Asymptotically Optimal Policies for Hard-deadline Scheduling over Fading Channels

Juyul Lee and Nihar Jindal

Department of Electrical and Computer Engineering
University of Minnesota
E-mail: {juyul, nihar}@umn.edu

***Abstract*—A hard-deadline, opportunistic scheduling problem in which B bits must be transmitted within T time-slots over a time-varying channel is studied: the transmitter must decide how many bits to serve in each slot based on knowledge of the current channel but without knowledge of the channel in future slots, with the objective of minimizing expected transmission energy. In order to focus on the effects of delay and fading, we assume that no other packets are scheduled simultaneously and no outage is considered. We also assume that the scheduler can transmit at capacity where the underlying noise channel is Gaussian such that the energy-bit relation is a Shannon-type exponential function. No closed form solution for the optimal policy is known for this problem, which is naturally formulated as a finite-horizon dynamic program, but three different policies are shown to be optimal in the limiting regimes where T is fixed and B is large, T is fixed and B is small, and where B and T are simultaneously taken to infinity. In addition, the advantage of optimal scheduling is quantified relative to a non-opportunistic (i.e., channel-blind) equal-bit policy.**

I. Introduction

Although the basic tenants of opportunistic communication over time-varying channels are well understood, much less is known when short-term delay constraints are imposed. Given the increasing importance of delay constrained communication, e.g., multimedia transmission, it is critical to understand how to optimize communication performance in delay-limited settings. Thereby motivated, we consider the discrete-time causal scheduling problem of transmitting a packet of B bits within a hard deadline of T slots over a time-varying channel. At each time slot the scheduler determines how many bits to transmit based on the current channel state information (CSI), but without future CSI, and the number of unserved bits, with the objective of minimizing the expected total energy cost. In order to focus on the interplay between opportunistic communication and delay, it is assumed that no other packets are simultaneously transmitted, and the hard deadline must always be met.

This basic problem was formulated as a finite-horizon dynamic program in [2], but an analytic form for the optimal scheduling policy cannot be found for most energy-bit relationships. Indeed, such a problem is difficult to solve because the transmitter only has *causal* CSI and because a particular rate must be guaranteed over a *finite* time-horizon. In our earlier work [3], we studied this problem in the setting where transmission occurs at the capacity of the underlying Gaussian noise channel and proposed different suboptimal scheduling policies.

Building upon [3], in this work we prove the *optimality* of certain scheduling policies in different asymptotic regimes. In particular, we show that:

- When the number of bits B is large, the optimal scheduling policy is a linear combination of a delay-associated term and an opportunistic-term. The opportunistic term depends on the logarithm of the channel quality, and the weight of this term decreases as the deadline approaches.
- When the number of bits B is small, a one-shot threshold policy where all B bits are transmitted in the first slot in which the channel quality is above a specified threshold is optimal.
- When the number of bits B and the time horizon T are both large, a waterfilling-like policy is optimal.

These results are particularly important in light of the fact that the general optimal solution appears intractable. In addition, the different asymptotically optimal schedulers provide an understanding of how the conflicting objectives of opportunistic communication (i.e. transmit only when the channel is strong) and delay-limited communication are optimally balanced, and how this balance depends on the time-horizon and the packet size.

In addition to showing asymptotic optimality, we also quantify the power benefits of optimal channel- and delay-aware scheduling relative to non-opportunistic equal-bit/rate transmission. These results identify that the largest benefits are obtained for severe fading, small packet size, and large time horizon.

The basic scheduling problem was first proposed and formulated as a finite-horizon dynamic program (DP) in [2]. In that work a closed-form solution for the optimal scheduler is provided for the special case where the number of transmitted bits is linear in the transmit energy/power and the channel quality is restricted to integer multiples of some constant. In [4], the formulation is extended to continuous time; closed-form descriptions of the optimal policies for some specific models are found, but these do not directly apply to the discrete-time problem considered here. In our earlier work

A longer version of this work, including detailed proofs, has been submitted to the *IEEE Transactions on Information Theory* [1]. The work of J. Lee is supported by a Motorola Partnership in Research Grant.

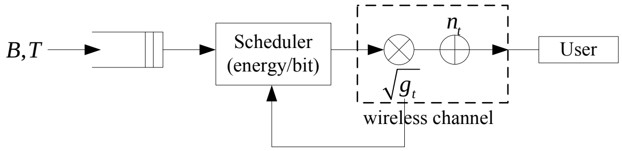

Fig. 1: Point-to-point delay constrained scheduling

[3], we specialized [2] to the setting where the energy-bit relationship is dictated by AWGN channel capacity and proposed several different suboptimal policies. Two of these policies are shown to be asymptotically optimal in the present work.

Because transmission scheduling corresponds to power allocation, it is also useful to put the present work in the context of prior work on optimal power allocation in fading channels, with and without delay constraints. In [5] it is established that waterfilling maximizes the long-term average transmitted rate; analogously, the long-term average power needed to achieve a particular long-term average rate is minimized by waterfilling. At the other extreme, channel inversion is known to be the optimal policy when a constant rate is desired in every fading state [6]. The current setting lies between these two extremes, because our objective is to find a power allocation policy (based on causal CSI) such that a particular rate (i.e. B/T) is guaranteed over T fading slots. The case $T = 1$ clearly corresponds to zero-outage/delay-limited capacity in [6], while we intuitively expect $T \to \infty$ to correspond to the long-term average rate scenario of [5]. The latter correspondence is made precise in Section IV-C.

II. Problem Setup

This section summarizes the scheduling problem introduced in [3], which is a discrete-time delay constrained scheduling problem over a wireless fading channel as illustrated in Fig. 1. A packet of B bits[1] is to be transmitted within a deadline of T slots. The scheduler determines the number of bits to allocate at each time slot using the fading realization/statistics to minimize the total expected transmit energy while satisfying the delay deadline constraint. We assume no other packets are to be scheduled simultaneously and that no outage is allowed.

The discrete-time slots are indexed by t in descending order (i.e., starting at $t = T$ down to $t = 1$), and thus t represents the number of remaining slots to the deadline. The channel state (at slot t) is denoted by g_t in power units. We assume that $g_T, g_{T-1}, \cdots, g_1$ are independently and identically distributed (i.i.d.) and the probability density function (PDF) and the cumulative distribution function (CDF) are denoted by f and F, respectively[2] The scheduler is assumed to have only *causal* knowledge of channel states (at time t, $g_T, \cdots, g_t$ are known but $g_{t-1}, \cdots, g_1$ are unknown). Assuming unit variance Gaussian additive noise and transmission at capacity, if energy E_t is used under channel state g_t, the number of transmitted bits is given by:

$$b_t = \log(1 + g_t E_t) \tag{1}$$

By inverting this formula, the required energy E_t to transmit b_t bits with channel state g_t is:

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

The queue state is denoted by β_t, which is the number of unserved bits at the beginning of slot t. Thus, the number of bits to allocate at slot t is determined by the queue state β_t and the channel state g_t. That is, a scheduling policy is a sequence of functions, indexed by the time step, that map from the current queue and channel state to the bit allocation: $\{b_T(\beta_T, g_T), b_{T-1}(\beta_{T-1}, g_{T-1}), \cdots, b_1(\beta_1, g_1)\}$. As for terminology, the entire set $\{b_T(\cdot,\cdot), b_{T-1}(\cdot,\cdot), \cdots, b_1(\cdot,\cdot)\}$ is referred to as a *policy* or a *scheduler*, and each element of it is referred to as a *policy function* or a *scheduling function*.

III. Optimal & Suboptimal Schedulers

In this section we describe the optimal scheduling policy, two suboptimal policies introduced in [3], and a heuristic modification of the ergodic (infinite-horizon) policy.

A. The Optimal Scheduler

The optimal scheduler for the hard-deadlined causal scheduling problem described in Section II can be found by solving the sequential optimization:

$$b_t^{\text{opt}}(\beta_t, g_t) = \begin{cases} \arg \min\limits_{0 \le b_t \le \beta_t} \left\{ E_t(b_t, g_t) + \mathbb{E}\left[\sum\limits_{s=1}^{t-1} E_s(b_s, g_s) \middle| b_t \right] \right\}, & \\ & t = T, \ldots, 2, \\ \beta_1, & t = 1. \end{cases} \tag{3}$$

where $\mathbb{E}$ denotes the expectation operator. Equivalently, this can be formulated as a finite-horizon dynamic program (DP):

$$J_t^{\text{opt}}(\beta_t, g_t) = \begin{cases} \min\limits_{0 \le b_t \le \beta_t} \left(\frac{e^{b_t}-1}{g_t} + \bar{J}_{t-1}^{\text{opt}}(\beta_t - b_t) \right), & t \ge 2 \\ \frac{e^{\beta_1}-1}{g_1}, & t = 1, \end{cases} \tag{4}$$

where $\bar{J}_{t-1}^{\text{opt}}(\beta) = \mathbb{E}_g[J_{t-1}^{\text{opt}}(\beta, g)]$ is the cost-to-go function, i.e., the expected cost to serve β bits in $t-1$ slots if the optimal policy is used.

At the final step ($t = 1$) all β_1 remaining bits must be served because outage is not allowed. At all other steps the optimal bit allocation is determined by balancing the current energy cost $\frac{e^{b_t}-1}{g_t}$ and the expected energy expenditure in future slots $\bar{J}_{t-1}^{\text{opt}}(\beta_t - b_t)$. Although the optimal scheduler can be found in closed form for $T = 2$ (Section III-A in [3]), it is not possible to do the same for $T > 2$ because no close-form expression for the cost-to-go function is known for $T \ge 2$. Nevertheless, the

[1] We operate in "nats" instead of "bits" since we adopt log-base e expression in the capacity formula to avoid constant factors in the analysis. We use "bits" and "bit allocation" as generic terms.

[2] The fading distribution must have a non-zero delay-limited capacity, i.e., $\mathbb{E}[1/g] < \infty$, for this problem to be feasible.

$$b_t^{\text{opt}}(\beta_t, g_t) = \begin{cases} 0, & g_t \le \frac{1}{(\bar{J}_{t-1}^{\text{opt}})'(\beta_t)}, \\ \arg_b \left\{ \frac{e^b}{g_t} = (\bar{J}_{t-1}^{\text{opt}})'(\beta_t - b) \right\}, & \frac{1}{(\bar{J}_{t-1}^{\text{opt}})'(\beta_t)} < g_t < \frac{e^\beta}{(\bar{J}_{t-1}^{\text{opt}})'(0)}, \\ \beta_t, & g_t \ge \frac{e^\beta}{(\bar{J}_{t-1}^{\text{opt}})'(0)}, \end{cases} \quad (5)$$

optimal scheduling functions can be described as (5), where $\arg_b\{\cdot\}$ represents the solution[3] of the argument equation [3]. The differentiability of $\bar{J}_{t-1}^{\text{opt}}$ can be verified by the properties of convexity and infimal convolution (pp. 254-255 in [7]).

B. The Boundary-relaxed Scheduler

The first suboptimal scheduler is derived by relaxing the boundary constraints (we no longer require $0 \le b_t \le \beta_t$), while maintaining the deadline constraint $\sum_{t=1}^T b_t = B$. The relaxed version of the original optimization (4) is given by

$$U_t(\beta_t, g_t) = \begin{cases} \min\limits_{b_t} \left(\frac{e^{b_t}-1}{g_t} + \bar{U}_{t-1}(\beta_t - b_t) \right), & t \ge 2, \\ \frac{e^{\beta_1}-1}{g_1}, & t = 1, \end{cases} \quad (6)$$

where $\bar{U}_{t-1}(\beta) = \mathbb{E}_g[U_{t-1}(\beta, g)]$ and can be calculated by induction [3]:

$$\bar{U}_t(\beta) = te^{\frac{\beta}{t}}\, \mathbb{G}(\nu_t, \nu_{t-1}, \cdots, \nu_1) - t\nu_1, \quad (7)$$

where $\mathbb{G}$ denotes the geometric mean operator (i.e., $\mathbb{G}(x_1, \cdots, x_n) = (\prod_{k=1}^n x_k)^{1/n}$) and $\nu_1, \nu_2, \cdots, \nu_t$ are the fractional moments of the fading distribution defined as:

$$\nu_m = \left(\mathbb{E}_g \left[\left(\frac{1}{g} \right)^{\frac{1}{m}} \right] \right)^m, \quad m = 1, 2, \cdots. \quad (8)$$

Due to the simple form of the cost-to-go function $\bar{U}_t$, by substituting (7) into (6) and solving the minimization we obtain the following closed-form description of the optimal policy for the relaxed problem [3]:

$$b_t(\beta_t, g_t) = \frac{1}{t}\beta_t + \frac{t-1}{t} \log \left(\frac{g_t}{\eta_t^{\text{relax}}} \right) \quad (9)$$

where η_t^{relax} serves as a channel threshold given by

$$\eta_t^{\text{relax}} = \frac{1}{\mathbb{G}(\nu_{t-1}, \nu_{t-2}, \cdots, \nu_1)}. \quad (10)$$

The policy function in (9) solves the boundary-relaxed problem but does not guarantee $0 \le b_t \le \beta_t$ in each slot.

To obtain a policy for the actual unrelaxed problem, we simply truncate at 0 and β_t, and reach what we refer to as the *boundary-relaxed* scheduler[4]:

$$b_t^{\text{relax}}(\beta_t, g_t) = \left\langle \frac{1}{t}\beta_t + \frac{t-1}{t} \log \frac{g_t}{\eta_t^{\text{relax}}} \right\rangle_0^{\beta_t} \quad (11)$$

where $\langle \cdot \rangle_0^{\beta_t}$ denotes truncation below 0 and above β_t. Notice that this policy function is optimal for $t = 2$, i.e., $b_2^{\text{relax}} = b_2^{\text{opt}}$ for all β_2 and g_2 since $(\bar{U}_1)' = (\bar{J}_1^{\text{opt}})'$.

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

[4]This is referred to as the suboptimal II scheduler in [3].

C. The One-shot Scheduler

The second scheduler is derived by modifying the boundary constraint into a stronger constraint $b_t \in \{0, \beta_t\}$ (equivalently, $b_t \in \{0, B\}$), i.e., in each slot either the entire packet is transmitted or nothing is transmitted. Then, the dynamic program is given by

$$J_t^{\text{one}}(\beta_t, g_t) = \begin{cases} \min\limits_{b_t \in \{0, \beta_t\}} \left(\frac{e^{b_t}-1}{g_t} + \bar{J}_{t-1}^{\text{one}}(\beta_t - b_t) \right), & t \ge 2, \\ \frac{e^{\beta_1}-1}{g_1}, & t = 1, \end{cases} \quad (12)$$

where $\bar{J}_t^{\text{one}}(\beta) = \mathbb{E}_g[J_t^{\text{one}}(\beta, g)]$. Equivalently, we can express the above DP as an optimal stopping problem [8] (this can be shown inductively with $\beta_T = B$):

$$J_t^{\text{one}}(B, g_t) = \begin{cases} \min \left\{ \frac{e^B-1}{g_t}, \;\; \bar{J}_{t-1}^{\text{one}}(B) \right\}, & t \ge 2, \\ \frac{e^B-1}{g_1}, & t = 1. \end{cases} \quad (13)$$

The optimal solution is a *sequential* threshold policy [3]:

$$b_t = \begin{cases} B, & \text{first } t \text{ such that } g_t > 1/\omega_t, \\ 0, & \text{otherwise}, \end{cases} \quad (14)$$

where $1/\omega_t$ is the channel threshold in slot t, and is recursively computed as:

$$\omega_t = \begin{cases} \mathbb{E} \left[\min \left(\frac{1}{g}, \omega_{t-1} \right) \right], & t = T, \cdots, 3, \\ \mathbb{E} \left[\frac{1}{g} \right], & t = 2, \\ \infty, & t = 1. \end{cases} \quad (15)$$

Notice that the thresholds depend only on the channel statistics and are independent of B, and that the thresholds decrease as the deadline approaches (i.e., as t decreases) [3].

D. The Delay-constrained Ergodic Scheduler

The above two suboptimal policies are developed to solve the DP, formulated in (4), by simplifying the cost-to-go function. Unlike these two policies, we now consider a policy by modifying the ergodic scheduling policy to meet the hard deadline constraint. The ergodic policy is the optimal solution to a problem of minimizing the average energy to transmit a certain *average* number of bits (i.e., no hard deadline constraint). If we denote this average rate constraint as $\bar{b}$, the ergodic scheduling policy function $b(g)$, which does not depend on t and determines how many bits to transmit based only upon the channel state g, is determined by solving:

$$\begin{aligned} \bar{E}^{\text{erg}}(\bar{b}) = \quad & \min_{b(g)} \quad \mathbb{E}_g \left[\frac{e^{b(g)}-1}{g} \right] \\ & \text{subject to} \quad \mathbb{E}_g[b(g)] \ge \bar{b}, \quad b(g) \ge 0. \end{aligned} \quad (16)$$

This optimization is readily solvable by standard waterfilling [9] and the solution is given by

$$b^{\text{erg}}(\bar{b}, g) = \left\langle \log\left(\frac{g}{\eta^{\text{erg}}}\right)\right\rangle_0^{\infty} = \begin{cases}\log\left(\frac{g}{\eta^{\text{erg}}}\right), & g \geq \eta^{\text{erg}}, \\ 0, & \text{else},\end{cases} \tag{17}$$

where η^{erg} serves as a channel threshold and is the solution to:

$$\mathbb{E}[b^{\text{erg}}(\bar{b}, g)] = \bar{b}. \tag{18}$$

When the time-horizon T is large, we intuitively expect the ergodic policy to perform well in the delay-limited setting considered here. In order to meet the deadline constraint, we utilize the ergodic policy, with $\bar{b} = \frac{B}{T} + \delta$ for some $\delta > 0$,[5] at each time step with the exception that all remaining unserved bits are transmitted in the final step:

$$b_t^{\text{constrained-erg}}\left(\frac{B}{T}, g_t; \delta\right) = \begin{cases} b^{\text{erg}}\left(\frac{B}{T} + \delta, g_t\right), & t = T, T-1, \cdots, 2, \\ \beta_1, & t = 1,\end{cases} \tag{19}$$

which is referred to as the *delay-constrained ergodic scheduler.*

IV. Asymptotic Optimality

This section investigates the optimality of the suboptimal schedulers introduced in the previous section. The optimality can be analyzed in two ways: optimality in policy and optimality in the associated energy cost. Both forms of optimality are shown for the boundary-relaxed scheduler and the one-shot scheduler, whereas energy optimality is shown for the delay-constrained ergodic scheduler.

A. Large B and Finite T: Asymptotic Optimality of Boundary-relaxed Scheduler

We first prove that the boundary-relaxed scheduler converges to the optimal policy when T is fixed and the number of bits B is taken to infinity. When B is large, we intuitively expect that the optimal policy will allocate strictly positive bits to all T time slots with high probability due to the nature of the Shannon energy-bit function. Thus, we expect the boundary-relaxed scheduler to coincide with the optimal policy when the number of bits to serve is large. The following theorem makes this relationship precise:

Theorem 1: Let the PDF f of g_t be continuous on $[g_{\min}, g_{\max}]$ with Support$(f) = [g_{\min}, g_{\max}]$, where $g_{\min} > 0$ and $g_{\max} < \infty$. For every time step t, the boundary-relaxed policy function in (11) converges to the optimal scheduling policy function uniformly on $[g_{\min}, g_{\max}]$ as the number of unserved bits β goes to infinity: for every given $\epsilon > 0$, there exists $\mathfrak{B}_0$ such that

$$\left|b_t^{\text{relax}}(\beta, g_t) - b_t^{\text{opt}}(\beta, g_t)\right| < \epsilon, \quad \forall g_t \in [g_{\min}, g_{\max}]. \tag{20}$$

for $\beta > \mathfrak{B}_0$.

Proof: By construction, $b_2^{\text{relax}} \equiv b_2^{\text{opt}}$ and $b_1^{\text{relax}} \equiv b_1^{\text{opt}}$. We assume that $b_{t-1}^{\text{relax}}(\cdot, g_{t-1}) \to b_{t-1}^{\text{opt}}(\cdot, g_{t-1})$ uniformly on $[g_{\min}, g_{\max}]$ holds as an induction hypothesis. From (4) and (5), we have

$$\bar{J}_{t-1}^{\text{opt}}(\beta) = \int_0^{\frac{1}{(\bar{J}_{t-2}^{\text{opt}})'(\beta)}} \bar{J}_{t-2}^{\text{opt}}(\beta) dF(x) + \int_{\frac{e^{\beta}}{(\bar{J}_{t-2}^{\text{opt}})'(\beta)}}^{\infty} \frac{e^{\beta}-1}{x} dF(x) + \int_{\frac{1}{(\bar{J}_{t-2}^{\text{opt}})'(\beta)}}^{\frac{e^{\beta}}{(\bar{J}_{t-2}^{\text{opt}})'(\beta)}} \left[\frac{e^{b_{t-1}^{\text{opt}}}-1}{x} + \bar{J}_{t-2}^{\text{opt}}(\beta - b_{t-1}^{\text{opt}})\right] dF(x) \tag{21}$$

and

$$\lim_{\beta\to\infty} (\bar{J}_{t-1}^{\text{opt}})'(\beta) = \lim_{\beta\to\infty} \int_{\frac{1}{(\bar{J}_{t-2}^{\text{opt}})'(\beta)}}^{\frac{e^{\beta}}{(\bar{J}_{t-2}^{\text{opt}})'(\beta)}} \frac{e^{b_{t-1}^{\text{opt}}}}{x} f(x)dx \tag{22}$$

since $\lim_{\beta\to\infty} \frac{1}{(\bar{J}_{t-2}^{\text{opt}})(\beta)} < g_{\min}$ and $\lim_{\beta\to\infty} \frac{e^{\beta}}{(\bar{J}_{t-2}^{\text{opt}})'(\beta)} > g_{\max}$, where $[g_{\min}, g_{\max}]$ is the support of f. With (7) and the induction hypothesis, we have

$$\lim_{\beta\to\infty}\left[(\bar{U}_{t-1})'(\beta) - (\bar{J}_{t-1}^{\text{opt}})'(\beta)\right] = \lim_{\beta\to\infty}\left[e^{\frac{\beta}{t-1}}\,\mathbb{G}(\nu_{t-1}, \cdots, \nu_1) - \int_{\frac{1}{(\bar{J}_{t-2}^{\text{opt}})'(\beta)}}^{\frac{e^{\beta}}{(\bar{J}_{t-2}^{\text{opt}})'(\beta)}} \frac{e^{b_{t-1}^{\text{relax}}}}{x} f(x)dx\right] \tag{23}$$

By substituting (11) into b_{t-1}^{relax} and re-writing $\mathbb{G}(\nu_{t-1}, \cdots, \nu_1)$ as

$$\mathbb{G}(\nu_{t-1}, \cdots, \nu_1) = (\mathbb{G}(\nu_{t-2}, \cdots, \nu_1))^{\frac{t-2}{t-1}} \int \left(\frac{1}{x}\right)^{\frac{1}{t-1}} f(x)dx, \tag{24}$$

we have $\lim_{\beta\to\infty}\left[(\bar{U}_{t-1})'(\beta) - (\bar{J}_{t-1}^{\text{opt}})'(\beta)\right] = 0$. By continuity, monotonicity, and convexity of $\bar{U}_{t-1}$ and $\bar{J}_{t-1}^{\text{opt}}$, we have $\lim_{\beta\to\infty}\left[(\bar{U}_{t-1})'(\beta - b_t^{\text{relax}}) - (\bar{J}_{t-1}^{\text{opt}})'(\beta - b_t^{\text{opt}})\right] = 0$. Since b_t^{relax} and b_t^{opt} are completely dependent on $(\bar{U}_{t-1})'$ and $(\bar{J}_{t-1})'$, respectively (β and g are common), we have the result by induction. See [1] for technical details. ■

We now compare the incurred energy costs of the two polices. We first define the incurred energy with the boundary-relaxed scheduler as:

$$J_t^{\text{relax}}(\beta_t, g_t) = \begin{cases}\frac{e^{b_t^{\text{relax}}}-1}{g_t} + \bar{J}_{t-1}^{\text{relax}}(\beta_t - b_t^{\text{relax}}), & t = T, T-1, \cdots, 2, \\ \frac{e^{\beta_1}-1}{g_1}, & t = 1,\end{cases} \tag{25}$$

where $\bar{J}_{t-1}^{\text{relax}}(\beta) = \mathbb{E}_g[J_{t-1}^{\text{relax}}(\beta, g)]$. Notice that (25) is not an optimization but is instead a calculation based upon the definition of b_t^{relax} in (11). Also notice that $\bar{J}_t^{\text{relax}}$ denotes the

[5] This policy is motivated by Theorem 3 of [10], where a modified version of the ergodic rate-maximizing policy is shown to maximize the expected transmitted rate over a finite time-horizon when the transmitter is subject to a finite energy constraint (which is the dual of the problem considered here).

cost for the actual un-relaxed problem (the energy cost with a policy satisfying $0 \le b_t \le \beta_t$ for all t), while the function $\bar{U}_t$ defined in Section III-B denotes the cost for the relaxed problem (the energy cost with a policy that may not satisfy $0 \le b_t \le \beta_t$).

Theorem 2: Let the PDF f of g_t be continuous on $[g_{\min}, g_{\max}]$ with Support$(f) = [g_{\min}, g_{\max}]$, where $g_{\min} > 0$ and $g_{\max} < \infty$. For any number of time slots T, the energy cost of the boundary-relaxed scheduler converges to the optimal energy cost as the number of bits B goes to infinity:

$$\lim_{B\to\infty} \left[\bar{J}_T^{\text{relax}}(B) - \bar{J}_T^{\text{opt}}(B)\right] = 0. \tag{26}$$

Proof: We will prove this by showing that $\lim_{B\to\infty} \left[\bar{U}_T(B) - \bar{J}_T^{\text{opt}}(B)\right] = 0$ and $\lim_{B\to\infty} \left[\bar{J}_T^{\text{relax}}(B) - \bar{U}_T(B)\right] = 0$. First, we show that $\lim_{B\to\infty} \left[\bar{U}_T(B) - \bar{J}_T^{\text{opt}}(B)\right] = 0$ by induction. At $t = 1$, all the bits are to be served and thus: $\bar{U}_1 \equiv \bar{J}_1^{\text{opt}}$. As an induction hypothesis, we assume that $\lim_{\beta\to\infty} \left[\bar{U}_{t-1}(\beta) - \bar{J}_{t-1}^{\text{opt}}(\beta)\right] = 0$. From (4) and (5), we write the expected cost-to-go as:

$$\begin{aligned}\bar{J}_t^{\text{opt}}(\beta) = &\int_0^{\frac{1}{(\bar{J}_{t-1}^{\text{opt}})'(\beta)}} \bar{J}_{t-1}^{\text{opt}}(\beta) f(x) dx \\ &+ \int_{\frac{e^\beta}{(\bar{J}_{t-1}^{\text{opt}})'(0)}}^{\infty} \frac{e^\beta - 1}{x} f(x) dx \\ &+ \int_{\frac{1}{(\bar{J}_{t-1}^{\text{opt}})'(\beta)}}^{\frac{e^\beta}{(\bar{J}_{t-1}^{\text{opt}})'(0)}} \left[\frac{e^{b_t^{\text{opt}}} - 1}{x} + \bar{J}_{t-1}^{\text{opt}}(\beta - b_t^{\text{opt}})\right] dF(x),\end{aligned} \tag{27}$$

where b_t^{opt} is a function of β (and x). From Theorem 1 and the induction hypothesis,

$$\lim_{\beta\to\infty} \left[b_t^{\text{relax}}(\beta, g) - b_t^{\text{opt}}(\beta, g)\right] = 0 \quad \text{uniformly } \forall g \in [g_{\min}, g_{\max}], \tag{28}$$

$$\lim_{\beta\to\infty} \left[\bar{U}_{t-1}(\beta - b_t^{\text{relax}}(\beta, g)) - \bar{J}_{t-1}^{\text{opt}}(\beta - b_t^{\text{opt}}(\beta, g))\right] = 0 \quad \text{uniformly } \forall g \in [g_{\min}, g_{\max}], \tag{29}$$

and thus,

$$\lim_{\beta\to\infty} \bar{J}_t^{\text{opt}}(\beta) = \lim_{\beta\to\infty} \int_{\frac{1}{(\bar{J}_{t-1}^{\text{opt}})'(\beta)}}^{\frac{e^\beta}{(\bar{J}_{t-1}^{\text{opt}})'(0)}} \left[\frac{e^{b_t^{\text{relax}}} - 1}{x} + \bar{U}_{t-1}(\beta - b_t^{\text{relax}})\right] dF(x) \tag{30}$$

Therefore, we have (31). By substituting (11) into b_t^{relax}, we have $\lim_{\beta\to\infty} [\bar{U}_t(\beta) - \bar{J}_t^{\text{opt}}(\beta)] = 0$ as desired. Thus, the induction holds. Similarly, we can prove $\lim_{B\to\infty} \left[\bar{J}_T^{\text{relax}}(B) - \bar{U}_T(B)\right] = 0$ by induction. See [1] for details. ∎

Although the analytic form of the optimal scheduler is not available, the above two theorems tell us that the boundary-relaxed scheduler, which has a very simple form that can be easily implemented, is asymptotically optimal when the number of bits to transmit (B) is sufficiently large. Furthermore, the scheduling function (11) provides intuition on the interplay between the channel quality and the deadline. When the deadline is far away (large t), the bit allocation is almost completely determined by the channel quality; on the other hand, as the deadline approaches (small t), the policy becomes less opportunistic.

B. Small B and Finite T: Asymptotic Optimality of One-shot Scheduler

We now show that the one-shot scheduling policy is asymptotically optimal when T is fixed and B is taken to zero. We first show convergence in terms of the policy function, and then in terms of the energy cost.

Theorem 3: For arbitrary time step t, the one-shot policy function in (14) converges to the optimal scheduling policy function as the number of unserved bits β tends to zero, i.e., the optimal policy becomes a threshold policy and the threshold coincides with the threshold of the one-shot policy:

$$\lim_{\beta\to 0} \sup\{g : b_t^{\text{opt}}(\beta, g) = 0\} = \lim_{\beta\to 0} \inf\{g : b_t^{\text{opt}}(\beta, g) = \beta\} = \frac{1}{\omega_t}, \tag{32}$$

where $1/\omega_t$ is the threshold of the one-shot policy as in (14) and (15).

Proof: The main idea is

$$\lim_{\beta\to 0} \left[\frac{e^\beta}{(\bar{J}_{t-1}^{\text{opt}})'(0)} - \frac{1}{(\bar{J}_{t-1}^{\text{opt}})'(\beta)}\right] = 0, \tag{33}$$

which implies that the case of $1/(\bar{J}_{t-1}^{\text{opt}})'(\beta) < g_t < e^\beta/(\bar{J}_{t-1}^{\text{opt}})'(0)$ occurs with vanishing probability as $\beta \to 0$. Thus, the optimal policy is a threshold policy. The threshold becomes identical to the threshold of the one-shot scheduler as $B \to 0$, which can be proved by induction. See [1] for details. ∎

Furthermore, we claim that the costs of the two policies also converge to one another. Since the average costs for the two policies converge to zero as $B \to 0$, cost convergence is investigated by studying the ratio, rather than the absolute difference, between the two costs:

Theorem 4: For arbitrary delay deadline T, the energy cost of the one-shot scheduler converges to the optimal energy cost as the number of bits B goes to zero:

$$\lim_{B\to 0} \frac{\bar{J}_T^{\text{one}}(B)}{\bar{J}_T^{\text{opt}}(B)} = 1. \tag{34}$$

Proof: Since $\lim_{B\to 0} \bar{J}_T^{\text{opt}}(B) = \lim_{B\to 0} \bar{J}_T^{\text{one}}(B) = 0$, by L'Hopital's rule, we have $\lim_{B\to 0} (\bar{J}_T^{\text{opt}})'(B) = \lim_{B\to 0} (\bar{J}_T^{\text{one}})'(B) > 0$. See [1] for details. ∎

$$
\begin{aligned}
\lim_{\beta\to\infty}[\bar{U}_t(\beta)-\bar{J}_t^{\text{opt}}(\beta)] &= \lim_{\beta\to\infty}\left[te^{\frac{\beta}{t}}\,\mathbb{G}\left(\nu_t,\nu_{t-1},\cdots,\nu_1\right)-t\nu_1-\int_{\frac{1}{(\bar{J}_{t-1}^{\text{opt}})'(\beta)}}^{\frac{e^{\beta}}{(\bar{J}_{t-1}^{\text{opt}})'(0)}}\left[\frac{e^{b_t^{\text{relax}}}-1}{x}+\bar{U}_{t-1}(\beta-b_t^{\text{relax}})\right]dF(x)\right] \\
&= \lim_{\beta\to\infty}\left[te^{\frac{\beta}{t}}\,\mathbb{G}\left(\nu_t,\nu_{t-1},\cdots,\nu_1\right)-\int_{\frac{1}{(\bar{J}_{t-1}^{\text{opt}})'(\beta)}}^{\frac{e^{\beta}}{(\bar{J}_{t-1}^{\text{opt}})'(0)}}\left[\frac{e^{b_t^{\text{relax}}}-1}{x}+(t-1)e^{\frac{\beta-b_t^{\text{relax}}}{t-1}}\,\mathbb{G}\left(\nu_{t-1},\cdots,\nu_1\right)\right]dF(x)\right]
\end{aligned} \tag{31}
$$

The optimality of one-shot scheduling can also be seen by upper and lower bounding the energy-bit function by linear functions. Using $x \le e^x - 1 \le xe^B$ for $0 \le x \le B$, we have:

$$\frac{b_t}{g_t} \le E_t(b_t, g_t) \le \frac{b_t e^B}{g_t}. \tag{35}$$

If we solve the DP using either of these bounds on the energy-bit function, the optimization in (4) becomes a linear program and thus a one-shot policy is optimal because a constrained linear program has a solution at a boundary of the constraint. Furthermore, the one-shot policy based on the upper and lower bounds converge to the one-shot policy described in Section III-C as $B \to 0$ because the bounds themselves converge.

C. Large T: Asymptotic Optimality of Causal Delay-constrained Ergodic Scheduler

When B and T are simultaneously taken to infinity at a particular ratio (i.e., $B, T \to \infty$ with $B = \bar{b}T$ for some constant $\bar{b} > 0$), we can show the energy-cost optimality of the ergodic policy in Section III-D.

The average energy cost of the delay-constrained ergodic scheduler is given by

$$
\begin{aligned}
\bar{J}_T^{\text{constrained-erg}}(\bar{b}T;\delta) &= \mathbb{E}\left[\sum_{t=1}^{T}\frac{e^{b_t^{\text{constrained-erg}}}-1}{g_t}\right] \\
&= \mathbb{E}\left[\sum_{t=2}^{T}\frac{e^{b_t^{\text{erg}}(\bar{b}+\delta,g_t)}-1}{g_t}\right]+\mathbb{E}\left[\frac{e^{\beta_1}-1}{g_1}\right],
\end{aligned} \tag{36}
$$

where β_1 denotes the remaining bits at the final slot and the value of δ is chosen such that

$$\bar{J}_T^{\text{constrained-erg}}(\bar{b}T) = \inf_{\delta>0}\bar{J}_T^{\text{constrained-erg}}(\bar{b}T;\delta). \tag{37}$$

Theorem 5: For any given average rate $\bar{b}(>0)$, the per-slot energy cost of the delay-constrained ergodic policy converges to the optimal ergodic energy cost as T tends to infinity:

$$\lim_{T\to\infty}\frac{1}{T}\bar{J}_T^{\text{constrained-erg}}(\bar{b}T) = \lim_{T\to\infty}\frac{1}{T}\bar{J}_T^{\text{opt}}(\bar{b}T) = \bar{E}^{\text{erg}}(\bar{b}). \tag{38}$$

Proof: See [1]. ■

The effect of the hard-deadline becomes inconsequential for large T because the channel realizations over the deadline horizon closely match the fading distribution. As a result, the delay-constrained ergodic scheduler performs similar to the ergodic scheduler when T is large. Moreover, the delay-constrained ergodic scheduler becomes causal optimal since any causal policy cannot be better than the ergodic policy.

D. Numerical Results: Policy Comparison

In order to compare the different asymptotically optimal policies, we compare their respective energy costs for different time-horizons (T). Since the analytical expression for the optimal policy is not available for $T > 2$, we solve the dynamic programming (4) numerically by the discretization method [11]. In Fig. 2 the per-slot energy consumption of the suboptimal schedulers is plotted for $T = 5$ and $T = 50$ assuming that the fading $\{g_t\}_{t=1}^T$ are i.i.d. truncated exponential with a support of $[0.001, 10^6]$, i.e.,

$$f(g) = \begin{cases} ce^{-(g-0.001)}, & \text{if } 0.001 \le g \le 10^6, \\ 0, & \text{otherwise}, \end{cases} \tag{39}$$

where c is a normalization factor. As can be seen, the one-shot scheduler is near-optimal only when B is small. The other schedulers performs close to the optimal through all ranges of B. When $T = 5$, as in Fig. 2a, the delay-constrained ergodic scheduler performs worse than the boundary-relaxed for all B. This is because $T = 5$ is too small for the delay-constrained ergodic scheduler to perform like the optimal. When $T = 50$, as in Fig. 2b (given the number of bits are in logarithmic scale), there exists a range of B such that the delay-constrained ergodic scheduler outperforms the boundary-relaxed scheduler. As can be seen in Fig. 2b, the one-shot scheduler performs best for small B (region A) and the boundary-relaxed scheduler outperforms when B is very large (region C). In the middle range (region B), the delay-constrained ergodic scheduler performs better than the other two.

V. Scheduling Gain

We have shown that the boundary-relaxed and the one-shot schedulers are asymptotically optimal as $B \to \infty$ and $B \to 0$, respectively. Another interesting issue is quantifying the advantage these schedulers provide compared to a non-opportunistic equal-bit scheduler that simply transmits B/T bits during each time slot.

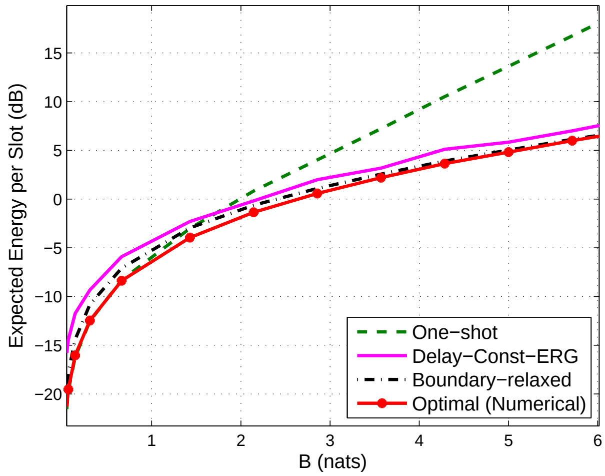

(a) $T = 5$

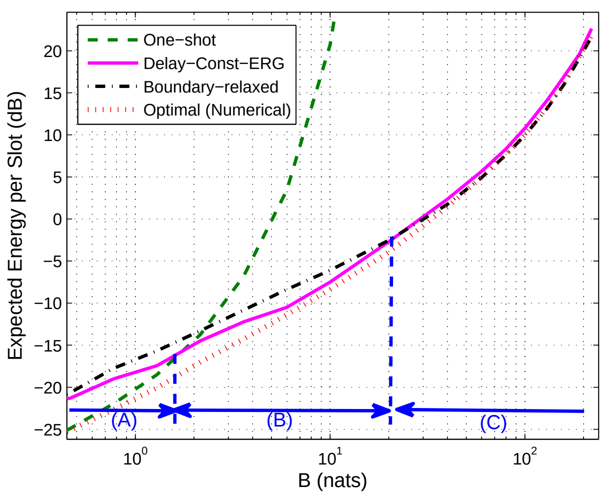

(b) $T = 50$

Fig. 2: Per slot energy cost for $T = 5$ and $T = 50$ when g is a truncated exponential variable with support $[0.001, 10^6]$

To compare energy performance, we first calculate the expected energy cost of the equal-bit scheduler, which is

$$\bar{J}_T^{\text{equal}}(B) = \mathbb{E}\left[\sum_{t=1}^{T} \frac{e^{\frac{B}{T}} - 1}{g_t}\right] = T\left(e^{\frac{B}{T}}\nu_1 - \nu_1\right), \tag{40}$$

since the equal-bit scheduler chooses $b_t = B/T$ for all t. Notice that the equal-bit scheduler achieves the delay-limited capacity [6] [12] (i.e., zero-outage capacity) with rate B/T.

We define the scheduling gain as the ratio between the expected energy expenditures:

$$\Delta_T^{\text{opt}}(B) \triangleq \frac{\bar{J}_T^{\text{equal}}(B)}{\bar{J}_T^{\text{opt}}(B)} \tag{41}$$

and quantify its behavior in the following theorem:

Theorem 6: For any T, the scheduling gain $\Delta_T^{\text{opt}}(B)$ is monotonically decreasing with respect to B. Furthermore, the limiting scheduling gains are given by:

$$\lim_{B\to 0} \Delta_T^{\text{opt}}(B) = \lim_{B\to 0} \frac{\bar{J}_T^{\text{equal}}(B)}{\bar{J}_T^{\text{one}}(B)} = \frac{\nu_1}{\omega_{T+1}}, \tag{42}$$

and if the PDF of the fading distribution is compactly supported and continuous,

$$\lim_{B\to\infty} \Delta_T^{\text{opt}}(B) = \lim_{B\to\infty} \frac{\bar{J}_T^{\text{equal}}(B)}{\bar{J}_T^{\text{relax}}(B)} = \frac{\nu_1}{\mathbb{G}(\nu_T, \cdots, \nu_1)}. \tag{43}$$

Proof: See [1]. ■

Since the boundary-relaxed scheduler is optimal as $B \to \infty$, the scheduling gain of the optimal scheduler and that of the boundary-relaxed scheduler are the same as $B \to \infty$; the same is true for the optimal and the one-shot scheduler as $B \to 0$. Intuitively, scheduling delivers a larger power gain for small B because in such scenarios one can be very opportunistic and transmit the entire packet once a sufficiently good channel state is realized. For larger B, however, it is inefficient to transmit the entire packet in a single slot (because energy increases exponentially with the number of bits) and thus transmissions must be spread across many slots (in fact, all slots are used as $B \to \infty$), which reduces the channel quality during those transmissions and thus reduces the benefit of scheduling.

VI. Conclusion

We have shown the asymptotic optimality of three different scheduling policies for delay-constrained transmission over a fading channel. When only a small number of bits need to be served, a one-shot threshold policy is optimal: once a sufficiently good channel state is experienced, the entire packet is transmitted. On the other hand, when the number of bits is large, the number of transmitted bits at each time step should be a weighted sum of the unserved bits and a channel state-related term, where the weight is proportional to the time to deadline. In each of these two policies, the scheduler is opportunistic while also being cognizant of the deadline. Furthermore, a modification of the ergodic waterfilling policy is shown to be optimal when the number of bits and the time horizon are both large.

Although problems involving delay-limited communication are of great practical importance and have been the subject of considerable research, such problems generally do not have closed-form solutions. In this work, however, we are able to circumvent this general difficulty by considering different asymptotic regimes. It would be interesting to see if the asymptotically optimal policies identified here, which admit a very simple analytical form, can be extended to other more general settings. For example, to scheduling with time-varying channels and randomly arriving packets [13] [14] and possibly to multi-user channels [15].

REFERENCES

[1] J. Lee and N. Jindal, "Asymptotically optimal policies for hard-deadline scheduling over fading channels," submitted to *IEEE Trans. Inform. Theory* (preprint available at http://arxiv.org/abs/0906.5397), June 2009.

[2] 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.

[3] J. Lee and N. Jindal, "Energy-efficient scheduling of delay constrained traffic over fading channels," *IEEE Trans. Wireless Commun.*, vol. 8, no. 4, pp. 1866–1875, Apr. 2009.

[4] M. Zafer and E. Modiano, "Delay constrained energy efficient data transmission over a wireless fading channel," in *Workshop on Inf. Theory and Appl.*, La Jolla, CA, Jan./Feb. 2007, pp. 289–298.

[5] A. Goldsmith and P. Varaiya, "Capacity of fading channels with channel side information," *IEEE Trans. Inform. Theory*, vol. 43, pp. 1986–1992, Nov. 1997.

[6] G. Caire, G. Taricco, and E. Biglieri, "Optimum power control over fading channels," *IEEE Trans. Inform. Theory*, vol. 45, no. 5, pp. 1468–1489, Jul. 1999.

[7] R. T. Rockafellar, *Convex Analysis*. Princeton Univ. Press, 1970.

[8] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 3rd ed. Mass.: Athena Scientific, 2005, vol. 1.

[9] A. Goldsmith, *Wireless Communications*. New York, NY: Cambridge Univ. Press, 2005.

[10] 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.

[11] D. P. Bertsekas, "Convergence of discretization procedures in dynamic programming," *IEEE Trans. Automat. Contr.*, vol. AC-20, no. 3, pp. 415–419, Jun. 1975.

[12] V. Hanly and D. Tse, "Multiaccess fading channels. Part II: Delay-limited capacities," *IEEE Trans. Inform. Theory*, vol. 44, pp. 2816–2831, Nov. 1998.

[13] 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.

[14] 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.

[15] E. M. Yeh and A. S. Cohen, "Throughput and delay optimal resource allocation in multiaccess fading channels," in *Proc. IEEE Int. Symp. on Inform. Theory (ISIT)*, Yokohama, Japan, Jun./Jul. 2003, p. 245.