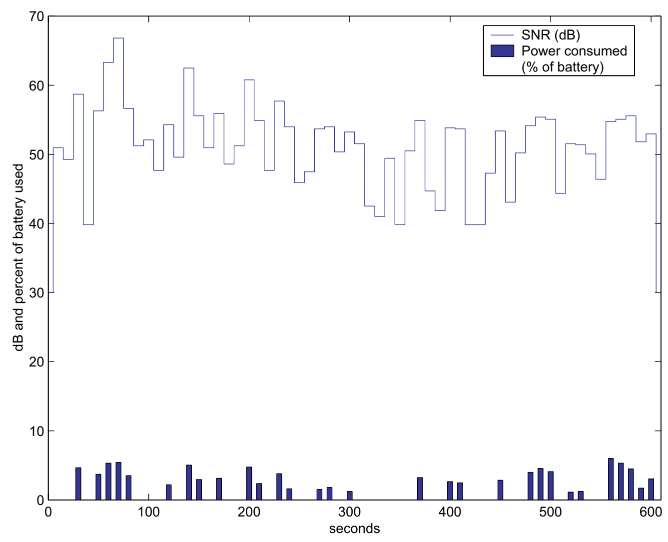

Fig. 3. Channel Quality and Energy Consumption: SNR = 50 dB, W = 100 KHz.

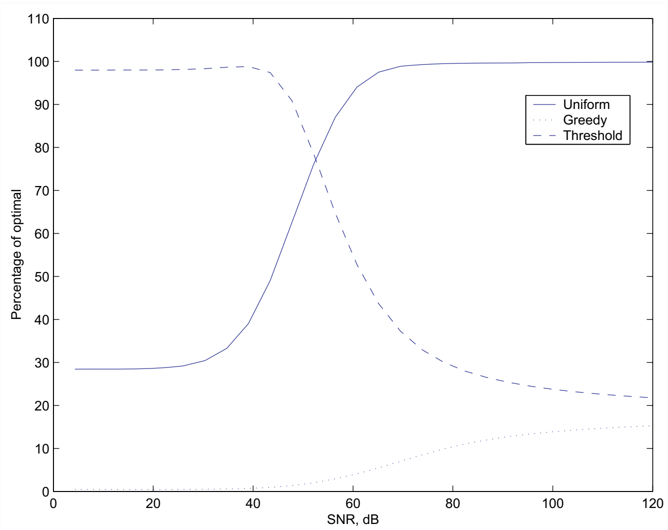

Fig. 4. Performance of heuristics vs SNR.

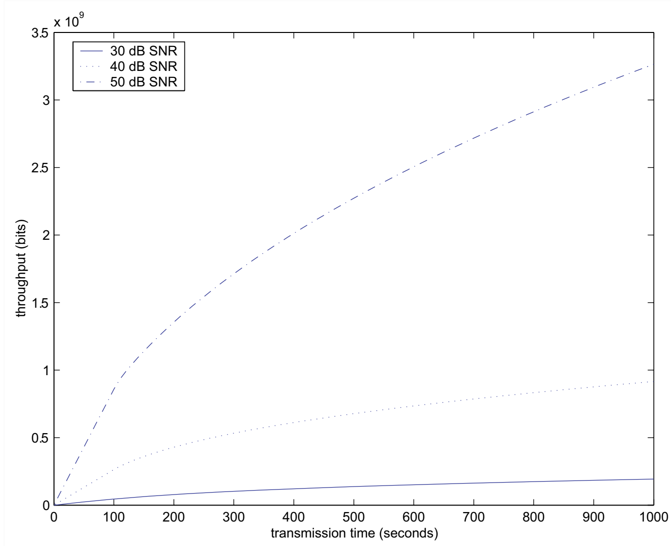

Fig. 5. Average transmitted bits vs transmission time.

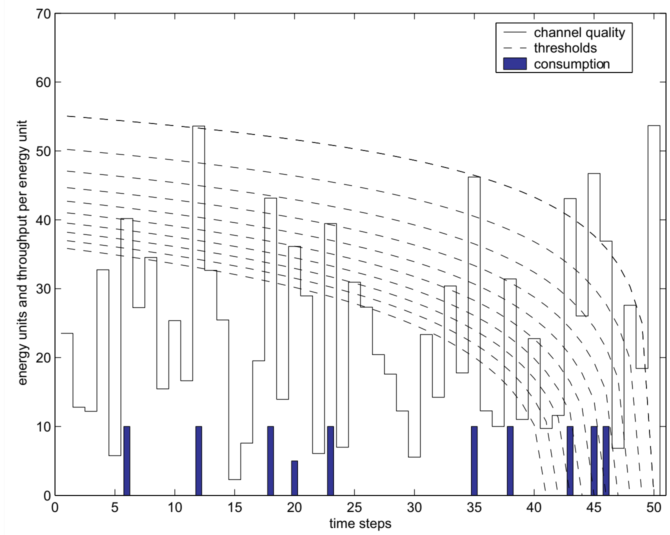

Fig. 6. Channel quality, consumption, and thresholds.

time slots, and with a lesser amount of power. In the limit, the optimal policy is to use an equal amount of energy at every timeslot. To demonstrate this point, Figs. 2 and 3 show the optimal policy in a higher SNR regime where average SNR at maximum power is set at 40 dB and 50 dB, while bandwidth is set at 1 MHz and 100 KHz respectively. All other parameters remain constant. As can be seen from the figures, the optimal algorithm uses more time-slots as the SNR increases.

In Fig. 4, we plot average SNR versus the throughput obtained by three different heuristics as a percentage of the throughput obtained by the optimal policy generated by our dynamic programming recursion. The "uniform" heuristic spreads the available energy uniformly across all available timeslots, while the threshold heuristic transmits (at full power) only if the current channel quality is in the top 20% of possible channel qualities. The greedy heuristic simply transmits at full power until it runs out of energy. Note that the threshold heuristic is more appropriate at low SNRs, while the uniform heuristic is asymptotically optimal as SNR goes to infinity.

Fig. 5 is a graph that shows the effect of increasing transmission time on average throughput for three different average received signal-to-noise ratios. As might be expected, increasing the available time for transmission invariably increases the average throughput.

B. Throughput Maximization, Piecewise Linear $f(c, q)$

In this subsection, we examine the performance of an optimal policy for throughput maximization in the case where $f(c, q)$ is piecewise linear. We compare the performance of an optimal policy to a threshold heuristic that transmits whenever the channel quality is above a fixed threshold. We find that no matter what threshold is used in the heuristic, we are able to obtain superior average performance by using our optimal policy.

The scenario consists of 50 time steps where the channel

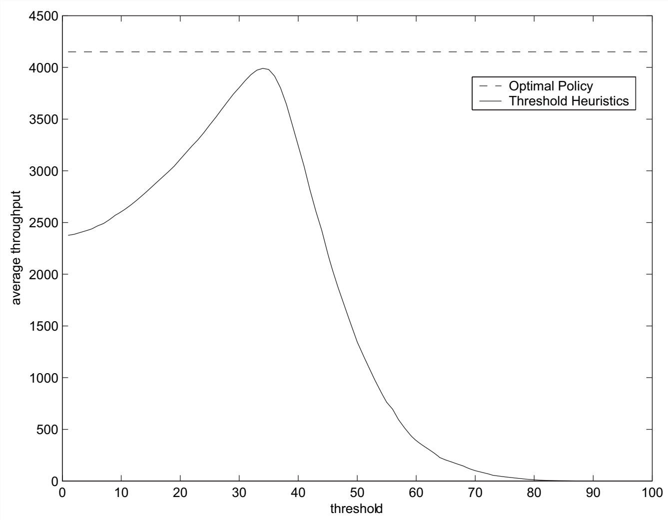

Fig. 7. Throughput for optimal and threshold policies.

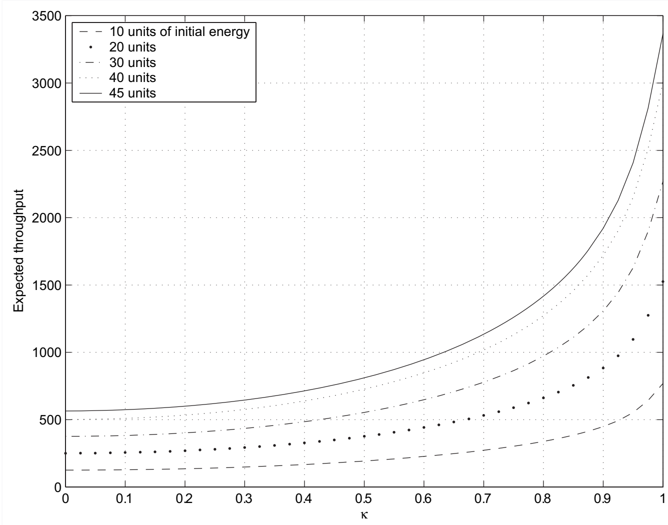

Fig. 8. Expected throughput under an optimal policy for different κ and initial energy.

throughput q_k per unit energy is integer valued and Rayleigh distributed with a mean of 20 during each time step. It is assumed that consuming c units of energy yields $q_k \min(c, P)$ units of throughput, where the power limit P for each time step is 10 units of energy. The initial energy is 95 energy units. Fig. 6 shows a set of randomly generated channel qualities and the consumption schedule as determined by the optimal policy. The figure also shows a set of thresholds corresponding to values of γ_j^i generated by the optimal policy. This allows one to gain an idea of how the optimal policy functions. The topmost dashed line is the value of γ_{k+1}^{k+1} at each time step k. This represents the expected throughput that can be obtained per unit energy for the first ten units of energy saved. The dashed line just below the top is the value of γ_{k+1}^{k+2}. Unsurprisingly, this represents the expected throughput per unit energy for the next ten units of energy saved. The pattern continues for the rest of the dashed lines.

The lines represent thresholds between consuming and saving energy. With the battery full, at energy state 95, the optimal policy consumes energy when channel quality is higher than the bottom-most threshold line. This line represents the expected throughput that can be obtained by the 91st to 100th unit of energy saved. Whenever the current possible throughput is higher than the expected future throughput, the optimal policy consumes.

After the first transmission, the battery only has 85 units of energy. At this point, the threshold line second from the bottom becomes relevant because it represents the expected throughput from the 81st through 90th energy units saved. The optimal policy consumes when the current channel quality is greater than this threshold. Notice also that the optimal policy will only consume five energy units if the current channel quality is greater than this threshold but less than the threshold just above it.

Fig. 7 shows the average throughput obtained by the optimal policy and different fixed threshold policies. The fixed threshold policies always consume as much energy as possible when the channel state is better than or equal to the threshold, and save energy otherwise. The average throughput for each policy was obtained by generating 500 different channel state trajectories and applying the policies to each trajectory. The horizontal dashed line represents the average throughput obtained by the optimal threshold policy, and the solid line plots the throughput obtained by a fixed threshold policy as a function of the threshold. The leftmost point on the curve (threshold = 1) represents a greedy heuristic that transmits no matter what the channel quality, while the rightmost points represent heuristics that transmit only for the very best channel states. As can be seen from the figure, the optimal policy obtained a higher average throughput than any possible simple fixed threshold policy. The advantage of the optimal policy is further enhanced by the fact that finding the best simple threshold is often nontrivial. Moreover, Fig. 7 shows a large sensitivity to error: a poorly chosen threshold will result in a rapid decrease in performance.

C. Throughput Maximization, Markov Channel Model

In this section, we use a version of the Markovian fading channel model detailed in [24]. We assume the energy/transmission tradeoff is piecewise linear and given by $f(c, q) = q \min(c, P)$. The channel quality at time k, q_k, is generated by the model

$$q_k = |\kappa q_{k-1} + v_k|$$

where the noise v_n is zero-mean complex Gaussian with standard deviation σ per unit dimension, and κ is a constant. Thus, q_k follows a Rician distribution with parameters κq_{k-1} and σ. The constant κ is given by $\kappa = .5^{\frac{T_c}{T_s}}$, where T_c is the coherence time of the channel and T_s is the length of the time slot.

For this example, we assume a channel coherence time of 10ms. A timeslot of length 10ms (used by third generation W-CDMA systems) would result in $\kappa = .5$. We also assume a standard deviation of $\sigma = 10$, 50 time slots, an initial channel quality of $q_0 = 20$ and a power limit P of 10 energy units per time slot. Finally, we limit q_k and energy states to integer values.

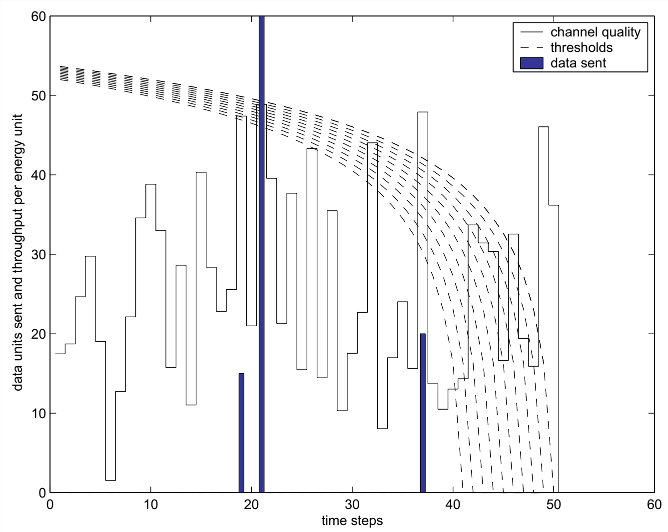

Fig. 9. Channel quality, data sent, and thresholds.

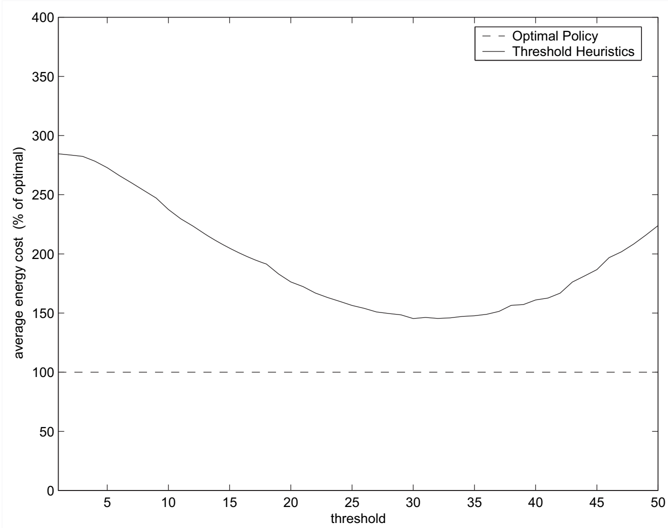

Fig. 10. Energy consumed for optimal and threshold policies.

Fig. 8 shows the expected throughput obtained by the optimal policy developed in section II-E for different values of κ and different initial energies. In the degenerate case when $\kappa = 0$, channel quality for each time slot follows a Rayleigh distribution and is completely independent. The expected throughput is simply the expended energy multiplied by the mean of the Rayleigh distribution, which is $\sigma\sqrt{\frac{\pi}{2}}$. As κ increases, the transmitter is increasingly able to predict future channel states, and can hence obtain increasing throughput.

D. Energy Minimization

We now present a similar example of an energy minimization problem. We consider a scenario where the transmitter has 50 time steps to send 95 units of data. Channel quality q_k is integer and Rayleigh distributed with a mean of 20, and a power limit of 10 energy units is imposed. Sending s units of data requires s/q_k units of energy.

Fig. 9 shows the channel qualities and the data transmission schedule as determined by our optimal policy. The figure also shows a set of thresholds corresponding to values of η_{k+1}^i generated by the optimal policy. The topmost dashed line is the value of η_{k+1}^1 for each time step k. This represents the expected data that can be transmitted per unit energy for the first ten units of energy saved. The pattern continues; the dashed line just below the top is the value of η_{k+1}^2 and represents the expected throughput per unit energy for the next ten units of energy saved. These threshold lines are used in the same fashion as those of Fig. 6.

Unlike the problem of throughput maximization, a policy that uses a fixed threshold at all times would not be appropriate. This is because unless the threshold is below q_{min}, the expected cost would be infinite, as there is a positive probability that the channel quality would be equal to q_{min} at all times. We consider instead a threshold policy of the following type: For times k such that

$$n - k \leq \frac{d_0}{Pq_{min}}$$

(that is, using the parameters in this example, for $k \geq 41$), the threshold is at zero and we always transmit at full power. For earlier times, we transmit if and only if the channel quality is above a threshold.

Fig. 10 shows the average energy consumed by different fixed threshold policies as a percentage of that consumed by the optimal policy. The results were obtained by applying the policies to 500 randomly generated channel state trajectories. The optimal policy obtained a significantly lower energy cost than any possible threshold policy of the type described above.

V. Conclusion

This paper used dynamic programming to develop strategies for transmission optimization over a wireless fading channel with energy, power and deadline constraints. Throughput maximization and energy minimization strategies were developed, first for channels with known fade states, and then for channels with fade states unknown until just before transmission. For the general case, the concave form of the value function was shown, and a method of finding the optimal policy was provided. In addition, closed-form optimal policies were derived for the special case of a piecewise linear energy-throughput relationship. Furthermore, the problem of throughput maximization in the case where fade states are not known before transmission and evolve according to a Markov process was also examined. Again, a general methodology was developed, along with a closed-form optimal policy for the piecewise linear special case.

References

[1] T. Cover and J. Thomas, *Elements of Information Theory*. New York: John Wiley and Sons, Inc., 1991.

[2] L. Li and A. Goldsmith, "Capacity and optimal resource allocation for fading broadcast channels," in *Proc. Thiry-Sixth Annual Allerton Conference on Communication, Control, and Computing*, 1998, pp. 516-525.

[3] A. Goldsmith, "Capacity and dynamic resource allocation in broadcast fading channels," in *Proc. Thirty-Third Annual Allerton Conference on Commun., Control, and Comp.*, 1995, pp. 915-924.

[4] D. Tse and S. Hanly, "Multiaccess fading channels: Polymatroid structure, optimal resource allocation and throughput capacities," *IEEE Trans. Inform. Theory*, vol. 44, no. 7, pp. 2796-2815, Nov. 1998.
[5] E. Biglieri, G. Caire, and G. Taricco, "Coding for the block-fading channel: optimum and suboptimum power-allocation schemes," in *Proc. 1998 Information Theory Workshop*, 1998.
[6] S. Hanly and D. Tse, "Min-max power allocation for successive decoding," in *Proc. 1998 Information Theory Workshop*, 1998 pp. 56-57.
[7] A. Goldsmith and P. Varaiya, "Capacity of fading channels with channel side information," *IEEE Trans. Inform. Theory*, vol. 43, pp. 1986-1992, Nov. 1997.
[8] M. Ferracioli, V. Tralli, and R. Verdone, "Channel based adaptive resource allocation at the MAC layer in UMTS TD-CDMA systems," in *Proc. IEEE VTS Fall VTC2000, 52nd Vehicular Technology Conference*, vol. 2, pp. 2549-2555.
[9] C. Wong, C. Tsui, R. Cheng, and K. Letaief, "A real-time sub-carrier allocation scheme for multiple access downlink OFDM transmission," in *Proc. IEEE VTS 50th Vehicular Technology Conference*, 1999, vol. 2, pp. 1124-1128.
[10] B. Tsybakov, "File transmission over wireless fast fading downlink," *IEEE Trans. Inform. Theory*, vol. 48, no. 8, Aug. 2002.
[11] R. Berry and R. Gallager, "Communication over fading channels with delay constraints," *IEEE Trans. Inform. Theory*, vol. 48, no. 5, pp. 1135-1149, May 2002.
[12] D. Zhang and K. Wasserman, " Transmission schemes for time-varying wireless channels with partial state observations," in *Proc. IEEE Infocom 2002*, vol. 2, pp. 467-476.
[13] M. Filip and E. Vilar, "Optimum utilization of the channel capacity of a satellite link in the presence of amplitude scintillations and rain attenuation," *IEEE Trans. Commun.*, vol. 38, no. 11, Nov. 1990.
[14] I. Koutsopoulos and L. Tassiulas, "A unified framework for handover prediction and resource allocation in non-geostationary mobile satellite networks," in *Proc. IEEE VTS 50th Vehicular Technology Conference*, 1999, vol. 4, pp. 2106-2110.
[15] R. Negi and J. Cioffi, "Transmission over fading channels with channel side information and delay constraint," in *Proc. Globecom 1999*, pp. 2550-53.
[16] E. Papapetrou, I. Gragopoulos, and F. Pavlidou, "Performance evaluation of LEO satellite constellations with inter-satellite links under self-similar and poisson traffic," *International J. Satellite Commun.*, vol. 17, pp. 51-64, 1999.
[17] P. D. Shaft, "Unconstrained allocation of communication satellite traffic," in *Proc. 1977 International Conference on Communications*, part 3, pp. 26-30.
[18] A. El Gamal, E. Uysal, and B. Prabhakar, "Energy-efficient transmission over a wireless link via lazy packet scheduling," in *Proc. Infocom 2001*, vol. 1, pp. 386-394.
[19] B. Collins and R. Cruz, "Transmission policies for time varying channels with average delay constraints," in *Proc. 1999 Allerton Conf. on Commun., Control, and Comp.*, 1999.
[20] Y. Karasawa, K. Yasukawa, and J. Allnutt, "A new prediction method for tropospheric scintillation on Earth-space paths," *IEEE Trans. Antennas Propagat.*, vol. 36, no. 11, pp. 1608-1614, Nov. 1988.
[21] Y. Karasawa, K. Yasukawa, and M. Yamada, "Tropospheric scintillation in the 14/11-GHz bands on Earth-space paths with low elevation angles," *IEEE Trans. Antennas Propagat.*, vol. 36, no. 4, pp. 563-569, Apr. 1988.
[22] R. Rockafellar, *Convex Analysis*. Princeton, NJ: Princeton University Press, 1970.
[23] A. Fu, "Optimal Energy Allocation for Space and Wireless Communications," Ph.D. thesis, MIT, 2003.
[24] D. Wu and R. Negi, "Effective capacity: a wireless link model for support of quality of service," *IEEE Trans. Wireless Commun.*, vol. 2, no. 4, July 2003.

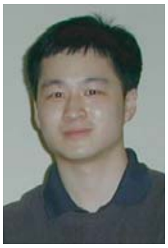

Alvin C. Fu (S '02) received a B.S. Electrical Engineering (1994) and a B.S. in Biology (1994) from Stanford University. He received his S.M. (1996) and Ph.D. (2003) in Electrical Engineering and Computer Science from MIT. His master's thesis was on the topic of music recognition, and his doctoral thesis covered optimal resource allocation for satellite and wireless networks. He is currently employed at Fidelity Investments.

Eytan Modiano received his B.S. degree in Electrical Engineering and Computer Science from the University of Connecticut at Storrs in 1986 and his M.S. and PhD degrees, both in Electrical Engineering, from the University of Maryland, College Park, MD, in 1989 and 1992 respectively. He was a Naval Research Laboratory Fellow between 1987 and 1992 and a National Research Council Post Doctoral Fellow during 1992-1993 while he was conducting research on security and performance issues in distributed network protocols. Between 1993 and 1999 he was with the Communications Division at MIT Lincoln Laboratory where he designed communication protocols for satellite, wireless, and optical networks and was the project leader for MIT Lincoln Laboratory's Next Generation Internet (NGI) project. He joined the MIT faculty in 1999, where he is presently an Associate Professor in the Department of Aeronautics and Astronautics and the Laboratory for Information and Decision Systems (LIDS). His research is on communication networks and protocols with emphasis on satellite, wireless, and optical networks. He is currently an Associate Editor for Communication Networks for IEEE Transactions on Information Theory and for The International Journal of Satellite Communications. He had served as a guest editor for IEEE JSAC special issue on WDM network architectures; the Computer Networks Journal special issue on Broadband Internet Access; the Journal of Communications and Networks special issue on Wireless Ad-Hoc Networks; and for IEEE Journal of Lightwave Technology special issue on Optical Networks. He is the Technical Program co-chair for Wiopt 2006 and co-chair for Infocom 2007.

John N. Tsitsiklis (F'99) received the B.S. degree in mathematics in 1980 and the B.S., M.S., and Ph.D. degrees in electrical engineering in 1980, 1981, and 1984, respectively, all from the Massachusetts Institute of Technology (M.I.T.), Cambridge. He is a Professor of Electrical Engineering and Computer Science and a past Codirector of the Operations Research Center at M.I.T. He has coauthored four books. His research interests are in the fields of systems, optimization, communications, control, and operations research. He is currently a member of the editorial board for the Springer-Verlag Lecture Notes in Control and Information Sciences series, and an Associate Editor of Mathematics of Operations Research.