Energy-Efficient Transmission Scheduling for Wireless Media Streaming with Strict Underflow Constraints

David Shuman and Mingyan Liu
Electrical Engineering and Computer Science Department
University of Michigan
Ann Arbor, MI 48109-2122
{dishuman,mingyan}@umich.edu

***Abstract*— We consider a single source transmitting media streams to multiple users over a shared wireless channel. The channel for each user is time-varying, and each user has a buffer to store received packets before they are decoded and played. At each time step, the source determines how much power to use for transmission to each user. The objective is for the source to allocate power in a manner that minimizes an expected cost measure, while satisfying strict buffer underflow constraints and a total power constraint in each slot. The expected cost measure is composed of costs associated with power consumption from transmission and packet holding costs. The buffer underflow constraints prevent the user buffers from emptying, so as to maintain playout quality. In the case of a single user, we show that a modified base-stock policy is optimal under the finite and infinite horizon discounted expected cost criteria. We present the sequences of critical numbers that characterize the optimal control laws in each of these two problems. We also discuss the structure of the optimal policy in the multi-user case.**

I. Introduction

Transporting multimedia over wireless networks is a promising application that has seen recent advances [1]. At the same time, a number of resource allocation issues need to be addressed in order to provide high quality and efficient media over wireless. First, streaming is in general bandwidth-demanding; therefore, simultaneously satisfying the different bandwidth needs of multiple users/mobiles sharing the same channel can be a challenging task. Second, streaming applications tend to have stringent quality of service (QoS) requirements (e.g., they can be delay and jitter intolerant). Third, it is desirable to operate the wireless system in an energy-efficient manner. This is obvious when the source of the media streaming (the sender) is a mobile. When the media comes from a base station that is not power-constrained, it is still desirable to conserve power in order to limit potential interference to other base stations and their associated mobiles.

In this paper, we focus on the problem of reducing both system-wide power consumption and playout interruptions to end users. We consider a single source transmitting audio/video sequences to multiple users over a shared wireless channel. Each user has a buffer to store received packets before they are decoded and played. The available data rate of the channel varies with time and from user to user, due to random fading. The transmitter's goal is to minimize total power consumption by exploiting the temporal and spatial variation of the channel, while preventing any user's buffer from emptying, thus reducing playout interruptions.

This problem falls into the general class of multi-user variable channel scheduling problems [2], also called opportunistic scheduling problems. The common theme amongst these problems is that the transmission scheduler has two competing interests. The first is to use power efficiently, by allocating more power for transmission to those users with the strongest channels at each time. The second is to meet the data rate or packet delay constraints of all users, maintaining a certain notion of *fairness* across the network.

The idea of increasing system throughput and reducing total power consumption through such a joint resource allocation policy is commonly referred to as *multiuser diversity* [3]. It was introduced in the context of the analogous *uplink* problem where multiple sources transmit to a single destination (e.g., the base station) [4]. Since, there has been a wide range of literature on opportunistic scheduling problems in wireless networks. Many of these studies have tried to maximize system throughput subject to some fairness and/or power constraints. For example, [5] and [6] consider *temporal fairness*; [3] and [7] consider *proportional fairness*; [8] considers *weighted max-min fairness*; [6] considers a more general *utilitarian fairness*; [9] considers a statistical delay constraint; and [10] considers individual user throughput and total power constraints simultaneously. [11] and [12] present algorithms to keep all users' queues stable while balancing throughput and delay. For a recent overview of opportunistic scheduling studies in wireless networks, see [13].

In the context of streaming media, the notion of fairness is often linked to playout quality, as measured by delay or probability of buffer underflow. Of the related work, [14] has the closest setup to our model. The main differences are that [14] features a loose constraint on underflow (i.e., it is allowed, but at a cost), as opposed to our tight constraint, and the two studies adopt different wireless channel models. In the extension [15], the receiver may slow down its playout rate (at some cost) to avoid underflow. In this setting, the authors investigate the tradeoffs between power consumption and playout quality, and examine joint power/playout rate control policies. In our model, the receiver does not have the option to adjust the playout speed. Our model also bears resemblance to [16]. The first difference here is that [16] aims to minimize transmission energy subject to a constant end-to-end delay constraint on each video frame. A second difference is that the controller in [16] must assign various source coding parameters such as quantization step size and coding mode, whereas our model assumes a fixed encoding/decoding scheme.

The remainder of this paper is organized as follows. In the next section, we describe the multi-user system model, relate it to models in inventory theory, and formulate two optimization problems (finite and infinite horizons). In Section III, we examine the case of a single receiver, and characterize the optimal scheduling policy for both problems. In Section IV, we discuss the structure of the optimal policy in the multi-user case.

II. Problem Description

In this section, we present an abstraction of the transmission scheduling problem outlined in the previous section and formulate the optimization problems. While most of the results we present are for the single user case, the formulation in this section is for the more general multi-user (multi-receiver) case. This allows us to precisely define the power constraint (in the general case), and to discuss the multi-user case in Section IV.

A. System Model and Assumptions

We consider a single source transmitting media sequences to M users/receivers over a shared wireless channel. The sender maintains a separate buffer for each receiver and is assumed to always have data to transmit. Each user has a playout/receiver buffer at the receiving end, assumed to be infinite. While in reality this cannot be the case, it is nevertheless a reasonable assumption considering the decreasing cost and size of memory, and the fact that our system model penalizes holding packets in the buffer. See Fig. 1 for a diagram of the system.

We consider time evolution in discrete steps indexed backwards by $n = N, N-1, \ldots, 1$, with n representing the number of slots remaining in the time horizon. N is the length of the time horizon, and slot n refers to the time interval $[n, n-1)$.

At the beginning of each time slot, the scheduler allocates some amount of power (possibly zero) for transmission to each user. The total power consumed in any one slot must not exceed the stationary power constraint, P. Following transmission and reception in each slot, a certain number of packets are removed/purged from the receiver buffer for playing. We assume that packets transmitted in slot n arrive in time to be used for playing in slot n. A per packet per slot holding cost of h^m is assessed on all packets remaining in user m's receiver buffer after playout consumption.

We assume that the transmitter (or scheduler) knows precisely the packet requirements of each user in each time slot.

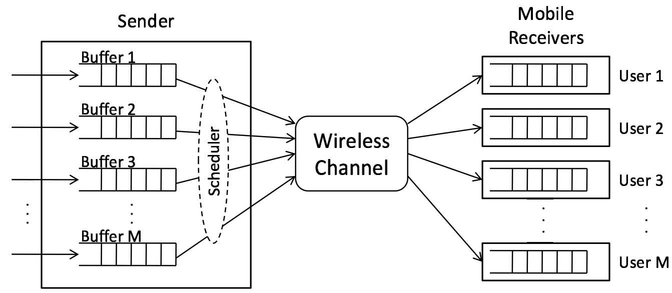

Fig. 1. System model.

This is justified by the assumption that the transmitter knows the encoding and decoding schemes used. We further assume that a user's consumption of packets in each slot is constant, denoted by d^m, $m = 1, 2, \cdots, M$. This assumption is less realistic, but may be justified if the receiving buffer is drained at a constant rate at the MAC layer, before packets are decoded by the media player at the application layer. It is also worth noting that the same techniques we use in this paper to analyze the stationary demand case can be used to examine the non-stationary demand case.

In general, wireless channel conditions are time-varying, and differ from user to user. Adopting a block fading model, we assume that the slot duration is within the channel coherence time such that the channel conditions within a single slot are constant. A user m's channel condition in slot n is modelled as a random variable, S_n^m, that can be in a finite number of states, indexed by the set $\mathcal{S}^m$. We assume that a given channel's state is independent and identically distributed from slot to slot, and is also independent of all other channels and scheduling decisions. If user m's channel condition is in state s, then the transmission of r units of data incurs a power consumption of $c_s^m(r)$. This power-rate function $c_s^m(\cdot)$ is commonly assumed to be convex (in the high SNR regime) or linear (in the low SNR regime). In this paper, we only consider the linear case; the more general convex case will be dealt with in a future paper. We will subsequently simplify our notation and use c_s^m or $c^m(S_n^m)$ to denote the power consumption per unit of data transmitted when user m's channel is in state s or given by random variable S_n^m at time n, respectively. Similarly, we denote $\left(c^1(s^1), c^2(s^2), \ldots, c^M(s^M)\right)^{\mathrm{T}}$ by $\mathbf{c_s}$. The transmitter is assumed to learn each channel's state through a feedback channel at the beginning of each time slot, prior to making the scheduling decision.

We are primarily concerned with two objectives in deriving a good transmission policy. One is to avoid underflow at the receiver buffer, thus avoiding disruption to the user playout. The other is to minimize system-wide power consumption. The goal of this study is to characterize the control laws that minimize the transmission power costs over a finite or infinite time horizon, subject to tight underflow constraints and a maximum power constraint in each time slot.

The model outlined above corresponds closely to models used in inventory theory. Borrowing that field's terminology, our abstraction is a multi-period, multi-item, discrete time inventory model with random ordering prices, a budget constraint, and deterministic demands. The items correspond to the different users receiving data, the random ordering prices to the random channel conditions, the budget constraint to the power available in each time slot, and the deterministic demands to each user's packet requirements for playout.

To the best of our knowledge, this particular problem has not been studied in the context of inventory theory, but similar problems have been examined. [17] - [23] all consider single item inventory models with random ordering prices. The key result for the case of deterministic demand of a single item with no resource constraint is that the optimal policy is a modified base-stock policy. Specifically, for each possible ordering price (translates into channel condition in our context), there exists a critical number such that the optimal policy is to fill the inventory (buffer) up to that critical number if the the current level is lower than the critical number, and not to order (transmit) anything if the current level is above the critical number. Of the prior work, Kingsman [18], [19] is the only author to consider a resource constraint, and he imposes a maximum on the number of items that may be ordered in each slot. The resource constraint we consider is of a different nature in that we limit the total amount of power available in each slot. This is equivalent to a limit on the total per slot budget (regardless of the stochastic price realizations), rather than a limit on the number of items that can be ordered. In Section III, we discuss the above single item results further, and leverage some of their techniques and solutions in deriving our results. We are not aware of any prior work regarding multiple item inventory models with stochastic ordering prices.

In addition to the assumptions described above, we also make the following assumptions for technical reasons:

1) We assume all user buffers have an initial queue size of 0.
2) We assume that even when all users are experiencing their worst possible channel conditions, the maximum power constraint P is sufficient to transmit enough packets to satisfy one time slot's playout demand for each user, i.e., $\sum_{m=1}^{M}\{d^m \cdot c^m(S^m)\} \leq P$ for every possible realization of the random vector $(S^1, S^2, \cdots, S^M)^{\mathrm{T}}$.
3) For the special case of a single user with constant demand d that is considered in Section III, we assume that for every possible channel condition s, $P/c_s = l_s \cdot d$ for some $l_s \in \mathbb{N}$; i.e., the maximum number of packets that can be transmitted in any slot covers exactly the playout requirements of some integer number of slots.

In general, vectors are in boldface and are taken as column vectors unless otherwise noted.

B. Problem Formulation

We consider two problems. The first, Problem (**P1**), is the finite horizon discounted expected cost problem. The second, Problem (**P2**), is the infinite horizon discounted expected cost problem. The two problems feature the same information state, action space, system dynamics, and cost structure, but different optimization criteria.

The information state at slot n is the pair $(\mathbf{X}_n, \mathbf{S}_n)$, where the random vector $\mathbf{X}_n = (X_n^1, X_n^2, \cdots, X_n^M)^{\mathrm{T}}$ denotes the current receiver buffer queue length of each user and $\mathbf{S}_n = (S_n^1, S_n^2, \cdots, S_n^M)^{\mathrm{T}}$ denotes the channel conditions in slot n (recall that n is the number of steps remaining until the end of the horizon). The queue dynamics for user m are governed by the simple equation $X_{n-1}^m = Y_n^m - d^m$ at all times $n = N, N-1, \ldots, 1$. Y_n^m is the queue length *after* transmission in the n^{th} slot takes place, but *before* playout in the n^{th} slot has occurred. $\mathbf{Y}_n$ is a controlled random vector chosen by the scheduler at each time n. It must satisfy the power constraint: $\sum_{m=1}^{M} c^m(S_n^m) \cdot (Y_n^m - X_n^m) \leq P$, and the underflow constraints: $Y_n^m \geq d^m, \ \forall m$. Clearly, the scheduler cannot transmit a negative number of packets, so it must also be true that $Y_n^m \geq X_n^m, \ \forall m$.

We now present the optimization criterion for each problem. In addition to the cost associated with power consumption from transmission, we introduce for technical purposes a holding cost on each packet stored in a user's playout buffer at the end of a time slot. In the finite horizon case, we can exclude this cost from the optimization (set $h^m = 0, \ \forall m$), but we use the fact that all holding cost rates are strictly positive (albeit arbitrarily small) in our infinite horizon solution. In Problem (**P1**), we wish to find a transmission policy $\boldsymbol{\pi}$ that minimizes J_N^{π}, the expected cost under policy $\boldsymbol{\pi}$, defined as:

$$J_N^{\pi} := \mathbb{E}^{\pi}\left\{\sum_{m=1}^{M}\sum_{t=1}^{N}\alpha^{N-t}\left[\begin{array}{c} c^m\big(S_t^m\big)\cdot\big(Y_t^m - X_t^m\big) \\ +h^m\cdot\big(Y_t^m - d^m\big)\end{array}\right] \mid \mathcal{F}\right\},$$

where $0 \leq \alpha < 1$ is the discount factor and $\mathcal{F}$ denotes all information available at the beginning of the time horizon. For Problem (**P2**), the cost function for minimization is defined as $J_{\infty}^{\pi} := \lim_{N\to\infty} J_N^{\pi}$. In both cases, we allow the transmission policy $\boldsymbol{\pi}$ to be chosen from the set of all randomized and deterministic control laws, $\boldsymbol{\Pi}$.

Combining the constraints and criteria, we present the optimization formulations for Problem (**P1**) (or (**P2**)):

$$\begin{aligned} & \min_{\pi\in\Pi} J_N^{\pi} \quad \Big(\text{or} \quad \min_{\pi\in\Pi} J_{\infty}^{\pi}\Big) \\ \text{s.t.} \quad & \sum_{m=1}^{M} c^m(S_n^m)\cdot(Y_n^m - X_n^m) \leq P, \quad \forall n \quad \text{and} \\ & Y_n^m \geq \max(X_n^m, d^m), \quad \forall m, \forall n. \end{aligned}$$

The above problem may be solved using standard dynamic programming (see, e.g., [24]). The recursive DP equations in the finite horizon case are given by:

$$\begin{aligned} V_n(\mathbf{x},\mathbf{s}) &= \min_{\mathbf{y}\in\mathcal{A}^{\mathbf{d}}(\mathbf{x},\mathbf{s})}\left\{\begin{array}{l} \mathbf{c}_{\mathbf{s}}^{\mathrm{T}}(\mathbf{y}-\mathbf{x}) + \mathbf{h}^{\mathrm{T}}(\mathbf{y}-\mathbf{d}) \\ +\alpha\cdot\mathbb{E}\big[V_{n-1}(\mathbf{y}-\mathbf{d},\mathbf{S})\big]\end{array}\right\} \\ V_0(\mathbf{x},\mathbf{s}) &= 0, \quad \forall \mathbf{x}, \forall \mathbf{s} \end{aligned} \qquad (1)$$

where $V(\cdot,\cdot)$ is the value function or expected cost-to-go. The relevant infinite horizon functional equation is:

$$V_\infty(\mathbf{x},\mathbf{s}) = \min_{\mathbf{y}\in\mathcal{A}^{\mathbf{d}}(\mathbf{x},\mathbf{s})} \left\{ \begin{array}{l} \mathbf{c}_\mathbf{s}^{\mathrm{T}}(\mathbf{y}-\mathbf{x}) + \mathbf{h}^{\mathrm{T}}(\mathbf{y}-\mathbf{d}) \\ +\alpha\cdot I\!\!E\big[V_\infty(\mathbf{y}-\mathbf{d},\mathbf{S})\big] \end{array} \right\}. \quad (2)$$

In both functional equations (1) and (2), the action space is defined as:

$$\mathcal{A}^{\mathbf{d}}(\mathbf{x},\mathbf{s}) := \left\{ \mathbf{y}\in I\!\!R_+^M : \begin{array}{l} \max(\mathbf{d},\mathbf{x}) \preceq \mathbf{y} \text{ and} \\ \mathbf{c}_\mathbf{s}^{\mathrm{T}}(\mathbf{y}-\mathbf{x}) \leq P \end{array} \right\}, \quad (3)$$

where the maximum in (3) is taken element-by-element $\big(\text{i.e.}, \max(d^m,x^m) \leq y^m \quad \forall m\big)$. In the single user case, this action space reduces to:

$$\mathcal{A}^d(x,s) := \left\{ y\in I\!\!R_+ : \max(d,x) \leq y \leq \frac{P}{c_s} + x \right\}.$$

III. Analysis of the Single User Case

In this section, we analyze the finite and infinite horizon discounted expected problems when there is only a single user ($M = 1$). Kingsman [18], [19] considers a similar problem with a restriction on the maximum quantity that can be ordered (transmitted) at the offered price (rate) in any time slot. In his problem, however, the restriction is the same in every slot, regardless of the realization of the stochastic ordering price (channel condition). Using a similar approach, we analyze the model described in Section II-A with a fixed power constraint rather than a fixed limit on the maximum number of packets that can be transmitted in a slot. Throughout this section, we drop the superscript referring to the item index, as there is only one item under consideration.

A. Finite Horizon Discounted Expected Cost Problem

We now present the optimal transmission policy by first defining a set of thresholds recursively, and then using them to determine the optimal transmission level in each state. The threshold $\gamma_{n,j}$ may be interpreted as the per packet power cost at which, with n slots remaining in the horizon, the expected cost-to-go of transmitting packets to cover the user's playout requirements for the next $j-1$ slots is the same as the expected cost-to-go of transmitting packets to cover the user's requirements for the next j slots.

Theorem 1: Define the thresholds $\gamma_{n,j}$ for $n \in \{1,2,\ldots,N\}$ and $j \in I\!\!N$ recursively, as follows:

(i) If $j = 1$, $\gamma_{n,j} = \infty$;
(ii) If $j > n$, $\gamma_{n,j} = 0$;
(iii) If $2 \leq j \leq n$,

$$\gamma_{n,j} = -h+\alpha\cdot \left[\begin{array}{c} \displaystyle\sum_{\{s:\ c_s>\gamma_{n-1,j-1}\}} p_s\cdot\gamma_{n-1,j-1} \ + \\ \displaystyle\sum_{\{s:\ \gamma_{n-1,j+L(s)-1}\leq c_s\leq\gamma_{n-1,j-1}\}} p_s\cdot c_s \ + \\ \displaystyle\sum_{\{s:\ c_s<\gamma_{n-1,j+L(s)-1}\}} p_s\cdot\gamma_{n-1,j+L(s)-1} \end{array} \right],$$

where p_s is the probability of the channel being in state s in a time slot, and $L(s) := \frac{P}{d\cdot c_s}$. For each $n \in \{1,2,\ldots,N\}$ and $s \in \mathcal{S}$, if $\gamma_{n,j+1} < c_s \leq \gamma_{n,j}$, define $b_n(s) := j\cdot d$. We refer to $b_n(s)$ as a *critical number*. The optimal control strategy is then given by $\pi^* = \{y_N^*, y_{N-1}^*, \ldots, y_1^*\}$, where

$$y_n^*(x,s) := \begin{cases} x, & \text{if } x \geq b_n(s) \\ b_n(s), & \text{if } b_n(s)-\frac{P}{c_s} \leq x < b_n(s) \\ x+\frac{P}{c_s}, & \text{if } x < b_n(s)-\frac{P}{c_s} \end{cases}.$$

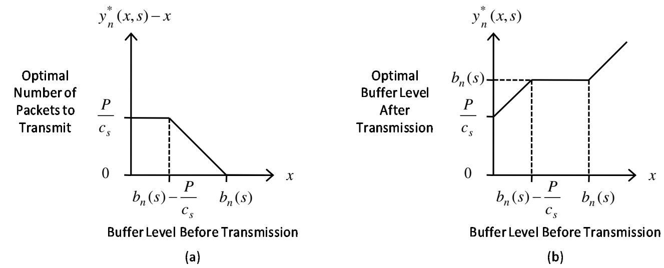

Fig. 2. Optimal policy in slot n when the state is (x,s). (a) depicts the optimal transmission quantity, and (b) depicts the resulting number of packets available for playout in slot n.

Note that with n slots remaining, $0 \leq \gamma_{n,n+1} \leq \gamma_{n,n} \leq \gamma_{n,n-1} \leq \ \ldots \ \leq \gamma_{n,2} \leq \gamma_{n,1} = \infty$, so $b_n(s)$ is well-defined. The proof of Theorem 1 follows a similar technique to the proof of Golabi's Theorem 1 [23]. Specifically, we show by backwards induction that it is worse to transmit either fewer or more packets than the number suggested by the policy π^*. The detailed proof is included in a forthcoming paper.

The optimal transmission policy in Theorem 1 is a modified base-stock policy. At time n, for each possible channel condition realization s, the critical number $b_n(s)$ describes the ideal number of packets to have in the user's buffer after transmission in the n^{th} slot. If that number of packets is already in the buffer, then it is optimal to not transmit any packets; if there are fewer than ideal and the available power is enough to transmit the difference, then it is optimal to do so; and if there are fewer than ideal and the available power is not enough to transmit the difference, then the sender should use the maximum power to transmit. See Fig. 2 for diagrams of the optimal policy.

Compared to using the dynamic program to compute the optimal policy, the above result not only sheds more insight on the structural properties of the problem and its optimal solution, but also offers a computationally simpler method. In particular, the optimal policy is completely characterized by the thresholds $\gamma_{n,j}$. Calculating these thresholds recursively, as described in Theorem 1, is considerably simpler from a computational standpoint than solving the full dynamic program.

Comparing (iii) to the corresponding thresholds in the unrestricted (no power constraint) single user problem [18], [23], we see that the difference is:

$$\sum_{\{s:\ c_s<\gamma_{n-1,j+L(s)-1}\}} p_s\cdot\left[\gamma_{n-1,j+L(s)-1}-c_s\right].$$

For all $n \in \{1,2,\ldots,N\}$ and $j \in I\!\!N$, this term is non-negative. Thus, for a fixed n and j, the threshold in the restricted case is at least as high as the corresponding threshold in the unrestricted case. It follows that the optimal stock-up level $b_n(s)$ is also at least as high in the restricted case for all $n \in \{1, 2, \ldots, N\}$ and $s \in \mathcal{S}$. The intuition behind this difference is that the sender should transmit more packets under the same (medium) conditions, because it is not able to take advantage of the best channel conditions to the same extent due to the power constraint.

B. Infinite Horizon Discounted Expected Cost Problem

In this section, we show that the infinite horizon optimal policy is the natural extension of the finite horizon optimal policy; namely, it is a modified base-stock policy of the same form, characterized by *stationary* target buffer levels $\{b_\infty(s)\}_{s\in\mathcal{S}}$.

Theorem 2:

(a) $\lim_{n\to\infty} V_n(x,s)$ exists and is finite, $\forall x \in \mathbb{R}_+, \forall s \in \mathcal{S}$.

(b) Define $f(x,s) := \lim_{n\to\infty} V_n(x,s)$. Then $\lim_{x\to\infty} f(x,s) = \infty, \forall s \in \mathcal{S}$, and $f(x,s)$ is convex in x for any fixed $s \in \mathcal{S}$.

(c) $f(x,s)$ satisfies the infinite horizon functional equation (2), and the minimum is achieved by:

$$y^*_\infty(x,s) := \begin{cases} x, & \text{if } x \geq b_\infty(s) \\ b_\infty(s), & \text{if } b_\infty(s) - \frac{P}{c_s} \leq x < b_\infty(s) \\ x + \frac{P}{c_s}, & \text{if } x < b_\infty(s) - \frac{P}{c_s} \end{cases},$$

where $b_\infty(s) = \lim_{n\to\infty} b_n(s)$.

(d) The optimal stationary policy is given by $\pi^*_\infty = y^*_\infty$.

The proof of Theorem 2 follows the logic conveyed in the statement of the theorem, and leverages techniques similar to [25] and [26]. Details are included in a forthcoming paper.

IV. Analysis of the Multiple Users Case

Based on numerical experiments, it appears that the optimal policies discussed in Section III extend to the case of multiple receivers in the following sense:

1) For each vector of channel conditions, $\mathbf{s}$, at time n, there exists a vector of critical numbers (corresponding to $b_n(s)$ in the single user case), with one critical number for each user. These critical numbers satisfy:

$$G_n(\mathbf{b}_n(\mathbf{s}), \mathbf{s}) = \min_{\mathbf{y}\in\mathbb{R}^n_+} \{G_n(\mathbf{y}, \mathbf{s})\}, \text{ where}$$
$$G_n(\mathbf{y}, \mathbf{s}) := \mathbf{c}^{\mathrm{T}}_{\mathbf{s}} \cdot \mathbf{y} + \mathbf{h}^{\mathrm{T}} \cdot \mathbf{y} + \alpha \cdot \mathbb{E}\big[V_{n-1}(\mathbf{y} - \mathbf{d}, \mathbf{S})\big]$$

2) Each user's critical number b^i_n depends only on its current channel condition s^i, and is independent of its own current buffer level, other users' current buffer levels, and other users' current channel conditions.
3) It is optimal for the transmitter to not transmit any packets to any user whose current buffer level is greater than or equal to its critical number.
4) If it is possible for the transmitter to schedule packet transmissions to bring all other users' buffer levels up to their respective critical numbers without exceeding the power constraint, then it is optimal for it to do so.
5) If the power constraint prevents the transmitter from doing so, then it should allocate the full power P for transmission to different receivers in some manner yet to be determined.

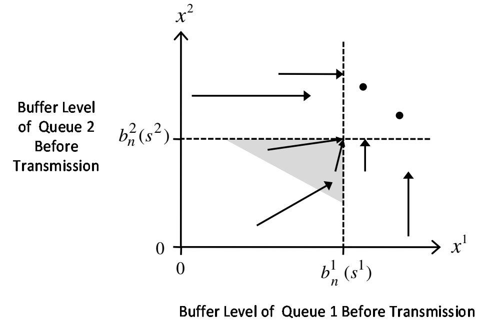

Fig. 3. Structure of the optimal policy for the two-user case in slot n, with a fixed vector of channel conditions, $\boldsymbol{s}$. The tails of the arrows represent the buffer levels before transmission, and the heads represent buffer levels after transmission.

The above structure is analogous to the finite and infinite horizon optimal policies in the case of multiple users with stochastic playout requirements (demands) and deterministic channel conditions (ordering prices) [27], [25].

Figure 3 illustrates an optimal policy of this structure in a two-user problem, for a fixed vector of channel conditions. In the upper right region, both users' buffer levels are above their critical numbers, so it is optimal to not transmit any packets. In the lower right and upper left regions, it is optimal to only transmit to user 2 or user 1, respectively. If the power budget is sufficient to bring that user up to the critical number, then it is optimal to do so; otherwise, the full power is allocated to the user whose buffer level is below its critical number. In the shaded region, both users start below their critical numbers, but there is sufficient power to transmit enough packets to make up the difference for both. In the unshaded portion of the lower left region, there is not sufficient power to bring both users to their critical numbers, so the sender should transmit at full power. In this case, the distribution of power between the two users is not specified.

V. Conclusion

In this paper, we considered the problem of transmitting media streams over a wireless channel in a manner that prevents the receiver's buffer from emptying. In the case of a single receiver, we showed that under both the finite and infinite horizon discounted expected cost criteria, the optimal transmission schedule is a modified base-stock policy. We also presented a series of thresholds that completely characterize this optimal policy. In the case of multiple receivers using a shared channel, we discussed the structure of the optimal policy.

In addition to proving the structure of the optimal policy, future work on the multi-user case includes identifying low complexity sub-optimal policies that perform well. We also plan to explore how the single user results presented in this paper may be leveraged in the special multi-user case of M identical users.

REFERENCES

[1] B. Girod, M. Kalman, Y. J. Liang, and R. Zhang, "Advances in channel-adaptive video streaming," in *Proceedings of the International Conference on Image Processing*, Rochester, NY, September 2002, vol. 1, pp. 9–12.

[2] M. Andrews, K. Kumaran, K. Ramanan, A. Stolyar, P. Whiting, and R. Vijayakumar, "Providing quality of service over a shared wireless link," *IEEE Communications Magazine*, vol. 39, no. 2, pp. 150–154, February 2001.

[3] P. Viswanath, D. N. C. Tse, and R. Laroia, "Opportunistic beamforming using dumb antennas," *IEEE Transactions on Information Theory*, vol. 48, no. 6, pp. 1277–1294, June 2002.

[4] R. Knopp and P. A. Humblet, "Information capacity and power control in single-cell multiuser communications," in *Proceedings of the International Conference on Communications*, Seattle, WA, June 1995, vol. 1, pp. 331–335.

[5] F. Berggren and R. Jäntti, "Asymptotically fair transmission scheduling over fading channels," *IEEE Transactions on Wireless Communications*, vol. 3, no. 1, pp. 326–336, January 2004.

[6] X. Liu, E. K. P. Chong, and N. B. Shroff, "A framework for opportunistic scheduling in wireless networks," *Computer Networks*, vol. 41, no. 4, pp. 451–474, March 2003.

[7] J. M. Holtzman, "CDMA forward link waterfilling power control," in *Proceedings of the IEEE Vehicular Technology Conference*, Tokyo, Japan, May 2000, vol. 3, pp. 1663–1667.

[8] S. Borst and P. Whiting, "Dynamic channel-sensitive scheduling algorithms for wireless data throughput optimization," *IEEE Transactions on Vehicular Technology*, vol. 52, no. 3, pp. 569–586, May 2003.

[9] D. Wu and R. Negi, "Utilizing multiuser diversity for efficient support of quality of service over a fading channel," in *Proceedings of the IEEE International Conference on Communications*, Anchorage, AK, May 2003, vol. 3, pp. 2202–2207.

[10] Y. Liu and E. Knightly, "Opportunistic fair scheduling over multiple wireless channels," in *Proceedings of IEEE INFOCOM*, San Francisco, CA, April 2003, vol. 2, pp. 1106–1115.

[11] M. Andrews, K. Kumaran, K. Ramanan, A. Stolyar, R. Vijayakumar, and P. Whiting, "Scheduling in a queueing system with asynchronously varying service rates," *Probability in the Engineering and Informational Sciences*, vol. 18, pp. 191–217, 2004.

[12] S. Shakkottai, R. Srikant, and A. Stolyar, "Pathwise optimality of the exponential scheduling rule for wireless channels," *Advances in Applied Probability*, vol. 36, no. 4, pp. 1021–1045, 2004.

[13] X. Liu, E. K. P. Chong, and N. B. Shroff, "Optimal opportunistic scheduling in wireless networks," in *Proceedings of the Vehicular Technology Conference*, Orlando, FL, October 2003, vol. 3, pp. 1417–1421.

[14] Y. Li and N. Bambos, "Power-controlled wireless links for media streaming applications," in *Proceedings of the Wireless Telecommunications Symposium*, Pomona, CA, May 2004, pp. 102–111.

[15] Y. Li, A. Markopoulou, N. Bambos, and J. Apostolopoulos, "Joint power-playout control for media streaming over wireless links," *IEEE Transactions on Multimedia*, vol. 8, no. 4, pp. 830–843, August 2006.

[16] C. E. Luna, Y. Eisenberg, R. Berry, T. N. Pappas, and A. K. Katsaggelos, "Joint source coding and data rate adaptation for energy efficient wireless video streaming," *IEEE Journal on Selected Areas in Communications*, vol. 21, no. 10, pp. 1710–1720, December 2003.

[17] T. Fabian, J. L. Fisher, M. W. Sasieni, and A. Yardeni, "Purchasing raw material on a fluctuating market," *Operations Research*, vol. 7, no. 1, pp. 107–122, January-February 1959.

[18] B. G. Kingsman, "Commodity purchasing," *Operational Research Quarterly*, vol. 20, pp. 59–80, 1969.

[19] B. G. Kingsman, *Commodity Purchasing in Uncertain Fluctuating Price Markets*, Ph.D. thesis, University of Lancaster, 1969.

[20] B. Kalymon, "Stochastic prices in a single-item inventory purchasing model," *Operations Research*, vol. 19, no. 6, pp. 1434–1458, October 1971.

[21] V. F. Magirou, "Stockpiling under price uncertainty and storage capacity constraints," *European Journal of Operational Research*, vol. 11, pp. 233–246, 1982.

[22] K. Golabi, "A single-item inventory model with stochastic prices," in *Proceedings of the Second International Symposium on Inventories*, Budapest, Hungary, 1982, pp. 687–697.

[23] K. Golabi, "Optimal inventory policies when ordering prices are random," *Operations Research*, vol. 33, no. 3, pp. 575–588, May-June 1985.

[24] D. Bertsekas and S. E. Shreve, *Stochastic Optimal Control: The Discrete-Time Case*, Athena Scientific, 1996.

[25] G. A. DeCroix and A. Arreola-Risa, "Optimal production and inventory policy for multiple products under resource constraints," *Management Science*, vol. 44, no. 7, pp. 950–961, July 1998.

[26] A. Federgruen and P. Zipkin, "An inventory model with limited production capacity and uncertain demands II. The discounted-cost criterion," *Mathematics of Operations Research*, vol. 11, no. 2, pp. 208–215, May 1986.

[27] R. Evans, "Inventory control of a multiproduct system with a limited production resource," *Naval Research Logistics Quarterly*, vol. 14, pp. 173–184, 1967.