Optimal Energy Allocation for Delay-Constrained Data Transmission over a Time-Varying Channel

Alvin Fu, Eytan Modiano, and John Tsitsiklis
Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
alfu@mit.edu, modiano@mit.edu, jnt@mit.edu

***Abstract*— We seek to maximize the data throughput of an energy and time constrained transmitter sending data over a fading channel. The transmitter has a fixed amount of energy and a limited amount of time to send data. Given that the channel fade 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 closed-form transmission schedule. We then extend our approach to the problem of minimizing the energy required to send a fixed amount of data over a fading channel given deadline constraints.**

I. Introduction

For many mobile 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.

Unfortunately, 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.

It is not difficult to think of mobile transmitters that face such issues. A laptop computer uploading data to the internet, a personal digital assistant sending email, or a military sensor in a remote location, must all operate with finite battery energy, a limit on transmission power, and a delay constraint. For all these applications, increased throughput and energy efficiency would in the very least result in longer operational life or lowered cost.

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 Shannon capacity for a channel is a logarithmic function of 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 linear relation to energy expended. 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.

With the relationship between throughput and power in mind, one may envision a transmitter trying to send information over a channel whose fade state and throughput/power tradeoff is constantly changing. In this paper, we seek to maximize the throughput of the transmitter when there is only a finite amount of energy available. Conversely, we also seek to minimize the energy required to send a certain amount of data. For both problems, it is assumed that transmission must be completed by a deadline.

Resource allocation for fading multi-user broadcast channels is a popular topic in information theory. However, the resource being allocated is usually average power or bandwidth, and the quantity to be maximized is most often Shannon capacity. Goldsmith and Li [12] [9] and Tse and Hanly [15] found capacity limits and optimal resource allocation policies for such channels. Biglieri et al. [2] examined power allocation schemes for the block-fading Gaussian channel. Tse and Hanly [11] also studied channel allocations in multi-access fading channels that minimize power consumption. Of particular relevance is the paper of Goldsmith and Varaiya [10], 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 [13], 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, however, explore the effects of scheduling transmissions with a finite amount of available energy.

More recently, as interest in mobile communications has increased, transmission scheduling for fading channels has attracted more interest. Ferracioli et al. [5] 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. [17] 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 examined the problem of transmission time minimization between cellular base stations and mobile users [16], but did not consider situations with limited energy. Berry and Gallager [1] analyzed schemes that trade off transmission power and buffer delay, while Zhang and Wasserman [18] used dynamic programming and partially observable Markov models to study the tradeoff between throughput and energy efficiency under incomplete channel state information.

Perhaps the work closest to this paper is that of El Gamal et al. [4] and Collins and Cruz [3]. Both papers studied the problem of minimizing expended energy for a transmitter with a buffer accepting packets arriving according to a random process. El Gamal et al. postulated a hard deadline constraint for all data packets, and increased energy efficiency with slower transmission rates. The problem was to choose transmission rates for each data packet that would allow transmission after arrival and before the deadline, while minimizing expended energy. The paper did not include the effects of a fading channel. Collins and Cruz used 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 assumed energy expenditure that is linear with transmitted data and 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 throughput given limited energy, and of minimizing energy given minimum throughput constraints. We solve both problems in the presence of a fading channel and hard deadline constraints. Furthermore, we provide tractable numerical methods for the general case where data throughput is concave in expended energy, and closed form optimal solutions for special cases.

II. Throughput Maximization

A. System Model

We consider a transmitter operating over a fading channel. 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 it 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

$$a_{k+1} = a_k - c_k$$

At each time slot k the transmitter can consume up to a_k units of energy.

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 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

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

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

$$\sum_{k=1}^{n} c_k \leq a_1 \tag{2}$$

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 general concave function. Then, 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 examine additional variations of the throughput maximization problem. The dynamic programming algorithm is extended to the case where the channel quality q_k is revealed only at time $k+1$, and evolves as a Markov process. The case where the transmitter receives additional incoming energy before the deadline is also discussed.

B. Known Channel Quality

Let us start by examining the throughput maximization problem in the simple case where channel quality is completely known for all time. In other words, q_k is known at time 1 for all k. 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

$$f_k(c) = f(c, q_k)$$

The problem (1) can then be restated as maximizing

$$\sum_{k=1}^{n} f_k(c_k) \tag{3}$$

subject to the same constraints as before, and where every 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 (2) 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,

$$\begin{aligned} f_k'(c_k) - \lambda - \mu_k &= 0 \\ \mu_k c_k &= 0 \\ \mu_k &\geq 0 \\ c_k &\geq 0 \\ \sum_{k=1}^{n} c_k &= a_1 \end{aligned}$$

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 $f_k'(c_k) = \lambda$ 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.

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.

In this case, 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:

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

and

$$J_k(a, q) = \max_{0 \leq c \leq a} \left[f(c, q) + \overline{J}_{k+1}(a - c) \right] \tag{4}$$

where $\overline{J}_k(a) = E[J_k(a, q_k)]$.

The first term in the right hand side of equation (4), $f(c, q)$, 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 (4), is an infimal convolution of two concave functions and is therefore concave [14].

We now observe that the maximization in equation (4) 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.

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

We now assume that throughput is a piecewise linear function of expended energy, and energy consumption at each time step is subject to a power limit. Then

$$f(c, q) = q \min(c, P)$$

where P is the power limit. Substituting into (4), the dynamic programming recursion becomes

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

and

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

It is possible to precisely identify an optimal policy and obtain a closed-form formula for this value function.

Theorem 1:

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

$$\begin{aligned} \overline{J}_k(a) = & \gamma_k^k \min(a, P) \\ & + \gamma_k^{k+1}[\min(a, 2P) - \min(a, P)] \\ & + \gamma_k^{k+2}[\min(a, 3P) - \min(a, 2P)] \\ & \vdots \\ & + \gamma_k^{n}[\min(a, (n-k+1)P) \\ & \quad - \min(a, (n-k)P)] \end{aligned} \tag{6}$$

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 computed according to the following recursion: 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

$$\begin{aligned}\gamma_k^k &= E[\max(q_k, \gamma_{k+1}^{k+1})] \\ \gamma_k^n &= E[\min(q_k, \gamma_{k+1}^n)]\end{aligned}$$

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:

$$\begin{array}{ll} \min(P, a_k) & \text{for } \gamma_{k+1}^{k+1} < q_k \\ \min(P, \max(a_k - P, 0)) & \text{for } \gamma_{k+1}^{k+2} < q_k \leq \gamma_{k+1}^{k+1} \\ \vdots & \\ \min(P, \max(a_k - (n-k)P, 0)) & \text{for } q_k \leq \gamma_{k+1}^n \end{array} \quad (7)$$

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

Proof:

Given in the appendix.

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

$$\begin{aligned}\gamma_k^k &= \max_\tau E[q_\tau] \\ \gamma_k^n &= \min_\tau E[q_\tau]\end{aligned}$$

where the optimization is over all non-anticipative stopping times 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 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. For instance, one may eliminate the assumption that the channel is known just before the decision to transmit. Instead, it can be assumed that the current channel is unknown, but that the channel in the previous stage is known and that the state of the channel in the current time step is dependent on the state of the channel in the previous time step.

This channel dependency can be modeled as a Markov chain and we extend the earlier results to this case. 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

$$\begin{aligned} J_k(a, q) = \max_{0 \leq c \leq a} \{ & E[f(c, q_k))|q_{k-1} = q] \\ & + E[J_{k+1}(a - c, q_k)|q_{k-1} = q]\} \end{aligned} \quad (8)$$

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

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

It can be shown that $J_k(a, q_{k-1})$ is concave in a for any fixed q_{k-1} [6]. Since the weighted sum of a concave function is again concave, it is evident that both terms on the right hand side of equation (8) are also concave. An optimal policy can thus be obtained using our earlier techniques for the case of independent channels. More precisely, in the Markovian model, the expectation $E[q_k|q_{k-1}]$ 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, the results for section II-C easily extend to the Markov case.

Other problem variations may be obtained by allowing the transmitter to receive additional energy input at each stage, and to have a battery of finite size. Unused energy that cannot be stored in the battery is assumed to be lost.

Let b_k be the incoming energy for each stage and let $E_{\max}$ be the battery capacity. Then the expression for available energy at each stage evolves according to

$$a_{k+1} = \min(a_k - c_k, E_{\max}) + b_{k+1}$$

Let us suppose b_k is known at time 1 for all k. Then the dynamic programming equation (4) becomes, for $k < n$,

$$J_k(a,q) = \max_{0\le c\le a} \Big[f(c,q) + \overline{J}_{k+1}(\min(a-c, E_{\max}) + b_{k+1}) \Big] \quad (9)$$

where, again,

$$\overline{J}_k(a) = E[J_k(a,q)]$$

and for $k = n$,

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

Assuming that $\overline{J}_{k+1}(a)$ is concave in a, it can be shown that $\overline{J}_{k+1}(\min(a-c, E_{\max}) + b)$ is also concave in a. Then, we can again reduce the maximization in equation (9) to that of the two-stage known channel problem of equation (4).

When there is no power limit and $f(c,q)$ is linear, i.e.

$$f(c,q) = q\ c$$

the dynamic programming recursion becomes

$$J_k(a,q) = \max_{0\le c\le a} \big[q\ c + \overline{J}_{k+1}(\min(a-c, E_{\max}) + b_{k+1}) \big]$$

and an optimal policy can be obtained in closed form. Furthermore when $f(c,q)$ is of the form

$$f(c,q) = q \min(c, P)$$

fast numerical methods can be used to rapidly obtain an optimal policy [6].

It is also possible to handle the case where the energy input is random with known distribution $p_{b_k}(b)$ and where b_k is revealed to the transmitter just before transmission. In this case, we define

$$\hat{J}_k(a) = E[J_k(a + b_k, q_k)]$$

The dynamic programming equation becomes

$$J_k(a,q) = \max_{0\le c\le a} \Big[f(c,q) + \hat{J}_{k+1}(a-c) \Big] \quad (10)$$

the proof of concavity still goes through, and the maximization is again the same as that of the two-stage known channel problem.

Unfortunately, we are unable to obtain a closed-form formula for $J_k(a,q)$, for the case of a general function $f(c,q)$. Moreover, the numerical evaluation of the value function can become difficult because the expectation over q_k and b_k imposes a heavy computational burden. However, when $f(c,q)$ is linear or piecewise linear, it is possible to apply the numerical techniques outlined in [6] to obtain an optimal policy.

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 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 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

$$E\left[\sum_{k=1}^{n} g(s_k, q_k)\right] \quad (11)$$

subject to the constraints that $s_k \ge 0$ for all k and

$$\sum_{k=1}^{n} s_k \ge d_1$$

where d_1 is the amount of data to be sent.

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. Then, we study the case where q_k is revealed to the transmitter just before transmission at time k. We assume throughout that the random variables q_k are independent, with known probability distribution $p_{q_k}(q_k)$. We present a dynamic programming algorithm that can be used to obtain an optimal policy. Furthermore, when $g(s_k, q_k)$ 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. Finally, several additional variants of the problem are analyzed, including a case where additional data packets with fixed time to live can arrive before the transmission deadline.

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 its 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:

$$J_k(d, q) = \min_{0 \leq s \leq d} \left[g(s, q) + \overline{J}_{k+1}(d - s) \right] \qquad (12)$$

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 can be interpreted as 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

$$J_k(d, q) = \min_{0 \leq s \leq \min(d, Pq)} \{ \frac{s}{q} + \overline{J}_{k+1}(d - s) \} \qquad (13)$$

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

$$J_{n+1}(d, q) = \begin{cases} 0 & \text{for } d \leq 0 \\ \infty & \text{for } d > 0 \end{cases} \qquad (14)$$

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

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

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 (13) 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

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

attains the minimum in the right hand side of (13).

The proof is omitted for brevity but can be found in [6].

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 remaining amount of 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 [6].

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

$$\left(\frac{1}{\eta_{k-1}^1}, \ldots, \frac{1}{\eta_{k-1}^{n-k+2}}\right) = E\left[\theta\left(\frac{q_k}{q_{\min}}, \frac{1}{q_k}, \frac{1}{\eta_k^1}, \ldots, \frac{1}{\eta_k^{n-k+1}}\right)\right] \quad (15)$$

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

$$\theta\left(\frac{q_{k+1}}{q_{\min}}, \frac{1}{q_{k+1}}, \frac{1}{\eta_k^1}, \ldots, \frac{1}{\eta_k^{n-k+1}}\right) \quad (16)$$

for $0 \le d \le (n-k+1)Pq_{\min}$. Furthermore, the slopes for the expected value function $\overline{J}_{k-1}(a)$ at time $k-1$ are given by equation (15).

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 $1 \le k \le n-1$) is to set s_k as follows: For $q_k > \eta_{k+1}^1$,

$$s_k = \min(d_k, Pq_k)$$

for $\eta_{k+1}^2 < q_k \le \eta_{k+1}^1$

$$s_k = \min(\max(d_k - Pq_{\min}, 0), Pq_k)$$

for $\eta_{k+1}^3 < q_k \le \eta_{k+1}^2$

$$s_k = \min(\max(d_k - 2Pq_{\min}, 0), Pq_k)$$

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 given in [6] and is similar to the proof of Theorem 1 for throughput maximization, and is omitted for brevity. 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 Kuhn-Tucker conditions dictate the optimal policy.

E. Additional Problem Variations

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 (which all must be sent by time n) that arrive after time $k = 0$.

One problem variation of interest consists of having a linear $g(s, q)$, but with a limit on throughput, i.e.

$$g(s, q) = \begin{cases} s/q & \text{for } s \le T_{\max} \\ \infty & \text{for } s > T_{\max} \end{cases}$$

where $T_{\max}$ is the throughput limit per time slot. Such a problem may arise when a fixed modulation scheme is employed or when data protocols impose a limitation on throughput (e.g. the TCP window size limit).

A closed-form optimal transmission schedule for this case can be obtained using a procedure analogous to the one provided in Theorem 1. This is because a throughput limit applied to the energy minimization problem has exactly the same effect as a power (energy) limit applied to the throughput maximization problem. (In the same fashion, a throughput limit in the throughput problem may be handled by an analog of Theorem 3.)

We now focus our attention on the case where the transmitter must deal with data that arrives after time $k = 0$, and where data packets must be sent l time steps after arrival. In this situation, each data packet is said to have a time to live, or expiration time, of l time steps.

Consider a transmitter that is sending data and that receives b_k additional units of data at each time k, which must be sent within a time to live l. We assume that b_k is known at time 0 for all k. Data that has not yet arrived cannot be sent, and data

must be sent before it expires. The objective is to minimize the expected energy required to send all the data.

It is clear that the data with the shortest time to live can be transmitted first. This is because sending data with a shorter time to live permits greater flexibility in arranging transmission times. In [6] it is formally shown that the class of first-in-first-out (FIFO) policies contains an optimal policy when all data packets arrive with the same time to live l.

With this property in mind, we may therefore envision the data packets arriving in a queue, each labeled with its own time to live and with the restriction that each packet must be serviced by its expiration time. This expiration time constraint may be transformed into an equivalent constraint where the queue buffer size changes during each time slot, and any packets that do not fit in the buffer must be serviced immediately.

Define d_k as the total amount of data to be transmitted, regardless of time to live. Then d_k evolves according to the recursion

$$d_{k+1} = d_k + b_{k+1} - s_k$$

where s_k is the amount of data transmitted at time k, and the base case is $d_1 = b_1$ We further define the buffer size at time k to be B_k, and we have the following relationship between incoming data and buffer size:

$$B_k = \sum_{i=k-l+1}^{k} b_i$$

This definition of B_k insures that packets arriving at time k are transmitted at time $k + l$ or before.

The objective is to choose s_k so as to minimize

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

subject to the constraints that $d_k - s_k \leq B_k$ and $0 \leq s_k \leq d_k$.

When b_k (and hence B_k) is known for all k, the dynamic programming recursion becomes

$$J_k(d,q) = \min_{0 \leq u \leq \min(d, B_k)} \Big\{ g_k(d-u, q) + E\big[J_{k+1}(u, q_{k+1})\big] \Big\}$$

where u is unsent data. This problem can be solved using the methods of section III-C.

Now consider the case where each b_k is revealed only at time k. Here, the state space for the dynamic programming recursion becomes $l + 2$ dimensional. For all but small values of time to live l, the problem becomes intractable and approximate methods are necessary.

IV. Numerical Examples

A. Throughput Maximization

We consider a specific instance of the throughput maximization problem and 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.

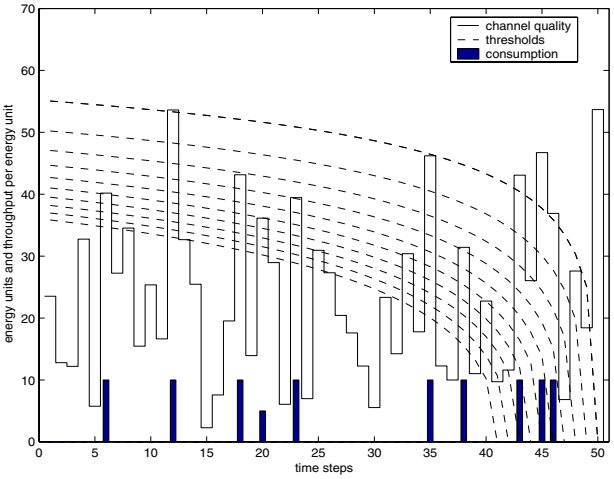

Fig. 1. Channel Quality, Consumption, and Thresholds

The scenario consists of 50 time steps where the channel 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.

Figure 1 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} for 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

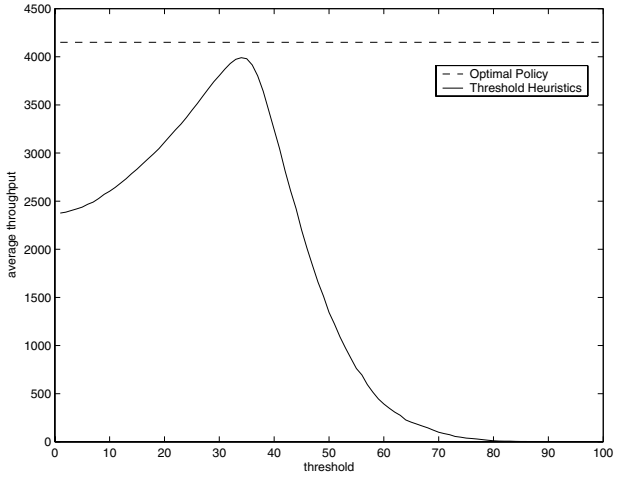

Fig. 2. Throughput for Optimal and Threshold Policies

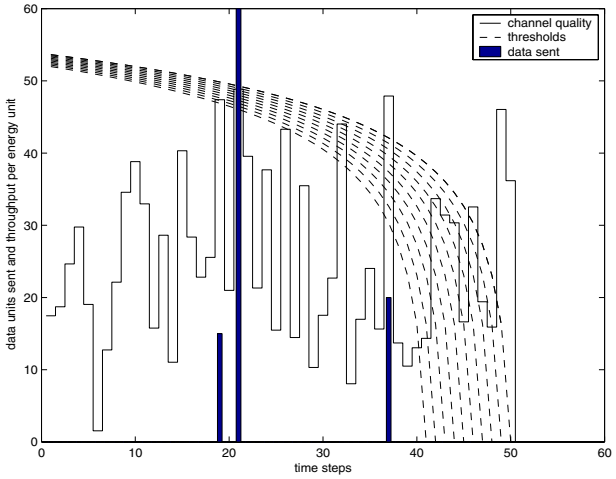

Fig. 3. Channel Quality, Data Sent, and Thresholds

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.

Figure 2 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, Figure 2 shows a large sensitivity to error: a poorly chosen threshold will result in a rapid decrease in performance.

The calculation of the optimal threshold policy described above was compared with an earlier-developed numerical algorithm that directly solves the dynamic programming recursion [8]. The numerical algorithm, although capable of handling a much broader array of situations (such as random power limits and battery recharges) was significantly slower: on a Sun Ultra 10 computer running Matlab 6.0, the numerical algorithm was slower by a factor of 10. Both the optimal policy and the numerical algorithm are orders of magnitude faster than a brute-force calculation of the value function.

B. 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.

Figure 3 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 Figure 1.

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 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

$$n - k \leq \frac{d_0}{P q_{min}}$$

(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.

Figure 4 shows the average energy consumed by different fixed threshold policies as a percentage of that consumed by

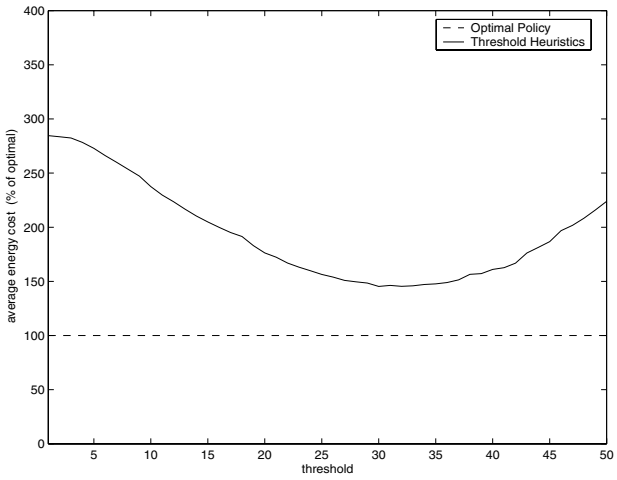

Fig. 4. Energy Consumed for Optimal and Threshold Policies

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 developed strategies for transmission optimization over a fading channel with energy and time 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. Furthermore, closed form optimal policies were derived for a number of cases, and several problem variations were examined. These variations included throughput maximization in the presence of additional incoming energy, energy minimization in the presence of additional incoming data, and a scenario where the channel fade state evolves according to a Markov process.

There are several areas that require further investigation. First, a tractable solution for the important scenario of energy minimization with random additional data inputs needs to be developed. In addition, it may be interesting to explore transmission optimization under the additional constraints imposed by network protocols, such as those that might arise under explicit TCP time-out and window size limits.

Acknowledgement

This research was supported by DARPA under the Next Generation Internet Initiative and by the ARO under grant DAAD10-00-1-0466.

Appendix

Proof of Theorem 1:

We first show that $\overline{J}_k(a_k)$ satisfies equation (6). From the base case of the dynamic programming recursion, we have

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

which establishes equation (6) for the case $k=n$.

We now assume that $\overline{J}_{k+1}(a)$ satisfies equation (6), and show that $\overline{J}_k(a)$ has the same property. Substituting equation (6) into equation (5), we obtain

$$\begin{aligned}J_k(a,q) = \max_{0\le c\le a}\{&q\min(c,P)+\gamma_{k+1}^{k+1}\min(a-c,P)\\ &+\gamma_{k+1}^{k+2}[\min(a-c,2P)-\min(a-c,P)]\\ &\vdots\\ &+\gamma_{k+1}^{n}[\min(a-c,(n-k)P)\\ &\quad-\min(a-c,(n-k-1)P)]\end{aligned} \tag{17}$$

Using the above expression, one may employ an algebraic approach to prove the theorem [6].

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 (5) 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 Kuhn-Tucker conditions (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.

More precisely, the function $q\min(c,P)$ has derivative

$$\begin{array}{ll} q & \text{for} \quad 0\le c<P\\ 0 & \text{for} \quad P>c\end{array} \tag{18}$$

and $\overline{J}_{k+1}(a-c)$ has derivative

$$\begin{array}{ll} \gamma_{k+1}^{k+1} & \text{for} \quad 0\le a-c<P\\ \gamma_{k+1}^{k+2} & \text{for} \quad P<a-c<2P\\ & \vdots\\ \gamma_{k+1}^{n} & \text{for} \quad (n-k-1)P<a-c<(n-k)P\\ 0 & \text{for} \quad (n-k)P>a-c\end{array} \tag{19}$$

To summarize (18) and (19), we say that $q\min(c,P)$ has marginal throughput (slopes) $q,0$, and $\overline{J}_{k+1}(a-c)$ has marginal throughput $\gamma_{k+1}^{k+1},\ldots,\gamma_{k+1}^{n}$. Note that due to concavity, the elements of both marginal throughput expressions are in descending order.

Using the Kuhn-Tucker conditions and applying the waterfilling process, we see that $J_k(a,q)$ has almost the same marginal throughput as $\overline{J}_{k+1}(a-c)$, except that q is inserted into it so as to maintain a descending order. The resulting marginal throughput, $\gamma_{k+1}^{k+1},\ldots,q,\ldots\gamma_{k+1}^{n}$ can be exactly expressed as

$$\begin{aligned}
&\max(q,\gamma_{k+1}^{k+1}),\\
&\min(q,\gamma_{k+1}^{k+1})-\min(q,\gamma_{k+1}^{k+2})+\gamma_{k+1}^{k+2},\\
&\min(q,\gamma_{k+1}^{k+2})-\min(q,\gamma_{k+1}^{k+3})+\gamma_{k+1}^{k+3},\\
&\qquad\vdots\\
&\min(q,\gamma_{k+1}^{n})
\end{aligned}\tag{20}$$

Since $J_k(a,q)$ is piecewise linear and equation (20) provides an exact expression for its slope in consecutive intervals of length P, we may obtain $J_k(a,q)$ in closed form:

$$\begin{aligned}
J_k(a,q) &= \max(q,\gamma_{k+1}^{k+1})\min(a,P)\\
&\quad+[\min(q,\gamma_{k+1}^{k+1})-\min(q,\gamma_{k+1}^{k+2})+\gamma_{k+1}^{k+2}]\\
&\quad\cdot[\min(a,2P)-\min(a,P)]\\
&\qquad\vdots\\
&\quad+\min(q,\gamma_{k+1}^{n})\\
&\quad\cdot[\min(a,(n-k+1)P)-\min(a,(n-k)P)]
\end{aligned}$$

Taking the expected value with respect to q, and using the definition of γ_k^i, we obtain equation (6). The expected value function $\overline{J}_k(a)$ is thus a concave piecewise linear function with marginal throughput $\gamma_k^k,\ldots,\gamma_k^n$.

Now that the form of the expected value function has been ascertained, the optimal policy readily follows. The value function $\overline{J}_{k+1}(a-c)$ is identical to the one that would be obtained if, at time k, future channel qualities $q_{k+1},\ldots,q_n$ were known to be respectively equal to $\gamma_{k+1}^{k+1},\ldots,\gamma_{k+1}^{n}$. As seen in the discussion subsequent to the statement of Theorem 1, an optimal policy for this situation simply consists of picking the best time slots. The resulting consumption at time k is then given by (7).

REFERENCES

[1] R. Berry, R. Gallager, "Communication over fading channels with delay constraints." *IEEE Transactions on Information Theory* vol.48, no.5, pp.1135-1149, May 2002.

[2] E. Biglieri, G. Caire, G. Taricco, "Coding for the block-fading channel: optimum and suboptimum power-allocation schemes," in *1998 Information Theory Workshop*. (New York, NY, USA, 1998).

[3] B. Collins, R. Cruz, "Transmission policies for time varying channels with average delay constraints," in *Proceedings, 1999 Allerton Conf. on Commun., Control, and Comp.*. (Monticello, IL, USA, 1999).

[4] A. El Gamal, E. Uysal, and B. Prabhakar, "Energy-efficient transmission over a wireless link via lazy packet scheduling," *Infocom 2001 Proceedings*. (2001, vol.1, pp.386-94).

[5] M. Ferracioli, V. Tralli, and R. Verdone, "Channel based adaptive resource allocation at the MAC layer in UMTS TD-CDMA systems," in *IEEE VTS Fall VTC2000. 52nd Vehicular Technology Conference*. (2000, vol.2, pp.2549-55).

[6] A. Fu, *Optimial Energy Allocation for Space and Wireless Communications*, Ph.D. thesis, MIT, 2002.

[7] A. Fu, E. Modiano, and J.N. Tsitsiklis, "Optimal energy allocation and admission control for communications satellites," in *Infocom 2002 Proceedings*. (2002, vol.2, pp.648-56).

[8] A. Fu, E. Modiano, and J.N. Tsitsiklis, "Transmission scheduling over a fading channel with energy and deadline constraints," *2002 Conference on Information Sciences and Systems*. (Princeton, NJ, 2002).

[9] A. Goldsmith, "Capacity and dynamic resource allocation in broadcast fading channels," in *Proceedings, Thirty-Third Annual Allerton Conference on Commun., Control, and Comp.*. (Urbana-Champaign, IL, USA, 1995, pp.915-24).

[10] A. Goldsmith, P. Varaiya, "Capacity of fading channels with channel side information." *IEEE Trans. Inform. Theory*, vol.43, pp.1986-92, Nov. 1997.

[11] S. Hanly, D. Tse, "Min-max power allocation for successive decoding", in *1998 Information Theory Workshop*. (New York, NY, USA, 1998 pp.56-7).

[12] L. Li, A. Goldsmith, "Capacity and optimal resource allocation for fading broadcast channels," in *Proceedings, Thiry-Sixth Annual Allerton Conference on Communication, Control, and Computing*. (1998, pp.516-25).

[13] R. Negi, J. Cioffi, "Transmission over fading channels with channel side information and delay constraint," in *Globecom 1999 Proceedings*. (1999, pp. 2550-53).

[14] R. Rockafellar, *Convex Analysis*. Princeton, NJ: Princeton University Press, 1970.

[15] D.N. Tse, S.V. Hanly, "Multiaccess fading channels: Polymatroid structure, optimal resource allocation and throughput capacities." *IEEE Transactions on Information Theory* vol.44, no.7, pp.2796-815, Nov. 1998.

[16] B. Tsybakov, "File Transmission Over Wireless Fast Fading Downlink." *IEEE Transactions on Information Theory*, vol.48, no.8, August 2002.

[17] C. Wong, C. Tsui, R. Cheng, and K. Letaief, "A real-time sub-carrier allocation scheme for multiple access downlink OFDM transmission," in *IEEE VTS 50th Vehicular Technology Conference.* (1999, vol.2, pp.1124-8).

[18] D. Zhang and K. Wasserman, " Transmission schemes for time-varying wireless channels with partial state observations," in *IEEE Infocom 2002*. (2002, vol.2, pp.467-476).