

# Joint Power-Playout Control for Media Streaming Over Wireless Links

Yan Li, *Member, IEEE*, Athina Markopoulou, *Member, IEEE*, Nick Bambos, and John Apostolopoulos, *Member, IEEE*

**Abstract**—Media streaming applications over wireless links face various challenges, due to both the nature of the wireless channel and the stringent delivery requirements of media traffic. In this paper, we seek to improve the performance of media streaming over an interference-limited wireless link, by using appropriate transmission and playout control. In particular, we choose both the power at the transmitter and the playout scheduling at the receiver, so as to minimize the power consumption and maximize the media playout quality. We formulate the problem using a dynamic programming approach, and study the structural properties of the optimal solution. We further develop a justified, low-complexity heuristic that achieves significant performance gain over benchmark systems. In particular, our joint power-playout heuristic 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 (playout) quality 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 the investigation.

**Index Terms**—Adaptive playout, cross-layer optimization, dynamic programming, multimedia streaming, network control, quality-of-service, wireless multimedia.

## I. INTRODUCTION

RECENT advances in video compression and streaming, and also in wireless networking technologies (next-generation cellular networks and high-throughput wireless LANs), are rapidly opening up opportunities for media streaming over interference-limited and erratic wireless channels. Supporting high-quality media applications over wireless channels introduces technical challenges 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.

Manuscript received April 12, 2005; revised October 17, 2005. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Anna Hac.

Y. Li was with the Electrical Engineering Department, Stanford University, Stanford, CA 94304-9510 USA. He is now with Qualcomm, Campbell, CA 95008 USA (e-mail: liyan@stanfordalumni.org).

A. Markopoulou was with the Electrical Engineering Department, Stanford University, Stanford, CA 94304-9510 USA. She is now with the Electrical Engineering and Computer Science Department, University of California at Irvine, Irvine, CA 92697-2625 USA (e-mail: amarko@stanfordalumni.org).

N. Bambos is with the Electrical Engineering Department, Stanford University, Stanford, CA 94304-9510 USA (e-mail: bambos@stanford.edu).

J. Apostolopoulos is with Hewlett-Packard Laboratories, Palo Alto, CA 94304 USA (e-mail: japos@hpl.hp.com).

Digital Object Identifier 10.1109/TMM.2006.876236

We briefly note the following intuitive points regarding the core tradeoff of transmitter power versus 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., 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. To mitigate the risk of playout buffer underflow, the system has the following options:

- 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, deplete the battery power of mobile nodes and thus should be very carefully exercised.
- 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 avoided.

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 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 competing goals define 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 personal digital assistant (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, but does not require speeding up beyond the nominal speed.

The rest of the paper is structured as follows. In Section II, we present related work. In Section III, we develop a parsimonious formulation/model that captures the key performance tradeoff and provides a framework for computing and evaluating efficient dynamic power/playout control algorithms. This formulation is 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 IV, we study interesting and practical special cases, when a subset of the control variables is used, e.g., power-only control or playout-only control. In Section V, we design practical heuristics based on the analysis of the optimal policies and we evaluate their performance in a markovian interference environment, as well as in a responsive interference environment; we demonstrate that the proposed heuristics achieve substantial performance gains over benchmark systems, yet they are simple and have low-complexity. In Section VI, we discuss how our framework could be applied/extended to address real-world streaming scenarios. Section VII concludes the paper.

## II. RELATED WORK

Wireless video is a large problem space with several aspects and control parameters; a nice discussion can be found in [5]. In this paper, we focus on controlling power at the transmitter and playout at the receiver and on exploring the tradeoff between the two. In the past, these problems have been studied separately. Power control offers potentially many benefits: it can increase network capacity (via interference mitigation), maintain link quality (via adaptation to variations), and conserve 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], [14], [16], [17]. 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 [20], [21], [26] and video [12], [13]. Our work focuses on the joint power-playout control, and the improvements that result from the interaction between the two.

There is a considerable body of related work on optimizing media transmission over a network and on cross-layer optimization over wireless in particular, e.g., [23]. Reference [24] uses dynamic programming to obtain the best quantization levels and perform rate control. A discussion on recent advances on network-adaptive video can be found in [6].




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  $r$ .

Among those, [3] develops a framework for rate-distortion optimized streaming: in every transmission opportunity, a decision is made as to which media units to transmit and which to discard, so as to maximize the expected quality of received video subject to a constraint in the transmission rate. In contrast, we take a network control perspective, with power and playout being the main control variables of interest in this paper. In particular, we allocate power and playout rate to hide the channel variations, while RaDiO allocates bandwidth efficiently across video frames. In other words, our control policy “rides” on channel state fluctuations while RaDiO “rides” on frame size/transmission status fluctuations. In this paper, we only consider scenarios where the sender is not allowed to drop any packets and focus on control policies that exploit channel fluctuations, as opposed to media units dependencies. In summary, both works try to improve video transmitted over networks, and provide analytical frameworks based on dynamic programming. However: 1) the two modeling approaches are different and 2) they have been applied so far to address orthogonal axes of the problem (namely power/playout control versus rate-distortion optimized scheduling). Our follow-up work, [18], extends our framework to include packet-scheduling control and distortion costs.

An early version of this work appeared in the IEEE Packet Video Workshop [19]. This journal version is extended to include: 1) additional results, including illustrative examples of the joint control, the re-buffering special case, the responsive interference environment, and a range of channel parameters and 2) a discussion of applications to real-world scenarios and possible extensions.

## III. BASIC MODEL AND PROBLEM FORMULATION

In this section, we describe the streaming scenario under study and introduce the basic model, performance tradeoffs and control issues. We formulate 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  $t = 0, 1, 2, 3, \dots$ . The transmitter is equipped with a buffer where  $N$  media packets are 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.

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 PVR 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 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 buffer underflow duration (freezes) and jitter as low as possible; and this must be achieved while minimizing the power requirements at the transmitter. The formulation provides a method to specify and tradeoff transmitter power for playout quality.

### A. Wireless Channel Model

The condition of a wireless channel varies with time. This is typically due to both the varying quality of the wireless channel (path-loss, shadowing and fast fading) and to the varying level of interference caused by other users (e.g., as users join and leave dynamically the channel). Without loss of generality, we denote with  $i$  the channel interference and/or impairment. In this paper, we focus on interference-limited environments, therefore  $i$  will refer simply to the interference level itself, and  $p/i$  will be the signal-to-interference ratio (SIR). However,  $i$  could be extended to also incorporate channel gain variations and noise:  $i = (\text{interference} + \text{noise})/\text{gain}$ , and  $p/i$  would then be the signal-to-noise and interference-ratio (SINR). This ratio  $p/i$  ultimately determines the probability of successful transmission of a packet.<sup>1</sup>

To capture the time-varying nature of the channel, we consider that  $i$  fluctuates according to a time-homogeneous Markov chain, taking values in the finite set  $\mathcal{I}$  of all attainable interference states.  $i$  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. For most of the paper, we assume that  $i$  is not responsive to transmitter actions—like power variations—but is driven by an agent extraneous to the system in this model. For example, each state could correspond to the number of users using the channel, and will change as users dynamically join and leave the system; this is out of the control of the single connection under study. The case of responsive interference is explicitly investigated in Section V-C; it is also implicitly incorporated into the current model via the channel stress cost  $\Phi(p, i)$  discussed below.

### B. 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

<sup>1</sup>Throughout the paper, the unit of power and interference is considered to be the minimum level of interference  $i_{\min} = 1$ . In all figures, the values of both  $p$  and  $i$  refer to the above units, i.e., they are relative to  $i_{\min}$  (but not translated to decibels). This is because the successful transmission of a packet depends on the ratio  $p/i$  and not the absolute values of  $p$  and  $i$ , as discussed in Section III-C.

more generally implementing some transmission mode  $m \in \{0, 1, 2, 3, \dots, M\}$  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  $\mathcal{M} = \{0, m_1, m_2, \dots, m_k, \dots, m_K\}$  with  $0 < m_1 < m_2 < \dots < m_k < \dots < m_K < M$ . The available modes are typically system-specific.

### C. 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

$$s(m, p, i) \quad (1)$$

is the probability that all  $m$  packets will be successfully recovered at the receiver and, hence, removed from the transmitter buffer.<sup>2</sup> Note that the function  $s(m, p, i)$  should be decreasing in  $m$ , increasing in  $p$ , 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 structural properties, we do not assume any specific formula for the success probability, which is ultimately system-specific, as discussed in Section VI-C.

Transmitting power  $p$  in a time slot when the interference is  $i$ , introduces a cost  $\Phi(p, i)$  paid in that slot. This cost first captures the power spent in that times slot, which may be an issue in battery-constrained devices. It also 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 power 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  $i$ ; the intuition is that the more congested the channel is (the less the available bandwidth resource), the more power 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.<sup>3</sup>

<sup>2</sup>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 (perhaps 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.

<sup>3</sup>There are also several other mitigation pressures that may be incorporated into  $\Phi(p, i)$ , e.g., 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.

#### D. Playout Rate Control and Costs

Let us 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 all packets have the same size (the variable bit rate (VBR) case is treated as an extension in Section VI-C). Deviations from  $R$ , as well as rapid rate variations in consecutive slots, are perceived by the user as degradation in media quality. 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  $\mathcal{R} = \{0, r_1, r_2, \dots, r_l, \dots, r_L\}$  with  $0 < r_1 < r_2 < \dots < r_l < \dots < r_L = R$ . In each time slot, the playout buffer can choose to play  $r \in \mathcal{R}$  packets (assuming they are available in the buffer), or may choose to not play out any packet at all ( $r = 0$ ), in order to re-buffer against future underflows.

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 slowdown 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 function ( $f(\cdot)$ ) as  $r$  deviates from  $R$ . A significantly higher cost is incurred as the playout rate becomes too slow. At the extreme  $r = 0$ , the playout is completely interrupted (frozen). In general,  $f(\cdot)$  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 [13].

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,  $\tau$ . In this manner, the media quality can also be evaluated based on the number and duration of the playout freezes. This cost could be 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  $g(\cdot)$ . In the later sections, we use the quadratic cost  $C_v(r) = C_2(r - r')^2$  similarly to [13].<sup>4</sup>

#### E. 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

<sup>4</sup>Other simple cost functions could be  $C_v(r; r') = \delta|r - r'|^{2\epsilon}$  or even  $C_v(r; r') = \delta(e^{|r-r'|} - 1)$  for positive real  $\delta, \epsilon$ . Ultimately, the functional form  $g(\cdot)$  for  $C_v$  is determined by situation-specific perceptual considerations. We can further consider play rate variation costs  $C_v(r; r', r'', \dots)$  which track the rates  $r', r'', \dots$  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.

cost incurred in the process. The system state to be tracked in each time slot is

$$(n, i, b; r') \quad (2)$$

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  $n \in \{1, 2, 3, \dots, N\}$ ,  $b \in \{0, 1, 2, 3, \dots, B\}$ ,  $i \in \mathcal{I}$ ,  $r' \in \mathcal{R}$

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

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  $m \leq n$  and  $m \leq B - b + r$ . The  $m \leq n$  constraint caps the number of transmitted packets in the current slot below the available packets  $n$  in the transmitter buffer. The  $m \leq B - b + r$  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)  $p$  is optimized over a continuous and bounded power range  $[0, P_{max}]$  (or a corresponding discrete one  $\{0, P_1, \dots, P_{max}\}$ ).<sup>5</sup>
- 3)  $r$  is optimized over  $\mathcal{R}$  with the additional constraint that  $r \leq b$ , 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).

After the transmitter buffer has been emptied, the playout buffer can be emptied at zero cost, using the natural constant playout rate  $R$ . This is captured by the boundary conditions

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

<sup>5</sup>One could even optimize  $p$  over  $[0, \infty)$ . (Even then, the optimal policy will choose  $p^* \leq \infty$ .) A power ceiling  $P_{max} \leq \infty$  simply results in clipping the Tx power at the ceiling.

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 as follows.

- 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  $j$ .
- 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 (3).

### F. 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 [2], as the following proposition asserts.

*Proposition 1:* There exists a stationary optimal control solution of (3) and (4) obtainable by value iteration.

*Sketch of Proof:* The DP terminates when the Tx buffer empties, i.e.,  $n = 0$ . Any policy that does not empty the transmitter buffer in finite time will incur an infinite cost. And there exist policies under which with positive probability the transmitter buffer will be emptied in finite time. (For example, fix transmission rate and power so that  $s(m, p, i) > 0$ ). These two conditions guarantee the existence of a stationary optimal control [2]. The full proof is omitted 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 the dynamic programming recursive (3), (4) to obtain the new  $J$  function. Let  $J^k(n, i, b; r')$  be the function obtained after the  $k^{\text{th}}$  iteration. Then

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

The iteration continues until convergence, i.e., until  $J_k - J_{k+1}$  is within a desired error margin (we used 0.001  $J_k$ ).

Let us consider the simple 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  $s(p, i) = (p/(\alpha \cdot p + \beta \cdot i)) = (p/i)/((\alpha \cdot p)/(i + \beta))$ , which is increasing with the signal-to-interference ratio  $p/i$ . Similar

functional forms are used in [15] for EDGE systems and in [25] for 802.11a Wireless LANs [7], [8]. Then, we get

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

where

$$\begin{aligned} X_r^k(n, i, b; r') &= \sum_{j \in \mathcal{I}} q_{ij} J^k(n, j, b-r; r) \\ &\quad - \sum_{j \in \mathcal{I}} q_{ij} J^k(n-m, j, b+m-r; r) \\ Y_r^k(n, i, b; r') &= C_s(r) + C_v(r; r') \\ &\quad + \sum_{j \in \mathcal{I}} q_{ij} J^k(n, j, b-r; r). \end{aligned} \quad (7)$$

$X_r^k(n, i, b; r')$  can be interpreted as the cost of sending  $m$  packets instantaneously when the current playout rate is  $r$ , which decreases the number of remaining packets to  $n-m$  and increases the client buffer level by  $m$  simultaneously. Furthermore,  $X$  and  $Y$  are functions of the cost-to-go functions  $J$  and can only be obtained through value iteration. Introducing  $X$  and  $Y$  simplifies the DP formulation. 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:  $p_r^k(n, i, b; r') = (1/\alpha) \{ \sqrt{\beta X_r^k(n, i, b; r') i} - \beta i \}$  for  $i < X_r^k(n, i, b; r')/\beta$  and 0 otherwise. At large  $k$ , there is convergence

$$\begin{aligned} p_r(n, i, b; r') &= \begin{cases} \frac{1}{\alpha} \left\{ \sqrt{\beta X_r(n, i, b; r') i} - \beta i \right\}, & i < \frac{X_r(n, i, b; r')}{\beta} \\ 0, & \text{otherwise} \end{cases} \end{aligned} \quad (8)$$

This power form is proved useful later in the design of heuristics. Substituting  $p_r^k$  into (6), we can now optimize over  $r$  and find  $r^*$ . Notice that this is not a closed form, because  $X_r$  in the right-hand side depends on the cost-to-go  $J_k$ , and we still need to use value iteration.

*Properties of Optimal Control:* It is interesting to explore the structural properties of the optimal control  $(p^*, r^*)$ . In Fig. 2, we plot the optimal power  $p^*$  versus the interference level  $i$  for fixed  $b$  and  $r$ , as given by (8) and assuming i.i.d interference. For a fixed playout rate, the transmitter chooses the transmission power  $p^*(i)$  according to one of the curves in Fig. 2. The shape of the  $p^*(i)$  curve, for a fixed  $r$ , is similar to what was observed in [1]: for low  $i$  values,  $p^*$  increases aggressively to overcome interference; for medium  $i$ ,  $p^*$  backs-off; for large  $i$ ,  $p^* = 0$  i.e., the transmitter stops completely. The additional control of playout  $r$ , provides a “cushion” to the transmitter, which can now rely on the playout buffer for overcoming short bad channel periods, without spending much power. In Fig. 2, this results in the  $p^*(i)$  curves moving lower and to the left as  $r$  decreases. The cooperation between the control variables leads to significant gains in terms of both performance and power, as we show in the results section. An illustrative example that demonstrates the interaction between  $p^*$  and  $r^*$  for a specific sample path of interference will be discussed in detail in a later section (Fig. 4).




Fig. 2. Optimal power  $p^*(i)$  versus interference level  $i$ , for fixed Rx buffer level  $b$  and for various playout rates  $r$ . (Color version available online at <http://ieeexplore.ieee.org>.)

#### IV. ANALYZING INSIGHTFUL SPECIAL CASES

In this section, we consider special cases of the general DP formulation, where only a subset of the control variables is used (e.g., power-only or playout-only) which correspond to practical scenarios and also provide insight into the structural properties of the joint control policy. For the rest of the paper, and among the variety of scenarios we can address in this framework, we choose to focus on the simple -yet interesting and natural- one that we solved in Section III-F (the power cost is simply  $\Phi(p, i) = p$  and there is a single mode with a  $m = R$ ; hence, we do not optimize over  $m$ ). For brevity, we also normalize the playout rate  $r$ , client buffer level  $b$  and size  $B$ , number of packet to be sent  $n$  and total number of packet  $N$  with respect to the media's natural rate  $R$ . Therefore,  $0 \leq r \leq 1$ , where  $r = 1$  corresponds to natural rate and  $r = 0$  corresponds to freezing.

##### A. Power-Controlled Streaming (Adapt $p$ , Fix $r = 1$ )

In this scenario, similarly to [16], [17] the power  $p$  is the only control parameter that can be varied to combat the bad channel periods;  $r$  and  $r'$  no longer belong to the control parameters and the system state, respectively. The Rx initially waits for  $T$  time slots time, and then starts playing always at natural rate 1, except during underflow that it freezes ( $r = 0$ ). The optimality equations then become:<sup>6</sup>

For  $b > 1, N \geq n > 0$

$$J(n, i, b) = \min_p \left( p + s(p, i) \sum_j q_{ij} J(n-1, j, b) + (1 - s(p, i)) \sum_j q_{ij} J(n, j, b-1) \right) \quad (9)$$

<sup>6</sup>The recursive dynamics of (9), (10), and (11) are similar to the general optimality equation. One subtle change concerns the additional cost  $C_v(0; 1)$  in (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  $b = 1$ .




Fig. 3. Backlog pressure  $X(n, i, b)$  versus receiver buffer level  $b$ , for different number of packets  $n$  at the Tx. (Color version available online at <http://ieeexplore.ieee.org>.)

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

$$J(n, i, 1) = \min_p \left( p + s(p, i) \sum_j q_{ij} J(n-1, j, 1) + (1 - s(p, i)) \left( C_v(0; 1) + \sum_j q_{ij} J(n, j, 0) \right) \right) \quad (10)$$

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

$$J(n, i, 0) = \min_p \left( p + C_s(0) + s(p, i) \sum_j q_{ij} J(n-1, j, 1) + (1 - s(p, i)) \sum_j q_{ij} J(n, j, 0) \right) \quad (11)$$

Using value iteration, we can obtain the optimal Tx power policy, which turns out a function only of  $X$  [defined similarly to (7)] and  $i$ :  $p(n, i, b) = (1/\alpha)(\sqrt{\beta X(n, i, b)}i - \beta i)$  if  $i < X(n, i, b)/\beta$  and 0 otherwise. However, notice that this is not a closed form, as  $X$  depends on  $J_k$ , and we still use value iteration to compute it. We refer the interested reader to [16], [17], and [19] for details on  $X$ . Here, we only briefly illustrate its structural properties, which are important for the design of the power heuristic in the next section. In Fig. 3, we plot  $X$  versus the Rx buffer level  $b$ , for various values of  $n = 1, \dots, 100$ . We can see that the backlog pressure  $X$  is: 1) decreasing in  $b$ ; 2) increasing in  $n$ ; and 3)  $X(n, i, b) \rightarrow X^*(i, b)$  as  $n \rightarrow \infty$ .

##### B. Adaptive Playout Control (Fix $p$ , Vary $r$ )

If the Tx side uses fixed power  $p$ , then the playout rate  $r$  at the Rx is the only control parameter to deal with bad channel conditions. Intuitively, when the channel is good, the optimal policy tends to choose the natural playout rate ( $r = 1$ ); when the channel is bad, the optimal playout policy should slow down, and at the extreme freeze and re-buffer ( $r < 1$ ) to prevent buffer

underflows ( $r < 1$ ). The optimality equations now become as follows (notice that  $p$  is constant, therefore  $s(p, i) = s(i)$ ):

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

$$J(n, i, b, r') = \min_r (C_s(r) + C_v(r; r') + s(i) \sum_j q_{ij} J(n-1, j, b-r+1, r) + (1-s(i)) \sum_j q_{ij} J(n, j, b-r, r)) \quad (12)$$

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

$$J(n, i, b, r') = \left( C_s(0) + C_v(0; r') + s(i) \sum_j q_{ij} J(n-1, j, 1, 0) + (1-s(p, i)) \sum_j q_{ij} J(n, j, 0, 0) \right). \quad (13)$$

The boundary condition for  $n = 0$  is the same. There is no closed-form solution, and we use again the value iteration method to obtain the optimal policy.

*Re-Buffering*: An interesting special case is re-buffering, in which  $r$  can take only one of two values: either  $r = 1$  (natural rate) or  $r = 0$  (freeze). The optimality equations are the same as in (12) and (13). We now list the structural properties of the optimal re-buffering policy, together with their intuitive explanations.

- 1) The optimal rate  $r^*$  is nondecreasing in  $b$ : Intuitively, the optimal re-buffering should playout packets when  $b$  is high and rebuffered when it is low. This property allows us to represent the optimal re-buffering control  $r^*(n, i, b; r')$  using a threshold  $\theta(n, i, r')$ :  $r^*(n, i, b; r') = 0$  if  $b < \theta(n, i, r')$  and  $r^*(n, i, b; r') = 1$  if  $b \geq \theta(n, i, r')$ .
- 2)  $\theta(n, i, r')$  is nondecreasing in  $n$ : Indeed, the more packets need to be transmitted, the more likely a failed transmission will cause buffer underflow; thus, the more re-buffering is needed.
- 3)  $\theta(n, i, r')$  is nondecreasing in  $i$ : For worse channel condition, we need more re-buffering.
- 4)  $\theta(n, i, r' = 1) \leq \theta(n, i, r' = 0)$ : This is essentially a hysteresis effect in  $r(n, i, b; r')$ : to avoid excessive jitter, the playout control tends to rebuffered more packets if it is already in the re-buffering mode ( $r' = 0$ ) and vice versa.

Inspired by these properties, in the next section we define a practical threshold-based heuristic.

### C. Performance Comparison of the Optimal Policies

In this section, we use simulation to compare the optimal policies to each other and to the no-control case.

1) *Policies Under Comparison*: We now compare the performance of the optimal policies summarized in Table I. The name  $T_x R_y$  describes the control available at the Transmitter ( $T$ ) and at the Receiver  $R$ . The subscript indicates the kind of control: absence of subscript means *no control*; *a* means *adaptive*, using the full range of the control variable; *rb* in particular

TABLE I  
SUMMARY OF POLICIES UNDER COMPARISON

Control	Name	Control at Tx	Control at Rx
No Control	$TR$	none ( $p = const$ )	none ( $r = 1$ )
Power Only	$T_a R$	adaptive ( $0 \leq p \leq P_{max}$ )	none ( $r = 1$ )
Power and Re-buffering	$T_a R_{rb}$	adaptive ( $0 \leq p \leq P_{max}$ )	re-buffering ( $r \in \{0, 1\}$ )
Playout Only	$TR_a$	none ( $p = const$ )	adaptive ( $0 \leq r \leq 1$ )
Joint Control	$T_a R_a$	adaptive ( $0 \leq p \leq P_{max}$ )	adaptive ( $0 \leq r \leq 1$ )

means *re-buffering*. In summary, we compare the following optimal control policies.

- *No control exercised ( $TR$ ) at neither side*. This will be our benchmark for comparison.
- *Power Control Only ( $T_a R$ )*, as in Section IV-A. This corresponds to the case where the receiver is not able or is not willing to do playout control; then power is the only control the system designer can use to combat channel fluctuations.
- *Power Control with Re-buffering ( $T_a R_{rb}$ )*. At the Tx side, the power is optimally controlled as in  $T_a R$ . At the Rx side, there is limited adaptivity: the Rx can either play at full speed ( $r = 1$ ) or freeze to re-buffer ( $r = 0$ ). Re-buffering is the commonly used playout policy today in practical systems.
- *Playout Control Only ( $TR_a$ )*, as in Section IV-B. Tx power is constant, but the playout is optimally controlled within the set of available rates ( $0 \leq r \leq 1$ ). This may be the case when power control is not enabled by the wireless technology in use, or when transmission power is not an issue (the transmitters are powered and do not need to save on battery and/or interference is low and does not limit network capacity).
- *Joint Power-Playout Control ( $T_a R_a$ )*, as in Section III. In this case, the system designer has control over the operation of both the Tx and the Rx. By appropriately choosing the weight  $W$ , i.e., the relative importance to power versus playout quality, he/she can indicate a preference for saving power or improving viewing quality. The optimal policy will then exercise the appropriate mix of power control versus playout control to reach an operating point that satisfies this preference.

2) *Performance Metrics*: We are interested in two performance metrics. The first is the degradation in media (playout) quality, due to variation of the playout rate  $r$ . 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 a 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  $w$ :  $C_{total} = P_{avg} + w \cdot C_{avg}$ . By varying  $w$ , we can choose to give more importance to media quality or to power, which will result to different optimal control policies.

TABLE II  
EXAMPLE (OF 2-STATE MARKOV) CHANNELS USED IN SIMULATION

Example Channel	Good State		Bad State	
	$i$ (units)	mean duration (slots)	$i$ (units)	mean duration (slots)
1	1	7	100	3.5
2	1	7	10	3.5
3	1	7	100	14

Therefore, the performance of a control policy is captured by the quality-power tradeoff.

3) *Interference Model:* We use a two-state Markov model for the channel interference  $i$ . As our main setup, we use  $i = 1$  for the good state and  $i = 100$  for the bad state; the average durations of the good and the bad state are 7 and 3.5 time slots, respectively. We chose to discuss this example, among many simulation results, because it demonstrates key-points of our approach; e.g., there is a significant difference between the values of interference in the good and the bad state which makes power-control necessary. The control policies make a difference when the system is under stress; on the contrary, if power is not an issue, the Tx should transmit at as high power as required to deliver the packets. In addition to the above channel, which we call Example Channel 1, we also simulated two different two-state Markov channels, summarized in Table II.

The values of power shown in all figures are relative to the interference level. Thus, no units are mentioned for  $i$  or  $p$ , as we are mainly interested in the ratio  $p/i$ .

4) *Other Simulation Parameters:* 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 play-out rates  $r \in \{0, 0.5, 0.6, 0.7, 0.8, 0.9, 1\}$

5) *Simulation Results:* In Fig. 4, we discuss a scenario that demonstrates the intuitive behavior of the joint control. The top plot shows a sample-path of interference  $i$ . The second plot shows the resulting Rx buffer level  $b$ . The joint optimal control (power in the third and play-out rate in the fourth plot) reacts to the variations of the channel to provide a smooth play-out. One can make the following observations by looking at Fig. 4. Initially, the Rx buffers some packets before start playing at full speed ( $r/R = 1$ ). During the first bad channel period, the Tx power backs-off since the buffer level is high due to the initial slow play-out. During the longer second bad channel period, the Tx power initially backs-off; however as the Rx buffer level decreases, the Tx power increases to help; the play-out also slows down to prevent the Rx buffer from underflow.

Fig. 4 showed the optimal joint control for a single sample of  $i$ , and for a single choice of  $W$ , i.e., the relative value of media quality versus power spent. The next step is to simulate a large number (we chose 1000) of realizations of  $i$  according to the statistics of Example Channel 1, and average the results ( $P_{avg}$ ,  $C_{avg}$  for each policy) over all realizations. The results are shown in Fig. 5. Each curve corresponds one of the policies under comparison:  $TR$ ,  $T_aR$ ,  $T_aR_{rb}$ ,  $TR_a$ ,  $T_aR_a$ . For policies  $T_aR$ ,  $T_aR_{rb}$ ,  $T_aR_a$ , the weight  $W$  is varied in order to obtain the entire curve, i.e., the quality versus power tradeoff. The more weight on media quality, the more willing we are to spend

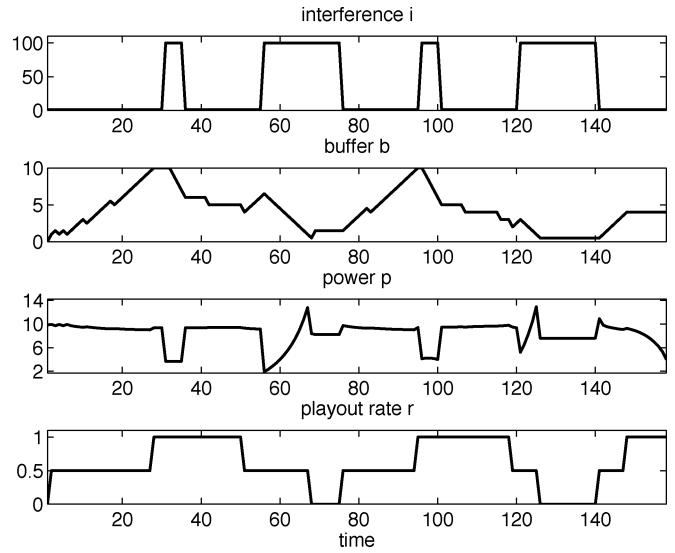


Fig. 4. Top plot shows a sample path of interference  $i$ :  $i$  oscillates between low ( $i = 1$ ) and high ( $i = 100$ ) levels. The second plot shows the resulting Rx buffer level. The optimal joint control policy adjusts the Tx power (third plot) and the play-out rate (4th plot) to overcome the variations of  $i$  and provide a smooth play-out.




Fig. 5. Quality versus power tradeoff for optimal policies, in the range where power matters (i.e., the power cost is weighted more than the degradation in media quality). (Color version available online at <http://ieeexplore.ieee.org>.)

power, the further to the right of the curves we operate. For policies  $TR$  and  $TR_a$ , the constant transmission power is varied to obtain the quality-power curve.

Fig. 5 compares the policies in various operational ranges (i.e., various relative weights on the power cost); in particular we varied the weight  $W$  from 0.1 to 10. In the regime where power matters more (left part of the figure), the power controlled algorithms,  $T_aR$  and  $T_aR_{rb}$ , outperform the adaptive play-out algorithm  $TR_a$  significantly, by roughly 50% in terms of media quality. As less and less weight is given to the power cost, we move to the right of the figure, and the adaptive play-out algorithm  $TR_a$  starts to outperform the power controlled algorithms  $T_aR$ ,  $T_aR_{rb}$ . Comparing  $T_aR$  to  $T_aR_{rb}$ , we observe that allowing the play-out controller to adaptively re-buffer ( $T_aR_{rb}$ ) improves the performance substantially. The adaptive play-out



policy  $TR_a$  outperforms policies  $T_aR$  to  $T_aR_{rb}$ : for the same average power, it reduces the quality degradation by 50%. The joint optimal control  $T_aR_a$  further reduces the quality degradation by another 30% more than the adaptive playout alone, which demonstrates the merit of joint power-playout control.


In both the low- and high-power regimes, the absence of any control leads to the highest degradation and power consumption, thus performs the worst. Allowing power and/or playout control improves the performance. As expected, joint control outperforms each individual control and achieves the lowest degradation-power tradeoff; e.g., for the same average power spent, the joint optimal policy reduces the quality degradation by 20% to 75% compared to all other policies.

In addition to the results obtained above for Channel 1, we also varied the parameters of the 2-state Markov model and looked at how the benefit from the joint control varies. First, we considered Example Channel 2 with smaller difference between the interference values in the good and bad state than in Example Channel 1 ( $i_{bad} : i_{good} = 1 : 10$  instead of 1:100). Then, we considered Example Channel 3 with the same interference levels as in Channel 1, but with longer duration of bad periods (twice the average duration of Channel 1). Table II summarizes the characteristics of all three channels. Simulation results showed at least 50% benefit when using joint control, in terms of both power and media quality, and for all three channels.

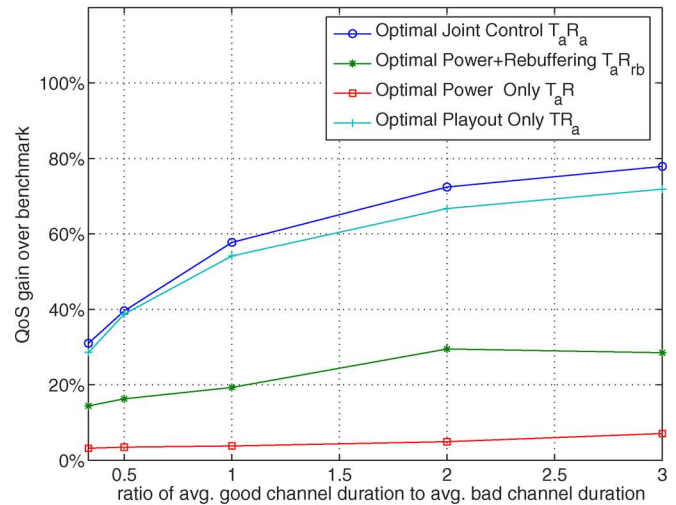
Generalizing the properties of Channels 2 and 3, we looked at a wide range of channels by varying the difference in interference levels and the ratio between the average duration, in the good and bad states. Intuitively: 1) the higher the interference in the bad state (compared to the interference in the good state) and/or 2) the larger the average duration of the bad state (compared to the average duration of the good state), the worse the channel, and the worse the performance gains of all control policies compared with benchmark. For a fair comparison, we consider the same average power consumption (25 power units, which is in the high-power regime) and we look at the improvement in playout quality achieved by the control policies. This improvement is calculated as  $(C_{benchmark} - C_{policy})/C_{benchmark}$ , where  $C_{policy}$  is the playout quality achieved by this policy.

In Fig. 6(a), we varied the interference value in the bad state ( $i_{bad}$ ) from 10 to 100, while keeping the level in the good state the same ( $i_{good} = 1$ ). As expected, the performance gain decreases as the gap between  $i_{bad}$  and  $i_{good}$  increases. Interestingly, this gain decays slowly. All the power control policies ( $T_aR_a$ ,  $T_aR_{rb}$ ,  $T_aR$ ) attempt to use high power to combat the high interference. The higher the interference is, the less efficient the power control policies. This explains the losses in performance gains. The gain from the playout also decreases as the channel gets worse. In Fig. 6(b), we vary the ratio between the good and the bad duration from 1/3 to 3, while keeping their sum constant to 10; this can also be interpreted as decreasing the frequency of bad state occurrences. As expected, the performance gains of all policies increase as the ratio increases.

Over all the channels considered, the joint control brings the largest performance benefit over the no-control case. In the high-power regime considered in Fig. 6, the playout control also performs close to the joint control.



(a)



(b)

Fig. 6. Quality gains for all policies, considering various channels. (a) Increasing the gap in interference levels between bad and good state ( $i_{bad} : i_{good}$ ) decreases quality. (b) Increasing the ratio of average duration in good versus bad state ( $avg\ good : avg\ bad$ ) improves quality. (Color version available online at <http://ieeexplore.ieee.org>.)

## V. HEURISTICS: DESIGN AND PERFORMANCE

The optimal policies studied in the previous section achieve the performance limits of streaming over an interference-limited wireless link. However, they require computational power and knowledge of the channel statistics, which may be difficult or impossible to obtain. In this section, we design practical, low-complexity heuristics that achieve near-optimal performance without the above limitations. To achieve these desired properties (i.e., both good performance and low complexity), we design our heuristics to mimic the structural properties of the optimal control. Thus, we name the heuristics after the optimal policies that inspires them; e.g.,  $T_xR_yH$  is the heuristic based on the optimal policy  $T_xR_y$ .

### A. Design

Let us now describe each heuristic in detail.

1) *Power Control Only Heuristic— $T_aRH$* : One can construct a power control heuristic by imitating the structural prop-

erties observed in IV-A; in particular, we analytically approximate  $X(n, i, b)$ . The interested reader is referred to [16], [19] for details.

2) *Fixed Threshold Playout Heuristic*: A playout-only heuristic widely used in practice is to choose the playout rate looking only at the playout buffer level; e.g., a single threshold may be used, as in [12] 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  $\{r_0 = 0, r_1, r_2, \dots, r_L\}$

$$r = r_l, \text{ when } \theta_l < b \leq \theta_{l+1} \quad (14)$$

where  $\theta_0 = -1$  (just a notation to include the case  $b = 0$  when  $\theta_1 \geq b \geq \theta_0$ ) and  $\theta_k = B$ . For our simulations, we first used six equally spaced buffer thresholds  $\theta_1, \theta_2 \dots \theta_6$  and seven 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.

3) *Adaptive Threshold Playout Heuristic,  $TR_aH$* : In Section IV-B, we observed that when  $p$  is constant, the optimal playout policy  $TR_a$  uses several thresholds to divide the playout buffer occupancy into intervals, each corresponding to a different playout rate. Intuitively, these thresholds should vary according to the quality of the channel. e.g., when the channel is bad, the playout should be more conservative, i.e., increase the thresholds and slow down even if there are some packets in the buffer, to prevent underflow. Conversely, if the channel is good, then the playout policy can afford to be more optimistic, i.e., decrease the thresholds and avoid slowing down, unless the buffer is almost empty. We obtained this intuition by looking at the structural properties of the optimal control. We now design an adaptive threshold policy based on the above intuition. Our heuristic observes the channel condition (captured by the probability of successful transmission  $s = s(p, i)$ ). Instead of having a fixed threshold  $\theta_l$ , we define  $\theta_l(s)$  as a function of  $s$

$$\theta_l(s) = \theta_l^* - \frac{L - l + 1}{L} \delta(s) \quad (15)$$

where  $\theta_l^*$  are the thresholds that divide buffer into uniform  $L+1$  levels. Equation (15) describes the following simple operation. As  $s$  increases from 0, threshold  $\theta_l$  is reduced by  $((L+1-l)/L)\delta(s)$ . For example, with  $L = 6$  thresholds,  $\theta_1$  is reduced by  $\delta(s)$ ,  $\theta_2$  is reduced by  $(5/6)\delta(s)$  and so on. This allows the Rx to play the received packets at a faster rate. We pick  $\delta(s) = (B/(L+1))/(s/(s+\gamma))$  to be the initial width  $B/(L+1)$  multiplying a function of  $s$ .  $\gamma$  is a tuning parameter to adjust the sensitivity of these thresholds to variations in the channel conditions. Given these thresholds, we can use (14) to find the corresponding playout rate.

In simulation, we found that the fixed and adaptive threshold heuristics performed very close to the optimal playout policy. Furthermore, the adaptive threshold heuristic reduced the performance gap between the optimal policy and the fixed threshold heuristic by 50%, thanks to its ability to adapt the thresholds to the channel conditions. We further investigated the sensitivity of

our heuristic ( $TR_aH$ ) to the choice of: 1) playout rate range (excluding rate 0) and 2) playout rate resolution. We found that our scheme is robust to these choices. The performance degrades only slightly as the range decreases and the resolution becomes coarser. We again have to omit the figures for lack of space.

4) *Joint Power-Playout Heuristic,  $T_aR_aH$* : In Section IV-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 III-F. We further notice that given  $X$ , the optimal power can be calculated as

$$p = \begin{cases} \frac{1}{\alpha}(\sqrt{\beta X i} - \beta i), & i < \frac{X}{\beta} \\ 0, & \text{else} \end{cases} \quad (16)$$

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 properties of the optimal power control, observed in Section IV-A. Specifically

$$\begin{aligned} X(n, b) &= K \exp\left((B+1-b)^{(m)^\theta}\right) + \epsilon; \\ K &= C \exp\left(-B^{(m)^\theta}\right) \end{aligned} \quad (17)$$

where the parameters  $\eta$  and  $\theta$  adjust the sensitivity of  $X$  to the backlog pressure  $n$ .  $\epsilon$  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 (16) and (17), 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 to choose the playout rate based on the calculated transmission power  $p$ .

## B. Performance Evaluation Under Markovian Interference

For the performance evaluation of these heuristics, we use the simulation setup of Section IV-C; in particular, we show results obtained for Example Channel 1. Fig. 7 shows the performance of the joint power-playout heuristic under tight power constraints. The simulation setup is the same as in Fig. 5 for the optimal policies which are re-plotted here for comparison. We observe that the proposed joint control heuristic  $T_aR_aH$  performs quite close to the joint optimal  $T_aR_a$ . Furthermore, it outperforms substantially the individual policies,  $T_aR$  and  $TR_a$ , even in the regimes that these policies are designed for (i.e., low-power regime for  $T_aR$ , high-power regime for  $TR_a$ ). Note that the optimal policies require the solution of the optimality equation, and hence are computationally intensive, while the heuristic has minimal computational complexity. Furthermore, the optimal policies require knowledge of the channel statistics (basically the transition matrix  $Q$ ), while the joint-control

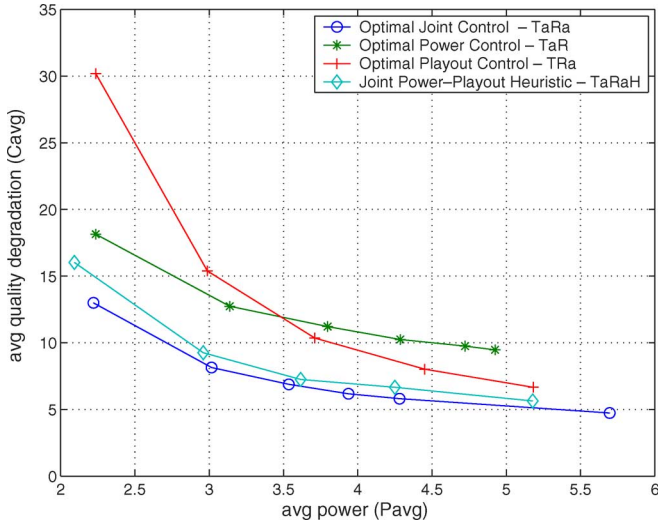


Fig. 7. Comparing the joint control heuristic  $T_a R_a H$  to the optimal policies, in low-moderate power regimes. (Color version available online at <http://ieeexplore.ieee.org>.)

TABLE III  
PERCENT (%) IMPROVEMENT IN PLOUT QUALITY OVER BASELINE,  
FOR THE SAME POWER CONSUMPTION

Optimal Policies and Heuristic Algorithms	Policy Abbr.	Low power range	High power range
Optimal Power (only) Control	$T_a R$	50%	31%
Optimal Power Control with Re-buffering	$T_a R_r b$	57%	40%
Optimal Playout (only) Control	$T R_a$	29%	70%
<b>Optimal Joint Power-Playout Control</b>	$T_a R_a$	<b>69%</b>	<b>83%</b>
<b>Joint Power-Playout Heuristic</b>	$T_a R_a H$	<b>64%</b>	<b>78%</b>

heuristic requires only a good estimate of the current channel interference.

We also performed simulations in the high-power regime (where power is not an issue, and high power is used to achieve high media quality). We again found that the proposed joint control heuristic outperforms all other heuristics, and performs close to the optimal joint policy. In Table III, we summarize the gains achieved over the no-control baseline  $TR$ , by the optimal policies and heuristic, for the operational regimes of low power consumption (2.2 units per packet) and high power consumption (60 units per packet). Clearly, our proposed joint heuristic  $T_a R_a H$  can achieve near-optimal playout quality over all power regimes, at significantly lower complexity.

### C. Performance Under Power-Responsive Interference

So far, we assumed Markovian, nonresponsive interference. The next important methodological step was to simulate our algorithms in a realistic wireless environment, where all links perform power control, resulting in 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. Interference can no longer be modeled as a Markov chain, and the optimality equations can no longer be formulated. We

simulated our heuristics in this environment and we found that they still perform very well. The details of the simulation setup are as follows.

- We used the OMNeT++ simulator [22] and simulated our single media-streaming link in a pool of a large number of data packet links. The latter are power-controlled, so when the primary media-streaming link raises its power, they will also raise their powers and induce more interference on it. All background data transmitters are assumed to use the power controlled multiple access (PCMA) [1].
- The simulation considers a network space of 500 by 500 space units, wrapped around in a torus. Data or media transmission sessions are established between pairs of transmitters and receivers. The locations of transmitters and receivers are chosen at random. We assume a free-space path loss model where the power gain  $G = (1/d^4)$  between any transmitter and receivers where  $d$  is the distance between the pair. Thermal noise is assumed  $10^{-12}$ . All transmissions operate at the same channel.
- The media streaming session, between the same pair of transmitter and receiver, lasts for the entire simulation time. It consists of shorter sessions of  $N = 100$  packets each, initiated continuously one after the other. The buffer size is  $B = 10$  and the initial buffering is  $T = 5$  time slots.
- Data transmission sessions are randomly generated throughout the simulation between randomly chosen pairs of nodes. Packets arrive at a transmitter following a Bernoulli distribution. The lifetime of each session is geometrically distributed.
- Since the statistics of the interference are unknown, we use the interference observed in the previous time slot as the estimate of the interference in the current time slot.

We simulate three heuristics: power control only heuristic  $T_a R H$ , adaptive threshold playout  $T R_a H$  and joint power-playout heuristic  $T_a R_a H$ , and compare their performance to the benchmark no-control case. Fig. 8 shows the results of the simulation for an illustrative example. The figure plots the performance metrics of interest over simulation time, i.e.: 1) the average quality-of-service (QoS) (in terms of playout quality) and 2) the average power spent. We are mainly interested in the steady-state performance, reached at the right of the plots. We observe that all three heuristics outperform the benchmark, achieving a lower quality degradation and power consumption. The power control heuristic consumes less power compared to the adaptive threshold playout control, but results in a higher quality degradation. The joint power-playout control outperforms all other schemes (the dotted line corresponding to the joint heuristic is lower than all other lines in both plots). Indeed, it reduces the power consumption by more than 60% and reduces the quality degradation by 66% compared to the benchmark.

## VI. PRACTICAL ISSUES AND EXTENSIONS

In this work, we formulated the problem of media streaming over a fluctuating channel as a stochastic control problem. So far we focused on analyzing the power-playout tradeoff. In this section, we discuss how the proposed framework can be applied/extended in real-world multimedia data types and wireless transmission systems. However, we would like to stress that this brief

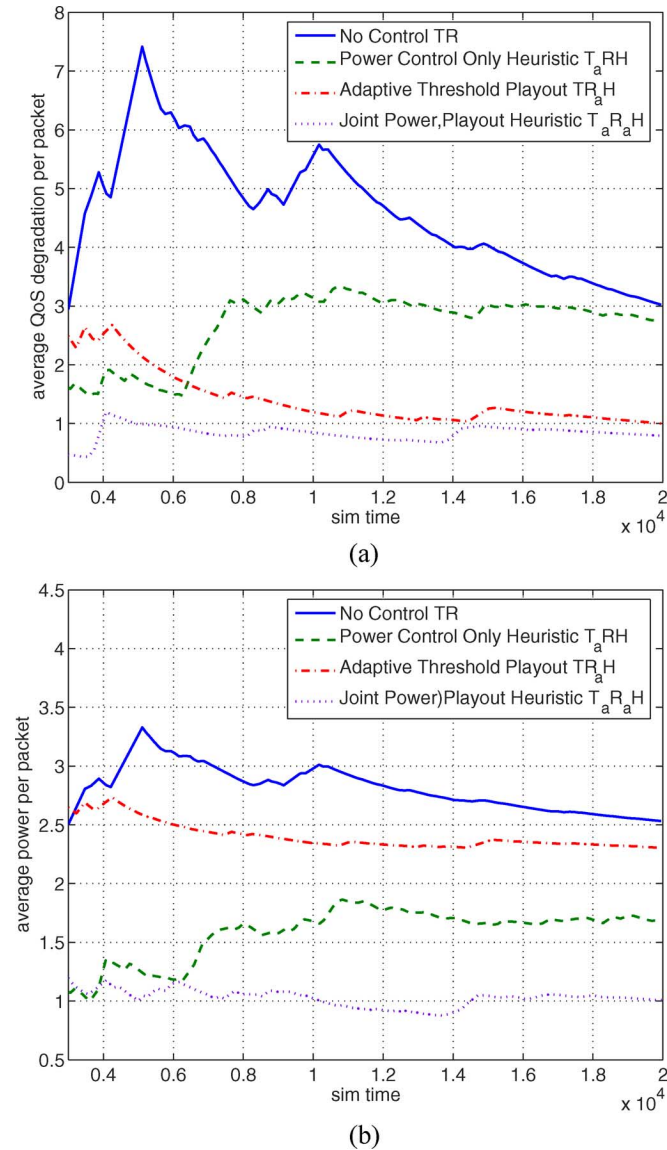


Fig. 8. Simulation for responsive interference: power and quality costs achieved by various policies versus simulation time. (a) Average playout quality degradation. (b) Average transmission power. (Color version available online at <http://ieeexplore.ieee.org>.)

discussion only provides high-level guidance, and does not consider the complex tradeoffs that must be considered in each area when designing a practical, high-performance system.

#### A. Adapting Video, Speech and Audio Playout Rates

One of the control variable we used in this work was the video playout rate. Let us now discuss some practical issues related to the playout rate adaptation of different media types.

*Video playout* rate can be adapted in a straightforward manner by adapting the frame rate of the displayed video. The key point, as covered by the cost metrics, is that the playout variation should be smooth so that a smooth movement in the video scene is still mapped after playout rate variation to a smooth movement. For example, if a person in the video runs from left to right across the display or moves a painting across the display, the movement should be a temporally scaled

version of the original movement—without any confusing warping or abrupt stops/starts.

*Speech playout* rate adaptation should be performed in a manner to preserve the pitch of the speaker; if this is not done, then the speaker may acquire a chipmunk voice (for speedup) or an unnaturally deep voice (for slowdown). This playout rate adaptation can be achieved by identifying the pitch in the speech and extending or reducing the speech length by adding or deleting an integer number of pitch periods, e.g., see the WSOLA algorithm in [20], [27]. Speech is generally decomposed in time into three types of segments: 1) voiced segments, where the speech has a periodic pattern driven by the pitch; 2) unvoiced segments, where the speech is modeled as filtered white noise; and 3) silence periods, which occupy  $1/2$ – $2/3$  of the time. Voiced segments are adapted as mentioned above to preserve the pitch, while unvoiced segments can be adapted much more flexibly, and silence segments can also be shortened or expanded.

*Audio playout* rate adaptation is a complicated topic which depends on the specific audio content. We distinguish audio, which refers to a generic audio signal, from speech, which is an audio waveform produced by a human speaker. If the audio consists of speech, then speech playout rate adaptation methods can be applied as discussed above. If the audio consists of music then the situation is much more complex; for example there may be sharp transients such as drum beats which should not be adapted, and periodic signals such as from a violin which may or may not be adapted without distorting the perceived quality. Furthermore, audio may consist of a wide range of events, from explosions to the sound of raindrops. Audio playout rate adaptation is an active area of research.

#### B. Media-Specific Extensions

1) *Adapting a Compressed and Packetized Media Stream:* In this work, we focused on reliable delivery of video packets. However, in many practical scenarios, it is desirable to adapt a media stream to meet constraints on delay (e.g., for low-delay applications) and/or transmission bandwidth (e.g., in wireless environments with limited bandwidth resources); in those cases reliable delivery is not necessary and media packets may be dropped. Such scenarios can be addressed in our framework by using additional controls (e.g., packet scheduling) and associated costs (e.g., distortion values due to missing media units).

For example, one well-known technique to adapt the transmission of a packetized stream is by intelligent discarding of packets, e.g., discard B frames if they exist, alternatively discard P and I frames using a rate-distortion framework, see, e.g., [3], [6].<sup>7</sup> We are currently working on formulating the problem of packet scheduling for low-delay streaming within our framework [18]. Another popular technique to adapt a compressed media stream is transcoding, which can reduce the required bit rate by, for example, coding with a coarser quantizer or reducing the frame rate. For an overview of video transcoding techniques, see [4].

<sup>7</sup>This requires that the video be coded and packetized using application level framing (ALF) principles so that if a packet is dropped the remaining packets can still be independently decodable. For example, video coded with the H.263, MPEG-4, or H.264 video compression standards may be packetized and sent over RTP/UDP/IP, where the RTP payload format (e.g., RFC 2190 or RFC 2429 for H.263, RFC 3016 for MPEG-4, and RFC 3984 for H.264 [11]) specifies recommended modes of application level framing.

2) *Media Quality Metrics*: Mapping network-level metrics to appropriate subjective-quality costs is an interesting open research problem. In this work, we considered some simple, widely used functions to capture the cost of playout rate variation. In our recent work [18], we also defined metrics that capture the motion intensity of different scenes, in order to perform content-aware playout control. In general, our framework is able to incorporate any new and improved quality metrics.

3) *Multiple Packets per Frame*: Typically, video frames are of different sizes and are also split into multiple packets each. It is straightforward to extend our model to include this case [18], by defining an additional state variable to capture the number of remaining packets of the same frame, remaining in the transmitter's buffer.

### C. Wireless Technology-Specific Issues

1) *Power (and QoS) Control depending on the Wireless Environment*: Transmission power control (TPC) has been used in cellular networks. In WCDMA networks, power control consists of two loops. The outer loop power control is used to set a target signal-to-noise and/or interference-ratio (which we refer to as "SINR", see Section III-A for a discussion) to maintain a desired packet loss rate. The outer loop power control is updated every 20 ms. The inner loop power control is updated at a much faster rate (of 1500 Hz) to maintain the instantaneous SINR at its target value. In the downlink direction (from the base station to cellular phones), a lower SINR target value can increase the network capacity at the expense of higher packet error rate. The joint control scheme proposed in this paper can be applied to the outer loop power control to set the optimal SINR target value, which improves the network capacity while maintaining the desired video viewing quality.

TPC has also been proposed for wireless LANs (standardized by [7] in IEEE 802.11 h [10]) to reduce the co-channel interference. Due to the small number of non overlapping channels (three channels in the 2.4-GHz band and 8–12 channels in the 5.8-GHz band), in a dense-deployed environment, such as apartment buildings, co-channel interference is of a great concern. As in cellular networks, TPC can reduce interference and increase network capacity. Besides TPC, the MAC protocol of wireless LANs has also been enhanced with *Quality-of-Service* capabilities (see IEEE 802.11e [9]), which can further aid multimedia streaming services. Particularly the HCF function can allocate fixed transmission opportunities/rate to each multimedia stream, which is an improvement over random access. Under the HCF function, our proposed schemes can also be used to reduce the channel interference via TPC and improve the view quality via playout adaptation.

2) *The effect of Coding and Modulation*: In some cases, it may be possible to control not only the transmission power but also the transmission mode, e.g., by changing modulation and coding scheme. Then, we can control the number of packets sent in one time slot and the probability of their successful transmission. The modes available are ultimately system/technology specific; e.g., IEEE 802.11 h [10] defines the power and rate adaptation capabilities for wireless LANs. We have already included this control in the general model of Section III, although we have not provided analytical and simulation results in this paper.

Depending on the type of wireless network, as well as on the coding and modulation used, the probability of successful transmission as a function of SINR,  $s(p, i)$ , takes a specific functional form. This affects the optimal power control policy but not its structural properties. The power control heuristic may need tuning to maximize performance gains. The design of playout control policies is less affected, since it is based only on the packet error rate and not on the underlying SINR.

3) *Channel Model Extensions*: One could naturally extend the model used in this paper, to express additional properties of the channel, other than the probability of successful transmission in each state. First, we could model the delay of the forward and/or the feedback channel. Second, we could consider multiple channels, and add an additional control variable to choose over which channel(s) to transmit.

The aforementioned are some of the streaming scenarios that can be addressed within our framework.

## VII. CONCLUSION

In this paper, we examined the problem of joint power and playout control for media streaming over an interference-limited wireless link. We formulated the problem in a dynamic programming framework; we studied the optimal control for joint power-playout, as well as for the special cases of power-only and playout-only control. Furthermore, we designed and evaluated practical, low-complexity heuristics that achieve near-optimal performance. Our framework can be naturally extended to address a large range of scenarios for media streaming over wireless.

## REFERENCES

- [1] N. Bambos and S. Kandukuri, "Power Controlled Multiple Access (PCMA) in wireless communication networks," in *Proc. of IEEE INFOCOM 2000*, New York, June 2000, pp. 368–395.
- [2] D. Bertsekas, *Dynamic Programming and Optimal Control*. Belmont, MA: Athena Scientific, 1995, vol. 2.
- [3] P. Chou and Z. Miao, "Rate-distortion optimized streaming of packetized media," *IEEE Trans. Multimedia*, vol. 8, no. 2, pp. 390–404, Apr. 2006.
- [4] A. Vetro, C. Christopoulos, and H. Sun, "Video transcoding architectures and techniques: an overview," *IEEE Signal Process. Mag.*, vol. 20, no. 2, pp. 18–29, Mar. 2003.
- [5] N. Färber and B. Girod, "Wireless video," in *Compressed Video Over Networks*, A. Reibman and M.-T. Sun, Eds. New York: Marcel Dekker, 2000.
- [6] B. Girod, J. Chakareski, M. Kalman, Y. J. Liang, E. Setton, and R. Zhang, "Advances in network-adaptive video streaming," in *Proc. IWDC 2002*, Capri, Italy, Sept. 2002, pp. 1–8.
- [7] *IEEE 802.11 Working Group for Wireless Local Area Networks (WLAN) Standards*, [Online]. Available: <http://www.grouper.ieee.org/groups/802/11/>
- [8] *IEEE Standard for Information Technology—LAN/MAN—Specific Requirements—Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, IEEE Std. 802.11, 1999/8802-11 (ISO/IEC 8802-11:1999), 1999.
- [9] *Draft Amendment to Standard for Information Technology—Telecommunications and Information Exchange Between Systems—LAN/MAN Specific Requirements—Part 11 Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Amendment 7: Medium Access Control (MAC) Quality of Service (QoS) Enhancements*, IEEE P802.11E (D13), 2005.
- [10] *IEEE Standard for IT—Telecommunications and Information Exchange Between Systems—LAN/MAN—Specific Requirements Part 11: Wireless MAC and PHY Specifications—Spectrum and Transmit Power Management Extensions in the 5 GHz Band in Europe*, IEEE Std. 802.11H-2003, 2003.

- [11] Internet RFC (Request for Comments) Index [Online]. Available: <http://www.faqs.org/rfcs/>
- [12] M. Kalman, E. Steinbach, and B. Girod, "Adaptive media playout for low delay video streaming over error-prone channels," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 14, no. 6, pp. 841–851, Jun. 2004.
- [13] —, "Rate-distortion optimized video streaming with adaptive playout," in *Proc. IEEE ICIP 2002*, Rochester, NY, Sept. 2002, vol. 3, pp. 189–192.
- [14] S. Kandukuri and N. Bambos, "Multimodal Dynamic Multiple Access in Wireless Packet Networks," in *Proc. of IEEE INFOCOM 2001*, Anchorage, AK, Apr. 2001, pp. 199–208.
- [15] D. Krishnaswamy, "Network-assisted link adaptation with power control and channel re-assignment in wireless networks," in *Proc. 3G Wireless Conference 2002*, San Francisco, CA, Sept. 2002, pp. 165–170.
- [16] Y. Li and N. Bambos, "Power-controlled wireless links for media streaming applications," in *Proc. of Wireless Telecommunications Symp.*, Pomona, CA, May 2004, pp. 102–111.
- [17] —, "Power-controlled streaming in interference-limited wireless networks," in *Proc. IEEE Broadband Networks*, San Jose, CA, Oct. 2004, pp. 560–568.
- [18] Y. Li, A. Markopoulou, J. Apostolopoulos, and N. Bambos, "Packet transmission and content-dependent playout variation for video streaming over wireless networks," in *Proc. IEEE MMSP 2005 (Special Session on Content Aware Video Coding and Transmission)*, Shanghai, China, Oct. 2005.
- [19] Y. Li, A. Markopoulou, N. Bambos, and J. Apostolopoulos, "Joint power-playout control schemes for media streaming over wireless links," in *Proc. IEEE Packet Video Workshop 2004*, Irvine, CA, Dec. 2004.
- [20] Y. Liang, N. Färber, and B. Girod, "Adaptive playout scheduling and loss concealment for voice communication over IP networks," *IEEE Trans. Multimedia*, vol. 5, no. 4, pp. 532–543, Dec. 2003.
- [21] S. Moon, J. Kurose, and D. Towsley, "Packet audio playout delay adjustment: performance bounds and algorithms," *ACM/Springer Multimedia Syst.*, vol. 6, pp. 17–28, Jan. 1998.
- [22] OMNeT++ Discrete Event Simulation System [Online]. Available: <http://www.omnetpp.org/>
- [23] S. Shankar, Z. Hu, and M. Van der Schaar, "Cross layer optimized transmission of wavelet video over IEEE 802.11a/e WLANs," in *Proc. IEEE Packet Video 2004*, Irvine, CA, Dec. 2004.
- [24] J. Cabrera, A. Ortega, and J. I. Ronda, "Stochastic rate-control of video coders for wireless channels," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 12, no. 6, pp. 496–510, Jun. 2002.
- [25] K. B. Song and A. Mujtaba, "On the code-diversity performance of bit-interleaved coded OFDM in frequency-selective fading channels," in *Proc. IEEE VTC*, Orlando, FL, Sep. 2003, vol. 1, pp. 572–576.
- [26] D. Towsley, H. Schulzrinne, R. Ramjee, and J. Kurose, "Adaptive playout mechanisms for packetized audio applications in wide-area networks," in *Proc. IEEE Infocom 1994*, Ontario, Canada, June 1994, vol. 2, pp. 680–688.
- [27] W. Verhelst and M. Roelands, "An overlap-add technique based on waveform similarity (WSOLA) for high quality time-scale modification of speech," in *Proc. ICASSP '93*, Apr. 1993, vol. II, pp. 554–557.



**Yan Li** (M'02) received the B.S. degree in electrical engineering and computer science from University of California at Berkeley in 2001 and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford CA, in 2003 and 2006, respectively.

He is currently working in the Residential Gateway and Embedded Systems Group at Texas Instruments, San Jose CA. His research interests include supporting multimedia streaming over wireless networks and performance engineering in ad hoc networks.



**Athina Markopoulou** (SM'98–M'02) received the Diploma degree in electrical and computer engineering from the National Technical University of Athens, Athens, Greece, in 1996 and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1998 and 2002, respectively.

She has been a Postdoctoral Fellow at Sprint Advanced Technologies Labs (2003) and Stanford (2004–2005), a Member of the Technical Staff at Arastra, Inc. (2005), and an Assistant Professor with

the Electrical Engineering and Computer Science Department, University of California at Irvine (2006–present). Her research interests are in the general area of networking, including network reliability and security, support of multimedia traffic over wired and wireless networks, network measurement and control.

Dr. Markopoulou is the recipient of Mathematical Olympiad prizes (1987–1991) and the author of one of the eight best papers in IEEE INFOCOM'02.



**Nick Bambos** received the Ph.D. degree in electrical engineering and computer science from the University of California at Berkeley in 1989 and the Diploma in electrical engineering from the National Technical University of Athens, Athens, Greece, in 1984.

He is a Professor of electrical engineering and of management science at Stanford University, Stanford, CA, where he heads the Network Architecture and Performance Engineering group and the Networking Research Lab. He has been the Cisco

Systems Faculty Scholar and the Morgenthaler Scholar at Stanford University. His current research interests are in wireless network architectures, high-speed switching, robust networking, queueing, and scheduling processes.

Dr. Bambos has been awarded the IBM Faculty Award, the National Science Foundation National Young Investigator Award, and Research Initiation Awards, among many others.



**John Apostolopoulos** (SM'91–M'97) received his B.S., M.S., and Ph.D. degrees from Massachusetts Institute of Technology, Cambridge, MA.

He joined Hewlett-Packard Laboratories, Palo Alto, CA, in 1997, where he is currently a Principal Research Scientist and Project Manager for the Streaming Media Systems Group. He also teaches at Stanford University, Stanford, CA, where he is a Consulting Assistant Professor of Electrical Engineering. He contributed to the U.S. Digital Television and JPEG-2000 Security (JPSEC) standards.

His research interests include improving the reliability, fidelity, scalability, and security of media communication over wired and wireless packet networks.

Dr. Apostolopoulos received a Best Student Paper award for part of his Ph.D. dissertation, the Young Investigator Award at VCIP 2001 for his paper on multiple description video coding and path diversity for reliable video communication over lossy packet networks, and in 2003 was named "one of the world's top 100 young (under 35) innovators in science and technology" (TR100) by *Technology Review*. He served as an Associate Editor of the IEEE TRANSACTIONS ON IMAGE PROCESSING and the IEEE SIGNAL PROCESSING LETTERS, a member of the IEEE Image and Multidimensional Digital Signal Processing (IMDSP) Technical Committee, and he is currently Co-Guest Editor of a special issue of *IEEE Network* on "Multimedia over Broadband Wireless Networks" and General Co-Chair of VCIP'06.