Performance Evaluation 106 (2016) 1–18

Contents lists available at ScienceDirect

Performance Evaluation

journal homepage: www.elsevier.com/locate/peva

Power control for packet streaming with head-of-line deadlines

Neal Master[a,*], Nicholas Bambos[a,b]

[a] Department of Electrical Engineering, Stanford University, Stanford, CA, 94305, USA

[b] Department of Management Sciences & Engineering, Stanford University, Stanford, CA, 94305, USA

ARTICLE INFO

Article history:
Received 21 October 2014
Received in revised form 30 November 2015
Accepted 14 October 2016
Available online 20 October 2016

Keywords:
Service rate control
Deadlines
Wireless streaming
Power control

ABSTRACT

We consider a mathematical model for streaming media packets (as the motivating key example) from a transmitter buffer to a receiver over a wireless link while controlling the transmitter power (hence, the packet/job processing rate). When each packet comes to the head-of-line (HOL) in the buffer, it is given a deadline D which is the maximum number of times the transmitter can attempt retransmission in order to successfully transmit the packet. If this number of transmission attempts is exhausted, the packet is ejected from the buffer and the next packet comes to the HOL. Costs are incurred in each time slot for holding packets in the buffer, expending transmitter power, and ejecting packets which exceed their deadlines. We investigate how transmission power should be chosen so as to minimize the total cost of transmitting the items in the buffer. We formulate the optimal power control problem in a dynamic programming framework and then hone in on the special case of fixed interference. For this special case, we are able to provide a precise analytic characterization of how the power control should vary with the backlog and how the power control should react to approaching deadlines. In particular, we show monotonicity results for how the transmitter should adapt power levels to the backlog and approaching deadlines. We leverage these analytic results from the special case to build a power control scheme for the general case. Monte Carlo simulations are used to evaluate the performance of the resulting power control scheme as compared to the optimal scheme. The resulting power control scheme is sub-optimal but it provides a low-complexity approximation of the optimal power control. Simulations show that our proposed schemes outperform benchmark algorithms. We also discuss applications of the model to other practical operational scenarios.

1. Introduction

Packet streaming over wireless channels is a ubiquitous technology today. Indeed, the rise of smartphones, tablets, wearable electronics, and the Internet of Thing (IoT) has led to an increasing interest in wireless multimedia applications. However, real-time mobile multimedia streaming poses challenges in both theory and practice. Mobile devices have strict power limitations, and as they become more compact, energy efficiency becomes increasingly important. In addition, multimedia streaming is time-sensitive with both latency and jitter constraints. These constraints are complicated by the fact that wireless channel quality fluctuates stochastically in both time and space. These limitations and constraints of a wireless

* Corresponding author.
E-mail addresses: nmaster@stanford.edu (N. Master), bambos@stanford.edu (N. Bambos).

http://dx.doi.org/10.1016/j.peva.2016.10.003

streaming system are typically at odds with one another which makes practical yet effective control schemes difficult to formulate.

In the literature, wireless streaming over cellular networks has been studied in the context of downlink packet scheduling [1,2]. In such systems, a base station performs the task of transmitting multimedia streams over various wireless channels to different end users. Since the base station has limited resources, there is a need for scheduling algorithms for supporting the multitude of streams. In this case, channel quality and time-sensitivity require consideration but the base station does not have the same stringent power restrictions as a mobile device.

A different model is the case of device-to-device (D2D) transmission [3]. In these systems, devices communicate directly over local wireless channels without intermediate base stations or access points. D2D networks have been implemented and studied in the context of LTE [4] as well as Wi-Fi [5]. D2D communication networks are flexible and can be used for applications as diverse as content distribution and vehicular communication. As a result, a subset of a D2D network could likely be a device streaming video to another device. For example, if a particular user has already downloaded a video, streaming the content directly to a nearby user avoids complexity of downlink scheduling and hence reduces the system of interest to a point-to-point communication model. In this paper we consider such a point-to-point system and develop schemes which incorporate power constraints, varying wireless channel quality, as well as latency and jitter constraints. Rather than focusing on a specific set of technologies, we use an abstract mathematical model and prove structural results about the optimal transmitter power control. These properties are used to develop low-complexity schemes and the design principles of these schemes are demonstrated via simulation. The abstract nature of the model allows for the principles to be applied to a variety of technologies.

1.1. Related work

There has been considerable work on the case of streaming multimedia over a wireless link. Li et al. [6] formulated the problem as a joint transmitter power and receiver playout rate control problem. They used the optimal control to motivate useful heuristics for jointly choosing the transmitter power and the receiver playout rate when communicating over an interference-limited wireless link. These heuristics can be used to ensure "smooth" multimedia playout but they do not explicitly capture jitter constraints. In particular, the heuristics do not make use of any form of packet deadlines.

Xing et al. [7] conducted an experimental study of D2D video delivery using Wi-Fi ad-hoc mode. This work gives empirical insights into the technological details of multimedia streaming over Wi-Fi. Unfortunately, power control is not yet part of the IEEE 802.11 standard (which is the basis for Wi-Fi) [8], so the study was unable to incorporate any investigation of transmitter power control.

There has been substantial work on power allocation at the network level; a thorough survey is presented by Chiang et al. [9]. Le [10] studied fair resource allocation in D2D Orthogonal Frequency Division Multiple Access (OFDMA)-based wireless cellular networks. Power was included as a resource and the focus was on general traffic rather than specifically multimedia. Yu et al. [11] studied how to share resources between a D2D network and traditional underlying cellular network. In a general setting (not necessarily D2D), Kandukuri and Boyd [12] used a convex programming framework to optimally allocate power across a network so as to minimize outage probabilities. These works address the question of how to assign a power budget to each transmitter in a network of transmitters. Our work builds on this idea by investigating how a transmitter should dynamically choose to use the allocated power budget.

Dynamic power control and delay control have been considered in the network utility maximization literature. For example, O'Neill et al. [13] used a convex optimization framework to derive power and rate control policies for several modulation schemes. Although rate is implicitly related to the delay of information, this work did not explicitly consider deadlines or delay bounds. In another line of research, Neely [14] used Lyapunov based techniques to approximately maximize throughput utility based on explicit head-of-line packet delays. Neely [15] also investigated the asymptotic trade-off between energy and delay in a wireless downlink system. Our work is different not only in the approach (we use a dynamic programming framework) but also in that our results are not asymptotic.

Outside of the communication engineering community, queueing models with deadlines have been considered in the scheduling literature. Delay-sensitive scheduling has applications in networking of computers [16] and patient triage [17]. More recently, there has been work on scheduling with "impatient" users [18]. In such a model, jobs do not have hard deadlines but users have preferences which create soft deadlines. In all of these works, the focus is on scheduling jobs under different notions of delay sensitivity where each server completes jobs with a fixed rate. In the current work, we are concerned with a complementary issue. The schedule for the packets is fixed (they are served sequentially) but we control the service rate (via the transmitter power) and are interested in understanding how the service rate should vary with the number of jobs (packets) and the approaching deadlines. Service rate control for a single Markovian queue has been previously considered [19]. Besides the focus on wireless streaming, our model differs in a few ways with the most notable difference being the head-of-line deadlines.

Early limited results of this work were presented in our previous conference publication [20], where a baseline model was developed and some preliminary analysis done. The current paper presents a detailed mathematical treatment of the full system model (including random interference), resulting control schemes and emerging design principles, and incorporates additional key metrics into the performance evaluation.