Some dynamic resource allocation problems in wireless networks

Randall Berry

Dept. of ECE, Northwestern University, Evanston, IL

ABSTRACT

We consider dynamic resource allocation problems that arise in wireless networking. Specifically transmission scheduling problems are studied in cases where a user can dynamically allocate communication resources such as transmission rate and power based on current channel knowledge as well as traffic variations. We assume that arriving data is stored in a transmission buffer, and investigate the trade-off between average transmission power and average buffer delay. A general characterization of this trade-off is given and the behavior of this trade-off in the regime of asymptotically large buffer delays is explored. An extension to a more general utility based quality of service definition is also discussed.

Keywords: wireless networks, resource allocation, power control, quality of service

1. INTRODUCTION

Two fundamental aspects of typical mobile wireless communication systems are that the channel is time-varying and the channel must be shared between many geographically separate users. One approach for addressing these issues is to dynamically allocate communication resources, such as transmission power, data rate and bandwidth, based in part on any channel knowledge that is available at the transmitter. Such knowledge can be gained in a variety of ways including measurements of a reverse channel in a duplex system, explicit feedback from the receiver, or measurements of a pilot signal. A variety of techniques can be used for allocating resources including transmission scheduling, power control, variable rate spreading, and adaptive modulation and coding.

The allocation of communication resources influences not only physical layer performance measures such as the bit error rate or signal-to-noise ratio, but can also effect network level Quality of Service (QoS) parameters such as throughput, delay and loss. In this paper, we discuss some dynamic resource allocation problems for wireless networks in which both network layer and physical layer concerns are jointly considered. The resource allocation problems we consider arise in the situation shown in Figure 1; this depicts a single mobile user communicating over a fading channel. Here a higher layer application running on a mobile platform generates data which is placed into a transmission buffer. Periodically, the transmitter removes some of the data from the buffer, encodes the data and transmits the encoded data over a fading channel. Eventually this data will be received, decoded, and passed to a corresponding higher layer application running at the receiver. We consider situations where the transmitter can allocate resources based on both knowledge of the channel state and the buffer occupancy. Our focus here is on single user situations where the main issue is the allocation of resources in response to time varying channel conditions and user requirements. Though we will not address it here, many of these ideas can be extended to multi-user settings, where the allocation of resources between users must also be considered.[1]

In the above situation we consider two conflicting objectives. One objective is the physical layer goal of minimizing the long term average transmission power required to reliably transmit the data. In a wireless network, mobile users often rely on a battery with a limited supply of energy; minimizing the average transmission power leads to a more efficient utilization of battery energy. We are interested here in long term average power consumption rather than short term averages of interest in, say, regulatory constraints. Such short term considerations may be modeled as a constraint on each codeword sent, while the long term average power depends on the sequence of codewords that are sent. The second objective is to minimize a quantity related to the required network layer QoS. Initially we consider the case where this quantity is the average delay incurred by arriving data, then we generalize this to allow for a larger class of QoS specifications. Note there is a clear trade-off between these objectives - transmitting at a higher rate requires more power but reduces the average delay. Our emphasis in this paper is on quantifying the trade-off between these objectives. We give some general characteristics of this trade-off, and we explore the behavior of this trade-off in the regime of asymptoticly large buffer delays. In this regime, we show that the required power converges

email:rberry@ece.nwu.edu, phone: (847) 491-7074, fax: (847) 491-4455

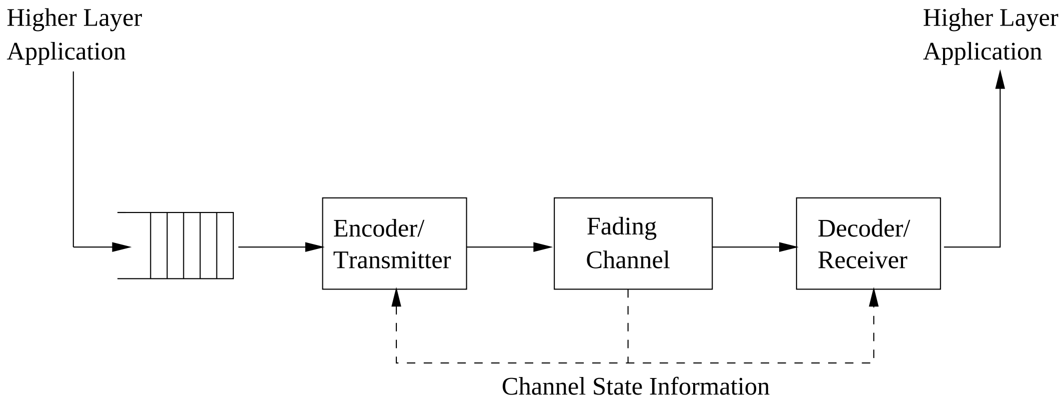

Figure 1. System Model.

to an asymptotic value and give the rate at which convergence occurs. We also give a sequence of policies which attain this convergence rate; these policies depend only weakly on the buffer occupancy and have a relatively simple structure.

The above approach requires a tighter coupling between the physical layer and higher network layers than is found in typical wired network architectures. This is in the same spirit as a number of papers[2–5] where approaches that allow for greater integration across traditional layer boundaries are considered in wireless settings. Several of the characteristics of wireless channels make such approaches appealing; indeed some of the information theoretic results for wireless channels provide some motivation for a tighter integration of layers.[6]

The outline of the remainder of the paper is as follows. In the next section, we describe the model of the system in Figure 1. In Sect. 3, we analyze the trade-off between average power and average buffer delay for these models. To accomplish this we consider a buffer control problem which can be analyzed using ideas from Markov decision theory. We characterize the "optimal power/delay trade-off curve" for such problems. In Sect. 4 we analyze this trade-off in the asymptotic regime of large delay. In this regime, we can characterize simple buffer control strategies which have some optimal properties. Section 5 contains an extension of the results in Sect. 4 to a case where instead of average delay, the user's QoS is expressed in a more general utility framework. Finally, in Sect. 6 some concluding remarks are made.

2. PROBLEM FORMULATION

In this section we describe a simple model for the situation shown in Figure 1. We note many of the following assumption may be relaxed, but we will not consider these generalizations here.[7] We consider a discrete time model where time samples occur every Δ seconds. At the start of each Δ second interval, the transmitter removes some number of bits from the buffer and encodes them. The resulting codeword is then transmitted to the receiver during this interval. We assume that at the end of this interval, the codeword is either successfully received or found to be in error and discarded. In the later case, we assume the bits are returned to the transmission buffer (An alternative to this assumption is discussed in another paper[8]). Let S_n denote the buffer occupancy at the start of the nth time interval and let R_n denote the transmission rate during the nth interval. Thus the number of bits transmitted during the nth interval is $R_n\Delta$ where $R_n\Delta \leq S_n$. Let q_n denote the probability a codeword is not received in the nth interval. The number of bits received during this interval will be a random variable U_n where

$$U_n = \begin{cases} R_n\Delta & \text{with probability } 1 - q_n \\ 0 & \text{with probability } q_n. \end{cases} \tag{1}$$

Denote the number of bits that arrive to the buffer during the nth interval by A_n. The dynamics of the buffer are then are given by:

$$S_{n+1} = S_n + A_n - U_n. \tag{2}$$

For this paper we will assume that $\{A_n\}$ is a sequence of i.i.d. random variables which take values in a compact set $\mathcal{A} \subset \mathbb{R}^+$ and have mean $\mathbb{E}A_1 = \bar{A}$. Let $A_{min} = \inf \mathcal{A}$ and $A_{max} = \sup \mathcal{A}$.

Next the channel model is described. Let H_n be a random variable which denotes the available channel knowledge at the the start of the nth time interval; this will be referred to as the channel state at time n. We assume that given H_n the channel behavior during the nth block is independent of H_m, for all $m < n$ and also independent of all previous transmissions and receptions. For a given transmission rate $R_n = r$ and a given value $H_n = h$, let $P(h, r, q)$ denote the minimum transmission power required so that the probability a codeword is not received is at most q. At times we will denote $P(h, r, q)$ by $P(h, U)$, where U denotes the random variable representing the received rate as in (1). The following analysis will only depend on a few general characteristics of the function $P(h, r, q)$; specifically we assume that this function has the following characteristics: (*i*) it is increasing in r for all h and all q, (*ii*) it is decreasing in q for all r and all h, (*iii*) it is strictly convex in (r, q) for all h, and (*iv*) it is a continuous function of h for all r and all q. In this paper we focus on the case where $\{H_n\}$ is an i.i.d. sequence of random variables, and we assume that these random variable take values in a compact set $\mathcal{H} \subset \mathbb{R}$ and that H_i has a continuous probability density over this set.

The following illustrates one example which satisfies the above assumptions. Suppose the channel is modeled as a additive white Gaussian noise channel with narrow-band Rayleigh fading, and suppose H_n represents an estimate of the expected fading level during the nth time interval. Assume that for a given rate R_n, a packet will be correctly received if the actual realization of the fading results in a channel with a capacity greater than R_n, otherwise an outage occurs.[9] In this case, it can be shown that

$$P(h, r, q) = \frac{N_o W}{-\ln(1-q)h}\left(2^{2r/W} - 1\right), \tag{3}$$

where N_o denotes the power spectral density of the noise and W is the channel bandwidth. This is an increasing and strictly convex function function of $U = (r, q)$. Other possibilities for $P(h, U)$ can be derived from considering particular adaptive coding/modulation schemes or from considering information theoretic bounds on the required power.[6]

In the above setting we are interested in both minimizing the long term transmission power and the average buffer delay. At time n, we assume that the transmitter chooses a rate distribution U_n (or equivalently R_n and q_n) based on the buffer state S_n and the channel state H_n. Let $\mathcal{S} = [0, \infty)$ denote the buffer state space* and let $\mathcal{U}$ denote the set of possible rate distributions. Assume that U_n at each time n is specified by a stationary policy, $\mu : \mathcal{S} \times \mathcal{H} \mapsto \mathcal{U}$. The expected long term average power under such a policy is

$$\limsup_{m\to\infty} \frac{1}{m}\sum_{n=1}^{m} \mathbb{E}(P(H_n, \mu(S_n, H_n))). \tag{4}$$

We denote this by $\bar{P}^\mu$. Similarly, define $\bar{D}^\mu$ to be

$$\limsup_{m\to\infty} \frac{1}{m}\sum_{n=1}^{m} \frac{\mathbb{E}(S_n)}{\bar{A}}. \tag{5}$$

Note that if the Markov chain induced by the policy μ is ergodic then we have $\bar{P}^\mu = \mathbb{E}P(H, \mu(S, H))$ and $\bar{D}^\mu = \frac{\mathbb{E}S}{\bar{A}}$. Thus from Little's law, $\bar{D}^\mu$ is the expected time average delay in the buffer (measured in units of Δ time intervals).

3. OPTIMAL POWER/DELAY TRADE-OFFS

We are interested in understanding the trade-off between the two objectives in (4) and (5). In the following we begin to characterize this trade-off. Consider minimizing a weighted combination of the two criteria. Specifically, for $\beta > 0$, we seek to to find the policy μ which minimizes:

$$\limsup_{m\to\infty} \frac{1}{m}\sum_{n=1}^{m} \mathbb{E}(P(H_n, \mu(S_n, H_n)) + \beta\frac{S_n}{\bar{A}}). \tag{6}$$

*Note we allow the buffer to be an arbitrary real value. This is done primarily for mathematical convenience.

This is an average cost Markov decision problem with state space $\mathcal{S} \times \mathcal{H}$. At each time step the transmitter chooses an action, namely the transmission rate, and incurs a per stage cost of $(P(H_n, \mu(S_n, H_n)) + \beta \frac{S_n}{\bar{A}})$. Such problems can be solved via dynamic programming techniques.[10] For the problem at hand, it can be shown that there always exists a stationary policy μ which is optimal and satisfies Bellman's equation for the average cost problem.

Assume that μ^* is an optimal policy for a given β. Let $\bar{P}^{\mu^*}$ and $\bar{D}^{\mu^*}$ be the corresponding average power and delay, as given in (4) and (5). Note that $\bar{P}^{\mu^*}$ must be the minimum average power such that the average delay is less than $\bar{D}^{\mu^*}$. For any $D \geq 1$, define $P^*(D)$ to be the minimum average power such that the average delay is less than D. Thus, by the above argument, $P^*(D^{\mu}) = P^{\mu^*}$. We refer to $P^*(D)$ as the (optimum) power/delay curve. Varying β and finding the optimal policy for each value can provide different points on the power/delay curve. It is natural to then ask if all values of $P^*(D)$ can be found in this way, with an appropriate choice of β. This problem can be viewed as a *multi-objective optimization* problem.[11] By this we mean an optimization problem with a vector valued objective function $f : X \mapsto \mathbb{R}^n$. In our case f has two components corresponding to the average delay and average power. For such problems, a feasible solution, x is defined to be *Pareto optimal* if there exists no other feasible $\hat{x}$ such that $f(\hat{x}) < f(x)$, where the inequality is to be interpreted component-wise. It can be seen that the points $\{(P^*(D), D) : D \geq 1\}$ are a subset of the Pareto optimal solutions for this problem. If $\{(P^*(D), D) : D \geq 1\}$ is not the entire set of Pareto optimal solutions, then from Prop. 1 below, for any remaining Pareto optimal point $(\tilde{P}, \tilde{D})$, it must be that $P^*(\tilde{D}) \leq \tilde{P}$. Furthermore, $\inf\{D : P^*(D) < \infty\}$ is only value of delay such points could have. Thus these other Pareto optimal solutions are not very interesting to us. For a general multi-objective optimization problem, not every Pareto optimal solution can be found by considering problems with scalar objectives $k'f$ where $k \in \mathbb{R}^n$. For the problem at hand, except in the degenerate case where the channel and arrival processes are both constant, every point on $P^*(D)$ (and thus every interesting Pareto optimal solution) can be found by solving the minimization in (6) for some choice of β. This also follows directly from the characterization of $P^*(D)$ given in the following proposition.

Proposition 1. The optimum power/delay curve, $P^*(D)$, is a non-increasing, convex function of $D \geq 1$. Except in the degenerate case where channel and arrival processes are both constant, it is a decreasing and strictly convex function of D.

The following definition will be useful in proving this proposition. Let

$$P_{min}(h, \bar{u}) = \inf\{P(h, r, q) : r\Delta(1 - q) \geq \bar{u}\}. \tag{7}$$

This is the minimum power required for an expected transmission rate greater than or equal to $\bar{u}/\Delta$ when the channel state is h. Using the convexity of $P(h, r, q)$ it can be shown that $P_{min}(h, \bar{u})$ is strictly convex in $\bar{u}$ for all h.

Proof of Proposition 1. That $P^*(D)$ is non-increasing is obvious. We show that it is convex. Let D^1 and D^2 be two values of delay with corresponding values $P^*(D^1)$ and $P^*(D^2)$. We want to show that for any $\lambda \in [0, 1]$,

$$P^*(\lambda D^1 + (1 - \lambda)D^2) \leq \lambda P^*(D^1) + (1 - \lambda)P^*(D^2). \tag{8}$$

We prove this using a sample path argument. Let $\{H_n(\omega)\}_{n=1}^{\infty}$ and $\{A_n(\omega)\}_{n=1}^{\infty}$ be a given sample path of channel states and arrival states. Let $\{U_n^1(\omega)\}$ be a sequence of control actions corresponding to the policy which attains $P^*(D^1)$. Let $\{S_n^1(\omega\}$ be the corresponding sequence of buffer states. Likewise define $\{U_n^2(\omega\}$ and $\{S_n^2(\omega)\}$ corresponding to $P^*(D^2)$. Now consider the new sequence of control actions, $\{U_n^{\lambda}(\omega)\}$, where for all n, $U_n^{\lambda}(\omega)$ is the control action which achieves $P_{min}(H_n(\omega), \lambda\mathbb{E}(U_n^1(\omega)) + (1 - \lambda)\mathbb{E}(U_n^2(\omega))$. Let $\{S_n^{\lambda}(\omega)\}$ be the sequence of buffer states using this policy. Assume at time $n = 0$, $S_0^{\lambda}(\omega) = S_0^1(\omega) = S_0^2(\omega) = 0$ for all sample paths, ω. Using that $S_{n+1}^i(\omega) = S_n^i(\omega) + A_n(\omega) - U_{n+1}^i(\omega)$ for $i = 1, 2$ and all $n \geq 0$, and recursion, we have $\mathbb{E}(S_n^{\lambda}(\omega)) \leq \lambda\mathbb{E}(S_n^1(\omega)) + (1 - \lambda)\mathbb{E}(S_n^2(\omega))$ for all n. Thus, summing and taking the expectation over all sample paths yields

$$\lim_{m\to\infty} \frac{1}{m} \sum_{n=1}^{m} \frac{\mathbb{E}S_n^{\lambda}}{\bar{A}} \leq \lambda D^1 + (1 - \lambda)D^2. \tag{9}$$

From the convexity of $P_{min}(h, \bar{u})$ in $\bar{u}$, we have for all n

$$\begin{aligned} P(H_n(\omega), U_n^{\lambda}(\omega)) &= P_{min}(H_n(\omega), \lambda\mathbb{E}(U_n^1(\omega)) + (1 - \lambda)\mathbb{E}(U_n^2(\omega))) \\ &\leq \lambda P_{min}(H_n(\omega), \mathbb{E}(U_n^1(\omega))) + (1 - \lambda)P_{min}(H_n(\omega), \mathbb{E}(U_n^2(\omega))) \\ &\leq \lambda P(H_n(\omega), (U_n^1(\omega)) + (1 - \lambda)P(H_n(\omega), U_n^2(\omega)). \end{aligned}$$

Again, summing and taking expectations over sample paths we have

$$\lim_{m\to\infty}\frac{1}{m}\sum_{n=1}^{m}\mathbb{E}P(H_n(\omega),U_n^{\lambda}(\omega))\leq \lambda P^*(D^1)+(1-\lambda)P^*(D^2). \tag{10}$$

Thus we must have $P^*(\lambda D_1+(1-\lambda)D_2)\leq \lambda P^*(D_1)+(1-\lambda)P^*(D_2)$ as desired.

The final statement in the proposition follows directly from the above and the results in the next section. ■

Define $\mathcal{P}_d(a)=\mathbb{E}_H P(H,a/\Delta,0)$, this is the minimum average power required to transmit at rate a/Δ with no probability of loss for every channel state. For the discrete-time model we have formulated the delay in the buffer must be at least one time interval, thus the average buffer delay can be no less than 1. The only way this bound can be attained is if for all n, $R_n\Delta=A_{n-1}$ and $q_n=0$. Thus we have :

$$P^*(1)=\int_{\mathcal{A}}\mathcal{P}_d(a)\,d\pi_A(a). \tag{11}$$

For most realistic channel models $P^*(1)$ will be infinite.

Define $\mathcal{P}_a(\bar{A})$ be the solution to

$$\begin{aligned}\underset{\Psi:\mathcal{H}\mapsto\mathcal{U}}{\text{minimize}}\ \ &\mathbb{E}P(H,\Psi(H))\\ \text{subject to}:\ &\mathbb{E}(\Psi(H))\geq\bar{A}\end{aligned} \tag{12}$$

We have restricted Ψ to be only a function of the channel state H in this optimization. Clearly for each $h\in\mathcal{H}$, $P(h,\Psi(h))=P_{min}(h,\mathbb{E}\Psi(h))$. From this it follows that $\mathcal{P}_a(\bar{A})$ is an increasing and strictly convex function of $\bar{A}$. Let $\Psi^{\bar{A}}(h)$ denote a feasible rate allocation that achieves $\mathcal{P}_a(\bar{A})$. It can be seen that this rate allocation will be almost surely unique and is only a function of h. Furthermore, it is a continuous function of h for all $\bar{A}>0$. Likewise, for any $h\in\mathcal{H}$, $\Psi^{\bar{A}}(h)$ is a continuous and non-decreasing function of $\bar{A}$. The quantity $\mathcal{P}_a(\bar{A})$ is the minimum average power needed to transmit at average rate $\bar{A}$ with no other constraints. Thus $\mathcal{P}_a(\bar{A})$ is a lower bound to $P^*(D)$ for all $D\geq 1$. If both the channel and arrival processes are constant, then $\mathcal{P}_a(\bar{A})=\mathcal{P}_d(\bar{A})$; in this case, the power delay curve is a horizontal line. Assume that the channel and arrival processes are not both constant. In this case, the only way a stationary policy μ can have $\bar{P}^{\mu}=\mathcal{P}_a(\bar{A})$ is if $\mu(s,h)=\Psi^{\bar{A}}(h)$ for all $(s,h)\in\mathcal{S}\times\mathcal{H}$, except a set with measure zero. Such a policy results in $\bar{D}^{\mu}=\infty$. In other words, $P^*(D)>\mathcal{P}_a(\bar{A})$ for all finite D. In the next section we shall see that this bound can be approached as $D\to\infty$.

An example of the power/delay curve is shown in Figure 2. This was calculated using dynamic programming techniques for a channel model with two states and a constant arrival rate.

4. ASYMPTOTIC ANALYSIS

In this section we characterize the behavior of the tail of the power/delay curve, $P^*(D)$, as the buffer delay $D\to\infty$. In this limit, we shown that $P^*(D)\to\mathcal{P}_a(\bar{A})$. We also look at the rate† at which this limit is approached and show that $P^*(D)-\mathcal{P}_a(\bar{A})=\Theta(\frac{1}{D^2})$. First we bound the rate of approach. Then we show that this bound is tight. Furthermore we show that a sequence of "simple" policies can be found which have this rate of convergence. The approach in this section is closely related to Tse's work[12] on buffer control for variable rate lossy compression. In this work the input rate into a buffer is controlled by changing the quantizer used to compress blocks of real valued data. The goal is to optimally trade-off distortion and buffer overflow probability. In the problem at hand, the buffer is controlled by varying the output rate and we interested in trading off power and average delay. There are many similarities between the mathematical structure underlying these problems.

To characterize the behavior of this tail, we will consider a sequence of policies $\{\mu_k\}$, such that as $k\to\infty$, $\bar{D}^{\mu_k}\to\infty$ and $\bar{P}^{\mu_k}\to\mathcal{P}_a(\bar{A})$. We restrict our attention to the class of policies that satisfy the following technical assumptions.

†The following notation is used to compare the rates of growth of two real-valued sequences $\{a_n\}$ and $\{b_n\}$: $a_n=o(b_n)$ if $\lim_{n\to\infty}\frac{a_n}{b_n}=0$; $a_n=O(b_n)$ if $\limsup_{n\to\infty}\frac{|a_n|}{|b_n|}<\infty$; $a_n=\Omega(b_n)$ if $b_n=O(a_n)$; and $a_n=\Theta(b_n)$ if $a_n=O(b_n)$ and $a_n=\Omega(b_n)$.

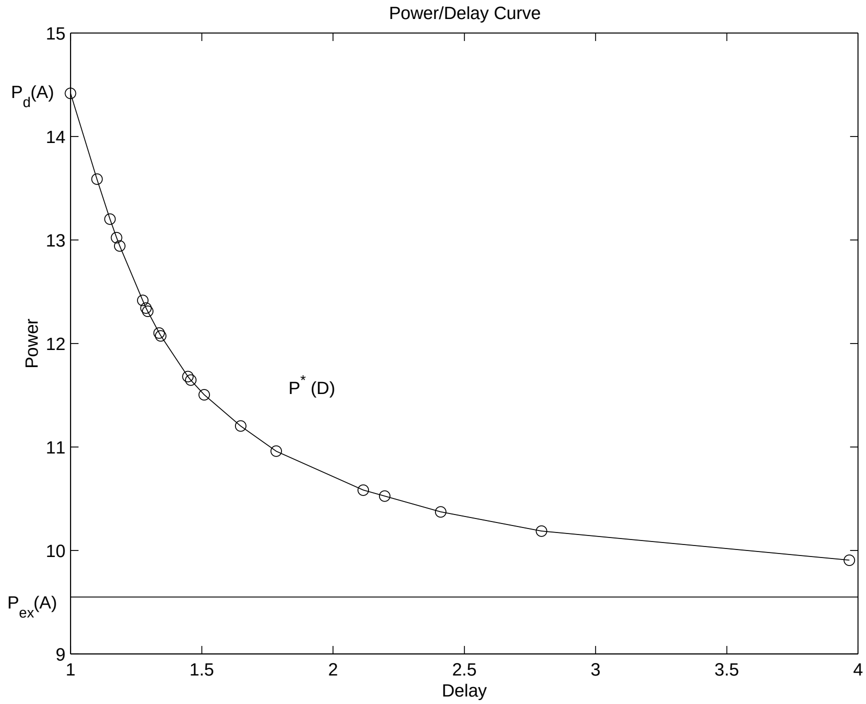

Figure 2. Example of power/delay curve.

Definition: A sequence of buffer control policies $\{\mu_k\}$ is *admissible* (for a given fading process, $\{H_n\}$ and arrival process $\{A_n\}$) if it satisfies the following assumptions:

1. For all k, $\bar{D}^{\mu_k} < \infty$, and $\lim_{k\to\infty} \bar{D}^{\mu_k} = \infty$.
2. Under each policy, μ_k, $\{S_n\}$ forms an ergodic Markov chain; we denote the steady state distribution under the kth policy by $\pi_S^{\mu_k}$.
3. There exists an $\epsilon > 0$, a $\delta > 0$ and a $M > 0$ such that for all $k > M$ and for all $s \leq 2\mathbb{E}(S^{\mu_k})$,

$$\Pr(A - \mu_k(S^{\mu_k}, H) > \delta | S^{\mu_k}{=}s) > \epsilon$$

where S^{μ_k}, A and H are random variables with respective state spaces $\mathcal{S}$, $\mathcal{A}$, and $\mathcal{H}$ and whose joint distribution is the steady state distribution of (S_n, A_n, H_n) under policy μ_k.

We are interested in sequences of policies which characterize the tail behavior of $P^*(D)$ as $D \to \infty$. The first assumption says a sequence of policies is admissible only if the average delay of these policies has the desired behavior. Under any stationary policy, the sequence of buffer states is a Markov chain. The policy, along with the fading process and arrival process, determines the transition kernel for this Markov chain. By the second assumption, for each policy in an admissible sequence, this Markov chain is ergodic. This will be true if the transition kernel is "well-behaved".[13] The third assumption means that for large k and any buffer state $s < 2\mathbb{E}(S^{\mu_k})$, there is a positive steady state probability that the next buffer state is bigger than $s + \delta$. If $A_{min} > \delta$ and $\Pr(U = 0) > \epsilon$ for all $U \in \mathcal{U}$, then this assumption must be satisfied by any policy. If this is not the case, then this is a restriction on the allowable policies.[‡]

We also assume that at $a = \bar{A}$, the first and second derivatives of $\mathcal{P}_a(a)$ exist and are non zero. Recall, $\mathcal{P}_a(a)$ is a strictly convex and increasing function of a. For such a function, the first and second derivatives of $\mathcal{P}_a(a)$ exist and are non-zero at every point except for a set with measure zero (This follows from Lebesgue's theorem which states that a monotonic function is differentiable almost everywhere[14]). Thus, this is not a very restrictive assumption.

[‡]It can be argued that for any sequence of policies satisfying the first condition and such that $\bar{P}^{\mu_k} \to \mathcal{P}_a(\bar{A})$, then provided that both A_n and H_n are not constant for all n, assumption 3 must hold, except not necessarily uniformly over s.

Let $\Delta^\mu(s) = \mathbb{E}(A - \mu(S^\mu, H)|S^\mu{=}s)$ denote the expected drift given that the buffer is in state s under policy μ. For any admissible sequence of policies, the average drift over the tail of the buffer will be negative. This is stated in the following lemma.

Lemma 1. Let M, δ and ϵ be as given in the definition of an admissible sequence. For any admissible sequence of buffer control schemes $\{\mu_k\}$, for each $k > M$, there exists an $s_k \in \mathcal{S}$ such that

$$\int_{s>s_k} \Delta^{\mu_k}(s)\, d\pi_S^{\mu_k}(s) \leq \frac{-\epsilon\delta^2}{16\mathbb{E}(S^{\mu_k})}$$

The proof of this lemma is omitted due to space limitations. We will use this result to establish a bound on the rate of convergence of $P^*(D)$. This bound is stated in the following proposition, which implies that the "tail" of $P^*(D)$ converges to $\mathcal{P}_a(\bar{A})$ at least as slowly as $\frac{1}{D^2}$.

Proposition 2. Any admissible sequence of policies $\{\mu_k\}$ must satisfy

$$\bar{P}^{\mu_k} - \mathcal{P}_a(\bar{A}) = \Omega((1/\bar{D}^{\mu_k})^2).$$

Proof. For the kth policy, let $\Delta^{\mu_k}(s)$ denote the expected drift in state s as above. The average transmission rate conditioned on being in state s is $\mathbb{E}(\mu_k(S^{\mu_k}, H)|S^{\mu_k}{=}s) = \bar{A} - \Delta^{\mu_k}(s)$. Recall that $\mathcal{P}_a(x)$ is the minimum average power required to transmit at average rate x. Thus the average power used under policy μ_k when the buffer is in state s is lower bounded by $\mathcal{P}_a(\bar{A} - \Delta^{\mu_k}(s))$. Averaging over the buffer state space we have:

$$\bar{P}^{\mu_k} \geq \int_{\mathcal{S}} \mathcal{P}_a(\bar{A} - \Delta^{\mu_k}(s))\, d\pi_S(s) \tag{13}$$

Via a first order Taylor expansion around $x = \bar{A}$, $\mathcal{P}_a(x)$ can be written as:

$$\mathcal{P}_a(x) = \mathcal{P}_a(\bar{A}) + \mathcal{P}_a'(\bar{A})(x - \bar{A}) + G(x - \bar{A}) \tag{14}$$

where the remainder term $G(x)$ has the following properties: (i) $G(x)$ is strictly convex, (ii) for $x \neq 0$, $G(x) > 0$ and $G(0) = 0$, and (iii) $G'(x) > 0$ for $x > 0$ and $G'(0) = 0$. These all follow from the strict convexity and monotonicity of $\mathcal{P}_a$. Substituting this into (13) yields:

$$\bar{P}^{\mu_k} - \mathcal{P}_a(\bar{A}) \geq \mathcal{P}_a'(\bar{A}) \int_{\mathcal{S}} (-\Delta^{\mu_k}(s))\, d\pi_S(s) + \int_{\mathcal{S}} G(-\Delta^{\mu_k}(s))\, d\pi_S(s) \tag{15}$$

$$= \int_{\mathcal{S}} G(-\Delta^{\mu_k}(s))\, d\pi_S(s). \tag{16}$$

Here we have used that

$$\int_{\mathcal{S}} \Delta^{\mu_k}(s)\, d\pi_S(s) = 0 \tag{17}$$

for any policy μ_k which has $\mathbb{E}S^{\mu_K} < \infty$. This follows from the fact that the buffer size is infinite and thus no bits are lost due to overflow. Let s_k be as defined in Lemma 1 and assume that $k > M$ so that the lemma applies. Then we have

$$\begin{aligned} &\bar{P}^{\mu_k} - \mathcal{P}_a(\bar{A}) \\ &\geq \int_{s>s_k} G\left(-\Delta^{\mu_k}(s)\right)\, d\pi_S(s) \\ &= \int_{s>s_k} G\left(-\Delta^{\mu_k}(s)\right)\, d\pi_S(s) + \pi_S([0, s_k))G(0) \\ &\geq G\left(\int_{s>s_k} -\Delta^{\mu_k}(s)\, d\pi_S(s)\right) \\ &\geq G\left(\frac{\epsilon\delta^2}{16\mathbb{E}S^{\mu_k}}\right) \end{aligned} \tag{18}$$

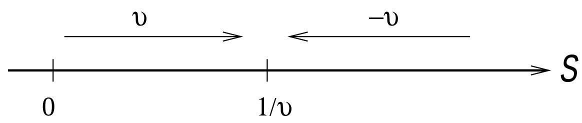

Figure 3. A simple policy with drift υ.

We have again used the concavity and monotonicity properties of G along with the result of Lemma 1. Finally, expanding G in a Taylor series around 0, and using $G'(0) = 0$ we have:

$$\bar{P}^{\mu_k} - \mathcal{P}_a(\bar{A}) \geq \frac{1}{2}G''(0)\left(\frac{\epsilon\delta^2}{16\mathbb{E}S^{\mu_k}}\right)^2 + o\left(\left(\frac{\epsilon\delta^2}{16\mathbb{E}S^{\mu_k}}\right)^2\right). \tag{19}$$

That $G''(0)$ exists and is non-zero follows from the assumption that the second derivative of $\mathcal{P}_a(x)$ exists and is non zero at $x = \bar{A}$. Thus we have $\bar{P}^{\mu_k} - \mathcal{P}_a(\bar{A}) = \Omega((\frac{1}{\mathbb{E}(S^{\mu_k})})^2)$. Using Little's law, this gives us $\bar{P}^{\mu_k} - \mathcal{P}_a(\bar{A}) = \Omega((1/\bar{D}^{\mu_k})^2)$ as desired. ■

Next we show that this bound is tight. To do this we give a sequence of policies, which achieve the rate of convergence given by the bound, *i.e.* we show that there exists a sequence of policies μ_k, such that $\bar{P}^{\mu_k} - \mathcal{P}_a(\bar{A}) = O((1/\bar{D}^{\mu_k})^2)$. Moreover the sequence of policies that we use have a relatively simple structure to them. First we describe the type of policies to be used. Then, the convergence rate of these policies is demonstrated.

Definition: For a given $\upsilon > 0$, partition the buffer state space into two distinct sets: $[1/\upsilon, \infty)$ and $[0, 1/\upsilon)$. Recall, $\Psi^a : \mathcal{H} \mapsto \mathcal{U}$ denotes the policy with average rate a/Δ which achieves $\mathcal{P}_a(a)$. Under such a policy, the transmission rate depends only on the channel state. Define a *simple policy with drift* υ, to be a policy μ with the form:[§]

$$\mu(s,h) = \begin{cases} \Psi^{\bar{A}+\upsilon}(h) & \text{if } s \in [1/\upsilon, \infty) \\ \Psi^{\max(\bar{A}-\upsilon,0)}(h) & \text{if } s \in [0, 1/\upsilon). \end{cases}$$

Under a simple policy the transmission rate now depends on the buffer state as well as the channel state, but it depends only weakly on the buffer state. For such a policy, the drift in any buffer state $s \geq 1/\upsilon$ will be $-\upsilon$ and the drift in any state $s \leq 1/\upsilon$ will be υ provided that $\upsilon < \bar{A}$ (otherwise the drift will be $\bar{A}$). Thus these policies tend to regulate the buffer towards the state $1/\upsilon$ as illustrated in Figure 3. A sequence of such policies can achieve the optimal convergence rate; this is proved next.

Proposition 3. Let $\{\mu_k\}$ be a sequence of simple policies with drifts $\{\upsilon_k\}$, where $\{\upsilon_k\}$ is a nonnegative decreasing sequence such that $\upsilon^k \to 0$ as $k \to \infty$. Then we have $\bar{P}^{\mu_k} - \mathcal{P}_a(\bar{A}) = O((\frac{1}{\bar{D}^{\mu_k}})^2)$.

Proof. We show that $\bar{D}^{\mu_k} = O(\frac{1}{\upsilon_k})$ and $\bar{P}^{\mu_k} - \mathcal{P}_a(\bar{A}) = O((\upsilon_k)^2)$. The desired result then follows directly. First we show that $\bar{D}^{\mu_k} = O(\frac{1}{\upsilon_k})$. This relies on the following lemma.

lemma 2. For a simple policy μ with drift υ, the average delay satisfies:

$$\bar{D}^{\mu} \leq \frac{1/\upsilon}{\bar{A}} + \frac{e^{r^*(\upsilon)\eta(\upsilon)}}{\bar{A}r^*(\upsilon)} \tag{20}$$

where $\eta(\upsilon)$ is a nonnegative function such that $\eta(\upsilon) \to 0$ as $\upsilon \to 0$, and $r^*(\upsilon)$ is the unique positive root of $\gamma(r) = \ln(\mathbb{E}[e^{(A-\Psi^{\bar{A}+\upsilon}(H))r}])$ (Note the expected value is taken with respect to both A and H).

We omit a proof of this lemma but briefly mention two key ideas which are used in the proof. First, Little's law is used to relate the average delay to the average buffer occupancy. Second, for the memoryless case, while the buffer

[§]More generally, we could partition the buffer into the sets $[0, K/\upsilon)$ and $[K/\upsilon, \infty)$ where $K > 0$. These sets could then be used in the definition of a simple policy. The following results still hold with such a generalization.

process stays in $[1/\upsilon, \infty)$ it behaves as a random walk with a negative drift. Thus the steady-state probability that the buffer is in state s, decays exponentially for s large enough.

From Lemma 2 we have a bound on $\bar{D}^{\mu_k}$ given by (20). The first term on the right hand side of (20) is clearly $O(1/\upsilon_k)$. We focus on the second term of (20). To bound the growth of this term, we need to know how $r^*(\upsilon_k)$ changes with υ_k. This is given in the following lemma whose proof is also omitted.

Lemma 3. Let $r^*(\upsilon)$ denote the unique nonzero root of the semi-invariant moment generating function of $A - \Psi^{\bar{A}+\upsilon}(H)$ (for $\upsilon \neq 0$). Assume that for all υ in a neighborhood of 0, that $\frac{d^2}{d\upsilon^2}\mathbb{E}e^{r^*(\upsilon)(A-\Psi^{\bar{A}+\upsilon}(H))}$ exists and that¶

$$\frac{d^2}{d\upsilon^2}\mathbb{E}e^{r^*(\upsilon)(A-\Psi^{\bar{A}+\upsilon}(H))} = \mathbb{E}\frac{d^2}{d\upsilon^2}e^{r^*(\upsilon)(A-\Psi^{\bar{A}+\upsilon}(H))}.$$

Then, $r^*(0) = 0$ and

$$\left.\frac{dr^*(\upsilon)}{d\upsilon}\right|_{\upsilon=0} = \frac{2}{\mathrm{Var}(A - \Psi^{\bar{A}}(H))}.$$

Taking the Taylor series of $r^*(\upsilon)$ around $\upsilon = 0$ and using this lemma we have

$$r^*(\upsilon) = 0 + \Lambda\upsilon + o(|\upsilon|) \tag{21}$$

where $\Lambda = \frac{2}{\mathrm{Var}(A-\Psi^{\bar{A}}(H))}$. Recall in Lemma 2 it was shown that $\eta(\upsilon) \to 0$ as $\upsilon \to 0$. From this it follows that $r^*(\upsilon)\eta(\upsilon) = \Lambda\eta(\upsilon)\upsilon + o(|\upsilon|)$. With these expansions we have

$$\frac{e^{r^*(\upsilon_k)\eta(\upsilon_k)}}{\bar{A}r^*(\upsilon_k)} = \frac{e^{\Lambda\eta(\upsilon_k)\upsilon_k + o(\upsilon_k)}}{\bar{A}(\Lambda\upsilon_k + o(\upsilon_k))}. \tag{22}$$

Now since:

$$\lim_{k\to\infty} \frac{\upsilon_k e^{\Lambda\eta(\upsilon_k)\upsilon_k + o(\upsilon_k)}}{\bar{A}(\Lambda\upsilon_k + o(\upsilon_k))} = \frac{1}{\bar{A}\Lambda} \tag{23}$$

it follows that:

$$\frac{e^{r^*(\upsilon_k)\eta(\upsilon_k)}}{\bar{A}r^*(\upsilon_k)} = O(1/\upsilon_k) \tag{24}$$

and therefore $\bar{D}^{\mu_k} = O(1/\upsilon_k)$ as desired.

Next we show that $\bar{P}^{\mu_k} - \mathcal{P}_a(\bar{A}) = O((\upsilon_k)^2)$. For the simple policy μ_k, the average power is

$$\bar{P}^{\mu_k} = \pi_S^{\mu_k}((1/\upsilon_k, \infty))\mathcal{P}_a(\bar{A} + \upsilon_k) + \pi_S^{\mu_k}([0, 1/\upsilon_k])\mathcal{P}_a(\bar{A} - \upsilon_k) \tag{25}$$

Taking the Taylor series of $\mathcal{P}(x)$ around $x = \bar{A}$ we have

$$\bar{P}^{\mu_k} = \mathcal{P}_a(\bar{A}) + \mathcal{P}'_a(\bar{A})(\pi_S^{\mu_k}((1/\upsilon, \infty))\upsilon_k - \pi_S^{\mu_k}([0, 1/\upsilon])\upsilon_k) + O((\upsilon_k)^2) \tag{26}$$

Now $\pi_S^{\mu_k}((1/\upsilon, \infty))\upsilon_k - \pi_S^{\mu_k}([0, 1/\upsilon])\upsilon_k \geq 0$ and thus $\bar{P}^{\mu_k} - \mathcal{P}_a(\bar{A}) = O((\upsilon_k)^2)$ as desired. ■

The simple policies used in the proof of Proposition 3 involved splitting the buffer into two regions. In each region a policy was used that depended only on the current channel state. When the transmission rate is a function of the buffer state, this information may need to be relayed to the receiver. Conveying this information to the receiver requires some overhead. When a simple policy is used, the receiver only needs to know in which region of the buffer the current buffer state lies; this requires only one bit of overhead. An even simpler policy would be to use one policy for the entire buffer; the policy just depending on the channel gain. With such a policy, the receiver would require

¶As an example of when these assumptions will hold, assume that $|\mathcal{A}| < \infty$ and $|\mathcal{H}| < \infty$. In this case if the second derivative of $\Psi^{\bar{A}+\upsilon}(h)$ with respect to υ exists and is continuous at $\upsilon = 0$ for all h, then the above assumptions hold.

no information about the transmitter's buffer state. Proposition 4 below shows that a sequence of such policies can not achieve the optimal convergence rate. Before stating this proposition some preliminary notation is established.

We want to consider a sequence of policies $\{\mu_k\}$ which depend only on the channel state. Let $\upsilon_k = \bar{A} - \mathbb{E}\mu_k(H)$; the average transmission rate in every buffer state $s \in \mathcal{S}$ is then $\bar{A} + \upsilon_k$. For the buffer to be stable under policy μ_k it must be that $\upsilon_k > 0$. To prove Proposition 4, we will use a result similar to Lemma 3. However, we do not want to restrict the policy μ_k to be a policy of the form Ψ^x as in Lemma 3. Instead we assume that that each policy μ_k is determined by an arbitrary parameterized function Φ^x. Specifically, for every $x \geq \bar{A}$, let $\Phi^x : \mathcal{H} \mapsto \mathbb{R}^+$ be an arbitrary policy which depends only on the channel gain such that $\mathbb{E}\Phi^x(H) = x$. Assume that each policy μ_k is given by $\mu_k = \Phi^{\bar{A}+\upsilon_k}$. Let $r^*(\upsilon)$ denote the unique nonzero root of the semi-invariant moment generating function of $A - \Phi^{\bar{A}+\upsilon}(H)$. Assume that for all υ in a neighborhood of 0, that $\frac{d^2}{d\upsilon^2}\mathbb{E}e^{r^*(\upsilon)(A-\Phi^{\bar{A}+\upsilon}(H))}$ exists and that

$$\frac{d^2}{d\upsilon^2}\mathbb{E}e^{r^*(\upsilon)(A-\Phi^{\bar{A}+\upsilon}(H))} = \mathbb{E}\frac{d^2}{d\upsilon^2}e^{r^*(\upsilon)(A-\Phi^{\bar{A}+\upsilon}(H))}.$$

Note this is the same set of assumptions used in Lemma 3 and they serve a similar purpose here. Any sequence of policies μ_k satisfying the above assumptions can not achieve the optimal convergence rate; this is stated in the following proposition, which is given without proof..

Proposition 4. Let $\{\upsilon_k\}$ be a nonnegative decreasing sequence such that $\upsilon^k \to 0$ as $k \to \infty$. Let $\{\mu_k\}$ be a sequence of policies such that for each k, $\mu_k = \Phi^{\bar{A}+\upsilon_k}$, where Φ^x satisfies the above assumptions. Then $\bar{P}^{\mu_k} - \mathcal{P}_a(\bar{A}) = \Omega(\frac{1}{\bar{D}^{\mu_k}})$.

Thus using more than one policy allows the rate of convergence to be squared. Some intuition as to why two policies are needed is given by the following argument. With two policies we regulate the buffer towards the point $\frac{1}{\upsilon}$, while with one policy (with finite average delay) the buffer is regulated towards the empty state. When considering average delay, keeping the buffer empty appears more desirable. However, when considering the average power, there is a disadvantage to keeping the buffer nearly empty– when the buffer is nearly empty, one can not take advantage of a good channel by transmitting at a high rate, which is desirable for minimizing power. By using two policies and regulating the buffer towards the point $\frac{1}{\upsilon}$, a better balance is obtained between these two considerations.

5. A GENERALIZATION TO UTILITY BASED QOS

In the previous section we assumed that a user's QoS was given by the average delay experienced by the data. In this section we give a generalization this formulation which allows for a larger classes of QoS specifications. We consider the case where during the nth time interval the user receives a utility $U_t(\mu(S_n, H_n))$ which is a function of the transmission rate distribution during the time slot. Additionally, assume the user incurs a penalty $Q_d(S_n)$ which is a function of the buffer occupancy at the start of the time interval. In this setting instead of (5), we consider minimizing the objective

$$\limsup_{m\to\infty} \frac{1}{m}\sum_{n=1}^{m} \mathbb{E}(Q_d(S_n) - U_t(\mu(S_n, H_n)))$$

This corresponds to the time average penalty minus utility received. For a particular stationary policy μ, we denote the above quantity by $\bar{Q}^\mu$

Assume that $Q_d(s)$ is a non-negative, increasing, convex function of s. Likewise, assume that $U_t(u)$ is a non-negative, non-decreasing, concave function of u. In the following we show that with some additional assumptions the asymptotic analysis in the previous section can be generalized to this case. The bound on the rate of convergence in Proposition 2, is generalized in the following proposition.

Proposition 5. Any admissible sequence of policies $\{\mu_k\}$ must satisfy:

$$\bar{P}^{\mu_k} - \mathcal{P}_a(\bar{A}) = \Omega\left(1/(Q_d^{-1}(\bar{Q}^{\mu_k}))^2\right).$$

Proof. Note that $Q_d(s) - U_t(u)$ is convex in (s, u) thus for all m,

$$\frac{1}{m}\sum_{n=1}^{m}\mathbb{E}(Q_d(S_n) - U_t(\mu_k(S_n, H_n))) \geq Q_d\left(\frac{1}{m}\sum_{n=1}^{m}\mathbb{E}S_n\right) - U_t\left(\frac{1}{m}\sum_{n=1}^{m}\mathbb{E}\mu_k(S_n, H_n)\right).$$

Thus letting $m \to \infty$, we have $\bar{Q}^{\mu_k} \geq Q_d(\mathbb{E}S^{\mu_k}) - U_t(\bar{A})$. Since Q_d is nondecreasing it has a well defined inverse, Q_d^{-1} which is also nondecreasing, thus

$$\begin{aligned} Q_d^{-1}(\bar{Q}^{\mu_k}) &\geq Q_d^{-1}(Q_d(\mathbb{E}S^{\mu_k}) - U_t(\bar{A})) \\ &\geq Q_d^{-1}(Q_d(\mathbb{E}S^{\mu_k})) \\ &= \mathbb{E}S^{\mu_k} \\ &= \bar{D}^{\mu_k}\bar{A}. \end{aligned}$$

Using this the proposition then follows directly from Proposition 2. ■

If $Q_d(D)$ satisfies an additional assumption then once again this bound can be shown to be tight and achieved by a sequence of simple policies. This is summarized in the following proposition; here Q_d' indicates the first derivative of Q_d.

Proposition 6. Assume that there exist M_1 and M_2 such that for all $s \in \mathcal{S}$, $0 \leq M_1 \leq Q_d'(s) \leq M_2 \leq \infty$. Let $\{\mu_k\}$ be a sequence of simple policies, with drifts v_k, where v_k is a nonnegative decreasing sequence such that $v_k \to 0$ as $k \to \infty$. Then we have $\bar{P}^{\mu_k} - \mathcal{P}_a(\bar{A}) = O\left(\left(\frac{1}{Q_d^{-1}(\bar{Q}^{\mu_k})}\right)^2\right)$.

Proof. By assumption we can find a $M_2 < \infty$ such that $Q_d'(s) < M$ for all s. Thus

$$Q_d(S_n) \leq Q_d(\mathbb{E}S^{\mu_k}) + M_2 S_n \quad \forall n.$$

From this it follows that for all n,

$$\begin{aligned} \mathbb{E}(Q_d(S_n) - U_t(\mu_k(S_n, H_n))) &\leq Q_d(\mathbb{E}S^{\mu_k}) + M_2\mathbb{E}S_n - \mathbb{E}U_t(\mu_k(S_n, H_n)) \\ &\leq Q_d(\mathbb{E}S^{\mu_k}) + M_2\mathbb{E}S_n \end{aligned}$$

Thus averaging over n and taking the limit, we have

$$\bar{Q}^{\mu_k} \leq Q_d(\mathbb{E}S^{\mu_k}) + M_2\mathbb{E}S^{\mu_k}$$

Since Q_d is increasing and convex, Q_d^{-1} will be increasing and concave. From this fact it follows that:

$$\begin{aligned} Q_d^{-1}(\bar{Q}^{\mu_k}) &\leq Q_d^{-1}(Q_d(\mathbb{E}S^{\mu_k}) + M_2\mathbb{E}S^{\mu_k}) \\ &\leq Q_d^{-1}(Q_d(\mathbb{E}S^{\mu_k})) + M_2/M_1\mathbb{E}S^{\mu_k} \\ &= (1 + M_2/M_1)\mathbb{E}S^{\mu_k} \\ &= (1 + M_2/M_1)\bar{A}\bar{D}^{\mu_k} \end{aligned}$$

Combining this with Proposition 3 we get the desired result. ■

We also note that by a similar argument to that used in the proof of Proposition 5, Proposition 4 can also be generalized to this setting.

6. CONCLUSIONS

In this paper we have looked at a simple model of transmission scheduling over a time-varying channel which takes into account both physical layer and network layer concerns. We investigated the trade-off between long term average power and average buffer delay. In the asymptotic region of large delay, we bounded the rate of convergence of the required transmission power, and we found simple buffer control policies which exhibit the asymptotic convergence rate. Next we extended this analysis to a case where the user's quality of service is specified in a more general utility framework. In this paper we only considered a single user setting, clearly extensions to multi-user situations would be desirable. Also in this paper our emphasis was on users with a long term average power constraint. Such a model is appropriate for the reverse link in a cellular system, but less appropriate for the forward link. However in the forward link setting other resource allocation issues arise, extensions to these situations is also of interest.

ACKNOWLEDGMENTS

Many of the ideas behind this work are derived from the author's PhD. thesis supervised by R. G. Gallager, whose guidance and support are gratefully acknowledged.

REFERENCES

1. R. Berry, "Communication over fading channels with finite buffer constraints - single user and multiple access cases," in *Proc. 2001 IEEE International Symposium on Info. Th.*, (Washington, DC), June 24-29 2001.
2. N. Bambos, "Toward power-sensitive network architectures in wireless communications: concepts, issues, and design aspects," *IEEE Personal Communications Magazine* , June 1998.
3. B. Collins and R. Cruz, "Transmission Policies for Time Varying Channels with Average Delay Constraints," in *Proc. 1999 Allerton Conf. on Commun. Control, & Comp.*, (Monticello, IL), 1999.
4. R. Leelahakriengkrai and R. Agrawal, "Scheduling in multimedia cdma wireless networks," tech. rep., ECE Dept., Univ. of Wisconsin-Madison, July 1999.
5. A. Ephremides and B. Hajek, "Information theory and communication networks: An unconcummated union," *IEEE Trans. on Information Theory* **44**, pp. 2416–2434, Oct 1998.
6. R. Berry and R. Gallager, "Communication over fading channels with delay constraints." *submitted to IEEE Trans. on Info. Th.*, 2000.
7. R. Berry, *Power and Delay Trade-offs in Fading Channels.* PhD thesis, Massachusetts Institute of Technology, 2000.
8. R. Berry and R. Gallager, "Buffer control for communication over fading channels," in *Proc. 2000 IEEE International Symposium on Info. Th.*, p. 409, (Sorrento, Italy), June 25-30 2000.
9. L. Ozarow, S. Shamai, and A. Wyner, "Information Theoretic Considerations for Cellular Mobil Radio," *IEEE Tranactions on Vehicular Technology, Vol. 43, No. 2* , pp. 359–378, May 1994.
10. D. Bertsekas, *Dynamic Programing and Optimal Control, Vol. I and II*, Athena Scientific, 1995.
11. Y. Sawaragi, *Theory of Multiobjective Optimization*, Academic Press, Orlando, 1985.
12. D. Tse, *Variable-rate Lossy Compression and its Effects on Communication Networks.* PhD thesis, Massachusetts Institute of Tech., Cambridge, MA 02139, Sep 1994. Also available as LIDS-TH-2269.
13. W. Feller, *An Introduction To Probability Theory and Its Applications*, vol. 2, Wiley, New York, 1957.
14. F. Riesz and B. Nagy, *Functional Analysis*, Ungar, New York, 1955.