

Minimum Energy Transmission Scheduling subject to Deadline Constraints

Alessandro Tarello¹
Dipartimento di Elettronica
Politecnico di Torino
Torino, Italy
e-mail: tarello@mail.tlc.polito.it

Jun Sun¹, Murtaza Zafer and Eytan Modiano
Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
Cambridge, MA
e-mail: {junsun,murtaza,modiano}@mit.edu

Abstract

We consider the problem of transmission scheduling of data over a wireless fading channel with hard deadline constraints. Our system consists of N users, each with a fixed amount of data that must be served by a common deadline. Given that, for each user, the channel fade state determines the throughput per unit of energy expended, our objective is to minimize the overall expected energy consumption while satisfying the deadline constraint. We consider both a linear and a strictly convex rate-power curve and obtain optimal solutions, based on dynamic programming (DP), and tractable approximate heuristics in both cases. For the special non-fading channel case with convex rate-power curve, an optimal solution is obtained based on the Shortest Path formulation. In the case of a linear rate-power curve, our DP solution has a nice “threshold” form; while for the convex rate-power curve we are able to obtain a heuristic algorithm with comparable performance with that of the optimal scheduling scheme.

1 Introduction

Increasing data transmission rates provides benefits in terms of efficiency of the bandwidth utilization and in terms of range and quality of services offered to the users. However, sending data at the maximum rate often decreases energy efficiency. Since many mobile

wireless devices are battery powered, and the energy for transmission is often a major source of energy consumption, saving energy during transmission can lead to a significant performance improvement, resulting in smaller batteries or longer battery lifetimes. Thus, a well designed mobile transmitter must not only maximize data throughput but also optimize the use of resources, effectively cope with fading channel, and meet operational constraints.

In this paper, we consider the problem of minimizing the energy expended by a transmitter in a wireless network serving multiple users within a given amount of time. A number of realistic wireless networking scenarios can be modeled in this way. For example, the transmitter can be a “backbone” node, which serves as an access point for a set of end nodes, with limit amount of energy and finite life time in an *ad-hoc network*, or a satellite with time-critical data (i.e., information that must be sent before certain deadline) as described by the authors in [1].

The communication system considered in this paper consists of a single transmitter sending data to N users, each user has a fixed amount of data that must be received by a deadline. Time is assumed to be discrete. During each time slot, at most one user can be chosen for transmission. Our objective is to select the user to transmit, and the associated data rate, so that the total energy consumption is minimized subject to a deadline constraint. The key drivers behind energy savings in this setting are : i) *varying transmission rate* over the time horizon since, as addressed in [2] and [3], it is possible to reduce the energy consumption by lowering transmission power and transmitting data over a longer period of time and ii) *the use of oppor-*

¹This research was supported in part by NASA under the Space Communication Project, grant number NAG3-2835, and in part by Politecnico di Torino, Torino, Italy. The research was conducted while A. Tarello was visiting LIDS at MIT.

tunistic scheduling to exploit the channel variability, as addressed in [4].

The problem of resource allocation in wireless networks has received much attention in recent years. Resource allocation for fading multi-user broadcast channels is a popular topic in information theory. However, the resource being allocated is usually average power or bandwidth, and the quantity to be maximized is most often Shannon Capacity. The problems of maximizing throughput given a finite amount of energy, or minimizing energy consumption given a short-time deadline, have not been well-studied, to the best of our knowledge. More recently, in [4], the authors try to maximize the data throughput of an energy and time constrained transmitter communicating to a single user over a fading channel. They also explored the problem of using the minimum amount of energy to send a fixed amount of data to a single user by a deadline (i.e., the single user version of the problem that we address here). In [2], the authors use a calculus approach to obtain energy efficient transmission policies with arbitrary arrival and deadline constraints over a time-invariant channel. In [3], the authors consider the problem of sending a random number of packets that arrive in a time interval by a deadline using the minimum amount of energy. The transmission channel is assumed to be time invariant in their setup. Other works that address the similar problem include [9] and [10].

This paper is organized as follows. In section II, we present the details of the problem formulation. In section III, the minimum energy scheduling scheme is obtained as the solution of a shortest path problem when channels are known and time invariant. Cases where all users have the same deadline and where different users has different deadline are both studied. In section IV, we present the minimum energy scheduling scheme for the case where users' channels are unknown and the rate-power curve is linear. Heuristic algorithm is given for the case where user's rate-power curve is convex. Section V concludes this paper.

2 Model Description

We consider a system with a single transmitter sending data to N users through N time-varying channels. The channel for each user c , $c \in \{1, \dots, N\}$, can be in

one of a finite set S_c of states. The channel process is represented by a vector $\bar{Q}(t) = (q_1(t), \dots, q_N(t))$, where $\bar{Q}(t) \in S_1 \times \dots \times S_N$. The time axis is discretized: channels hold their states for time slots of length T , with transitions on the boundaries $t = kT$. The channels' states are assumed to be known at the beginning of each time slot, either through direct measurement or through a combination of measurement and prediction. Furthermore, we assume no correlation between channels of different users. Each user has an amount of data that must be transmitted by a deadline. The transmitter controls the consumed energy by adjusting the rate allocation vector $\bar{\Gamma}(t) = (\mu_1(t), \dots, \mu_N(t))$, subject to the constraints that only one user can transmit in each slot and that all of the data must be transmitted by the deadline. The rate per time slot assigned to the user j at time slot k is μ_{jk} . For any given state q_{jk} of the channel j at time slot k , there is a rate-power curve $f(\mu_{jk}, q_{jk})$ representing the amount of energy required to transmit at rate μ_{jk} when the channel is in state q_{jk} .

In this paper we address two cases of rate-power curves. The first one is when $f(\mu_{jk}, q_{jk})$ is linear in $\frac{\mu_{jk}}{q_{jk}}$, so that q_{jk} can be interpreted as the rate obtained per unit of consumed energy. This kind of curve is an accepted assumption in the low SNR regime or for ultra-wideband transmissions. A linear rate-power curve implies that the service of each user can be completed within a single time slot. Theoretically, there is no limit on the energy that might be consumed during a single time slot; practically, situations characterized by small amount of data in short time intervals make this assumption more realistic (the concave Shannon curve can be approximated by a linear curve around the origin).

The second case is when $f(\mu_{jk}, q_{jk})$ is strictly convex in μ_{jk} (see Fig. 1). The convexity is a reasonable assumption due to the following two factors. First, the Shannon capacity for an AWGN channel is a logarithmic function of the expended energy; second, under a fixed modulation scheme, throughput has a linear relation to the expended energy, and since the curve could represent a set of coding schemes, with a power limitation, the curve becomes piecewise linear and convex. The convexity reflects increasing costs in signal power with each incremental increase in transmission rate: this makes the spreading of the service over more

time slots less energy costly than concentrating it on a single slot.

The goal of this paper is to find a transmission schedule that minimizes the expected consumed energy, subject to a constraint on the minimum amount of data to serve for each user and a deadline by which it must be transmitted. We consider the time window constraint composed by K time slots, $K \geq N$. The optimization problem becomes

$$\min E \left[\sum_{j=1}^N \sum_{k=1}^K f(\mu_{jk}, q_{jk}) \cdot \tau_{jk} \right] \quad (1)$$

subject to the constraint that at least the initial amount of data d_j for each queue is served within a finite time window:

$$\sum_{k=1}^K \mu_{jk} \tau_{jk} \geq d_j \quad \forall j = 1, \dots, N \quad (2)$$

$$\sum_{j=1}^N \tau_{jk} \leq 1 \quad \forall k = 1, \dots, K \quad (3)$$

where τ_{jk} is equal to 1 if the queue j has been served during the time slot k , 0 otherwise. The inequality (2) expresses that the service of all the users has to be completed within the frame of K time slots, while (3) that at most one user per slot can be served. For the remainder of this paper, if not otherwise specified, we will also assume $d_j = d$ for all users.




Figure 1. Set of rate-power curves $f(\mu_{jk}, q_{jk})$ for users 1, 2 and 3 at time k .

3 Known Channel Quality

Let us assume that the channels' quality is completely known, namely q_{jk} is known for all the users

and all the time slots. For each user j , the rate-power curve $f(\mu_{jk}, q_{jk})$ is then a family of convex curves in μ_{jk} indexed by time:

$$f(\mu_{jk}, q_{jk}) = f_{jk}(\mu_{jk})$$

The problem (1) can be restated as minimizing

$$\sum_{j=1}^N \sum_{k=1}^K f_{jk}(\mu_{jk}) \cdot \tau_{jk} \quad (4)$$

subject to the constraints (2) and (3). For the optimal energy minimization policy, we can see that constraint (3) will be met with equality due to the convexity of the rate-power function. Note that the target of the optimization problem is now the minimization of the energy effectively consumed, instead of its expectation, since a complete knowledge of the channel in the future is assumed.

While future knowledge of the channel states is usually not realistic, the optimal solution to the scheduling problem with full knowledge provides a lower-bound on the "on-line" policies presented later. It also shows the additional energy cost associated with lack of knowledge of the channel state.

3.1 Time-Invariant Known Channel

Let us first consider the case when the channel is known and time invariant, i.e. the channel quality for each user j is equal to a constant value Q_j . Knowing the fade state is equivalent to knowing the rate-power curve for each user; a constant fade state implies that each user maintains the same rate-power curve for all the time considered. This is a realistic assumption for the slow fading wireless channel [5].

To solve the optimization problem stated in Eq.(4) with the time invariant channel assumption, the scheduler must decide: 1) to which user a particular time slot should be allocated and 2) how much power to spend. Due to the time invariant channel, the time at which a single user is served, and therefore the order with which users are served, becomes irrelevant; note that this holds only if all the users have the same deadline. Also, given that m slots were assigned to a particular user with rate-power curve $f(\mu)$, to send d amount of data in m slots with minimum energy, the optimal policy is to consume an equal amount of power in each

of these m time slots. This can be shown easily using the convexity of the $f(\mu)$. Hence, to find the optimal energy minimization policy, we need to consider only one factor: the number of time slots allocated to each user.

Since the time horizon considered here is finite, we can formulate this energy minimization problem into a deterministic shortest path problem. Due to the irrelevance of the service order, we assume that the scheduler serves a user in consecutive time slots until all data of that user are transmitted (i.e., user 2 will not be served until all the data of user 1 are sent). The state variable x_n is then defined to be the time at which the scheduler finishes serving user n . The state space, denoted here as X_n , is therefore a finite set: indeed, x_n can only assume integer values within the set $\{n, \dots, K - (N - n)\}$. The value of x_n cannot be less than n because n users cannot be served using less than n time slots, which is result of the fact that at most one user per time slot can be served; similarly, x_n cannot be greater than $K - (N - n)$ since sufficient number of time slots (i.e., $N - n$ time slots) must be reserved to finish the data transmission for the remaining $N - n$ users. At any state x_n a control $^1 w_n$ decides how many time slots will be used to serve the next user, the $(n + 1)$ st, and can be associated with a transition from the state x_n to the state x_{n+1} . This transition has a cost. If each user has an amount of data d to transmit, the cost of a transition from the node $x_n = i$ to the node $x_{n+1} = j$ is the energy required to serve the $(n + 1)$ st user over $j - i$ time slots, i.e. $(j - i)f(\frac{d}{j-i}, Q_{n+1})$, where the function f is the considered rate-power curve, the first component is the rate required to transmit the amount d of data over $j - i$ time slots and Q_{n+1} is the channel quality of the $(n + 1)$ st user². Notice that, given a state x_n , the effect of using a certain number of slots to serve the next user is perfectly predictable because the channel is constant. Such a finite-state deterministic problem can be equivalently represented by a graph G as in Fig. 2, with the following properties:

- The graph is composed of N stages; the generic stage n corresponds to the state space X_n , namely

¹In the time-varying channel formulation the decision variable is u_n and represents which user is served

²Note that, due to the convexity of the rate-power curve, in general $(j - i)f(\frac{d}{j-i}, Q_{n+1}) < f(d, Q_{n+1})$

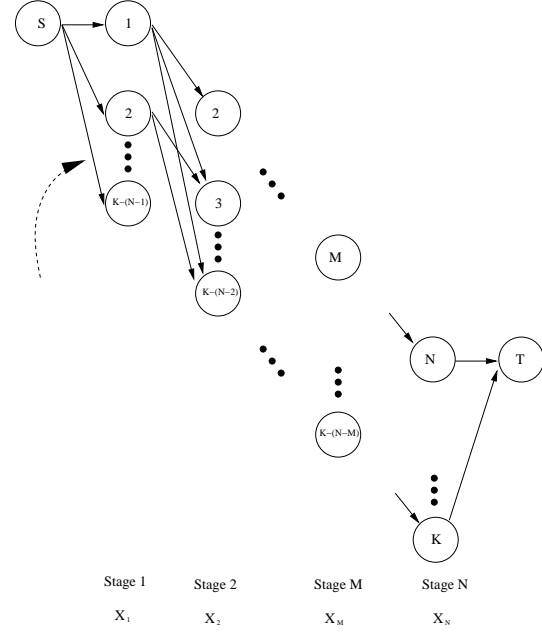


Figure 2. A graph G generated considering N users with a deadline of K time slots. Consider the arc pointed by the dash arrow: $(K - (N - 1))f(\frac{d}{K-(N-1)}, Q_1)$ is the cost associated to that transition, namely it is the energy consumption to serve the first user over the first $K - (N - 1)$ time slots.

the nodes of the stage n represent the set of feasible values that x_n may take.

- The arcs correspond to transitions between states at successive stages, and each arc has an associated cost equal to the cost of the represented transition; the cost of transition towards the n th stage, from a node $x_{n-1} = i$ to a node $x_n = j$, represents the additional energy consumption to serve the n th user in $j - i$ time slots. At the state $x_n = j$, n users have been served over j time slots.
- An initial state at which no user has been served and no time slots has been used is connected to every state of the first stage. The cost of each arc is the energy required to serve the first user using t time slots, with $t \in \{1, \dots, K - N + 1\}$.

Theorem 1. Given a graph G built according to the presented procedure, the shortest path from the initial state to the last stage represents the optimal solution to

the optimization problem (1) in case of time-invariant channel.

Proof. The graph includes all of the feasible solutions to the scheduling problem. Each arc from one stage to the next one represents the service of an additional user using a certain amount of time, which determines the energy cost of that transition. A path is a sequence of arcs, each one associating a certain amount of time slots to the service of each user: at the end of each path, all the N users have been served using all the K time slots available³, and the sum of the arc weights is the total cost of that particular scheduling associated to that path. The optimal cost is clearly achieved applying the shortest path algorithm to the graph.

Notice that i) the procedure used to generate the graph G includes all the feasible solutions and excludes all the infeasible ones, namely all the solutions that do not fulfill the deadline constraint; the generic i^{th} user cannot be served in such a way that the total amount of time spent to serve i users exceeds $K - (N - i)$ time slots, otherwise the service of the remaining users could not be completed ii) the particular order with which users are served does not affect the result of the optimization problem, since all the users transmit on a constant channel and have the same deadline; the optimal scheduling is a set of frames associated with the users, and the service of each user over the number of time slots specified by a frame does not change depending on the particular time slot at which the service is offered. Point ii) is no longer true if users have different deadlines, because in this case the set of feasible solution is affected by the order of the nodes. \square

The scheduling algorithm presented in Theorem 1 deals with a single deadline and the same amount of data d for all the users. It is possible to extend this result to the case of multiple deadlines. Here, we consider only multiple deadlines that are feasible (i.e., there exists at least one scheduling scheme that will meet all users' deadlines). This multiple deadlines problem can be formulated into the shortest path problem in the following two steps:

1. Order the service of the users on a earliest deadline first basis
2. Build a graph G according to the procedure previously proposed, and associate a weight equal to ∞ to those arcs of the graph G that do not fulfill the deadline constraints. Specifically, consider the transition from the node $x_n = i$ to the node $x_{n+1} = j$: if the deadline for user $(n + 1)$, say t_{n+1} , is less than j (i.e., user $(n + 1)$'s deadline is not met), we will assign an infinite weight to that transition.

We call the obtained graph with the above modified procedure G_m . It is clear that the shortest path algorithm will avoid all the paths in which one or more arcs has infinite weight, namely the paths that do not meet the deadlines. While for the scheduling algorithm of Theorem 1 the order with which users are served does not affect the performance in terms of energy consumption, in the case of multiple deadlines the service order matters. However, constraint (3) will still be met with equality (i.e., the transmitter will not be idle during any time slot). Also, given that m time slots were allocated to user i , the transmitter will consume an equal amount of energy in each of the m time slots. The following theorem shows the existence of an optimal policy using the procedure described above.

Theorem 2. *In the case of multiple deadlines and time-invariant channel:*

- a) *the set of feasible solutions to the optimization problem (1) obtained ordering users on a earliest deadline first basis still includes the optimal solution*
- b) *the optimal solution is achieved applying the Shortest Path formulation on the graph G_m as stated in Theorem 1.*

Proof. The ordering operation on a earliest deadline first basis generates the set of all feasible solutions. It is straightforward to see that any other feasible scheduling scheme that meets all deadline constraints can be transformed to the earliest deadline scheme. Since we are considering only feasible multiple deadlines and the number of states is finite, there exists an optimal solution to the shortest path problem. \square

In this section, for simplicity, we assume that each user has an equal amount of data that needs to be

³Since we consider the general case of convex rate-power curve, exploiting all the available time is a necessary condition of optimality

served by the transmitter. However, the previous two theorems hold even when different users have different amount of data to be served by the transmitter.

3.2 Time-Varying Known Channel

We, now, consider the case of time-varying but known channel quality. Specifically, the channel state q_{jk} is different over the K time slots but known in advance. We specialize to the case of linear rate power curve, i.e. $f(\mu_{jk}, q_{jk}) = \frac{\mu_{jk}}{q_{jk}}$. Suppose that there is only one user, then, the optimal policy is to simply select the slots with the best channel quality and transmit the data in these slots. Since the rate-power curve is linear, transmitting the data in a single slot or multiple slots with the best channel does not affect the energy cost. Clearly, transmitting in any other slot that does not have the best channel state is not optimal. Thus, for a single user the optimal policy is to simply select the slot with the best channel state and empty the entire data. Now, extending this result to multiple users is not trivial, since, slots with the best channel for the different users may overlap in time. In this case, the optimal policy can be obtained as follows.

Consider a bipartite graph with one set of nodes (inputs) representing the N users and the other set of nodes (outputs) representing the K time slots. As the channel state of each user is known for all the time slots, we can obtain the transmission energy required to serve the entire data of a user in a particular time slot. Let the edge connecting the input node j (user j) with the output node k (time slot k) have weight equal to c_{jk} , where c_{jk} is the energy required to serve all the data of user j in time slot k . Thus for each user j we have K edges connecting to each time slot and the weight (or the energy cost) of each edge is also known. Clearly, the optimal policy now is simply a *Minimum Weight Matching* applied to the above graph.

4 Minimum Energy Scheduling with a Stochastic Channel Process

We now examine problem (1) with a stochastic channel process. Here, the channel state of each user is a random process over time but is assumed to be known at the beginning of each time slot; i.e. q_{jk} is assumed known at the beginning of time slot k for each

user j but is unknown for future time slots. The channel state of each user is assumed independent of other users and is also independent over time. First we consider a linear rate-power curve and then in the next section investigate the more general case of a convex rate-power relationship.

4.1 Linear Rate-Power Curve

The power expenditure is assumed to be a linear function of the transmission rate, i.e. $f(\mu_{jk}, q_{jk}) = \frac{\mu_{jk}}{q_{jk}}$. With a linear rate-power curve it can be shown that it is optimal for the transmitter to transmit the entire data of a user in just one slot. Intuitively, this stems from the fact that the transmitter will only decide to serve a user during a given time slot if the channel quality during that time slot is better than the expected channel quality that the user will experience in the future. Due to the linearity of the rate power curve, the cost per unit data is constant during the time slot. Hence, if the transmitter decided to serve a given user, then it follows that it should transmit all of the data belonging to that user. Mathematically, this follows from the fact that the dynamic programming value function, which measure the expected energy cost to serve that user in the future, can be shown to be concave in the amount of data that remains to be served. This concavity, when combined with the linear rate power relationship, is sufficient to show that the optimal solution is either not to serve the user or to transmit all of the user's data at once. We treat the convex rate-power curve case in Section 4.2.

We begin by first presenting an optimal on-line policy based on dynamic programming (DP) and then compare the performance of various simple heuristic policies. For simplicity we assume a constant amount of data d for each user. As each user is served within a single time slot, the value function for the DP recursion just depends on the number and the channel conditions of the remaining users. In the more general case where each user has a different amount of data, a similar approach can be applied, however that would involve keeping track of the amount of data that each user has to be served (or at least the order of the remaining users in terms of how much data they each have to be served). Let $\mathcal{J}_k(n_k, \underline{q}_k)$ be the cost of having n_k users remaining to be served at time slot k ,

where \underline{q}_k is the vector whose component $(q_j)_k$ is the channel state of user j at time slot k . The recursive DP equation for this problem can be written as,

$$\mathcal{J}_k(n_k, \underline{q}_k) = \min_{u_k} \left\{ \sum_{j=1}^N \frac{d}{q_{jk}} \mathbf{I}_j(u_k) + \overline{\mathcal{J}}_{k+1}(n_{k+1}) \right\} \quad (5)$$

where u_k is the decision taken by the server at time k (i.e., if a user j is served $u_k = j$), n_{k+1} is the number of remaining users after the decision ($n_{k+1} = n_k$ if no user is served, otherwise it equals $n_k - 1$), \mathbf{I} is the indicator function whose form is

$$\mathbf{I}_{j^*}(u_k) = \begin{cases} 1, & \text{if } u_k = j^* \\ 0, & \text{otherwise} \end{cases}$$

and

$$\overline{\mathcal{J}}_{k+1}(n_{k+1}) = E \left[\mathcal{J}_{k+1}(n_{k+1}, \underline{q}_{k+1}) \right] \quad (6)$$

is the expected future cost of the decision. The summation is, finally, the actual cost of the decision, which is either 0 if no user is served, or $\frac{d}{q_{j^*k}}$ if the user j^* is served. To complete the recursion, the *termination condition* is given as,

$$\overline{\mathcal{J}}_{K+1}(n_{K+1}) = \begin{cases} 0, & n_{K+1} = 0 \\ \infty, & n_{K+1} > 0 \end{cases}$$

where an infinite cost is imposed if there are unserved users remaining after time slot K .

Theorem 3. *For the multi-user stochastic channel scenario with a linear rate-power curve, the optimal on-line policy that minimizes the expected energy expenditure is a threshold policy of the form:*

$$u_k = \begin{cases} \text{idle} & \text{if } \frac{d}{q_{max,k}} > \alpha_k \\ U_{q_{max,k}} & \text{otherwise} \end{cases}$$

where

$$\alpha_k = \overline{\mathcal{J}}_{k+1}(n_k) - \overline{\mathcal{J}}_{k+1}(n_k - 1) \quad (7)$$

is a threshold that can be easily computed off-line, and $U_{q_{max,k}}$ is the selected user, whose channel is the best among all the channels associated with the remaining users at time k .

Proof. During each time slot the policy either serves one user or does not serve any user. This implies that the value function (5) can be written as,

$$\min_{u_k \in U} \left\{ \min_{j^*} \left(\sum_{j=1}^N \frac{d}{q_{jk}} \mathbf{I}_j(u_k) \right) + \overline{\mathcal{J}}_{k+1}(n_k - 1), \right. \\ \left. \overline{\mathcal{J}}_{k+1}(n_k) \right\} \quad (8)$$

where U is the set of users still waiting to be served. $\overline{\mathcal{J}}_{k+1}(n_k)$ is the expected future cost when all the users are delayed ($n_{k+1} = n_k$) while $\overline{\mathcal{J}}_{k+1}(n_k - 1)$ is the expected future cost when all but one user is delayed ($n_{k+1} = n_k - 1$). On the other hand, the cost for consuming energy at time k is $\frac{d}{q_{j^*k}}$, where j^* is the selected user. The user that minimizes the quantity $\frac{d}{q_{j^*k}}$ is simply the one with the best channel state among the set of remaining users. This choice is also justified by the fact that users have iid channel conditions and the only objective is to reduce the number of remaining users using the minimum amount of energy. Thus, the total cost of serving this user at time k is $\frac{d}{q_{max,k}} + \overline{\mathcal{J}}_{k+1}(n_k - 1)$. To satisfy (8), the optimal policy transmits on the best channel among the remaining users if

$$\frac{d}{q_{max,k}} + \overline{\mathcal{J}}_{k+1}(n_k - 1) \leq \overline{\mathcal{J}}_{k+1}(n_k) \quad (9)$$

where the left side represents the cost for serving one user, the one with the best channel, while the right side represents the cost of delaying all the services. The inequality (9) is equivalent to

$$\frac{d}{q_{max,k}} \leq \alpha_k \quad (10)$$

where α_k is defined as in (7). Thus the threshold is equal to the difference between the expected future cost of delaying all n_k transmissions or serving one user and delaying only $n_k - 1$ transmissions. Clearly, if the cost of serving a user during the current time slot is less than this difference, then it is better to serve a user during this slot, otherwise it is better to delay all transmissions. \square

The optimal policy is a threshold policy that requires the (pre) computation of the threshold for each time-step. To compare the performance of the optimal solution, we now present simple heuristics that

are based on the solution of the *optimal stopping time* problem as described in Appendix A. Notice that the problem we considered with only a single user reduces to a simple optimal stopping time problem. With multiple users, however, optimal stopping time solutions cannot be directly applied as it is possible that more than one user would have the same stopping time. The heuristic policies that we consider are as follows.

- *OptStop Max*: This policy calculates the optimal stopping time independently for each user, fixing the last time slot (the K^{th} slot) as the deadline for all the users. If two or more users have the same optimal stopping time the policy serves the one with the best channel. When the number of the remaining users is equal to the time slots remaining until the K^{th} slot, a greedy algorithm is applied that transmits in each slot the user with the best channel among the remaining users.
- *OptStop Dyn*: This policy is a modified version of *OptStop Max*. It is called "Dynamic" because it updates dynamically the deadline after the service of a user. Initially it calculates the optimal stopping time independently for each user by fixing the slot $K - N - 1$ as the first deadline for the N users. After the service of the first user (eventual collisions are always solved selecting the user with the best channel), the policy recalculates the optimal stopping times of the remaining $N - 1$ users moving the time deadline ahead by one slot and so on each time a user is served. As the previous heuristic, when the number of remaining users is equal to the remaining slots the same greedy algorithm is applied.
- *OptStop Rand*: This policy is similar to *OptStop Max* except that it solves the collisions randomly.

For comparison, we can also consider the optimal algorithm, obtained in Section 3.2, that assumes complete future knowledge of the channel state evolution. Comparison with a policy that assumes channel knowledge helps understand the additional energy cost of the online policies due to lack of such knowledge.

Simulation Results

In Figure 3 we compare the performance of the optimal on-line algorithm with the heuristic algorithms

previously explained. The average is obtained generating 500 different channel state trajectories for each value of the window constraint from 16 to 76. It can be noticed that the optimal off-line policy is obviously a lower bound. The best performance possible is represented by the optimal on-line policy which turns out to be very close to the ideal performance of the system represented by the lower bound. When K is large, the performance of the stopping-time heuristics is very close to the optimum since the probability of collision between the calculated optimal stopping times is low. A worse performance is expected when the probability of collision increases, namely when the time window constraint is shorter. In this case, the strategies to solve collisions characterize the heuristics' sub-optimality. Notice that the performance of *OptStop Max* and *OptStop Dyn* is very good even when the deadline is small. However the performance of *OptStop Rand* worsens but this is not surprising for a policy where collisions are randomly solved instead of choosing the best channel.

Let us focus now on a single channels' state trajectory for the previous setting. The scenario consists of $K = 24$ time slots. In Figures 4, 5 and 6 we compare the transmission schedule obtained by the optimal on-line policy and the two heuristics *OptStop Max* and *OptStop Rand*. On each figure the line representing the threshold between saving or consuming energy and the sequence of best channels among those associated to the remaining users are reported. Notice that the two heuristics have exactly the same threshold, since it is calculated in the same way (optimal stopping time strategy), and the difference in the policies is determined by the way collisions are solved. While *OptStop Max* solves collisions transmitting on the best channel among those eligible to be served, the *OptStop Rand* chooses randomly. At each time slot, all those users whose channel is above the threshold are eligible to be served: in the case of the optimal on-line policy and *OptStop Max* the user with the best channel is selected; in the case of *OptStop Rand* the user to serve is randomly selected. If all the channels associated with the remaining users are under the threshold no user is served. Note that, since *OptStop Rand* does not serve the user with the best channel among those eligible to be served, we show the quality of the used channel in Figure 6, which may not be the best one (only when-

ever no user is served the best channel is shown). It can be noticed that *OptStop Rand* immediately under performs at the first time slot where it does not transmit using the best channel.

Finally notice, that the threshold of the optimal on-line policy is a function of the number of the remaining users and this intrinsically forces the system to respect the delay constraint, thereby, serving all the users by the short term deadline. In contrast, the threshold of the two heuristics is calculated for each user in isolation. While this leads to an easier pre-computation phase it may lead to situations in which there are not enough slots remaining to exploit the channel variation over time. To satisfy the deadline constraint, if the residual time in terms of number of time slots is equal to the number of remaining users, the threshold is no longer considered and a greedy service is applied. In Figure. 6, this can be observed for the policy *OptStop Rand* in the last two time slots.

The computation of the optimal on-line policy turns out to be quite efficient. However, when the number of users is large (more than 100) it can be substantially slower than the heuristics. In particular, the pre-computation phase to calculate the threshold is the most critical compared to the simpler heuristic algorithms, where a threshold for each user in isolation or a search for the best channel over few time slots has to be implemented. In contrast, the heuristics are slightly more complicated during service because they have to manage the collisions.

4.2 Convex Rate-Power Curve

In this section, we consider the case of a convex rate-power curve, i.e. $P_{jk} = f(\mu_{jk}, q_{jk})$, where $f()$ is convex in rate. While the linear rate-power curve implies that there is no limit on the amount of energy that can be consumed during a single time slot and so there is no limit on the amount of data of a single user that can be served, the convex curve introduces a smooth power limitation. Theoretically any user may still be served within a single time slot but each increment in the data rate produces an increasing additive cost of signal power. Thus it might be better for the server to spread the service of a single user over multiple slots. The problem is now clearly more complex as we must also keep track of the amount of data that remains to




Figure 3. Average consumed energy vs time window constraint, fading Rayleigh distributed channels with mean 20, 16 users, linear rate-power curve, $d = 450$ units of data

be transmitted by each user. This means a state space whose vectorial component at time k is $(\underline{d}_k, \underline{q}_k)$, where each component j of \underline{d}_k is the residual amount of data at time k for user j and each component of \underline{q}_k is the corresponding channel state.

A dynamic programming formulation as in (4.1) can be formulated. The cost function satisfies the following recursion,

$$\mathcal{J}_k(\underline{d}_k, \underline{q}_k) = \min_{\substack{u_k \\ 0 \leq \mu_{u_k k} \leq d_{u_k k}}} \left\{ \sum_{j=1}^N f(\mu_{jk}, q_{jk}) \mathbf{I}_j(u_k) + \bar{\mathcal{J}}_{k+1}(\underline{d}_{k+1}) \right\} \quad (11)$$

where u_k is the decision taken by the server at time k , $\mu_{u_k k}$ is the transmitted amount of data of the selected user, the summation is the actual cost of the decision (0 if no user is served and $f(\mu_{jk}, q_{jk})$ if user j is served) and

$$\bar{\mathcal{J}}_{k+1}(\underline{d}_{k+1}) = E \left[\mathcal{J}_{k+1}(\underline{d}_{k+1}, \underline{q}_{k+1}) \right] \quad (12)$$

is the expected future cost of the decision. The *termination condition* of the recursion imposes an infinite cost if one or more users still have data to be sent after the deadline:

$$\bar{\mathcal{J}}_{K+1}(\underline{d}_{K+1}) = \begin{cases} 0, & \underline{d}_{K+1} = \underline{0} \\ \infty, & \text{otherwise} \end{cases}$$




Figure 4. Transmission schedule obtained by the optimal on-line policy




Figure 5. Transmission schedule obtained by *Opt-Stop Max*




Figure 6. Transmission schedule obtained by *Opt-Stop Rand*

This problem does not allow a semi-analytical solution as in (4.1) where we are able to calculate a threshold and then implement an optimal on-line policy. In this case, we can numerically solve the recursion. The set of rate-power curves that we consider here are the standard curves for Shannon capacity

$$\mu_{jk}(c_{jk}, q_{jk}) = \log(1 + q_{jk}c_{jk}) \quad (13)$$

where, as usual, c_{jk} is the power per time slot required for user j at time slot k to transmit μ_{jk} units of data and q_{jk} is the corresponding channel state depending on the received signal-to-noise level for channel j . Since the optimal solution is the result of an intensive numerical computation, we also propose a practical and computationally less intensive heuristic and compare its performance with the optimal solution. The heuristic we propose is inspired by the single-user infinite-horizon version of the minimization problem addressed in this section.

Infinite Horizon Heuristic

For the moment, let us ignore the deadline constraint and focus on a *single user*. We want to minimize the expected consumed energy guaranteeing a certain long term throughput to the user. Therefore, during each time slot, the transmitter need to choose a data rate based on the channel state of that time slot. The formulation of this problem is given by the following:

$$\begin{aligned} \min & \sum_q P(q) \frac{\exp(\mu_q) - 1}{q} \\ \text{s.t.} & \sum_q P(q) \mu_q = LTRG \end{aligned} \quad (14)$$

where the optimization is taken over the values of μ_q , the data rate chosen when the channel is in state q . The function to be minimized represents the expected energy cost averaged over all possible discrete channel values, and the constraint represents the average long-term rate obtained. $P(q)$ is the probability mass function of the channel state, and $LTRG$ is the Long Term Rate Guarantee. The sum in the objective function is a simple expectation of the consumed energy: for each possible value of the channel state q , the energy consumption to transmit the amount of data μ_q is evaluated by inverting the convex rate-power curve of (13);

each amount of energy is then weighted with the probability of the channel state, and the sum gives the expected energy consumption. The solution of this optimization problem will give us, for any channel state q , the value μ_q of the data rate to offer, in order to meet the long term rate guarantee, subject to minimizing the expected consumed energy. The nonlinear optimization problem can be computed efficiently because the objective function is convex in the energy variable for every fixed channel state q . Using standard Lagrange multiplier techniques as in [8], it can be shown that a solution is optimal if and only if the rate is allocated according to the constraints so that the derivatives $P(q) \frac{d(\frac{\exp(\mu_q)-1}{q})}{d\mu_q}$ are equalized to some value γ^* for all channel values q that receive some service, while all channel values receiving zero rate have derivatives greater than γ^* . Given the convex rate-power curve (13), the solution to (14) is:

$$\mu_q = \frac{LTRG - \sum_{q' \in Q} P(q') \log\left(\frac{q'}{P(q')}\right)}{\sum_{q' \in Q} P(q')} + \log\left(\frac{q}{P(q)}\right) \quad (15)$$

for channel state values $q \in Q$. Initially the set Q includes all channel state values. The above equation produces the optimal rate allocation whenever the resulting μ_q are nonnegative. If any value is negative, it is set to zero, and the corresponding channel state is removed from Q . The calculation is repeated at most $N_q - 1$ times, being N_q the number of possible channel state values.

Now, to solve the multi-users problem with convex rate-power curve, the transmitter has to make the following two decisions at the beginning of each time slot: 1) to whom the time slot should be allocated. 2) the transmission rate for a particular time slot given this time slot is already assigned to a user. The idea behind our heuristic algorithm for the multiple users problem is to address the first question by selecting the user with the best channel state, and address the second question by applying the rates obtained for the single user problem (i.e., the optimal solution to (14) with appropriately chosen LTRG parameter). The LTRG parameter is updated at each time slot on the basis of the amount of data already served and the remaining time before the deadline. Specifically, our heuristic algorithm consists of the following three steps: i) At the

beginning of each time slot, among users with remaining data, the one with best channel state will be chosen to be served ii) The selected user is considered in isolation in computing its allocated data rate according to (15). The LTRG parameter can be obtained from the ratio of the amount of data that remains to be served and the number of time slots left iii) If the number of remaining time slots is equal to the number of remaining users to be served, a greedy policy is applied.

In step (ii) above, the allocated rate is the optimal solution to the optimization problem (14). However, a minor subtlety arises here. The channel state distribution used in (ii) is not simply that of a single user in isolation. Rather, since each user is chosen for service only when it has the best channel among all users, the distribution used in (14) must be adjusted accordingly. In particular, at each time slot the observed channel is the maximum among n channels, where n is the number of users remaining to be served, with probability $\frac{1}{n}$, and is zero with probability $\frac{n-1}{n}$, since the channel is not considered for transmission in that time slot.




Figure 7. Average consumed energy vs time window constraint; Numerical Solution of the DP formulation and Infinite Horizon Heuristic for the multi-user, fading channel scenario with nonlinear rate-power curve.

Comparing our heuristic with the optimal solution, we see that the two results are very close as the time horizon increases. As an example, the numerical solution of the optimal on-line problem is obtained for $N = 3$ users. The channel throughput per unit of

energy q_{jk} is integer valued and Rayleigh distributed with a mean of 2 during each time slot. The amount of data to serve is integer valued with the unit as minimum granularity. Their initial values are 2, 1 and 3, respectively for the three users. The average consumed energy is presented in Fig. 7, obtained by generating 1000 different channel state trajectories for each value of the time window constraint. The time window constraint length is moved from 6 to 450 time slots. The results of the optimal policy are compared with the results of the infinite horizon heuristic. As expected, as the deadline constraint increases, the better is the heuristic's performance. In Fig. 7 a third curve is reported. The second heuristic is an approximation of the infinite horizon problem: since the optimal problem cannot be solved without discretizing data, the performance of the infinite horizon problem is also obtained discretizing data, in order to make the comparison more fair. In this last case, the performance is suboptimal even in the long term.

The numerical approach, which is a brute-force solution of the dynamic programming formulation, is orders of magnitude slower than the heuristic. For this reason we limited the above example to just 3 users. However the heuristic could be easily implemented for any number of users as its complexity is linear in the number of users.

Another interesting heuristic can be inspired by the Shortest Path approach of section 3.1. The constant channel assumption is realistic for non-fading or slowly fading channels. However, this approach can be a reasonable starting point to develop an algorithm for the fading case.

Framed Shortest Path

Consider N users with a common deadline K . Each user j had an amount of data d to transmit. A reasonable heuristic may be

1. divide the time window in M frames $\frac{K}{M}$ time slot long. Then, on each frame:
2. assume the channel of each user constant and equal to the first channel quality sample
3. apply the shortest path formulation to the graph G generated according to the procedure specified in section 3.1, scaling the amount of data d of each user of an M factor.

5 Conclusions

In this paper we consider the problem of scheduling transmissions over a wireless channel subject to a deadline constraint. In principle, it may be possible to satisfy deadline constraints by simply increasing the data rate. However, higher data rates consume more energy; and typically the required energy is a convex function of the data rate. Therefore, a key objective is to minimize the total energy consumed subject to satisfying the deadline constraints. To that end, we exploit two aspects of wireless transmission. First, since energy is convex in data rate, we seek transmission schemes that use the lowest possible data rate while satisfying the deadline constraints. Second, since the wireless channel is time-varying, we exploit opportunistic scheduling by attempting to transmit at times that the channel conditions are good.

Specifically, we consider a single wireless transmitter with N users. Each user has a fixed amount of data that must be transmitted by a deadline. We consider both a linear and a strictly convex rate-power curves and obtain optimal solutions, based on dynamic programming, and tractable approximate heuristics in both cases. When the channel is time-invariant, we obtain an optimal solution based on a Shortest Path formulation. When the channel is time-varying and the rate-power curve is linear, our dynamic programming solution has a nice "threshold" form. While, for the strictly convex rate-power curve we are able to obtain a heuristic algorithm that performs well when compared to a "brute-force" implementation of the dynamic programming solution.

The problem of scheduling wireless transmissions with a deadline constraint is very important for supporting both wireless data services as well as traditional voice services; yet, understanding of this complex scheduling problem remains very limited. In this paper we address a simple version of the problem with fixed amounts of data and known deadlines. Important extensions to this work include scheduling transmissions with stochastic traffic; obtaining "optimal" solutions when the channel is time-varying and the rate-power curve is convex or including a power limitation; supporting different classes of traffic (e.g., with different deadlines or priorities); as well as assuming fading channels with correlation. Finally, this investi-

gation could be headed to distributed versions of the proposed algorithms, in order to apply these solutions to problems in which all the elements of the networks (servers and terminals) have complete mobility.

References

- [1] A.Tarello, J.Gao, E.Modiano, "Energy Efficient Transmission Scheduling over Mars Proximity Links", Proceedings of *IEEE Aerospace Conference 2006*, Big Sky, MT, USA, 2006.
- [2] M.Zafer, E.Modiano, "A Calculus Approach to Minimum Energy Transmission Policies with Quality of Service Guarantees", Proceedings of *IEEE INFOCOM 2005*, Miami, FL, USA.
- [3] A.El Gamal, E. Uysal, B.Prabhakar, "Energy-Efficient transmission over a wireless link via lazy packet scheduling", Proceedings of *IEEE INFOCOM 2001*.(2001, vol.1, pp.386-94).
- [4] A.Fu, E.Modiano, J.Tsitsiklis, "Optimal Energy Allocation for Delay-Constrained Data Transmission over a Tyme-varying Channel", Proceedings of *IEEE INFOCOM 2003*.(2003, vol.2, pp.1095-1105).
- [5] T.Rappaport, *Wireless Communication, Principles and Practice*, Prentice-Hall Inc., N.J., 1996.
- [6] D.Bertsimas, J.N.Tsitsiklis, *Introduction to Linear Optimization*, Belmont, MA, USA: Athena Scientific, 1997.
- [7] D.P.Bertsekas, *Dynamic Programming and Optimal Control*, Belmont, MA, USA: Athena Scientific, 2000.
- [8] D.P.Bertsekas, *Nonlinear Programming*, Belmont, MA, USA: Athena Scientific, 1999.
- [9] D. Zhang and K. Wasserman, "Transmission schemes for time-varying wireless channels with partial state observations", Proceedings of *IEEE INFOCOM 2002*.(2003, vol.2, pp.467-476).
- [10] B. Collins and R. Cruz, "Transmission policies for time varying channels with average delay constraints", *Proceedings, 1999 Allerton Conference on Communication, Control and Computing*, Monticello, IL, USA, 1999.

A Optimal Stopping Time

Optimal stopping problems are characterized by the availability, at each state, of a control stopping the evolution of the system. At each stage the decision maker observes the current state of the system and decides whether to continue the process, perhaps at a certain cost, or stop the process and incur another cost. The target is developing a policy to maximize the revenue, stopping the process when the *actual* cost of doing it is more convenient than continuing the process itself.

Let us consider a single user scenario of the problem discussed in this paper. The user j is waiting to transmit on a radio channel a fixed amount d of data. The power-rate curve is linear, so all the amount of data can be transmitted within a single time slot. The server has just to decide the best time slot for transmitting; the transmission will be allowed when the actual cost of transmission is less than the expected cost of delaying that transmission. We consider an horizon of K time slots. If the server has not transmitted after the $K - 1$ th time slot, it must necessarily transmit in the K th slot. The objective is to find a policy that minimizes the expected cost of transmission.

We assume that the channel samples, denoted q_0, q_1, \dots, q_{K-1} are random and independent. A Dynamic Programming approach can be developed. The control space consist of two elements, u_1 and u_2 , denoting "transmission" and "idle" respectively. Let q_k denote the channel state and d_k the amount of data left in the queue at time k . The data is either completely transmitted in a slot or the system idles in that slot. We can write the state evolution as,

$$d_{k+1} = f_k(d_k, u_k, q_k) \quad k = 0, 1, \dots, K - 1$$

defined as

$$d_{k+1} = \begin{cases} 0, & d_k = 0 \text{ or } (d_k \neq 0 \text{ and } u_k = u_1) \\ d, & \text{otherwise} \end{cases}$$

The corresponding cost function may be written as

$$E_{q_k, k=0,1,\dots,K-1} \left[g_K(d_K) + \sum_{k=0}^{K-1} g_k(d_k, u_k, q_k) \right]$$

where

$$g_K(d_K) = \begin{cases} \frac{d}{q_K}, & d_K \neq 0 \\ 0, & \text{otherwise,} \end{cases}$$

$$g_k(d_k, u_k, q_k) = \begin{cases} \frac{d}{q_k}, & d_k \neq 0 \text{ and } u_k = u_1 \\ 0, & \text{otherwise.} \end{cases}$$

are respectively the cost at the last time slot and the cost of the whole path excluding the last time slot.

The DP formulation is then

$$\overline{\mathcal{J}}_K(d_K) = \begin{cases} \frac{d}{q_K}, & d_K \neq 0 \\ 0, & d_K = 0, \end{cases} \quad (16)$$

$$\begin{aligned} & \mathcal{J}_k(d_k, q_k) = \\ & = \begin{cases} \min \left[\frac{d}{q_k}, \overline{\mathcal{J}}_{k+1}(d_{k+1}) \right], & d_k \neq 0 \\ 0, & d_k = 0. \end{cases} \end{aligned}$$

When the remaining amount of data to serve at time k is d_k , the cost resulting from transmission is $\frac{d_k}{q_k}$, while $\overline{\mathcal{J}}_{k+1}(d_{k+1})$ is the expected cost of waiting. Clearly, the optimal policy, then, is to transmit when the transmission cost is smaller than the expected future cost of delaying the transmission, i.e., if

$$\frac{d_k}{q_k} \leq \alpha_k \quad (17)$$

where

$$\alpha_k = E[\mathcal{J}_{k+1}(d_{k+1}, q_{k+1})]$$

is a threshold representing the expected future cost associated with delaying the transmission.

The policy is thus characterized by a backward and forward process. A sequence of α_k is initially pre-computed backward from the termination state. Then, during the system operation, the optimal stopping time is the one satisfying (17). The termination condition is (16) and forces the server to transmit by the last time slot in order to respect the deadline constraint. Further discussions on the Optimal Stopping Time Strategy can be found in [7].