

Optimal Rate Allocation for Adaptive Wireless Video Streaming in Networks with User Dynamics

Vinay Joseph[†], Sem Borst^{*}, Martin I. Reiman^{*}

[†]University of Texas at Austin, Austin, TX 78712

^{*}Alcatel-Lucent Bell Labs, Murray Hill, NJ 07974

Abstract—We consider the problem of optimal rate allocation and admission control for adaptive video streaming sessions in wireless networks with user dynamics. The central aim is to achieve an optimal tradeoff between several key objectives: maximizing the average rate utility per user, minimizing the temporal rate variability, and maximizing the number of users supported. We identify the structure of algorithms that achieve asymptotically optimal performance in large-capacity systems, and exploit the insight into this structure to devise parsimonious and robust online algorithms. Extensive simulation experiments demonstrate that the proposed online algorithms perform well, even in systems with relatively small capacity.

I. INTRODUCTION

Video traffic is experiencing tremendous growth, fueled by the proliferation of online video content and the steady expansion in transmission bandwidths. The amount of video traffic is forecast to double annually in the next several years, and is expected to account for the dominant share of wireline as well as wireless Internet traffic soon [7].

The huge growth in traffic volume goes hand in hand with a shift towards HTTP-based adaptive streaming mechanisms, which allow the video rate to be adjusted to the available network bandwidth. Thus, it is critical to design video streaming mechanisms that use the available network resources efficiently and provide an adequate quality-of-experience (QoE) to the users. While a comprehensive video quality perception metric is hard to define, it is widely agreed that the QoE improves not only with a higher time-average rate, but also with a smoother temporal evolution [9], [11], [22], [28]. Typically a fundamental trade-off arises between these two criteria, because a higher time-average rate entails a more responsive scheme that makes more aggressive rate adjustments, compromising smoothness. Measurement experiments indicate that currently deployed schemes do not necessarily perform well in that regard, and may induce high variability and even oscillations as a result of interactions among several rate-adaptive users [2], [3], [8], [15]. This has triggered numerous proposals for enhancements, see for instance [1], [14], [15], [18], [21], [23], [26], involving a variety of techniques, ranging from multi-path solutions, network caching, traffic shaping and layered coding to improved bandwidth estimation at the client side in conjunction with dynamic rate selection at server side.

The above challenges are particularly pertinent in wireless networks where the available bandwidth is not only relatively

limited, but also inherently uncertain and time-varying due to fading and user mobility [24]. On the other hand, resource allocation mechanisms in cellular networks offer greater capabilities for controlling the user throughput, especially since the wireless link is commonly the bottleneck in the end-to-end path. In particular, scheduling algorithms at the base stations tend to provide support for specifying *suitable target throughputs* for the various users, and thus ensuring a *smooth streaming rate* [10], [27]. The proposed enhancements of HTTP-based adaptive streaming mechanisms mentioned above have exploited a wide range of approaches for mitigating the impact of variations in network bandwidth at the client and/or server, but have *not* leveraged the option to exercise direct control over the user throughput at a network element.

Motivated by the above considerations, we examine in the present paper the problem of how to set target rates for various video users in a wireless network subject to given feasibility constraints. The objective is to achieve an optimal trade-off between maximizing the aggregate time-average rate utility and minimizing the rate variability over time.

As an inherent characteristic of wireless networks, we allow the users to have heterogeneous channel conditions, and hence different resource requirements for a given bit rate. Since the key aim of the target rates is to ensure smoothness over longer time scales, we do not explicitly account for fast variations in channel conditions, and focus on the impact of slower dynamics arising due to the arrivals (service requests) and departures (service completions) of users. We assume that the target rates for the various users can be adapted at any point in time based on the user dynamics, and also incorporate admission control, by allowing a user to be assigned a zero target rate upon arrival.

We first establish a sample path upper bound for the objective function (the aggregate average utility reduced by a rate variation penalty). The upper bound is characterized through the optimal value of a relatively simple linear program. We then proceed to demonstrate that the upper bound is asymptotically achievable (in a ratio sense) in a regime where the offered traffic and capacity grow large in proportion. Asymptotic optimality is achieved by ‘static’ policies which either deny service to a user upon arrival or assign it a fixed target rate (in accordance with the optimal solution of the above-mentioned linear program), without requiring advance knowledge of the duration.

The asymptotic optimality of static rate assignment policies

provides a natural paradigm for the design of sound, lightweight approaches, only involving the calculation of the asymptotically optimal rate assignments. The latter task is elementary when all the relevant system parameters are known, but in practice several parameters will be subject to uncertainty and variation over time, raising all sorts of estimation challenges and accuracy issues. Instead, we therefore exploit the fact that the asymptotically optimal rate assignment policy exhibits a specific structure, allowing it to be represented through just a single variable, and has the further property that the load induced on the system is unity. Leveraging these two properties, we develop an approach for learning and tracking the asymptotically optimal rate assignments by adapting that single variable so as to drive the observed load to unity. The proposed approach only relies on observations of the occupancy state, without any explicit knowledge or estimates of the system parameters. Numerical experiments demonstrate that the approach performs well, even in scenarios with a relatively moderate amount of capacity and offered traffic.

Somewhat related problems have been considered in several recent papers, see for instance [4], [13], [16], [17], [19], [25]. While these papers account for deadlines associated with real-time streaming ([13], [17], [25]), general video quality models ([19]), general network settings ([4]) etc, the impact of user dynamics is ignored. A further recent paper [5] does consider user dynamics in evaluating the performance of real-time video streaming. However, it simply assumes that users request a fixed rate, leaving no flexibility for video rate adaptation. The restrictive scenarios in these studies reflect the difficulty of capturing the interaction between user dynamics and lower-level scheduling and adaptation mechanism. Yet, accounting for user dynamics is vital for algorithm design and performance evaluation, as demonstrated by an extensive body of work on user-level models of channel-aware scheduling policies and wireline bandwidth-sharing networks.

Finally, it is worth reiterating that setting target rates at a base station is commonly supported by existing wireless scheduling algorithms, and complementary with bandwidth estimation algorithms at the client and dynamic rate selection at the server. In fact, the resulting throughput smoothing makes the latter tasks substantially lighter, and the joint control reduces the potential for oscillations and harmful interactions among separately controlled competing adaptive players.

The remainder of the paper is organized as follows. In Section II we present a detailed model description and formal problem formulation. We derive upper bounds for the maximum QoE rate in Section III, and show in Section IV that these bounds are asymptotically achievable. In Section V we provide insights in the structure of the asymptotically optimal rate assignments, which are leveraged in Section VI for the design of adaptive, online algorithms. In Section VII we present the simulation experiments that we have conducted to examine the performance and corroborate the convergence of the proposed algorithms. In Section VIII we make brief concluding remarks and offer a few suggestions for further research.

II. MODEL DESCRIPTION AND PROBLEM FORMULATION

We consider a single wireless downlink shared by streaming users from a finite set of classes indexed by $s \in \mathcal{S}$. The various classes could reflect heterogeneity among users and video objects in terms of wireless channel characteristics, available compression rates, or Quality Rate (QR) dependencies. Let τ_i be the arrival time of the i -th user, i.e., the time when the user requests a particular video, with $\tau_i \geq 0$. Denote by $s(i)$ the index of the class to which the i -th user belongs. Let σ_i be the service time of the i -th user, i.e., the nominal duration of the video of interest.

Denote by $r_i(t)$ the rate assigned to the i -th user at time $t \geq 0$. As alluded to in the introduction, it is worth emphasizing that $r_i(t)$ should *not* be thought of as an instantaneous transmission rate. Rather, it represents a target for the average throughput over a time scale which is long compared to the slot level on which wireless schedulers operate, and commensurate with the buffering depth at the video client. The value of $r_i(t)$ only has significance for $t \in [\tau_i, \tau_i + \sigma_i]$, but for notational convenience we assume the value to be defined for all $t \geq 0$. The rate assignments for class- s users are constrained to a finite set \mathcal{R}_s :

$$r_i(t) \in \mathcal{R}_{s(i)} \cup \{0\} \quad \text{for all } t \geq 0, \quad (1)$$

which could correspond to the various compression rates available for videos streamed by class- s users. We do not explicitly impose the constraint $r_i(t) = 0$ for $t \notin [\tau_i, \tau_i + \sigma_i]$, since this will not be directly relevant in the utility optimization problem that we consider. The capacity requirement per unit bit rate for class- s users is c_s , which would depend on channel quality characteristics like the Signal-to-Interference-and-Noise Ratio. Assuming the total capacity to be normalized to unity, the rate assignments should then satisfy the further constraint:

$$\sum_{i=1}^{\infty} c_{s(i)} r_i(t) \leq 1 \quad \text{for all } t \geq 0. \quad (2)$$

The quantity $p_s = 1/c_s$ may be interpreted as the transmission rate of a class- s user if it were assigned the full resources.

As observed in the introduction, the QoE enjoyed by a streaming user depends on the time-average rate utility as well as the temporal variability, and we aim to achieve an optimal trade-off between these two criteria. We will specifically assume that the QoE metric for the i -th user during the time interval $[0, T]$, $\tau_i \leq T$, can be expressed as $A_i(0, T) - \gamma V_i(0, T)$, with $A_i(0, T)$ representing the time-average rate utility, $V_i(0, T)$ a penalty term for the rate variability, and γ a nonnegative weight factor. We will suppose

$$A_i(0, T) = \int_{t=\tau_i}^{\tau_i + \sigma_i(T)} U_{s(i)}(r_i(t)) dt,$$

with $\sigma_i(T) = \min\{\sigma_i, T - \tau_i\}$ and $U_s(\cdot)$ a class-based utility function with $U_s(0) \equiv 0$. The function $U_s(\cdot)$ (which can be obtained as in [16]) could be chosen to reflect class-related QR dependencies. These QR dependencies will be video-dependent, where a ‘high-detail’ video with a lot of scene

changes (e.g., an action movie) will have a ‘steeper’ QR dependency compared to a ‘slower’ movie.

We allow $V_i(0, T)$ to be an arbitrary nonnegative function of $\{r_i(t)\}_{t \in [0, T]}$, with the property that $V_i(0, T) = 0$ when $r_i(t) \equiv r_i$ for all $t \in [\tau_i, \tau_i + \sigma_i(T)]$. For example, $V_i(0, T)$ could represent the time standard deviation of quality over the interval $[0, T]$, i.e.,

$$V_i(0, T) = \left(\frac{1}{\sigma_i(T)} \int_{t=\tau_i}^{\tau_i + \sigma_i(T)} \left(U_{s(i)}(r_i(t)) - \frac{A_i(0, T)}{\sigma_i(T)} \right)^2 dt \right)^{1/2}$$

With this choice, the metric $A_i(0, T) - \gamma V_i(0, T)$ would fit into the video QoE model proposed in [28].

The maximum achievable net aggregate QoE during the time interval $[0, T]$ may thus be expressed as $\text{OPT}_\gamma(T) = \sup Q_\gamma(T)$, with

$$Q_\gamma(T) = \sum_{i=1}^{I(T)} [A_i(0, T) - \gamma V_i(0, T)],$$

with $I(T) = \sup\{i : \tau_i \leq T\}$, subject to the constraints:

$$r_i(t) \in \mathcal{R}_{s(i)} \cup \{0\} \quad \text{for all } t \in [0, T], \quad (3)$$

$$\sum_{i=1}^{I(T)} c_{s(i)} r_i(t) \leq 1 \quad \text{for all } t \in [0, T]. \quad (4)$$

Observe that $Q_\gamma(T)$ and hence $\text{OPT}_\gamma(T)$ is nonincreasing in γ , and in particular bounded from above by $\text{OPT}_0(T)$ (i.e., net aggregate utility) where the penalty for the rate variability is entirely dropped.

Remark 1. For non-trivial choices of $V_i(0, T)$ (such as time variance), determining the rate assignments that achieve $\text{OPT}_\gamma(T)$ would require a clairvoyant or off-line optimization scheme with full knowledge of the sample path in terms of the arrival times and service times. In practice, one would use online schemes that need to make rate assignment decisions without any prior knowledge of future arrival times or residual service times. As will be shown however, under mild assumptions, suitable online schemes asymptotically achieve the same long-term value for $\frac{1}{T} Q_\gamma(T)$ for any $\gamma > 0$ as an optimal clairvoyant scheme, even when exempt from a penalty for rate variability, can achieve for $\frac{1}{T} \text{OPT}_0(T)$.

Remark 2. The above formulation allows for zero rate assignments, i.e., $r_i(t) = 0$, for some time epoch $t \in [\tau_i, \tau_i + \sigma_i(T)]$, which may be interpreted as a service request being denied or a flow being aborted. Thus it may not be meaningful in practice to change the rate assignment from a zero to a non-zero value. In the above formulation, we do not explicitly rule out such rate changes, and achieving $\text{OPT}_\gamma(T)$ might require such actions in certain cases. However, the rate assignment policies that we develop never make such changes, and in fact never change the initial rate assignment at all. Yet, they asymptotically achieve the same long-term value for $\frac{1}{T} Q_\gamma(T)$ for any $\gamma > 0$ as schemes that are allowed to make such changes and are not penalized for rate variability can achieve for $\frac{1}{T} Q_0(T)$.

III. UTILITY UPPER BOUNDS

We now derive upper bounds for the maximum achievable net aggregate utility (which bounds the maximum net aggregate QoE). Some bounds hold in a sample path sense, but for other results we will rely on a mild statistical assumption:

A1: class- s users arrive as a renewal process with mean interarrival time α_s ,

or a somewhat stronger Markovian assumption:

A2: class- s users arrive as a Poisson process of rate λ_s , in conjunction with:

B: class- s users have i.i.d. service times with mean β_s .

We define the offered traffic of class $s \in \mathcal{S}$ under assumptions A1 or A2 as $\rho_s = \beta_s/\alpha_s$ or $\rho_s = \lambda_s\beta_s$, respectively.

Because of page constraints, we henceforth focus our attention on a scenario with discrete rate sets \mathcal{R}_s , but most of the results easily extend to continuous rate sets. For compactness, we denote by $K = \sum_{s \in \mathcal{S}} M_s$, with $M_s = |\mathcal{R}_s| < \infty$, the total number of class-rate tuples, and define the set $\mathcal{I}(s, T) = \{i : s(i) = s, \tau_i \leq T\}$ containing the indices of the users that belong to class s and arrive before time T .

Proposition 1. $\text{OPT}_0(T)$ is bounded from above sample path wise by the optimal value $\text{OPT-UB-D}(T)$ of the linear program $\text{LP-UB-D}(T)$:

$$\begin{aligned} & \max_{[(x_{sr})_{s \in \mathcal{S}, r \in \mathcal{R}_s}] \in \mathbb{R}_+^K} \sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{R}_s} x_{sr} U_s(r) \\ & \text{sub} \quad \sum_{r \in \mathcal{R}_s} x_{sr} \leq \sum_{i \in \mathcal{I}(s, T)} \sigma_i, \quad \forall s \in \mathcal{S} \\ & \sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{R}_s} c_s r x_{sr} \leq T \end{aligned}$$

Also, under Assumptions A1 and B, $\frac{1}{T} \text{OPT-UB-D}(T) \rightarrow \text{OPT-UB-D}$ in probability as $T \rightarrow \infty$, where OPT-UB-D is the optimal value of the linear program LP-UB-D :

$$\max_{[(f_{sr})_{s \in \mathcal{S}, r \in \mathcal{R}_s}] \in \mathbb{R}_+^K} \sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{R}_s} \rho_s f_{sr} U_s(r) \quad (5)$$

$$\text{sub} \quad \sum_{r \in \mathcal{R}_s} f_{sr} \leq 1, \quad \forall s \in \mathcal{S} \quad (6)$$

$$\sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{R}_s} c_s r \rho_s f_{sr} \leq 1 \quad (7)$$

Proof:

Let σ_{ir} be the amount of time during $[\tau_i, \tau_i + \sigma_i(T)]$ that the i -th user ($\tau_i \leq T$) is assigned rate $r \in \mathcal{R}_{s(i)}$. Let $x_{sr} = \sum_{i \in \mathcal{I}(s, T)} \sigma_{ir}$ be the total amount of time during $[0, T]$ that any class- s user is assigned rate $r \in \mathcal{R}_s$. Noting that $U_s(0) \equiv 0$, $A_i(0, T)$ may be written as $\sum_{r \in \mathcal{R}_{s(i)}} \sigma_{ir} U_s(r)$, and thus

$$\sum_{i=1}^{I(T)} A_i(0, T) = \sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{R}_s} x_{sr} U_s(r).$$

Invoking the rate constraints (3) and observing that $\sum_{r \in \mathcal{R}_{s(i)}} \sigma_{ir} \leq \sigma_i(T) \leq \sigma_i$, we obtain

$$\sum_{r \in \mathcal{R}_s} x_{sr} \leq \sum_{i \in \mathcal{I}(s, T)} \sigma_i, \quad s \in \mathcal{S}. \quad (8)$$

Further, integrating the capacity constraints (4) over $t \in [0, T]$, we obtain

$$\sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{R}_s} c_s r x_{sr} \leq T. \quad (9)$$

Thus, any rate assignment solution that satisfies constraints (3) and (4) will satisfy constraints (8) and (9) as well. We conclude that $\text{OPT}_0(T)$ is bounded from above by the maximum value of $\sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{R}_s} x_{sr} U_s(r)$ subject to the constraints (8) and (9), which concludes the proof of the first statement of the proposition.

In order to establish the second statement, we note that, introducing variables $f_{sr} = x_{sr}/(\rho_s T)$, the linear program LP-UB-D(T) may be equivalently stated as

$$\begin{aligned} & \max_{[(f_{sr})_{s \in \mathcal{S}, r \in \mathcal{R}_s}] \in \mathbb{R}_+^K} \sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{R}_s} \rho_s f_{sr} U_s(r) \\ & \text{sub} \quad \sum_{r \in \mathcal{R}_s} \rho_s f_{sr} \leq \frac{1}{T} \sum_{i \in \mathcal{I}(s, T)} \sigma_i, \quad \forall s \in \mathcal{S} \\ & \quad \sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{R}_s} c_s r \rho_s f_{sr} \leq 1. \end{aligned}$$

The key renewal theorem and the law of large numbers imply that $\frac{1}{T} \sum_{i \in \mathcal{I}(s, T)} \sigma_i \rightarrow \rho_s$ in probability as $T \rightarrow \infty$. Applying Theorem 2.1 in [12] then completes the proof. ■

As reflected in the proof of the above proposition, the linear program LP-UB-D may be interpreted as follows. The decision variables f_{sr} represent the fractions of class- s users that are assigned rate $r \in \mathcal{R}_s$, so their sum cannot exceed unity as expressed by the constraint (6). Further observe that ρ_s is the mean number of class- s users with an active request at an arbitrary epoch. Thus the objective function (5) represents the expected utility rate at an arbitrary epoch, while the constraint (7) stipulates that the ‘load’, i.e., the expected total capacity requirement, cannot exceed unity.

IV. ASYMPTOTIC OPTIMALITY

We now proceed to show that under Assumptions A2 and B the upper bound OPT-UB-D for the long-term average utility rate as provided in Proposition 1 is asymptotically achievable (in a ratio sense) in a regime where the offered traffic and transmission capacity grow large in proportion. More specifically, let us consider a sequence of systems, each indexed by a positive integer n . Each system has the same set of user classes \mathcal{S} , the same rate sets \mathcal{R}_s , and the same utility functions $U_s(\cdot)$. In the n -th system the class- s capacity requirement is $c_s^{(n)} = c_s/n$, and the arrival rate and mean service time are such that $\rho_s^{(n)} = \lambda_s^{(n)} \beta_s^{(n)} = n\rho_s$. Let $\text{OPT-UB-D}^{(n)}$ be the optimal value of the linear program LP-UB-D for the n -th system. Since the number of ‘potentially’ active users evolves as the number of customers in an infinite-server queue, it follows that the system is regenerative. The renewal-reward theorem then implies that the long-term average utility rate $U_\pi^{(n)}$ achieved in the n -th system by a stationary policy π exists.

Proposition 2. *Under Assumption A2 and B, there exist stationary policies π such that*

$$\lim_{n \rightarrow \infty} \frac{U_\pi^{(n)}}{\text{OPT-UB-D}^{(n)}} = 1,$$

and in particular suitably defined complete partitioning and complete sharing policies are asymptotically optimal in the above sense.

Proof:

Proposition 1 implies $U_\pi^{(n)} \leq \text{OPT-UB-D}^{(n)}$ for all n . Also, substituting $c_s^{(n)} = c_s/n$ and $\rho_s^{(n)} = n\rho_s$ in the linear program LP-UB-D for the n -th system yields $\text{OPT-UB-D}^{(n)} = n \times \text{OPT-UB-D}$. Thus it suffices to show that $\liminf_{n \rightarrow \infty} \frac{U_\pi^{(n)}}{n \times \text{OPT-UB-D}} \geq 1$ for some specific policy π . There are several policies π for which the latter can be shown, in particular complete partitioning and complete sharing policies.

(i) (Complete Partitioning)

For every class-rate pair (s, r) , define $C_{sr}^{(n)} = \lfloor n\rho_s f_{sr}^* \rfloor$, where $(f_{sr}^*)_{s \in \mathcal{S}, r \in \mathcal{R}_s}$ is an optimal solution to the linear program LP-UB-D. (Usually, only $|\mathcal{S}| + 1$ of these variables will in fact be non-zero, as will be shown later.) Consider a policy π which assigns an arriving class- s user a nominal rate $r \in \mathcal{R}_s$ with probability f_{sr}^* or rate 0 with probability $1 - \sum_{r \in \mathcal{R}_s} f_{sr}^*$. In the n -th system a user with a non-zero rate assignment is then admitted and provided that rate for its entire duration, if the total number of class- s users with that rate assignment remains below $C_{sr}^{(n)}$, or otherwise rejected. Note that the aggregate capacity requirement for admitted users in the n -th system is bounded from above at all times by

$$\sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{R}_s} c_s^{(n)} r C_{sr}^{(n)} \leq \sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{R}_s} c_s r \rho_s f_{sr}^* \leq 1,$$

by virtue of (7), ensuring that the capacity constraints are obeyed. In effect, this yields K (or typically only $|\mathcal{S}| + 1$) independent Erlang-B loss systems with offered loads $n\rho_s f_{sr}^*$ and capacities $C_{sr}^{(n)}$. Since each of these systems is critically loaded, it can be easily shown that the blocking probabilities for each of the class-rate pairs (s, r) tend to zero as $n \rightarrow \infty$. Thus, for any $\epsilon > 0$, the expected number of active class- s users with rate assignment $r \in \mathcal{R}_s$ in the n -th system will be larger than $(1 - \epsilon)n\rho_s f_{sr}^*$ for n sufficiently large. Noting that the utility enjoyed by each of these users is $U_s(r)$, we deduce

$$U_\pi^{(n)} \geq (1 - \epsilon)n \sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{R}_s} \rho_s f_{sr}^* U_s(r) = (1 - \epsilon)n \times \text{OPT-UB-D},$$

for n sufficiently large, implying $\liminf_{n \rightarrow \infty} \frac{U_\pi^{(n)}}{n \times \text{OPT-UB-D}} \geq 1 - \epsilon$. Since $\epsilon > 0$ can be taken arbitrarily small, the statement follows.

(ii) (Complete Sharing)

Consider a policy π which assigns an arriving class- s user a nominal rate $r \in \mathcal{R}_s$ with probability f_{sr}^* or rate 0 with probability $1 - \sum_{r \in \mathcal{R}_s} f_{sr}^*$. In the n -th system a user with a non-zero rate assignment is then admitted and provided that rate for its entire duration, if the total capacity required by all active users remains below 1, or otherwise rejected. In

effect, this yields an Erlang loss system with K (or typically only $|\mathcal{S}| + 1$) classes with offered loads $n\rho_s f_{sr}^*$ and per-user capacity requirements rc_s , and total available capacity n . Since $\sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{R}_s} c_s r \rho_s f_{sr}^* \leq 1$ by virtue of (7), this system is at most critically loaded, and it can be easily shown that the blocking probabilities for each of the class-rate pairs (s, r) , $s \in \mathcal{S}$, $r \in \mathcal{R}_s$, tend to zero as $n \rightarrow \infty$. The statement then follows in a similar way as in the previous case. ■

As reflected in the proof of the above proposition, asymptotic optimality is achieved by ‘static’ rate assignment policies, which either deny service to a user or assign it a fixed rate independent of the state of the system at the time of arrival, without requiring advance knowledge of the duration. The latter features make such policies straightforward to implement. While the specific complete partitioning and complete sharing policies may not be ideal from a practical perspective, the asymptotic optimality of static rate assignment policies yields a useful principle for the design of heuristic approaches. The only caveat is that solving the linear program LP-UB-D requires knowledge of the relevant channel and traffic parameters. In practice several of these parameters will be subject to uncertainty and variation over time, raising all sorts of estimation challenges. We will later revisit these issues and present parsimonious, online policies which retain the simplicity of a static rate assignment, but do not require any explicit knowledge or estimates of the traffic statistics.

V. STRUCTURAL PROPERTIES OF (ASYMPTOTICALLY) OPTIMAL POLICIES

As an initial step towards the design of parsimonious, online rate assignment algorithms we now proceed to examine the structural properties of the optimal solution to the linear program LP-UB-D. It will be convenient to use the representation $\mathcal{R}_s = \{r_s^{(1)}, r_s^{(2)}, \dots, r_s^{(M_s)}\}$, with $r_s^{(1)} < r_s^{(2)} < \dots < r_s^{(M_s)}$. In the case that $\sum_{s \in \mathcal{S}} c_s \rho_s r_s^{(M_s)} \leq 1$, the optimal solution is easily seen to be $f_{sr_s^{(M_s)}} = 1$ for all $s \in \mathcal{S}$, and we henceforth assume that $\sum_{s \in \mathcal{S}} c_s \rho_s r_s^{(M_s)} > 1$.

For compactness, we also introduce the notation $\Delta U_{sm} = U_s(r_s^{(m)}) - U_s(r_s^{(m-1)})$ and $\Delta R_{sm} = r_s^{(m)} - r_s^{(m-1)}$, with $r_s^{(0)} \equiv 0$. Denote $x_{sm} = f_{sr_s^{(m)}}$ and $y_{sm} = \sum_{n=m}^{M_s} x_{sn}$. Noting that $\sum_{r \in \mathcal{R}_s} r f_{sr} = \sum_{m=1}^{M_s} \Delta R_{sm} y_{sm}$ and $\sum_{r \in \mathcal{R}_s} f_{sr} U_s(r) = \sum_{m=1}^{M_s} \Delta U_{sm} y_{sm}$, it is easily verified that the linear program LP-UB-D may be equivalently stated as the ‘knapsack-like’ problem KP-D:

$$\begin{aligned} & \max_{\{(y_{sm})_{s \in \mathcal{S}, m \in \{1, \dots, M_s\}}\} \in \mathbb{R}_+^K} \sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{R}_s} \rho_s \sum_{m=1}^{M_s} \Delta U_{sm} y_{sm} \\ \text{sub} \quad & y_{s1} \leq 1 \quad \forall s \in \mathcal{S} \\ & y_{sm} \leq y_{sm-1} \quad m = 2, \dots, M_s, \forall s \in \mathcal{S} \\ & \sum_{s \in \mathcal{S}} c_s \rho_s \sum_{m=1}^{M_s} \Delta R_{sm} y_{sm} \leq 1 \end{aligned} \quad (10)$$

Define

$$Q_{sm} = \Delta U_{sm} / (c_s \Delta R_{sm}),$$

which may be interpreted as a marginal payoff coefficient, i.e., the ratio of the improvement in the rate utility and the increment in the required capacity when increasing the rate of a class- s user from $r_s^{(m-1)}$ to $r_s^{(m)}$. We assume that $Q_{sm-1} > Q_{sm}$, for all $m = 2, \dots, M_s$. This assumption means that there exists a strictly concave function that coincides with the utility function $U_s(\cdot)$ in all the points $0, r_s^{(1)}, \dots, r_s^{(M_s)}$. Let $K_{\mathcal{R}}$ be the size of the set $\{Q_{sm} : s \in \mathcal{S}, m = 1, \dots, M_s\}$. Thus, $K_{\mathcal{R}} \leq K$ is the number of distinct values of Q_{sm} . Let $\pi_{\mathcal{R}}(\cdot, \cdot)$ be the mapping from $(s, m)_{s \in \mathcal{S}, m \in \{1, \dots, M_s\}}$ to $\{1, 2, \dots, K_{\mathcal{R}}\}$ so that $\pi_{\mathcal{R}}(s_1, m_1) < \pi_{\mathcal{R}}(s_2, m_2) \iff Q_{s_1 m_1} > Q_{s_2 m_2}$. Denote $\mathcal{W} = \{w \in [0, 1]^{K_{\mathcal{R}}} : w_k = 1 \text{ if } w_{k+1} > 0\}$. For a given $w \in \mathcal{W}$, define $L_D(w) = \sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{R}_s} c_s \rho_s r f_{sr}$ as the ‘load’ associated with w , with f_{sr} defined in terms of w by

$$y_{sm} = w_{\pi_{\mathcal{R}}(s, m)}, \quad m = 1, \dots, M_s, s \in \mathcal{S}, \quad (11)$$

$$f_{sr_s^{(m)}} = y_{sm} - y_{sm+1}, \quad m = 1, \dots, M_s, s \in \mathcal{S}. \quad (12)$$

The next lemma provides a useful characterization of the optimal solution to the linear program LP-UB-D.

Lemma 1. *If $L_D(w) = 1$, then the variables f_{sr} constitute an optimal solution to the linear program LP-UB-D.*

Proof:

It is easily verified that the variables f_{sr} provide a feasible solution to problem LP-UB-D. In order to establish optimality, it may be checked that if $L_D(w) = 1$, then the variables y_{sm} saturate the capacity constraint (10) and constitute an optimal solution of a relaxed version of problem KP-D. Recalling that $\sum_{r \in \mathcal{R}_s} f_{sr} U_s(r) = \sum_{m=1}^{M_s} \Delta U_{sm} y_{sm}$, it then follows that the variables f_{sr} in fact constitute an optimal solution. ■

The above lemma shows that determining the asymptotically optimal rate assignment fractions is no harder than solving the equation $L_D(w) = 1$. The only caveat is that the coefficients of the latter equation involve various channel and traffic parameters. While the rate sets \mathcal{R}_s and the capacity requirements c_s can essentially be thought of as a matter of definition (specification of classes), the corresponding amounts of offered traffic ρ_s would in practice typically involve uncertainty and variation over time, raising measurement-related trade-offs. Instead we will therefore develop adaptive, online algorithms which do not require any explicit knowledge or estimates of the traffic statistics.

VI. DESIGN OF ADAPTIVE RATE ALLOCATION POLICIES

We now discuss how the insights in the structural properties of the optimal solution to the linear program LP-UB-D can be exploited to design adaptive, online rate assignment policies.

A. Preliminary results

We first use the structural properties obtained in Section V to establish that the asymptotically optimal rate assignments f_{sr}^* only depend on the offered traffic through a single scalar variable θ^* . This fact will play an instrumental role in the design of adaptive rate assignment algorithms.

Let $G(\cdot)$ be a function on \mathcal{W} (as defined in Section V) of the form

$$G(w) = h_0 + \sum_{k=1}^{K_{\mathcal{R}}} h_k(w_k),$$

where h_0 is a constant coefficient and each of the functions $h_k(\cdot)$ is strictly decreasing and continuous. Further denote $\theta_{\min} = G(\mathbf{1}_{K_{\mathcal{R}}}) = h_0 + \sum_{k=1}^{K_{\mathcal{R}}} h_k(1)$ and $\theta_{\max} = G(\mathbf{1}_0) = h_0 + \sum_{k=1}^{K_{\mathcal{R}}} h_k(0)$, where $\mathbf{1}_k$ is the $K_{\mathcal{R}}$ -dimensional vector whose first k components are 1 and whose remaining $K_{\mathcal{R}} - k$ components are 0, $k = 0, \dots, K_{\mathcal{R}}$.

Lemma 2. *The function $G(\cdot)$ is injective, i.e., for any $\theta \in [\theta_{\min}, \theta_{\max}]$, there exists a unique $w(\theta) = G^{-1}(\theta) \in \mathcal{W}$ such that $G(w(\theta)) = \theta$, which may be obtained as*

$$w_k(\theta) = \begin{cases} 1, & k = K_{\mathcal{R}}(\theta) \\ h_{K_{\mathcal{R}}(\theta)}^{-1}(\theta - G(\mathbf{1}_{K_{\mathcal{R}}(\theta)}) + h_{K_{\mathcal{R}}(\theta)}(1)), & k = K_{\mathcal{R}}(\theta) \\ 0, & k > K_{\mathcal{R}}(\theta) \end{cases}$$

with

$$K_{\mathcal{R}}(\theta) = \min\{k : 1 \leq k \leq K_{\mathcal{R}}, G(\mathbf{1}_k) \leq \theta\}.$$

Based on the above lemma, we introduce assignment fractions $f_{sr}(\theta)$ defined in terms of the variables $w_k(\theta)$ using (11)–(12). Also, define $L_D(\theta) = \sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{R}_s} c_s r \rho_s f_{sr}(\theta)$ as the ‘load’ associated with θ . Denote by $L_D^{\max} = \sum_{s \in \mathcal{S}} c_s \rho_s r_s^{(M_s)}$ the maximum possible load, i.e., the load if all users were assigned the highest possible rate for their respective classes. Further define

$$\theta^* = \min\{\theta \in [\theta_{\min}, \theta_{\max}] : L_D(\theta) \leq 1\}.$$

The next lemma collects a few useful properties of the function $L_D(\cdot)$.

Lemma 3. (i) *The function $L_D(\cdot)$ is strictly decreasing and continuous on $[\theta_{\min}, \theta_{\max}]$, with $L_D(\theta_{\min}) = L_D^{\max}$ and $L_D(\theta_{\max}) = 0$;*
(ii) *If $L_D^{\max} < 1$, then $\theta^* = \theta_{\min}$;*
(iii) *If $L_D^{\max} \geq 1$, then $L_D(\theta^*) = 1$.*

The next lemma provides a useful optimality result based on Lemmas 1 and 3.

Lemma 4. *The variables $f_{sr}(\theta^*)$ constitute an optimal solution to the linear program LP-UB-D.*

B. An adaptive rate allocation scheme: ARA-D

Lemma 4 reduces the problem of determining the asymptotically optimal rate assignments to that of finding the value of θ^* with associated load $L_D(\theta^*) = 1$. We now proceed to develop an adaptive scheme for performing the latter task in an online manner without any explicit knowledge of the amounts of offered traffic ρ_s . For convenience, we take the function $G(\cdot)$ to be affine,

$$h_k(w_k) = w_k(Q_{\mathcal{R}}(k) - Q_{\mathcal{R}}(k-1)),$$

where $Q_{\mathcal{R}}(k) = Q(s, m)$ for $(s, m) \in \pi_R^{-1}(k)$, $1 \leq k \leq K_{\mathcal{R}}$, and the coefficient h_0 is set to some arbitrary value

$Q_{\mathcal{R}}(0) > Q_{\mathcal{R}}(1)$. In particular, $\theta_{\min} = G(\mathbf{1}_{K_{\mathcal{R}}}) = Q_{\mathcal{R}}(K_{\mathcal{R}})$ and $\theta_{\max} = G(\mathbf{1}_0) = Q_{\mathcal{R}}(0)$. We henceforth make the further natural assumption that any rates that cannot possibly be supported have been removed from the rate sets, i.e., $r_s^{(M_s)} \leq 1/c_s$ for all $s \in \mathcal{S}$.

The main thrust of the adaptive scheme is to treat θ as a control parameter and use estimates of $L_D(\theta)$ to drive it towards the value θ^* for which $L_D(\theta^*) = 1$. There are obviously various potential ways in which the value of θ can be adapted and in which $L_D(\theta)$ can be estimated. The specific scheme that we develop below, termed ARA-D, has the key advantage that it naturally extends to scenarios where no discrete user classes are specified as will be briefly discussed in Subsection VI-D.

The ARA-D scheme is comprised of two closely coupled components: the rate assignment component ARA-DR(θ) and the θ -learning component ARA-DA. The ARA-DR(θ) component uses the θ values produced by the ARA-DA component to assign a nominal rate to each arriving user. In turn, the ARA-DA component adapts the θ values based on observations of the occupancy state generated by the rate decisions of the ARA-DR(θ) component, and serves to find the value θ^* for which $L_D(\theta^*) = 1$.

We now proceed to describe both components of the ARA-D scheme in more detail. Denote by $N_{sr}(t)$ the number of class- s users provided rate r in the system at time t . Define the free capacity at time t as

$$m(t) = 1 - \sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{R}_s} c_s r N_{sr}(t).$$

ARA-DR(θ)

Given the value of the control parameter θ , ARA-DR(θ) assigns an arriving class- s user a nominal rate $r \in \mathcal{R}_s$ with probability $f_{sr}(\theta)$ as defined previously. Denote by θ_i the value of θ used for the i th arriving user with arrival epoch τ_i , and denote by R_i the nominal rate assigned to that user. If R_i is either zero or larger than the capacity constraint allows for, i.e., $c_{s(i)} R_i > m(\tau_i^-)$, then the i th user is blocked, i.e., $r_i(t) = 0$ for all $t \in [\tau_i, \tau_i + \sigma_i]$. Otherwise, the i th user is admitted and provided rate R_i for its entire duration.

Note that the assignment fractions $f_{sr}(\theta)$ can be efficiently calculated given the value of θ for the above choice of $G(\cdot)$. Specifically, we have $K_{\mathcal{R}}(\theta) = \min\{1 \leq k \leq K_{\mathcal{R}} : Q_{\mathcal{R}}(k) \leq \theta\}$. Users of the class(es) $s \in \mathcal{S}$ for which there is an $m_s(\theta) \in \{1, \dots, M_s\}$ such that $(s, m_s(\theta)) \in \pi_{\mathcal{R}}^{-1}(K_{\mathcal{R}}(\theta))$, are either assigned a nominal rate $r_s^{(m_s(\theta))}$ with probability

$$\frac{Q_{\mathcal{R}}(K_{\mathcal{R}}(\theta) - 1) - \theta}{Q_{\mathcal{R}}(K_{\mathcal{R}}(\theta) - 1) - Q_{\mathcal{R}}(K_{\mathcal{R}}(\theta))}$$

or rate $r_s^{(m_s(\theta)-1)}$ otherwise. Users of all the other classes $s \in \mathcal{S}$ are always assigned a nominal rate $r_s^{(m_s(\theta))}$, where $m_s(\theta) = \max\{0\} \cup \{m : 1 \leq m \leq M_s, Q_{sm} > \theta\}$, i.e., the largest rate $r_s^{(m)}$ (or possibly zero) for which the marginal payoff Q_{sm} is no less than θ , or equivalently $Q_{\mathcal{R}}(K_{\mathcal{R}}(\theta))$.

ARA-DA

The role of ARA-DA is to use estimates of $L_D(\theta)$ to adapt the control parameter θ and drive it to the value θ^* such that $L_D(\theta^*) = 1$. The key challenge here is that the occupancy state only provides an indication of the *carried* traffic, and that blocking must be properly taken into account to obtain an (unbiased) estimate of the actual *offered* traffic. In order to achieve that, we will use a tagging procedure, and tag the i th user when the free capacity $m(\tau_i^-)$ at the arrival epoch of that user is larger than $m_{\min} = \max_{s \in \mathcal{S}} c_s r_s^{(M_s)} \leq 1$. This choice implies that every user getting tagged will either be admitted, or assigned a zero rate and be blocked. In other words, blocked users with non-zero nominal rate assignments will never be tagged. This ensures that the nominal rates assigned to tagged users can be tracked without having to monitor blocked users.

Let the 0–1 variable $\hat{T}_i = I_{\{m(\tau_i^-) \geq m_{\min}\}}$ indicate whether the i th user is tagged or not. Denote by $\hat{N}_{sr}(t)$ the number of tagged class- s users provided rate r at time t . Define the load associated with tagged users at time t as

$$\hat{n}(t) = \sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{R}_s} c_s r \hat{N}_{sr}(t),$$

which can be readily calculated without maintaining any detailed state information or tracking blocked users.

When the j th user arrives, ARA-DA computes the value θ_{j+1} as follows

$$\theta_{j+1} = \left[\theta_j + \epsilon \left(\hat{n}(\tau_j^-) - \hat{T}_j \right) \right]_{\theta_{\min}}^{\theta_{\max}}, \quad (13)$$

with step size $\epsilon > 0$, an arbitrary initial value θ_1 , and $[x]_a^b$ equal to a , x or b if $x < a$, $a \leq x \leq b$ or $x > b$, respectively. Although the value of θ_{j+1} is already known at the time of the arrival of the j th user, it is worth stressing that ARA-DR ‘lags’ by one time step, in the sense that the value of θ_j is used for its nominal rate assignment.

As may be recognized in the update rule (13), ARA-DA balances the load associated with the tagged users with the fraction of users which are tagged. This, in turn, steers the load of the system to unity and drives the value of θ to θ^* , as will be proved next under some suitable assumptions.

C. Convergence proof

Throughout we make Assumption A2 (Poisson arrival processes) and a slightly strengthened version of Assumption B, namely that the service times of the various users are independent and have phase-type distributions. The class of phase-type distributions is well-known to be dense in the set of all distributions for positive random variables.

In order to establish convergence, we need to consider the stationary behavior of the system if it operated under ARA-DR(θ) with a fixed value θ . Specifically, let $\hat{a}(\theta)$ be the fraction of users which get tagged and let $(N_{sr}^\theta)_{s \in \mathcal{S}, r \in \mathcal{R}_s}$ be a random vector with the stationary distribution of the occupancy state. Then the PASTA property implies

$$\hat{a}(\theta) = \lim_{J \rightarrow \infty} \frac{\sum_{j=1}^J \hat{T}_j}{J} = P \left[\sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{R}_s} c_s N_{sr}^\theta r \leq 1 - m_{\min} \right].$$

Now consider the following projected ordinary differential equation (see Section 4.3 of [20]) which captures the smoothened evolution of $\{\theta_j\}_{j \geq 1}$ under the update rule (13):

$$\frac{d\theta(t)}{dt} = \hat{a}(\theta(t)) (L_D(\theta(t)) - 1) + z(\theta(t)),$$

where

$$z(\theta) = \begin{cases} -\hat{a}(\theta_{\min}) (L_D(\theta_{\min}) - 1), & \text{if } L_D(\theta_{\min}) < 1, \theta = \theta_{\min}; \\ 0 & \text{otherwise.} \end{cases}$$

Note that $[L_D(\theta) - 1]^+$ may be interpreted as the excess load when operating under a fixed value θ . For each t , the function $z(t)$ is nonnegative, and guarantees that $\frac{d\theta(t)}{dt} \geq 0$ for $\theta(t) = \theta_{\min}$, thus ensuring that $\theta(t)$, like θ_j , cannot fall below θ_{\min} .

Lemma 5. *For any initial value $\theta(0) \in [\theta_{\min}, \theta_{\max}]$,*

$$\lim_{t \rightarrow \infty} \theta(t) = \theta^*.$$

The proof uses a Lyapunov function $V(\theta) = \frac{1}{2} (\theta - \theta^*)^2$. The asymptotic stability can then be established using that $\frac{dV(\theta(t))}{dt} \leq -\hat{a}(\theta(t)) (\theta(t) - \theta^*) (L_D(\theta(t)) - 1)$ in conjunction with Lemma 3 and continuity of the functions $L_D(\cdot)$ and $\hat{a}(\cdot)$. Because of page constraints, the details are omitted.

The next theorem states that for a small enough step size ϵ and sufficiently large number of iterates, virtually all of the iterates obtained by the ARA-D scheme will belong to an arbitrarily close neighborhood of θ^* .

Theorem 1. *For any $\delta_\theta > 0$, the fraction of iterates $(\theta_j)_{1 \leq j \leq T/\epsilon}$ that take values in $(\theta^* - \delta_\theta, \theta^* + \delta_\theta)$ goes to one in probability as $\epsilon \downarrow 0$ and $T \rightarrow \infty$.*

The proof relies on Theorem 4.1 in Chapter 8 of [20] and the asymptotic stability of θ^* established in Lemma 5. The main task is to verify that all the conditions of Theorem 4.1 are satisfied in the present setting, which is tedious but not difficult. Because of page limitations, the details are omitted.

D. Alternative schemes: ARA-A

As noted earlier, under ARA-DR(θ) the users of most classes are assigned the largest possible nominal rate for which the marginal payoff is strictly higher than θ , or equivalently $Q_{\mathcal{R}}(K_{\mathcal{R}}(\theta))$. Typically only some users of one particular class receive a larger nominal rate assignment so as to make full use of the available capacity. This is also borne out by the fact that $Q_{\mathcal{R}}(K_{\mathcal{R}}(\theta^*))$ can be shown to be the optimal dual variable corresponding to the capacity constraint (10). This suggests the following somewhat simpler alternative rate assignment scheme, which we will refer to as **ARA-AR**(θ). Given the value of the control parameter θ , ARA-AR(θ) assigns an arriving class- s user a nominal rate 0 if $Q_{s1} < \theta$ and $r_s(\theta) = \max\{r_s^{(m)} : Q_{sm} \geq \theta\}$ otherwise. The latter scheme has the key advantage that the rate assignment does not involve any parameters of other user classes, and can be applied in conjunction with ARA-DA as described before,

yielding the scheme **ARA-A**. In particular, interpreting Q_{sm} as a left-derivative, ARA-AR naturally extends to scenarios where no discrete user classes are specified, which can easily be handled in ARA-DA as well.

VII. SIMULATION RESULTS

In this section, we examine the performance of the ARA-D and ARA-A schemes through simulations. We are not aware of any benchmark schemes for allocating target rates in the presence of user dynamics, and a comparison with a default Proportional Fair mechanism as implemented in most cellular networks today would not seem fair. Also, the actual optimal policy is computationally intractable, even assuming the dynamics to be Markovian and all parameters fully known. Hence, we will compare the performance of the ARA-D and ARA-A schemes with the upper bound OPT-UB-D. The fact that the gap proves to be relatively small, demonstrates that no significant gains can be achieved by alternate strategies, even if we completely disregard the penalty for rate variability that such policies might incur in order to improve the average rate utility. In the simulation experiments we do not explicitly consider the impact of variability in the actual delivered throughput and playout rate due to opportunistic scheduling and adaptation at the video client and server. Since the ARA-A and ARA-D schemes never overallocate or change the initial target rate assigned to a given user, it is plausible however that they provide greater rate stability than any alternative scheme.

We consider a system with five classes of users indexed by $\mathcal{S} = \{1, 2, 3, 4, 5\}$. Class- s users arrive as a Poisson process of rate $\lambda_s^{(n)} = n\lambda$, and we vary the scaling coefficient n and λ for different simulations. The mean service time for each class is 300 seconds (hence, $\rho_s^{(n)} = 300\lambda n$ for each class $s \in \mathcal{S}$). Service times of users of classes 2 and 4 are deterministic (i.e., are equal to 300 seconds) and those of users of classes 1, 3 and 5 are exponentially distributed. Furthermore, $c_s^{(n)} = c_s/n = 1/(p_s n)$ for each class $s \in \mathcal{S}$, with $[p_1 \ p_2 \ p_3 \ p_4 \ p_5] = [7.5 \ 6 \ 4.5 \ 3 \ 1.5]$ Mbps. With these parameter settings, a representative value of the scaling coefficient would be around $n = 10$, and even $n = 20$ may not be unreasonable from the viewpoint of typical cell peak rates ranging from tens to hundreds of Mbps in 4G LTE systems. The respective rate sets are as follows: $\mathcal{R}_1 = [0.1 \ 0.2 \ 0.4 \ 0.8]$ Mbps, $\mathcal{R}_2 = [0.2 \ 0.3 \ 0.5 \ 0.8 \ 1]$ Mbps, $\mathcal{R}_3 = [0.3]$ Mbps, $\mathcal{R}_4 = [0.1 \ 0.3 \ 0.5 \ 0.7]$ Mbps, and $\mathcal{R}_5 = [0.1 \ 0.2 \ 0.3]$ Mbps. We let $U_s(r) = \log_e(1 + 10r)$ where r is measured in Mbps.

Choices for the value of ϵ were determined in an empirical way, and we chose $\epsilon = 0.1n$ throughout. Selecting a good ϵ involves a tradeoff between the tracking ability and the ability to guide θ_j towards θ^* . Small values of ϵ ensure that θ_j stays close to θ^* . However if the system parameters change over time, reasonably large values of ϵ allow the ARA-D and ARA-A schemes to track changes in the parameters. Also, note that good choices for ϵ will typically depend on system parameters like the amounts of offered traffic ρ_s , the utility functions $U_s(\cdot)$ and the capacity requirements c_s as well. This is due to the fact that ϵ controls the responsiveness of the update rule (13)

which guides θ_j towards θ^* , which in turn keeps the system load close to unity. However, the local sensitivity of the system load near θ^* depends on θ^* (for instance, see Fig. 1), and hence depends on these system parameters.




Fig. 1. Load $L_D(\theta)$ under ARA-D as function of θ for $n = 8$.

In Fig. 2, we compare the average user QoE of ARA-D and ARA-A with the upper bound $\text{OPT-UB-D} / \sum_{s \in \mathcal{S}} \rho_s^{(n)}$ for different values of the scaling coefficient n . The top and bottom charts correspond to $\lambda = 0.01$ and $\lambda = 0.02$, respectively. Each point is obtained by simulating 10 hours of system time. As the results indicate, for both schemes the performance is close to the upper bound, within 10% for realistic values of n as small as 5 to 10.




Fig. 2. Performance of ARA-D and ARA-A for $\lambda = 0.01$ (top figure) and $\lambda = 0.02$ (bottom figure).

The top and bottom figures in Fig. 3 depict the evolution

over time of θ_i and the estimate of the time-average QoE respectively associated with a sample path for $\lambda = 0.02$ and $n = 8$. These figures illustrate typical initial transient and steady-state behavior of θ_i under ARA-D. We observed similar behavior under ARA-A also.




Fig. 3. Evolution over time of θ_i (top figure) and time average QoE (bottom figure) under ARA-D for $\lambda = 0.02$, $n = 8$.

Next, we investigate the tracking ability of ARA-D. Fig. 4 depicts the evolution of θ_i under ARA-D for $n = 8$ and a time-varying arrival rate $\lambda(t) = 0.02\alpha(t)$, with $\alpha(t)$ as depicted in the lower half of the figure. We observe that ARA-D is able to track the time-varying loads well (refer to Fig. 1 also), and add that ARA-A exhibited similar tracking ability.




Fig. 4. ARA-D: Evolution of θ_i under dynamic offered traffic

VIII. CONCLUSION

We have addressed the problem of optimal rate allocation for video streaming in wireless networks with user dynamics, and developed online, parsimonious algorithms with provable convergence properties and performance guarantees. In future work we plan to extend the work to account for capacity variations due to slow fading as caused by user mobility.

REFERENCES

- [1] S. Akhshabi, L. Anantakrishnan, C. Dovrolis, A.C. Begen (2013). Server-based traffic shaping for stabilizing oscillating adaptive streaming players. In: *Proc. NOSSDAV 2013*.
- [2] S. Akhshabi, L. Anantakrishnan, C. Dovrolis, A.C. Begen (2012). What happens when HTTP adaptive streaming players compete for bandwidth? In: *Proc. NOSSDAV 2012*.
- [3] S. Akhshabi, A.C. Begen, C. Dovrolis (2012). An experimental evaluation of rate adaptation algorithms in adaptive streaming over HTTP. In: *Proc. ACM MMSys 2011*.
- [4] D. Bethanabhotla, G. Caire, M.J. Neely (2013). Joint transmission scheduling and congestion control for adaptive video streaming in small-cell networks. *arXiv:1304.8083 [cs.NI]*.
- [5] B. Blaszczyzyn, M. Jovanovic, M.K. Karray (2013). Quality of real-time streaming in wireless cellular networks - Stochastic modeling and analysis. *arXiv:1304.5034 [cs.NI]*.
- [6] G.E. Cho, C.D. Meyer (2001). Comparison of perturbation bounds for the stationary distribution of a Markov chain. *Lin. Alg. Appl.* **335** (1-3), 137–150.
- [7] Cisco Inc (2011). Cisco visual networking index.
- [8] T. Cloonan (2011). Competitive analysis of adaptive video streaming implementations. In: *SCTE Cable-Tec Expo*.
- [9] N. Cranley, P. Perry, L. Murphy (2006). User perception of adapting video quality. *Int. J. Human-Computer Studies*.
- [10] D. De Vleeschauwer, H. Viswanathan, A. Beck, S. Benno, G. Li, R.B. Miller (2013). Optimization of HTTP adaptive streaming over mobile cellular networks. In: *Proc. IEEE Infocom 2013*.
- [11] F. Dobrian, V. Sekar, A. Awan, I. Stoica, D.A. Joseph, A. Ganjam, J. Zhan, H. Zhang (2011). Understanding the impact of video quality on user engagement. In: *Proc. ACM SIGCOMM 2011*.
- [12] A.V. Fiacco, Y. Ishizuka (1990). Sensitivity and stability analysis for nonlinear programming. *Ann. Oper. Res.* **27** (1-4), 215–236.
- [13] F. Fu, M. van der Schaar (2010). A systematic framework for dynamically optimizing multi-user wireless video transmission. *IEEE J. Sel. Areas Commun.* **28**.
- [14] K.A.D. Havey, R. Chertov (2012). Receiver driven rate adaptation for wireless multimedia applications. In: *Proc. ACM MMSys 2012*.
- [15] R. Houdaille, S. Gouache (2012). Shaping HTTP adaptive streams for a better user experience. In: *Proc. ACM MMSys 2012*.
- [16] H. Hu, X. Zhu, Y. Wang, R. Pan, J. Zhu, F. Bonomi (2012). QoE-based multi-stream scalable video adaptation over wireless networks with proxy. In: *Proc. IEEE ICC 2012*.
- [17] J. Huang, Z. Li, M. Chiang, A. Katsaggelos (2008). Joint source adaptation and resource allocation for multi-user wireless video streaming. *IEEE Trans. Circ. Syst. Video Techn.*
- [18] J. Jiang, V. Sekar, H. Zhang (2012). Improving fairness, efficiency and stability in HTTP-based adaptive video streaming with FESTIVE. In: *Proc. CoNEXT 2012*.
- [19] V. Joseph, G. de Veciana (2012). Jointly optimizing multi-user rate adaptation for video transport over wireless systems: Mean-fairness-variability tradeoffs. In: *Proc. IEEE Infocom 2012*.
- [20] H.J. Kushner, G.G. Yin (2003). *Stochastic Approximation and Recursive Algorithms and Applications*. Springer.
- [21] K. Miller, E. Quacchio, G. Gennari, A. Wolisz (2012). Adaptation algorithm for adaptive streaming over HTTP. In: *Proc. Packet Video Workshop 2012*.
- [22] R.K.P. Mok, E.W.W. Chan, X. Luo, R.K.C. Chang (2011). Inferring the QoE of HTTP video streaming over HTTP. In: *Proc. ACM SIGCOMM W-MUST*.
- [23] R.K.P. Mok, X. Luo, E.W.W. Chan, R.K.C. Chang (2012). QDASH: A QoE-aware DASH system. In: *Proc. ACM MMSys 2012*.
- [24] ORTIVA (2010). Optimizing mobile video.
- [25] G.M. Su, Z. Han, M. Wu., K.J.R. Liu (2007). Joint uplink and downlink optimization for real-time multi-user video streaming over WLAN's. *IEEE Trans. Circ. Syst. Video Techn.*
- [26] G. Tian, Y. Liu (2012). Towards agile and smooth video adaptation in dynamic HTTP streaming. In: *Proc. CoNEXT 2012*.
- [27] H. Viswanathan, D. De Vleeschauwer, A. Beck, S. Benno, R.B. Miller, G. Li, M.M. Clougherty, D.C. Robinson (2013). Mobile video optimization at the base station: adaptive guaranteed bit rate for HTTP adaptive streaming. *Bell Labs Techn. J.* **18** (2), 159–174.
- [28] C. Yim, A.C. Bovik (2011). Evaluation of temporal variation of video quality in packet loss networks (2011). *Image Commun.*