

Adaptive Modulation with Smoothed Flow Utility

Ekine Akuiyibo * Stephen Boyd†

May 26, 2010

Abstract

We consider the problem of choosing the data flow rate on a wireless link with randomly varying channel gain, to optimally trade off average transmit power and the average utility of the smoothed data flow rate. The smoothing allows us to model the demands of an application that can tolerate variations in flow over a certain time interval; we will see that this smoothing leads to a substantially different optimal data flow rate policy than without smoothing. We pose the problem as a convex stochastic control problem. For the case of a single flow, the optimal data flow rate policy can be numerically computed using stochastic dynamic programming. For the case of multiple flows on a single link, we propose an approximate dynamic programming approach to obtain suboptimal data flow rate policies. We illustrate, through numerical examples, that nearly optimal performance can be obtained with these approximate policies.

1 Introduction

1.1 Average smoothed flow utility

A wireless communication link supports n data flows in a channel that varies with time, which we model using discrete-time intervals $t = 0, 1, 2, \dots$. We let $f_t \in \mathbf{R}_+^n$ be the data flow rate vector on the link, where $(f_t)_j$, $j = 1, \dots, n$, is the j th flow's data rate at time t and \mathbf{R}_+ denotes the set of nonnegative numbers. We let $F_t = \mathbf{1}^T f_t$ denote the total flow rate over all flows, where $\mathbf{1}$ is the vector with all entries one. The flows, and the total flow rate, will depend on the random channel gain (through the flow policy, described below) and so are random variables.

We will work with a *smoothed* version of the flow rates, which is meant to capture the tolerance of the applications using the data flows to time variations in data rate. This was introduced in [TZB08] using delivery contracts, in which the utility is a function of the total

*Electrical Engineering Department, Stanford University.

†Electrical Engineering Department, Stanford University.

flow over a given time interval; here, we use instead a very simple first order linear smoothing. At each time t , the smoothed data flow rate vector $s_t \in \mathbf{R}_+^n$ is given by

$$s_{t+1} = \Theta s_t + (I - \Theta)f_t, \quad t = 0, 1, \dots,$$

where $\Theta = \mathbf{diag}(\theta)$, and $\theta_j \in [0, 1)$, $j = 1, \dots, n$, is the smoothing parameter for the j th flow, and we take $s_0 = 0$. Thus we have

$$(s_t)_j = \sum_{\tau=0}^{t-1} (1 - \theta_j)\theta_j^{t-1-\tau}(f_\tau)_j,$$

where at time t , each smoothed flow rate $(s_t)_j$ is the exponentially weighted average of previous flow rates.

The smoothing parameter θ_j determines the level of smoothing on flow j . Small smoothing parameter values (θ_j close to zero) correspond to light smoothing; large values (θ_j close to one) correspond to heavy smoothing. (Note that $\theta_j = 0$ means that flow j is not smoothed; we have $(s_{t+1})_j = (f_t)_j$.) The level of smoothing can be related to the time scale over which the smoothing occurs. We define $T_j = 1/\log(1/\theta_j)$ to be the *smoothing time* associated with flow j . Roughly speaking, the smoothing time is the time interval over which the effect of a flow on the smoothed flow decays by a factor $1/e$. Light smoothing corresponds to short smoothing times, while heavy smoothing corresponds to longer smoothing times.

We associate with each smoothed flow rate $(s_t)_j$ a strictly concave nondecreasing differentiable utility function $U_j : \mathbf{R}_+ \rightarrow \mathbf{R}$, where the utility of $(s_t)_j$ is $U_j((s_t)_j)$. The average utility derived over all flows, over all time is

$$\bar{U} = \lim_{N \rightarrow \infty} \mathbf{E} \frac{1}{N} \sum_{t=0}^{N-1} U(s_t),$$

where $U(s_t) = U_1((s_t)_1) + \dots + U_n((s_t)_n)$. Here, the expectation is over the smoothed flows s_t , and we are assuming that the expectations and limit above exist.

While most of our results will hold for more general utilities, we will focus on the family of power utility functions, defined for $x \geq 0$ as

$$U(x) = \beta x^\alpha, \tag{1}$$

parametrized by $\alpha \in (0, 1)$ and $\beta > 0$. The parameter α sets the curvature (or risk aversion), while β sets the overall weight of the utility. (For small values of α , U approaches a log utility.)

Before proceeding, we make some general comments on our use of smoothed flows. The smoothing can be considered a type of time averaging; then we apply a concave utility function; and finally, we average this utility. The time averaging and utility function operations do not commute, except in the case when the utility is linear (or affine). Jensen's inequality

tells us that average smoothed utility is greater than or equal to the average utility applied directly to the flow rates, *i.e.*,

$$U((s_t)_j) \geq \frac{1}{t} \sum_{\tau=0}^{t-1} (1 - \theta_j) \theta_j^{t-1-\tau} U(f_\tau)_j.$$

So the time smoothing step does affect our average utility; we will see later that it has a dramatic affect on the optimal flow policy.

1.2 Average power

We model the wireless channel with time-varying positive gain parameters g_t , $t = 0, 1, \dots$, which we assume are independent identically distributed (IID), with known distribution. At each time t , the gain parameter affects the power P_t required to support the total data flow rate F_t . The power P_t is given by

$$P_t = \phi(F_t, g_t),$$

where $\phi : \mathbf{R}_+ \times \mathbf{R}_{++} \rightarrow \mathbf{R}_+$ is increasing and strictly convex in F_t for each value of g_t (\mathbf{R}_{++} is the set of positive numbers).

While our results will hold for the more general case, we will focus on the more specific power function described here. We suppose that the signal-to-interference-and-noise (SINR) of the channel is given by $g_t P_t$. (Here g_t includes the effect of time-varying channel gain, noise, and interference.) The channel capacity is then $\mu \log(1 + g_t P_t)$, where μ is a constant; this must equal at least the total flow rate F_t , so we obtain

$$P_t = \phi(F_t, g_t) = \frac{e^{F_t/\mu} - 1}{g_t}. \quad (2)$$

The total average power is

$$\bar{P} = \lim_{N \rightarrow \infty} \mathbf{E} \frac{1}{N} \sum_{t=0}^{N-1} P_t,$$

where again, we are assuming that the expectations and limit exist.

1.3 Flow rate control problem

The overall objective is to maximize a weighted difference between average utility and average power,

$$J = \bar{U} - \lambda \bar{P}, \quad (3)$$

where $\lambda \in \mathbf{R}_{++}$ is used to trade off average utility and power.

We require that the flow policy is causal, *i.e.*, when f_t is chosen, we know the previous and current values of the flows, smoothed flows, and channel gains. Standard arguments in

stochastic control (see, *e.g.*, [Ber05, Ber07, Åst70, Whi82, BS96]) can be used to conclude that, without loss of generality, we can assume that the flow control policy has the form

$$f_t = \varphi(s_t, g_t), \quad (4)$$

where $\varphi : \mathbf{R}_+^n \times \mathbf{R}_{++} \rightarrow \mathbf{R}_+^n$. In other words, the policy depends only on the current smoothed flows, and the current channel gain value.

The flow rate control problem is to choose the flow rate policy φ to maximize the overall objective in (3). This is a standard convex stochastic control problem, with linear dynamics.

1.4 Previous work

In the wireless communications literature, varying a link’s transmit rate (and power) depending on channel conditions is called *adaptive modulation* (AM); see, *e.g.*, [Hay68, Cav72, Hen74, WS95, GC97]. One drawback of AM is that it is a physical layer optimization technique with no knowledge of upper layer optimization protocols.

Maximizing a total utility function is also very common in various communications and networking problem formulations, where it is referred to as network utility maximization (NUM); see, *e.g.*, [KMT97, LL99, CLCD07, NML08]. In the NUM framework, performance of an upper layer protocol (*e.g.*, TCP) is determined by *utility* of flow attributes, *e.g.*, utility of link flow rate.

Our setup involves both adaptive modulation and utility maximization, but is nonstandard in several respects: We consider the utility of the smoothed flows, and we consider multiple flows over the same wireless link [OABG10].

1.5 Our results

We let J^* be the optimal overall objective value, and φ^* be an optimal policy. We will show that in the general (multiple flow) case, the optimal policy includes a “no-transmit” zone, *i.e.*, a region in the (s_t, g_t) space in which the optimal flow rate is zero. Not surprisingly, the optimal flow policy can be roughly described as waiting until the channel gain is large, or until the smoothed flow has fallen to a low level, at which point we transmit (*i.e.*, choose nonzero f_t). Roughly speaking, the higher the level of smoothing, the longer we can afford to wait for a large channel gain before transmitting. The average power required to support a given utility level decreases, sometimes dramatically, as the level of smoothing increases.

We show that the optimal policy for the case of a single flow is readily computed numerically, working from Bellman’s characterization of the optimal policy, and is not particularly sensitive to the details of the utility functions, smoothing levels, or power functions.

For the case of multiple flows, we cannot easily compute (or even represent) the optimal policy. For this case we propose an approximate policy, based on approximate dynamic programming [BT96, Pow07]. By computing an upper bound on J^* , by allowing the flow control policy to use future values of channel gain (*i.e.*, relaxing the causality requirement [BSS10]), we show in numerical experiments that such policies are nearly optimal.

2 Optimal policy characterization

2.1 No smoothing

We first consider the special case $\Theta = 0$, in which case there is no smoothing. Then we have $s_t = f_{t-1}$, so the average smoothed utility is then the same as the average utility, *i.e.*,

$$\bar{U} = \lim_{N \rightarrow \infty} \mathbf{E} \frac{1}{N} \sum_{t=0}^{N-1} U(f_t),$$

In this case the optimal policy is trivial, since the stochastic control problem reduces to a simple optimization problem at each time step. At time t , we simply choose f_t to maximize $U(f_t) - \lambda P_t$. Thus we have

$$\varphi(s_t, g_t) = \operatorname{argmax}_{f_t \geq 0} (U(f_t) - \lambda P_t),$$

which does not depend on s_t . A simple and effective approach is to pre-solve this problem for a suitably large set of values of the channel gain g_t , and store the resulting tables of individual flow rates $(f_t)_i$ versus g_t ; online we can interpolate between points in the table to find the (nearly) optimal policy. Another option is to fit a simple function to the optimal flow rate data, and use this function as our (nearly) optimal policy.

For future reference, we note that the problem can also be solved using a waterfilling method (see, *e.g.*, [BV04, §5.5]). Dropping the time index t , and using j to denote the flow index, we must solve the problem

$$\begin{aligned} & \text{maximize} && \sum_{j=1}^n U_j(f_j) - \lambda \phi(F, g) \\ & \text{subject to} && F = \mathbf{1}^T f, \quad f \geq 0, \end{aligned} \tag{5}$$

with variables f_j and F . Introducing a Lagrange multiplier ν for the equality constraint (which we can show must be nonnegative, using monotonicity of ϕ with F), we are to maximize

$$\sum_{j=1}^n U_j(f_j) - \lambda \phi(F, g) + \nu (F - \mathbf{1}^T f)$$

over $f_j \geq 0$. This problem is separable in f_j and F , so we can maximize over f_j and F separately. We find that

$$f_j = \operatorname{argmax}_{w \geq 0} (U_j(w) - \nu w), \quad j = 1, \dots, n,$$

and

$$F = \operatorname{argmax}_{y \geq 0} (\nu y - \lambda \phi(y, g)).$$

(Each of these can be expressed in terms of conjugate functions; see, *e.g.*, [BV04, §3.3].) We then adjust ν (say, using bisection) so that $\mathbf{1}^T f = F$. An alternative is to carry out

bisection on ν , defining f_j in terms of ν as above, until $\lambda\phi'(\mathbf{1}^T f, g) = \nu$, where ϕ' refers to the derivative with respect to y .

For our particular power law utility functions (1) we can give an explicit formula for f_j in terms of ν :

$$f_j = \left(\frac{\alpha_j \beta_j}{\nu} \right)^{1/(1-\alpha_j)}.$$

For our particular power function (2), we use bisection to find the value of ν that yields

$$\mathbf{1}^T f = \mu \log(\nu \mu g / \lambda),$$

where the flow values come from the equation above. (The lefthand side is decreasing in ν , while the righthand side is increasing.)

2.2 General case

We now consider the more general case, with smoothing. We can characterize the optimal flow rate policy φ^* using stochastic dynamic programming [Put94, Ros83, Den82, WB09] and a form of Bellman's equation [Bel57]. The optimal flow rate policy has the form

$$\varphi^*(z, g) = \operatorname{argmax}_{w \geq 0} (V(\Theta z + (I - \Theta)w) - \lambda\phi(\mathbf{1}^T w, g)), \quad (6)$$

where $V : \mathbf{R}_+^n \rightarrow \mathbf{R}$ is the Bellman (relative) value function. The value function (and optimal value) is characterized via the fixed point equation

$$J^* + V = \mathcal{T}V, \quad (7)$$

where for any function $W : \mathbf{R}_+^n \rightarrow \mathbf{R}$, the Bellman operator \mathcal{T} is given by

$$(\mathcal{T}W)(z) = U(z) + \mathbf{E} \left(\max_{w \geq 0} (W(\Theta z + (I - \Theta)w) - \lambda\phi(\mathbf{1}^T w, g)) \right),$$

where the expectation is over g . The fixed point equation and Bellman operator are invariant under adding a constant; that is, we have $\mathcal{T}(W + a) = \mathcal{T}W + a$, for any constant (function) a , and similarly, V satisfies the fixed point equation if and only if $V + a$ does. So without loss of generality we can assume that $V(0) = 0$.

The value function can be found (in principle) by value iteration [Bel57, Ber07]. We take $V^{(0)} = 0$, and repeat the following iteration, for $k = 0, 1, \dots$

1. $\tilde{V}^{(k)} = \mathcal{T}V^{(k)}$. (Apply Bellman operator.)
2. $J^{(k)} = \tilde{V}^{(k)}(0)$. (Estimate optimal value.)
3. $V^{(k+1)} = \tilde{V}^{(k)} - J^{(k)}$. (Normalize.)

For technical conditions under which the value function exists, and can be obtained via value iteration, see, *e.g.*, [Der70, Bla62, ABFG⁺93]. We will simply assume here that the value function exists, and $J^{(k)}$ and $V^{(k)}$ converge to J^* and V , respectively.

The iterations above preserve several attributes of the iterates, which we can then conclude holds for V . First of all, concavity of $V^{(k)}$ is preserved, *i.e.*, if $V^{(k)}$ is concave, so is $V^{(k+1)}$. It is clear that normalization does not affect concavity, since we simply add a constant to the function. The Bellman operator \mathcal{T} preserves concavity since partial maximization of a function concave in two sets of variables results in a concave function (see, *i.e.*, [BV04, §3.2]), and expectation over a family of concave functions yields a concave function; finally, addition (of U) preserves concavity. So we can conclude that V is concave.

Another attribute that is preserved in value iteration is monotonicity; if $V^{(k)}$ is monotone increasing (in each component of its argument), then so is $V^{(k+1)}$. We conclude that V is monotone increasing.

2.3 No-transmit region

From the form of the optimal policy, we see that $\varphi(z, g) = 0$ if and only if $w = 0$ is optimal for the (convex) problem

$$\begin{aligned} & \text{maximize} && V(\Theta z + (I - \Theta)w) - \lambda\phi(\mathbf{1}^T w, g) \\ & \text{subject to} && w \geq 0, \end{aligned}$$

with variable $w \in \mathbf{R}^n$. This is the case if and only if

$$(I - \Theta)\nabla V(\Theta z) + \lambda\phi'(0, g)\mathbf{1} \leq 0$$

(see, *e.g.*, [BV04, p142]). We can rewrite this as

$$\frac{\partial V}{\partial z_i}(\Theta z) \leq \frac{\lambda\phi'(0, g)}{1 - \theta_i}, \quad i = 1, \dots, n.$$

Using the specific power function (2) associated with the log capacity formula, we obtain

$$\nabla V(\Theta z) \leq \frac{\lambda}{\mu g} \left(\frac{1}{1 - \theta_1}, \dots, \frac{1}{1 - \theta_n} \right) \quad (8)$$

as the necessary and sufficient condition under which $\varphi(z, g) = 0$. Since ∇V is decreasing (by concavity of V), we can interpret (8) roughly as follows: Do not transmit if the channel is bad (g small) *or* if the smoothed flows are large (z large).

3 Single flow case

3.1 Optimal policy

In the case of a single flow (*i.e.*, $n = 1$) we can easily carry out value iteration numerically, by discretizing the argument z and values of g , and computing the expectation and maximization numerically. For the single flow case, then, we can compute the optimal policy and optimal performance (up to small numerical integration errors).

3.2 Power law suboptimal policy

We replace the optimal value function (in the above optimal flow policy expression) with a simple analytic approximation of the value function to get the approximate policy

$$\hat{\varphi}(z, g) = \operatorname{argmax}_{w \geq 0} \left(\hat{V}(\theta z + (1 - \theta)w) - \lambda \phi(w, g) \right), \quad (9)$$

where $\hat{V}(z)$ is an approximation of the value function.

Since V is increasing, concave, and satisfies $V(0) = 0$, it is reasonable to fit it with a power law function as well, say $V(z) \approx \tilde{\beta} z^{\tilde{\alpha}}$, with $\tilde{\beta} > 0$, $\tilde{\alpha} \in (0, 1)$. For example, we can find the minimax (Chebyshev fit) by varying $\tilde{\alpha}$; for each $\tilde{\alpha}$ we choose $\tilde{\beta}$ to minimize

$$\max_i \left| V_i - \tilde{\beta} z_i^{\tilde{\alpha}} \right|,$$

where z_i are the discretized values of z , with associated value function values V_i . We do this by bisection on $\tilde{\beta}$.

Experiments show that these power law approximate functions are, in general, reasonable approximations for the value function. For our power law utilities, these approximations yield very good matches to the true value function. For other concave utilities, the approximation is not as accurate; but experiments show that the associated approximate policies still yield nearly optimal performance.

We can derive an explicit expression for the approximate policy (9) for our power function:

$$\hat{\varphi}(z, g) = \begin{cases} \kappa(z, g) - \gamma z & \kappa(z, g) - \gamma z > 0 \\ 0 & \kappa(z, g) - \gamma z \leq 0, \end{cases} \quad (10)$$

where

$$\kappa(z, g) = \mu(1 - \tilde{\alpha}) \mathcal{W} \left(\frac{(\mu(1 - \tilde{\alpha})(1 - \theta))^{-1} e^{\left(\frac{\gamma z}{\mu(1 - \tilde{\alpha})}\right)}}{\left(\frac{\lambda}{\tilde{\beta} \tilde{\alpha} (1 - \theta) g \mu}\right)^{1/(1 - \tilde{\alpha})}} \right), \quad \gamma = \theta/(1 - \theta),$$

and \mathcal{W} is the Lambert function, *i.e.*, $\mathcal{W}(a)$ is the solution of $we^w = a$ [CGH⁺96].

Note that this suboptimal policy is not needed in the single flow case since we can obtain the optimal policy numerically. However, we found that the difference between our power law policy and the optimal policy (see example value functions below) is small enough that in practice they are virtually the same. This approximate policy is needed in the case of multiple flows.

3.3 Numerical example

In this section we give simple numerical examples to illustrate the effect of smoothing on the resulting flow rate policy in the single flow case. We consider two examples, with different

levels of smoothing. The first flow is lightly smoothed ($T = 1$; $\theta = 0.37$), while the second flow is heavily smoothed ($T = 50$; $\theta = 0.98$). We use utility function $U(s) = s^{1/2}$, *i.e.*, $\alpha = 1/2$, $\beta = 1$ in our utility (1). The channel gains g_t are IID exponential variables with mean $\mathbf{E} g_t = 1$. We use the power function (2), with $\mu = 1$.

We first consider the case $\lambda = 1$. The value functions are shown in figure 1, together with the power law approximations, which are $1.7s^{0.6}$ (light smoothing) and $42.7s^{0.74}$ (heavy smoothing). Figure 2 shows the optimal policies for the lightly smoothed flow ($\theta = 0.37$), and the heavily smoothed flow ($\theta = 0.98$). We can see that the optimal policies are quite different. As expected, the lightly smoothed flow transmits more often, *i.e.*, has a smaller no-transmit region.

Average power vs average utility. We further illustrate the difference between the two flow rate policies using their average power-average utility, trade-off curves shown in figure 3. In this example, we computed, via simulation, the average power-average utility curve using $\lambda \in (0, 1]$, and an initial smoothed flow rate $s_0 = 1$. As expected, we can see that the heavily smoothed flow achieves more average utility than the lightly smoothed flow, for a given average power. (The heavily smoothed flow requires less average power to achieve a target average utility.)

Comparing average power. Here we compare the average power required by each flow to generate a given average utility. We can estimate this roughly from the trade-off curves in figure 3 or more precisely via simulation: choose a target average utility, and then run each controller, adjusting λ separately, until we reach the target utility. In our example, we ran simulations to find $\lambda = 0.29$ for the lightly smoothed flow, and $\lambda = 0.57$ for the heavily smoothed flow, to give $\bar{U} = 0.8$, for each policy.

The associated power trajectories are shown in figure 4, with the dashed (horizontal) line indicating the average power in each trajectory. Clearly the lightly smoothed flow requires more power than the heavily smoothed flow, by a factor of around 3: the heavily smoothed flow requires $\bar{P} = 0.3$, compared to $\bar{P} = 0.94$ for the lightly smoothed flow.

Average utility versus smoothing time. Lastly, we illustrate the relationship between smoothing time and average utility, for a fixed average power. For this figure we use $T = (1, 5, 10, 15, 20, 25, 30, 35, 40, 45, 50)$ and $\bar{P} = 0.5$. Figure 5 shows, as we expect, a gradual increase in average utility as the smoothing time increases.

4 A suboptimal policy for the multiple flow case

4.1 Approximate dynamic programming (ADP) policy

In this section we describe a suboptimal policy that can be used in the multiple flow case. Our proposed policy has the same form as the optimal policy, with the true value function

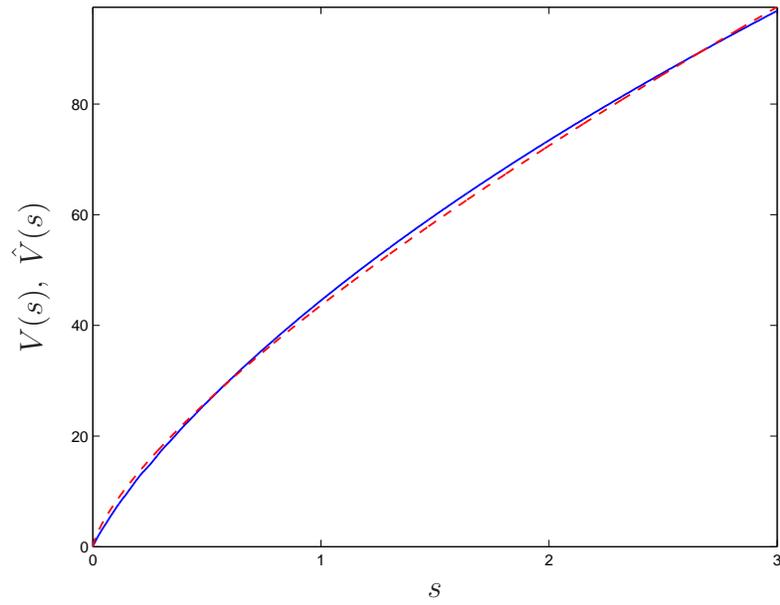
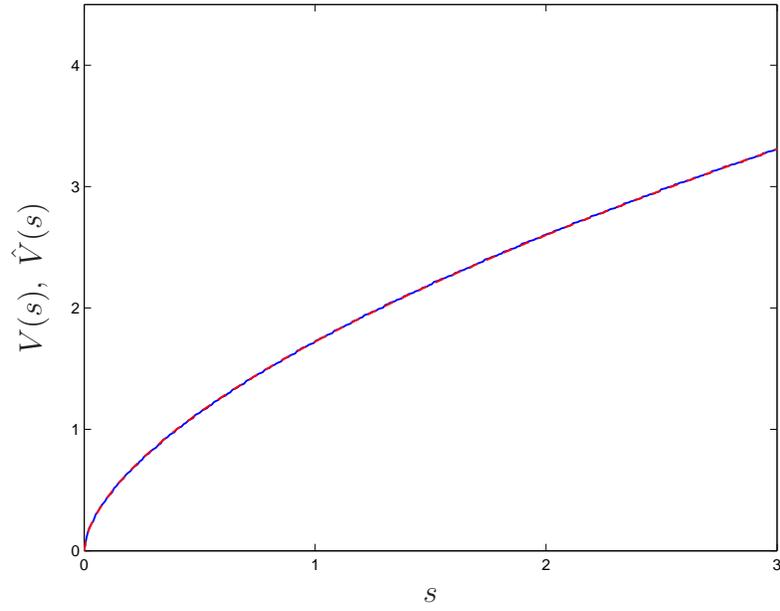


Figure 1: *Top.* Comparing V (blue) with \hat{V} (red, dashed) for the lightly smoothed flow. *Bottom.* Comparing V (blue) with \hat{V} (red, dashed) for the heavily smoothed flow.

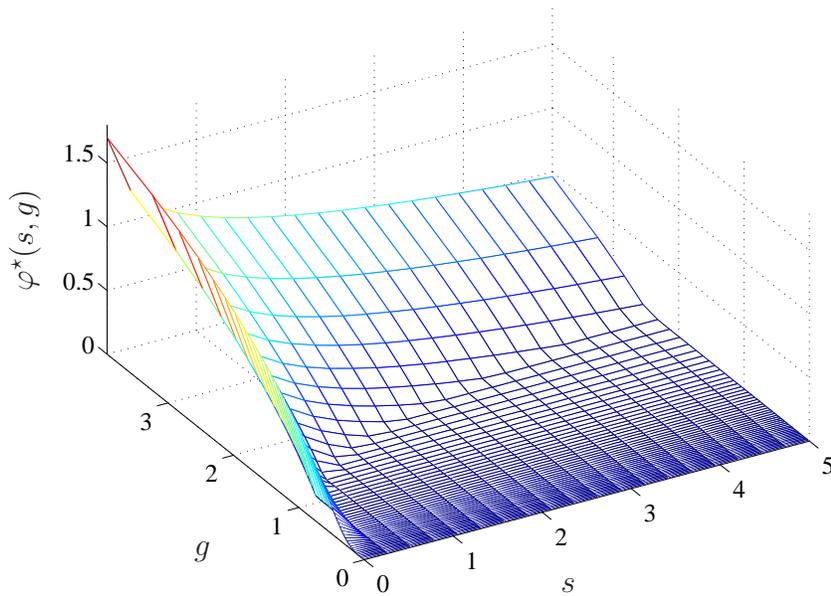
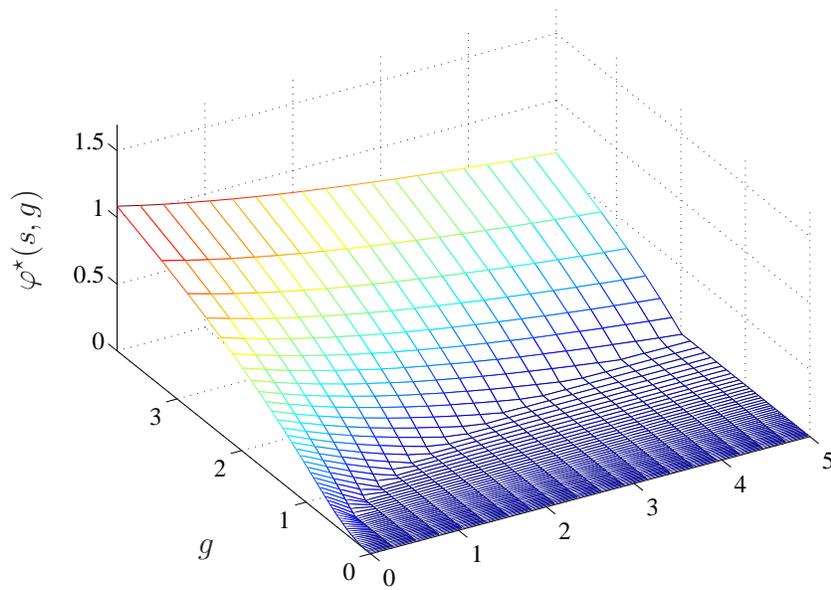


Figure 2: *Top.* Optimal policy $\varphi^*(s, g)$ for smoothing time $T = 1$ ($\theta = 0.37$).
Bottom. Optimal policy for $T = 50$ ($\theta = 0.98$).

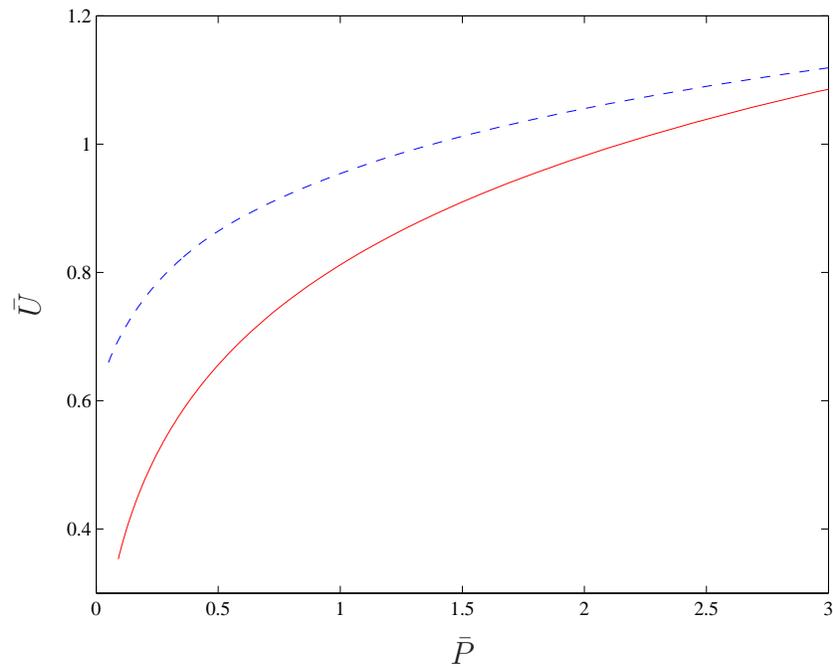


Figure 3: Average utility versus average power; heavily smoothed flow (top, dashed), lightly smoothed flow (bottom).

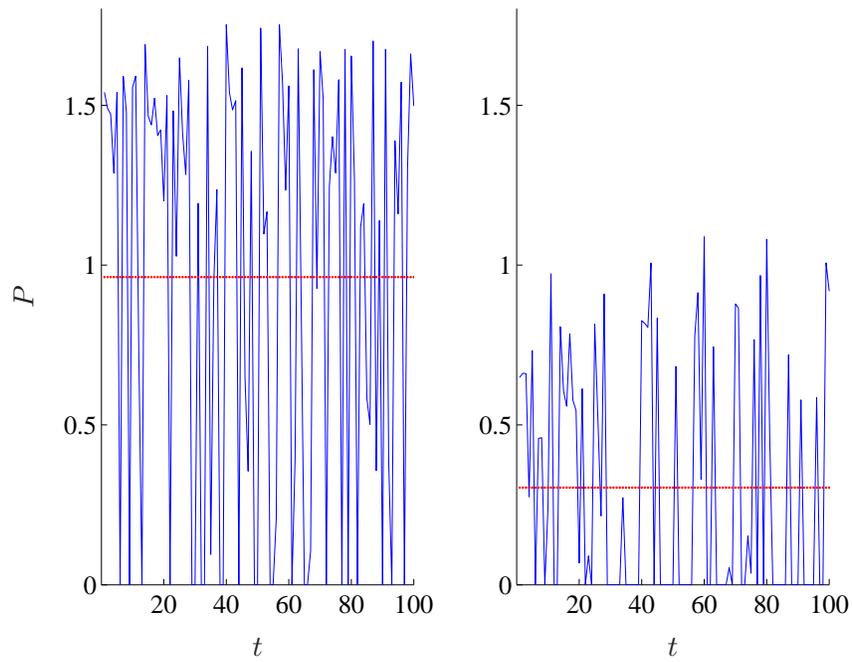


Figure 4: Sample power trajectories; lightly smoothed flow (left), heavily smoothed flow (right).

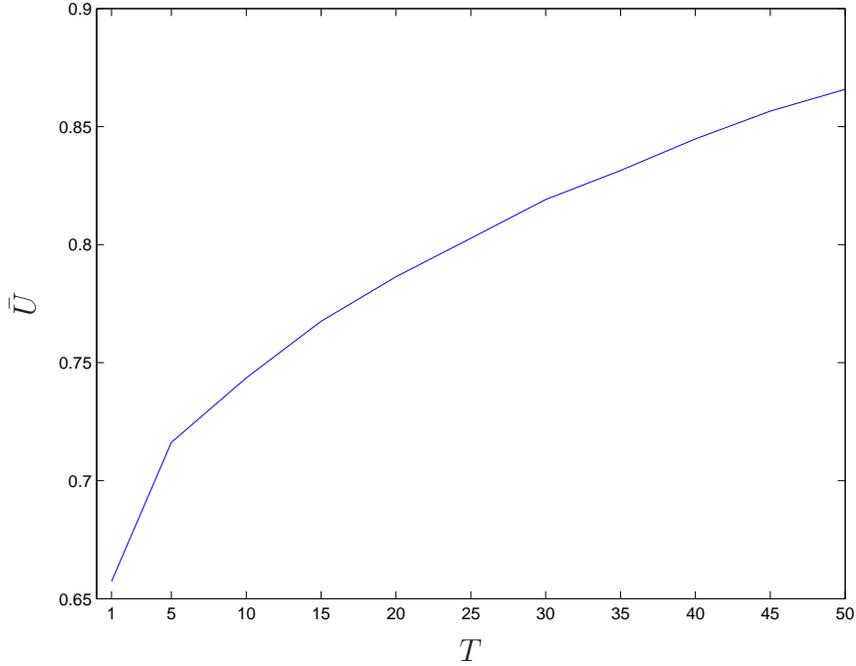


Figure 5: Average utility versus smoothing time.

V replaced with an approximation or surrogate V^{adp} :

$$\varphi^{\text{adp}}(z, g) = \underset{w \geq 0}{\operatorname{argmax}} \left(V^{\text{adp}}(\Theta z + (I - \Theta)w) - \lambda \phi(\mathbf{1}^T w, g) \right). \quad (11)$$

A policy obtained by replacing V with an approximation is called an approximate dynamic programming (ADP) policy [BT96, Pow07, Man60]. (Note that by this definition (9) is an ADP policy for $n = 1$.)

We construct V^{adp} in a simple way. Let $\hat{V}_j : \mathbf{R}_+ \rightarrow \mathbf{R}$ denote the power law approximate function for the associated single flow problem with only the j th flow. (This can be obtained numerically as described above.) We then take

$$V^{\text{adp}}(z) = \hat{V}_1(z_1) + \dots + \hat{V}_n(z_n).$$

This approximate value function is separable, *i.e.*, a sum of functions of the individual flows, whereas the exact value function is (in general) not. The approximate policy, however, is not separable; the optimization problem solving to assign flow rates couples the different flow rates.

In the literature on approximate dynamic programming, \hat{V}_j would be considered basis functions [SS85, TZ97, dFR03]; however, we fix the coefficients of the basis functions as one. (We have found that very little improvement in the policy is obtained by optimizing over the coefficients.)

Evaluating the approximate policy, *i.e.*, solving (11), reduces to solving the resource allocation problem

$$\begin{aligned} & \text{maximize} && \sum_{j=1}^n \hat{V}_j(\theta_j z_j - (1 - \theta_j) f_j) - \lambda \phi(F, g) \\ & \text{subject to} && F = \mathbf{1}^T f, \quad f_j \geq 0, \quad j = 1, \dots, n, \end{aligned} \tag{12}$$

with optimization variables f_j, F . This is a convex optimization problem; its special structure allows it to be solved extremely efficiently, via waterfilling.

4.2 Solution via waterfilling

We can solve (12) using the waterfilling method (described earlier). At each time t , we are to maximize

$$\sum_{j=1}^n \left(\hat{V}_j(\theta_j z_j + (1 - \theta_j) f_j) - \nu f_j \right) - (\lambda \phi(F, g) - \nu F),$$

over variables $f_j \geq 0$, where as before, $\nu > 0$ is a Lagrange multiplier associated with the equality constraint. For our particular power law approximate functions we can express f_j in terms of ν :

$$f_j = \frac{1}{1 - \theta_j} \left(\left(\frac{\tilde{\alpha}_j \tilde{\beta}_j (1 - \theta_j)}{\nu} \right)^{1/(1 - \tilde{\alpha}_j)} - \theta_j z_j \right)_+.$$

We then use bisection on ν to find the value of ν for which

$$\mathbf{1}^T f = \mu \log(\nu \mu g / \lambda).$$

Since our surrogate value function is only approximate, there is no reason to solve this to great accuracy; experiments show that around 5–10 bisection iterations is more than enough.

Each iteration of the waterfilling algorithm has a cost that is $\mathcal{O}(n)$ which means that we can solve (12) very fast. An interior point method that exploits the structure would also yield a very efficient method; see, *e.g.*, [MBL10].

4.3 Upper bound policies

In this section we describe two heuristic data flow rate policies: A steady-state flow policy and a prescient flow policy. We show that both policies result in upper bounds on J^* (the optimal objective value). These upper bounds give us a way to measure the performance of our suboptimal flow policy φ^{adp} : If we obtain a J from φ^{adp} that is close to an upper bound, then we know that our suboptimal flow policy is nearly optimal.

4.3.1 Steady-state policy

The steady-state policy is given by

$$\varphi^{\text{ss}}(s, g_t) = \operatorname{argmax}_{f_t \geq 0} (U(s) - \lambda P_t), \tag{13}$$

where g_t is channel gain at time t , and s is the *steady-state* flow rate vector (independent of time) obtained by solving the optimization problem

$$\begin{aligned} & \text{maximize} && U(s) - \lambda\phi(s, \mathbf{E}g) \\ & \text{subject to} && s \geq 0, \end{aligned}$$

with optimization variable s , and λ known. Let J^{ste} be our steady-state upper bound on J^* obtained using the policy (13) to solve (3). Note that in the above optimization problem, we ignore time (and hence, smoothing), and variations in channel gains, and so for each λ , s is the optimal (steady-state) flow vector. (This is sometimes called the certainty equivalent problem associated with the stochastic programming problem [BL97, Pre95].)

By Jensen's inequality (and convexity of the max) it is easy to see J^{ste} is an upper bound on J^* . Note that once s is determined, we can evaluate (13) using the waterfilling algorithm described earlier.

4.3.2 Prescient policy

To obtain a prescient upper bound on J^* we relax the causality requirement imposed earlier on the flow policy in (4) and assume complete knowledge of the channel gains for all t . (For more on prescient bounds, see, *e.g.*, [BSS10].) For each realization of channel gains, the flow rate control problem reduces to the optimization problem

$$\begin{aligned} & \text{maximize} && \frac{1}{N} \sum_{\tau=0}^{N-1} \left(\sum_{j=1}^n U_j(s_\tau)_j - \lambda\phi(\mathbf{1}^T f_\tau, g_\tau) \right) \\ & \text{subject to} && s_{\tau+1} = \Theta s_\tau + (I - \Theta)f_\tau, \quad F_\tau = \mathbf{1}^T f_\tau, \quad f_\tau \geq 0, \quad \tau = 0, 1, \dots, N-1, \end{aligned} \tag{14}$$

where the optimization variables are the flow rates f_0, \dots, f_{N-1} , and smoothed flow rates s_1, \dots, s_N . (The problem data are s_0 and g_0, \dots, g_{N-1} .) The optimal value of (14) is a random variable. Let J^{pre} denote our prescient upper bound on J^* . We obtain J^{pre} by Monte Carlo simulation: We take N large and solve (14) for multiple, independent realizations of channel gains. The mean optimal value is our prescient upper bound.

4.4 Numerical example

In this section we compare the performance of our ADP suboptimal policy to the above heuristic policies using a simple numerical example. Let J^{adp} be the objective value we obtain using the ADP suboptimal policy.

We construct our multiple flow example from previous problem instance except that here, the lightly and heavily smoothed flows share a single link, *i.e.*, $n = 2$. We use $\lambda = 1$.

We carried out Monte Carlo simulation (100 realizations, each with 1000 time-steps).

We found that $J^{\text{adp}} = 0.47$ ($\bar{U}^{\text{adp}} = 0.69, \bar{P}^{\text{adp}} = 0.22$), and $J^{\text{pre}} = 0.5$ ($\bar{U}^{\text{pre}} = 0.74, \bar{P}^{\text{pre}} = 0.24$), so we know that $0.47 \leq J^* \leq 0.5$. So for this example, J^{adp} is no more than 0.03-suboptimal.

5 Conclusion

In this paper we present a variation on a multi-period stochastic network utility maximization problem as a constrained convex stochastic control problem. We show that judging flow utilities dynamically, *i.e.*, with a utility function and a smoothing time scale, is a good way to account for network applications with heterogenous rate demands.

For the case of a single flow, our numerically computed value functions obtain flow policies that optimally trade-off average utility and average power. We show that simple power law functions are reasonable approximations of the optimal value functions, and that these simple functions obtain near optimal performance.

For the case of multiple flows on a single link (where the value function is not practically computable using dynamic programming), we approximate the value function with a combination of the simple one-dimensional power law functions. Simulations, and comparison with upper bounds on the optimal value, show that the resulting ADP policy obtains near-optimal performance.

Acknowledgment

We thank Yang Wang and Dan O'Neill for helpful discussions.

References

- [ABFG⁺93] A. Arapostathis, V. Borkarand, E. Fernández-Gaucherand, M. K. Ghosh, and S. I. Marcus. Discrete-time controlled Markov processes with average cost criterion: A survey. *SIAM J. on Control and Optimization*, 31(2):282–344, 1993.
- [Åst70] K. Åström. *Introduction to Stochastic Control Theory*. Dover, 1970.
- [Bel57] R. Bellman. *Dynamic Programming*. Courier Dover, 1957.
- [Ber05] D. Bertsekas. *Dynamic Programming and Optimal Control: Volume 1*. Athena Scientific, 2005.
- [Ber07] D. Bertsekas. *Dynamic Programming and Optimal Control: Volume 2*. Athena Scientific, 2007.
- [BL97] J. Birge and F. Louveaux. *Introduction to Stochastic Programming*. Springer, 1997.
- [Bla62] D. Blackwell. Discrete dynamic programming. *Ann. Math. Statist.*, 33:719–726, 1962.
- [BS96] D. Bertsekas and S. Shreve. *Stochastic optimal control: The discrete-time case*. Athena Scientific, 1996.
- [BSS10] D. Brown, J. Smith, and P. Sun. Information relaxations and duality in stochastic dynamic programs. *Operations Research*, 2010. To appear.
- [BT96] D. Bertsekas and J. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, 1996.
- [BV04] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [Cav72] J. Cavers. Variable-rate transmission for Rayleigh fading channels. *IEEE Transactions on Communication Technology*, 20(1):15–22, 1972.
- [CGH⁺96] R. Corless, G. Gonnet, D. Hare, D. Jeffrey, and D. Knuth. On the Lambert W function. *Advances in Computational Mathematics*, 5:329–359, 1996.
- [CLCD07] M. Chiang, S. Low, A. Calderbank, and J. Doyle. Layering as optimization decomposition: A mathematical theory of network architectures. *Proceedings of the IEEE*, 95:255–312, 2007.
- [Den82] E. Denardo. *Dynamic Programming: Models and Applications*. Prentice-Hall, 1982.
- [Der70] C. Derman. *Finite state Markovian decision processes*. Academic Press, 1970.

- [dFR03] D. de Farias and B. Van Roy. The linear programming approach to approximate dynamic programming. *Operations Research*, 51(6):850–865, 2003.
- [GC97] A. Goldsmith and S. Chua. Variable rate variable power MQAM for fading channels. *IEEE Transactions on Communications*, 45(10):1218–1230, 1997.
- [Hay68] J. Hayes. Adaptive feedback communications. *IEEE Transactions on Communication Technology*, 16(1):29–34, 1968.
- [Hen74] V. Hentinen. Error performance for adaptive transmission on fading channels. *IEEE Transactions on Communication Technology*, 22(9):1331–1337, 1974.
- [KMT97] F. Kelly, A. Maulloo, and D. Tan. Rate control for communication networks: Shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49:237–252, 1997.
- [LL99] S. Low and D. Lapsley. Optimization flow control I: Basic algorithms and convergence. *IEEE/ACM Transactions on Networking*, 7(6):861–874, 1999.
- [Man60] A. Manne. Linear programming and sequential decisions. *Management Science*, 6(3):259–267, 1960.
- [MBL10] R. Madan, S. Boyd, and S. Lall. Fast algorithms for resource allocation in wireless cellular networks. *IEEE Transactions on Networking*, 2010. To appear.
- [NML08] M. Neely, E. Modiano, and C. Li. Fairness and optimal stochastic control for heterogeneous networks. *IEEE/ACM Transactions on Networking*, 16(2):396–409, 2008.
- [OABG10] D. O’Neill, E. Akuiyibo, S. Boyd, and A. Goldsmith. Optimizing adaptive modulation in wireless networks via multi-period network utility maximization. In *Proceedings of International Conference on Communications*, 2010. To appear.
- [Pow07] W. Powell. *Approximate dynamic programming: Solving the curses of dimensionality*. John Wiley & Sons, 2007.
- [Pre95] A. Prekopa. *Stochastic Programming*. Kluwer Academic Publishers, 1995.
- [Put94] M. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 1994.
- [Ros83] S. Ross. *Introduction to Stochastic Dynamic Programming: Probability and Mathematical*. Academic Press, 1983.
- [SS85] P. Schweitzer and A. Seidmann. Generalized polynomial approximations in Markovian decision processes. *Journal of Mathematical Analysis and Applications*, 110:568–582, 1985.

- [TZ97] M. Trick and S. Zin. Spline approximations to value functions: A linear programming approach. *Macroeconomics Dynamics*, 1, 1997.
- [TZB08] N. Trichakis, A. Zymnis, and S. Boyd. Dynamic network utility maximization with delivery contracts. In *Proceedings IFAC World Congress*, pages 2907–2912, 2008.
- [WB09] Y. Wang and S. Boyd. Performance bounds for linear stochastic control. *Systems and Control Letters*, 58(3):178–182, 2009.
- [Whi82] P. Whittle. *Optimization Over Time: Dynamic Programming and Stochastic Control*. Wiley, 1982.
- [WS95] W. Webb and R. Steele. Variable rate QAM for mobile radio. *IEEE Transactions on Communications*, 43(7):2223–2230, 1995.