Joint Power/Playout Control Schemes for Media Streaming over Wireless Links

†Yan Li, †Athina Markopoulou, †Nicholas Bambos, ‡John Apostolopoulos
†Stanford University, ‡ HP Labs
{liyan,amarko,bambos}@stanford.edu, japos@hplabs.hp.com

***Abstract*— We investigate transmission and playout policies for streaming media over a wireless link. In particular, we choose both the power at the transmitter and the playout rate at the receiver, in order to minimize the power consumption and maximize the media quality. We formulate the problem using a dynamic programming approach, study the structural properties of the optimal solution, develop justified heuristics, and demonstrate significant performance gain over benchmark systems. In particular, we develop a low-complexity, practical joint power-playout heuristic that outperforms (1) the optimal power control policy in the regime where power is most important, and (2) the optimal playout control policy in the regime where media quality (playout) is most important; furthermore, this heuristic has only a slight performance loss as compared to the optimal joint power-playout control policy over the entire range of investigation.**

I. Introduction

Recent advances in video compression and streaming, and also in wireless networking technologies (next-generation cellular networks and more importantly high-throughput wireless LANs), are rapidly opening up opportunities for media streaming over interference-limited and erratic (due to mobility, channel variations, etc.) wireless channels. Of particular interest (due to emerging technical feasibility and potential commercial viability) are wireless digital home-entertainment video systems. The packetized video content is streamed to the home gateway (via satellite, cable, or Internet access) and forwarded to client video devices (TV screens, tablet PCs, etc.) via wireless links in the home LAN network.

Supporting high-quality media applications over wireless channels introduces technical faces at multiple levels, including the important problems of (1) transmitter power (and/or rate) control and (2) receiver playout rate control. In this paper, we formulate and investigate the joint power/playout control problem for media streaming over wireless, coupling the receiver and transmitter actions to achieve substantial performance gains.

We briefly note the following intuitive points regarding the core tradeoff of *transmitter power vs. receiver playout quality* for media streaming across a wireless channel of fluctuating quality. Suppose the video content should normally be played out at R packets (e.g. say frames) per time slot at the receiver. If the playout buffer runs empty, the video freezes and the user gets annoyed significantly. This may happen during periods when the channel has low-quality (high interference and/or other impairments) and packets may have to be transmitted several times before they are successfully received. To mitigate the risk of playout buffer underflow, the system has a few options, as follows:

1) The transmitter may increase its power to overcome the channel interference and/or other impairments and successfully push packets to the receiver buffer, which is otherwise at risk of underflow. In conjunction with power, the transmitter rate may also be increased to amplify the effect. Power (and/or rate) increases, however, stress the channel further and should be very carefully exercised. Moreover, for mobile transmitter nodes, increasing the power spends precious battery energy.
2) The receiver buffer may slow down its playout rate in order to extend the time until underflow if fresh packets do not arrive from the transmitter. In the extreme, it may even preemptively freeze playout for some time to accumulate a sufficient number of packets to increase the probability that it can provide a smoother playout later on. The penalty is that playout rate slow-down (and especially freezing) is noticed by the user as a media quality degradation and should be avoided if possible and exercised judiciously if not. Moreover, variations (jitter) of the playout rate also degrade the user-perceived video quality and should be suppressed.

On the other hand, during high-quality periods of the channel (interference and other impairments are low) the transmitter has an opportunity to push many packets to the receiver (depending on the space in the playout buffer) and replenish the playout buffer at low power.

In a nutshell, transmitter power increases may overcome channel interference/impairments and feed the playout buffer to prevent it from underflowing, but at the cost of stressing the channel more (and other users sharing it) and depleting battery energy (on mobile transmitters). In turn, the playout buffer 'cushions' the channel's quality lows and, hence, lowers the average transmission power and extends the battery life of mobile transmitters. The combination of power and playout control improves the key power-quality tradeoff.

In the past, the problems of power control at the transmitter and playout control at the receiver (mainly for wireline networks) have been studied separately. Power control can increase network capacity (via interference mitigation) and maintain link quality (via adaptation to variations), while conserving energy on mobile terminals.

It manages the key trade-off between delay experienced by individual packets and the power (energy) spent in their transmission [1], [2], [3], [4]. Playout control at the receiver can mitigate packet delay variation and provide better jitter suppression and smoother playout. Adaptive playout has been used in the past for media streaming over the Internet for both audio [6], [7] and video [8], [9], [10]. Wireless video is, of course, a large problem space with many different aspects, including rate control and selection of the appropriate error resilience mechanism. Various mechanisms can be used at the network and/or at the application layer or by combining the capabilities of the two. A nice discussion can be found in [5].

In this paper, we address the important question of how to jointly leverage both power and playout control dimensions for high-performance wireless video. The goal is to design a joint power-playout control technique which (1) supports a desired video playout quality at the (2) minimum stress to the wireless channel (interference) and the transmitter battery. These are competing goals defining the key power/media-quality tradeoff. In this paper, we primarily focus on streaming pre-stored media content over a wireless link, e.g. movies to wireless TVs or video/music-clips to portable computers or PDAs. These applications can typically tolerate some latency, as long as the entire content is delivered correctly and played out smoothly. This is an important example scenario where the benefits of joint power/playout control can be realized. An important attribute of this scenario is that the optimal solution requires slowing down the playout below the nominal speed, however it does not require speeding up beyond the nominal speed. A later extension of our approach will cover more delay-sensitive applications, such as live-streaming and interactive voice/video communications, where it may be necessary to both slow down and speed up the playout in order to keep the average delay at the desired level.

This paper continues in section II by developing a parsimonious formulation/model (the simplest possible, but not simplistic) that captures the key performance tradeoff(s) and provides a framework for computing and evaluating efficient dynamic power/playout control algorithms. This formulation is very general and flexible in terms of the performance costs it can incorporate. We leverage the dynamic programming methodology to obtain optimal controls and we characterize their structural properties. In section III, we study interesting and practical special cases, when a subset of the control variables are used, e.g. power-only control or playout-only control. In section IV, we leverage the structural properties of the optimal policies to design justified (based on the analysis of previous sections), thus near-optimal, yet practical (low-complexity) heuristics. We evaluate our heuristics in a Markovian interference environment and we show that they achieve substantial performance gains over conventional benchmarks.

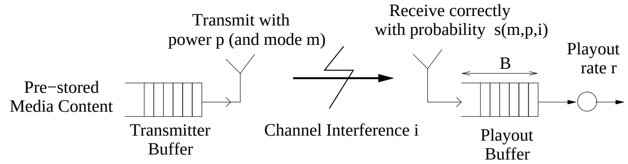

Fig. 1. Power-playout controlled streaming of pre-stored media over a wireless link: the system has control of the transmitter's transmit power p (and mode m) and the receiver's playout rate r.

In section V, we also briefly evaluate our schemes in a responsive interference environment, where the increase in power by one user may cause an increase in transmission power by other users in the environment leading to more interference. Finally, in section VI we conclude the paper and we discuss various important extensions and future research.

II. Basic Model and Problem Formulation

In this section we introduce the basic model, reflecting our problem formulation and capturing the performance tradeoffs and control issues. We embed the problem within a Markov decision process framework and use dynamic programming to compute the optimal control.

We consider a system shown in Fig. 1, which is comprised of a transmitter (Tx) and a receiver (Rx) communicating over a wireless communication link. Time is slotted and indexed by $t = 0, 1, 2, 3, ...$ The transmitter is equipped with a buffer where content comprised of N media packets is initially stored (at time 0). The receiver is equipped with a buffer of size B, where received packets are queued up while waiting to be played out.

The interference in the channel fluctuates according to a time-homogeneous Markov chain, taking values in the finite set $\mathcal{I}$ of all attainable interference states. It switches with probability q_{ij} from each state $i \in \mathcal{I}$ in a time slot to state $j \in \mathcal{I}$ in the next time slot. It is assumed for simplicity that the channel interference remains invariant within each time slot. Furthermore, it is assumed that the interference is not responsive to transmitter actions - like power variations - but is driven by an agent extraneous to the system. The case of responsive interference is explicitly investigated in section V by extending the analysis in this section and the next. It is also implicitly incorporated into the current model via the channel stress cost $\Phi(p, i)$ discussed below.

In this paper we focus on pre-encoded and stored content, where a small playout buffer is acceptable at the receiver. We focus on the case of reliable and timely delivery of all packets to the receiver. This is the requirement today for delivery of MPEG-2 coded content from a DVD over a wireless link to a wireless TV display (though it is not a requirement for streaming error-resilient MPEG-4 or H.264/MPEG-4 Advanced Video Coding (AVC) coded video over a wireless link). Because of the reliable delivery assumption the dependencies among packets, and error propagation that can result from losses, can be ignored. However the reliable delivery requirement means that if one falls behind in delivery of the media, one does not have the option of discarding packets to catch up. The viewer wants to view the entire media content played out as smoothly as possible with minimal freezes and minimal playout speed variation. High viewing quality is achieved by keeping underflow and (re)buffering delay (freezes) and jitter as low as possible. And of course this must be achieved while minimizing the power requirements at the transmitter. The formulation provides a method to specify and tradeoff the relative importance of transmitter power versus playout quality.

We now describe the general formulation, which is followed by specific incarnations for special cases.

A. Transmission Rate Control and Costs

The transmitter can transmit up to M packets in a single time slot, that is, it can change its transmission rate per slot. This is achieved by changing modulation and/or coding scheme, or more generally implementing some transmission mode $m \in \{0, 1, 2, 3, ..., M\}$ which allows for transmitting m packets in a single slot (say, by increasing the bit rate in it). This comes at an overhead cost $\Psi(m)$ to implement the mode, which reflects various stress factors, for example, the power drain and computation bandwidth at the transmitter needed for signal processing associated with mode m, etc. The cost $\Psi(m)$ is an increasing function of the number of packets m transmitted in the time slot. To cover the case where the transmitter can only transmit certain packet combinations (say, 0,2,4,16, etc. packets) in a time slot we introduce a mode set $\mathcal{M} = \{0, m_1, m_2, ..., m_k, ..., m_K\}$ with $0 < m_1 < m_2 < ... < m_k < ... < m_K < M$. The available modes are typically system-specific.

B. Power Control and Costs

When m packets at the head of the transmitter buffer are concurrently transmitted in a time slot, and the transmitter power used is p, while the channel interference is i (both assumed constant throughout the slot), then

$$s(m, p, i) \tag{1}$$

is the probability that all m packets will be successfully recovered at the receiver and, hence, removed from the transmitter buffer. For simplicity of this basic model, we assume that there is no partial recovery: either all or none of the m packets of a transmitted group are received correctly. We also assume that there is a fast and reliable feedback channel (perhpas a separate control channel) over which the receiver ACKs/NACKs the received packets at the end of each time slot. This may be of very low bandwidth compared to the forward channel.

Note that the function $s(m, p, i)$ should be decreasing in m, increasing in p and decreasing in i. It should be decreasing in m because packing more packets in a slot makes them more sensitive to interference and more difficult to recover at the transmitter. The functional dependence on p and i is obvious. Besides these general properties, we do not assume any specific formula for the success probability, which is ultimately system-specific.

Transmitting power p in a time slot when the interference is i, introduces a cost $\Phi(p, i)$ paid in that slot. This cost may reflect the interference stress that the transmission under consideration induces on the channel, e.g., interfering with 'background' transmissions sharing it. The latter may in turn stress the original 'foreground' transmitter in response to its power increases, by adjusting their powers accordingly and generating more interference on it. This entanglement effect is implicitly captured in the cost $\Phi(p, i)$. The cost $\Phi(p, i)$ should be increasing in both p and i. This is consistent with the intuition that the more congested the channel is (the less the available bandwidth resource), the more power (effort) should be spent to capture it and support the required success probability. Besides these general properties, we do not assume any specific formula for $\Phi(p, i)$, which is ultimately system-specific. [1]

C. Playout Rate Control and Costs

Let's consider the natural packet playout rate of the media content to be a constant rate of R packets per time slot, i.e. we assume constant bit rate (CBR) coding and we also assume all packets have the same size. Deviations from R in a slot, as well as rapid rate variations in consecutive slots, are perceived as media quality degradations from the user.

The playout buffer can slow down and play fewer packets $r \leq R$ in a time slot (for reasons explained below). That is, there is a set of possible (system-specific) playout rates $\mathcal{R} = \{0, r_1, r_2, ..., r_l, ..., r_L\}$ with $0 < r_1 < r_2 < ... < r_l < ... < r_L = R$. In each time slot the playout buffer can choose to play $r \in \mathcal{R}$ packets (assuming they are available in the buffer). It may even choose to playout no packet at all ($r = 0$), even though there may be available packets, in order to re-buffer against future underflows, see section III-B.

There are several pressures to be considered and captured into performance/operational costs. First, to capture the user-perceived quality degradation when the playout r is less than its natural rate R, we introduce a slow-down cost $C_s(r) = f(R - r)$ which is 0 under normal playout rate R (or $f(0) = 0$), but is positive and rapidly increasing as r deviates from R. A significantly higher cost is incurred as the playout rate becomes too slow.

[1] There are also several other mitigation pressures that may be incorporated into $\Phi(p, i)$, for example, interference suppression, electromagnetic 'pollution' avoidance, etc. Of particular interest is the case where the transmitter is a battery-limited mobile node, then $\Phi(p, i)$ also directly incorporates the transmission power.

At the extreme $r = 0$, the playout is completely interrupted (freezed). In general, $f(.)$ should be determined by situation-specific perceptual considerations of typical users. As a concrete example, we use the quadratic cost $C_s(r) = C_1(R-r)^2$, as proposed in [10].

The slowdown cost $C_s(r)$ implicitly includes the cost of playout interruption. However, since the perceptual effect of an interruption or playout freeze is different than a slowdown, it may be desirable to have a separate cost $C_f(\tau)$ which expresses the cost of the playout freeze as a function of its duration, τ. In this manner, the media quality can also be evaluated based on the number (or frequency) and duration of the playout freezes. This cost can be straightforwardly incorporated into the proposed formulation, however in this paper for simplicity we consider the freeze cost as a special case of the slowdown cost $C_s(r)$, where $r = 0$.

Another important effect to consider is the playout smoothness. Indeed, playout rate variations in consecutive time slots degrades the user-perceived media quality. Let r' be the playout rate used in the previous time slot and r in the current one. Then, the system incurs in the current time slot a playout variation (or jitter) cost $C_v(r; r') = g(r-r')$, which is 0 for no variation (or $g(0) = 0$) and increases rapidly as r deviates from r' according to a potentially general functional form. In the later sections we use the quadratic cost $C_s(r) = C_2(r-r')^2$ similarly to [10]. [2]

D. System State and Joint Optimal Control

The objective is to transfer to and play at the receiver (Rx) all the media packets of the transmitter (Tx), minimizing the overall cost incurred in the process. The system state to be tracked in each time slot is

$$(n, i, b; r') \tag{2}$$

that is, the current packet backlog n in the transmitter buffer, the current interference state i in the channel, the current packet backlog b in the receiver playout buffer, and the playout rate r' used in the previous time slot.

The controls applied (decisions made) in each time slot are simply (m, p, r), that is, the packet transmission rate m, the transmission power p, and the packet playout rate r in the current time slot. Hence, the decisions made per slot are: how many packets to jointly transmit and at what power, and how many packets to play out.

Given this formulation, the system simply becomes a controlled Markov chain and, hence, we can develop a Dynamic Programming (DP) recursion to compute the optimal control. Let $J(n, i, b; r')$ be the cost-to-go, that is the minimum cost incurred from now on until all the content is played out, given that the optimal control is used and the current state is $(n, i, b; r')$. The quantity $J(n, i, b; r')$ satisfies then the following functional recursive equations, with $n \in \{1, 2, 3, ..., N\}$, $b \in \{0, 1, 2, 3, ..., B\}$, $i \in \mathcal{I}$, $r' \in \mathcal{R}$:

$$\begin{aligned} J(n, i, b; r') = \min_{m,p,r} \{ &\Psi(m) + \Phi(p, i) + C_s(r) + C_v(r; r') \\ &+ s(m, p, i) \sum_{j \in \mathcal{I}} q_{ij} J(n-m, j, b+m-r; r) \\ &+ [1 - s(m, p, i)] \sum_{j \in \mathcal{I}} q_{ij} J(n, j, b-r; r) \} \end{aligned} \tag{3}$$

where the joint (m, p, r) minimization is performed over the following selection sets, when the system is in state $(n, i, b; r')$.

1) m is optimized over $\mathcal{M}$ with the additional constraints that $m \le n$ and $m \le B - b + r$. The $m \le n$ constraint caps the number of transmitted packets in the current slot below the available packets n in the transmitter buffer. The $m \le B - b + r$ constraint prevents the transmitter from transmitting in the time slot more packets than the available empty places $B - b + r$ at the receiver playout buffer given the chosen playout rate r.
2) p can be optimized over a continuous and bounded power range $[0, P_{max}]$ (or a corresponding discrete one $\{1, 2, 3, ..., P_{max}\}$). For simplicity, we assume later that the transmitter power ceiling P_{max} is high enough, so that p is optimized over $[0, \infty)$. Considering a lower power ceiling would simply result in clipping the Tx power at the ceiling.
3) r is optimized over $\mathcal{R}$ with the additional constraint that $r \le b$, since the packets played out in a slot have to be less than those b currently in the buffer (assumed to be store-and-forward as opposed to cut-through).

We must also consider the boundary conditions

$$J(0, i, b; r') = C_v(R; r'), \text{ for all } i, b, r' \tag{4}$$

Indeed, after the transmitter buffer has been emptied the playout buffer can be emptied at zero cost, using the natural constant playout rate R. $C_v(R; r')$ captures the cost of changing playout rate to R from r' initially. [3]

Equation (3) can be explained as follows. Starting from state $(n, i, b; r')$ and exercising optimal control (m, p, r), the cost-to-go $J(n, i, b; r')$ is comprised of the current cost $\Psi(m) + \Phi(p, i) + C_s(r) + C_v(r; r')$ plus the expected future cost. The future cost is

[2] Other simple cost functions could be $C_v(r; r') = \delta |r - r'|^{2\epsilon}$ or even $C_v(r; r') = \delta(e^{\epsilon|r-r'|} - 1)$ for positive real δ, ϵ. Ultimately, $g()$ is determined by situation-specific for typical systems/users. We can further consider play rate variation costs $C_v(r; r', r'', ...)$ which track the rates $r', r'', ...$ used in several past time slots and capture the degradation of user experience due to past play-rate jitter. For simplicity, we limit ourselves to the baseline case of consecutive slots, which is powerful enough to capture the effect and spotlight the relevant intuition.

[3] Depending on r', the receiver may decide to increase rate to R gradually and results a smaller cost than $C_v(R; r')$. This effect is insignificant and neglected here.

1) $J(n-m,j,b+m-r;r)$ with probability $s(m,p,i)q_{ij}$, corresponding to the m transmitted packets being successfully received, r packets being played out, and the interference switching to state j, and
2) $J(n,j,b-r;r)$ with probability $[1-s(m,p,i)]q_{ij}$, corresponding to the m transmitted packets *not* being successfully received, r packets being played out, and the interference switching to state j.

We must then sum the above terms over all possible interference states $j \in \mathcal{I}$ in the next time slot. This results in the sum appearing in equation (3).

E. Computing the Optimal Control

Solving the DP recursion (3), (4) results in the optimal controls of the transmission mode $m(n,i,b;r')$, transmission power $p(n,i,b;r')$ and playout rate $r(n,i,b;r')$, when the system state is $(n,i,b;r')$. The solution of (3), (4) can be obtained using the *value iteration* method [11], as the following proposition asserts.

Proposition 1. There exists a stationary optimal control solution of (3), (4) obtainable by value iteration.
Sketch of Proof: Note that the DP terminates when the transmitter buffer empties $n=0$. Any policy that does not empty the transmitter buffer in finite time will incur an infinite cost. However, using some fixed rate and power so that $s(m,p,i)>0$, we can empty the transmitter buffer in N finite steps with positive probability [11]. The full proof is ommitted here for lack of space.

Using standard value iteration, we start with an initial guess for the minimal cost-to-go function $J^0(n,i,b;r')$ and plug it into (3), (4) to obtain the new J function, iterating until convergence. Let $J^k(n,i,b;r')$ be the function obtained after the k^{th} iteration. Then,

$$J^{k+1}(n,i,b;r') = \min_{m,p,r}\{\Psi(m)+\Phi(p,i)+C_d(r) + C_v(r;r') + s(m,p,i)\sum_{j\in\mathcal{I}} q_{ij}J^k(n-m,j,b+m-r;r) + [1-s(m,p,i)]\sum_{j\mathcal{I}} q_{ij}J^k(n,j,b-r;r)\} \quad (5)$$

Let us consider the simple - yet natural and interesting - case, where the transmission rate is fixed $m=R$ (hence, we do not need to optimize over m), the processing cost is negligible $\Psi(m)=0$, and the power cost is simply the transmitted power $\Phi(p,i)=p$ for each i. For fixed mode m, we use a generally accepted functional form $s(p,i) = \frac{p}{\alpha\cdot p+\beta\cdot i} = \frac{p/i}{\alpha\cdot p/i+\beta}$, which is increasing with the signal-to-interference ratio p/i. Similar functional forms are used in [12] for EGPRS systems and in [13] for 802.11a wireless LANs. Then, we get

$$J^{k+1}(n,i,b;r') = \min_{p,r}\{p - s(p,i)X_r^k(n,i,b;r') + Y_r^k(n,i,b;r')\} \quad (6)$$

where $X_r^k(n,i,b;r') = \sum_{j\in\mathcal{I}} q_{ij}J^k(n-m,j,b+m-r;r) - \sum_{j\in\mathcal{I}} q_{ij}J^k(n,j,b-r;r)$ and $Y_r^k(n,i,b;r') = C_s(r)+C_v(r_p,r)+\sum_{j\in\mathcal{I}} q_{ij}J^k(n,j,b-r;r)$

Minimizing (6) first over p we can obtain the optimal power $p_r^k(n,i,b;r')$ for fixed r in each step k, as follows:

$$p_r^k(n,i,b;r') = \frac{1}{\alpha}\left\{\sqrt{\beta X_r^k(n,i,b;r')i} - \beta i\right\} \quad (7)$$

for $i < \frac{X_r^k(n,i,b;r')}{\beta}$ and 0 otherwise. Substituting (7) then into (6), we can now optimize over r. At large k, as convergence occurs, we get:

$$p_r(n,i,b;r') = \frac{1}{\alpha}\left\{\sqrt{\beta X_r(n,i,b;r')i} - \beta i\right\} \quad (8)$$

for $i < \frac{X_r(n,i,b;r')}{\beta}$ and 0 otherwise. This power form is proved useful later in the design of heuristics.

III. Analyzing Insightful Special Cases

For the rest of the paper, we consider special cases of the general DP formulation, where only a subset of the control variables are used (e.g. individual controls of power-only or playout-only) which correspond to practical scenarios and also provide insight into the structural properties of the joint control policy. In particular, we focus on the simple -yet interesting and natural- case of section II-E: the power cost is simply $\Phi(p,i)=p$ and there is a single mode with a $m=R$, hence we do not need to optimize over m. For brevity, we also normalize the playout rate r, client buffer level b and size B, number of packets to be sent n and total number of packets N, with respect to the media's natural rate R. Therefore, $0\le r\le 1$, where $r=1$ corresponds to perfect quality (natural rate) while $r=0$ corresponds to freezing.

A. Power-Controlled Streaming (adapt p, fix $r=1$)

In this scenario, similar to [3], [4], the power p is the only control parameter that can be varied to combat the bad channel periods and provide smooth playout. The Rx initially waits for T time slots time, and then starts playing always at its natural rate 1. In case of underflow the playout will freeze until packets arrive. We slightly modify the optimality equations to describe this scenario. Notice that r and r' no longer belong to the control parameters and the system state, respectively.
for $b>1, N\ge n>0$

$$\begin{aligned} J(n,i,b) &= \min_p(p+s(p,i)\textstyle\sum_j q_{ij}J(n-1,j,b) \\ &+ (1-s(p,i))\textstyle\sum_j q_{ij}J(n,j,b-1))); \end{aligned} \quad (9)$$

for $b=1, N\ge n>0$

$$\begin{aligned} J(n,i,1) &= \min_p(p+s(p,i)\textstyle\sum_j q_{ij}J(n-1,j,1) \\ &+ (1-s(p,i))(C_v(1,0) \\ &+ \textstyle\sum_j q_{ij}J(n,j,0))); \end{aligned} \quad (10)$$

for $b=0, N\ge n>0$

$$\begin{aligned} J(n,i,0) &= \min_p (p + C_s(0) \\ &+ s(p,i)\textstyle\sum_j q_{ij} J(n-1,j,1) \\ &+ (1-s(p,i))\textstyle\sum_j q_{ij} J(n,j,0)); \end{aligned} \quad (11)$$

The recursive dynamics of Eq. (9)(10)(11) are similar to the general optimality equation. One subtle change concerns the additional cost $C_v(1,0)$ in Eq. (10): r changes from 1 to 0 only when the buffer level changes from $b = 1$ to $b = 0$, i.e. if the transmission is unsuccessful and $b = 1$. As in the general case, we can re-arrange the equation (9) by defining

$$\begin{aligned} X(n,i,b) &= \textstyle\sum_j q_{ij} J(n,j,b-1)) \\ &- \textstyle\sum_j q_{ij} J(n-1,j,b) \\ Y(n,i,b) &= \textstyle\sum_j q_{ij} J(n,j,b-1)) \end{aligned} \quad (12)$$

Similar X and Y can be defined for equations (10)(11).

Then, we can apply the value iteration method and obtain the optimal transmission power policy as:

$$p(n,i,b) = \begin{cases} \frac{1}{\alpha}(\sqrt{\beta X(n,i,b) i} - \beta i), & i < \frac{X(n,i,b)}{\beta} \\ 0, & \text{else} \end{cases} \quad (13)$$

An interesting observation from Eq. (13) is that the optimal power is merely a function of X and i. From (12), we can interpret $X(n,i,b) = \sum_j q_{ij} J(n,j,b-1) - \sum_j q_{ij} J(n-1,j,b)$ as the cost of sending one packet instantaneously, which decreases the number of remaining packets to $n-1$ and increases the client buffer level to $b+1$. For $b = 1$, X also captures the potential cost of rate variation. Thus X can be interpreted as the expected cost of an unsuccessful transmission.

We now list the structural properties of X and provide some intuitive explanations.

Proposition 2. $X(n,i,b) \geq 0$.
Intuitively, the expected cost should be non-negative.

Proposition 3. $X(n+1,i,b) \geq X(n,i,b)$.
$X(n,i,b)$ is increasing in n. With a fixed buffer level, the more packets need to be transmitted, the more likely a failed transmission will cause buffer underflow, which leads to playout disruptions and freezes.

Proposition 4. $X(n,i,b) \geq X(n,i,b+1)$ for $b \geq 1$.
$X(n,i,b)$ is decreasing in b. With a smaller playout buffer level, a failed transmission will more likely cause future buffer underflow.

Proposition 5. $\forall b$, $X(n,i,b) \to X^*(i,b)$ as $n \to \infty$.
Fig. 2 plots the function X vs. the playout buffer level b for $n = 1$ to 100 where the above structural properties are demonstrated. This convergence property for large n is interesting for the continuous streaming scenario where $n = \infty$. It means, we can use the optimal power obtained at large n values as the power control policy.

B. Power-Control and Re-buffering (vary p, $r \in \{0,1\}$)

In this practical and widely used scenario, we assume a slightly more complicated playout rate control at the client.

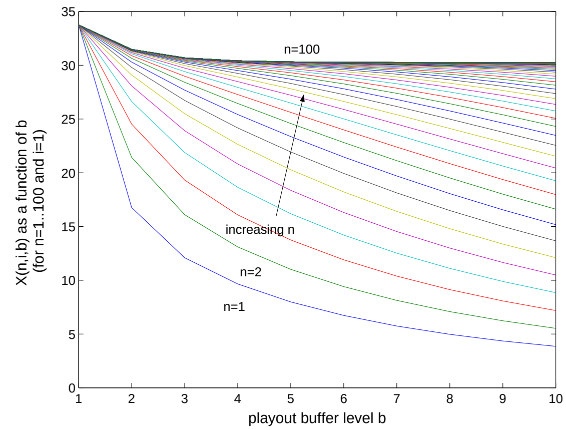

Fig. 2. $X(n,i,b)$ vs. playout buffer b (for $i = 1$ and $n = 1, ..100$.

The Rx can re-buffer but cannot vary the playout speed. At each time slot, the rate controller can either choose to play out the media packet from the buffer at the natural rate $r = 1$, or remain idle $r = 0$ to build up the buffer level and provide continuous playout for the remaining media sequence. However, rebuffering happens at the cost of $C_s + C_v$. This scenario is clearly a special case of the general model introduced in section II. The optimality equations remains the same, and the set of available playout rates r is limited to $\{0,1\}$.

C. Adaptive Playout Control (fix p, vary r)

In this scenario, we assume a primitive wireless transmitter where power is fixed. However, the Rx is able to perform re-buffering and playout rate variation, in order to deal with bad channel conditions. Intuitively, when channel is good, the optimal policy tends to choose the natural playout rate; when the channel is bad, the optimal playout policy should slow down the playout to prevent buffer underflows. We modify the optimality equation to capture this effect mathematically. p is constant therefore $s(p,i) = s(i)$.

for $b > 0, N \geq n > 0$:

$$\begin{aligned} J(n,i,b,r') &= \min_r (C_s(r) + C_v(r',r) \\ &+ s(i)\textstyle\sum_j q_{ij} J(n-1,j,b-r+1,r) \\ &+ (1-s(i))\textstyle\sum_j q_{ij} J(n,j,b-r,r)); \end{aligned} \quad (14)$$

for $b = 0, N \geq n > 0$:

$$\begin{aligned} J(n,i,b,r') &= (C_s(0) + C_v(r',0) \\ &+ s(i)\textstyle\sum_j q_{ij} J(n-1,j,1,0) \\ &+ (1-s(p,i))\textstyle\sum_j q_{ij} J(n,j,0,0)); \end{aligned} \quad (15)$$

The boundary condition for $n = 0$ is the same. Using the value iteration method, we obtain the optimal policy.

D. Performance Comparison of the Optimal Policies

We now compare the performance of the optimal policies, for all the above-mentioned scenarios:

(a) No power - no playout control: a baseline to illustrate the benefit of adding controls.

(b) Power Control, Constant Playout: as in section III-A,

(c) Power Control with Rebuffering: as in section III-B.

(d) Constant Power, Playout Control: as in section III-C.

(e) Joint Power-Playout Control: as in section II, but assuming a single transmission mode. It takes full advantage of controlling power and playout rate.

The setup for the performance comparison of the optimal policies is the following.

- The interference i takes values according to a 2-state Markov chain. When the channel is good, $i = 1$; when the channel is bad, $i = 100$. The transition matrix between these states is $Q = [0.86\ 0.14;\ 0.28\ 0.72]$ and leads to average durations of the good and the bad state 7 and 3.5 time slots respectively.[4] We solve the optimality equation off-line and store the optimal policy for online use. Then, we simulate multiple sample paths of the interference and for each one, we use the optimal control policy.
- The values of power shown in all figures are relative to the interference level. Thus, no units are mentioned for i or p, and we are mainly interested in the signal-to-interference ratio p/i. We assume there are $N = 100$ pre-stored media packets. The Rx buffer size B is set to 10. Recall that r is normalized to 1, and we use possible playout rates $r \in \{0, 0.5, 0.6, 0.7, 0.8, 0.9, 1\}$
- We are interested in two performance metrics. The first is the media quality degradation due to variation of r. We compute the 'average degradation per packet' as the sum of $C_s + C_v$ over all time slots needed to transfer all packets and divided by the number of packets N: $C_{avg} = 1/N \cdot \sum_{t=1}^{t=t_{max}} (C_s(t) + C_v(t))$. The second is the transmission power spent. We calculate the 'average power per packet' in similar way: $P_{avg} = 1/N \cdot \sum_{t=1}^{t=t_{max}} p(t)$. In the DP formulation, we weight the media quality cost by a factor w: $C_{total} = P_{avg} + W \cdot C_{avg}$. Depending on w, more importance can be given to media quality or to power, leading to different optimal control. The objective is to improve the power-quality tradeoff.

The results of the comparison are shown in figures 3 and 4. Each curve corresponds to an optimal policy. For policies (b)(c)(e), the weight factor for the media quality cost is varied in order to obtain the different operation points on the curve. The more weight on media quality, the more willing we are to spend power, the further to the right of the tradeoff curves we operate. For policies (a)(d), the constant transmission power is varied to demonstrate this tradeoff. Fig. 3 compares the policies in operational ranges where power is important, therefore power consumption is kept in low-to-moderate ranges. Fig. 4 compares the policies when media quality is more important while power consumption can be high.

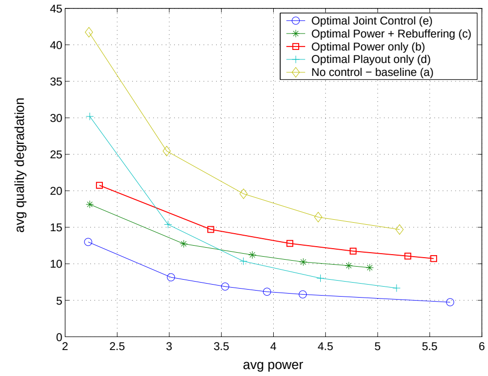

Fig. 3. Comparison of optimal policies in low-moderate power ranges, in terms of the tradeoff between media quality degradation and power consumption.
In these scenarios, it is important to preserve power, while some degradation in media quality can be tolerated. Notice that the power-controlled algorithms (b)(c) outperform the adaptive playout algorithm (d) significantly. Without power control, even the optimal playout policy (d) performs almost 2 times worse than power-controlled algorithms (b)(c). As the power constraints are relaxed, and we move to the right of the figure, the adaptive playout algorithm (d) starts to outperform the power controlled algorithms (b)(c). Comparing (b) to (c), we observe that allowing the playout controller to adaptively rebuffer improves the performance substantially.

In both figures, the absence of any control leads to high degradation and high power consumption, thus performs the worst. Allowing power control and/or playout control improves the performance. As expected, joint control of these two parameters clearly outperforms each individual control and achieves the lowest degradation-power tradeoff. E.g. for the same average power spent, the joint optimal policy (e) reduces the quality degradation by 20% to 75% compared to all other policies. Detailed comments can be found in the figures' captions.

IV. Heuristics: Design and Performance

The optimal policies studied in the previous section show the performance limits of streaming over interference-limited wireless links. However, they require computational power and knowledge of the channel statistics, which may be difficult or impossible to obtain. In this section, we use the structural properties of the optimal solutions to design practical, low-complexity heuristics that achieve near-optimal performance (because they share the same structural properties with the optimal solution).[5] For the performance evaluation of these heuristics, we use the setup of section III-D.

(b1) Power Control (only) Heuristic: One can construct a power control by immitating the structural properties of

[4]We present results for this scenario, for which the channel is stressed by the significant difference between the good and the bad state. In that case the control policies make a difference and key-aspects of our approach are demonstrated. On the contrary, if power is not an issue, it can be as high as required to deliver the packets.

[5]Thus, we name the heuristics after the optimal policies that inspire them. E.g. (d2) is a playout heuristic based on the optimal playout policy (d). (e1) is based on the joint optimal policy (e).

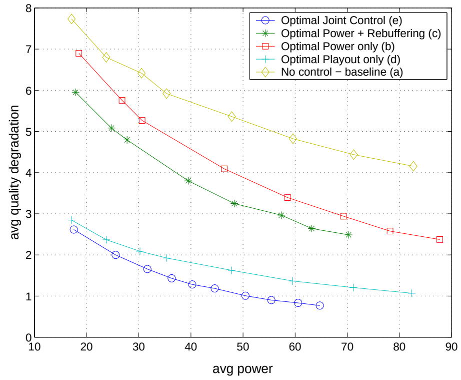

Fig. 4. Comparison of optimal policies in high power ranges, in terms of the tradeoff between media quality degradation and power consumption. This scenario is the opposite of Fig.3: the power constraints are relaxed and it is important to preserve media quality. Notice that the adaptive playout algorithms become necessary in this case. Policy (d) with adaptive playout outperforms the fixed playout policies (b)(c): for the same average power, it reduces the quality degradation by 50%. The joint optimal control (e) further reduces the quality degradation by another 30% more than the adaptive playout (d) alone, which demonstrates the merit of joint power-playout control.

III-A and in particular by analytically approximating $X()$. The interested reader is referred to [3] for details.

(d1) Fixed Power, Threshold Playout Heuristic. A playout-only heuristic, widely used in practice is to look only at the playout buffer level and decide the playout rate. E.g. a single threshold may be used, as in [9] and the playout slows down/speeds up whenever the buffer level is below/above that threshold. In general, more than one thresholds can be defined and different playout rates can be assigned to different ranges of playout buffer occupancy. E.g.: given a buffer level b, and a set of available playout rate $\{r_0 = 0, r_1, r_2, ..., r_L\}$

$$r = r_l, \text{ when } B_l < b \leq B_{l+1} \quad (16)$$

where $B_0 = -1$ (just a notation to include the case $b = 0$ when $B_1 \geq b \geq B_0$) and $B_k = B$. [6]

(d2) Fixed Power, Adaptive Threshold Heuristic. In section III-C, we observed that when p is constant, the optimal playout policy (d) uses several thresholds to divide the playout buffer occupancy into intervals, each corresponding to a different playout rate. Intuitively, when i and p vary, these thresholds should vary accordingly. E.g. when i increases, the thresholds should increase, to slow down playout and prevent underflow.

We design an adaptive threshold policy that observes the channel condition (captured by the probability of successful transmission $s = s(p, i)$) and varies the playout thresholds $B_1, ..., B_L$ accordingly:

$$B_l = \frac{B}{L+1}(l - \frac{L-l+1}{L}\frac{s}{s+\gamma}) \quad (17)$$

Eq. (17) describes the following simple operation. At $s = 0$, the thresholds evenly divide the buffer into $K+1$

[6]For our simulations, we first used 6 equally-spaced buffer thresholds B_1,B_2...B_6 and 7 possible playout rates $\{0,\ 0.5\ 0.6\ 0.7\ 0.8\ 0.9\ 1\}$. We then varied the number and the granularity of thresholds.

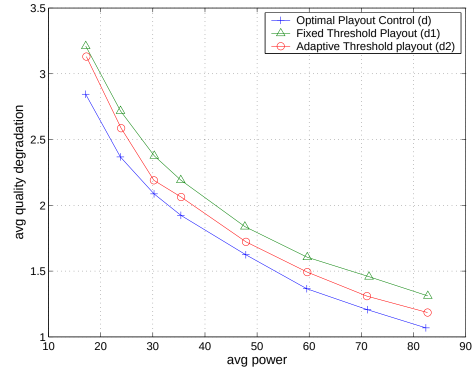

Fig. 5. Comparison of playout control policies, via tradeoff of media quality vs. power consumption.

regions with width $W = \frac{B}{L+1}$. As s increases, threshold B_l is reduced by $(L+1-l)\delta(s)$. For example, with $L = 6$ thresholds, B_1 is reduced by $6\delta(s)$, B_2 is reduced by $5\delta(s)$ and so on. This allows the Rx to play the received packets at a faster rate. We pick $\delta(s) = \frac{W}{L}\frac{s}{s+\gamma}$ to be the $\frac{1}{L}$ of the initial region width W multiplying a function of s. γ is a tuning parameter to adjust the sensitivity of these thresholds to variations in the channel conditions. Given these thresholds, we can use equation (16) to find the corresponding playout rate.

In Fig. 5, we compare the performance of the Adaptive Threshold heuristic (d2) to the Optimal Playout Policy (d) and the Fixed Threshold heuristic (d1). Clearly, both heuristics perform close to the optimal. Furthermore, (d2) reduces the performance gap between (d1) and (a) by 50%, thanks to its ability to adapt the thresholds to the channel conditions.

We further investigated the sensitivity of our heuristic (d2) to the choice of (i) playout rate range (excluding rate 0) and (ii) playout rate resolution. We found that our scheme is robust to these choices. In the simulations of Fig.5, we set the range of the non-zero possible rates from 0.5 to 1 with a resolution of 0.1. Intuitively, the performance should become worse as the range decreases and the resolution becomes coarser. Indeed, Fig.6 confirms this intuition through simulations.

In Fig.6(a), we vary the range and keep the resolution constant at 0.05. For example, $range = 0.2$ leads to possible playout rates $r \in \{0, 0.8, 0.85, 0.9, 0.95, 1\}$. As expected, the performance gain decreases as the range of non-zero playout rates becomes narrow. [7] We further observe that the proposed heuristic (d2) outperforms the

[7]At the extreme case that the range becomes 0, $r \in \{0, 1\}$, and the playout degenerates to re-buffering. In this case, the fixed threshold policy (d1) actually performs worse than having no control, with a 10% performance loss! This is because with constant power $p = 21$ units, the average channel throughput $s(p, i)$ is high but the fixed threshold playout (d1) does not adapt to this good condition. E.g. it unnecessarily re-buffers when Rx buffer level is half empty. In contrast, the adaptive threshold playout (d2) and the optimal playout policy (d) can adapt to the good channel condition and improve the media quality.

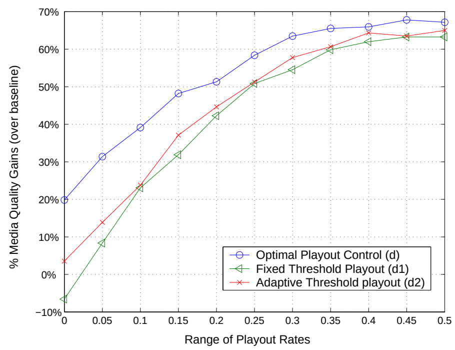

(a) Sensitivity to the range of playout rate

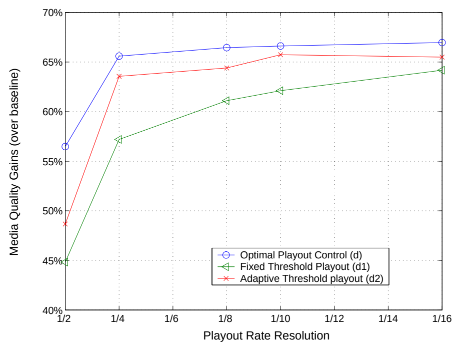

(b) Sensitivity to the resolution of rate

Fig. 6. Sensitivity of three playout control schemes to the range and resolution of available playout rates. The Y-axis is the gain (reduction in media quality degradation compared to the no-control approach (a)), achieved by the three playout control policies (d)(d1)(d2). The transmission power is fixed at $p = 21$ units.

fixed threshold policy (d1) in all parameter ranges. In Fig. 6(b), we vary the resolution and we keep the range the same: $\{0.5..1\}$. E.g. $resolution = 1/4$ means that $r \in \{0, 0.5, 0.75, 1\}$. As expected, the performance gain increases with the resolution of the playout rates. However, the gains for all three policies saturate at resolution finer than $1/8$. Thus, in designing a good playout control policy, one does not need to have an extremely fine resolution in playout rates.

(e1) Joint Power-Playout Heuristic. In section III-A, we observed the structural properties of the power-only optimal policy. Recall that the auxiliary parameter X captured the expected cost of an unsuccessful transmission. X was a function of number of packets to be transmitted n (backlog pressure) and playout buffer level b (buffer underflow pressure). We observe that a similar X is defined in the optimality equation for the joint power-playout control in section II-E. We further notice that given X, the optimal power can be calculated as

$$p = \begin{cases} \frac{1}{\alpha}(\sqrt{\beta X i} - \beta i), & i < \frac{X}{\beta} \\ 0, & \text{else} \end{cases} \qquad (18)$$

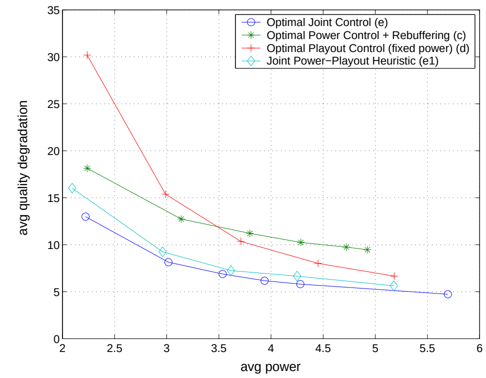

Fig. 7. Comparing the Joint Control Heuristic (e1) to optimal policies, in low-moderate power regimes.

We now exploit this property to design a good joint power-playout heuristic. We approximate X by an analytical function of n, b, while satisfying the key structural Propositions 2-5. Specifically:

$$\begin{aligned} X(n,b) &= K\exp((B+1-b)^{(\eta n)^{\theta}}) + \epsilon; \\ K &= C\exp(-B^{(\eta n)^{\theta}}); \end{aligned} \qquad (19)$$

where the parameters η and θ adjust the sensitivity of X to the backlog pressure n. ϵ is chosen as the minimum of all interference levels and is added to ensure that the system won't enter a deadlock state. The heuristic works as follows: With equations (19)(18), we can obtain the transmission power given the number of packets to be transmitted n, the channel interference i and the playout buffer level b. Then, the Tx uses this power level to send the current packet. Since we assume that the same state information (n, i, b) is also known at the Rx, the Rx is able to calculate the same transmission power using exactly the same formula. Given the channel interference i and the playout buffer level b, we can use the adaptive threshold playout heuristic (d2) to choose the playout rate based on the calculated transmission power p.

Fig.7 shows the performance of the joint power-playout heuristic under tight power constraints. The simulation setup is the same as in Fig. 3 for the optimal policies (e),(c),(d) which are re-plotted here for comparison. We observe that the proposed joint control heuristic (e1) performs quite close to the joint optimal (e) and outperforms substantially all other policies. Note that the optimal policies require the solution of the optimality equation, hence are computational intensive, while the heuristic (e1) has minimal computational complexity. Furthermore, the optimal policies require knowledge of the channel statistics (basically the transition matrix Q), while the joint-control heuristic requires only a good estimate of the current channel interference.

Fig.8 shows the performance of the joint power-playout heuristic under relaxed power consumption constraints. The simulation setup is the same as in Fig. 4. Again, for performance comparison, we re-plot the curves for the optimal joint control (e), and the three playout control

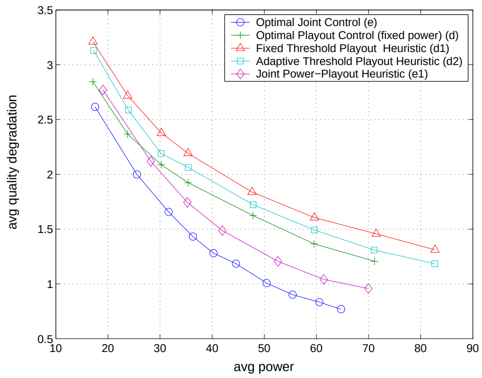

Fig. 8. Comparing the Joint Control Heuristic (e1) to the optimal policy as well as to other heuristics, in high power regimes.

TABLE I

% IMPROVEMENT IN MEDIA QUALITY (FOR THE SAME POWER CONSUMPTION), ACHIEVEDR BY THE VARIOUS POLICIES.

Optimal Policies and Heuristic Algorithms		Low power	High power
Optimal Power (only) Control	(b)	50%	31%
Optimal Power Control with Rebuffering	(c)	57%	40%
Optimal Playout (only) Control	(d)	29%	70%
Optimal Joint Power-Playout Control	**(e)**	**69%**	**83%**
Fixed Threshold Playout Heuristic	(d1)	24%	64%
Adaptive Threshold Playout Heuristic	(d2)	25%	68%
Joint Power-Playout Heuristic	**(e1)**	**64%**	**78%**

policies (d)(d1)(d2). We do not show neither the power control policies (b)(c) nor the baseline (a), due to their relative worse performance under this scenario. The proposed joint control heuristic (e1) again outperforms all other heuristics, including the playout-only heuristic (d2).

In Table IV, we summarize the gains in media quality by the optimal policies and their corresponding heuristics, first for the operational regimes of low power consumption (2.2 units per packet) and then for the high power consumption (60 units per packet). The gains are computed with respect to the no-control baseline (a). Clearly, our proposed heuristic (e1) can achieve near-optimal media quality (i.e. close to the optimal policy (e)) over all power regimes, at significantly lower complexity.

V. POWER-RESPONSIVE INTERFERENCE

So far, we assumed Markovian, non-responsive interference. The next important methodological step was to simulate our algorithms in a realistic wireless environment, where all links perform power control, resulting to responsive interference. When our 'primary' media-streaming link raises its power, other wireless 'background' links, transmitting in the same neighborhood, will also raise their powers and induce more interference on it, etc. Due to the responsive nature of interference, the optimality equations can no longer be formulated.

We used OMNeT++ [14] and simulated our heuristics in an environment with a large number of 'background' links implementing the PCMA power-control algorithm, [1]. We found that the heuristics inspired by the optimal solutions perform extremely well. The joint power-playout heuristic (e1) outperforms again the other schemes: it reduces the power consumption by more than 50% and achieves a high media quality. Due to space limitations, we do not present details of this study.

VI. CONCLUSION

In this paper, we examined the important problem of joint power-playout control for media streaming over interference-limited wireless links. We formulated the problem in a dynamic programming framework, and identified the optimal policies for power-only, playout-only, and joint power-playout control. Finally, we designed and evaluated practical, low-complexity heuristics that achieve near-optimal performance.

In this basic scenario, we considered reliable delivery of all packets, we varied the playout rate, but we did not allow packets to be dropped. However, our approach can be naturally applied to a variety of streaming scenarios, performance costs and channel characteristics. We are currently working on applying the same approach to delay-sensitive streaming, where frames can occasionally be dropped or playout can speed-up, in order to catch up with live content.

REFERENCES

[1] N.Bambos and S.Kandukuri, "Power Controlled Multiple Access (PCMA) in Wireless Communication Networks," *INFOCOM 2002*, pp. 368-395.
[2] S.Kandukuri and N.Bambos, "Multimodal Dynamic Multiple Access in Wireless Packet Networks," *INFOCOM 2001*, pp.199-208.
[3] Y. Li and N. Bambos, "Power-controlled wireless links for media streaming applications", *Wireless Telecom. Symposium*, May 2004.
[4] Y. Li and N. Bambos, "Power-controlled streaming in interference limited wireless networks", *to appear in Broadband Networks*.
[5] N.Färber, B.Girod, "Wireless Video," in A. Reibman, M.-T. Sun (eds.), *Compressed Video over Networks, Marcel Dekker, 2000*
[6] D. Towsley, H. Schulzrinne, R. Ramjee, J. Kurose, "Adaptive playout mechanisms for packetized audio applications in wide-area networks," *Infocom 1994*, Jun 1994.
[7] S. Moon, J. Kurose, and D. Towsley, "Packet audio playout delay adjustment: performance bounds and algorithms," *ACM/Springer Multimedia Systems*, vol. 6, pp. 17-28, Jan. 1998.
[8] Y. Liang, N. Farber, and B. Girod, "Adaptive playout scheduling and loss concealment for voice communications over the networks," *IEEE Trans. on Multimedia*, Apr. 2001.
[9] M. Kalman, E. Steinbach, and B. Girod, "Adaptive media playout for low delay video streaming over error-prone channels," *IEEE Transactions on Circuits and Systems for Video Technology*.
[10] M. Kalman, E. Steinbach, and B. Girod, "Rate-distortion optimized video streaming with adaptive playout," in *IEEE ICIP-2002*, vol. 3, pp. 189-192, Sep 2002.
[11] D. Bertsekas, "Dynamic Programming and Optimal Control", vol. 2, Athena Scientific, 1995.
[12] D. Krishnaswamy. "Network-assisted link adaptation with power control and channel re-assignment in wireless networks," in *3G Wireless 2002 conference*, pp. 165-170, Sep 2002.
[13] K.B. Song and A. Mujtaba, "On the code-diversity performance of bit-interleaved coded OFDM in frequency-selective fading channels," in *Proc. of Vehi. Tech. Conf.* , Fall 2003.
[14] OMNeT++ Discrete Evemt Simulation System, publicly available at: *http://www.omnetpp.org/*