

Energy-Efficient Scheduling of Delay Constrained Traffic over Fading Channels

Juyul Lee, *Student Member, IEEE*, and Nihar Jindal, *Member, IEEE*

Abstract—A delay-constrained scheduling problem for point-to-point communication is considered: a packet of B bits must be transmitted by a hard deadline of T slots over a time-varying channel. The transmitter/scheduler must determine how many bits to transmit, or equivalently how much energy to transmit with, during each time slot based on the current channel quality and the number of unserved bits, with the objective of minimizing expected total energy. In order to focus on the fundamental scheduling problem, it is assumed that no other packets are scheduled during this time period and no outage is allowed. Assuming transmission at capacity of the underlying Gaussian noise channel, a closed-form expression for the optimal scheduling policy is obtained for the case $T = 2$ via dynamic programming; for $T > 2$, the optimal policy can only be numerically determined. Thus, the focus of the work is on derivation of simple, near-optimal policies based on intuition from the $T = 2$ solution and the structure of the general problem. The proposed bit-allocation policies consist of a linear combination of a delay-associated term and an opportunistic (channel-aware) term. In addition, a variation of the problem in which the entire packet must be transmitted in a single slot is studied, and a channel-threshold policy is shown to be optimal.

Index Terms—Delay effects, dynamic programming, fading, scheduling.

I. INTRODUCTION

A Time-varying channel is a fundamental feature of wireless communication. In this context, opportunistic scheduling refers to the idea of transmitting with more power/higher rate when the channel quality is good and less power/lower rate when the channel is in a poor state. While this strategy is efficient from the perspective of long-term average rate, it is not necessarily appropriate for delay-constrained traffic which requires guaranteed short-term performance.

In this paper we consider the problem of transmitting a packet of B bits over T time slots, where the channel fades independently from slot to slot and the transmitter has perfect *causal* channel information (i.e., knowledge of the current channel, but not of the future channel). During each slot, the transmitter (or scheduler hereafter) determines how many bits to transmit based on the current channel quality and the number of bits yet to be served. The scheduler must balance the desire to be opportunistic, i.e., wait to serve many of the bits when the channel is in a good state, with the hard deadline. We investigate the setting where there is a single packet to be transmitted (i.e., no other packets are scheduled during the

T slot delay horizon), the packet must be transmitted by the deadline, and transmission occurs at capacity of the underlying Gaussian noise channel. In this framework our objective is to design a scheduling policy that minimize the expected energy consumed. This setup reasonably models delay-constrained applications such as VoIP, where packets arrive regularly and each must be received within a short delay window. In such a setting perhaps the most important design objective is to minimize the resources (in our case, energy) needed to meet the delay requirements. In the cellular uplink, for example, an energy-minimizing policy would extend the battery life of mobile terminals.

Delay constrained scheduling in wireless communication systems has been actively studied in various network settings under different traffic models and delay constraints (see for example [1][2][3][4][5][6][7][8] and references therein). In [1][2][3], power/rate control policies that minimize *average delay* are studied for a fading channel with random packet arrivals. In [4][5][6][7][8] systems with random packet arrivals, hard delay constraints, and general energy-rate relationships are studied, but the emphasis is on “offline” algorithms in which the scheduler has non-causal knowledge of the packet arrivals and the channel states; heuristic variations of the optimal “offline” algorithms are also proposed for the more challenging “online” (i.e., *causal*) setting.

In this paper, we rather focus on the interplay between fading, hard deadlines, and causal channel information by studying transmission of only a single packet, and thus do not consider random arrivals. Not only is this model more tractable, but it also more reasonably models applications with deterministic packet arrivals, e.g., VoIP or video streaming. To emphasize our treatment of physical-layer issues, we use the terms *causal* and *non-causal* rather than online and offline to indicate whether the scheduler has knowledge of future channel states. Recently, Fu et al. [9] considered this problem (single packet transmission over a block fading channel, subject to a hard deadline) and formulated it as a finite-horizon dynamic program (DP). For general energy-bit functions this DP can only be solved numerically, but in [9] a closed-form description of the optimal policy is derived for the special case where the energy-bit relationship is linear and the channel state is restricted to be an integer multiple of some constant. In this work we specialize the framework of [9] to the case where the energy-bit relationship is governed by the AWGN channel capacity formula, and derive closed-form descriptions of the optimal policy for $T = 2$ and sub-optimal policies for $T > 2$. In [10] the work of [9] is extended to a setting where the channel evolves according to a continuous Markov process, and the optimal scheduler is derived for the case where the energy-bit relationship is given by the AWGN capacity

Manuscript received January 9, 2008; revised July 21, 2008 and August 1, 2008; accepted August 5, 2008. The associate editor coordinating the review of this paper and approving it for publication was S. Hanly.

The authors are with the ECE Department, University of Minnesota (e-mail: {juyul, nihar}@umn.edu).

The work of J. Lee is supported by a Motorola Partnership in Research Grant. Initial results were presented at IEEE ISIT 2008.

Digital Object Identifier 10.1109/T-WC.2008.080037

formula under particular assumptions on the channel model (channels with *drift*). However, these results do not apply to the block fading model considered here and the policies are rather different in structure from those developed here.

In an earlier work, Negi and Cioffi [11] studied the dual problem of maximizing the expected number of transmitted bits in a finite number of slots subject to a finite energy constraint (with the energy-bit relationship described by the AWGN capacity formula). The optimal policy can generally only be found by numerical methods (although a threshold policy is found to be optimal at low SNR; in [12] this result is extended to multiple-access channels), and thus the solutions give little insight into how the scheduling parameters (e.g., channel state, number of bits to serve, number of slots remaining toward the deadline, and the like) affect the scheduling process. Although we deal primarily with suboptimal scheduling policies, we are able to deduce the effect of these parameters on the optimal policy.

A. Summary of Contribution

In this paper, we develop low-complexity and near-optimal scheduling policies for delay-constrained causal scheduling. Our main result is the following scheduler: a time-dependent weighted sum of a delay associated term and an opportunistic term as

$$b_t = \underbrace{\frac{1}{t}\beta_t}_{\text{delay associated}} + \underbrace{\frac{t-1}{t}\log\frac{g_t}{\eta_t}}_{\text{opportunistic}}, \quad (1)$$

where b_t is the number of bits to serve (from the remaining β_t bits) at time slot t (t is in descending order and thus represents the number of remaining slots), g_t denotes the current channel state, and η_t denotes a channel threshold determined by the channel statistics and the particular policy. If the current channel quality is equal to the threshold level, then a fraction $\frac{1}{t}$ of the remaining bits are transmitted. If the channel quality is better/worse than the threshold, then additional/fewer bits are transmitted. The scheduler acts very opportunistically when the deadline is far away (t large) but less so as the deadline approaches. The motivation of this form was raised from the simple $T = 2$ case, for which this form is shown to be optimal.

Two different suboptimal policies in the form of (1) are proposed, one through a simple extension of the optimal $T = 2$ scheduler and the other by solving a relaxed version of the optimization. Numerical results are presented to illustrate that these policies provide a significant advantage over a naive equal-bit policy, and that they perform quite close to the optimal for moderate/large values of B . In addition, we consider the case of one-shot allocation where the entire packet must be transmitted in only one of the slots. This is an optimal stopping problem, from which it follows that a simple channel threshold policy is optimal.

Notations: The operation $\mathbb{E}[X]$ for a random variable X denotes the expected value. The operation $\mathbb{G}[X]$ for a random variable X denotes $e^{\mathbb{E}[\ln X]}$ and the function $\mathbb{G}(x_1, \dots, x_m)$ for deterministic quantities x_1, \dots, x_m denotes the geometric mean $(\prod_{i=1}^m x_i)^{1/m}$. The operation $\langle \cdot \rangle_x^y$ denotes truncation from below at x and truncation from above at y . The function $\mathbb{1}_{\{\cdot\}}$ denotes the indicator function, i.e., its value is 1 if the argument is true and 0 otherwise. The sets \mathbb{R}_+ and \mathbb{R}_{++}

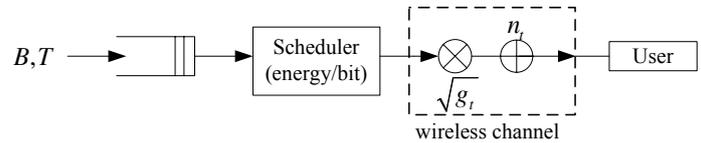


Fig. 1: Single-user delay constrained scheduling

denote the set of non-negative numbers and the set of positive numbers, respectively.

II. PROBLEM FORMULATION

We consider a single-user delay constrained scheduling problem as illustrated in Fig. 1: a packet of B bits must be transmitted within T time slots through a fading channel, in which T is referred to as the *delay-limit* or *deadline*. We assume no other packet is scheduled during the T time slots, and that the packet must be transmitted by the deadline (i.e., no outage is allowed). Although these two assumptions may not be entirely realistic, even for relatively deterministic traffic (e.g., in VoIP, the next packet generally arrives before the deadline of the previous has expired; furthermore, a small percentage of packets are allowed to miss their deadlines), these set of assumptions allow for a relatively tractable problem and allow us to focus on the central issue of meeting deadlines based upon causal channel information. The purpose of the scheduler is to determine the energy, or equivalently the number of bits, to be served during each time slot such that the expected energy is minimized and the bits are served by the deadline T .

Time is indexed in descending order, i.e., $t = T$ is the initial slot, $t = T - 1$ is the 2nd slot, \dots , and $t = 1$ is the final slot before the deadline; in doing so, t represents the number of remaining slots. The channel state, in power units, is denoted by g_t . We assume that the channel states $\{g_t\}_{t=1}^T$ are independently and identically distributed (i.i.d.) and the scheduler has causal knowledge of these channel states (i.e., at time t , g_T, g_{T-1}, \dots, g_t are known but g_{t-1}, \dots, g_1 are unknown). In this context, we refer to this type of scheduler as a *causal scheduler*. The channel state g is assumed to be a non-degenerate positive continuous random variable.

Assuming unit variance Gaussian additive noise and transmission at capacity, the number of transmitted bits, denoted as b_t , if E_t energy is used, is given by $b_t = \log_2(1 + g_t E_t)$. By solving for E_t we arrive at a formula for the energy cost in terms of the channel state g_t , and the number of bits¹ served b_t :

$$E_t(b_t, g_t) = \frac{2^{b_t} - 1}{g_t}. \quad (2)$$

We use β_t to denote the queue state; i.e., the remaining bits at time slot t . Then, β_t can be calculated recursively as $\beta_t = \beta_{t+1} - b_{t+1}$. Given this setup, a scheduler is a sequence of functions $\{b_t\}_{t=1}^T$ that maps from the remaining bits and

¹An implicit assumption is that each slot spans n channel symbols, for n reasonably large, and that powerful coding allows for transmission of $n b_t$ bits in the t -th slot. Thus, the quantity b_t should be thought of as the number of bits transmitted per channel symbol during the t -th scheduling slot.

the current channel state² to the number of bits served, i.e., $b_t : \mathbb{R}_+ \times \mathbb{R}_{++} \rightarrow [0, \beta_t]$. Then, the optimal energy-efficient scheduler is the set of scheduling functions $\{b_t^{\text{opt}}(\cdot, \cdot)\}_{t=1}^T$ that minimizes the total expected energy cost (summed over the T slots): i.e.,

$$\min_{b_T, \dots, b_1} \mathbb{E} \left[\sum_{t=1}^T E_t(b_t, g_t) \right] \quad (3)$$

subject to $\sum_{t=1}^T b_t = B$ and $b_t \geq 0$ for all t .

The optimization in (3) can be formulated sequentially (via dynamic programming) with the remaining bits β_t as a state variable that summarizes the bit allocation up until the previous time step.

$$b_t^{\text{opt}}(\beta_t, g_t) = \begin{cases} \arg \min_{0 \leq b_t \leq \beta_t} \left\{ E_t(b_t, g_t) + \mathbb{E} \left[\sum_{s=1}^{t-1} E_s(b_s, g_s) \middle| b_t \right] \right\}, & t = T, \dots, 2, \\ \beta_1, & t = 1. \end{cases} \quad (4)$$

This is the standard backward iteration: we first determine the optimal action at $t = 1$, then find the optimal policy at $t = 2$ by taking into account the optimal policy to be used at $t = 1$, and so forth. Since g_t is known but future channel states g_{t-1}, \dots, g_1 are unknown, the quantity E_t is not random but the future energy costs E_{t-1}, \dots, E_1 are random. Note also that the optimization (4) should be performed for all possible values of β_t and g_t . In other words, deriving the optimal scheduling function b_t^{opt} is equivalent to finding the optimal decision rule for all possible pairs (β_t, g_t) .

III. OPTIMAL SCHEDULING

In this section we attempt to derive the optimal (causal) scheduler using the conventional dynamic programming technique [13]. Unfortunately, an analytic expression is obtained only when $T = 2$ (besides the $T = 1$ trivial case). For $T > 2$, we discuss the difficulty in obtaining an analytic expression. When the scheduler has non-causal knowledge of the future channel states, however, deriving an optimal scheduler is possible; the optimal non-causal scheduler provides useful intuition and is derived in Appendix A. In the present section, however, we assume causal channel state information at the scheduler.

A. Optimal Scheduler for $T = 2$

In the final time slot ($t = 1$), the scheduler is required to transmit all β_1 unserved bits regardless of the channel state g_1 , due to the hard delay constraint. Thus, the energy cost is given by $E_1(\beta_1, g_1) = (2^{\beta_1} - 1)/g_1$ for all g_1 , and the expected cost to serve β_1 bits in the final slot is $\mathbb{E}_{g_1} [E_1(\beta_1, g_1)] = \mathbb{E} \left[\frac{1}{g} \right] (2^{\beta_1} - 1)$.

At $t = 2$, g_2 is known but g_1 is unknown. The scheduler needs to determine b_2 , based on g_2 and B , while balancing

²Because the channel states are assumed to be i.i.d., it is sufficient to make scheduling decisions based only on the current channel (while ignoring past channels). If channels are correlated across time slots, then the past and present channel should be used to compute the conditional distributed of future channel states and all expected future energy costs should be computed with respect to these conditional distributions.

the current energy cost (of serving b_2 bits in the current slot) and the *expected* future cost (of deferring $B - b_2$ bits to the last slot). Thus, the optimum scheduler is the solution to the following minimization:

$$b_2^{\text{opt}}(B, g_2) = \arg \min_{0 \leq b_2 \leq B} \left(\underbrace{\frac{1}{g_2} (2^{b_2} - 1)}_{\text{current energy cost}} + \underbrace{\mathbb{E} \left[\frac{1}{g_1} \right] (2^{B-b_2} - 1)}_{\text{expected future cost}} \right). \quad (5)$$

The objective function in (5) is convex, and therefore the minimizer is found by setting the derivative to zero while taking into account the constraints on b_2 :

$$b_2^{\text{opt}}(B, g_2) = \left\langle \frac{1}{2}B + \frac{1}{2} \log_2 (g_2 \nu_1) \right\rangle_0^B, \quad (6)$$

where $\nu_1 \triangleq \mathbb{E} [1/g]$ is a constant that depends only on the distribution of the channel state g (see Appendix B for the definition of constants ν_m for $m = 1, 2, \dots$). Note that this policy depends only on the unserved bits and the current channel state. This policy is only meaningful when ν_1 is finite; this rules out Rayleigh fading, in which case g is exponentially distributed and thus $\mathbb{E} [1/g]$ is not finite.

Notice that the optimal scheduling function (6) has two additive terms: (a) $\frac{1}{2}B$ corresponds to an equal distribution to time slots $t = 1$ and $t = 2$, and (b) $\frac{1}{2} \log_2 (g_2 \nu_1)$ associated with a measure of the channel quality at $t = 2$. That is, if the channel quality g_2 is bigger than a threshold $1/\nu_1$, then more bits are allocated than $\frac{1}{2}B$; if g_2 is smaller than the threshold then fewer bits are allocated and more bits are deferred to the final slot.

B. Optimal Scheduler for $T > 2$

From (4), the optimization that the scheduler solves at each time step is:

$$J_t^{\text{opt}}(\beta_t, g_t) = \begin{cases} \min_{0 \leq b_t \leq \beta_t} \left(\frac{2^{b_t} - 1}{g_t} + \bar{J}_{t-1}^{\text{opt}}(\beta_t - b_t) \right), & t \geq 2 \\ E_1(\beta_1, g_1), & t = 1, \end{cases} \quad (7)$$

where $\bar{J}_{t-1}^{\text{opt}}(\beta) = \mathbb{E}_g [J_{t-1}^{\text{opt}}(\beta, g)]$ denotes the *cost-to-go* function, which is the expected cost to serve β bits in $(t - 1)$ slots if the optimal control policy is used at each step. This is a one-dimensional convex optimization (pp. 87-88 in [14]) over b_t and the optimal solution satisfies (8) assuming $\bar{J}_{t-1}^{\text{opt}}$ is differentiable (pp. 254-255 in [15]), where $\arg_b \{ \cdot \}$ represents the solution³ of the argument equation.

When $t = 2$, the cost-to-go function $\bar{J}_1^{\text{opt}}(\beta) = (2^\beta - 1)\nu_1$ (as well as its derivative) takes on a very simple form and thus (8) can be solved in closed form as in (6). However, the same is not true for $t > 2$. Because the optimal policy for $t = 2$ is known, the cost-to-go $\bar{J}_2^{\text{opt}}(\beta)$ can be written in closed form. The derivative $(\bar{J}_2^{\text{opt}})'(\beta)$ can also be written in closed form but cannot be analytically inverted; thus, the optimal policy for $t = 3$ can only be written in the form of (8) with the

³Because of the convexity, the solution exists uniquely if it exists.

$$b_t^{\text{opt}}(\beta, g_t) = \begin{cases} 0, & g_t \leq \frac{\ln 2}{(\bar{J}_{t-1}^{\text{opt}})'(\beta)}, \\ \arg_b \left\{ \frac{2^b}{g_t} = \frac{1}{\ln 2} (\bar{J}_{t-1}^{\text{opt}})'(\beta - b) \right\}, & \frac{\ln 2}{(\bar{J}_{t-1}^{\text{opt}})'(\beta)} < g_t < \frac{2^\beta \ln 2}{(\bar{J}_{t-1}^{\text{opt}})'(0)}, \\ \beta, & g_t \geq \frac{2^\beta \ln 2}{(\bar{J}_{t-1}^{\text{opt}})'(0)}, \end{cases} \quad (8)$$

second condition given by the following fixed point equation:

$$\frac{2^{b_3}}{g_3} = 2^{\beta-b_3} \int_0^{\frac{2^{-(\beta-b_3)}}{\nu_1}} \nu_1 dF(x) + 2^{\frac{\beta-b_3}{2}} \nu_1^{\frac{1}{2}} \int_{\frac{2^{-(\beta-b_3)}}{\nu_1}}^{\frac{2^{\beta-b_3}}{\nu_1}} \left(\frac{1}{x}\right)^{\frac{1}{2}} dF(x) + 2^{\beta-b_3} \int_{\frac{2^{\beta-b_3}}{\nu_1}}^{\infty} \frac{1}{x} dF(x), \quad (9)$$

where F is the cumulative distribution function of the channel state g . As a result, no analytical characterization of $\bar{J}_3^{\text{opt}}(\beta)$ is possible, and thus neither $b_t^{\text{opt}}(\cdot, \cdot)$ nor $\bar{J}_t^{\text{opt}}(\beta)$ can be found in closed form for $t \geq 4$.

Alternately, we can numerically find the optimal scheduler by the discretization method [16]. However, large complexity and memory is required for sufficiently fine discretization. More importantly, this numerical method gives little insight on how the delay constraint and channel state affect the scheduling function.

IV. SUBOPTIMAL SCHEDULING POLICIES

Because the optimal scheduler cannot be written in closed form, it is of interest to develop suboptimal schedulers. The first scheduler is based on the intuition from the optimal $T = 2$ policy, and the second is found by solving a relaxed version of the optimization.

A. Suboptimal I Scheduler

If we compare the optimal causal scheduler for $T = 2$ (Section III-A) to the non-causal scheduler (see Appendix A), we can immediately notice that the optimal scheduler determines b_2^{opt} by inverse-waterfilling over channels g_2 and $1/\nu_1$, where the non-causal scheduler inverse waterfills over g_2 and the actual value of g_1 ⁴. This is because of the particularly simple form of the expected future cost. Although the expected future cost does not take on such a simple form for $T > 2$, we can get a suboptimal scheduler by simply applying this inverse-waterfilling at every time slot t . In other words, at time step t , perform inverse-waterfilling over the following t channels:

$$g_t, \underbrace{\frac{1}{\nu_1}, \dots, \frac{1}{\nu_1}}_{t-1}$$

to determine how many of the unserved β_t bits to serve now. We denote this bit allocation policy as $b_t^{(I)}$. Since $t-1$ of the

⁴When both g_2 and g_1 are known at $t = 2$, the optimal non-causal scheduling policy is given by $b_2^{\text{WF}}(B, g_2) = \left\langle \frac{1}{2}B + \frac{1}{2} \log_2 \left(\frac{g_2}{g_1} \right) \right\rangle_0^B$ from (31), in which ‘‘IWF’’ stands for *inverse waterfilling* (see Appendix A for detail).

t channels are equal, the inverse-waterfilling operation is very simple and the policy is given by

$$b_t^{(I)}(\beta_t, g_t) = \left\langle \frac{1}{t} \beta_t + \frac{t-1}{t} \log_2 \frac{g_t}{\eta_t^{(I)}} \right\rangle_0^{\beta_t}, \quad (10)$$

where $\eta_t^{(I)} = 1/\nu_1$ serves as the channel threshold. Notice that this threshold value depends only on the channel statistics and is constant with respect to t .

When the deadline is far away (large t), the first term in (10) is negligible and the bit allocation is almost completely dependent on the instantaneous channel quality. As the deadline approaches (t decreases toward 1), the weight of the channel-dependent second term decreases and the weight of the delay-associated first term increases.

B. Suboptimal II Scheduler

The inability to find a general analytic solution to the original optimization (7) is due to complications caused by the constraint $0 \leq b_t \leq \beta_t$ (for each t) in the dynamic optimization. However, if we relax this constraint (i.e., allow $b_t < 0$ and $b_t > \beta_t$ while maintaining the constraint $\sum_{t=1}^T b_t = B$) we can derive the optimal policy in closed form.

If we define the function L_t as below, then we can show inductively that L_t represents the cost-to-go function for the relaxed optimization:

$$L_t(\beta_t) = t 2^{\frac{\beta_t}{t}} \mathbb{G}(\nu_t, \nu_{t-1}, \dots, \nu_1) - t \nu_1 \quad (11)$$

where ν_1, ν_2, \dots are the fractional moments defined in Appendix B and $\mathbb{G}(\cdot)$ represents the geometric mean operation defined in Section I. When $t = 1$, (11) holds trivially. If we assume (11) holds for $t-1$, then the relaxed optimization for the next time step is given by

$$\min_{b_t} \left(\frac{2^{b_t} - 1}{g_t} + L_{t-1}(\beta_t - b_t) \right) \quad (12)$$

and the solution (i.e., the optimum scheduler for the relaxed problem) is found by setting the derivative of the objective to zero:

$$b_t = \frac{1}{t} \beta_t + \frac{t-1}{t} \log_2 (g_t \mathbb{G}(\nu_{t-1}, \dots, \nu_1)). \quad (13)$$

By plugging in the optimum value of b_t in (13) into (12) and taking expectation with respect to g_t , we reach (11). By truncating the policy in (13) at 0 and β_t we get a policy, referred to as Suboptimal II, for the original (un-relaxed) problem:

$$b_t^{(II)} = \left\langle \frac{1}{t} \beta_t + \frac{t-1}{t} \log_2 \frac{g_t}{\eta_t^{(II)}} \right\rangle_0^{\beta_t}, \quad (14)$$

where

$$\eta_t^{(\text{II})} = \frac{1}{\mathbb{G}(\nu_{t-1}, \nu_{t-2}, \dots, \nu_1)} \quad (15)$$

denotes the threshold that depends only on the statistics not the realizations.

C. Remarks on the Suboptimal Schedulers

From (10) and (14), we can see that the two schedulers have a very similar form with the only difference term being the threshold η_t . Based on the policy formulations, this subsection investigates the common and different characteristics of the suboptimal schedulers.

1) *General Framework*: The two schedulers can be cast into a single framework:

$$b_t(\beta_t, g_t) = \left\langle \frac{1}{t}\beta_t + \frac{t-1}{t} \log_2 \frac{g_t}{\eta_t} \right\rangle^{\beta_t}, \quad (16)$$

where η_t is the channel threshold determined by the individual algorithms. This simple allocation strategy reveals how the delay constraint affects the scheduling algorithms: at time step t serve a fraction $1/t$ of the remaining bits plus/minus a quantity that depends on the strength of the current channel compared to a channel threshold. If the current channel is good (i.e., g_t is bigger than the threshold η_t), additional bits are served (up to β_t), while fewer bits are served when the current channel is poorer than the threshold. Furthermore, note that when t is large (i.e., far from the deadline), the first term β_t/t is very small and the number of bits served is almost completely determined by the current channel conditions. This agrees with intuition that we should make aggressive, almost completely channel dependent (and deadline independent) decisions when the deadline is far away, while we should make more conservative (more deadline dependent, less channel dependent) decisions near the deadline (small t).

Using $\log_2 10 \approx 3$ we can rewrite the policy in dB units as:

$$b_t(\beta_t, g_t) \approx \left\langle \frac{1}{t}\beta_t + \left(\frac{t-1}{t} \right) \left(\frac{g_t^{\text{dB}} - \eta_t^{\text{dB}}}{3} \right) \right\rangle^{\beta_t}. \quad (17)$$

For large t , approximately one bit is allocated for every 3 dB by which the channel exceeds the threshold.

2) *Channel Thresholds*: The difference between the two policies is in the threshold values, which are illustrated in Fig. 2 for a particular channel distribution. Suboptimal I has a constant threshold $\eta_t^{(\text{I})} = 1/\nu_1$ for all t , whereas Suboptimal II has a threshold that increases with t (by Proposition 1). It is intuitive to use a larger threshold when the deadline is far away (large t), as the scheduler can be more selective because many different channels remain to be seen before the deadline is reached.

By using a constant threshold, Suboptimal I is not selective enough and transmits too many bits when the deadline is far away. To see this, consider the average number of bits transmitted in slot t (ignoring truncation):

$$\begin{aligned} \mathbb{E}_{g_t}[b_t(\beta_t, g_t)] &= \mathbb{E}_{g_t} \left[\frac{1}{t}\beta_t + \frac{t-1}{t} \log_2 \frac{g_t}{\eta_t} \right] \\ &= \frac{1}{t}\beta_t + \frac{t-1}{t} \mathbb{E} \left[\log_2 \frac{g_t}{\eta_t} \right]. \end{aligned} \quad (18)$$

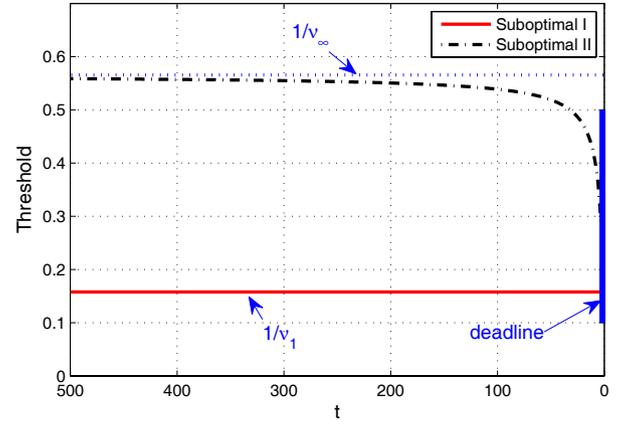


Fig. 2: Thresholds $\eta_t^{(\text{I})}$ for the suboptimal I scheduler and $\eta_t^{(\text{II})}$ for the suboptimal II scheduler when the channel state has the truncated exponential with $\gamma_0 = 0.001$.

Because $\eta_t^{(\text{I})} = 1/\nu_1 = 1/\mathbb{E}[1/g]$, by Jensen's inequality $\mathbb{E} \left[\log_2 \frac{g_t}{\eta_t^{(\text{I})}} \right] = \mathbb{E} [\log_2 g_t] + \log_2 \mathbb{E} [1/g] > 0$. Thus, Suboptimal I transmits more than $\frac{B}{T}$ bits on average when scheduling begins, which is in some sense overly aggressive. On the other hand, the quantity $\mathbb{E}_{g_t} \left[\log_2 \left(\frac{g_t}{\eta_t^{(\text{II})}} \right) \right]$ decreases as t increases and the limit is given by (see Proposition 1 in Appendix B)

$$\lim_{t \rightarrow \infty} \mathbb{E}_{g_t} \left[\log_2 \frac{g_t}{\eta_t^{(\text{II})}} \right] = 0. \quad (19)$$

This implies that Suboptimal II allocates B/T bits on the average when the deadline is far away and thus, unlike Suboptimal I, is not overly biased, nor overly aggressive. Numerical results given later support the fact that Suboptimal II generally performs better than Suboptimal I.

D. Equal-bit Scheduler

For comparison purposes, we consider one of the simplest causal schedulers: Equal-bit scheduler. This policy allocates B/T bits in each time slot, regardless of channel conditions, i.e.,

$$b_t^{\text{eq}}(\beta, g_t) = \frac{B}{T} = \frac{1}{t}\beta_t. \quad (20)$$

The corresponding expected energy is given by

$$\bar{J}_t^{\text{eq}}(\beta) = t(2^{\frac{\beta}{t}} - 1) \mathbb{E} \left[\frac{1}{g} \right] = t(2^{\frac{\beta}{t}} - 1)\nu_1. \quad (21)$$

Although equal-power scheduling is asymptotically optimal for the dual problem of maximizing rate over T slots when given a finite energy budget in the high power regime [11], it will be seen that equal-bit scheduling is suboptimal even when B is large.

E. Inverse Waterfilling Interpretation

If Suboptimal I and II and Equal-bit schedulers are compared to the optimal non-causal policy (inverse waterfilling), one can see that each of the algorithms mimics inverse waterfilling using either the current channel or channel statistics for the future channels, as summarized in Table I.

TABLE I: Waterfilling interpretation.

	At each t , perform IWF over the following
Equal-bit scheduler	$\underbrace{g_t, g_t, g_t, \dots, g_t}_{t-1}$
Suboptimal I scheduler	$g_t, \underbrace{\frac{1}{\nu_1}, \frac{1}{\nu_1}, \dots, \frac{1}{\nu_1}}_{t-1}$
Suboptimal II scheduler	$g_t, \frac{1}{\nu_{t-1}}, \frac{1}{\nu_{t-2}}, \dots, \frac{1}{\nu_1}$
Non-causal IWF	$g_t, g_{t-1}, g_{t-2}, \dots, g_1$

TABLE II: Average energy offsets for $T = 2$

distribution of channel state g	$\bar{J}_2^{\text{eq}}(B)/\bar{J}_2^{\text{opt}}(B)$	
	$B \rightarrow 0$	$B \rightarrow \infty$
truncated exponential with $\gamma_0 = 0.1$	1.96 dB	0.44 dB
truncated exponential with $\gamma_0 = 0.01$	3.26 dB	1.04 dB
truncated exponential with $\gamma_0 = 0.001$	4.32 dB	1.68 dB
1×2 Rayleigh fading ($g \sim \chi_2^2$)	1.99 dB	0.52 dB
1×3 Rayleigh fading ($g \sim \chi_3^2$)	1.37 dB	0.27 dB
1×4 Rayleigh fading ($g \sim \chi_4^2$)	1.10 dB	0.18 dB

V. ANALYSIS & NUMERICAL RESULTS

In this section, we compare the performance of the optimal, Suboptimal I and II, and Equal-bit schedulers. For $T = 2$ we are able to quantify the advantage of optimal scheduling relative to equal bit scheduling in two extreme cases, while for $T > 2$ we can only consider numerical results.

A. Asymptotic Analysis for $T = 2$

From the optimal scheduling expression for $T = 2$ given in (6), we can see that the packet is split over both time slots (i.e., $0 < b_2 < B$) if and only if $2^{-B}/\nu_1 < g_2 < 2^B/\nu_1$. As $B \rightarrow 0$, the probability of this event goes to zero: if $g_2 < 1/\nu_1$ then all bits are deferred to the final slot, while if $g_2 > 1/\nu_1$ all bits are served at $t = 2$. As a result, the expected energy cost takes on a rather simple form as $B \rightarrow 0$ (the derivation is provided in Appendix C):

$$\bar{J}_2^{\text{opt}}(B) \cong (2^B - 1) \mathbb{E} \left[\min \left(\frac{1}{g_2}, \nu_1 \right) \right], \quad (22)$$

where \cong represents equivalence in the limit (i.e., the ratio between both sides converges to 1 as $B \rightarrow 0$). This implies that the corresponding effective channel is $\max(g_2, 1/\nu_1)$. On the other hand, when $B \rightarrow \infty$ the probability of only utilizing one slot goes to zero and the limiting expected cost can be derived. The following theorem quantifies the power advantage of optimal scheduling:

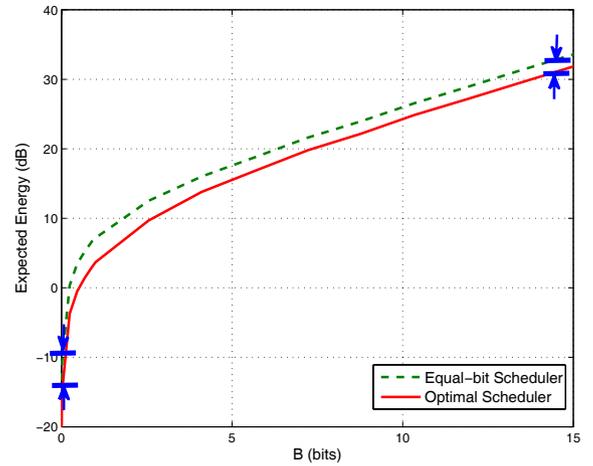
Theorem 1: The energy savings of optimal scheduling with respect to equal bit scheduling in extremes of $B \rightarrow 0$ and $B \rightarrow \infty$ is given by:

$$\lim_{B \rightarrow 0} \frac{\bar{J}_2^{\text{eq}}(B)}{\bar{J}_2^{\text{opt}}(B)} = \frac{\nu_1}{\mathbb{E} \left[\min \left(\frac{1}{g}, \nu_1 \right) \right]}, \quad (23)$$

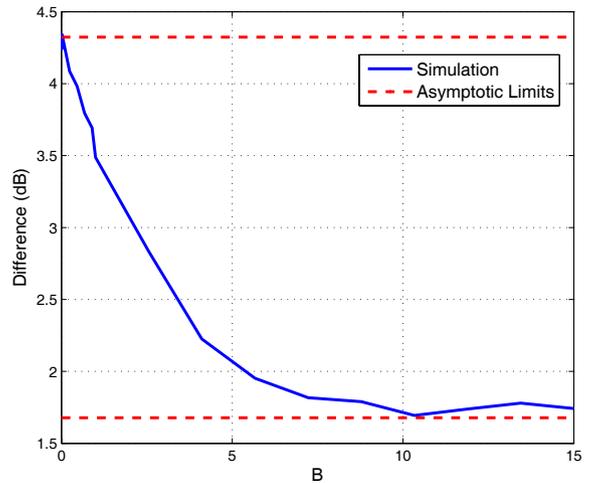
$$\lim_{B \rightarrow \infty} \frac{\bar{J}_2^{\text{eq}}(B)}{\bar{J}_2^{\text{opt}}(B)} = \sqrt{\frac{\nu_1}{\nu_2}}. \quad (24)$$

Proof: See Appendix C. ■

Table II summarizes typical values of the energy savings (at the extremes of $B \rightarrow 0$ and $B \rightarrow \infty$) for several fading distributions, as given by Theorem 1. As intuitively



(a) Average total energy



(b) Energy advantage of optimal relative to equal-bit scheduling (difference in dB)

Fig. 3: Average total energy consumptions for $T = 2$ and average energy offset when g is a truncated exponential variable with threshold $\gamma_0 = 0.001$

expected, the energy advantage is larger for more severe fading distributions. In other words, optimal scheduling is more beneficial in more severe fading environments.

Figure 3 contains a plot of expected energy versus B for the optimal and Equal-bit schedulers as well as a plot of the energy difference between the two schedulers as a function of B , for channel state g distributed as a truncated exponential with the threshold $\gamma_0 = 0.001$. The energy advantage is seen to decrease from its $B \rightarrow 0$ advantage of 4.32 dB to the large B asymptote of 1.68 dB.

B. Numerical Results for $T > 2$

Throughout the simulations, we assume that the channel state g_t is a truncated exponential with parameter $\lambda = 1$ and threshold $\gamma_0 = 0.001$. The fractional moments of this truncated

exponential variable can be calculated as:

$$\nu_m = \begin{cases} \lambda e^{\lambda\gamma_0} E_1(\lambda\gamma_0), & m = 1, \\ \lambda [e^{\lambda\gamma_0} \Gamma(\frac{m-1}{m}, \lambda\gamma_0)]^m, & m > 1, \end{cases}$$

where $E_1(\cdot)$ and $\Gamma(\cdot, \cdot)$ denote the exponential integral and the incomplete gamma function, respectively, and its limit is given by $\nu_\infty = \frac{1}{\gamma_0} e^{-e^{\lambda\gamma_0} E_1(\lambda\gamma_0)}$.

Figures 4a and 4b compare the energy consumption of the four different algorithms (Equal-bit, Suboptimal I and II, optimal causal) for $T = 5$ and $T = 50$, in which the optimal scheduler is calculated by numerical methods. The x -axis denotes the total number of bits B transmitted in T time slots, and thus B/T can be thought of as the average bits per channel use. The y -axis denotes the average total energy cost \bar{J}_T^{eq} , $\bar{J}_T^{(I)}$, $\bar{J}_T^{(II)}$, and \bar{J}_T^{opt} .

From Fig. 4a we see that both Suboptimal I and II perform nearly as well as the optimal scheduler, although Suboptimal II performs better than I. There are significant differences between Equal-bit and optimal schedulers, which is to be expected given the time diversity available over the five time slots. In Fig. 4b we see even larger differences between equal-bit and optimal causal, which can be explained by the even larger degree of time diversity ($T = 50$). Furthermore, Suboptimal II significantly outperforms Suboptimal I for $T = 50$ due to the over-aggressive nature of Suboptimal I. Suboptimal II performs nearly as well as the optimal scheduler when B is approximately 50 or larger (i.e., $B/T \geq 1$), but is sub-optimal for smaller values of B .

Figure 5 shows the expected bit allocation $\mathbb{E}[b_t]$ for the different algorithms for $T = 10$ slots when B is large ($B = 50$, upper) and small ($B = 2$, lower). While the optimal causal scheduling policy allocates roughly an equal number of bits (averaged across different realizations, and not for each particular realization) to each of the slots, Suboptimal I is immediately seen to allocate too many bits (on average) to early time slots which agrees with our earlier claim that this algorithm is often overly-aggressive as explained in Section IV-C2. For $B = 50$ the bit allocation of Suboptimal II is very similar to that of the optimal policy. However, for $B = 2$ Suboptimal II is also overly-aggressive as compared to the optimal. We suspect that the performance of Suboptimal II could be further improved by performing some heuristic modifications to the algorithm, but this is beyond the scope of the paper and is left to future work.

To summarize, the numerical results indicate that (a) Suboptimal II is nearly optimal for moderate to large values of B , (b) Suboptimal II outperforms Suboptimal I, and (c) neither suboptimal algorithm is near optimal for small values of B . In the next section, we will consider a policy that performs close to the optimal when B is small.

VI. ONE-SHOT ALLOCATION

In some settings it may be undesirable to split the packet across multiple time slots, e.g., because there is a large overhead associated with each slot used for transmission. In this scenario we may wish to find only one time slot among the T slots for the transmission of B bits; i.e., the action b_t can be either 0 or B .

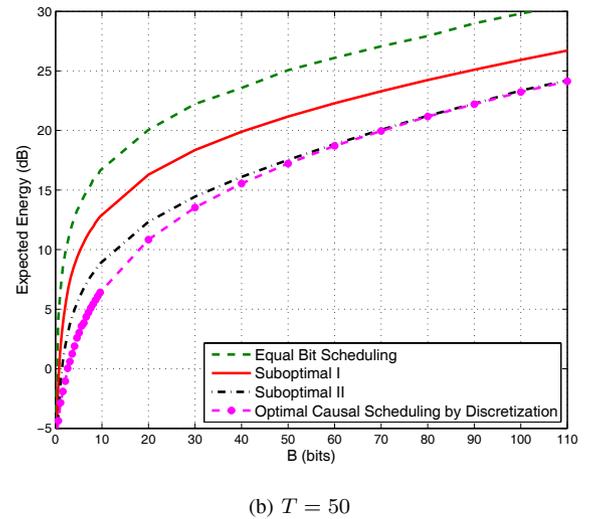
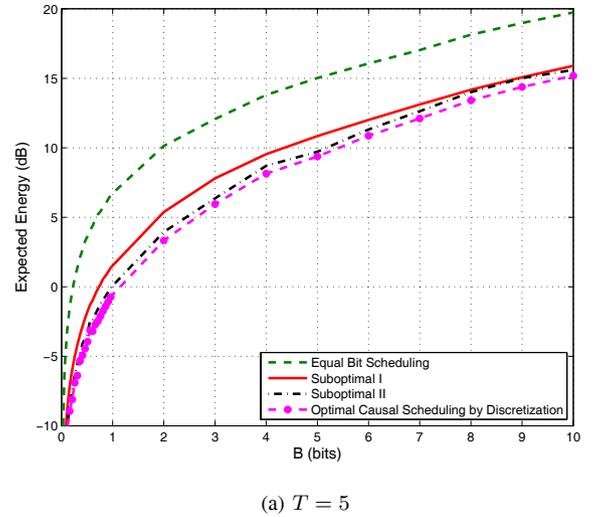


Fig. 4: Average total energy consumption for $T = 5$ and $T = 50$

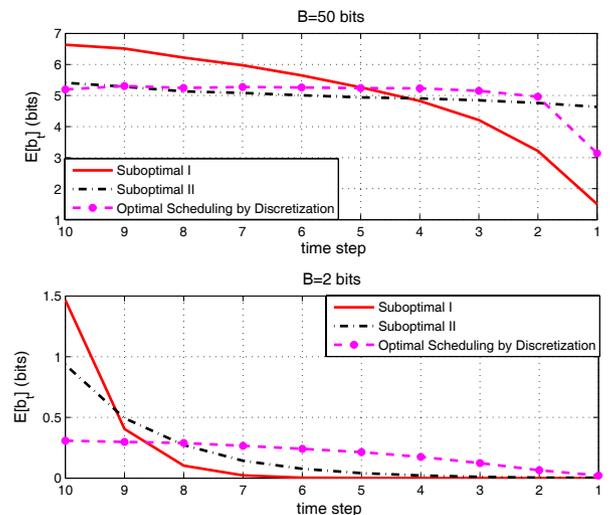


Fig. 5: Bit allocation profiles for $T = 10$ when $B = 50$ (upper) and $B = 2$ (lower)

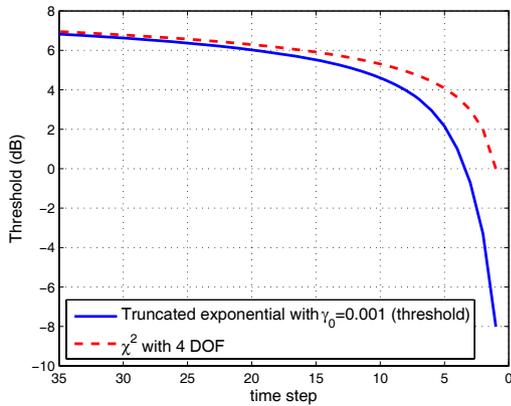


Fig. 6: Thresholds for the one-shot allocation

The dynamic program in this setting can be written as

$$\bar{J}(B, g_1) = \frac{2^B - 1}{g_1}, \quad (25)$$

$$J_t(B, g_t) = \min \left\{ \frac{2^B - 1}{g_t}, \bar{J}_{t-1}(B) \right\}, \quad t = 2, \dots, T, \quad (26)$$

where $\bar{J}_{t-1}(B) = \mathbb{E}_g[J_{t-1}(B, g)]$, which is precisely an optimal stopping problem [13]. Thus, a threshold policy is optimal: allocate all B bits at the first slot t such that $g_t > 1/\omega_t$, where $1/\omega_t$ is the threshold. That is,

$$b_t = \begin{cases} B, & t = \max \{s : g_s > 1/\omega_s\}, \\ 0, & \text{elsewhere.} \end{cases} \quad (27)$$

At $t = 1$ a packet must be served and thus ω_1 is infinite. Because the expected cost-to-go decreases as t increases, the threshold also decreases with t . In Appendix D we show the thresholds are given by the following recursive formula.

$$\omega_t = \begin{cases} \mathbb{E} \left[\frac{1}{g} \right], & t = 2, \\ \mathbb{E} \left[\frac{1}{g} \mid \frac{1}{g_t} < \omega_{t-1} \right] \Pr \left\{ \frac{1}{g_t} < \omega_{t-1} \right\} + \omega_{t-1} \Pr \left\{ \frac{1}{g_t} \geq \omega_{t-1} \right\}, & t = 3, \dots, T. \end{cases} \quad (28)$$

Notice that the threshold $1/\omega_t$ depends only on the channel statistics and does not depend on B .

Figure 6 illustrates the thresholds for the truncated exponential g (with $\lambda = 1$ and $\gamma_0 = 0.001$) and the chi-squared g (with 4 degrees of freedom). Figure 7 illustrates the energy usage (normalized by T) of the optimal one-shot allocation policy and the multiple slot policies. The $B/T = 0.1$ and $B/T = 1$ curves illustrate performance for relatively small and large values of B , respectively. When B is small, the energy of the one-shot allocation is nearly the same as the optimal policy that allows for multiple slots to be used. However, this one-shot allocation is not appropriate when B is relatively large because the required energy of the one-shot policy grows exponentially with B .

VII. CONCLUSION

In this paper we considered the problem of bit/energy allocation for transmission of a finite number of bits over a

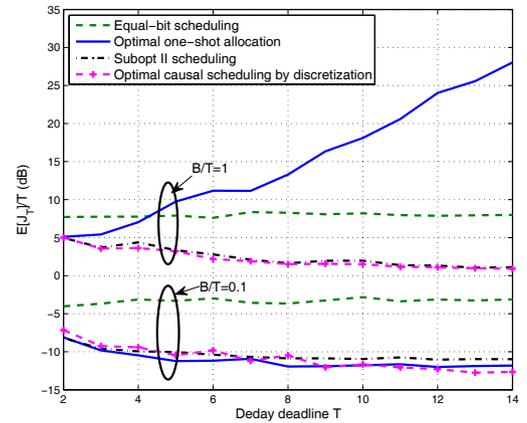


Fig. 7: Performance of the optimal one-shot allocation compared with multi-slot allocation algorithms

finite delay horizon, assuming perfect instantaneous channel state information is available to the transmitter and that the energy and rate are related by the Shannon-type (exponential) function. We derived the optimal scheduling policy when the deadline spans two time slots, and derived two near-optimal policies for general deadlines. The proposed schedulers have a simple and intuitive form that gives insight into the optimal balance between channel-awareness (i.e., opportunism) and deadline-awareness in a delay-limited setting. We also considered the same problem under the additional constraint that only a single of the available time slots can be used, and in this case found the optimal threshold-based policy. Based upon the policy constructions and the numerical results, we observed that the suboptimal II scheduler is near-optimal for large/moderate values of B while the one-shot policy is near-optimal for small values of B .

Given the increasing volume of delay-limited traffic over packet-switched wireless networks (e.g., VoIP or multimedia transmission in 3G systems), we expect problems of this sort to become increasingly important. Of course, the problem considered here represents only a particular instance of the rich space of delay-limited scheduling problems. Interesting extensions include consideration of discrete code rates, peak power constraints, and multi-user issues, and we hope this work provides useful insight for some of these other formulations.

APPENDIX A NON-CAUSAL SCHEDULING

If the channel states are known non-causally, i.e., g_T, \dots, g_1 are known at $t = T$, the optimal scheduling/allocation is determined by waterfilling because each time slot serves as a parallel channel. While conventional waterfilling maximizes rate subject to a power constraint, this is the dual of minimizing power/energy subject to a rate/bit constraint and is referred to as *inverse-waterfilling* (IWF):

$$J_T^{\text{IWF}}(B, \{g_t\}_{t=1}^T) = \min_{b_T, \dots, b_1} \sum_{t=1}^T \frac{2^{b_t} - 1}{g_t}, \quad (29)$$

subject to $\sum_{t=1}^T b_t = B$ and $b_t \geq 0$. This is a convex optimization problem and can be easily solved using the

standard Lagrangian method:

$$b_t^{\text{IWF}} = \left\langle \log_2 \left(\frac{g_t}{g_{\text{th}}} \right) \right\rangle_0^\infty, \quad (30)$$

where g_{th} is the solution to $\sum_{i=1}^T \left\langle \log_2 \left(\frac{g_i}{g_{\text{th}}} \right) \right\rangle_0^\infty = B$. A time slot t is called *utilized* if a positive bit is scheduled at t , i.e., $b_t > 0$ or equivalently $g_t > g_{\text{th}}$. With algebraic manipulations, we can express this IWF policy in (30) sequentially like other causal scheduling policies as

$$b_t^{\text{IWF}}(\beta_t, g_t) = \frac{1}{t'} \beta_t + \frac{t' - 1}{t'} \log_2 \frac{g_t}{\eta_t^{\text{IWF}}}, \quad \text{if } g_t > g_{\text{th}}, \quad (31)$$

otherwise $b_t^{\text{IWF}}(\beta_t, g_t) = 0$, where $t' = \sum_{i=1}^t 1_{\{g_i \geq g_{\text{th}}\}}$ and $\eta_t^{\text{IWF}} = \left(\prod_{i=1}^{t-1} g_i^{1_{\{g_i > g_{\text{th}}\}}} \right)^{1/(t'-1)}$. Notice that g_{t-1}, \dots, g_1 are relatively future quantities at slot t .

Like causal scheduling, the bit allocation process is described in two stages: first the remaining bits are divided equally amongst the utilized slots and then bits are added/subtracted depending on the channel state.

APPENDIX B

CHANNEL CHARACTERIZATION BY FRACTIONAL MOMENTS

We characterize the statistics of the channel states by using the fractional moments of the inverse of the channel states g . We define the following quantity for $m = 1, 2, \dots$,

$$\nu_m = \left(\mathbb{E} \left[\left(\frac{1}{g} \right)^{\frac{1}{m}} \right] \right)^m. \quad (32)$$

Then, the properties of these quantities are summarized as follows:

Proposition 1: The channel statistics defined according to (32) for a non-degenerate⁵ positive random variable have the following properties:

- (a) the sequence $\{\nu_m\}$ is strictly decreasing and the limit exists (denote the limit as ν_∞),
- (b) the sequence $\{(\nu_m \nu_{m-1} \dots \nu_1)^{1/m}\}$ is also strictly decreasing and its limit is also ν_∞ .

Proof:

- (a) First, we show the sequence $\{\nu_m\}$ is monotonically decreasing. Let $Y = 1/g$ and $f_Y(y)$ be the pdf of Y . By the Hölder's inequality [17],

$$\begin{aligned} \mathbb{E} \left[Y^{\frac{1}{m+1}} \right] &= \int_0^\infty y^{\frac{1}{m+1}} f_Y(y) dy \\ &= \int_0^\infty \left(y^{\frac{1}{m}} f_Y(y) \right)^{\frac{m}{m+1}} \left(f_Y(y) \right)^{\frac{1}{m+1}} dy \\ &< \left(\int_0^\infty y^{\frac{1}{m}} f_Y(y) dy \right)^{\frac{m}{m+1}} \left(\int_0^\infty f_Y(y) dy \right)^{\frac{1}{m+1}} \\ &= \left(\mathbb{E} \left[Y^{\frac{1}{m}} \right] \right)^{\frac{m}{m+1}}. \end{aligned}$$

The strict inequality is due to the fact that Y is not a point-mass density. Raising both sides to the power $(m+1)$ gives $\nu_{m+1} < \nu_m$.

⁵This eliminates a delta-type density (point-mass) function.

Second, we show convergence of the sequence. Let $\phi_m(y) = y^{\frac{1}{m}}$ for $y > 0$ and $\psi(y) = 1 + y$ for $y > 0$. Then, it is clear that $\lim_{m \rightarrow \infty} \phi_m(y) = 1$ for all $y > 0$, and $0 < \phi_m(y) \leq \psi(y)$ for all $y > 0$. Additionally, $\int_0^\infty \psi(y) f_Y(y) dy < \infty$. By the dominated convergence theorem [17], we have

$$\begin{aligned} \lim_{m \rightarrow \infty} \mathbb{E} \left[Y^{\frac{1}{m}} \right] &= \lim_{m \rightarrow \infty} \int_0^\infty \phi_m(y) f_Y(y) dy \\ &= \int_0^\infty 1 \cdot f_Y(y) dy = 1. \end{aligned}$$

Let x be a positive real number. By the continuity of the logarithmic function, we have $\lim_{x \rightarrow 0} \ln \mathbb{E} \left[Y^x \right] = 0$. By L'Hospital rule,

$$\lim_{x \rightarrow 0} \frac{\ln \mathbb{E} \left[Y^x \right]}{x} = \lim_{x \rightarrow 0} \frac{\mathbb{E} \left[Y^x \ln Y \right]}{\mathbb{E} \left[Y^x \right]} = \mathbb{E} \left[\ln Y \right]$$

since $\lim_{x \rightarrow 0} \mathbb{E} \left[Y^x \right] = 1$ and $\lim_{x \rightarrow 0} \mathbb{E} \left[Y^x \ln Y \right] = \mathbb{E} \left[\ln Y \right]$ (due to the dominated convergence theorem). By the continuity of the exponential function, $\lim_{x \rightarrow 0} e^{\frac{1}{x} \ln \mathbb{E} \left[Y^x \right]} = e^{\mathbb{E} \left[\ln Y \right]}$. Since the above limit exists and x is in the superset of integers, we have the result.

- (b) The monotonicity of the sequence $\{(\nu_m \nu_{m-1} \dots \nu_1)^{1/m}\}$ follows immediately from the monotonicity of the sequence $\{\nu_m\}$ and its positivity.

By the property of the exponential function, we have $(\nu_m \nu_{m-1} \dots \nu_1)^{\frac{1}{m}} = e^{\frac{1}{m} \ln(\nu_m \nu_{m-1} \dots \nu_1)} = e^{\frac{1}{m} \sum_{n=1}^m \ln \nu_n}$. Since $\lim_{m \rightarrow \infty} \nu_m = \nu_\infty$ and \log is continuous, $\lim_{m \rightarrow \infty} \ln \nu_m = \ln \nu_\infty$. By Cesàro mean, $\lim_{m \rightarrow \infty} \frac{1}{m} \sum_{n=1}^m \ln \nu_n = \ln \nu_\infty$. From the continuity of the exponential function, we have the result. ■

Notice that ν_1 and ν_∞ represent the arithmetic mean and the geometric mean of random variable $1/g$, respectively. All other values in the sequence $\{\nu_m\}$ lie between these two means.

APPENDIX C

PROOF OF THEOREM 1

For simple derivation, we work in units of nats rather than bits. From (6), the energy cost can be derived as

$$J_2^{\text{opt}}(g_2, B) = \begin{cases} (e^B - 1)\nu_1, & g_2 \leq \frac{e^{-B}}{\nu_1}, \\ 2e^{\frac{B}{2}} \left(\frac{1}{g_2} \nu_1 \right)^{1/2} - \frac{1}{g_2} - \nu_1, & \frac{e^{-B}}{\nu_1} < g_2 < \frac{e^B}{\nu_1}, \\ \frac{e^B - 1}{g_2}, & g_2 \geq \frac{e^B}{\nu_1}. \end{cases} \quad (33)$$

Thus,

$$\begin{aligned} \bar{J}_2^{\text{opt}}(B) &= \mathbb{E}_{g_2} \left[J_2^{\text{opt}}(g_2, B) \right] \\ &= \int_0^{\frac{e^{-B}}{\nu_1}} (e^B - 1)\nu_1 dF(x) \\ &\quad + \int_{\frac{e^{-B}}{\nu_1}}^{\frac{e^B}{\nu_1}} \left[2e^{\frac{B}{2}} \left(\frac{1}{x} \nu_1 \right)^{1/2} - \frac{1}{x} - \nu_1 \right] dF(x) \\ &\quad + \int_{\frac{e^B}{\nu_1}}^\infty \frac{e^B - 1}{x} dF(x), \end{aligned} \quad (34)$$

where F is the cumulative distribution function (CDF) of the channel state g .

By the limit rules,

$$\begin{aligned} & \lim_{B \rightarrow 0} \frac{\bar{J}_2^{\text{opt}}(B)}{e^B - 1} \\ &= \lim_{B \rightarrow 0} \frac{\left\{ \int_0^{\frac{e^{-B}}{\nu_1}} (e^B - 1) \nu_1 dF(x) + \int_{\frac{e^{-B}}{\nu_1}}^{\infty} \frac{e^B - 1}{x} dF(x) \right\}}{e^B - 1} \\ &= \lim_{B \rightarrow 0} \frac{\int_0^{\infty} (e^B - 1) \min\left(\frac{1}{x}, \nu_1\right) dF(x)}{e^B - 1} \\ &= \mathbb{E} \left[\min\left(\frac{1}{g}, \nu_1\right) \right] \end{aligned} \quad (35)$$

and

$$\lim_{B \rightarrow 0} \frac{e^B - 1}{2(e^{B/2} - 1)} = 1. \quad (36)$$

With (21), we obtain (23). Likewise,

$$\begin{aligned} & \lim_{B \rightarrow \infty} \frac{\bar{J}_2^{\text{opt}}(B)}{2e^{\frac{B}{2}} (\nu_2 \nu_1)^{1/2}} \\ &= \lim_{B \rightarrow \infty} \frac{\int_{\frac{e^{-B}}{\nu_1}}^{\frac{e^{-B}}{\nu_1}} \left[2e^{\frac{B}{2}} \left(\frac{1}{x} \nu_1\right)^{1/2} - \frac{1}{x} - \nu_1 \right] dF(x)}{2e^{\frac{B}{2}} (\nu_2 \nu_1)^{1/2}} \\ &= \lim_{B \rightarrow \infty} \frac{2e^{\frac{B}{2}} (\nu_2 \nu_1)^{1/2} - 2\nu_1}{2e^{\frac{B}{2}} (\nu_2 \nu_1)^{1/2}} = 1 \end{aligned} \quad (37)$$

and

$$\lim_{B \rightarrow \infty} \frac{\bar{J}_2^{\text{eq}}(B)}{2e^{\frac{B}{2}} \nu_1} = 1 \quad (38)$$

Thus, we have shown (24).

APPENDIX D DERIVATION OF (28)

From (26) the threshold ω_t is related to the expected cost-to-go by $\omega_t = \frac{1}{2^B - 1} \mathbb{E}[J_{t-1}(B, g)]$. The one-step cost-to-go is $\mathbb{E}[J_1(B, g)] = (2^B - 1) \mathbb{E}\left[\frac{1}{g}\right]$ and therefore $\omega_2 = \mathbb{E}\left[\frac{1}{g}\right]$. For $t > 2$, we expand the cost-to-go in terms of ω_{t-1} to give:

$$\begin{aligned} \omega_t &= \frac{1}{2^B - 1} \mathbb{E}[J_{t-1}(B, g)] \\ &= \frac{1}{2^B - 1} \left(\mathbb{E} \left[\frac{2^B - 1}{g_{t-1}} \middle| \frac{1}{g_{t-1}} < \omega_{t-1} \right] \Pr \left\{ \frac{1}{g_{t-1}} < \omega_{t-1} \right\} + \right. \\ &\quad \left. \mathbb{E} \left[\mathbb{E}[J_{t-2}(B, g)] \middle| \frac{1}{g_{t-1}} > \omega_{t-1} \right] \Pr \left\{ \frac{1}{g_{t-1}} > \omega_{t-1} \right\} \right) \end{aligned}$$

By substituting $\mathbb{E}[J_{t-2}(B, g)] = (2^B - 1)\omega_{t-1}$, we have the result.

REFERENCES

- [1] R. A. Berry and R. G. Gallager, "Communication over fading channels with delay constraints," *IEEE Trans. Inform. Theory*, vol. 48, no. 5, pp. 1135–1149, May 2002.
- [2] D. Rajan, A. Sabharwal, and B. Aazhang, "Delay-bounded packet scheduling of bursty traffic over wireless channels," *IEEE Trans. Inform. Theory*, vol. 50, no. 1, pp. 125–144, Jan. 2004.
- [3] B. E. Collins and R. L. Cruz, "Transmission policies for time varying channels with average delay constraints," in *Proc. 1999 Allerton Conf. on Commun., Control, & Comp.*, Monticello, IL, 1999.
- [4] B. Prabhakar, E. Uysal-Biyikoglu, and A. E. Gamal, "Energy-efficient transmission over a wireless link via lazy packet scheduling," in *Proc. IEEE INFOCOM*, Anchorage, AK, Apr. 2001, pp. 386–394.
- [5] E. Uysal-Biyikoglu and A. E. Gamal, "On adaptive transmission for energy efficient in wireless data networks," *IEEE Trans. Inform. Theory*, vol. 50, 2004.
- [6] M. J. Neely, "Optimal energy and delay tradeoffs for multiuser wireless downlinks," *IEEE Trans. Inform. Theory*, vol. 53, no. 9, pp. 3095–3113, Sept. 2007.
- [7] W. Chen, M. J. Neely, and U. Mitra, "Energy efficient scheduling with individual packet delay constraints: offline and online results," in *Proc. IEEE INFOCOM*, Anchorage, AK, May 2007, pp. 1136–1144.
- [8] —, "Energy efficient scheduling with individual delay constraints over a fading channel," in *Proc. WiOpt*, 2007.
- [9] A. Fu, E. Modiano, and J. N. Tsitsiklis, "Optimal transmission scheduling over a fading channel with energy and deadline constraints," *IEEE Trans. Wireless Commun.*, vol. 5, no. 3, pp. 630–641, Mar. 2006.
- [10] M. Zafer and E. Modiano, "Delay constrained energy efficient data transmission over a wireless fading channel," in *Proc. Workshop on Inf. Theory and Appl.*, La Jolla, CA, Jan./Feb. 2007, pp. 289–298.
- [11] R. Negi and J. M. Cioff, "Delay-constrained capacity with causal feedback," *IEEE Trans. Inform. Theory*, vol. 48, no. 9, pp. 2478–2494, Sept. 2002.
- [12] G. Caire, D. Tuninetti, and S. Verdú, "Variable-rate coding for slowly fading Gaussian multiple-access channels," *IEEE Trans. Inform. Theory*, vol. 50, no. 10, pp. 2271–2292, Oct. 2004.
- [13] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 3rd ed. Athena Scientific, 2005, vol. 1.
- [14] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, UK: Cambridge Univ. Press, 2004.
- [15] R. T. Rockafellar, *Convex Analysis*. Princeton Univ. Press, 1970.
- [16] D. P. Bertsekas, "Convergence of discretization procedures in dynamic programming," *IEEE Trans. Automat. Contr.*, vol. AC-20, no. 3, pp. 415–419, Jun. 1975.
- [17] W. Rudin, *Real and Complex Analysis*, 3rd ed. McGraw-Hill, 1987.



Juyul Lee (S'07) received the B.E. degree in electronic engineering from Hongik University, Korea, in 1996, and the M.S. degree in information and communications from Gwangju Institute of Science and Technology, Korea, in 1998.

From 1998 to 2000, he was a Research Engineer with the Agency for Defense Development, Korea. From 2000 to 2004, he was with the Electronics and Telecommunications Research Institute (ETRI), Korea, as a Member of the Technical Staff. Since 2004, he has been working toward the Ph.D. degree in electrical engineering at the University of Minnesota, MN. His research lies in the area of resource allocation in multi-user/multiple-antenna communications.



Nihar Jindal (S'99-M'04) received the B.S. degree in Electrical Engineering and Computer Science from U.C. Berkeley in 1999, and the M.S. and Ph.D. degrees in Electrical Engineering from Stanford University in 2001 and 2004. He is an assistant professor in the Department of Electrical and Computer Engineering at the University of Minnesota. His industry experience includes internships at Intel Corporation, Santa Clara, CA in 2000 and at Lucent Bell Labs, Holmdel, NJ in 2002. Dr. Jindal currently serves as an Associate Editor for IEEE

TRANSACTIONS ON COMMUNICATIONS, and was a guest editor for a special issue of the EURASIP JOURNAL ON WIRELESS COMMUNICATIONS AND NETWORKING on the topic of multiuser communication.

Dr. Jindal's research spans the fields of information theory and wireless communication, with specific interests in multiple-antenna/multi-user channels, dynamic resource allocation, and sensor and ad-hoc networks. Dr. Jindal was the recipient of the 2005 IEEE Communications Society and Information Theory Society Joint Paper Award, the University of Minnesota McKnight Land-Grant Professorship Award in 2007, and the NSF CAREER award in 2008.