Communication Over Fading Channels with Delay Constraints

Randall A. Berry and Robert G. Gallager

Abstract

We consider a user communicating over a fading channel with perfect channel state information. Data is assumed to arrive from some higher layer application and is stored in a buffer until it is transmitted. We study adapting the user's transmission rate and power based on the channel state information as well as the buffer occupancy; the objectives are to regulate both the long-term average transmission power and the average buffer delay incurred by the traffic. Two models for this situation are discussed; one corresponding to fixed-length/variable-rate codewords and one corresponding to variable-length codewords. The trade-off between the average delay and the average transmission power required for reliable communication is analyzed. A dynamic programming formulation is given to find all Pareto optimal power/delay operating points. We then quantify the behavior of this trade-off in the regime of asymptotically large delay. In this regime we characterize simple buffer control policies which exhibit optimal characteristics. Connections to the delay-limited capacity and the expected capacity of fading channels are also discussed.

Keywords

Fading channels, power control, resource allocation, wireless networks

I. Introduction

In mobile wireless networks, communication typically takes place over time-varying channels. This time-variation or fading is due to several effects such as variations in multi-path interference and shadowing. One technique to compensate for the channel's fading is to dynamically allocate communication resources, such as the transmission power or bit rate, based upon knowledge of the channel's state. Various methods for allocating transmission resources are part of most third-generation (3G) cellular standards (see e.g., [1]). These methods include adjusting the transmission power, changing the constellation size and coding rate, and varying the spreading gain in CDMA based systems. In this paper, we are concerned with such resource allocation

R. Berry is with the Department of Electrical and Computer Engineering, Northwestern University, 2145 Sheridan Rd., Evanston, IL 60208 USA, (e-mail: rberry@ece.nwu.edu).

R. Gallager is with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, 77 Massachusetts Ave. Rm 35-206, Cambridge, MA 02139 USA (e-mail: gallager@mit.edu).

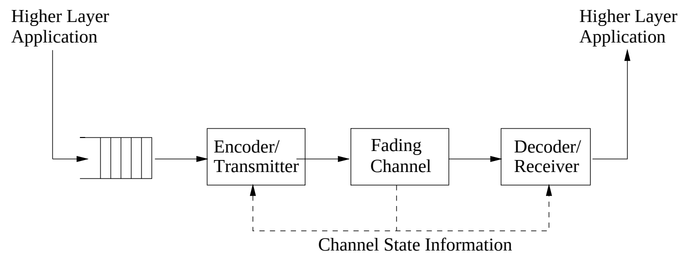

Fig. 1. System Model.

problems. Specifically, we consider the situation depicted in Figure 1. In this situation, data arrives from some higher layer application and is placed into a transmission buffer. Periodically the transmitter removes some of the data from the buffer, encodes it and transmits the encoded data over a fading channel. After sufficient information is received, the data is decoded and sent to a higher layer application at the receiver. We assume that the transmitter can allocate communication resources based on both the buffer occupancy and its knowledge of the channel.

In the above situation, we consider two conflicting objectives. One objective is to minimize the average transmission power required to reliably transmit the data. In a wireless network, mobile users often rely on a battery with a limited amount 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 the average delay incurred by the data. This objective can be viewed as arising from the Quality of Service (QoS) desired by the user. There is a clear trade-off between these objectives - transmitting at a higher rate requires more power but reduces the average delay. There are many aspects of the above description that need to be more precisely defined; this will be done in Section 2.

The delay experienced by data in the system of Fig. 1 is the sum of two components – the time spent in the buffer and the time from when data is encoded until it is decoded. The issue of reliably communicating data over a fading channel falls mainly within the province of information theory. Indeed, there has been much work in this area; see [2] for a recent survey. Information theoretic treatments typically either ignore delay completely or only consider the second component of delay. Buffer delay is usually considered a network layer problem and divorced from physical layer considerations. Generally it is assumed that when data leaves the buffer it is delivered with a fixed rate and fixed delay to the destination.

From a practical point of view, the above separation of physical layer coding delay and network layer buffer delay is very reasonable in a wired, point-to-point link. We give two reasons.[1] First, one can often send at rates near the information theoretic limits with an acceptable probability of error and with moderate delay relative to application requirements. For bursty traffic, the required coding delays are often on a much smaller time-scale than the traffic variations which are addressed by higher layer buffer management. Second, in a wired network, there is little reason to consider varying the transmission rate and power. The channel is typically not time-varying and users do not rely on a battery. Thus when transmitting, one should always transmit at the peak rate and power.

For the wireless situations we are interested in, neither of the above arguments need be true. With fading channels, varying the transmission power or coding rate can be useful in approaching capacity. Indeed, in many cases it is required. Additionally, approaching the capacity of a fading channel often requires the use of codewords long enough to "average over" the fading – the time required for this may be much longer than the acceptable delay. If such long codewords can not be used, then capacity may not be "meaningful". By this we mean that capacity does not give a good indication of the rate at which data can be sent with acceptable performance. This is the motivation behind the work on capacity vs. outage [5] and delay-limited capacity [6]. We look at these concepts in the next section and discuss their relation to the model in this paper. Regarding other related work, situations similar to that in Figure 1 have been looked at in [7] and [8], but not in the information theoretic context we take here.

The outline of the remainder of the paper is as follows. In Section II, a precise description of several models for the system in Fig. 1 is given. A model of the channel as well as two different models of the buffer dynamics are discussed. We also review several related capacity definitions for the channel model. In Section III, the trade-off between average power and average delay for these models is analyzed. We view this as a multi-objective optimization problem and give a Markov decision formulation for finding Pareto optimal solutions. The "optimal power/delay trade-off curve" for such problems is also characterized. In Section IV, this trade-off is analyzed in the asymptotic regime of large delay. In this regime the limiting required power is found and we provide the rate of convergence to this limit as a function of the average delay. Simple buffer control strategies are also given, which exhibit the optimal convergence rates. Section V contains some concluding remarks. Detailed proofs are given in the appendices.

[1] These arguments are for a single user channel. In the case of a multi-user channel, there is an additional coupling between delay and physical layer issues that arises in trying to allocate resources between many bursty users ([3], [4]).

II. Model and Problem Description

In this section we describe two different models for the situation in Fig. 1. In both cases we consider a block-fading model for the channel. This channel model is described next; we also review several notions of capacity for this channel, such as capacity vs. outage and delay-limited capacity. We then discuss two different approaches for modeling how the transmission rate and power are allocated over time. In the first approach we consider fixed-length, variable-rate codewords, while in the second approach we consider a fixed number of codewords with a variable length. Both of these models lead to buffer control problems that can be analyzed in a common framework.

A. Block-fading Channel

We consider a user communicating over a discrete-time, block-fading channel with additive Gaussian noise. This channel has been used to model a slowly-varying, flat-fading channel [5], [9] and is a generalization of the block interference channel introduced by McEliece and Stark [10]. In such a channel the transmitted signal is multiplied by a time-varying gain that models the fading. Over each block of N consecutive channel uses, the gain stays fixed. Let H_m denote the (baseband) complex channel gain during the mth block. Let $\mathbf{X}_m = (X_{m,1}, \ldots, X_{m,N})$ and $\mathbf{Y}_m = (Y_{m,1}, \ldots, Y_{m,N})$ be vectors in $\mathbb{C}^N$ which denote, respectively, the channel inputs and outputs over the mth block. These are related by:

$$\mathbf{Y}_m = H_m \mathbf{X}_m + \mathbf{Z}_m, \tag{1}$$

where the additive noise $\mathbf{Z}_m$ is a complex, circularly symmetric Gaussian random vector with zero mean and covariance matrix $\sigma^2 I$. Furthermore, the sequence $\{\mathbf{Z}_m\}$ is i.i.d. We assume that the sequence of channel gains, $\{H_n\}$, is a stationary ergodic Markov chain with state space $\mathcal{H}$. Conditioned on the current channel state, the next state, H_{m+1}, is independent of previous inputs and outputs, *i.e.* for all measurable $B \subset \mathcal{H}$, all $\mathbf{x}^m$, $\mathbf{y^m}$ and all $m \geq 1$,

$$\begin{aligned}&\Pr(H_{m+1} \in B | H^m{=}h^m, \mathbf{X}^m{=}\mathbf{x}^m, \mathbf{Y}^m{=}\mathbf{y}^m)\\ &\quad = \Pr(H_{m+1} \in B | H_m{=}h_m).\end{aligned}$$

Here we have denoted the sequence $(x_1, \ldots, x_m)$ by x^m. Let π_H denote the steady-state distribution of $\{H_m\}$ (by the above assumptions such a distribution exists and is unique). For technical reasons we also assume that $\mathcal{H}$ is a compact subset of $\mathbb{C}$.[2]

It is worth discussing the appropriateness of a such a model for a wireless channel. Clearly, if we intend to model a channel in which there are r channel uses per second, then N/r, the number of seconds per block, must be less than the coherence time of the

[2]This assumption is used in the proof of Lemma 4.3.

channel. Since we allow the fading process to have memory, N/r may be strictly less than the coherence time; for a memoryless fading process N/r must be approximately equal to the coherence time. If the underlying system we are modeling uses frequency hopping or TDMA, where the dwell time is N/r seconds, this is a good model for the channel variation. Otherwise, this model can be considered an approximation of a more physically motivated continuously-varying channel model as in [11]. A better approximation of such a model would be to choose $N = 1$ and account for all of the channel memory with the underlying Markov chain. We do not rule out such a choice of N in the above definition, and indeed for the model in Section II-C this assumption may be appropriate. For the model in Section II-B, having $N >> 1$ is more appropriate; using a block fading model also facilitates drawing connections with previous work on outage capacity and delay-limited capacity.

The assumption of flat fading is reasonable for a narrow-band system in which the bandwidth of a user is less than the channel's coherence bandwidth. The model we describe can easily be modified for a wide-band system with block-memoryless fading. Such a model would assume no ISI between blocks, but allow ISI within a block. This more general model would not provide any additional insights and would further complicate our notation, so we focus on the narrow-band case in the following.

Assume that both the transmitter and receiver have perfect CSI, meaning that during the mth block, both the transmitter and receiver know the value of H_m.[3] Several different notions of capacity appear in the literature that are applicable to the block-fading channels with perfect CSI. In the remainder of this section, we review these capacity definitions and discuss there significance for the problem at hand. Let C denote the solution to the following optimization problem:

$$\begin{aligned} &\underset{P:\mathcal{H}\mapsto\mathbb{R}^+}{\text{maximize}}\ \mathbb{E}_H \log\left(1 + \frac{|H|^2 P(H)}{\sigma^2}\right) \\ &\text{subject to: } \mathbb{E}_H P(H) \leq \bar{P}, \end{aligned} \tag{2}$$

where H is a random variable with the steady-state distribution π_H and $P : \mathcal{H} \mapsto \mathbb{R}^+$ is a power allocation, *i.e.* a function which indicates the average power used for each channel state $h \in \mathcal{H}$. In [12] a coding theorem and converse are proved showing that C is the capacity of this fading channel. We emphasize that in this case the capacity, C, has the "usual" operational significance that for any rate $R < C$, there exists a sequence of rate R codes of increasing block length such that the error probability goes to zero with increasing block length. This is to be contrasted with other notions of capacity defined in the following. The optimizing power allocation P in (2) is given

[3]This is clearly an idealized assumption which will be more appropriate the longer the channel's coherence time.

by

$$P(h) = \left(\frac{1}{\lambda} - \frac{\sigma^2}{|h|^2}\right)^+ \quad \text{for all } h \in \mathcal{H}, \tag{3}$$

where λ is a constant chosen so that the average power constraint is met. This is the well-known "water-filling" allocation over the channel state space [13]. It has been shown that C can be achieved by using either a "single-codebook, variable-power" transmission scheme [14] or a "multiplexed multi-rate, variable-power" scheme [12]. In either case approaching capacity requires one to use codewords long enough to take advantage of the ergodic properties of the fading process $\{H_m\}$. Delay constraints can prohibit the use of such long codewords, in which case this capacity does not provide a useful performance indication in the above operational sense.

While delay considerations may prohibit code-lengths long enough to average over the fading process, in many cases code-lengths are long enough for sufficient averaging of the additive noise. For example, suppose that each codeword must be sent in one block of N channel uses; in other words the delay constraint is less than the coherence time of the channel. If $N >> 1$, then reliable communication may still be possible during that block. In such situations, a *composite channel* model may be more appropriate.[4] Specifically, consider a family of channels, one channel corresponding to each possible realization of H_m. Assume that each of these channels occurs with the steady-state probability π_H. A codeword is then sent over one channel from the family; the channel staying fixed for the entire codeword. In this context, several notions of capacity have been defined, including capacity vs. outage, delay-limited capacity and expected capacity. Each of these notions of capacity is intended to operationally correspond to a different notion of rate. We define these quantities next.

In [5] the capacity versus outage probability ϵ or ϵ-capacity of the composite channel is defined to be the solution to the optimization problem

$$\begin{aligned} &\underset{P:\mathcal{H}\mapsto\mathbb{R}^+,R}{\text{maximize}} \quad R \\ &\text{subject to: } \Pr\left(\log\left(1+\frac{|H|^2P(H)}{\sigma^2}\right)\le R\right)\le\epsilon \\ &\qquad\qquad \mathbb{E}P(H)\le\bar{P}. \end{aligned} \tag{4}$$

The event $\log\left(1+\frac{|H|^2P(H)}{\sigma^2}\right) < R$ is referred to as an outage. Capacity versus outage probability ϵ, is the maximum mutual information rate that can be transmitted in every channel realization except a subset whose probability is less than ϵ. The capacity

[4]A composite channel is a compound channel where each sub-channel has an *a priori* probability associated with it. [15].

versus outage probability 0 is also referred to as the *delay-limited capacity* [6]. The delay-limited capacity can be shown to be given by [14]

$$\log\left(1+\frac{\bar{P}}{\mathbb{E}(1/|H|^2)}\right)$$

for any channel in which $\mathbb{E}(1/|H|^2)$ is finite; otherwise the delay-limited capacity will be zero.

Finally the *expected capacity* of the composite channel is defined to be the solution to the same optimization problem as in (2) above.[5] Now this quantity is given a different interpretation. A variable rate of mutual information (per codeword) is transmitted depending on the channel state. The expected capacity is the maximum expected rate.

The above capacities are all defined to be the maximum mutual information "rate" per codeword, where rate is interpreted differently in each case. For example, in the case of delay-limited capacity one is interested in the maximum constant rate per codeword; in the case of expected capacity, one is interested in the maximum expected rate per codeword. These quantities are intended to have the synonymous operational significance, that is they are meant to be the maximum "rate" for which there exists a sequence of block codes with that rate whose error probability goes to zero with increasing block length. To prove such a statement, a coding theorem and converse are needed. Recall we modeled the channel as a composite channel due to a delay constraint, N, which was assumed to be less than the coherence time of the channel. The usual type of converse via Fano's inequality holds with finite delay, *i.e.* for arbitrarily small probability of error, the rate must be less than the corresponding capacity. On the other hand, with finite delay, we can not get arbitrarily small probability of error and thus prove a coding theorem for the above capacity definitions. If we consider arbitrarily long codewords, then the assumption that N is less than the coherence time ceases to hold; thus the composite channel model is no longer appropriate. One way to prove an achievability result for these models, as in [14], is to consider the sequence of composite channels indexed by the block length $N = 1, 2, \ldots$. As N increases, it is assumed that the coherence time of the corresponding channel also increases. Letting $N \to \infty$ a coding theorem can be proved for the limiting channel. Of course, in the actual channel, the coherence time is fixed; thus this limiting operation has no physical significance, as opposed to the "usual" cases, such as a Gaussian channel without fading.

From a practical point of view, the above quantities can be useful if N is large enough relative to the block length required for reliable communication, but is still small relative to the coherence time of the channel. Again, by useful, we mean that these quantities give a good indication of the rates that are achievable with acceptably

[5]For the case where the transmitter has no CSI, approaching the expected capacity requires a broadcast coding strategy [15], [16]. With perfect CSI at the transmitter, a broadcast approach is not required.

small probability of error. If N is large enough, then a given probability of error can be achieved by transmitting at rates near the corresponding capacity. How large N must be depends on the error exponents for the composite channel.

The above ideas can be extended to delay constraints of more than one channel block *i.e.* more than one coherence time. This is done in [14] under the assumption that the transmitter has non-causal CSI for the entire channel realization over which each codeword is to be sent; in this case these ideas extend directly. Some discussion of the case where the transmitter has causal CSI is discussed in [17]; this situation is somewhat more problematic. In the following we will focus on the single block case, mainly to simplify notation.

We defined the above capacities as the maximum of a mutual information rate for a given power constraint. In the following it will be more useful to think of the inverse problem of finding the minimum power for a given rate of mutual information. We can define an analogous "power" formulation of both delay limited capacity and expected capacity for the block fading channel with delay constraint of one block. Corresponding to delay-limited capacity, the minimum power for rate R is given by:

$$\begin{aligned} &\underset{P:\mathcal{H}\mapsto\mathbb{R}^+}{\text{minimize}}\ \mathbb{E}P(H) \\ &\text{subject to: } \log\left(1+\frac{|h|^2 P(h)}{\sigma^2}\right) > R\ \forall h\in\mathcal{H} \end{aligned} \tag{5}$$

Likewise, corresponding to expected capacity, the minimum power for average rate R is given by:

$$\begin{aligned} &\underset{P:\mathcal{H}\mapsto\mathbb{R}^+}{\text{minimize}}\ \mathbb{E}_H P(H) \\ &\text{subject to: } \mathbb{E}_H \log\left(1+\frac{|H|^2 P(H)}{\sigma^2}\right) > R, \end{aligned} \tag{6}$$

where the solution to (6) corresponds to a water-filling power allocation. These quantities have an analogous interpretation to the corresponding capacities above. For a given delay constraint, N, they represent a lower bound on the required power to achieve arbitrarily small probability of error. Likewise as $N\to\infty$ these bounds are approachable.

The only difference between (5) and (6) is the constraint set for the minimization; this set corresponds to the particular mutual information rate of interest. In both cases, this set is determined by a requirement on the mutual information of a single codeword. Next, we will consider more complicated constraint sets, which depend on the entire sequence of codewords. These constraints will involve a buffer as in Fig. 1. Also note that both of the above formulations depend only on the steady-state distribution of the fading process; the memory in the fading process has no effect. This will no longer hold when we consider buffer constraints. Finally, we again emphasize that the above quantities are only meaningful if the time-scale of the delay constraint is small relative to the coherence time, but large relative to the error exponents of the component channels. In the next section we will consider models which allow for these assumptions to be relaxed in various degrees.

We will look at two different models of the situation depicted in Fig. 1. In both cases we consider a discrete-time model of the buffer where a time sample corresponds to a single block of N channel uses of a block fading channel. In the first case we will assume that all codewords are sent over the same number of channel uses, but that the rate, *i.e.* the number of possible codewords can vary. In the second case we assume a fixed number of codewords, but allow the number of channel uses over which a codeword is sent to vary. The first model is closely related to the composite channel models discussed above; we refer to this as the mutual information model. The second model is related to a model for multiple access communication introduced by Telatar in [4].

B. Mutual information model

As noted above, we use a discrete-time model of the buffer, where the time between adjacent samples corresponds to a block of N channel uses. Once again, assume that each codeword is sent in one block of N uses, and thus the length of time to send a codeword is less than the coherence time of the channel.[6] Let $\{A_n\}$ be an ergodic Markov chain with state space $\mathcal{A} \subset \mathbb{R}^+$ which represents the number of *bits* arriving at the buffer input between time $n-1$ and n. We assume that $\{A_n\}$ is independent of the channel fading and noise processes. Let $\bar{A} = \lim_{n\to\infty} \mathbb{E}A_n$ be the average arrival rate in steady-state. Assume that at the start of the nth block the transmitter removes U_n bits from the buffer and encodes these into a rate U_n/N code word which will be transmitted over the next N channel uses. Let S_n denote the buffer occupancy at the start of the nth block. The dynamics of the buffer are then given by

$$S_{n+1} = S_n + A_{n+1} - U_n. \tag{7}$$

This is illustrated in Fig. 2. Note as described above, the U_n bits to be transmitted are removed from the buffer before the next A_{n+1} bits arrive. Thus $U_n \leq S_n$, and $S_{n+1} > A_{n+1}$ for all n. We assume that the transmitter can choose U_n based on the buffer state S_n, the channel gain H_n, and the source state A_n.[7]

Let $P(h,u)$ be the required transmission power during a block when the channel gain is h and the transmitter chooses to transmit u bits. We assume that $P(h,u)$ is the power required so that the mutual information rate over the N channel uses

[6] Since a codeword is sent in one channel block, then clearly we must have $N >> 1$. This assumption may be relaxed, allowing for codewords that span $K > 1$ blocks. To do this requires a careful consideration of how one selects the rate of a codeword[17].

[7] More generally, U_n could be chosen based on the sequence of buffer, channel and source states up to time n, but for the Markov decision problem considered below there is no benefit in this.

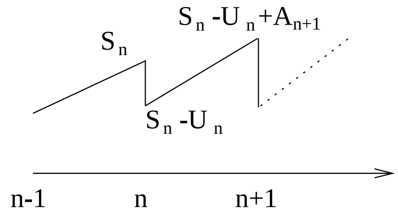

Fig. 2. Buffer dynamics.

is equal to u/N. We assume that the receiver knows the current buffer state at the transmitter, and thus knows the current transmission rate and power. Of course this requires some added overhead (unless the arrival rate is constant, in which case, the receiver can calculate the current buffer state). In this case we have

$$\log\left(1+\frac{|h|^2P(h,u)}{\sigma^2}\right)=u/N, \tag{8}$$

and thus

$$P(h,u)=\frac{\sigma^2}{|h|^2}(2^{u/N}-1). \tag{9}$$

For all h with $|h|>0$, $P(h,u)$ is an increasing and strictly convex function of $u\geq 0$. As with the composite channel model in the previous section, this model is sensible when N is large enough so that $P(h,u)$ is a reasonable indication of the power required to transmit at rate u/N with acceptable probability of error. For any N, $P(h,u)$ lower bounds the required power, for arbitrarily small probability of error. This bound is approachable as $N\to\infty$. The results in Sections 3 and 4 will only depend on the strict convexity and monotonicity of $P(h,u)$ and thus apply to any model that allows for a variable transmission rate and has a required power with these characteristics. For example, $P(h,u)$ could be the power required to transmit u bits for a particular modulation scheme such as the variable rate trellis coded M-QAM scheme in [18]. In this case, an approximation on the amount of transmitted power needed is given by

$$P(h,u)=\frac{\sigma^2}{|h|^2}\left(2^{\frac{u+2r}{N}}\right)K_c$$

where r is related to the rate of the convolutional coder used and K_c is a constant that depends on the coding gain and the required bit error rate. This is clearly a convex and increasing function of u. Another possibility is to let $P(h,u)$ be a bound on the power needed to transmit at a given rate over a fixed number of channel uses with a given probability of error. For example $P(h,u)$ could be derived from a random coding bound. This idea will be explored in more detail for the model described in the next section.

Recall we are interested in the average total delay[8] experienced by a bit in the system in Fig. 1. The total delay is the sum of the delay in the buffer plus the time from when a bit leaves the buffer until it is decoded. Once a bit leaves the buffer it is encoded into a codeword which takes 1 block of N channel uses to transmit. Assuming that a codeword is not decoded until it is entirely received, the second component of delay is $D_p + 1$ blocks, where D_p accounts for the propagation delay and processing time. We assume that this quantity is fixed for every codeword. From the above, the overall average delay is the average delay in the buffer plus $D_p + 1$. Thus we can ignore this constant factor and focus on the average delay in the buffer.

Let $\mathcal{S} = [0, \infty)$ denote the buffer state space[9]. Assume that the transmission rate, U_n at each time n is specified by a stationary policy, $\mu : \mathcal{S} \times \mathcal{H} \times \mathcal{A} \mapsto \mathbb{R}^+$ which specifies U_n as a function of the channel state H_n, the buffer occupancy S_n and the source state A_n. Under such a policy the sequence of combined buffer, channel, and source states, $\{(S_n, H_n, A_n)\}$ forms a Markov chain. The expected long term average power with such a policy is

$$\limsup_{m\to\infty} \frac{1}{m} \sum_{n=1}^{m} \mathbb{E}(P(H_n, \mu(S_n, H_n, A_n))). \tag{10}$$

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

$$\limsup_{m\to\infty} \frac{1}{m} \sum_{n=1}^{m} \frac{\mathbb{E}(S_n)}{\bar{A}}. \tag{11}$$

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

C. Telatar Model

Now we look at a different model of the situation in Fig. 1. In the previous model each codeword took a fixed amount of time to transmit, namely one block. The number of possible codewords per block varied according to the chosen rate. In this section we look at a model where one of a fixed number of codewords is chosen, but the length of time to transmit each codeword is variable. This can be considered a simple model of a hybrid ARQ situation [19].

We still consider a discrete-time model for the buffer, where each time slot corresponds to N channel uses of a block fading channel. As noted earlier, we do not need to assume that $N >> 1$ for the model in this section and indeed may assume

[8] Note we calculate delay for the discrete time model formulated above, if one assumes that this is a discretized model of a continuous time system, then this will upper bound the delay in the continuous time system.

[9] We allow the buffer to be an arbitrary real value. This is done primarily for mathematical convenience.

that $N = 1$. In the following, we develop a model in which the buffer occupancy corresponds to the *reliability* required by the data in the buffer plus the remaining reliability required by the data currently being transmitted. We assume that data arrives in fixed size packets of $\log M$ bits.[10] In this section we denote the number of packets that arrive between time $n-1$ and n by $\{\tilde{A}_n\}$, where once again $\{\tilde{A}_n\}$ is an ergodic Markov chain.[11] The transmitter takes a packet and encodes it into one of M codewords of infinite length and begins transmitting the message. While transmitting the message, the transmitter can adjust the transmission power by scaling the input symbol by an adjustable gain. Once the receiver can decode the message with acceptable probability of error, the transmitter stops transmitting the current codeword. The transmitter then proceeds to encode and transmit the next packet in the buffer.

We assume a random coding ensemble in which the codewords are chosen from a Gaussian ensemble. Each input symbol is chosen independently from a $\mathcal{N}(0,1)$ distribution. We allow the transmitter to adjust the transmission power at the start of each block. Let $\sqrt{P_i}$ be the gain used during the ith block. Thus the transmitted signal for each channel use during the ith block appears to be chosen from a $\mathcal{N}(0,P_i)$ distribution. As in the previous section, we assume that the receiver knows the current gain used by the transmitter. To model the amount of service time required by a codeword, we use a model derived from Telatar's model for multi-access communication in [4]. Specifically, if a given codeword is decoded after K blocks, there is the following random coding bound on the probability of error, for any $\rho \in (0,1]$:

$$P_e \leq \exp\left(\rho \ln M - N \sum_{i=1}^{K} E_o(\rho, |h_i|^2 P_i)\right), \tag{12}$$

where

$$E_o(\rho, |h_i|^2 P_i) = \rho \ln\left(1 + \frac{|h_i|^2 P_i}{\sigma^2(1+\rho)}\right) \tag{13}$$

and $\{h_i\}$ is the sequence of channel gains during the K blocks. Suppose there is a maximal allowable error probability of η. This error probability is achieved if the codeword is decoded after K blocks where

$$N \sum_{i=1}^{K} E_o(\rho, |h_i|^2 P_i) \geq \rho \ln M - \ln \eta \tag{14}$$

[10]This assumption is made primarily for mathematical convenience; if we allowed an arbitrary number of bits to arrive at each time, we would have to deal with the situation where fewer than $\log M$ bits remained in the buffer.

[11]The reason for this change in notation is that A_n will now be used to denote the amount of "reliability" required by the arriving packets.

for some fixed[12] $\rho \in (0, 1]$. Thus once (14) is satisfied, the transmitter can stop transmitting the current codeword. Since the transmitter has perfect CSI, it will know when this occurs. Without perfect CSI, some form of feedback from the receiver is needed to notify the transmitter when to stop transmitting. As in [4], $(\rho \ln M - \ln \eta)$ can be considered the demand of a codeword once it enters the encoder and $NE_o(\rho, |h_i|^2 P_i)$ as the service given to that codeword in the ith time step.

Let S_n be $(\rho \ln M - \ln \eta)$ times the number of packets in the buffer at time n plus the remaining amount of "service" required by the current codeword. We make the approximation that when a codeword receives its service, the next codeword immediately begins service. Practically, one would wait to begin transmitting the next codeword until the next channel use. If the typical service time of a codeword is many channel uses this effect will be small. With this approximation, the process $\{S_n\}$ evolves according to[13]:

$$S_{n+1} = S_n + A_{n+1} - U_n \tag{15}$$

where $A_n = (\rho \ln M - \ln \eta)\tilde{A}_n$ and $U_n = NE_o(\rho, |H_n|^2 P_n)$. We think of (15) as the dynamics of a new buffer with arrival process $\{A_n\}$ and departure process $\{U_n\}$. Once again this is a discrete time buffer where the buffer occupancy can take on any nonnegative real value.

As in the previous section, we assume that at each time n, the transmitter can choose U_n based on the current channel state, H_n, buffer state, S_n, and source state A_n. Since $U_n = NE_o(\rho, |H_n|^2 P_n)$, a given choice of $U_n = u$ when $H_n = h$ requires $P_n = \frac{\sigma^2(1+\rho)}{|h|^2}\left(e^{(\frac{u}{N\rho})} - 1\right)$. Motivated by this, we define $P(h, u)$ to be:

$$P(h, u) = \frac{\sigma^2(1+\rho)}{|h|^2}\left(e^{(\frac{u}{N\rho})} - 1\right). \tag{16}$$

As in the previous section we note that this is an increasing and strictly convex function of u for any h such that $|h| > 0$.

We again assume that the transmission rate U_n is specified by a stationary policy $\mu : \mathcal{S} \times \mathcal{H} \times \mathcal{A} \mapsto \mathbb{R}^+$. Under policy μ, the expected long term average power, $\bar{P}^\mu$, and the expected time average delay in the buffer, $\bar{D}^\mu$, are again given by (10) and (11) respectively.

This completes the description of the two models for the buffer dynamics. In this next section we begin an analysis of these models.

[12]The relation in (14) holds for any fixed $\rho \in (0, 1]$. One would naturally like to choose the ρ which is "optimum". For the Markov decision problem in the next section, this corresponds to the ρ which yields the minimum weighted combination of average delay and average power. Note varying ρ changes both the arrival process and the amount of energy needed; it is not clear that this optimization can be done analytically.

[13]As in the previous section note that $U_n \leq S_n$ and $S_{n+1} \geq A_{n+1}$ for all n.

III. Optimal Power/Delay Solutions

In the previous section we formulated two models for the situation in Fig. 1. These models have many characteristics in common. In both cases, we are interested in controlling a buffer with dynamics given by (7). At each time n, the transmission rate u is chosen based on the buffer occupancy S_n, the channel gain H_n and the arrival state A_n via a stationary policy μ. The sequences $\{A_n\}$ and $\{H_n\}$ are independent and both are stationary ergodic Markov chains. The power required to transmit at rate u when the channel gain is h is denoted by $P(h,u)$; this is a strictly convex and increasing function of u for all $h \in \mathcal{H}$. Finally, we are interested the trade-off between minimizing the average power and minimizing the average delay, as given by (10) and (11). In this section we will begin to characterize this trade-off. The following analysis will only rely on these general characteristics and thus applies to both of the previous models as well as any other model with these characteristics.

We are interested in two objectives, minimizing the average delay and minimizing the average power. Both of these criteria can not be minimized at the same time (except in the degenerate case, where the arrival rate and channel state are fixed for all time). Consider minimizing a weighted combination of the two criteria. Specifically, for $\beta > 0$, we seek to to find the policy μ which minimizes:[14]

$$\limsup_{m\to\infty} \frac{1}{m}\sum_{n=1}^{m} \mathbb{E}(P(H_n, \mu(S_n, H_n, A_n)) + \beta\frac{S_n}{\bar{A}}). \tag{17}$$

The weighting factor β indicates the relative importance of the buffer delay over the average power; larger values of β correspond to more placing more importance on delay. For the above models, the problem of finding the policy which minimizes (17) is an average cost Markov decision problem with state space $\mathcal{S} \times \mathcal{H} \times \mathcal{A}$. At each time step the transmitter chooses an action, namely the transmission rate, and incurs a per stage cost of $(P(H_n, \mu(S_n, H_n, A_n)) + \beta\frac{S_n}{\bar{A}})$. Such problems can be solved via dynamic programming techniques [20]. 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.[15]

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 (10) and (11). 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, $P^*(D^{\mu}) = \bar{P}^{\mu^*}$. 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

[14] The weighting factor β can be thought of as a Lagrange multiplier on an average delay constraint.

[15] One can also show several structural properties of optimal policies for this problem. We refer the reader to [17] for more details of this line of analysis.

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 [21]. 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.[16] 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 (17) for some choice of β. This also follows directly from the characterization of $P^*(D)$ given in the following proposition.

Proposition 3.1: 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.

Proof: 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 $\lambda \in [0, 1]$,

$$P^*(\lambda D^1 + (1-\lambda)D^2) \leq \lambda P^*(D^1) + (1-\lambda)P^*(D^2). \tag{18}$$

We will prove this using sample path arguments. 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 $P^*(D^2)$. As noted previously, $U_n^i(\omega) \leq S_n^i(\omega)$ for $i = 1, 2$, for all ω, and for all n. Now consider the new sequence of control actions, $\{U_n^\lambda(\omega)\}$, where for all n,

$$U_n^\lambda(\omega) = \lambda U_n^1(\omega) + (1-\lambda)U_n^2(\omega).$$

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 $S_{n+1}^i(\omega) = S_n^i(\omega) + A_{n+1}(\omega) - U_{n+1}^i(\omega)$ for $i = 1, 2$ and all $n \geq 0$, and recursion, we have

[16] Assume $\{(P^*(D), D) : D \geq 1\}$ is not the entire set of Pareto optimal solutions. From Prop. 3.1, 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.

$S_n^\lambda(\omega) = \lambda S_n^1(\omega) + (1-\lambda)S_n^2(\omega)$ for all n. Thus,

$$\lim_{m\to\infty} \frac{1}{m}\sum_{n=1}^{m} \frac{\mathbb{E}S_n^\lambda(\omega)}{\bar{A}} \leq \lambda D^1 + (1-\lambda)D^2, \tag{19}$$

where the expectation is taken over all sample paths. From the convexity of $P(h,u)$ in u, we have for all n

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

Again, summing and taking expectations we have

$$\begin{aligned} \lim_{m\to\infty} \frac{1}{m}\sum_{n=1}^{m} & \mathbb{E}P(H_n(\omega), U_n^\lambda(\omega)) \\ & \leq \lambda P^*(D^1) + (1-\lambda)P^*(D^2). \end{aligned} \tag{20}$$

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 $\mathcal{P}_d(a) = \mathbb{E}_H P(H,a)$ for any $a \geq 0$. For the model of Section II-B, $\mathcal{P}_d(a)$ corresponds to the solution of (5) with $R = a/N$. This is the minimum average power required to transmit at rate a in every channel state. As formulated, the delay in the buffer must be at least one time unit. The only way for the average delay to be equal to one is if $U_n = A_n$ for all n. Thus we have :

$$P^*(1) = \int_{\mathcal{A}} \mathcal{P}_d(a)\, d\pi_A(a). \tag{21}$$

For the mutual information model, if we have a constant arrival rate of $\bar{A}$ and an average delay constraint of 1, then the minimum average power is $\mathcal{P}_d(\bar{A})$, which corresponds to the delay-limited capacity of the channel. For channels whose delay-limited capacity is zero, $\mathcal{P}_d(\bar{A})$ must then be infinite for any $\bar{A} > 0$.

For any $a \geq 0$, define $\mathcal{P}_a(a)$ to be the solution to

$$\begin{aligned} \underset{\Psi:\mathcal{H}\mapsto\mathbb{R}^+}{\text{minimize}}\ & \mathbb{E}P(H,\Psi(H)) \\ \text{subject to :}\ & \mathbb{E}(\Psi(H)) \geq a. \end{aligned} \tag{22}$$

We have restricted Ψ to be a function only of the channel state H in this optimization. For the model of Section II-B, $\mathcal{P}_a(a)$ corresponds to the solution to (6), *i.e.* this corresponds to the expected capacity of the channel. Note that $\mathcal{P}_a(a)$ is an increasing

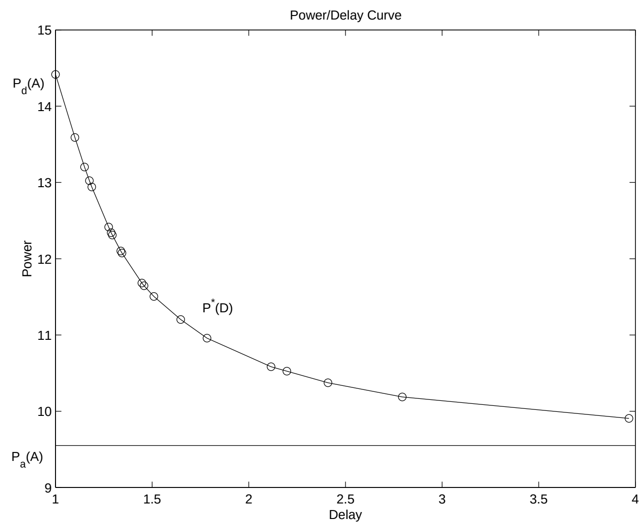

Fig. 3. Example of power/delay curve.

and strictly convex function of a; this follows directly from the strict convexity and monotonicity of $P(h,u)$. Let $\Psi^a(h)$ be a feasible rate allocation which achieves $\mathcal{P}_a(a)$. It can be seen that this rate allocation will be almost surely unique and is a function of only $|h|$. Furthermore, it is a continuous and non-decreasing function of $|h|$ for all $a > 0$. Likewise, for any $h \in \mathcal{H}$, $\Psi^a(h)$ is a continuous and non-decreasing function of 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. Assuming that the channel and arrival processes are not both constant, the only way a stationary policy μ can have $\bar{P}^\mu = \mathcal{P}_a(\bar{A})$ is if $\mu(s,h,a) = \Psi^{\bar{A}}(h)$ for all $(s,h,a) \in \mathcal{S} \times \mathcal{H} \times \mathcal{A}$, except possibly 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 $D \to \infty$.

Example: Figure 3 shows an example of the power/delay curve for a channel with memoryless fading and two states ($|\mathcal{H}| = 2$); in one state $|h|^2/\sigma^2 = 0.03$ and in the other state $|h|^2/\sigma^2 = 0.09$. For this example, the sequence of channel states is i.i.d. and each state is equally likely. The arrival process has a constant rate of $\bar{A}/N = 0.5$ and the power needed to transmit u bits is given by $P(h,u)$ in (9), corresponding to the mutual information model. To calculate the optimal policy, we discretized the buffer state space and allowable control actions. Using dynamic programming techniques, $P^*(D)$ was obtained computationally (within a small error margin) for various choices of β; the computed values of $P^*(D)$ are indicated in the figure. For this example $\mathcal{P}_d(\bar{A}) = 14.42$ and $\mathcal{P}_a(\bar{A}) = 9.55$. $\mathcal{P}_a(\bar{A})$ is indicated by a horizontal line in the figure.

IV. Asymptotic Analysis

In this section we characterize the behavior of the tail of the power/delay curve, $P^*(D)$, as the buffer delay $D \to \infty$. This corresponds to the solution of (17) as $\beta \to 0$. Throughout this section we restrict ourselves to the case of memoryless arrivals and memoryless fading, *i.e.* both $\{A_n\}$ and $\{H_n\}$ are sequences of i.i.d. random variables. This restriction is made primarily to simplify the following exposition. We also assume that the arrival state space $\mathcal{A}$ is a compact subset of $\mathbb{R}^+$. Let $A_{min} = \inf \mathcal{A}$ and $A_{max} = \sup \mathcal{A}$. With these assumptions, we show that $P^*(D) \to \mathcal{P}_a(\bar{A})$ as $D \to \infty$. We look at the rate[17] 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 in proving that the bound is tight we provide a sequence of policies with a relatively simple structure which exhibit the optimal rate of convergence. The approach in this section is closely related to Tse's work [22] on buffer control for variable-rate lossy compression. In [22] 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 $\bar{P}^{\mu_k} \to \mathcal{P}_a(\bar{A})$. Since in this section the arrival process is assumed to be memoryless, a stationary policy μ will only depend on the buffer state and the channel state, *i.e.* $\mu : \mathcal{S} \times \mathcal{H} \mapsto \mathbb{R}^+$. We restrict our attention to the class of policies that satisfy the following technical assumptions.

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 $s \leq 2\mathbb{E}(S^{\mu_k})$,

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

[17] The following standard 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 $a_n = \Omega(b_n)$.

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" [23]. 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 $\Pr(H=0) > \epsilon$, then this assumption must be satisfied by any policy that uses finite power. If this is not the case, then this is a restriction on the allowable policies.[18]

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.[19] Thus, this is not a very restrictive assumption.

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 $\{\mu_k\}$, the average drift over the tail of the buffer must be negative when k is large enough. This is stated in the following lemma.

Lemma 4.1: 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

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

Proof: Appendix A.

We use this result to establish the following bound on the rate of convergence:

Proposition 4.2: 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: Appendix B.

[18] 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.

[19] This follows from Lebesgue's theorem which states that a monotonic function is differentiable almost everywhere [24].

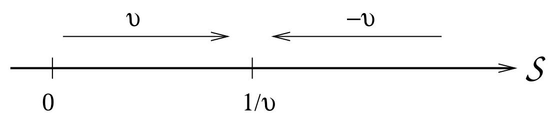

Fig. 4. A simple policy with drift υ.

In other words, the "tail" of $P^*(D)$ converges to $\mathcal{P}_a(\bar{A})$ at least as slowly as $\frac{1}{D^2}$. 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 $\bar{P}^{\mu_k} - \mathcal{P}_a(\bar{A}) = O((1/D^{\mu_k})^2)$. 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. We are still considering the case where the arrival process and the fading are memoryless.

Definition: For a given $\upsilon > 0$, partition the buffer state space into two distinct sets: $[1/\upsilon, \infty)$ and $[0, 1/\upsilon)$. Recall, $\Psi^a : \mathcal{H} \mapsto \mathbb{R}^+$ denotes the policy with average rate a which achieves $\mathcal{P}_a(a)$. Such a policy depends only on the channel state. Define a *simple policy with drift* υ, to be a policy μ with the form:[20]

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

In other words, with a simple policy the only dependency of the transmission rate on the buffer occupancy is through a simple threshold rule. Under 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 4. Lemma 4.3 below gives an upper bound on the average buffer delay under a simple policy. This bound depends on the semi-invariant moment generating function, $\gamma(r)$, of $A - \Psi^{\bar{A}+\upsilon}(H)$. This is defined as $\gamma(r) = \ln(\mathbb{E}[e^{(A-\Psi^{\bar{A}+\upsilon}(H))r}])$, where the expected value is taken with respect to both A and H. Since $\mathbb{E}A - \Psi^{\bar{A}+\upsilon}(H) < 0$, $\gamma(r)$ will have a unique positive r^* (where $r^* = \infty$ when no finite root exists) [25].

Lemma 4.3: For a simple policy μ with drift υ, the average delay satisfies:

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

[20] 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.

where $\eta(v)$ is a nonnegative function such that $\eta(v) \to 0$ as $v \to 0$, and $r^*(v)$ is the unique positive root of the semi-invariant moment generating function of $A - \Psi^{\bar{A}+v}(H)$.

The proof of this lemma can be found in Appendix C. There are two key ideas in this 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 process stays in $[1/v, \infty)$ it behaves as a random walk with a negative drift. Thus the steady-state probability that the buffer is in state s can be bounded by a function which decays exponentially, with an exponent given by $r^*(v)$. To show that simple policies have the desired convergence rate it is useful to characterize how $r^*(v)$ changes with v. The is given in the following lemma whose proof can be found in Appendix D.

Lemma 4.4: Let $r^*(v)$ denote the unique nonzero root of the semi-invariant moment generating function of $A - \Psi^{\bar{A}+v}(H)$ (for $v \neq 0$). Assume that for all v in a neighborhood of 0, that $\frac{d^2}{dv^2}\mathbb{E}e^{r^*(v)(A-\Psi^{\bar{A}+v}(H))}$ exists and that[21]

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

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

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

Using the above two lemmas it can be shown that a sequence of simple policies can achieve the bound given in Proposition 4.2.

Proposition 4.5: 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 $\bar{P}^{\mu_k} - \mathcal{P}_a(\bar{A}) = O((\frac{1}{D^{\mu_k}})^2)$.

Proof: Appendix E.

A simple policy as defined above requires splitting the buffer into two regions. In each region a policy was used that depended only on the current channel state. We have assumed that in addition to the current channel state, the receiver knows the current buffer state of the transmitter, so it would know the transmission rate and power used. 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

[21] 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 $\Psi^{\bar{A}+v}(h)$ with respect to v exists and is continuous at $v = 0$ for all h, then the above assumptions hold. When $P(h, u)$ corresponds to transmitting at capacity as in (9), this will be true for all but a finite number of values of $\bar{A}$. These values correspond to those rates $\bar{A}$ for which the "water level" $\frac{1}{\lambda}$ in some state h is exactly equal to $\frac{\sigma^2}{|h|^2}$ (*cf.* (3)).

simpler policy would be one with no dependence on the buffer state, *i.e.* a policy which only depended on the channel gain. With such a policy, the receiver would require no information about the transmitter's buffer state. Proposition 4.6 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 4.6, we will use a result similar to Lemma 4.4. However, we do not want to restrict the policy μ_k to be a policy of the form Ψ^x as in Lemma 4.4. 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

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

This is the same set of assumptions used in Lemma 4.4; by examining the proof of that lemma, it is apparent that the lemma also applies here. Specifically,

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

Any sequence of policies μ_k satisfying the above assumptions can not achieve the optimal convergence rate; this is stated in the following proposition.

Proposition 4.6: 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 $\bar{P}^{\mu_k} - \mathcal{P}_a(\bar{A}) = \Omega(\frac{1}{D^{\mu_k}})$.

Proof: Appendix F.

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.

V. Conclusions

In this paper we have looked at several simple models of communication over time-varying channels that incorporate buffer constraints. These models were chosen to illustrate the possible trade-offs between average power and average delay. To accomplish this we formulated a buffer control problem which was analyzed using ideas from Markov decision theory. We provided several characteristics of the optimal power/delay trade-off curve. In particular we characterized the asymptotic behavior of this trade-off in the regime of large buffer delay. In this asymptotic regime, we gave simple buffer control policies which exhibit the optimal convergence rate.

In conclusion we mention several directions in which this work can be extended. Instead of average delay, one can consider other network level quality of service indicators. For example with a finite buffer the probability of buffer overflow could be considered. If the arrival rate is constant, then the overflow probability corresponds to the probability of a maximum delay constraint being violated. Similar results can be shown in this setting. In this work we assumed that the transmitter has perfect channel state information. One can consider models that relax this assumption. Finally, we only considered single user channels. Models with multiple users can be considered. With more than one user, issues of allocating resources between users becomes important as does the coordination of the users.

Finally we mention some architectural issues related to this work. The problem formulation in this paper addresses issues which lie at the boundary of physical layer issues and higher layer network issues. From an architectural point of view there are many advantages to separating these layers. But as we have shown, in the context of mobile wireless communication it is not clear that the boundary between these layers should have the same characteristics as in a fixed wire-line network. One way to think about this is to ask what is a good "black box" abstraction for higher layers to have of the physical layer in such a network. In a wired point-to-point network, this abstraction is typically that the physical layer is a "packet pipe" that can deliver packets at a fixed rate, fixed delay, and some small probability of error. In wireless network, one has the potential to make a pipe with a variable rate, a variable delay and/or a variable probability or error. Furthermore one may even think of these as parameters which the next layer can adjust along with the transmission power. In this context, there are clearly many issues that extend beyond the simple models addressed here.

Appendices

A. Proof of Lemma 4.1

Proof: Let M, δ and ϵ be as in the definition of admissibility and assume that $k > M$. Let $F_n = A_n - U_{n-1}$; this represents the net change in the buffer occupancy between time $n-1$ and n. Thus, assuming the buffer is empty at time 0, we have

$$S_n = \sum_{m=1}^{n} F_m. \tag{23}$$

By assumption, the buffer process, S_n, reaches a steady state as $n \to \infty$. Thus the Markov inequality implies:

$$\lim_{n\to\infty} \Pr(S_n \geq 2\mathbb{E}(S^{\mu_k})) \leq \frac{1}{2}, \tag{24}$$

and so

$$\lim_{n\to\infty} \Pr(S_n < 2\mathbb{E}(S^{\mu_k})) > \frac{1}{2}. \tag{25}$$

Let $m = 4\mathbb{E}(S^{\mu_k})/\delta$ where δ divides $2\mathbb{E}(S^{\mu_k})$. Consider partitioning $[0, 2\mathbb{E}(S^{\mu_k}))$ into the following m segments: $[0,\delta/2), [\delta/2,\delta), \ldots, [(m-1)\delta/2, 2\mathbb{E}(S^{\mu_k}))$, where each segment has a length of $\delta/2$. Let $[(c-1)\delta/2, c\delta/2)$ be one of these segments which has the maximal probability with respect to π^{μ_k}. Thus,

$$\pi_S^{\mu_k}([(c-1)\delta/2, c\delta/2)) \geq \frac{1}{2m} = \frac{\delta}{8\mathbb{E}(S^{\mu_k})}. \tag{26}$$

Let $s_k = c\delta/2$ and define the process $\{\hat{S}_n\}$ by

$$\hat{S}_n = \max\{S_n, s_k\}. \tag{27}$$

Thus $\hat{S}_n$ is equal to S_n restricted to $[s_k, \infty)$. Let $\hat{F}_n = \hat{S}_n - \hat{S}_{n-1}$ be the net change in $\hat{S}_n$, so that

$$\hat{S}_n = \sum_{m=1}^{n} \hat{F}_m. \tag{28}$$

Thus

$$\lim_{n\to\infty} \mathbb{E}\frac{1}{n}\hat{S}_n = \lim_{n\to\infty} \mathbb{E}\frac{1}{n}\sum_{m=1}^{n} \hat{F}_m. \tag{29}$$

By assumption $\bar{D}^{\mu_k} < \infty$; therefore $\lim_{n\to\infty} \mathbb{E}S_n < \infty$. Furthermore, $\hat{S}_n \leq S_n + s_k$ for all n, which implies that $\mathbb{E}(\hat{S}_n) \leq \mathbb{E}(S_n) + s_k < \infty$. Thus,

$$\lim_{n\to\infty} \mathbb{E}\frac{1}{n}\hat{S}_n = 0. \tag{30}$$

The quantity $\hat{F}_n$ can be considered a reward gained at time $n-1$ by the original ergodic Markov chain $\{S_n\}$. Thus we have

$$\lim_{n\to\infty} \mathbb{E}\frac{1}{n}\sum_{m=1}^{n}\hat{F}_m = \int_{\mathcal{S}} \lim_{l\to\infty} \mathbb{E}(\hat{F}_l|S_{l-1}{=}s)\, d\pi_S^{\mu}(s). \tag{31}$$

Here $\lim_{l\to\infty} \mathbb{E}(\hat{F}_l|S_{l-1}{=}s)$ is the steady-state expected reward in state s. Using (30), (31), and (28) yields:

$$\int_{\mathcal{S}} \lim_{l\to\infty} \mathbb{E}(\hat{F}_l|S_{l-1}{=}s)\, d\pi_S^{\mu}(s) = 0. \tag{32}$$

Next we relate $\mathbb{E}(\hat{F}_l|S_{l-1}{=}s)$ to expected changes in the original process. We consider three cases:
1. First when $S_{l-1} \geq s_k$, then $\hat{F}_l \geq F_l$ and thus

$$\begin{aligned}&\lim_{l\to\infty} \mathbb{E}(\hat{F}_l|S_{l-1}{=}s)\\ &\geq \lim_{l\to\infty} \mathbb{E}(F_l|S_{l-1}{=}s) = \Delta^{\mu}(s), \quad \forall\ s \geq s_k.\end{aligned} \tag{33}$$

2. Next when $(c-1)\delta/2 \leq S_{l-1} < c\delta/2 = s_k$, $\hat{F}_l$ is nonnegative. Thus,

$$\mathbb{E}(\hat{F}_l|S_{l-1}{=}s) \geq \delta/2\,\Pr(\hat{F}_l > \delta/2|S_{l-1}{=}s) \tag{34}$$

$$\geq \delta/2\,\Pr(F_l > \delta|S_{l-1}{=}s). \tag{35}$$

Here (34) follows from the Markov inequality; (35) follows from the the fact that $\hat{F}_l \geq F_l - \delta/2$ for $(c-1)\delta/2 \leq S_{l-1} \leq c\delta/2$. Next taking the limit and using the admissibility of μ, we have:

$$\begin{aligned}\lim_{l\to\infty} \mathbb{E}(\hat{F}_l|S_{l-1}{=}s) &\geq \lim_{l\to\infty} \delta/2\,\Pr(F_l > \delta|S_{l-1}{=}s)\\ &\geq \frac{\epsilon\delta}{2}, \quad \forall\ s \in [(c-1)\delta/2, c\delta/2).\end{aligned} \tag{36}$$

3. Finally, when $S_{l-1} < (c-1)\delta/2$, $\hat{F}_l$ is also non-negative, and thus

$$\lim_{l\to\infty} \mathbb{E}(\hat{F}_l|S_{l-1}{=}s) \geq 0, \quad \forall\ s < (c-1)\delta/2. \tag{37}$$

Combining (33), (36), and (37) into (32) yields:

$$\int_{((c-1)\delta/2, c\delta/2]} \frac{\epsilon\delta}{2}\, d\pi_S^{\mu}(s) + \int_{s>s_k} \Delta^{\mu}(s)\, d\pi_S^{\mu}(s) \leq 0. \tag{38}$$

The first term can be bounded as follows using (26):

$$\int_{((c-1)\delta/2, c\delta/2]} \frac{\epsilon\delta}{2}\, d\pi_S^{\mu}(s) \geq \frac{\epsilon\delta^2}{16\mathbb{E}(S^{\mu_k})}. \tag{39}$$

Substituting this into (38) yields the desired result. ■

B. Proof of Proposition 4.2

Proof: For the kth policy, let $\Delta^{\mu_k}(s)$ denote the expected drift in state s. Thus 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 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:

$$\bar{P}^{\mu_k} \geq \int_{\mathcal{S}} \mathcal{P}_a(\bar{A} - \Delta^{\mu_k}(s))\, d\pi_S(s) \tag{40}$$

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

$$\mathcal{P}_a(x) = \mathcal{P}_a(\bar{A}) + \mathcal{P}_a'(\bar{A})(x - \bar{A}) + G(x - \bar{A}) \tag{41}$$

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 (40) yields:

$$\begin{aligned} \bar{P}^{\mu_k} - \mathcal{P}_a(\bar{A}) &\geq \mathcal{P}_a'(\bar{A}) \int_{\mathcal{S}} (-\Delta^{\mu_k}(s))\, d\pi_S(s) \\ &\qquad + \int_{\mathcal{S}} G(-\Delta^{\mu_k}(s))\, d\pi_S(s) \qquad (42) \\ &= \int_{\mathcal{S}} G(-\Delta^{\mu_k}(s))\, d\pi_S(s). \qquad (43) \end{aligned}$$

Here we have used that

$$\int_{\mathcal{S}} \Delta^{\mu_k}(s)\, d\pi_S(s) = 0 \tag{44}$$

for any policy μ_k which has $\mathbb{E}S < \infty$. 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

4.1 and assume that $k > M$ so that the lemma applies. Then we have

$$\begin{aligned}
&\bar{P}^{\mu_k} - \mathcal{P}_a(\bar{A}) \\
&\geq \int_{s>s_k} G\left(-\Delta^{\mu_k}(s)\right)\, d\pi_S(s) && (45)\\
&= \int_{s>s_k} G\left(-\Delta^{\mu_k}(s)\right)\, d\pi_S(s) + \pi_S([0, s_k))G(0) && (46)\\
&\geq G\left(\int_{s>s_k} -\Delta^{\mu_k}(s)\, d\pi_S(s) + \pi_S([0, s_k))G(0)\right) && (47)\\
&= G\left(\int_{s>s_k} -\Delta^{\mu_k}(s)\, d\pi_S(s)\right) && (48)\\
&\geq G\left(\frac{\epsilon\delta^2}{16\mathbb{E}S^{\mu_k}}\right). && (49)
\end{aligned}$$

In (45), (46) and (48) we have used that $G(x) \geq 0$ and $G(0) = 0$. Eq. (47) follows from Jensen's inequality and (49) follows from Lemma 4.1. Finally, expanding G in a Taylor series around 0, and using that $G'(0) = 0$ we have:

$$\bar{P}^{\mu_k} - \mathcal{P}_a(\bar{A}) \geq \frac{1}{2}G''(0)\left(\frac{\epsilon\delta^2}{16\mathbb{E}S^{\mu_k}}\right)^2 + o\left(\left(\frac{\epsilon\delta^2}{16\mathbb{E}S^{\mu_k}}\right)^2\right). \quad (50)$$

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

C. Proof of Lemma 4.3

Proof: From Little's law we have:

$$\bar{D}^{\mu} = \frac{\mathbb{E}(S)}{\bar{A}}, \quad (51)$$

where $\mathbb{E}(S)$ is the expected buffer occupancy in steady-state. This can be written as the integral of the complimentary distribution function of S, *i.e.*

$$\mathbb{E}(S) = \int_0^{\infty} \Pr(S > s)\, ds. \quad (52)$$

Upper bounding $\Pr(S > s)$ by 1 for $s \leq 1/\upsilon$, yields:

$$\mathbb{E}(S) \leq 1/\upsilon + \int_0^{\infty} \Pr(S > s + 1/\upsilon)\, ds. \quad (53)$$

For all $\upsilon \geq 0$, let

$$\eta(\upsilon) = \sup\{\Psi^{\bar{A}+\upsilon}(h) - \Psi^{\bar{A}-\upsilon}(h) : h \in \mathcal{H}\}.$$

We show that $\eta(\upsilon)$ is non-negative and converges to zero as $\upsilon \to 0$. As noted in Sect. III, $\Psi^a(h)$ is a continuous function of $|h|$ for all $a \geq 0$. Recall $\mathcal{H}$ is assumed to be compact; thus $\Psi^a(h)$ will be bounded for all a. Therefore, $\eta(\upsilon)$ is also bounded. Likewise, since $\Psi^a(h)$ is non-decreasing in a for all h, $\eta(\upsilon)$ will be non-negative. Finally, for all h, $\Psi^a(h)$ is continuous in a; thus, for all h, $\{\Psi^{\bar{A}+\upsilon}(h) - \Psi^{\bar{A}-\upsilon}(h)\}$ converges monotonically to 0 as $\upsilon \to 0$. Thus, by Dini's theorem [26], $\lim_{\upsilon\to 0} \eta(\upsilon) = 0$.

Next we bound $\Pr(S > s + 1/\upsilon)$. Consider a second buffer process $\{\breve{S}_n\}$ defined as follows. This second process only uses the policy $\Psi^{\bar{A}+\upsilon}$ and is restricted to stay in $[1/\upsilon, \infty)$ for all time. Specifically, let $\breve{U}_n = \Psi^{\bar{A}+\upsilon}(H_n)$ and let $\breve{S}_{n+1} = \max\{\breve{S}_n + A_{n+1} - \breve{U}_n, A_{n+1}, 1/\upsilon\}$. We assume that this buffer process and the original buffer process observe the same sequence of channel and source states. Furthermore assume that at time 0, $\breve{S}_0 = \max\{S_0, 1/\upsilon\}$. We claim that for all $n \geq 0$, $\breve{S}_n \geq S_n - \eta(\upsilon)$. This will be shown by induction on n. By assumption $\breve{S}_0 \geq S_0 \geq S_0 - \eta(\upsilon)$. Assume at time n, $\breve{S}_n \geq S_n - \eta(\upsilon)$, we will show that this holds for time $n+1$. Consider the following two cases:

Case 1: $S_n > 1/\upsilon$. In this case $\breve{U}_n = U_n$, and thus,

$$\begin{aligned}
\breve{S}_{n+1} &\geq \max\{\breve{S}_n - \breve{U}_n + A_{n+1}, A_{n+1}\} \\
&\geq \max\{S_n - \eta(\upsilon) - U_n + A_{n+1}, A_{n+1}\} \\
&\geq \max\{S_n - U_n + A_{n+1}, A_{n+1}\} - \eta(\upsilon) \\
&= S_{n+1} - \eta(\upsilon)
\end{aligned}$$

Case 2: $S_n \leq 1/\upsilon$. In this case $\breve{S}_n \geq 1/\upsilon \geq S_n$ and $\breve{U}_n \leq U_n + \eta(\upsilon)$. Thus

$$\begin{aligned}
\breve{S}_{n+1} &\geq \max\{\breve{S}_n - \breve{U}_n + A_{n+1}, A_{n+1}\} \\
&\geq \max\{S_n - (U_n + \eta(\upsilon)) + A_{n+1}, A_{n+1}\} \\
&\geq \max\{S_n - U_n + A_{n+1}, A_{n+1}\} - \eta(\upsilon) \\
&= S_{n+1} - \eta(\upsilon).
\end{aligned}$$

Thus we have $\breve{S}_n \geq S_n - \eta(\upsilon)$ for all $n \geq 0$. From this it follows that for all $n \geq 0$ and all s, $\Pr(S_n > 1/\upsilon + s) \leq \Pr(\breve{S}_n > 1/\upsilon + s - \eta(\upsilon))$. Letting $n \to \infty$ we have

$$\Pr(S > 1/\upsilon + s) \leq \Pr(\breve{S} > 1/\upsilon + s - \eta(\upsilon))$$

where S and $\breve{S}$ are random variables with the steady-state distributions for the respective processes. Note, the process $\{\breve{S}_n\}$ is a random walk restricted to $[1/\upsilon, \infty)$.

Therefore[22].

$$\Pr(\check{S} > 1/\upsilon + s - \eta(\upsilon)) \leq e^{-r^*(\upsilon)(s-\eta(\upsilon))}$$

and thus,

$$\Pr(S > 1/\upsilon + s) \leq e^{-r^*(\upsilon)(s-\eta(\upsilon))}.$$

Substituting this into (53) and carrying out the integration yields:

$$\mathbb{E}(S) \leq 1/\upsilon + \int_0^\infty e^{-r^*(\upsilon)(s-\eta(\upsilon))}\, ds \tag{54}$$

$$= 1/\upsilon + \frac{e^{r^*(\upsilon)\eta(\upsilon)}}{r^*(\upsilon)} \tag{55}$$

Finally, substituting this into (51) gives the desired result. ■

D. Proof of Lemma 4.4

Proof: From the definition of $r^*(\upsilon)$ we have, for all υ,

$$\mathbb{E}e^{r^*(\upsilon)(A-\Psi^{\bar{A}+\upsilon}(H))} = 1. \tag{56}$$

Differentiating this equation twice with respect to υ, and using the assumption in the lemma, we have, for all υ in a neighborhood of 0,

$$\mathbb{E}\frac{d^2}{d\upsilon^2}e^{r^*(\upsilon)(A-\Psi^{\bar{A}+\upsilon}(H))} = 0.$$

Letting $S(\upsilon) = \frac{dr^*(\upsilon)}{d\upsilon}$ then,

$$\begin{aligned}
&\mathbb{E}\frac{d^2}{d\upsilon^2}e^{r^*(\upsilon)(A-\Psi^{\bar{A}+\upsilon}(H))} \\
&= \mathbb{E}e^{r^*(\upsilon)(A-\Psi^{\bar{A}+\upsilon}(H))}\Bigg\{\Bigg((A-\Psi^{\bar{A}+\upsilon}(H))S(\upsilon) - r^*(\upsilon) \\
&\quad \cdot \left(\frac{d}{d\upsilon}\Psi^{\bar{A}+\upsilon}(H)\right)\Bigg)^2 + (A-\Psi^{\bar{A}+\upsilon}(H))\left(\frac{d}{d\upsilon}S(\upsilon)\right) \\
&\quad - 2S(\upsilon)\left(\frac{d}{d\upsilon}\Psi^{\bar{A}+\upsilon}(H)\right) - r^*(\upsilon)\left(\frac{d^2}{d\upsilon^2}\Psi^{\bar{A}+\upsilon}(H)\right)\Bigg\} \\
&= 0
\end{aligned}$$

[22]This inequality is referred to as the Kingman bound when applied to G/G/1 queues [25]

Next we evaluate this at $\upsilon = 0$. In doing this, note that for $\upsilon = 0$, the random variable $A - \Psi^{\bar{A}}(H)$ is zero mean, and thus $r^*(0) = 0$. Additionally note that since

$$\mathbb{E}\Psi^{\bar{A}+\upsilon}(H) = \bar{A} + \upsilon$$

then $\frac{d}{d\upsilon}\Psi^{\bar{A}+\upsilon}(H) = 1$ and $\frac{d^2}{d\upsilon^2}\Psi^{\bar{A}+\upsilon}(H) = 0$. Thus we have

$$S(0)^2\text{Var}(A - \Psi^{\bar{A}}(H)) - 2S(0) = 0. \tag{57}$$

This equation has two roots, corresponding to the two roots of $\ln(\mathbb{E}e^{r(A-\Psi^{\bar{A}}(H))}) = 0$. The root $S(0) = 0$ corresponds to the root of the log moment generating function that is always at zero, and the root at $\frac{2}{\text{Var}(A-\Psi^{\bar{A}}(H))}$ corresponds to the non-zero root, as desired. ■

E. Proof of Proposition 4.5

Proof: Let $\{\mu_k\}$ be a sequence of simple policies with drifts $\{\upsilon_k\}$ as in the statement of the proposition. 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 $\bar{D}^{\mu_k} = O(\frac{1}{\upsilon_k})$. From Lemma 4.3 we have

$$\bar{D}^{\mu_k} \leq \frac{1/\upsilon_k}{\bar{A}} + \frac{e^{r^*(\upsilon_k)\eta(\upsilon_k)}}{\bar{A}r^*(\upsilon_k)} \tag{58}$$

The first term on the right hand side of this bound is clearly $O(1/\upsilon_k)$. We focus on the second term of (58).

Taking the Taylor series of $r^*(\upsilon)$ around $\upsilon = 0$ and using Lemma 4.4 we have

$$r^*(\upsilon) = 0 + \Lambda\upsilon + o(|\upsilon|) \tag{59}$$

where $\Lambda = \frac{2}{\text{Var}(A-\Psi^{\bar{A}}(H))}$. Recall in Lemma 4.3 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

$$\frac{e^{r^*(\upsilon_k)\eta(\upsilon_k)}}{\bar{A}r^*(\upsilon_k)} = \frac{e^{\Lambda\eta(\upsilon_k)\upsilon_k + o(\upsilon_k)}}{\bar{A}(\Lambda\upsilon_k + o(\upsilon_k))}. \tag{60}$$

Now since:

$$\lim_{k\to\infty} \frac{\upsilon_k e^{\Lambda\eta(\upsilon_k)\upsilon_k + o(\upsilon_k)}}{\bar{A}(\Lambda\upsilon_k + o(\upsilon_k))} = \frac{1}{\bar{A}\Lambda} \tag{61}$$

it follows that:

$$\frac{e^{r^*(\upsilon_k)\eta(\upsilon_k)}}{\bar{A}r^*(\upsilon_k)} = O(1/\upsilon_k) \tag{62}$$

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

$$\begin{aligned}\bar{P}^{\mu_k} = \pi_S^{\mu_k}((1/\upsilon_k, \infty))\mathcal{P}_a(\bar{A} + \upsilon_k) \\ + \pi_S^{\mu_k}([0, 1/\upsilon_k])\mathcal{P}_a(\bar{A} - \upsilon_k)\end{aligned} \tag{63}$$

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

$$\begin{aligned}\bar{P}^{\mu_k} = \mathcal{P}_a(\bar{A}) + \mathcal{P}_a'(\bar{A})(\pi_S^{\mu_k}((1/\upsilon, \infty))\upsilon_k \\ - \pi_S^{\mu_k}([0, 1/\upsilon])\upsilon_k) + O((\upsilon_k)^2)\end{aligned} \tag{64}$$

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. ∎

I. F. Proof of Proposition 4.6

Proof: The average power under the kth policy, $\bar{P}^{\mu_k}$ is lower bounded by $\mathcal{P}_a(\bar{A} + \upsilon_k)$ thus

$$\bar{P}^{\mu_k} - \mathcal{P}_a(\bar{A}) \geq \mathcal{P}_a(\bar{A} + \upsilon_k) - \mathcal{P}_a(\bar{A}) \tag{65}$$

$$\geq \upsilon_k \mathcal{P}_a'(\bar{A}) \tag{66}$$

where the last step follows from the convexity of $\mathcal{P}_a$.

Now we show that $\mathbb{E}D^{\mu_k} = \Omega(\frac{1}{\upsilon_k})$. As in the proof of Lemma 4.3, using Little's law we have $\bar{D}^{\mu} = \mathbb{E}(S^{\mu_k})/\bar{A}$ where $\mathbb{E}S^{\mu_k}$ is the expected buffer occupancy in steady-state under policy μ_k. This can be written as

$$\mathbb{E}(S^{\mu_k}) = \int_0^{\infty} \Pr(S^{\mu_k} > s)\, ds$$

Since the transmission rate depends only on the channel state and the sequence of channel states are i.i.d., the buffer process is a random walk restricted to $[0, \infty)$. Therefore, $\Pr(S^{\mu_k} > s)$ can be lower bounded as follows:

$$\Pr(S^{\mu_k} > s) \geq e^{-r^*(\upsilon_k)(s - A_{max})}.$$

Here $r^*(\upsilon)$ is the unique nonzero root of the semi-invariant moment generating function of $A - \mu_k(H)$. A proof of this bound can be found in [13, Appendix 6B]. Thus

$$D^{\mu_k} \geq \frac{e^{r^* \upsilon_k A_{max}}}{\bar{A} r^* \upsilon_k}$$

By assumption, Lemma 4.4 still applies to $r^*(\upsilon)$. Thus we have $r^*(\upsilon_k) = \Lambda \upsilon_k + o(\upsilon_k)$ where $\Lambda = \frac{2}{Var(A - \Phi^A(H))}$. It follows that $D^{\mu_K} = \Omega(1/\upsilon)$. Combining this with the above bound for $\bar{P}^{\mu_k} - \mathcal{P}_a(\bar{A})$ we have $\bar{P}^{\mu_k} - \mathcal{P}_a(\bar{A}) = \Omega(1/D^{\mu_K})$ as desired. ∎

References

[1] S. Nanda, K. Balachandran, and S. Kumar, "Adaption techniques in wireless packet data services," *IEEE Communication Magazine*, vol. 38, pp. 54–64, January 2000.
[2] E. Biglieri, J. Proakis, and S. Shamai, "Fading Channels: Information-Theoretic and Communications Aspects," *IEEE Trans. Inf. Th. vol. 44, no. 6*, pp. 2619–2692, Oct 1998.
[3] R. Gallager, "A Perspective on Multiaccess Channels," *IEEE Trans. Inf. Th.*, vol. IT-31, pp. 124–142, Mar 1985.
[4] E. Telatar and R. Gallager, " Combining Queueing Theory with Information Theory for Multiaccess," *IEEE Journal on Selected Areas in Commun. Vol. 13, No. 6*, pp. 963–969, Aug. 1995.
[5] 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.
[6] S. Hanly and D. Tse, "Multi-access Fading Channels: Part II: Delay-Limited Capacities," *IEEE Transactions on Information Theory*, vol. 44, pp. 2816–2831, Nov. 1998.
[7] J. Rulnick and N. Bambos, "Mobile power management for wireless communication networks," *Wireless Networks*, vol. 3, pp. 3–14, 1997.
[8] 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.
[9] G. Caire, G. Taricco, and E. Biglieri, "Minimum Outage Probability for Slowly-Varying Fading Channels," *Proc. 1998 IEEE Int. Symp. Information Theory*, p. 7, Aug. 1998.
[10] R. McEliece and W. Stark, "Channels with Block Interference," *IEEE Trans. Inf. Th. vol. 30, no. 1*, pp. 44–53, Jan. 1984.
[11] W. Jakes, *Microwave Mobile Communications*. New York: John Wiley and Sons, 1974.
[12] A. Goldsmith and P. Varaiya, "Capacity of Fading Channels with Channel Side Information," *IEEE Trans. Inf. Th. vol. 43, no. 6*, pp. 1986–1992, Nov 1997.
[13] R. Gallager, *Information Theory and Reliable Communication*. New York: John Wiley and Sons, 1968.
[14] G. Caire, G. Taricco, and E. Biglieri, "Optimum power control over fading channels," *IEEE Trans. Inf. Th.*, vol. 45, pp. 1468–1489, July 1999.
[15] M. Effros and A. Goldsmith, "Capacity Definitions and Coding Strategies for General Channels with Receiver Side Information," in *1998 IEEE Int. Symp. Information Theory*, (Cambridge, MA), p. 39, Aug. 16-21 1998.
[16] P. Whiting and E. Yeh, "Optimal encoding over uncertain channels with decoding delay constraints," in *2000 IEEE Int. Symp. Information Theory*, (Sorrento, Italy), June 25-30 2000.
[17] R. Berry, *Power and Delay Trade-offs in Fading Channels*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, June 2000.
[18] A. Goldsmith, *Design and Performance of High-Speed Communication Systems over Time-Varying Radio Channels*. University of California at Berkeley: PhD dissertation Dept. Elec. Engin. Comput. Science, 1994.
[19] S. Lin and J. D.J. Costello, *Error Control Coding: Fundamentals and Applications*. Englewood Cliffs, NJ: Prentice-Hall, 1983.
[20] D. Bertsekas, *Dynamic Programing and Optimal Control, Vol. I and II*. Athena Scientific, 1995.
[21] Y. Sawaragi, *Theory of Multiobjective Optimization*. Orlando: Academic Press, 1985.
[22] 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.
[23] W. Feller, *An Introduction To Probability Theory and Its Applications*, vol. 2. New York: Wiley, 1957.
[24] F. Riesz and B. Nagy, *Functional Analysis*. New York: Ungar, 1955.
[25] R. Gallager, *Discrete Stocastic Processes*. Boston: Kluwer Academic Publishers, 1996.
[26] R. Dudley, *Real Analysis and Probability*. New York: Chapman and Hall, 1989.