

BOLA: Near-Optimal Bitrate Adaptation for Online Videos

Kevin Spiteri^{ID}, Member, IEEE, Rahul Urgaonkar, Senior Member, IEEE,
and Ramesh K. Sitaraman, Fellow, IEEE

Abstract—Modern video players employ complex algorithms to adapt the bitrate of the video that is shown to the user. Bitrate adaptation requires a tradeoff between reducing the probability that the video freezes (rebuffers) and enhancing the quality of the video. A bitrate that is too high leads to frequent rebuffering, while a bitrate that is too low leads to poor video quality. Video providers segment videos into short segments and encode each segment at multiple bitrates. The video player adaptively chooses the bitrate of each segment to download, possibly choosing different bitrates for successive segments. We formulate bitrate adaptation as a utility-maximization problem and devise an online control algorithm called BOLA that uses Lyapunov optimization to minimize rebuffering and maximize video quality. We prove that BOLA achieves a time-average utility that is within an additive term $O(1/V)$ of the optimal value, for a control parameter V related to the video buffer size. Further, unlike prior work, BOLA does not require prediction of available network bandwidth. We empirically validate BOLA in a simulated network environment using a collection of network traces. We show that BOLA achieves near-optimal utility and in many cases significantly higher utility than current state-of-the-art algorithms. Our work has immediate impact on real-world video players and for the evolving DASH standard for video transmission. We also implemented an updated version of BOLA that is now part of the standard reference player dash.js and is used in production by several video providers such as Akamai, BBC, CBS, and Orange.

Index Terms—Internet video, video quality, adaptive bitrate streaming, Lyapunov optimization, optimal control.

I. INTRODUCTION

ONLINE videos are the “killer” application of the Internet with videos currently accounting for more than half of the Internet traffic. Video viewership is growing at a torrid pace and videos are expected to account for more than 85% of all Internet traffic within a few years [2]. As all forms of traditional media migrate to the Internet, video providers face the daunting challenge of providing a good quality of experience (QoE) for users watching their videos. Video providers are diverse and include major media companies (e.g., NBC, CBS),

Manuscript received December 23, 2018; revised September 14, 2019; accepted April 25, 2020; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor S. Rao. This work was supported in part by the NSF under Grant CNS-1413998 and Grant CNS-1901137. A preliminary version of this article appeared at INFOCOM 2016. (Corresponding author: Kevin Spiteri.)

Kevin Spiteri and Ramesh K. Sitaraman are with the College of Information and Computer Sciences, University of Massachusetts at Amherst, Amherst, MA 01003 USA (e-mail: kspiteri@cs.umass.edu; ramesh@cs.umass.edu).

Rahul Urgaonkar is with Amazon Prime Video, Seattle, WA 98109 USA (e-mail: urgaonka@amazon.com).

Digital Object Identifier 10.1109/TNET.2020.2996964

news outlets (e.g., CNN), sports organizations (e.g., NFL, MLB), and video subscription services (e.g., Netflix, Hulu). Recent research has shown that low-performing videos that start slowly, play at lower bitrates, and freeze frequently can cause viewers to abandon the videos or watch fewer minutes of the videos, significantly decreasing the opportunity for generating revenue for the video providers [3]–[5], underscoring the need for a high-quality user experience.

Providing a high-quality experience for video users requires balancing two contrasting requirements. The user would like to watch the highest-quality version of the video possible, where video quality can be quantified by the bitrate at which the video is encoded. For instance, watching a movie in high definition (HD) encoded at 2 Mbps arguably provides a better user experience than watching the same movie in standard definition (SD) encoded at a bitrate of 800 kbps. In fact, there is empirical evidence that the user is more engaged and watches longer when the video is presented at a higher bitrate. However, it is not always possible for users to watch videos at the highest encoded bitrate, since the bandwidth available on the network connection between the video player on the user’s device and the video server constrains what bitrates can be watched. In fact, choosing a bitrate that is higher than the available network bandwidth¹ will lead to video freezes in the middle of the playback, since the rate at which the video is being played exceeds the rate at which the video can be downloaded. Such video freezes are called *rebuffers* and playing the video continuously *without rebuffers* is a key factor in the QoE perceived by the user [4]. Thus, balancing the contrasting requirements of playing videos at a high bitrate while at the same time avoiding rebuffers is central to providing a high-quality video watching experience.

A. Adaptive Bitrate (ABR) Streaming

Achieving a high QoE for video streaming is a major challenge due to the sheer diversity of video-capable devices that include smartphones, tablets, desktops, and televisions. Further, the devices themselves can be connected to the Internet in a multitude of ways, including cable, fiber, DSL, WiFi and mobile wireless, each providing different bandwidth characteristics. The need to adjust the video playback to the characteristics of the device and the network has led to the evolution of adaptive bitrate (ABR) streaming that is now the de facto standard for delivering videos on the Internet.

¹Throughout this paper, we say **bandwidth** when talking about network throughput and **bitrate** when talking about encoding quality.

ABR streaming requires that each video is partitioned into *segments*, where each segment corresponds to a few seconds of play. Each segment is then encoded in a number of different bitrates to accommodate a range of device types and network connectivities. When the user plays a video, the video player can download each segment at a bitrate that is appropriate for the available bandwidth of the network connection. Thus, the player can switch to a segment with a lower bitrate when the available bandwidth is low to avoid rebuffering. If more bandwidth becomes available at a future time, the player can switch back to a higher bitrate to provide a richer experience. The video player has a buffer that allows it to fetch and store segments *before* they need to be rendered on the screen. Thus, the video player can tolerate brief network disruptions without interrupting the playback of the user by using the buffered segments. A large disruption, however, will empty the buffer, resulting in rebuffering. The decision of which segments to download at what bitrates is made by a *bitrate adaptation algorithm* within the video player, the design of such algorithms being the primary focus of our work.

Several popular implementations of ABR streaming exist, including Apple's HTTP Live Streaming (HLS) [6], Microsoft's Live Smooth Streaming (Smooth) [7] and Adobe's Adaptive Streaming (HDS) [8]. Each has its own proprietary implementation and slight modifications to the basic ABR technique described above. A key recent development is a unifying open-source standard for ABR streaming called MPEG-DASH [9]. DASH is broadly similar to the other ABR protocols and is a particular focus in our empirical evaluation.

B. Our Contributions

Our primary contribution is a principled approach to the design of bitrate adaptation algorithms for ABR streaming. In particular, we formulate bitrate adaptation as a utility maximization problem that incorporates both key components of QoE: the average bitrate of the video experienced by the user and the duration of the rebuffer events. An increase in the average bitrate increases utility, whereas rebuffering decreases it. A strength of our framework is that utility can be defined in a very general manner, say, depending on the content, video provider, or user device.

Using Lyapunov optimization, we derive an online bitrate adaptation algorithm called BOLA (Buffer Occupancy based Lyapunov Algorithm) that provably achieves utility that is within an additive factor of the maximum possible utility in the large video regime. While numerous bitrate adaptation algorithms have been proposed [10]–[15] and implemented within video players, our algorithm is the first to provide a *theoretical guarantee* on the achieved utility. Further, BOLA provides an explicit knob for video providers to set the relative importance of a high video quality in relation to the probability of rebuffering.

While not an explicit part of the Lyapunov optimization framework, we also show how BOLA can be adapted to avoid frequent bitrate switches during video playback. Bitrate switches are arguably less annoying than rebuffering, but it is still of some concern to video providers and users alike if such switches occur too frequently.

Most algorithms implemented in practice use a *bandwidth-based* approach where the available bandwidth between the

server and the video player is predicted and the predicted value is used to determine the bitrate of the next segment that is to be downloaded. A complementary approach is a *buffer-based* approach that does not predict the bandwidth, but only uses the amount of data that is currently stored in the buffer of the video player. Recently, there has been empirical evidence that a buffer-based approach has desirable properties that bandwidth-based approaches lack and has been adopted by Netflix [11]. An intriguing outcome of our work is that the optimal algorithm within our utility maximization framework requires only knowledge of the amount of data in the buffer and no estimate of the available bandwidth. Thus, our work provides the first theoretical justification for why buffer-based algorithms perform well in practice and adds new insights to the ongoing debate [14] within the video streaming and DASH standards communities of relative efficacy of the two approaches. Further, since our algorithm BOLA is buffer-based, it avoids the overheads of more complex bandwidth prediction present in current video player implementations and is more stable under bandwidth fluctuations. Note that our results imply that the buffer level is a *sufficient statistic* that indirectly provides all information about past bandwidth variations required for choosing the next bitrate.

We also *empirically* evaluate BOLA on a wide set of network traces that include 12 test cases provided by the DASH industry forum [16] and 85 publicly-available 3G mobile bandwidth traces [17]. As a benchmark for comparison, we develop an optimal *offline* algorithm that uses dynamic programming and is guaranteed to produce the maximum achievable time-average utility for any given set of network traces. Unlike BOLA that works in an online fashion, the offline optimal algorithm makes decision based on perfect knowledge of future bandwidth variations. Remarkably, the utility achieved by BOLA is within 84–95% of offline optimal utility for all the tested traces.

Besides comparing BOLA with the offline optimal, we also empirically compared our algorithm with four state-of-the-art algorithms proposed in the literature. In *all* test cases, BOLA achieved a utility that is as good as or better than the best state-of-the-art algorithm.

We also implemented BOLA in dash.js, the open-source standard DASH reference player [18]. Deploying BOLA in production required a number of adjustments [19]. Through dash.js, BOLA is now being used in production by several major video providers and delivery networks such as Akamai, BBC, CBS and Orange. BOLA is available as an option to commercial video providers who often use the production dash.js reference implementation for building their own video players. Further, a second algorithm called DYNAMIC [19] is also available for commercial video providers. DYNAMIC is a hybrid algorithm that uses a simple throughput-estimation approach during the start-up phase of the video and then uses BOLA afterwards. Both algorithms can be evaluated in a web browser by clicking the “Show Options” button in the latest version of the dash.js reference player found at [18].

II. SYSTEM MODEL

Our system model closely captures how ABR streaming works on the Internet today. We consider a video player that downloads a video file from a server over the Internet and

plays it back to the user. The video file is segmented into segments that are downloaded in succession. The available bandwidth between the server and the player varies over time. This can be due to reasons such as network congestion and wireless fading among others. The viewing experience of the user is determined by both the video quality as quantified by the bitrates of the segments that are played back and the playback characteristics such as rebuffering. The objective of the player is to maximize a utility associated with the user's viewing experience while adapting to time-varying (and possibly unpredictable) changes in the available bandwidth.

Video Model: The video file is segmented into N segments indexed as $\{1, 2, \dots, N\}$ where each segment represents p seconds of the video. On the server, each segment is available in M different bitrates where a segment encoded at a higher bitrate has a larger size in bits and its playback provides a better user experience and higher utility. Suppose the size (in bits) of any² segment encoded at bitrate index m is S_m bits and suppose the utility derived by the user from viewing it is given by v_m where $m \in \{1, 2, \dots, M\}$. WLOG, let the segment bitrates be non-decreasing in index m . Then, the following holds.

$$v_1 \leq v_2 \leq \dots \leq v_M \iff S_1 \leq S_2 \leq \dots \leq S_M. \quad (1)$$

Note that the actual encoding bitrate for bitrate index m is given by S_m/p bits/second.

Video Player: The video player downloads successive segments of the video file from the server and plays back the downloaded segments to the user. Each segment must be downloaded in its entirety before it can be played back. We assume that the player sends requests to the server to download one segment at a time. Also, the segments are downloaded in the same order as they are played back. The video player has a finite buffer of size Q_{\max} segments³ to store the downloaded but yet-to-be-played-back segments. Measuring the buffer in segments is equivalent to measuring it in seconds since the segment duration p is fixed. If the buffer is full the player cannot download any new segments and waits for a fixed period of time given by Δ seconds before attempting to download a new segment. The segments that are fully downloaded are played back at a fixed rate of $1/p$ segments/second without any idling.

When sending a download request for a new segment, the player also specifies the desired bitrate for that segment. This enables the player to tradeoff the overall video quality with the likelihood of rebuffering that occurs when there are no segments in the buffer for playback. Note that while each segment has a fixed playback time of p seconds, the size of the segment (in bits) can be different depending on its bitrate. Thus, the choice of bitrate for a segment impacts its download time.

Network Model: The available bandwidth (in bits/second) between the server and player is assumed to vary continuously in time according to a stationary random process $\omega(t)$. We do not make any assumptions about knowing the statistical properties or probability distribution of $\omega(t)$ except that it has

²For simplicity, we assume that the segment size (in bits) is S_m for all segments of a given bitrate index m . However, our framework can be easily extended to the case where the segment size for the same bitrate can vary across segments.

³It is common practice for video players to measure the buffer in seconds of playback time rather than in bits.

finite first and second moments as well as a finite inverse second moment. Suppose the player starts to download a segment of bitrate index m at time t . Then the time t' when the download finishes satisfies the following:

$$S_m = \int_t^{t'} \omega(\tau) d\tau \quad (2)$$

Let $\mathbb{E}\{\omega(t)\} = \omega_{\text{avg}}$. Then, $\mathbb{E}\{t' - t\} = S_m / \omega_{\text{avg}}$.

III. PROBLEM FORMULATION

We consider two primary performance metrics⁴ that affect the overall QoE of the user: (1) time-average playback quality which is a function of the bitrates of the segments viewed by the user and (2) fraction of time spent not rebuffering. To formalize these metrics, we consider a time-slotted representation of our system model. The timeline is divided into non-overlapping consecutive slots of variable length and indexed by $k \in \{1, 2, \dots\}$. Slot k starts at time t_k and is $T_k = t_{k+1} - t_k$ seconds long. We assume that $t_1 = 0$. At the beginning of each slot, the video player makes a control decision on whether it should start downloading a new segment, and if yes, its bitrate. If a download decision is made, then a request is sent to the server and the download starts immediately.⁵ This download takes T_k seconds and is completed at the end of slot k . Note that T_k is a random variable whose actual value depends on the realization of the $\omega(t)$ process as well as the choice of segment bitrate. If the player decides not to download a new segment in slot k (for example, when the buffer is full), then this slot lasts for a fixed duration of Δ seconds.

We define the following indicator variable for each slot k :

$$a_m(t_k) = \begin{cases} 1 & \text{if the player downloads a segment} \\ & \text{of bitrate index } m \text{ in slot } k, \text{ and} \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

Then, for all k , we must have $\sum_{m=1}^M a_m(t_k) \leq 1$. Moreover, when $\sum_{m=1}^M a_m(t_k) = 0$, then no segments are downloaded.

Denote the buffer level (measured in number of segments) at the start of slot k by $Q(t_k)$. The dynamics of this queue can be expressed using the following equation:

$$Q(t_{k+1}) = \max[Q(t_k) - \frac{T_k}{p}, 0] + \sum_{m=1}^M a_m(t_k) \quad (4)$$

Here, the arrival value into this queue in slot k is given by $\sum_{m=1}^M a_m(t_k)$ which is 1 if a download decision is made in slot k and 0 otherwise. The departure value is T_k/p which represents the total number of segments (including fractional segments) that could have departed the buffer in slot k . Note that the actual value of T_k is revealed at the end of slot k . Also note that a segment that is downloaded in slot k becomes available for playback only from the next slot. We assume that the buffer level is initialized to 0, i.e., $Q(t_1) = 0$.

Let K_N denote the index of the slot in which the N^{th} (i.e., last) segment is downloaded. Also, denote the time at which the player finishes playing back the last segment by T_{end} . Then

⁴We do not include the secondary objective of avoiding frequent bitrate switches in our formulation, but we deal with it empirically in Section V-E.

⁵Any delays associated with sending the request can be added to the overall download time.

the first performance metric of interest is the time-average expected *playback utility* \bar{v}_N which is defined as

$$\bar{v}_N \triangleq \frac{\mathbb{E} \left\{ \sum_{k=1}^{K_N} \sum_{m=1}^M a_m(t_k) v_m \right\}}{\mathbb{E} \{T_{\text{end}}\}} \quad (5)$$

where the numerator denotes the expected total utility across all N segments. Note that a segment can only be played back after it has been downloaded entirely. Thus, T_{end} is greater than the last segment's download finish time, i.e., $T_{\text{end}} > t_{K_N} + T_{K_N}$.

The second performance metric of interest is the expected fraction of time \bar{s}_N that is spent not rebuffering and can be interpreted as a measure of the average playback "smoothness". This can be calculated by observing that the actual playback time for all N segments is Np seconds. Thus, the expected *playback smoothness* \bar{s}_N is given by

$$\bar{s}_N \triangleq \frac{Np}{\mathbb{E} \{T_{\text{end}}\}} = \frac{\mathbb{E} \left\{ \sum_{k=1}^{K_N} \sum_{m=1}^M a_m(t_k) p \right\}}{\mathbb{E} \{T_{\text{end}}\}} \quad (6)$$

where in the last step we use the relation that $Np = \sum_{k=1}^{K_N} \sum_{m=1}^M a_m(t_k) p$. Note that $T_{\text{end}} \geq Np$ (since at most one segment can be played back at any time), so that $\bar{s}_N \leq 1$.

A. Design Objective

We want to design a control algorithm that maximizes the joint utility $\bar{v}_N + \gamma \bar{s}_N$ subject to the constraint that $Q(t_k) \leq Q_{\max}$ for all k . Here, $\gamma > 0$ is an input weight parameter for prioritizing playback utility versus the playback smoothness.

This problem can be formulated as a stochastic optimization problem with a time-average objective over a finite horizon and dynamic programming (DP) based approaches can be used to solve it [20]. However, traditional DP based methods have two major disadvantages. First, they require knowledge of the distribution of the $\omega(t)$ process which may be hard to obtain. Second, even when such knowledge is available, the resulting DP can have a very large state space. This is because the state space for this problem under a DP formulation would consist of not only the timeslot index k and value t_k , but also the buffer size $Q(t_k)$. Further, an appropriate discretization of the $\omega(t)$ process would be required to obtain a tractable solution.

In order to overcome the above challenges associated with traditional DP based methods, we take an alternate approach in this paper. First, we consider the bitrate adaptation problem in the limiting regime when the video size becomes large, i.e., $N \rightarrow \infty$. Second, we replace the finite buffer constraint with a *rate stability* constraint (made precise in the next section). The reason for making these assumptions is that it results in simplifications to the original problem as discussed in the next section. This allows us to develop a bitrate adaptation algorithm that does not require *any* knowledge of the distribution of $\omega(t)$, yet offers provable theoretical performance guarantees in the large video size regime while satisfying the finite buffer constraint. As shown later in Section V-D, with slight modifications, this algorithm can be used for finite sized videos as well and offers close to optimal performance in our experiments.

B. Problem Relaxation

Consider the bitrate adaptation problem in the limiting regime when the video size becomes large, i.e., $N \rightarrow \infty$.

Then, the metrics \bar{v}_N and \bar{s}_N can be expressed as

$$\begin{aligned} \bar{v} &\triangleq \lim_{N \rightarrow \infty} \bar{v}_N = \lim_{N \rightarrow \infty} \frac{\mathbb{E} \left\{ \sum_{k=1}^{K_N} \sum_{m=1}^M a_m(t_k) v_m \right\}}{\mathbb{E} \{T_{\text{end}}\}} \\ &= \frac{\lim_{K_N \rightarrow \infty} \frac{1}{K_N} \mathbb{E} \left\{ \sum_{k=1}^{K_N} \sum_{m=1}^M a_m(t_k) v_m \right\}}{\lim_{K_N \rightarrow \infty} \frac{1}{K_N} \mathbb{E} \left\{ \sum_{k=1}^{K_N} T_k \right\}} \end{aligned} \quad (7)$$

$$\begin{aligned} \bar{s} &\triangleq \lim_{N \rightarrow \infty} \bar{s}_N = \lim_{N \rightarrow \infty} \frac{\mathbb{E} \left\{ \sum_{k=1}^{K_N} \sum_{m=1}^M a_m(t_k) p \right\}}{\mathbb{E} \{T_{\text{end}}\}} \\ &= \frac{\lim_{K_N \rightarrow \infty} \frac{1}{K_N} \mathbb{E} \left\{ \sum_{k=1}^{K_N} \sum_{m=1}^M a_m(t_k) p \right\}}{\lim_{K_N \rightarrow \infty} \frac{1}{K_N} \mathbb{E} \left\{ \sum_{k=1}^{K_N} T_k \right\}} \end{aligned} \quad (8)$$

This follows by noting that the difference between the expected total playback finish time $\mathbb{E} \{T_{\text{end}}\}$ and the expected total download finish time $\mathbb{E} \left\{ \sum_{k=1}^{K_N} T_k \right\}$ is upper bounded by a finite value due to the finite Q_{\max} . Specifically, this upper bound is given by $Q_{\max} p$. Therefore, instead of considering the total playback finish time, we can consider the total download finish time in the objective when the video size becomes large.

Next, replace the finite buffer constraint with a rate stability constraint [21]. This constraint only requires that the time-average arrival rate into the buffer cannot exceed the time-average playback rate. This is equivalent to requiring that

$$\lim_{K_N \rightarrow \infty} \frac{1}{K_N} \mathbb{E} \left\{ \sum_{k=1}^{K_N} \sum_{m=1}^M a_m(t_k) p \right\} \leq \lim_{K_N \rightarrow \infty} \frac{1}{K_N} \mathbb{E} \left\{ \sum_{k=1}^{K_N} T_k \right\} \quad (9)$$

The rate stability constraint is a relaxation of the finite buffer constraint since any policy that ensures finite buffers is always rate stable but not vice versa. Therefore, under this relaxation, the optimal time-average utility cannot be smaller than the optimal time-average utility with the finite buffer constraint.

With these relaxations, our performance objective for the bitrate adaptation problem is to maximize the joint utility $\bar{v} + \gamma \bar{s}$ subject to the rate stability constraint (9). Let us denote the optimal time-average utility for this problem by $v^* + \gamma s^*$. This problem fits in the framework of Lyapunov optimization for renewal systems [22]. Specifically, this framework extends the original Lyapunov optimization technique [21] to systems with variable length renewal frames and shows that minimizing a "drift-plus-penalty" ratio over every frame yields an optimal control algorithm. We refer to [22] for details on this method. In the context of our bitrate adaptation problem, the variable length slots represent the renewal frames.

The following characterization can be made about the optimality of *i.i.d. algorithms*.

Lemma 1: For the bitrate adaptation problem in the limiting regime when the video size becomes large, i.e., $N \rightarrow \infty$, there exists a *buffer-state-independent* stationary algorithm that makes *i.i.d.* control decisions in every slot and satisfies the rate stability constraint while achieving time-average utility no smaller than $v^* + \gamma s^*$.

Proof: This follows from Lemma 1 in [22] and uses the fact that the conditional expectations and conditional second

moments of the frame length and utility are bounded under any algorithm. The full proof is omitted for brevity. ■

Note that such a buffer-state-independent stationary algorithm is not necessarily feasible for our finite buffer system. Further, calculating it explicitly would require knowledge of the distribution of $\omega(t)$. However, instead of calculating this policy explicitly, we will use its existence and characterization per Lemma 1 to design an *online* control algorithm using the technique of Lyapunov optimization over renewal frames.

In the next section, we will present this algorithm and show that it meets the finite buffer constraint while achieving a time-average utility that is within $O(1/Q_{\max})$ of $v^* + \gamma s^*$ without requiring any knowledge of the distribution of $\omega(t)$.

IV. BOLA: AN ONLINE CONTROL ALGORITHM

We first give a high-level intuition of the Lyapunov optimization over renewals technique. This technique converts the problem of optimizing the time-average metrics in (7)–(8) subject to the time-average constraint in (9) into a series of per slot optimization problems. The problem to be solved in each slot involves minimizing a ratio of the expected drift-plus-penalty value in that slot to the expected length of the slot. As shown in the Appendix, this can be done without requiring any knowledge of the distribution of $\omega(t)$. The drift term consists of $\mathbb{E}\{(Q(t_{k+1})^2 - Q(t_k)^2)/2 \mid Q(t_k)\}$ and serves to meet the rate stability constraint (9). The penalty term consists of the playback utility and playback smoothness received in that slot. We keep the utility and smoothness as separate terms even though they can be folded into one metric. This allows us to tune the relative importance of increasing video bitrate and reducing rebuffering without changing the algorithm. The algorithm uses a control parameter $V > 0$ to allow a tradeoff between the buffer size and the performance objectives.

We now present the algorithm. In every slot k , given the buffer level $Q(t_k)$ at the start of the slot, our algorithm makes a control decision by solving the following deterministic optimization problem. Let

$$\rho(t_k, \mathbf{a}(t_k)) = \begin{cases} 0 & \text{if } \sum_{m=1}^M a_m(t_k) = 0, \\ \frac{\sum_{m=1}^M a_m(t_k)(Vv_m + V\gamma p - Q(t_k))}{\sum_{m=1}^M a_m(t_k)S_m} & \text{otherwise.} \end{cases} \quad (10)$$

Then determine $\mathbf{a}(t_k)$ by solving the optimization problem:

Maximize: $\rho(t_k, \mathbf{a}(t_k))$

Subject to: $\sum_{m=1}^M a_m(t_k) \leq 1, \quad a_m(t_k) \in \{0, 1\}$ (11)

The constraints of this problem result in a very simple solution structure. Specifically, the optimal solution is given by:

- 1) If $Q(t_k) > V(v_m + \gamma p)$ for all $m \in \{1, 2, \dots, M\}$, then the no-download option is chosen, i.e., $a_m(t_k) = 0$ for all m . Note that in this case $T_k = \Delta$.
- 2) Else, the optimal solution is to download the next segment at bitrate index m^* where m^* is the index that maximizes the ratio $(Vv_m + V\gamma p - Q(t_k))/S_m$ among all m for which this ratio is positive.

Notice that solving this problem does not require any knowledge of the $\omega(t)$ process. Further, the optimal solution depends only on the buffer level $Q(t_k)$. That's why we call our algorithm *BOLA: Buffer Occupancy based Lyapunov Algorithm*. These properties of BOLA should be contrasted with the bandwidth prediction based strategies that have been recently proposed for this problem that require explicit prediction of the available bandwidth for control decisions.

The following theorem characterizes the theoretical performance guarantees provided by BOLA.

Theorem 1: Suppose BOLA as defined by (11) is implemented in every slot using a control parameter $0 < V \leq \frac{Q_{\max}-1}{v_M + \gamma p}$. Assume $Q(0) = 0$. Then, the following hold.

- 1) The queue backlog satisfies $Q(t_k) \leq V(v_M + \gamma p) + 1$ for all slots k . Further, the buffer occupancy in segments never exceeds Q_{\max} .
- 2) The time-average utility achieved by BOLA satisfies

$$\bar{v}^{\text{BOLA}} + \gamma \bar{s}^{\text{BOLA}} \geq v^* + \gamma s^* - \frac{p^2 + \Psi}{2p^2V} \quad (12)$$

where Ψ is an upper bound on $\mathbb{E}\{T_k^2\}$ under any control algorithm and is assumed to be finite. ■

Proof: See the Appendix.

Remarks: The performance bounds in Theorem 1 show a $[O(1/V), O(V)]$ utility and backlog tradeoff that is typical of Lyapunov based control algorithms for similar utility maximization problems. Specifically, the time-average utility of BOLA is within an $O(1/V)$ additive term of the optimal utility and this gap may be made smaller by choosing a larger value of V . However, the largest feasible value of V is constrained by the buffer size and there is a linear relation between them.

A. Understanding BOLA With an Example

We now present a sample run to illustrate how BOLA works. We slice a 99-second video using 3-second segments and encode it at five different bitrates. While BOLA only requires the utilities to be a non-decreasing function of the segment bitrate, it is natural to consider concave utility functions with diminishing returns, e.g., a 1 Mbps increase in segment bitrate likely provides a larger utility gain for the user when that increase is from 0.5 Mbps to 1.5 Mbps than when it is from 5 Mbps to 6 Mbps. A natural choice for our example is the logarithmic utility function: let $v_m = \ln(S_m/S_1)$. Pick $\gamma = 5.0/p$ and $V = 0.93$. The bitrates and utilities are below.

bitrate (Mbps)	0.331	0.688	1.427	2.962	6.000
S (Mb)	0.993	2.064	4.281	8.886	18.00
v	0.000	0.732	1.461	2.192	2.897

For any slot we choose the segment bitrate to maximize $(Vv_m + V\gamma p - Q)/S_m$ for $1 \leq m \leq M$. Fig. 1 shows the relationship between the expression and the buffer level Q for different m . The line intersections mark the buffer levels that correspond to decision thresholds. Fig. 2 summarizes BOLA's bitrate choices as a function of the buffer level.

Fig. 3 shows how BOLA works. We use a synthetic network bandwidth profile as shown in Fig. 3(a). We can see the feedback loop involving the bitrate in (a) and the buffer level in (b). BOLA chooses the bitrate based directly on the buffer level using Fig. 2. The bitrate affects the download time, thus

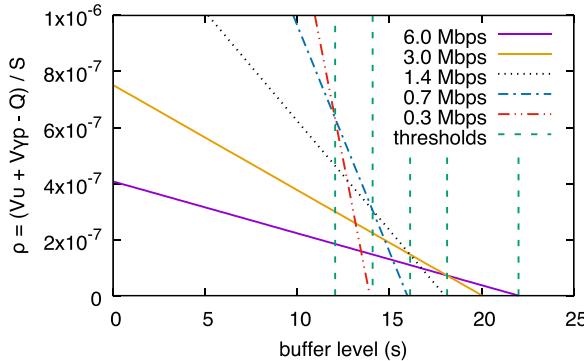


Fig. 1. The value of $(Vv_m + V\gamma p - Q)/S_m$ for different bitrates depends on buffer level. ($\gamma p = 5$ and $V = 0.93$.) Note that the buffer level is Qp seconds.

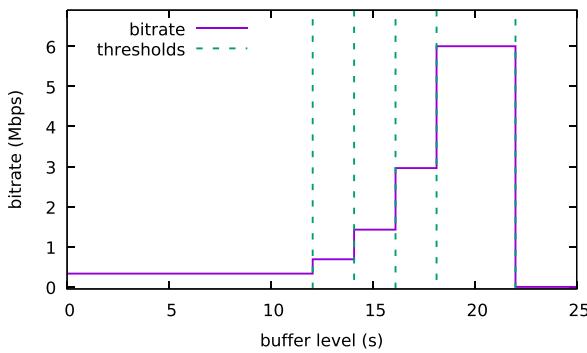


Fig. 2. BOLA's bitrate choice as function of buffer level. ($\gamma p = 5, V = 0.93$.) Note that the buffer level is Qp seconds.

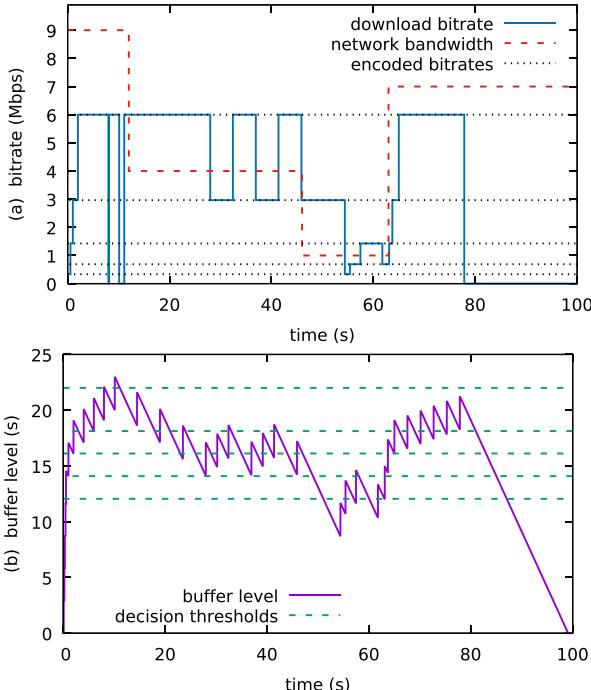


Fig. 3. Sample video download and playback using BOLA. (a) The video is encoded at 5 different bitrates. The network bandwidth varies from high to low and back to high. The downloaded segment bitrate adapts to the network bandwidth. (b) The buffer level variation triggers bitrate changes when it crosses the thresholds.

it indirectly affects the buffer level at the beginning of the following slot. Finally, when all the segments are downloaded,

the video player plays out the segments remaining in the buffer.

B. Choosing Utility and Parameters γ and V

While we chose a logarithmic utility function for the example, a video provider can use any utility function satisfying (1). The utility function might also take into account system characteristics such as the type of device a viewer is using.

γ corresponds to how strongly we want to avoid rebuffing. Increasing γ translates the graphs in Figs. 1 and 2 to the right, effectively shifting the thresholds higher without changing their relative distance. BOLA will thus download more low-bitrate segments to maintain a larger (and safer) buffer level.

Increasing V expands the graphs in Figs. 1 and 2 horizontally about the origin. If we have a maximum buffer level Q_{\max} we want to avoid downloading unless there is enough space for one full segment on the buffer, that is unless $Q \leq Q_{\max} - 1$. For a given Q_{\max} we can set $V = (Q_{\max} - 1)/(v_M + \gamma p)$.

While we showed how to choose reasonable values for γ and V , video providers are more familiar with choosing buffer level targets. A method to derive the parameters from buffer level targets is included in Section VI-B. Alternatively, video providers might choose γ and V by employing an approach such as Oboe [23] to auto-tune the BOLA parameters.

V. IMPLEMENTATION AND EMPIRICAL EVALUATION

We first implemented a basic version of BOLA, named BOLA-BASIC, directly from (11). Recall that when the buffer level is full BOLA does not download a segment but waits for Δ seconds. Rather than picking an arbitrary value for Δ , we use a dynamic wait until $Q(t_k) \leq V(v_M + \gamma p)$. This has the same effect as picking a fixed but very small Δ , so the theoretical analysis still holds. We also implemented other versions of BOLA, namely BOLA-FINITE, BOLA-O, and BOLA-U, that we describe later in this section.

A. Test Methodology

We simulated all versions of BOLA using the Big Buck Bunny movie [24]. The 10-minute movie was encoded at 10 different bitrates and sliced in 3-second segments. Although each quality index has a specified average bitrate, segments may have variable bitrate (VBR) because of the varying nature of the movie. We simulate playback times longer than 10 minutes by repeating the movie. Again we choose a logarithmic utility function: $v_m = \ln(S_m/S_1)$. Table I shows the mean and standard deviation of the bitrate and segment size for each quality index and the respective utility values.

The DASH Industry Forum provides benchmarks for various aspects of the DASH standard [16]. The benchmarks include twelve different network profiles. Profiles 1–6 have network bandwidths ranging from 1.5 to 5 Mbps while profiles 7–12 have bandwidths ranging from 1 to 9 Mbps. Different latencies are provided for each bandwidth, where the latency is half the round-trip time (RTT). Table II shows the odd-numbered bandwidth characteristics. Profile 1 spends 30s at each of 5, 4, 3, 2, 1.5, 2, 3 and 4 Mbps respectively, then starts back at the top. Even-numbered profiles are similar to the preceding odd-numbered profiles but start at the low bandwidth stage. For example, profile 2 starts at 1.5 Mbps.

TABLE I
BITRATES USED FOR BIG BUCK BUNNY TEST VIDEO

Bitrate Index m	Bitrate (Mbps)	Segment Size S (Mb)	Utility v
	Mean	Standard Deviation	$= \ln(S/S_1)$
1	0.230	0.038	0.690
2	0.331	0.054	0.993
3	0.477	0.096	1.431
4	0.688	0.120	2.064
5	0.991	0.182	2.973
6	1.427	0.275	4.281
7	2.056	0.394	6.168
8	2.962	0.564	8.886
9	5.027	0.891	15.08
$M = 10$	6.000	1.078	18.00
			3.232
			3.261

TABLE II
NETWORK PROFILES FOR THE DASH BENCHMARKS

1 Mbps (ms)	3 Mbps (ms)	5 Mbps (ms)	7 Mbps (ms)	9 Mbps (ms)	11 Mbps (ms)
5.0 (38)	5.0 (13)	5.0 (11)			
4.0 (50)	4.0 (18)	4.0 (13)	9.0 (25)	9.0 (10)	9.0 (6)
3.0 (75)	3.0 (28)	3.0 (15)	4.0 (50)	4.0 (50)	4.0 (13)
2.0 (88)	2.0 (58)	2.0 (20)	2.0 (75)	2.0 (150)	2.0 (20)
1.5 (100)	1.5 (200)	1.5 (25)	1.0 (100)	1.0 (200)	1.0 (25)
2.0 (88)	2.0 (58)	2.0 (20)	2.0 (75)	2.0 (150)	2.0 (20)
3.0 (75)	3.0 (28)	3.0 (15)	4.0 (50)	4.0 (50)	4.0 (13)
4.0 (50)	4.0 (18)	4.0 (13)			

In addition, we also tested our algorithms using a set of 86 3G mobile bandwidth traces that are publicly available [17]. One trace was excluded because it had an average bandwidth of 80 kbps; our lowest video bitrate is 230 kbps. Since the traces do not include latency measurements, we used 50 ms latency giving a RTT of 100 ms throughout. This is the median RTT measured empirically in [25].

B. Computing an Upper Bound on the Maximum Utility

In order to evaluate how well BOLA performs on the traces, it is important to derive an upper bound on the maximum utility that is obtainable by *any* algorithm on a given trace. We derive an *offline* optimal algorithm that provides the maximum achievable utility using dynamic programming. We define a table $r(n, t, b)$ that contains the maximum utility possible when we download the n^{th} segment and finish at time t with buffer level b . We initialize the table with $r(0, 0, 0) = 0$. Let $x(n, t, m)$ be the time to download the n^{th} segment at bitrate index m starting at time t . Note that the dependency of x on n is due to VBR. We quantize the time with granularity δ . While some accuracy is lost, we ensure the final result will still be an upper bound by rounding the download time down.

$$x_\delta(n, t, m) = \lfloor x(n, t, m) / \delta \rfloor \cdot \delta$$

We cap the buffer level at b_{\max} .

$$x'_\delta(n, t, b, m) = \max[x_\delta(n, t, m), b + p - b_{\max}]$$

Let $y(n, t, b, m)$ be the rebuffering time.

$$y(n, t, b, m) = \max[x'_\delta(n, t, b, m) - b, 0]$$

We generate entries for $r(n, \cdot, \cdot)$ from $r(n-1, \cdot, \cdot)$ using

$$r(n, t, b) = \max_{m, t', b'} \left(r(n-1, t', b') + v_m - \gamma y(n, t', b', m) \right)$$

such that $t = t' + x'_\delta(n, t', b', m)$ and $b = b' - x'_\delta(n, t', b', m) + y(n, t', b', m) + p$.

The dynamic programming algorithm is shown in Fig. 4.

```

1:  $r(0, t, b) \leftarrow \{0 \text{ for } t = b = 0, -\infty \text{ otherwise}\}$ 
2: for  $n$  in  $[1, N]$  do
3:   initialize  $r(n, t, b) \leftarrow -\infty$  for all  $t, b$ 
4:   for all  $(t', b')$  such that  $r(n-1, t', b') > -\infty$  do
5:     for  $m$  in  $[1, M]$  do
6:        $x \leftarrow \text{download time}(n, t', m)$ 
7:        $x_\delta \leftarrow \lfloor x / \delta \rfloor \cdot \delta$ 
8:        $x'_\delta \leftarrow \max[x_\delta, b' + p - b_{\max}]$ 
9:        $y \leftarrow \max[x'_\delta - b', 0]$ 
10:       $t \leftarrow t' + x'_\delta$ 
11:       $b \leftarrow b' - x'_\delta + y + p$ 
12:       $r' \leftarrow r(n-1, t', b') + v_m - \gamma y$ 
13:       $r(n, t, b) \leftarrow \max[r(n, t, b), r']$ 
14:    end for
15:  end for
16: end for
17:  $r^* \leftarrow \max_{(t, b)} \frac{r(N, t, b)}{(t + b)}$ 

```

Fig. 4. Calculating the offline optimal utility upper bound.

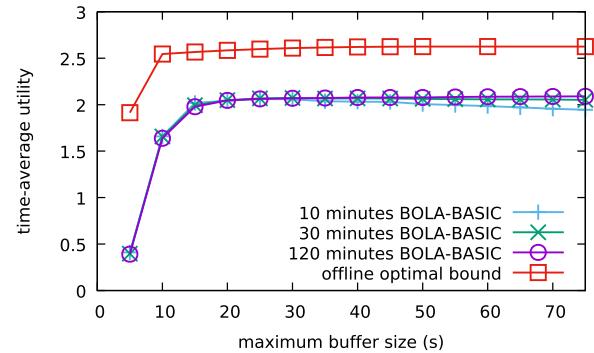


Fig. 5. Time-average utility for $\gamma p = 5$ using profile 1 for BOLA-BASIC.

C. Evaluating BOLA-BASIC

Fig. 5 shows the time-average utility of BOLA-BASIC when the video length is 10, 30 and 120 minutes. We set $\gamma p = 5$ and varied V for different buffer sizes. We compared the utility of BOLA-BASIC with the offline optimal bound described in Section V-B. The offline optimal gave nearly the same utility for the different video lengths. BOLA-BASIC only obtains about 80% of the offline optimal bound. Also, the utility of BOLA-BASIC decreases slightly when the buffer size is increased because it must download more lower-bitrate segments during startup before it can reach the buffer levels required to switch to higher-bitrate segments. Our results suggest that there is room to improve BOLA-BASIC that motivates our next version.

D. Adapting BOLA to Finite-Sized Videos

BOLA-BASIC was derived under the assumption that the videos are infinite. Thus, some adaptations are needed for BOLA to work effectively with smaller videos. Motivated by our initial experiments, we implemented two adaptations to BOLA-BASIC to derive a version we call BOLA-FINITE.

1) Dynamic V value for startup and wind down: A large buffer allows BOLA-BASIC to perform better but it has two drawbacks. First, it takes longer to prime a large buffer during startup. Lower bitrate segments are preferred until the buffer level reaches steady state. Second, at some late stage all downloads are complete and any remaining buffered video is played out. Any available bandwidth during this period is not utilized. Shortening this period would result in less unutilized available bandwidth. We mitigate these effects by introducing

```

1: for  $n$  in  $[1, N]$  do
2:    $t \leftarrow \min[\text{playtime from begin, playtime to end}]$ 
3:    $t' \leftarrow \max[t/2, 3p]$ 
4:    $Q_{\max}^D \leftarrow \min[Q_{\max}, t'/p]$ 
5:    $V^D \leftarrow (Q_{\max}^D - 1)/(v_M + \gamma p)$ 
6:    $m^*[n] \leftarrow \arg \max(V^D v_m + V^D \gamma p - Q)/S_m$ 
7:   if  $m^*[n] > m^*[n - 1]$  then
8:      $r \leftarrow \text{bandwidth measured when downloading segment } (n - 1)$ 
9:      $m' \leftarrow \max m \text{ such that } S_m/p \leq \max[r, S_1/p]$ 
10:    if  $m' \geq m^*[n]$  then
11:       $m' \leftarrow m^*[n]$ 
12:    else if  $m' < m^*[n - 1]$  then
13:       $m' \leftarrow m^*[n - 1]$ 
14:    else if some utility sacrificed for fewer oscillations then
15:      pause until  $(V^D v_{m'} + V^D \gamma p - Q)/S_{m'} \geq (V^D v_{m'+1} + V^D \gamma p - Q)/S_{m'+1}$   $\triangleright$  BOLA-O
16:    else
17:       $m' \leftarrow m' + 1$   $\triangleright$  BOLA-U
18:    end if
19:     $m^*[n] \leftarrow m'$ 
20:  end if
21:  pause for  $\max[p \cdot (Q - Q_{\max}^D + 1), 0]$ 
22:  download segment  $n$  at bitrate index  $m^*[n]$ , possibly abandoning
23: end for

```

Fig. 6. The BOLA Algorithm.

```

1: function SHALLABANDON( $m, S_m^R$ )
2:    $r_m \leftarrow (V^D v_m + V^D \gamma p - Q)/S_m^R$ 
3:    $r_{m'} \leftarrow (V^D v_{m'} + V^D \gamma p - Q)/S_{m'}^R$ 
4:   return true if  $r_{m'} > r_m$  for some  $m'$  subject to  $1 \leq m' < m$ 
5: end function

```

Fig. 7. BOLA-FINITE's Download Abandonment Heuristic: m is the current segment bitrate and S_m^R is the number of bits remaining to download in the current segment.

a dynamic V^D which corresponds to a dynamic buffer size Q_{\max}^D , shown in lines 2–5 in Fig. 6. BOLA-FINITE does not try to fill the whole buffer too soon and does not try to maintain a full buffer too long. We still need a minimum buffer size $3p$ for the algorithm to work effectively.

2) Download abandonment: BOLA-BASIC takes control decisions just before the download of each segment. Consider a scenario where the player is downloading high-bitrate 6 Mbps segments in good network conditions. The network bandwidth suddenly drops to 1 Mbps as the player has just started a new segment download. The segment will take $6p$ seconds to download, depleting the buffer and possibly causing rebuffering. BOLA-FINITE mitigates this problem by monitoring download progress and possibly abandoning a download. Fig. 7 shows how BOLA-FINITE decides whether or not to abandon the download. If a segment at bitrate index m is being downloaded, the remaining size S_m^R is less than S_m . The segment can be abandoned and downloaded at some bitrate index m' subject to $1 \leq m' < m$ when $(V^D v_m + V^D \gamma p - Q)/S_m^R < (V^D v_{m'} + V^D \gamma p - Q)/S_{m'}^R$. The control idea remains the same, but the current bitrate m has a smaller corresponding size S_m^R because part of the segment has already been downloaded. Fig. 3 illustrates a scenario where abandonment might help. At 46s a 3 Mbps segment download starts. Since there is a bandwidth drop at the time, the segment takes almost 9s to download. The buffer is depleted and BOLA-BASIC switches to downloading at a bitrate of 0.3 Mbps. BOLA-FINITE with abandonment logic would have detected the rapidly depleting buffer and stopped the long download, with the system only dropping to the 1.4 and 0.7 Mbps download bitrates in the low-bandwidth period.

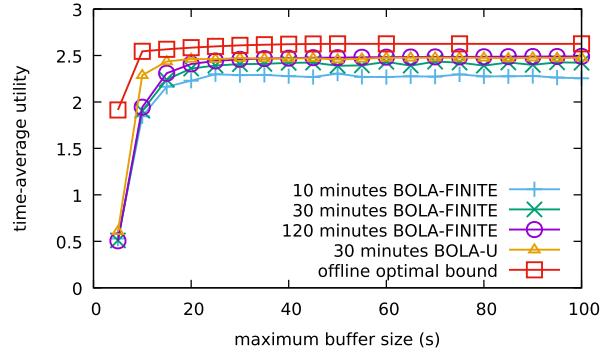
Fig. 8. Time-average utility for $\gamma p = 5$ using profile 1 for BOLA-FINITE and BOLA-U.

Fig. 8 shows the time-average utility of BOLA-FINITE for 10, 30 and 120 minutes of playback time with $\gamma p = 5$. Comparing with BOLA-BASIC in Fig. 5, we see that the time-average utility is much closer to the offline optimal bound. The benefit of the adjustments is also evident as the buffer grows larger, as there is no significant decrease in utility caused by filling the buffer with low-bitrate segments in the earlier stages of the video.

E. Avoiding Bitrate Oscillations

While our performance objective optimizes playback utility and playback smoothness, users are also sensitive to excessive bitrate switching. We discuss three causes of bitrate switches.

1) Bandwidth variation: As the network conditions change, the player varies the bitrate, tracking the network bandwidth. Such switches are acceptable; the player has no control on the bandwidth and should adapt to different network conditions.

2) Dense buffer thresholds: Either a larger number of bitrate levels and/or a smaller buffer size may push the threshold levels closer. If the differences between threshold levels are less than the segment duration p , adding one downloaded segment to the buffer may push the buffer level over several threshold levels at once. This might cause BOLA-FINITE to overshoot and choose a bitrate that is too high for the available bandwidth. Consequently, the segment download would take much more than p seconds, leading to excessive buffer depletion, causing BOLA-FINITE to switch down its bitrate by more than one level. In such a scenario BOLA-FINITE can oscillate between bitrates, even when the available bandwidth is stable.

3) Bitrate quantization: Having a stable network bandwidth and widely-spaced thresholds still does not avoid all bitrate switching. Suppose the bandwidth is 2.0 Mbps and it lies between two encoded bitrates of 1.5 and 3.0 Mbps. While the player downloads 1.5 Mbps segments, the buffer keeps growing. When the buffer crosses the threshold the player switches to 3.0 Mbps, depleting the buffer. After the buffer gets sufficiently depleted, the player switches back to 1.5 Mbps, and the cycle repeats. In this example, a viewer might prefer the video player to stick to the 1.5 Mbps bitrate, sacrificing some utility in order to have fewer oscillations. Or, a viewer might want to maximize utility and play a part of the video in the higher bitrate of 3.0 Mbps at the cost of more oscillations. We describe two variants of BOLA below to suit either viewer.

The first variant that we call BOLA-O mitigates oscillations by introducing bitrate capping (lines 7–20 in Fig. 6) when switching to a higher bitrate. BOLA-O verifies that the higher

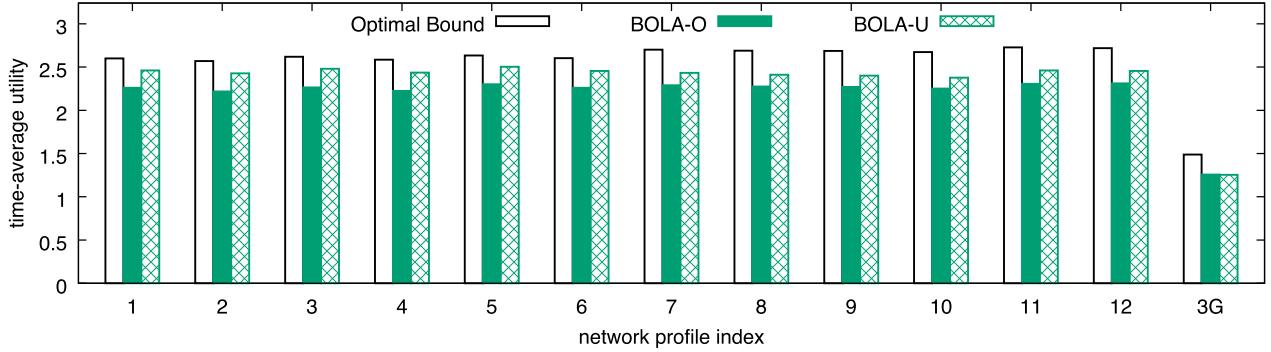


Fig. 9. The time-average utility of BOLA-O and BOLA-U with $\gamma p = 5$ and a 25-second buffer playing a 30-minute video for the DASH test network profiles 1–12 and mobile traces (3G). BOLA utility is within 84–95% of offline optimal utility.

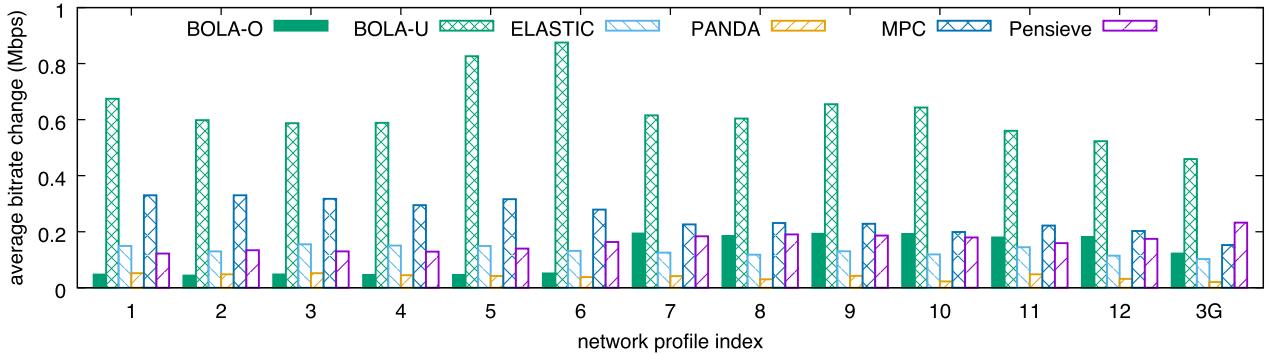


Fig. 10. The average bitrate change between adjacent segments was smaller for BOLA-O than for BOLA-U, but some bitrate change is needed to accurately track the network bandwidth. In our experiments, as an average across network profiles, ELASTIC and PANDA tracked the bandwidth with similar accuracy to BOLA-O, while MPC and Pensieve had more oscillations.

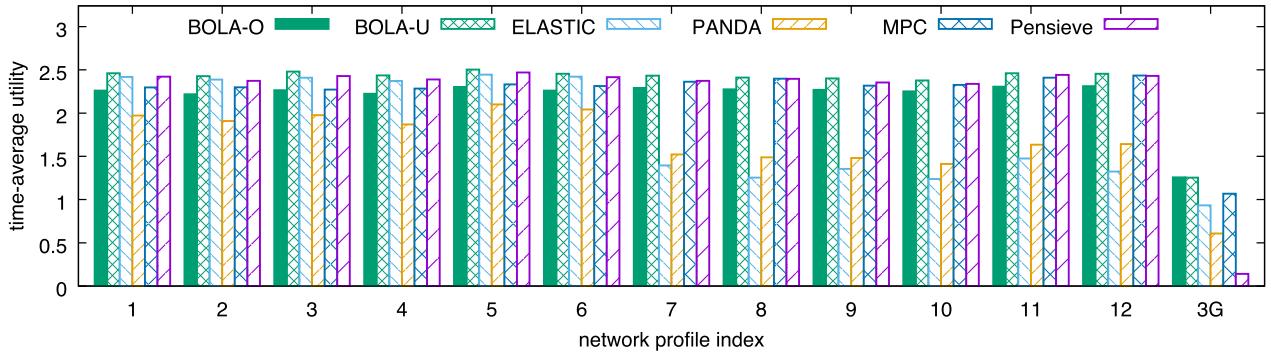


Fig. 11. The time-average utility of BOLA-O, BOLA-U, ELASTIC, PANDA, MPC and Pensieve with $\gamma p = 5$ playing a 30-minute video for the DASH test network profiles 1–12 and mobile traces (3G). Compared with ELASTIC and PANDA, BOLA-U has about 1.75 times the utility of the other algorithms in roughly half the cases. MPC has a utility between BOLA-O and BOLA-U. Pensieve has a utility between BOLA-O and BOLA-U for profiles 1–12 but performs worse for the mobile (3G) traces.

bitrate is sustainable by comparing it to the bandwidth as measured when downloading the previous segment (lines 8–11). Since the motive is to limit oscillations rather than to predict future bandwidth, this adaptation does not drop the bitrate to a lower level than in the previous download (lines 12–13). Continuous downloading at a bitrate lower than the bandwidth would cause the buffer to keep growing. BOLA-O avoids this by allowing the buffer to slip to the appropriate threshold before starting the download (line 15).

The second variant that we call BOLA-U does not sacrifice utility. Excessive buffer growth is avoided by allowing the

bitrate to be one level higher than the sustainable bandwidth (line 17). This allows the player to choose 3 Mbps in the example. While BOLA-U does not handle the third type of oscillations, it handles the more severe second type.

Looking back at Fig. 8, we see that the added stability of BOLA-U pays off when using a small buffer size and BOLA-U achieves a larger utility than BOLA-FINITE. Fig. 9 shows the time-average utility of BOLA-O and BOLA-U with $\gamma p = 5$ and $Q_{\max}p = 25$ s playing a 30-minute video. The utility lost by BOLA-O to avoid oscillations is clearly evident. In practice the lost utility is limited by the distance between encoded

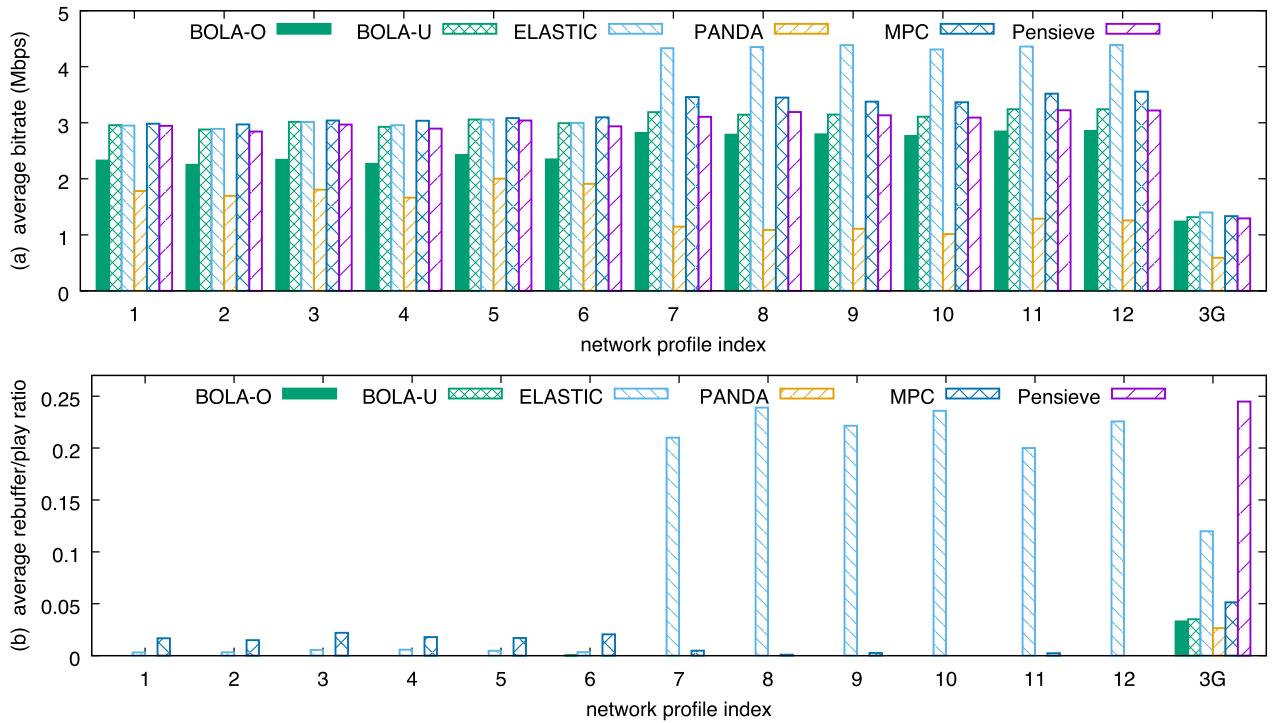


Fig. 12. Comparing BOLA with ELASTIC, PANDA, MPC and Pensieve using raw metrics: average bitrate and rebuffer-to-play ratio. BOLA, PANDA and Pensieve do not rebuffer for profiles 1–12. ELASTIC has almost no rebuffering for profiles 1–6, but it has a rebuffer-to-play ratio greater than 20% for profiles 7–12. MPC has some rebuffering for almost all profiles. Pensieve has no rebuffering for profiles 1–12. But, Pensieve has a 24% rebuffer-to-play ratio for the mobile (3G) traces, as it is unable to perform well for bandwidth conditions that are significantly different from its training set.

bitrates; if the next lower bitrate level is not far from the network bandwidth, then little utility will be lost.

We measure oscillations by comparing consecutive segments. The change in bitrate between a segment and the next is the absolute difference between bitrates (in Mbps) of the two segments. Fig. 10 shows the bitrate change averaged across all the segments. While BOLA-U has a high average bitrate change because of the quantization, BOLA-O only switches bitrate because of network bandwidth variations.

F. Comparison With State-of-the-Art Algorithms

We now compare BOLA with four state-of-the-art algorithms, ELASTIC [12], PANDA [13], MPC [14] and Pensieve [15]. We use the default design parameters in [12]–[15]. We test both BOLA-O and BOLA-U. Although BOLA performs better with larger buffers, we limited the buffer size to 25s for the tests to ensure fairness. ELASTIC targets a buffer level of 15s but the buffer level varies higher. PANDA targets a minimum buffer level of 26s. We use the RobustMPC variant of MPC with a buffer size of 25s. MPC relies on bandwidth estimation; we use the harmonic mean over the last five segment downloads to be consistent with the empirical evaluation method in [14]. We trained a Pensieve neural network model for the video with a buffer size of 25s. For training Pensieve, we used bandwidth traces generated using the tool provided in the Pensieve repository as recommended.

Fig. 11 compares the algorithms using each of the 12 network profiles and the mobile traces. BOLA-U consistently performs significantly better than PANDA. While BOLA-U and ELASTIC perform similarly for profiles 1–6, BOLA-U

performs significantly better for the other profiles that have larger bandwidth variations. MPC and Pensieve consistently obtains a utility between BOLA-O and BOLA-U for profiles 1–12, but perform worse for the mobile traces. We repeat the comparison using the average bitrate and rebuffering metrics in Fig. 12. This gives an insight into the strengths and weaknesses of the different algorithms.

Comparing BOLA-U with ELASTIC: For profiles 1–6, BOLA-U has approximately the same bitrate as ELASTIC. ELASTIC has a higher bitrate for profiles 7–12, but that comes at a significant cost in terms of rebuffering. For these profiles, the ratio of the rebuffering time to the play time is more than 20% for ELASTIC, while BOLA-U has no rebuffering. For the mobile traces, ELASTIC has marginally higher bitrate than BOLA-U but has a 12.0% rebuffer-to-play ratio compared with BOLA-U's 3.5%. ELASTIC rebuffers significantly more because it does not react in time when the bandwidth drops.

Comparing BOLA-U with PANDA: Both algorithms do not rebuffer for profiles 1–12. For the mobile traces, BOLA-U and PANDA have a rebuffer-to-play ratio of 3.5% and 2.6% respectively. However, PANDA has significantly lower bitrate than BOLA-U. The reason is that PANDA is more conservative and in some cases does not change to a higher bitrate even if it is sustainable.

Comparing BOLA-U with MPC: Both algorithms have similar average bitrates but MPC has slightly higher bitrate for some of the profiles. However, while BOLA-U does not rebuffer, MPC has some rebuffering for most of the profiles. While it is possible to tune the MPC parameters to avoid that rebuffering, it is not clear how to choose parameters that

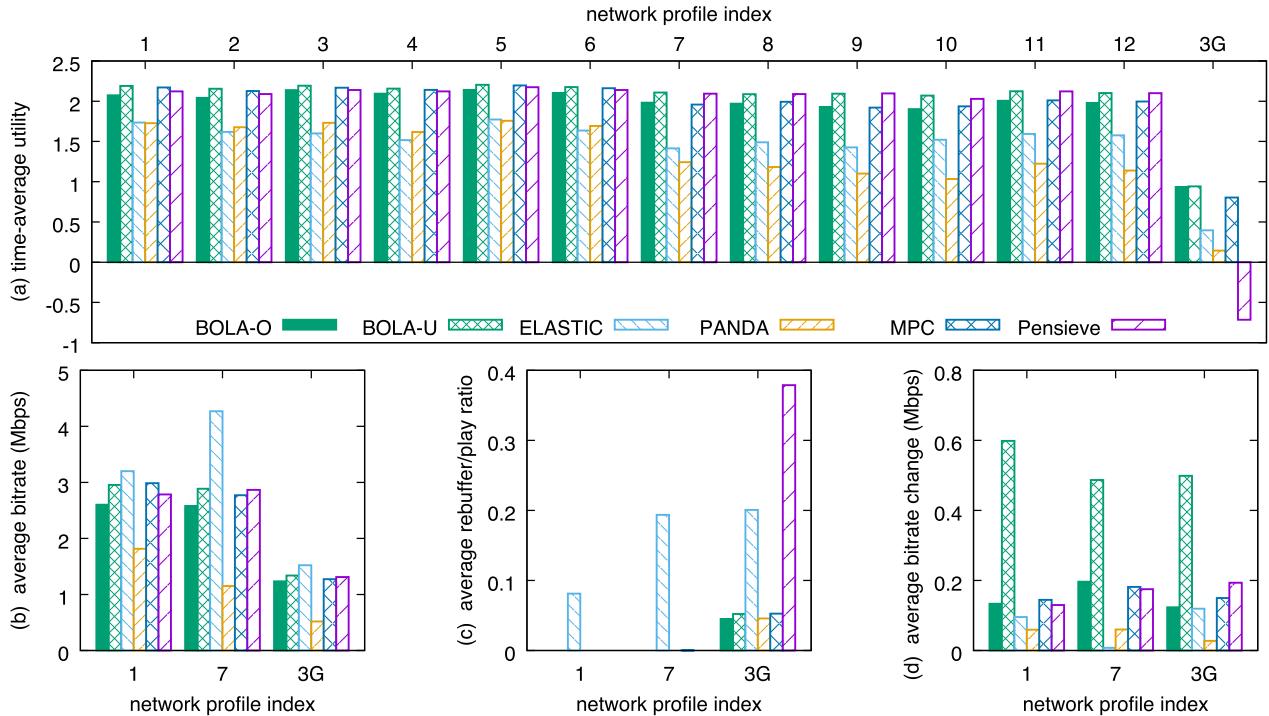


Fig. 13. The time-average utility of BOLA-O, BOLA-U, ELASTIC, PANDA, MPC and Pensieve with $\gamma p = 5$ playing a different video for 30 minutes, using the DASH test network profiles 1–12 and the mobile traces (3G). Note that Pensieve has negative 3G utility because of excessive rebuffing (the average rebuffer-to-play ratio is 38%). The raw metrics are also provided in the plots above.

consistently work for different network conditions. Another factor that might contribute to MPC rebuffing is the bandwidth estimation. When there is a large drop in bandwidth, the recommended harmonic mean bandwidth estimator takes a while to react. Even though RobustMPC factors in network estimation error, rebuffing is not totally eliminated.

Comparing BOLA-U with Pensieve: For profiles 1–12, Pensieve obtains utility between BOLA-O and BOLA-U, but consistently closer to BOLA-U. However, Pensieve has too much rebuffing in the mobile traces, resulting in much worse utility for these traces. While the network traces used to train Pensieve included periods with low bandwidth similar to the mobile traces, Pensieve did not learn a model that would perform well in relatively low bandwidth situations in a mobile setting. This points to a weakness in Pensieve as it is unable to adapt to bandwidth conditions that are significantly different from the training set.

In Fig. 10 we show our results for our secondary metric of bitrate oscillations. BOLA-U does not perform well in this metric, since it attempts to maximize utility at the cost of increased oscillations. Comparing BOLA-O with ELASTIC, PANDA, MPC, and Pensieve, ELASTIC has a lower average change than BOLA-O only in the cases where it has a slow reaction and excessive rebuffing. PANDA has a lower average change because it is more conservative and in some cases does not change to a higher bitrate even if that bitrate is sustainable. MPC has higher average change than BOLA-O for profiles 1–12. Pensieve has similar average change to BOLA-O for profiles 1–12.

We also tested the algorithms with more videos to investigate performance when changing characteristics such as content type, segment duration, and available bitrates. The tests

showed similar results. One example is the video provided with Pensieve. The video has 49 segments with a segment duration of 4s. It is encoded at six bitrates: 0.3, 0.75, 1.2, 1.85, 2.85 and 4.3 Mbps. Fig. 13 shows the utility and metrics for the same six algorithms with similar conclusions to Figs. 10–12. Note that Pensieve fails to perform well on mobile traces again, since it is significantly different from its training set.

Thus, from our empirical analysis, we can conclude that BOLA achieves higher utility, and performs more consistently across different scenarios in comparison with ELASTIC, PANDA, MPC and Pensieve. One reason for the consistency of BOLA is that it does not have a large number of parameters. BOLA has two design parameters γ and V , which have an intuitive significance as discussed in Section IV-B, and an option of whether or not to trade off some utility to reduce oscillations. Other algorithms have a number of different parameters and tuning the parameters for a particular scenario might make the system less suited for other scenarios. Also, BOLA's ability to abandon a segment during a download and start the download at a lower bitrate allows BOLA to achieve significantly less rebuffing than the other algorithms.

Note: A number of recent papers compare new algorithms with BOLA. Some of the prior work used the experimental version of BOLA in dash.js versions 2.0.0–2.5.0 that required bug fixes. We suggest that a stable version used in production (dash.js version 2.6.0 or later) be used for such comparisons for a more accurate evaluation of BOLA. Also, dash.js has a default buffer size of 12s, leading some researchers to compare a large-buffer algorithm with a small-buffer BOLA. Our work uses the correct BOLA implementation and the same buffer size for all the compared algorithms.

VI. DEPLOYMENT

A. The DASH Reference Player

After developing a theoretical foundation for BOLA and testing it by simulation, we deployed BOLA in a production setting. Particularly, we implemented BOLA in dash.js, the open-source standard DASH reference player [18]. Through dash.js, BOLA is now being used in production by several major video providers and delivery networks such as Akamai, BBC, CBS and Orange. Deployment in production presented a number of new challenges such as operating with even smaller buffer capacities, correctly handling events such as a user seek to a different point in the video, and tolerating delays caused by the video player unrelated to the network conditions. The techniques we implemented to handle these new challenges are described in [19].

B. BOLA Parameters

One deployment challenge involves choosing the BOLA parameters γ and V . We gave an intuition to pick the parameters in Section IV-B, but video providers are more familiar with choosing buffer level targets. For this purpose, we now discuss how to derive γ and V from intuitive requirements. Consider the following requirements:

1. We want a maximum buffer level Q_{\max} .
2. We want to download at the highest bitrate when the buffer level is Q_{\max} .
3. We want to download at the lowest bitrate when the buffer level is less than a threshold Q_{low} , and we want to download at a higher bitrate when the buffer level goes above the threshold. These requirements are easy to understand for video providers who might not be familiar with BOLA. In fact, video providers usually have some preferred maximum buffer level Q_{\max} . Further, they might have preference for Q_{low} such as 10s as described in [19].

To satisfy requirements 1–2, we want (11) to switch from choosing $a_M = 1$ to choosing $\sum a_m = 0$ at the threshold when the buffer level is Q_{\max} . This happens if

$$\frac{\rho_{a_M=1} - \rho_{a=0}}{V(v_M + \gamma p) - Q_{\max}} = 0 \quad (13)$$

Note that BOLA satisfies requirement 2 and downloads at the highest bitrate just before the Q_{\max} threshold because at that buffer level we get $\rho_{a_m=1} < \rho_{a_M=1}$ for $m < M$. This is illustrated in Fig. 1.

To satisfy requirement 3, we want (11) to switch from choosing $a_1 = 1$ to choosing $a_2 = 1$ at the threshold when the buffer level is Q_{low} . This happens if

$$\frac{\rho_{a_1=1} - \rho_{a_2=1}}{V(v_1 + \gamma p) - Q_{\text{low}}} = \frac{V(v_2 + \gamma p) - Q_{\text{low}}}{S_2} \quad (14)$$

Solving (13)–(14), we obtain

$$V = \frac{Q_{\max} - Q_{\text{low}}}{v_M - \alpha}, \quad \gamma p = \frac{v_M Q_{\text{low}} - \alpha Q_{\max}}{Q_{\max} - Q_{\text{low}}}$$

where

$$\alpha = \frac{S_1 v_2 - S_2 v_1}{S_2 - S_1}.$$

When calculating the BOLA parameters from Q_{low} and Q_{\max} , the previous intuition about γ and V still hold. If a video provider chooses a larger Q_{low} , γ will be larger and BOLA will give more weight to rebuffering. If a video provider chooses a larger Q_{\max} , V will be larger.

VII. RELATED WORK

There has been a lot of recent work on bitrate adaptation algorithms, much of which is based on estimating the bandwidth of the network connection. FESTIVE [10] uses a harmonic bandwidth estimator to predict future bandwidth from past downloads, limiting bitrate change to one level between successive segments for stability. Notably, FESTIVE attempts to find a tradeoff between efficiency and fairness with competing downloads. BBA [11] is a buffer-based algorithm. BOLA has a few similarities to BBA but the mapping function from buffer level to video bitrate is different. Also, BBA assumes that the buffer size is large (in the order of minutes), thereby making it not suitable for short videos. Further, it does not provide any theoretical guarantees for its buffer-based approach. A notable algorithm is ELASTIC [12] that uses control theory to adjust the bitrate so as to keep the buffer occupancy at a constant level. Another notable algorithm is PANDA [13] which also estimates the network bandwidth. PANDA drops the download bitrate as soon as low bandwidth is detected but only increases the bitrate slowly to probe the real capacity when a higher bandwidth is detected. Like FESTIVE, PANDA trades efficiency for fairness. In [14], an algorithm using model predictive control (MPC) is proposed to optimize a comprehensive set of metrics. In this approach, the bitrate for the current segment is chosen based on a network bandwidth prediction for the next few segments. But, its performance depends on the accuracy of such a prediction. The approach also requires significant offline optimization to be performed outside of the client for an exhaustive set of scenarios. Reference [15] presents Pensieve, a reinforcement-learning approach to ABR. A neural network model can be trained for a video using a particular buffer size, using a QoE function for reward. A set of bandwidth traces is used as training data. Unfortunately, a trained model does not transfer easily to a different video or, more importantly, to bandwidth conditions not represented in the training data. Unlike prior work, we derive a buffer-based algorithm with theoretical guarantees that is simple to implement within the client and we empirically show its efficacy on extensive network traces. In recent work [23], a method called Oboe for auto-tuning the parameters of BOLA and MPC was presented and shown to improve both algorithms. Further, the work showed that Oboe used in conjunction with traditional ABR algorithms performs better than reinforcement-learning based ABR such as Pensieve.

VIII. CONCLUSION

We formulated video bitrate adaptation for ABR streaming as a utility maximization problem and derived BOLA, an online control algorithm that is provably near-optimal. Further, we empirically demonstrated the efficacy of BOLA using extensive traces. In particular, we showed that our online algorithm achieves utility close to the optimal offline algorithm.

We showed that our algorithm performs better than state-of-the-art algorithms in a number of different test scenarios. We also implemented BOLA in dash.js, the open-source standard DASH reference player [18]. Through dash.js, BOLA is now being used in production by several major video providers and delivery networks such as Akamai, BBC, CBS and Orange.

ACKNOWLEDGMENT

The authors would like to thank Daniel Sparacio and Will Law of Akamai for their key insights on real-world player implementations. Further, Daniel was instrumental in helping them implement BOLA in the DASH reference player.

APPENDIX PROOF OF THEOREM 1

We first show part 1 using induction. Note that the bound $Q(t_k) \leq V(v_M + \gamma p) + 1$ holds for $k = 1$ since $Q(t_1) = Q(0) = 0$. Now suppose it holds for some k . We will show that it will also hold for $k + 1$. We have two cases.

Case 1: $Q(t_k) \leq V(v_M + \gamma p)$

From the queueing equation (4), it follows that the maximum that $Q(t_k)$ can increase in slot k is by 1. This implies that $Q(t_{k+1}) \leq V(v_M + \gamma p) + 1$.

Case 2: $V(v_M + \gamma p) < Q(t_k) \leq V(v_M + \gamma p) + 1$

We have $Q(t_k) > V(v_m + \gamma p)$ for all $m \in \{1, 2, \dots, M\}$ (using (1)). It follows from the structure of optimal solution to (11) that BOLA will choose the no-download option in this case. As a result, $Q(t_k)$ cannot increase and we have that $Q(t_{k+1}) \leq V(v_M + \gamma p) + 1$.

$Q(t_k)$ denotes the total number of segments in the buffer. This can be at most Q_{\max} using the relation

$$V \leq \frac{Q_{\max} - 1}{v_M + \gamma p}.$$

In part 2, we show the bound in (12) using the technique of Lyapounov optimization over variable size frames [21]. We first define a Lyapunov function $L(Q(t_k))$ as

$$L(Q(t_k)) = \frac{1}{2}Q^2(t_k)$$

and define the per-slot conditional Lyapunov drift $D(t_k)$ as

$$D(t_k) \triangleq \mathbb{E}\{L(Q(t_{k+1})) - L(Q(t_k))|Q(t_k)\}.$$

We use the queueing equation (4), to bound $D(t_k)$. We consider two cases for (4): $Q(t_k) \leq T_k/p$ and $Q(t_k) > T_k/p$. In the first case we have

$$D(t_k) = \mathbb{E}\left\{\frac{1}{2}\left(\sum_{m=1}^M a_m(t_k)\right)^2 - \frac{1}{2}Q^2(t_k)|Q(t_k)\right\}.$$

In the second case we have

$$D(t_k) = \mathbb{E}\left\{\frac{1}{2}\left(\frac{T_k}{p} - \sum_{m=1}^M a_m(t_k)\right)^2 - Q(t_k)\left(\frac{T_k}{p} - \sum_{m=1}^M a_m(t_k)\right)|Q(t_k)\right\}$$

In both cases, $D(t_k)$ is bounded by

$$D(t_k) \leq \frac{p^2 + \Psi}{2p^2} - Q(t_k)\mathbb{E}\left\{\frac{T_k}{p} - \sum_{m=1}^M a_m(t_k)|Q(t_k)\right\} \quad (15)$$

where Ψ is an upper bound on $\mathbb{E}\{T_k^2\}$ under any control algorithm and is assumed to be finite.

Following the methodology of the Lyapunov optimization technique, we subtract $V \times$ reward term from both sides of the above to get

$$\begin{aligned} D(t_k) - V\mathbb{E}\left\{\sum_{m=1}^M a_m(t_k)(v_m + \gamma p)|Q(t_k)\right\} \\ \leq \frac{p^2 + \Psi}{2p^2} - Q(t_k)\mathbb{E}\left\{\frac{T_k}{p} - \sum_{m=1}^M a_m(t_k)|Q(t_k)\right\} \\ - V\mathbb{E}\left\{\sum_{m=1}^M a_m(t_k)(v_m + \gamma p)|Q(t_k)\right\} \end{aligned} \quad (16)$$

Let us denote the control decisions (and resulting slot lengths) under our control algorithm by the superscript BOLA while those under the stationary policy of Lemma 1 by STAT. Since BOLA greedily maximizes over a frame, it ensures that

$$\begin{aligned} \mathbb{E}\left\{\sum_{m=1}^M a_m^{\text{BOLA}}(t_k)(Q(t_k) - V(v_m + \gamma p))|Q(t_k)\right\} \\ \leq \eta \times \mathbb{E}\left\{\sum_{m=1}^M a_m^{\text{STAT}}(t_k)(Q(t_k) - V(v_m + \gamma p))|Q(t_k)\right\} \end{aligned} \quad (17)$$

where $\eta = \frac{\mathbb{E}\{T_k^{\text{BOLA}}|Q(t_k)\}}{\mathbb{E}\{T_k^{\text{STAT}}|Q(t_k)\}}$. To see this, compare the ratio on the left hand side above with the objective in (11) while noting that we can express the denominator as $\mathbb{E}\{T_k^{\text{BOLA}}|Q(t_k)\} = (\sum_{m=1}^M a_m^{\text{BOLA}}(t_k)S_m)/\omega_{\text{avg}}$. It should be noted that this ratio can be minimized without requiring knowledge of ω_{avg} . Then we use (17) to express (16) as

$$\begin{aligned} D^{\text{BOLA}}(t_k) - V\mathbb{E}\left\{\sum_{m=1}^M a_m^{\text{BOLA}}(t_k)(v_m + \gamma p)|Q(t_k)\right\} \\ \leq \frac{p^2 + \Psi}{2p^2} - Q(t_k)\mathbb{E}\left\{\frac{T_k^{\text{BOLA}}}{p} - \eta \sum_{m=1}^M a_m^{\text{STAT}}(t_k)|Q(t_k)\right\} \\ - V\eta\mathbb{E}\left\{\sum_{m=1}^M a_m^{\text{STAT}}(t_k)(v_m + \gamma p)|Q(t_k)\right\} \end{aligned}$$

Substituting the time-average values for the stationary policy we get

$$\begin{aligned} D^{\text{BOLA}}(t_k) - V\mathbb{E}\left\{\sum_{m=1}^M a_m^{\text{BOLA}}(t_k)(v_m + \gamma p)|Q(t_k)\right\} \\ \leq \frac{p^2 + \Psi}{2p^2} - Q(t_k)\left(\frac{1}{p} - r^{\text{STAT}}\right)\mathbb{E}\{T_k^{\text{BOLA}}|Q(t_k)\} \\ - V(v^* + \gamma s^*)\mathbb{E}\{T_k^{\text{BOLA}}|Q(t_k)\} \end{aligned} \quad (18)$$

where r^{STAT} denotes the expected arrival rate under the stationary policy and cannot exceed $1/p$ since it is rate stable.

Thus we have

$$D^{\text{BOLA}}(t_k) - V\mathbb{E} \left\{ \sum_{m=1}^M a_m^{\text{BOLA}}(t_k)(v_m + \gamma p) | Q(t_k) \right\} \leq \frac{p^2 + \Psi}{2p^2} - V(v^* + \gamma s^*)\mathbb{E} \{ T_k^{\text{BOLA}} | Q(t_k) \} \quad (19)$$

Taking conditional expectation of both sides and summing over $k \in \{1, 2, \dots, K_N\}$, we get

$$\mathbb{E} \{ L(Q(t_{K_N+1})) \} - V\mathbb{E} \left\{ \sum_{k=1}^{K_N} \sum_{m=1}^M a_m^{\text{BOLA}}(t_k)(v_m + \gamma p) \right\} \leq \frac{(p^2 + \Psi)K_N}{2p^2} - V(v^* + \gamma s^*)\mathbb{E} \left\{ \sum_{k=1}^{K_N} T_k^{\text{BOLA}} \right\} \quad (20)$$

Dividing both sides by $V\mathbb{E} \left\{ \sum_{k=1}^{K_N} T_k^{\text{BOLA}} \right\}$ and taking the limit as $N \rightarrow \infty$ yields the bound in (12).

REFERENCES

- [1] K. Spiteri, R. Urgaonkar, and R. K. Sitaraman, "BOLA: Near-optimal bitrate adaptation for online videos," in *Proc. 35th Annu. IEEE Int. Conf. Comput. Commun. (INFOCOM)*, Apr. 2016, pp. 1–9.
- [2] Cisco. (2019). *Cisco Visual Networking Index: Forecast and Trends, 2017–2022 White Paper*. [Online]. Available: <https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white-paper-c11-741490.html>
- [3] F. Dobrian *et al.*, "Understanding the impact of video quality on user engagement," in *Proc. ACM SIGCOMM Conf.*, 2011, pp. 91–99.
- [4] S. S. Krishnan and R. K. Sitaraman, "Video stream quality impacts viewer behavior: Inferring causality using quasi-experimental designs," in *Proc. ACM Conf. Internet Meas. Conf. (IMC)*, 2012, pp. 211–224.
- [5] R. K. Sitaraman, "Network performance: Does it really matter to users and by how much?" in *Proc. COMSNETS*, 2013, pp. 1–10.
- [6] Apple HTTP Live Streaming. Accessed: Sep. 25, 2014. [Online]. Available: <https://developer.apple.com/resources/http-streaming/>
- [7] Microsoft Smooth Streaming. Accessed: Sep. 25, 2014. [Online]. Available: <http://www.iis.net/downloads/microsoft/smooth-streaming>
- [8] Adobe HTTP Dynamic Streaming. Accessed: Sep. 25, 2014. [Online]. Available: <http://www.adobe.com/products/hds-dynamic-streaming.html>
- [9] T. Stockhammer, "Dynamic adaptive streaming over HTTP—standards and design principles," in *Proc. ACM MMSys*, 2011, pp. 133–144.
- [10] J. Jiang, V. Sekar, and H. Zhang, "Improving fairness, efficiency, and stability in HTTP-based adaptive video streaming with festive," *IEEE/ACM Trans. Netw.*, vol. 22, no. 1, pp. 326–340, Feb. 2014.
- [11] T.-Y. Huang, R. Johari, N. McKeown, M. Trunnell, and M. Watson, "A buffer-based approach to rate adaptation: Evidence from a large video streaming service," in *Proc. ACM Conf. SIGCOMM*, 2014, pp. 187–198.
- [12] L. De Cicco, V. Caldaralo, V. Palmisano, and S. Mascolo, "Elastic: A client-side controller for dynamic adaptive streaming over HTTP (DASH)," in *Proc. Packet Video Workshop (PV)*, 2013, pp. 1–8.
- [13] Z. Li *et al.*, "Probe and adapt: Rate adaptation for HTTP video streaming at scale," *IEEE J. Sel. Areas Commun.*, vol. 32, no. 4, pp. 719–733, Apr. 2014.
- [14] X. Yin, A. Jindal, V. Sekar, and B. Sinopoli, "A control-theoretic approach for dynamic adaptive video streaming over HTTP," in *Proc. ACM SIGCOMM*, 2015, pp. 325–338.
- [15] H. Mao, R. Netravali, and M. Alizadeh, "Neural adaptive video streaming with pensieve," in *Proc. Conf. ACM Special Interest Group Data Commun.*, Aug. 2017, pp. 197–210.
- [16] DASH Industry Forum. (Jan. 2014). *Guidelines for Implementation: DASH-AVC/264 Test Cases and Vectors*. [Online]. Available: <http://dashif.org/guidelines/>
- [17] H. Riiser, P. Vigmstad, C. Griwodz, and P. Halvorsen, "Commute path bandwidth traces from 3G networks: Analysis and applications," in *Proc. 4th ACM Multimedia Syst. Conf. (MMSys)*, 2013, pp. 114–118.
- [18] DASH Reference Client. Accessed: Jun. 28, 2019. [Online]. Available: <https://reference.dashif.org/dash.js/>
- [19] K. Spiteri, R. Sitaraman, and D. Sparacio, "From theory to practice: Improving bitrate adaptation in the DASH reference player," *ACM Trans. Multimedia Comput. Commun. Appl.*, vol. 15, no. 2s, p. 67:1–67:29, 2019.
- [20] D. Bertsekas, *Dynamic Programming and Optimal Control*, vol. 1. Belmont, MA, USA: Athena Scientific, 1995.
- [21] M. J. Neely, "Stochastic network optimization with application to communication and queueing systems," *Synth. Lectures Commun. Netw.*, vol. 3, no. 1, pp. 1–211, Jan. 2010.
- [22] M. J. Neely, "Dynamic optimization and learning for renewal systems," *IEEE Trans. Autom. Control*, vol. 58, no. 1, pp. 32–46, Jan. 2013.
- [23] Z. Akhtar *et al.*, "Oboe: Auto-tuning video ABR algorithms to network conditions," in *Proc. Conf. ACM Special Interest Group Data Commun.*, Aug. 2018, pp. 44–58.
- [24] *Big Buck Bunny Movie*. Accessed: Jul. 31, 2015. [Online]. Available: <https://peach.blender.org/>
- [25] P. Romirer-Maierhofer, F. Ricciato, A. D'Alconzo, R. Franzan, and W. Karner, "Network-wide measurements of TCP RTT in 3G," in *Traffic Monitoring and Analysis*. Berlin, Germany: Springer, 2009, pp. 17–25.



Kevin Spiteri (Member, IEEE) received the B.Eng. degree (Hons.) in electrical engineering from the University of Malta, Msida, Malta, and the M.S. degree in electrical and computer engineering from Oakland University, Rochester, MI, USA. He is currently pursuing the Ph.D. degree with the College of Information and Computer Sciences, University of Massachusetts at Amherst, Amherst, MA, USA. His research interest includes algorithms for video delivery over the Internet.



Rahul Urgaonkar (Senior Member, IEEE) received the bachelor's degree from IIT Bombay and the master's and Ph.D. degrees from the University of Southern California, all in electrical engineering. He was with IBM Research, where he was a Task Leader with the U.S. Army Research Laboratory (ARL) funded Network Science Collaborative Technology Alliance (NS CTA) Program. He was also a Primary Researcher with the U.S./U.K. International Technology Alliance (ITA) research programs. He is currently a Senior Research Scientist with Amazon Prime Video (PV), where he works on optimizing PV's video delivery systems. His research interests include stochastic optimization, algorithm design and control theory with applications to communication networks, and cloud-computing systems.



Ramesh K. Sitaraman (Fellow, IEEE) received the B.Tech. degree from IIT Madras, Chennai, and the Ph.D. degree in computer science from Princeton University. He is currently a Professor with the College of Information and Computer Sciences, University of Massachusetts at Amherst. His research focuses on Internet-scale distributed systems, including algorithms, architectures, performance, energy efficiency, security, and economics. As a principal architect, he helped create the Akamai Content Delivery Network (CDN), the world's first major CDN that currently delivers a significant fraction of the Internet traffic. He retains a part-time role as the Akamai's Chief Consulting Scientist. He is a Fellow of the ACM. He was a recipient of the inaugural ACM SIGCOMM Networking Systems Award for his work on the Akamai CDN, the DASH-IF Excellence in DASH Award for his work on ABR algorithms, the NSF CAREER Award, the College of Natural Sciences Outstanding Teacher Award, and the UMass Distinguished Teaching Award.