NOVA: QoE-driven Optimization of DASH-based Video Delivery in Networks

Vinay Joseph and Gustavo de Veciana
Department of Electrical and Computer Engineering, The University of Texas at Austin

***Abstract*—We consider the problem of optimizing video delivery for a network supporting video clients streaming stored video. Specifically, we consider the joint optimization of network resource allocation and video quality adaptation. Our objective is to fairly maximize video clients' Quality of Experience (QoE) realizing tradeoffs among the mean quality, temporal variability in quality, and fairness, incorporating user preferences on rebuffering and cost of video delivery. We present a *simple asymptotically optimal online* algorithm, NOVA, to solve the problem. NOVA is *asynchronous*, and using minimal communication, *distributes* the tasks of resource allocation to network controller, and quality adaptation to respective video clients. Video quality adaptation in NOVA is also optimal for standalone video clients, and is well suited for use in the DASH framework. Further, NOVA can be extended for use with more general QoE models, networks shared with other traffic loads and networks using fixed/legacy resource allocation.**

I. Introduction

There has been tremendous growth in video traffic in the past decade. Current trends (see [1]) suggest that mobile video traffic will more than double each year till 2015, with two-thirds of mobile data traffic being video by 2015. It is unlikely that wireless infrastructure can keep up with such growth. Even brute force densification (e.g., using HetNets) would not resolve the problem since variability in throughput would likely worsen due to increased throughput sensitivity to the dynamic number of users sharing an access point and/or dynamic interference. Given these challenges, optimizing video delivery to make the best use of available network resources is one of the critical networking problems today.

We view the video delivery optimization problem for a network as that of fairly maximizing the video clients' QoE subject to network constraints. Here, QoE is a proxy for 'video client satisfaction'. A comprehensive solution to this problem requires two components- a network resource allocation component and a quality adaptation component. The allocation component decides how network resources (e.g., bandwidth, power etc) are allocated to the video clients. The adaptation component decides how the video clients adapt their video quality (or video compression rate) in response to the allocated resources, the nature of the video etc.

We develop a distributed algorithm, Network Optimization for Video Adaptation (NOVA), which jointly optimizes the two components. The adaptation component itself has strong optimality guarantees, and can also be used in standalone video clients. The adaptation component in NOVA can be used with video clients based on the DASH (Dynamic Adaptive Streaming over HTTP) framework ([2]). Under the DASH framework, video is stored as a sequence of short duration (e.g., secs) video segments. Various 'representations' for each segment may be made available by compressing it to different sizes by changing various parameters e.g., quantization, resolution, frame rate etc, where high quality representations of a segment are typically larger in size. Video clients can *adapt* their video quality across segments, i.e., can pick different representations for different segments. The choice of representation can be based on several factors such as the state of the playback buffer, current channel capacity, features of video content being downloaded etc. For instance, the video client can request representations of smaller size to *adapt* to poor channel conditions.

We identify the following four key factors determining the QoE of a video client: (a) average quality, (b) temporal variability in quality, (c) time spent rebuffering (including startup delay), and (d) cost to the video client and video content provider. Our technical focus is on solving the optimization problem given below (formally described in the sequel) which takes these key factors into account:

$$\max \quad \sum_{i\in\mathcal{N}} U_i^E\left(\text{Mean Quality}_i - \text{Quality Variability}_i\right) \qquad (1)$$
$$\text{subject to} \quad \text{Rebuffering}_i,\ \text{Cost}_i, \text{ and Network constraints,}$$

where $\mathcal{N}$ is the set of video clients supported by the network and U_i^E is a 'nice' concave function chosen in accordance with the fairness desired in the network. Network constraint captures *time varying* constraints on network resource allocation allowing us to model wide range variability in resource availability found in real networks.

Let us discuss the four key factors mentioned above. We measure mean quality for a video session as the average across Short Term Quality (STQ) associated with the downloaded representations of the video's segments. STQ of a downloaded segment should ideally capture the viewer's subjective evaluation of the quality of the downloaded representation. In practice, this subjective metric will be measured approximately using objective Video Quality Assessment (VQA) metrics (see [3] for a survey) like PSNR, SSIM, MSSSIM etc. In the sequel, we interchangeably use the terms STQ and quality.

While the benefit of high mean quality is clear, the detrimental impact of temporal variability on QoE (see [4], [5], [6]), and fundamental tradeoff between the average and temporal

- This is an extended version of a paper to appear in INFOCOM 2014.
- This research was supported in part by Intel and Cisco under the VAWN program, and by the NSF under Grant CNS-0917067. We thank Zheng Lu, Xiaoqing Zhu, Chao Chen and Sarabjot Singh for helpful discussions.

variability of quality is often ignored. Indeed [4] suggests that temporal variability in quality can result in a QoE that is *worse* than that of a constant quality video with *lower* average quality. Two prominent sources for such variability are the time varying nature of video content and time varying network capacity. The former can cause time variations in the dependence of STQ on parameters like compression rate, for instance, segments of the same size and duration could have very different STQ, for e.g., consider two such segments where the first segment is of an action scene (where there is a lot of changing visual content) and the second segment is of a slower scene (where things stay the same). Time varying network capacity is especially relevant when considering wireless networks where such variations can be caused by fast fading (on faster time scales, e.g., ms) and slow fading due to shadowing, dynamic interference, mobility, and changing loads (on slower time scales, e.g. secs).

Rebuffering happens when playback buffer of a video client empties, and video playback stalls. Rebuffering events have a significant impact on QoE. Indeed [7] points out that the total time spent rebuffering and the frequency of rebuffering events during a video session can significantly reduce video QoE. In our approach, we impose constraints on the fraction of total time spent rebuffering, and suggest simple ideas to reduce startup delay and the frequency of rebuffering events. We also provide flexibility to the video client in setting these constraints according to its preferences. For instance, a video client which is willing to tolerate rebuffering in return for higher mean quality (for e.g., to watch a movie in high definition over a poor network) can set these constraints accordingly. Such constraints driven by video client preferences will often be content and device dependent, and capture important tradeoffs for the video client. This heterogeneity, which is not really exploited in current solutions, can be a source of significant performance gains.

Client preferences concerning the cost of video delivery could be important when viewers wish to manage their wireless data costs. Note that video content providers may also pay Content Distribution Network operators for the delivery of video data. Thus, if the cost of data delivery is high, higher QoE often comes at higher cost, and the video client/content provider may want to tradeoff QoE versus delivery cost. In our framework, we allow each video client/content provider to set a constraint on the average cost per unit video duration which in turn reflects the desired tradeoff.

A. Main contributions

This paper presents a general optimization framework for stored video delivery optimization that factors heterogeneity in client preferences and QoE models, as well as capacity and video content variability. We develop a *simple online* algorithm NOVA (Network Optimization for Video Adaptation) to solve this multiuser joint resource allocation and quality adaptation problem. The algorithm has been both rigorously analyzed and validated through extensive simulations. NOVA's novelty lies in realizing a comprehensive set of features that meet the challenges of developing next-gen video transport protocols. Key features of NOVA, discussed in more detail in Subsection IV-B, are listed below:

1) *Strong optimality*: guaranteeing that NOVA performs as well as optimal offline scheme which is omniscient, i.e., knows everything about the evolution of channel and video ahead of time.
2) NOVA carries out 'cross-layer' joint optimization of resource allocation and quality adaptation.
3) NOVA is a *simple* and *online* algorithm.
4) NOVA is a *distributed* algorithm where network controller carries out resource allocation and video clients carry out their own quality adaptation.
5) NOVA is an *asynchronous* algorithm well suited for DASH-based video clients where the network controller and video clients operate 'at their own pace'. Value of this asynchrony (and consequential technical challenges) are discussed in Subsection I-B on Related Work.
6) *Suited for current networks:* The resource allocation in NOVA requires just a simple modification of legacy schedulers.
7) *Optimal Adaptation:* Quality adaptation proposed in NOVA is independently optimal and can even be used with a standalone video client, and this optimality is 'insensitive' to network resource allocation.

B. Related work

The problem of video delivery optimization in wireless networks has been studied in many works, for instance, see [8], [9], [10], [11], [12], [13], [14], [15] which utilize extensions of Network Utility Maximization (NUM) framework (see [16]). The main focus of [8] and [9] is real-time interactive video which present the challenge of meeting strict delivery deadlines. Papers [10] and [11] study video delivery optimization in wireless networks considering simpler QoE models, and do not explicitly incorporate rebuffering (nor cost) into their respective optimization frameworks, and instead control rebuffering through network congestion control. Using static QoE models, [13] and [14] study the resource allocation component of video delivery accounting for user dynamics. A major weakness of the aforementioned papers is the limited nature of the associated QoE models (that are essentially just the mean quality) and their lack of flexibility in managing/incorporating user preferences related to rebuffering and cost.

While [12] presents a novel algorithm for realizing mean-variability tradeoffs for video delivery (see [17] for genearalizations), the model involves a strong assumption of synchrony- the download of a segment of each video client starts at the beginning of a (network) slot and finishes by the end of the slot. This assumption on synchrony precludes any explicit control over rebuffering as it limits the ability of a video client to get ahead (by downloading more segments) during periods when channel is good and/or network is underloaded. Relaxed/different versions of this assumption can be found in the theoretical frameworks used in many previous papers (e.g., decision making in [15], [10], [11] is synchronous) as it facilitates an easier extension of tools from classical NUM framework. However, this assumption of synchrony is not ideal for DASH-based video clients in a wireless network that operate 'at their own pace'- downloading variable sized segments (with variable download times) one after the other. In this paper, we drop the assumption of synchrony which allows us to exploit opportunism across video clients' state of playback buffer (channels and features of video content like quality rate tradeoffs), and base our adaptation decision concerning a segment on network state information relevant to the download period of the segment. We also tackle the consequent novel technical challenges related to distributed asynchronous algorithms operating in a stochastic setting. Further, the rebuffering constraint in our asynchronous setting effectively induces a new type of constraint involving averages measured over two time scales.

C. Organization of the paper

Section II introduces the system model and assumptions. We formulate (1)-(2) as an offline optimization problem in Section III. In Section IV, we present an online algorithm NOVA which solves this optimization problem, and discuss its optimality properties. We present a sketch of the proof of optimality of NOVA in Subsection V. We discuss several useful extensions of NOVA in Section VI, present simulation results in VII, and conclude the paper in Section VIII.

II. System model

We first describe some notation used in this paper. We use bold letters to denote vectors. Given a T-length sequence $(a(t))_{1\leq t\leq T}$ or a (infinite) sequence $(a(t))_{t\in\mathbb{N}}$, we let $(a)_{1:T}$ denote the T-length sequence $(a(t))_{1\leq t\leq T}$. For e.g., consider a sequence $(\mathbf{a}(t))_{t\in\mathbb{N}}$ of vectors. Then $(\mathbf{a})_{1:T}$ denotes the T-length sequence containing the first T vectors of the sequence $(\mathbf{a}(t))_{t\in\mathbb{N}}$, and $(a_i)_{1:T}$ denotes the T-length sequence containing ith component of the first T vectors.

To develop our algorithmic framework, let us consider a network serving video to a fixed set of video clients $\mathcal{N}$ where $|\mathcal{N}| = N$. The network operates in a slotted manner with resources allocated for the duration of a slot τ_{slot} seconds. The slots are indexed by $k \in \{0, 1, 2...\}$.

We assume that resource allocation is subject to time varying constraints. In each slot k, a network controller (e.g., base station in a cellular network) allocates $\mathbf{r}_k = (r_{i,k})_{i\in\mathcal{N}} \in \mathbb{R}_+^N$ bits (or $\mathbf{r}_k/\tau_{slot}$ bits per second) to the video clients such that $c_k(\mathbf{r}_k) \leq 0$, where c_k is a real valued function modeling the current constraints on network resource allocation. We refer to c_k as the *allocation constraint* in slot k. This function could be determined by various parameters like video clients' SINR (Signal-to-Interference Noise Ratio). In the sequel, we refer to these functions as allocation constraints. Let C_k denote the random variable corresponding to the allocation constraint in slot k (and c_k is a realization of it). We make the following assumptions on these allocation constraints:

Assumptions C.1-C.3 (Time varying allocation constraints)

C.1 $(C_k)_{k\in\mathbb{N}}$ is a *stationary ergodic* process of functions selected from a set $\mathcal{C}$.

C.2 $\mathcal{C}$ is a (arbitrarily large) finite set of real valued functions on $\mathbb{R}_+^N$, such that each function $c \in \mathcal{C}$ is *convex* and continuously differentiable on an open set containing $[0, r_{\max}]^N$ with $c(\mathbf{0}) \leq 0$ and

$$\min_{\mathbf{r}\in[0,r_{\max}]^N} c(\mathbf{r}) < 0. \tag{2}$$

C.3 The feasible region for each allocation constraint is *bounded*: there is a constant $0 < r_{\max} < \infty$ such that for any $c \in \mathcal{C}$ and $\mathbf{r} \in \mathbb{R}_+^N$ satisfying $c(\mathbf{r}) \leq 0$, we have $r_i \leq r_{\max}$ for each $i \in \mathcal{N}$.

As indicated in Assumption C.1, we model the evolution of the allocation constraints as a stationary ergodic process. Hence, time averages associated with the allocation constraints will converge to their respective statistical averages, and the distribution of the random vector $(C_{k_1+s}, C_{k_2+s}, ..., C_{k_n+s})$ for any choice of indices $k_1, ..., k_n$ does not depend on the shift s, thus the marginal distribution of C_k does not depend on time. We denote the marginal distribution of this process by $(\pi(c))_{c\in\mathcal{C}}$. Without loss of generality, we assume that $\pi^{\mathcal{C}}(c) > 0$ for each $c \in \mathcal{C}$. Note that we are restricting ourselves to settings with convex capacity regions $\left\{(r_{i,k})_{i\in\mathcal{N}} \in \mathbb{R}_+^N : c_k(\mathbf{r}_k) \leq 0\right\}$ due to the convexity assumption in C.2. This model (along with the generalization mentioned in Subsection VI) captures a fairly general class of allocation constraints, including, for example, time-varying capacity constraints associated with bandwidth allocation in wireless networks. We impose an additional requirement on the resource allocation algorithm to ensure that the resource allocation to each video client $i \in \mathcal{N}$ in each slot should be at least $r_{i,\min}$ where $r_{i,\min}$ is a small positive constant. This technical requirement can be relaxed as long as we ensure that each video client can be guaranteed a strictly positive amount of resource allocation over a fixed (large) number of slots.

Segment dependent Quality Rate (QR) tradeoffs: The STQ of a downloaded representation of a segment typically increases with its effective compression rate, i.e., the ratio of the representation's size (which also includes overheads due to metadata etc.) to the duration of the segment. We abstract this relationship using a convex increasing function[1] referred to as a *QR tradeoff*. Note that we are assuming a continuous range of representations, and later address finiteness of the number of representations available in practice.

Each video client downloads segments of its video sequentially, and we index the segments using variables like s, s_i etc taking values in $\{0, 1, 2, ...\}$. Let l_i denote the length (or duration in seconds) of segments of video client i (see extensions to variable sized segments in [18]). Let $f_{i,s}$ denote QR tradeoff associated with the sth segment of video client i. Hence, QR tradeoffs can be user and device (screen size) dependent and further, can be segment dependent varying based on the nature of the segment's video content. For instance, a segment associated with a slow scene (where things stay the same) will typically have a 'steeper' QR tradeoff when

[1] Convexity is typically seen in QR tradeoffs except at very low compression rates, for e.g., see Fig. 1 in [12]. Also, for each segment and effective compression rate, we are implicitly restricting our attention to the representation with highest quality and ignoring less efficient representations

compared to that of an action scene (where there is a lot of changing visual content). Also, let $F_{i,s}$ denote the random variable corresponding to the QR tradeoff associated with the sth segment of video client i. Let $q_{i,s}$ denote the quality (i.e., STQ) associated with the segment s downloaded by video client i. Thus, to obtain a quality $q_{i,s}$ for the sth segment, the size of the segment that has to be downloaded by video client i is $l_i f_{i,s}(q_{i,s})$. Let $q_{\max}$ denote the maximum quality that can achieved in the given network setting which is assumed to be finite. For each video client $i \in \mathcal{N}$, we make the following assumptions on the QR tradeoffs associated with it:

Assumptions QR.1-QR.2 on QR tradeoffs

QR.1 $(F_{i,s})_{s\geq 0}$ is a *stationary ergodic* process taking values in a set $\mathcal{F}_i$.

QR.2 $\mathcal{F}_i$ is a finite set of differentiable *increasing convex* functions defined on an open set containing $[0, q_{\max}]$ such that $\min_{\{f_i \in \mathcal{F}_i\}} f_i(0) > 0$ and $\max_{\{f_i \in \mathcal{F}_i\}} (f_i)'(q_{\max})$ is finite.

As indicated in Assumption QR.1, we model the evolution of QR tradeoffs of each video client $i \in \mathcal{N}$ as a stationary ergodic process. Let $\left(\pi_i^{\mathcal{F}}(f_i)\right)_{f_i \in \mathcal{F}_i}$ denote the associated marginal distribution. Without loss of generality, we assume that $\pi_i^{\mathcal{F}}(f_i) > 0$ for each $f_i \in \mathcal{F}_i$. Let $f_{\min} := \min_{\{i, \in \mathcal{N}, f_i \in \mathcal{F}_i\}} f_i(0)$ which is strictly positive from QR.2, and this gives a lower bound on segment compression rates. Even at zero quality, there is usually overhead information associated with a representation of a segment which causes $f_{\min}$ to be positive. The constant $q_{\max}$ represents the maximum quality that can achieved in the given network setting. Let $f_{\max} := \max_{\{i, \in \mathcal{N}, f_i \in \mathcal{F}_i\}} f_i(q_{\max})$ denote an upper bound on segment compression rates.

QoE model: Our QoE model is a function of the quality of the segment representations, $(q_i)_{1:S}$, downloaded by a video client i on the condition that a rebuffering related constraint (discussed next) is met. While accurate QoE models are typically very complex, we use a simple model motivated by the discussion in Section I and the model proposed in [4]. Let $m_i^S(q_i)$ and $\mathrm{Var}_i^S(q_i)$ denote mean quality and temporal variance in quality respectively associated with the first S segments downloaded by the video client i, i.e.,

$$m_i^S(q_i) := \frac{\sum_{s=1}^S q_{i,s}}{S}, \ \mathrm{Var}_i^S(q_i) := \frac{\sum_{s=1}^S \left(q_{i,s} - m_i^S(q_i)\right)^2}{S}.$$

Note that the arguments of m_i^S and Var_i^S are actually $S-$length sequences $(q_i)_{1:S}$ (i.e., $(q_{i,s})_{1\leq s\leq S}$) although we are using a shorthand for simplicity. We model the QoE of video client i for these S segments as

$$e_i^S(q_i) = m_i^S(q_i) - \eta_i \mathrm{Var}^S(q_i), \tag{3}$$

where $\eta_i > 0$ scales penalty for temporal variability in quality. Also, see [18] for extensions to more general QoE models.

Our objective function capturing video clients' QoE is

$$\phi_S\left((\mathbf{q})_{1:S}\right) := \sum_{i\in\mathcal{N}} e_i^S(q_i). \tag{4}$$

Here, we have set $U_i^E(.)$ appearing in (1) as $U_i^E(e) = e$. In [18], we discuss extensions to concave $U_i^E(.)$ which provides more flexibility in imposing QoE fairness across users, and consider more general variability penalties involving non-linear functions of $\mathrm{Var}^S(q_i)$.

Rebuffering constraints: Let $\kappa > 0$ and let $K_S = \lceil \kappa S \rceil$. We obtain a good estimate for the fraction of time spent rebuffering by a video client under an additional assumption on resource allocation that for each video client i, $\frac{1}{K_S}\sum_{k=1}^{K_S} r_{i,k}$ converges, and hence provides an asymptotically accurate estimate for time-average resource allocation to video client i as S goes to infinity. Note that this condition is satisfied by alpha-fair resource allocation policies like proportionally fair allocation, max-min fair allocation etc under mild assumptions on allocation constraints, for e.g., under stationary ergodic evolution of allocation constraints. Next, note that the cumulative size of the first S segments is given by $\sum_{s=1}^S l_i f_{i,s}(q_{i,s})$. Thus, a good estimate (for large S) for the time required by video client i to download the first S segments is

$$\frac{\sum_{s=1}^S l_i f_{i,s}(q_{i,s})}{\frac{1}{\tau_{slot} K_S}\sum_{k=1}^{K_S} r_{i,k}}$$

which is the ratio of the cumulative size of S segments to the per slot resource allocation estimate. It can be shown (see [18]) that the following expression is an asymptotically (as S goes to infinity) accurate estimate for the percentage of time that video client i is rebuffering while watching the S segments:

$$\beta_{i,S}\left((q_i)_{1:S}, (r_i)_{1:K_S}\right) := \frac{\frac{\sum_{s=1}^S l_i f_{i,s}(q_{i,s})}{\frac{1}{\tau_{slot} K_S}\sum_{k=1}^{K_S} r_{i,k}}}{\sum_{s=1}^S l_i} - 1.$$

The first term in the right hand side is the ratio of the estimate for time required for download of the first S segments to the total duration $\sum_{s=1}^S l_i$ associated with the S segments. Note that $\beta_{i,S}\left((q_i)_{1:S}, (r_i)_{1:K_S}\right)$ can also take negative values which happens when segments are being downloaded at rate higher than the rate at which they are viewed. We express the rebuffering constraint as

$$\beta_{i,S}\left((q_i)_{1:S}, (r_i)_{1:K_S}\right) \leq \overline{\beta}_i, \ \forall\, i \in \mathcal{N}, \tag{5}$$

where each video client i specifies an upper bound $\overline{\beta}_i > -1$ on the fraction of time spent rebuffering. Though setting $\overline{\beta}_i = 0$ ensures that there is only an asymptotically negligible amount of rebuffering, we can enforce more stringent constraints on rebuffering by setting $\overline{\beta}_i$ to negative values. We also discuss simple ideas to reduce startup delay and frequency of rebuffering events after presenting NOVA in Section IV.

Cost constraints: The average compression rate associated with the first S segments of video client $i \in \mathcal{N}$ is $\frac{\sum_{s=1}^S l_i f_{i,s}(q_{i,s})}{\sum_{s=1}^S l_i}$. Let p_i^d denote the cost per unit of data (measured in dollar per bit) that video client $i \in \mathcal{N}$ (or the video content provider associated with the video client) has to pay. Then, the average cost per unit video duration the video client (/content provider) pays is

$$p_{i,S}\left((q_i)_{1:S}\right) := p_i^d \frac{\sum_{s=1}^S l_i f_{i,s}(q_{i,s})}{\sum_{s=1}^S l_i}.$$

We express the cost constraint as

$$p_{i,S}\left((q_i)_{1:S}\right) \leq \overline{p}_i, \ \forall\, i \in \mathcal{N},$$

where each video client i (or the video content provider associated with the video client) sets an upper bound $\overline{p}_i > 0$ on the amount of money per unit video duration.

III. Offline optimization formulation

We formulate the optimization problem in (1)-(2) formally as an offline optimization problem OPT(S) for jointly optimizing quality adaptation (i.e., finding $((q_i)_{1:S})_{i\in\mathcal{N}}$) and resource allocation (i.e., finding $(\mathbf{r})_{1:K_S}$). In the offline setting we assume $(c_k)_k$ and $(f_{i,s})_s$ for each video client $i \in \mathcal{N}$ are known ahead of time.

Based on the discussion in Section II, we rewrite (1)-(2) as the optimization problem OPT(S) given below:

$$\begin{aligned}
\max_{(\mathbf{q})_{1:S},(\mathbf{r})_{1:K_S}} \quad & \phi_S\left((\mathbf{q})_{1:S}\right) \\
\text{subject to} \quad & 0 \le q_{i,s} \le q_{\max} \ \forall\ s \in \{1,...,S\}, \forall\ i \in \mathcal{N}, \\
& r_{i,k} \ge r_{i,\min},\ \forall\ k \in \{1,...,K_S\}, \forall\ i \in \mathcal{N}, \\
& c_k\left(\mathbf{r}_k\right) \le 0, \quad \forall\ k \in \{1,...,K_S\}, \\
& \beta_{i,S}\left((q_i)_{1:S}, (r_i)_{1:K_S}\right) \le \overline{\beta}_i, \forall\ i \in \mathcal{N}, \qquad (6)\\
& p_{i,S}\left((q_i)_{1:S}\right) \le \overline{p}_i, \forall\ i \in \mathcal{N}. \qquad (7)
\end{aligned}$$

We need the following assumption to ensure strict feasibility which will be used in later sections.

Assumption-SF (Strict Feasibility): For each $c \in \mathcal{C}$, $c\left((r_{i,\min})_{i\in\mathcal{N}}\right) < 0$, and for each $i \in \mathcal{N}$, $\max_{\{f_i\in\mathcal{F}_i\}} \frac{\tau_{slot} f_i(0)}{r_{i,min}} < 1$, and $p_i^d \max_{\{f_i\in\mathcal{F}_i\}} f_i\left(0\right) < \overline{p}_i$.

This assumption[2] requires that the resource allocation $(r_{i,\min})_{i\in\mathcal{N}}$ is strictly feasible for any $c \in \mathcal{C}$, and that the maximum size of segments at zero quality is not too large.

We assume that the optimization problem OPT(S) is feasible (sufficient conditions are discussed in [18]). Let ϕ_S^{opt} denote the optimal value of objective function of OPT(S).

In practice, solving OPT(S) directly is impossible (except for trivial cases) since we need to know $(c_k)_k$ and $(f_{i,s})_s$ ahead of time. Further, it is also computationally prohibitive as the optimization would be over $O(NS)$ variables. Thus, from a practical point of view, the main challenge is to overcome these two hurdles and obtain a *simple* and *online* algorithm that performs as well as ϕ_S^{opt} asymptotically.

IV. A simple online algorithm for jointly optimizing allocation and adaptation

The algorithm NOVA comprises three components:

1) Allocate: Network resource allocation is done by the network controller at the beginning of each slot k by solving an optimization problem RNOVA($\mathbf{b}_k, c_k$) which depends on the parameter $\mathbf{b}_k$ (described below) and the allocation constraint c_k in the slot.

2) Adapt: When a video client $i \in \mathcal{N}$ completes downloading the s_ith segment, the video client selects the quality/representation for the next segment by solving an optimization problem QNOVA$_i(\boldsymbol{\theta}_{i,s_i}, f_{i,s_i+1})$ which depends on a parameter $\boldsymbol{\theta}_{i,s_i}$ (described later in the section) and the QR tradeoff f_{i,s_i+1} of the next segment.

3) Learn: involves learning parameters $(m_{i,s_i}, b_{i,k}, d_{i,s_i})_{i\in\mathcal{N}}$ used in the optimization problems RNOVA($\mathbf{b}_k, c_k$) and QNOVA$_i(\boldsymbol{\theta}_{i,s_i}, f_{i,s_i+1})$. Here s_i is the current segment index of video client i and k is the current slot index. The parameter m_{i,s_i} tracks mean quality of video client $i \in \mathcal{N}$. Parameters $b_{i,k}$ and d_{i,s_i} serve as indicators of risk of violation of rebuffering constraints (6) and cost constraints (7) respectively of video client $i \in \mathcal{N}$, and larger the parameter, larger the risk. We later see that, for $\overline{\beta}_i = 0$, the value of $b_{i,k}$ reflects the duration of video content in video client i's playback buffer (and is roughly a linear decreasing function of this duration). The parameters $(m_{i,s_i}, b_{i,k}, d_{i,s_i})$ are learnt/updated by video client i for each $i \in \mathcal{N}$, and the network controller only uses $\mathbf{b}_k$ for carrying out resource allocation in slot k.

For $\mathbf{b} \in \mathbb{R}^N$ and allocation constraint $c \in \mathcal{C}$, the (convex) optimization problem RNOVA($\mathbf{b}, c$) associated with network resource allocation is:

$$\max_{\mathbf{r}} \left\{ \sum_{i\in\mathcal{N}} h_i^B\left(b_i\right) r_i :\ c\left(\mathbf{r}\right) \le 0,\ r_i \ge r_{i,\min}\ \forall i \in \mathcal{N} \right\} \qquad (8)$$

where $h_i^B(.)$ is a non-negative valued Lipschitz continuous function such that $\lim_{b\to\infty} h_i^B(b) = \infty$, $h_i^B(b_i) = 0$ for all $b_i \le \underline{b}$ for some constant $\underline{b}$ (typically set as zero or small negative numbers), and is strictly increasing for $b_i \ge \underline{b}$. Simple examples of functions satisfying these conditions are $\max(b, 0)$, $\max(b^2, 0)$ etc. Let $\mathcal{R}^*\left(\mathbf{b}, c\right)$ denote the set of optimal solutions to RNOVA($\mathbf{b}, c$). When using RNOVA($\mathbf{b}, c$), we will set $\mathbf{b}$ as the current value of the rebuffering risk indicator $\mathbf{b}_k$. Hence, the objective function (8) gives more weight to video clients with a higher value of $b_{i,k}$ i.e., higher risk of violation of rebuffering constraints.

Let $m_i \in [0, q_{\max}],\ b_i, d_i \in \mathbb{R}$ and $\boldsymbol{\theta}_i = (m_i, b_i, d_i)$. For QR tradeoff f_i, let

$$\begin{aligned}
\phi^Q\left(q_i, \boldsymbol{\theta}_i, f_i\right) \quad = \quad & q_i - \eta_i\left(q_i - m_i\right)^2 \qquad (9)\\
& - \frac{h_i^B\left(b_i\right)}{\left(1 + \overline{\beta}_i\right)} f_i\left(q_i\right) - \frac{p_i^d h_i^D\left(d_i\right)}{\overline{p}_i} f_i\left(q_i\right),
\end{aligned}$$

where $h_i^D(.)$ satisfies conditions given for $h_i^B(.)$ with $\underline{b}$ replaced by $\underline{d}$ (also set as zero or a small negative number). The optimization problem QNOVA$_i(\boldsymbol{\theta}_i, f_i)$ associated with quality adaptation of video client i is given below:

$$\max_{q_i} \left\{ \phi^Q\left(q_i, \boldsymbol{\theta}_i, f_i\right) :\ 0 \le q_i \le q_{\max} \right\}.$$

When using QNOVA$_i(\boldsymbol{\theta}_i, f_i)$ in NOVA, we will use $\boldsymbol{\theta}_i = (m_{i,s}, b_{i,k+1}, d_{i,s})$ so that the objective function (9) includes a term $\left(q_i - m_{i,s}\right)^2$ ensuring that an optimal solution to QNOVA$_i(\boldsymbol{\theta}_i, f_i)$ is not too far from $m_{i,s}$ (current estimate of mean quality), and thus avoids high variance in quality. Further, the terms $\frac{h_i^B(b_{i,k+1})}{(1+\overline{\beta}_i)} f_i\left(q_i\right)$ and $\frac{p_i^d h_i^D(d_{i,s})}{\overline{p}_i} f_i\left(q_i\right)$ in (9) penalize quality choices leading to large segment sizes when $b_{i,k+1}$ or $d_{i,s}$ are high, and thus ensure that NOVA reacts to indicators of increased risk of violation of rebuffering constraints and cost constraints. Also note that we can control the

[2] The assumption requires a uniform upper bound on the size of the segments at zero quality which is used in Lemma 1. We conjecture that this per segment requirement can be replaced with a milder averaged version.

response of NOVA to these indicators by appropriately choosing $\left(h_i^B(.)\right)_{i\in\mathcal{N}}$ and $\left(h_i^D(.)\right)_{i\in\mathcal{N}}$. The optimization problem QNOVA$_i(\boldsymbol{\theta}_i, f_i)$ is convex with strictly concave objective function, and thus has a unique solution denoted as $q_i^*\left(\boldsymbol{\theta}_i, f_i\right)$.

Next, we present the algorithm NOVA. Let s_i be an indexing variable keeping track of the segment video client i is currently downloading. Let $\epsilon > 0$,

$$\mathcal{H}^{(i)} = \left\{(m_i, b_i, d_i) \in \mathbb{R}^3 : 0 \le m_i \le q_{\max},\ b_i \ge \underline{b},\ d_i \ge \underline{d}\right\},$$

and let $[x]_y = \max(x, y)$ for $x, y \in \mathbb{R}$. Also, assume that all video clients have already downloaded the 0th segment at the beginning of slot $k = 0$. The algorithm NOVA is given below.

NOVA

Initialization: Let $(m_{i,0}, b_{i,0}, d_{i,0}) \in \mathcal{H}^{(i)}$ for each $i \in \mathcal{N}$.

In each slot $k \ge 0$, carry out the following steps:

ALLOCATE: At the beginning of slot k, network controller allocates resources $\mathbf{r}_k^*$ choosing any solution to RNOVA$(\mathbf{b}_k, c_k)$. Update $\mathbf{b}_k$ as follows:

$$b_{i,k+1} = b_{i,k} + \epsilon\left(\frac{\tau_{slot}}{(1+\overline{\beta}_i)}\right). \tag{10}$$

ADAPT: In slot k, if any video client $i \in \mathcal{N}$ finishes download of s_i th segment, let $\boldsymbol{\theta}_{i,s_i} = (m_{i,s_i}, b_{i,k+1}, d_{i,s_i})$. For segment $s_i + 1$ of video client i, the video client selects representation with quality $q_i^*(\boldsymbol{\theta}_{i,s_i}, f_{i,s_i+1})$ (i.e., optimal solution to QNOVA$_i(\boldsymbol{\theta}_{i,s_i}, f_{i,s_i+1})$), denoted as q_{i,s_i+1}^* for brevity, and update parameters m_{i,s_i+1}, $b_{i,k+1}$, d_{i,s_i+1} and s_i as follows:

$$m_{i,s_i+1} = m_{i,s_i} + \epsilon\left(q_{i,s_i+1}^* - m_{i,s_i}\right), \tag{11}$$

$$b_{i,k+1} = \left[b_{i,k+1} - \epsilon\left(l_i\right)\right]_{\underline{b}}, \tag{12}$$

$$d_{i,s_i+1} = \left[d_{i,s_i} + \epsilon\left(p_i^d \frac{l_i f_{i,s_i+1}\left(q_{i,s_i+1}^*\right)}{\overline{p}_i} - l_i\right)\right]_{\underline{d}}, \tag{13}$$

$$s_i = s_i + 1.$$

For each $i \in \mathcal{N}$, parameters $(m_{i,s_i}, b_{i,k}, d_{i,s_i})$ are learnt/updated by video client i. The network controller only needs to know $\mathbf{b}_k$ for carrying out resource allocation in slot k and this can be achieved using minimal signaling as described in subsection IV-B. Under NOVA, allocation is done at the beginning of each slot whereas adaptation is asynchronous, i.e., adaptation related decisions about a segment are made by a video client only at the completion of download of previous segment. The update equation (11) associated with the parameter m_{i,s_i} is similar to update rules used for tracking EWMA (Exponentially Weighted Moving Averages), and ensures that m_{i,s_i} tracks the mean quality of video client i. Consider the evolution of the parameter $b_{i,k}$ which is updated in both (10) and (12) ignoring the operator $[.]_{\underline{b}}$ and setting initialization to zero. (10) ensures that $b_{i,k}$ is increased by fixed amount $\frac{\epsilon\tau_{slot}}{(1+\overline{\beta}_i)}$ at the beginning of each slot. (12) ensures that when a video client completes the download of a segment, $b_{i,k}$ is reduced by ϵ times the duration of the next segment. Hence, at some time t seconds (or $k = t/\tau_{slot}$ slots) after starting the video,

$$\frac{b_{i,k} - b_{i,0}}{\epsilon} \approx \frac{t}{(1+\overline{\beta}_i)} - L_i^D(t),$$

where $L_i^D(t)$ is the duration of video downloaded up to time t. This sheds light on the role of $b_{i,k}$ as an indicator of risk of violation of rebuffering constraint in (6) for video client i. In particular, we see that for $\overline{\beta}_i = 0$ and small enough $\underline{b}$, $(b_{i,k} - b_{i,0})/\epsilon$ is equal to $(t - L_i^D(t))$ which is equal to *negative* of the duration of video content in playback buffer (if there is any). Similarly, we can argue that d_{i,s_i} serves as an indicator of risk of violation of cost constraint (7) for video client i.

Note that a large value of $b_{i,k}$ results in the selection of a representation of smaller size (see (9)). This combined with the role of $b_{i,k}$ discussed above and the fact that NOVA satisfies the rebuffering constraint (5) asymptotically (see Theorem 1 (a)) suggests that NOVA aims to meet the rebuffering constraint (5) for finite S also. Further, start up delays can be reduced by appropriately choosing the initial conditions, e.g. pick large $b_{i,0}$ and small $m_{i,0}$ to encourage selection of representations with smaller size in the beginning so that they are downloaded quickly. Also, the frequency of rebuffering events can be reduced by forcing the video client to delay the resumption of playback after a rebuffering event until there is sufficient amounts of video content in the playback buffer. Also, note that although we have not explicitly incorporated the possibility of packet losses (in wireless networks, routers in wired networks etc) into our theoretical framework, the simplicity of quality adaptation in NOVA allows it operate in such settings as it does not rely on such 'low-level' network information and *only relies on a 'high level' view of the network encapsulated in segment download completions*.

It is interesting to note that the quality adaptation proposed in NOVA does not directly use any information about the allocation constraints. Neither does the resource allocation directly use any information about QR tradeoffs of the video clients. Yet, the joint resource allocation and quality adaptation under NOVA has strong optimality properties (which are presented later in this section). This is mainly due to the fact that the variables $(b_{i,k})_{i\in\mathcal{N}}$ carry almost all the information about the video clients' quality adaptation that is required by the network controller to carry out optimal resource allocation, and the variable $b_{i,k}$ carries almost all the information that the quality adaptation at video client i needs to know about the resource allocation (to the client). For e.g., consider a video client i in the network that has very few unwatched segments in the playback buffer, i.e., the video client is about to experience rebuffering. We see that the update rules for $b_{i,k}$ (and a large enough initialization) ensure that $b_{i,k}$ will be large in this scenario, and this forces the video client and the network controller to make the right moves, i.e., this forces the video client to switch to low quality representations (accounting for current QR tradeoffs), and forces the network controller to give higher priority to this video client in the resource allocation (accounting for allocation constraints).

A. Optimality of NOVA

The following theorem is the main optimality result for NOVA, and we discuss key steps of our proof in Section V.

Theorem 1. *Suppose* $(C_k)_{k\ge0}$ *and* $(F_{i,s})_{s\ge0}$ *are* stationary ergodic *processes for each* $i \in \mathcal{N}$*. Then,* *(a)* Feasibility: *NOVA asymptotically satisfies the constraints on rebuffering and cost, i.e., for each $i \in \mathcal{N}$*

$$limsup_{S\to\infty}\beta_{i,S}\left((q_i^*)_{1:S}, (r_i^*)_{1:K_S}\right) \leq \overline{\beta}_i, \quad (14)$$
$$limsup_{S\to\infty}p_{i,S}\left((q_i^*)_{1:S}\right) \leq \overline{p}_i. \quad (15)$$

(b) Optimality: *Let $S_\epsilon = \frac{S}{\epsilon}$. Then,*

$$\lim_{S\to\infty}\lim_{\epsilon\to 0}\left(\phi_{S_\epsilon}\left((\mathbf{q}^*)_{1:S_\epsilon}\right) - \phi_{S_\epsilon}^{opt}\right)$$

converges to zero in probability.

Here C_k and $F_{i,s}$ are random variables corresponding to c_k and $f_{i,s}$ respectively. Recall that, under NOVA, q_{i,s_i}^* is the quality associated with segment s_i of video client i (and the notation used in this result is described at the beginning of Section II). This result tells us that the difference in performance (according to definition (4)) of the *online* algorithm NOVA (i.e., $\phi_{S_\epsilon}\left((\mathbf{q}^*)_{1:S_\epsilon}\right)$) and that of the optimal *offline* scheme goes to zero for long enough videos and small enough ϵ. Recall that $\phi_{S_\epsilon}^{opt}$ is the optimal value of OPT(S_ϵ), i.e., the performance of the optimal omniscient offline scheme which knows all the allocation constraints $(c_k)_k$ and QR tradeoffs $(f_{i,s})_s$ ahead of time.

B. Key features and Implementation of NOVA

Next, we summarize the key features of NOVA.

***Optimality*:** NOVA carries out 'cross-layer' joint optimization of resource allocation and quality adaptation, with strong optimality guarantees (given in Theorem 1).

***Online*:** NOVA is an online algorithm as it only uses *current* information, i.e., network controller only needs to know the allocation constraint c_k to carry out resource allocation for slot k, and video client i only requires the QR tradeoff $f_{i,s}$ for quality adaptation of segment s.

***Simple*:** RNOVA$(\mathbf{b}, c)$ is an $N-$variable convex optimization problem, which becomes an even simpler linear program under linear allocation constraints (often this linear program has enough structure to allow for very efficient solution techniques). Also, note that QNOVA$_i(\boldsymbol{\theta}_i, f_i)$ is just a scalar convex optimization problem.

***Asynchronous and well suited for DASH*:** The asynchronous nature of NOVA ensures that the video clients can work at their own pace and the adaptation prescribed in NOVA is entirely *client driven* requiring no assistance from the network controller, and is thus well suited for DASH framework.

***Distributed implementation and information flow*:** NOVA can be implemented in a *distributed* manner with minimal signaling since quality adaptation is *client driven* and for the resource allocation, the network controller need only know $\mathbf{b}_k$. To ensure that the network controller knows the current value of rebuffering risk indicator vector $\mathbf{b}_k$, each video client can send a signal to the base station indicating the latest value of $b_{i,k}$ (just a signal indicating segment download completion is enough) at the end of each segment download which usually occurs at a low frequency (typically once a second). On receiving this signal from video client $i \in \mathcal{N}$, the network controller can then update $b_{i,k}$. Now, until the next signal from video client i, the network controller can update $b_{i,k}$ using (10) that requires only constant increments. The network controller could obtain information about allocation constraints through Channel Quality Information (CQI) feedback from the network, and video clients could obtain their respective QR tradeoffs using application layer information exchange. The flow of information across various layers of the network for this implementation of NOVA is depicted in Fig. 1. Note that we do not even need this signaling if the network controller could identify segment download completions on its own (for e.g., using deep packet inspection).

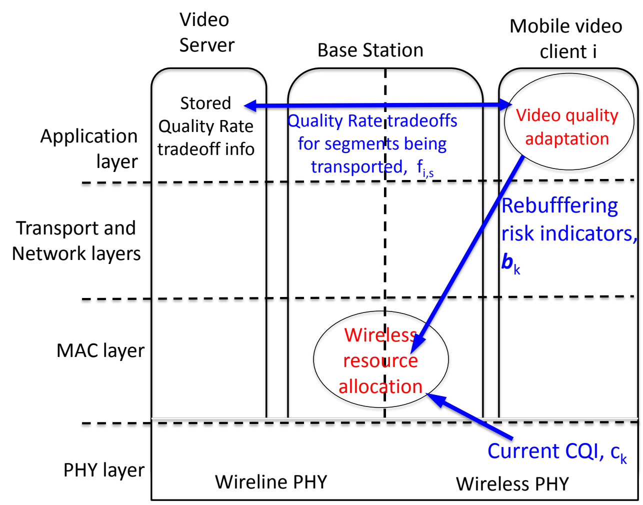

Fig. 1. Cross Layer Information Flow

***Optimal Adaptation*:** The adaptation proposed in NOVA is independently optimal, and the optimality properties of the adaptation component of NOVA is 'insensitive' to the resource allocation component, i.e., does not depend on detailed characteristics (for e.g., the specific resource allocation algorithm, time scale of operation etc) of the latter. See [18] for a detailed discussion of this property. As a corollary of this property, we have that the adaptation proposed in NOVA (which is well suited for DASH based video clients) is also optimal for standalone video clients.

***Well suited for legacy networks*:** Optimization algorithm for resource allocation, RNOVA$(\mathbf{b}, c)$ requires only a simple modification of legacy schedulers like proportionally fair schedulers (see [19]). This is clear on comparing (8) and (39) (which is discussed later).

V. Proof of Optimality of NOVA

This section is devoted to a discussion of the proof of the previously stated Theorem 1 related to optimality of NOVA focusing on the key intermediate results used in the proof. Due to space constraints, we have omitted detailed proofs of these intermediate results which can be found in [18]. We start this section with a discussion about some useful properties of NOVA. In Subsection V-A, we study an auxiliary optimization problem OPTSTAT and obtain Theorem 2 which suggests that we can prove the main optimality result Theorem 1 for NOVA if we establish an appropriate convergence result for NOVA's parameters. In Subsection V-B, we study an auxiliary differential inclusion (given in (31)-(36)) which evolves according to average dynamics of NOVA, and obtain a convergence result for the differential inclusion. In Subsection V-C, we view NOVA's update equations ((11)-(13) and (27)-(29)) as an asynchronous stochastic approximation update (see, e.g., [20] for reference), and relate this stochastic approximation update to the auxiliary differential inclusion (in (31)-(36)), and use this relationship to establish desired convergence of NOVA's parameters using the convergence result for the auxiliary differential inclusion established in Subsection V-B.

Next, we discuss some useful properties of NOVA. The optimization problem RNOVA$(\mathbf{b}, c)$ is convex, and using Assumption-SF, we can show that it satisfies Slater's condition (see [21] for reference). Thus, KKT conditions are necessary and sufficient for optimality. The optimization problem QNOVA$_i(\boldsymbol{\theta}_i, f_i)$ is also convex and satisfies Slater's condition (since the constraints are all linear), and thus, KKT conditions are necessary and sufficient for optimality.

The next result states that the parameters in NOVA stay in a compact set and in particular, points out that the parameters $b_{i,k}$ and $d_{i,s}$ can be uniformly bounded.

Lemma 1. *For any initialization $(m_{i,0}, b_{i,0}, d_{i,0})_{i\in\mathcal{N}} \in \prod_{i\in\mathcal{N}} \mathcal{H}^{(i)}$, the parameters evolving according to NOVA satisfy the following: for each $i \in \mathcal{N}$, $s \geq 1$ and $k \geq 1$, we have $0 \leq m_{i,s} \leq q_{\max}$, $\underline{b} \leq b_{i,k} \leq \overline{b}$ and $\underline{d} \leq d_{i,s} \leq \overline{d}$ for some finite constants $\overline{b}$ and $\overline{d}$ and for all k and s large enough.*

For the next two results, let $\boldsymbol{\theta}_i = (m_i, b_i, d_i)$ where $0 \leq m_i \leq q_{\max}$ and $b_i, d_i \in \mathbb{R}$. The next result provides smoothness properties for the optimal solutions of RNOVA$(\mathbf{b}, c)$ and QNOVA$_i(\boldsymbol{\theta}_i, f_i)$.

Lemma 2. *(a) For each $i \in \mathcal{N}$ and $f_i \in \mathcal{F}_i$, $q_i^*(\boldsymbol{\theta}_i, f_i)$ is a continuous function of $\boldsymbol{\theta}_i$.*
(b) For each $c \in \mathcal{C}$, $\mathcal{R}^*(\mathbf{b}, c)$ is a convex and compact set. Further, $\mathcal{R}^*(\mathbf{b}, c)$ is an upper semi-continuous set valued map of $\mathbf{b}$.
(c) For each $c \in \mathcal{C}$ and $\mathbf{r}^*(\mathbf{b}, c) \in \mathcal{R}^*(\mathbf{b}, c)$, $\phi^R(\mathbf{r}^*(\mathbf{b}, c), \mathbf{b})$ is a continuous function of $\mathbf{b}$.

In the next result, we discuss concavity and differentiability properties of the optimal value of QNOVA$_i(\boldsymbol{\theta}_i, f_i)$.

Lemma 3. *The following statements hold for each $i \in \mathcal{N}$ and $f_i \in \mathcal{F}_i$.*
(a) The optimal value of QNOVA$_i(\boldsymbol{\theta}_i, f_i)$, i.e., $\phi^Q(q_i^*(\boldsymbol{\theta}_i, f_i), \boldsymbol{\theta}_i, f_i)$, is a strictly concave function of m_i (with b_i and d_i fixed).
(b) The partial derivative of $\phi^Q(q_i^*(\boldsymbol{\theta}_i, f_i), \boldsymbol{\theta}_i, f_i)$ with respect of m_i is given by:

$$\frac{\partial \phi^Q(q_i^*(\boldsymbol{\theta}_i, f_i), \boldsymbol{\theta}_i, f_i)}{\partial m_i} = 2\eta_i (q_i^*(\boldsymbol{\theta}_i, f_i) - m_i). \tag{16}$$

(c) Let $\boldsymbol{\theta}_i^{(m)} = (m, b_i, d_i)$, i.e., $\boldsymbol{\theta}_i$ with the first component set to m. If $m \neq m_i$, the optimal value of QNOVA$_i(\boldsymbol{\theta}_i^{(m)}, f_i)$ satisfies

$$\begin{aligned}\phi^Q\left(q_i^*\left(\boldsymbol{\theta}_i^{(m)}, f_i\right), \boldsymbol{\theta}_i^{(m)}, f_i\right) < \phi^Q(q_i^*(\boldsymbol{\theta}_i, f_i), \boldsymbol{\theta}_i, f_i)\\ +2\eta_i(m - m_i)(q_i^*(\boldsymbol{\theta}_i, f_i) - m_i).\end{aligned}$$

A. NOVA, under stationary ergodic regime, is optimal if its parameters are picked from an optimal parameter set

In this section, we use the fact that the underlying allocation constraints and QR tradeoffs are drawn from stationary ergodic processes to show that the offline optimization problem OPT(S) has an 'asymptotically' optimal solution which corresponds to a *stationary* policy– a policy for which the allocation and quality adaptation decisions depend solely on the *current* state determined by the current allocation constraint and QR tradeoffs. Additionally, we establish a useful relationship (in Theorem 2) between such an 'optimal' stationary policy and NOVA that the former can be obtained by using RNOVA$(\mathbf{b}, c)$ for allocation and QNOVA$_i(\boldsymbol{\theta}_i, f_i)$ for quality adaptation if the parameters driving the allocation and adaptation (i.e., $\boldsymbol{\theta}_i$ for all i which also includes $\mathbf{b}$) are selected from an 'optimal' set of parameters.

The offline optimization formulation OPT(S) mainly involves time and segment averages of various quantities. By contrast, the formulation of OPTSTAT discussed in this section is based on the expected value of the corresponding quantities evaluated under the stationary distribution of $(C_k)_k$ and $(F_{i,s})_{s\geq 0}$ for each $i \in \mathcal{N}$. Recall (see Section II) that $(C_k)_k$ is a stationary ergodic random process with marginal distribution $\left(\pi^{\mathcal{C}}(c)\right)_{c\in\mathcal{C}}$. We let C^π denote a random variable with distribution $\left(\pi^{\mathcal{C}}(c) : c \in \mathcal{C}\right)$. Also, recall that for each $i \in \mathcal{N}$, $(F_{i,s})_{s\geq 0}$ is a stationary ergodic process with marginal distribution $\left(\pi_i^{\mathcal{F}}(f_i)\right)_{f_i\in\mathcal{F}_i}$. We let F_i^π denote a random variable with distribution $\left(\pi_i^{\mathcal{F}}(f_i)\right)_{f_i\in\mathcal{F}_i}$.

Consider a stationary policy with $(\mathbf{r}(c))_{c\in\mathcal{C}}$ being a vector (of vectors) representing the allocation $\mathbf{r}(c)$ $(\in \mathbb{R}^N)$ for each $c \in \mathcal{C}$. Though we are abusing earlier notation where $\mathbf{r}(t)$ denoted the allocation to the video clients in slot t, one can differentiate between the functions based on the context in which they are being discussed. Also, for the stationary policy, let $q_i(f)$ denote the quality associated with a segment of video client i with $f \in \mathcal{F}_i$. Mimicking the definition of $\phi_S((\mathbf{q})_{1:S})$, $m_i^S(q_i)$ and $\mathrm{Var}_i^S(q_i)$ in Section III, we let

$$\begin{aligned}\phi_\pi\left(\left((q_i(f_i))_{f_i\in\mathcal{F}_i}\right)_{i\in\mathcal{N}}\right) =\\ \sum_{i\in\mathcal{N}}\left(\mathrm{Mean}(q_i(F_i^\pi)) - \eta_i \mathrm{Var}(q_i(F_i^\pi))\right),\end{aligned} \tag{17}$$

where

$$\begin{aligned}\mathrm{Mean}(q_i(F_i^\pi)) &= \mathbb{E}[q_i(F_i^\pi)],\\ \mathrm{Var}(q_i(F_i^\pi)) &= \mathbb{E}\left[(q_i(F_i^\pi) - \mathrm{Mean}(q_i(F_i^\pi)))^2\right].\end{aligned}$$

Now, consider the optimization problem OPTSTAT given below:

$$\max_{\left((q_i(f_i))_{f_i\in\mathcal{F}_i}\right)_{i\in\mathcal{N}},(\mathbf{r}(c))_{c\in\mathcal{C}}} \phi_\pi\left(\left((q_i(f_i))_{f_i\in\mathcal{F}_i}\right)_{i\in\mathcal{N}}\right) \quad (18)$$

$$\begin{aligned}
\text{subject to } c(\mathbf{r}(c))) &\leq 0,\ \forall\ c\in\mathcal{C}, && (19)\\
0\leq q_i(f_i) &\leq q_{\max},\ \forall\ f_i\in\mathcal{F}_i,\ \forall\ i\in\mathcal{N}, &&\\
r_i(c) &\geq r_{i,\min},\ \forall\ c\in\mathcal{C},\ \forall\ i\in\mathcal{N}, &&\\
p_i^d\frac{\mathbb{E}\left[F_i^\pi(q_i(F_i^\pi))\right]}{\overline{p}_i} &\leq 1,\ \forall\ i\in\mathcal{N}, && (20)\\
\frac{\mathbb{E}\left[F_i^\pi(q_i(F_i^\pi))\right]}{\left(1+\overline{\beta}_i\right)} &\leq \frac{\mathbb{E}\left[r_i(C^\pi)\right]}{\tau_{slot}},\ \forall\ i\in\mathcal{N}. && (21)
\end{aligned}$$

We obtained the above formulation by replacing the time and segment averages of various quantities in OPT(S) with the expected value of the corresponding quantities. Note that in the constraint $c(\mathbf{r}(c))) \leq 0$ given in (19), c appearing as argument of $\mathbf{r}(c)$ is an index (for the corresponding element in $\mathcal{C}$) whereas the other c is the associated function. Similarly, in the term $F_i^\pi(q_i(F_i^\pi))$, the argument F_i^π serves as an index whereas $F_i^\pi(.)$ is the (random) function.

We can show that OPTSTAT is a convex optimization problem satisfying Slater's condition. Further, we can show that the optimal quality choices obtained by solving OPTSTAT are unique and we denote them by $\left((q_i^\pi(f))_{f\in\mathcal{F}_i}\right)_{i\in\mathcal{N}}$. Let $\left(\left((q_i^\pi(f))_{f\in\mathcal{F}_i}\right)_{i\in\mathcal{N}},(\mathbf{r}^\pi(c))_{c\in\mathcal{C}}\right)$ be an optimal solution to OPTSTAT, and let $\mathbf{b}^\pi$ and $\mathbf{d}^\pi$ denote the associated Lagrange multipliers for the constraints (20) and (21) respectively. Since OPTSTAT is a convex optimization problem satisfying Slater's condition, we can conclude that the KKT conditions are necessary and sufficient for optimality. For each $i\in\mathcal{N}$, let

$$\begin{aligned}
m_i^\pi &= \mathbb{E}\left[q_i^\pi(F_i^\pi)\right], && (22)\\
v_i^\pi &= \mathrm{Var}\left(q_i^\pi(F_i^\pi)\right), && (23)\\
\sigma_i^\pi &= \mathbb{E}\left[F_i^\pi(q_i^\pi(F_i^\pi))\right]. && (24)
\end{aligned}$$

Thus m_i^π, v_i^π and σ_i^π are the (statistical) mean quality, variance in quality and mean segment size for video client i associated with optimal solution to OPTSTAT. Also, let

$$\begin{aligned}
\mathcal{X}^\pi \;=\; &\{(\boldsymbol{\rho}^\pi,\mathbf{b}^\pi,\mathbf{d}^\pi) : \text{there is an optimal solution} \qquad (25)\\
&\left(\left((q_i^\pi(f))_{f\in\mathcal{F}_i}\right)_{i\in\mathcal{N}},(\mathbf{r}^\pi(c))_{c\in\mathcal{C}}\right) \text{ to OPTSTAT with}\\
&\rho_i^\pi=\mathbb{E}\left[r_i^\pi(C^\pi)\right] \text{ for each } i\in\mathcal{N}, \text{ and with}\\
&\mathbf{b}^\pi \text{ and } \mathbf{d}^\pi \text{ as the associated optimal Lagrange multipliers}\\
&\text{for constraints (20) and (21) respectively}\}.
\end{aligned}$$

In the next result, we present three useful properties of any optimal solution to OPTSTAT. The result in part (a) below provides a video client level optimality result which essentially suggests that we can decouple the quality adaptation of the video clients. It states that the component $(q_i^\pi(f))_{f\in\mathcal{F}_i}$ of the optimal solution to OPTSTAT associated with video client $i\in\mathcal{N}$ is itself an optimal solution to an optimization problem which can be solved by the video client i. This result hints at the possibility of distributing the task of quality adaptation across the video clients so that each video client manages its own adaptation. The result in part (b) points out that we only need to know a few parameters (specifically, the optimal Lagrange multipliers associated with the rebuffering constraints) associated with the quality adaptation to carry out optimal resource allocation. This suggests that we could potentially decouple the task of optimal resource allocation from quality adaptation. Part (c) states that that when NOVA parameter $\boldsymbol{\theta}_{i,s}$ of video client i is in the set $\mathcal{H}_i^*$ defined below

$$\mathcal{H}_i^* := \left\{\left(m_i^\pi,\left(h_i^B\right)^{-1}(b_i^\pi),\left(h_i^D\right)^{-1}(d_i^\pi)\right) : (\boldsymbol{\rho}^\pi,\mathbf{b}^\pi,\mathbf{d}^\pi)\in\mathcal{X}^\pi\right\}, \quad (26)$$

NOVA can provide optimal quality choices for OPTSTAT.

Lemma 4. *For parts (a) and (b) of this result, suppose* $(\boldsymbol{\rho}^\pi,\mathbf{b}^\pi,\mathbf{d}^\pi)\in\mathcal{X}^\pi$ *and let the associated optimal solution be* $\left(\left((q_i^\pi(f))_{f\in\mathcal{F}_i}\right)_{i\in\mathcal{N}},(\mathbf{r}^\pi(c))_{c\in\mathcal{C}}\right)$.

(a) For each $i\in\mathcal{N}$, $(q_i^\pi(f))_{f\in\mathcal{F}_i}$ *is the unique optimal solution to the following optimization problem*

$$\begin{aligned}
\max_{\left((q_i(f))_{f\in\mathcal{F}_i}\right)} \quad & \mathbb{E}\left[q_i(F_i^\pi)\right]-\eta_i \mathit{Var}\left(q_i(F_i^\pi)\right)\\
& -\sum_{i\in\mathcal{N}} d_i^\pi\left(\frac{p_i^d}{\overline{p}_i}\right)\mathbb{E}\left[F_i^\pi(q_i(F_i^\pi))\right]\\
& -\sum_{i\in\mathcal{N}}\frac{b_i^\pi}{\left(1+\overline{\beta}_i\right)}\mathbb{E}\left[F_i^\pi(q_i(F_i^\pi))\right],\\
\text{s.t.} \quad & 0\leq q_i(f)\leq q_{\max},\ \forall\ f\in\mathcal{F}_i.
\end{aligned}$$

(b) $(\mathbf{r}^\pi(c))_{c\in\mathcal{C}}$ *is an optimal solution to the following optimization problem*

$$\begin{aligned}
&\mathbb{E}\left[\sum_{i\in\mathcal{N}} b_i^\pi r_i(C^\pi)\right],\\
\text{s.t. } c(\mathbf{r}(c))) &\leq 0,\ \forall\ c\in\mathcal{C},\\
r_i(c) &\geq r_{i,\min},\ \forall\ c\in\mathcal{C},\ \forall\ i\in\mathcal{N}.
\end{aligned}$$

(c) The following holds for each $i\in\mathcal{N}$*: If* $\boldsymbol{\theta}_i^\pi\in\mathcal{H}_i^*$*, then* $q_i^*(\boldsymbol{\theta}_i^\pi,f)=q_i^\pi(f)$ *for each* $f\in\mathcal{F}_i$.

We use the observation in part (c) and properties of OPTSTAT to prove the next result which is an important intermediate result used in the proof of optimality result for NOVA given in Theorem 1. The result states that the performance of NOVA (measured in terms of $\phi_S(.)$ defined in (4)) with its parameters $\boldsymbol{\theta}_{i,s}$ picked from the set $\mathcal{H}_i^*$ for each $i\in\mathcal{N}$ is asymptotically optimal. Further, this result suggests that we can prove Theorem 1 if we can show that the updates (11)-(13) of NOVA guide the parameters $(\boldsymbol{\theta}_{i,s})_{s\geq 1}$ of video client i to $\mathcal{H}_i^*$ for each video client $i\in\mathcal{N}$. This motivates the study of convergence behavior of NOVA which is the main focus of the rest of this section.

Theorem 2. *Suppose* $\boldsymbol{\theta}_i^\pi\in\mathcal{H}_i^*$ *for each* $i\in\mathcal{N}$*. Then, for almost all sample paths*

$$\lim_{S\to\infty}\left(\phi_S\left(\left((q_i^*(\boldsymbol{\theta}_i^\pi,f_{i,s}))_{i\in\mathcal{N}}\right)_{1\leq s\leq S}\right)-\phi_S^{opt}\right)=0.$$

B. NOVA parameters also converge to the optimal parameter set, and proving Theorem 1

The next key step is to show that NOVA's 'learning component' (i.e., updates (10)-(13)) is able to guide its parameters to the optimal set (i.e., $\prod_{i\in\mathcal{N}}\mathcal{H}_i^*$). Instead of directly studying the (asynchronous) discrete time evolution of NOVA's parameters, we will first study a related set of 'fluid' NOVA parameters and (in Theorem 3) show that these converge to the optimal set. To this end, we study an auxiliary differential inclusion which evolves according to average dynamics of NOVA. The main goal of this subsection is to study the convergence of the differential inclusion which in turn will help us establish the desired convergence result for NOVA parameters in the next subsection.

For the rest of this section, we also consider the evolution of auxiliary parameters $(v_{i,s_i})_{s_i\geq 1}$, $(\sigma_{i,s_i})_{s_i\geq 1}$ and $(\rho_{i,k})_{k\geq 1}$ associated with NOVA. We update v_{i,s_i} and σ_{i,s_i} based on the quality q^*_{i,s_i+1} (shorthand for $q_i^*(\boldsymbol{\theta}_{i,s_i}, f_{i,s_i+1})$ where $\boldsymbol{\theta}_{i,s_i} = (m_{i,s_i}, b_{Q,i,s_i}, d_{i,s_i})$) chosen by NOVA for (s_i+1)th segment of video client $i\in\mathcal{N}$ as follows:

$$v_{i,s_i+1} = v_{i,s_i} + \epsilon\left(\left(q^*_{i,s_i+1} - m_{i,s_i}\right)^2 - v_{i,s_i}\right), \quad (27)$$

$$\sigma_{i,s_i+1} = \sigma_{i,s_i} + \epsilon\left(f_{i,s_i}\left(q^*_{i,s_i+1}\right) - \sigma_{i,s_i}\right). \quad (28)$$

Thus, the auxiliary parameters v_{i,s_i} and σ_{i,s_i} track the variance (roughly) and the mean segment size respectively of the segments downloaded by video client $i\in\mathcal{N}$. We update the parameter $\boldsymbol{\rho}_k$ based on the resource allocation $\mathbf{r}_k^* \in \mathcal{R}^*(\mathbf{b}_k, c_k)$ in slot k as described below

$$\rho_{i,k+1} = \rho_{i,k} + \epsilon\left(r^*_{i,k} - \rho_{i,k}\right) \ \forall\, i\in\mathcal{N}. \quad (29)$$

Thus, the auxiliary parameter $\boldsymbol{\rho}_k$ tracks the mean resource allocation to video clients. Note that the auxiliary parameters do not affect the allocation or adaptation in NOVA.

Next, let

$$\begin{aligned}\mathcal{H} = \ &\{(\mathbf{m},\mathbf{v},\mathbf{b},\mathbf{d},\boldsymbol{\sigma},\boldsymbol{\rho})\in\mathbb{R}^{6N}: \text{ for each } i\in\mathcal{N}, \quad (30)\\ &0\leq m_i\leq q_{\max}, 0\leq v_i\leq q^2_{\max},\ \underline{b}\leq b_i\leq\overline{b},\ \underline{d}\leq d_i\leq\overline{d},\\ &l_{\min}f_{\min}\leq\sigma_i\leq l_{\max}f_{\max},\ r_{i,\min}\leq\rho_i\leq r_{\max}\}.\end{aligned}$$

Using Lemma 1 and assumptions discussed in Section II, we can show that the parameters $(\mathbf{m}_s,\mathbf{v}_s,\mathbf{b}_k,\mathbf{d}_s,\boldsymbol{\sigma}_s,\boldsymbol{\rho}_k)_{s,k}$ remain in $\mathcal{H}$. For each video client $i\in\mathcal{N}$, we use the variables $\widehat{m}_i(t)$, $\widehat{v}_i(t)$, $\widehat{b}_i(t)$, $\widehat{d}_i(t)$, $\widehat{\sigma}_i(t)$ and $\widehat{\rho}_i(t)$ to track the average dynamics of the parameters m_{i,s_i}, v_{i,s_i}, $b_{i,k}$, d_{i,s_i}, σ_{i,s_i} and $\rho_{i,k}$ respectively associated with NOVA (explained in detail in the sequel). Let $\widehat{\boldsymbol{\Theta}}(t) = \left(\widehat{\mathbf{m}}(t),\widehat{\mathbf{v}}(t),\widehat{\mathbf{b}}(t),\widehat{\mathbf{d}}(t),\widehat{\boldsymbol{\sigma}}(t),\widehat{\boldsymbol{\rho}}(t)\right)\in\mathcal{H}$ and $\widehat{\boldsymbol{\theta}}_i(t) = (\widehat{m}_i(t),\widehat{b}_i(t),\widehat{d}_i(t))$ for each $i\in\mathcal{N}$, i.e., $\widehat{\boldsymbol{\theta}}_i(t)$ includes the components in $\widehat{\boldsymbol{\Theta}}(t)$ that affect the quality adaptation of video client $i\in\mathcal{N}$.

The main focus of this subsection is the following differential inclusion which describes the evolution of $\left(\widehat{\boldsymbol{\Theta}}(t)\right)_{t\geq 0}$:

Auxiliary differential inclusion related to NOVA

$\widehat{\boldsymbol{\Theta}}(0)\in\mathcal{H}$ and for almost all $t\geq 0$ and each $i\in\mathcal{N}$,

$$\dot{\widehat{m}}_i(t) = \frac{1}{u_i\left(\widehat{\boldsymbol{\Theta}}(t)\right)}\left(E\left[q_i^*\left(\widehat{\boldsymbol{\theta}}_i(t),F_i^\pi\right)\right]-\widehat{m}_i(t)\right), \quad (31)$$

$$\dot{\widehat{v}}_i(t) = \frac{1}{u_i\left(\widehat{\boldsymbol{\Theta}}(t)\right)}\left(E\left[\left(q_i^*\left(\widehat{\boldsymbol{\theta}}_i(t),F_i^\pi\right)-\widehat{m}_i(t)\right)^2\right] -\widehat{v}_i(t)\right), \quad (32)$$

$$\dot{\widehat{b}}_i(t) = \frac{1}{(1+\overline{\beta}_i)} - \frac{l_i}{u_i\left(\widehat{\boldsymbol{\Theta}}(t)\right)} + \widehat{z}_i^b\left(\widehat{\boldsymbol{\Theta}}(t)\right), \quad (33)$$

$$\dot{\widehat{d}}_i(t) = \frac{1}{u_i\left(\widehat{\boldsymbol{\Theta}}(t)\right)}\left(\frac{p_i^d E\left[l_iF_i^\pi\left(q_i^*\left(\widehat{\boldsymbol{\theta}}_i(t),F_i^\pi\right)\right)\right]}{\overline{p}_i} -l_i\right) + \widehat{z}_i^d\left(\widehat{\boldsymbol{\Theta}}(t)\right), \quad (34)$$

$$\dot{\widehat{\sigma}}_i(t) = \frac{1}{u_i\left(\widehat{\boldsymbol{\Theta}}(t)\right)}\left(E\left[F_i^\pi\left(q_i^*\left(\widehat{\boldsymbol{\theta}}_i(t),F_i^\pi\right)\right)\right]- \widehat{\sigma}_i(t)\right), \quad (35)$$

$$\dot{\widehat{\rho}}_i(t) = \frac{1}{\tau_{slot}}\left(\frac{\overline{r}_i^*\left(\widehat{\mathbf{b}}(t)\right)}{\tau_{slot}} - \widehat{\rho}_i(t)\right), \quad (36)$$

where

$$u_i\left(\widehat{\boldsymbol{\Theta}}(t)\right) = \tau_{slot}\frac{E\left[l_iF_i^\pi\left(q_i^*\left(\widehat{\boldsymbol{\theta}}_i(t),F_i^\pi\right)\right)\right]}{E\left[r_i^*\left(\widehat{\mathbf{b}}(t),C^\pi\right)\right]}, \quad (37)$$

and $\mathbf{r}^*\left(\widehat{\mathbf{b}}(t),c\right) \in \mathcal{R}^*\left(\widehat{\mathbf{b}}(t),c\right)$ for each $c\in\mathcal{C}$.

Here $\widehat{z}_i^b\left(\widehat{\boldsymbol{\Theta}}(t)\right)$ and $\widehat{z}_i^d\left(\widehat{\boldsymbol{\Theta}}(t)\right)$ are terms mimicking the role of the operators $[.]_{\underline{b}}$ and $[.]_{\underline{d}}$ in (12) and (13), and ensure that $\left(\widehat{\boldsymbol{\Theta}}(t)\right)_{t\geq 0}$ stays in $\mathcal{H}$ (see [18] for a more detailed discussion and see Section 4.3 of [20] for a discussion about projected stochastic approximation). Note that $u_i(.)$ is a set valued map (and hence (31)-(36) describes a differential inclusion) since the denominator $E\left[r_i^*\left(\widehat{\mathbf{b}}(t),C^\pi\right)\right]$ in (37) is a set valued map. Finally, note that the above definition only requires that $\left(\widehat{\boldsymbol{\Theta}}(t)\right)_{t\geq 0}$ is differentiable for *almost* all $t\geq 0$, i.e., we are considering the class of absolutely continuous functions $\left(\widehat{\boldsymbol{\Theta}}(t)\right)_{t\geq 0}$ that satisfy (31)-(36). We can show that the differential inclusion (31)-(36) is well defined, i.e., there exists an absolutely continuous function that solves (31)-(36) for any $\widehat{\boldsymbol{\Theta}}(0)\in\mathcal{H}$. Further, we can show that these solutions are Lipschitz continuous and stay in $\mathcal{H}$ and hence are bounded.

Although we will rigorously establish the relationship between the evolution of parameters of NOVA and (31)-(36) in the next subsection, we can see that the differential inclusion (31)-(36) reflects the average dynamics of the evolution of parameters in NOVA by comparing (31)-(36) against the update rules (11)-(13) and (27)-(29) in NOVA. For instance, this is apparent when we compare the update rule

$$m_{i,s_i+1} - m_{i,s_i} = \epsilon\left(q^*_{i,s_i+1} - m_{i,s_i}\right)$$

for NOVA parameter m_{i,s_i+1} given in (11), against (31) describing the evolution of the parameter $\widehat{m}_i(t)$. Note that the rate of change of $\widehat{m}_i(t)$ given in (31) has a scaling term $\frac{1}{u_i(\widehat{\boldsymbol{\Theta}}(t))}$ which corresponds to the segment download rate of video client i at time t (and $u_i\left(\widehat{\boldsymbol{\Theta}}(t)\right)$ defined in (37) corresponds to expected segment download time of video client i at time t). This scaling by segment download rate is naturally expected for the rate of change of parameters $\widehat{m}_i(t)$, $\widehat{v}_i(t)$, $\widehat{d}_i(t)$ and $\widehat{\sigma}_i(t)$ which correspond to NOVA parameters that are updated when a segment download is completed, and thus we can view $\frac{1}{u_i(\widehat{\boldsymbol{\Theta}}(t))}$ as the update rate associated with these parameters. Similarly, we can view the constant scaling term $\frac{1}{\tau_{slot}}$ in (36) describing the evolution of $\widehat{\rho}_i(t)$ as the corresponding update rate by noting that the associated (auxiliary) NOVA parameter $\rho_{i,k}$ is updated at the beginning of every slot, i.e., once every τ_{slot} seconds. Finally, note that (33) describing the evolution of $\widehat{b}_i(t)$ can be rewritten as

$$\dot{\widehat{b}}_i(t) = \frac{1}{\tau_{slot}}\left(\frac{\tau_{slot}}{(1+\overline{\beta}_i)}\right) - \frac{1}{u_i\left(\widehat{\boldsymbol{\Theta}}(t)\right)}(l_i) + \widehat{z}_i^b\left(\widehat{\boldsymbol{\Theta}}(t)\right),$$

and presence of the two scaling terms $\frac{1}{\tau_{slot}}$ and $\frac{1}{u_i(\widehat{\boldsymbol{\Theta}}(t))}$ reflects the fact that the corresponding NOVA parameter $b_{i,k}$ is updated at the beginning of every slot (using (10)) and when a segment download of video client i is completed (using (12)). Thus, we can expect that (31)-(36) captures the average dynamics of NOVA, and the presence of the video client dependent update rates $\left(\frac{1}{u_i(\widehat{\boldsymbol{\Theta}}(t))}\right)_{i\in\mathcal{N}}$ reflects the *asynchronous* nature of the evolution of NOVA parameters where different video clients are updating their parameters at their own (possibly time varying) rates.

Next we define certain classes of adaptation and allocation policies.

Definition 1. *Stationary resource allocation policy:* Let $(\mathbf{r}(c))_{c\in\mathcal{C}}$ be a $|\mathcal{C}|$ length vector (of vectors) where $\mathbf{r}(c) \in \mathbb{R}_+^N$. We refer to $(\mathbf{r}(c))_{c\in\mathcal{C}}$ as a stationary resource allocation policy as we can associate $(\mathbf{r}(c))_{c\in\mathcal{C}}$ with a resource allocation policy that allocates resource $\mathbf{r}(c)$ in a slot k when $C_k = c$, and thus the policy carries out the resource allocation in a slot based only on the allocation constraint in the slot.

Definition 2. *Feasible stationary resource allocation policy:* We say that a stationary resource allocation policy $\left((\mathbf{r}(c))_{c\in\mathcal{C}}\right)$ is feasible if

$$\mathbf{r}(c) \geq \mathbf{r}_{\min} \text{ and } c(\mathbf{r}(c)) \leq 0, \ \forall\, c \in \mathcal{C}.$$

Definition 3. *Stationary quality adaptation policy for video client i:*
Let $(q_i(f_i))_{f_i\in\mathcal{F}_i} \in \mathbb{R}_+^{\mathcal{F}_i}$. We refer to $(q_i(f_i))_{f_i\in\mathcal{F}_i}$ as a stationary quality adaptation policy for video client $i \in \mathcal{N}$ as we can associate $(q_i(f_i))_{f_i\in\mathcal{F}_i}$ with a quality adaptation policy for video client i that chooses quality $q_i(f_i)$ for each segment s with QR tradeoff f_i, and thus the policy carries out quality adaptation for a segment based only on the QR tradeoff of that segment.

Definition 4. *Feasible stationary quality adaptation policy for video client i:* We say that a stationary quality adaptation policy $(q_i(f_i))_{f_i\in\mathcal{F}_i}$ for video client i is feasible if $0 \leq q_i(f_i) \leq q_{\max}$ for each $f_i \in \mathcal{F}_i$.

Next, we define the set $\widetilde{\mathcal{H}} \subset \mathbb{R}^{6N}$ as

$$\begin{aligned}\widetilde{\mathcal{H}} \;=\; & \Big\{(\mathbf{m},\mathbf{v},\mathbf{b},\mathbf{d},\boldsymbol{\sigma},\boldsymbol{\rho}) \in \mathcal{H} : \ \exists \text{ a feasible stationary} \\ & \text{resource allocation policy } (\mathbf{r}(c))_{c\in\mathcal{C}} \text{ s.t. } \frac{\mathbb{E}\left[r_i(C^\pi)\right]}{\tau_{slot}} = \rho_i \\ & \forall\, i \in \mathcal{N}; \text{ for each } i \in \mathcal{N}, \ \exists \text{ there is a feasible stationary} \\ & \text{quality adaptation scheme } \left((q_i(f_i))_{f_i\in\mathcal{F}_i}\right) \text{ such that} \\ & \mathrm{Var}\left(q_i(F_i^\pi)\right) \leq v_i \leq q_{\max}^2,\ \mathbb{E}\left[F_i^\pi\left(q_i(F_i^\pi)\right)\right] \leq \sigma_i \leq f_{\max}\Big\}.\end{aligned}$$

We can view $\widetilde{\mathcal{H}}$ as the set of 'achievable' parameters in $\mathcal{H}$, i.e., for any element $(\mathbf{m},\mathbf{v},\mathbf{b},\mathbf{d},\boldsymbol{\sigma},\boldsymbol{\rho}) \in \mathcal{H}$ there is some feasible stationary resource allocation policy with mean resource allocation per unit time $\boldsymbol{\rho}$, and there is some feasible stationary quality adaptation policy for each i that has a variance in quality which is at least v_i and mean segment size which is at least σ_i.

It can be verified that $\widetilde{\mathcal{H}}$ is a bounded, closed and convex set (using an approach similar to Lemma 5 (b) in [17]). Hence, we conclude that for any $\boldsymbol{\Theta} \in \mathcal{H}$, there exists a unique projection of $\widetilde{\boldsymbol{\Theta}} \in \mathcal{H}$ onto the set $\widetilde{\mathcal{H}}$. Let $\widetilde{\ }$ denote this projection operator. Hence, for any $\boldsymbol{\Theta} \in \mathcal{H}$, $d_{6N}\left(\boldsymbol{\Theta}, \widetilde{\mathcal{H}}\right) = d_{6N}\left(\boldsymbol{\Theta}, \widetilde{\boldsymbol{\Theta}}\right)$. The next result states that, irrespective of the initialization, the differential inclusion converges to the bounded, closed and convex set $\widetilde{\mathcal{H}}$ of achievable parameters.

Lemma 5. *There exists a finite constant $\chi_0 > 0$ such that for any initialization $\widehat{\boldsymbol{\Theta}}(0) \in \mathcal{H}$,*

$$\frac{d}{dt} d_{6N}\left(\widehat{\boldsymbol{\Theta}}(t), \widetilde{\mathcal{H}}\right) \leq -\chi_0 d_{6N}\left(\widehat{\boldsymbol{\Theta}}(t), \widetilde{\mathcal{H}}\right).$$

Hence,

$$\lim_{t\to\infty} d_{6N}\left(\widehat{\boldsymbol{\Theta}}(t), \widetilde{\mathcal{H}}\right) = 0.$$

In the next result, we provide the main convergence result for the differential inclusion (31)-(36) which states that $\widehat{\boldsymbol{\Theta}}(t)$ converges to the following set

$$\begin{aligned}\mathcal{H}^* \;=\; & \{(\mathbf{m},\mathbf{v},\mathbf{b},\mathbf{d},\boldsymbol{\sigma},\boldsymbol{\rho}) \in \mathcal{H} : \\ & \left(\boldsymbol{\rho}, \left(h_i^B(b_i)\right)_{i\in\mathcal{N}}, \left(h_i^D(d_i)\right)_{i\in\mathcal{N}}\right) \in \mathcal{X}^\pi, \\ & \text{and for each } i \in \mathcal{N},\ m_i = m_i^\pi,\ v_i = v_i^\pi\} \quad (38)\end{aligned}$$

Recall that Theorem 2 suggested that we can prove Theorem 1, if we can show that the updates (11)-(13) guide NOVA parameters $(\boldsymbol{\theta}_{i,s})_{s\geq 1}$ of video client i to the set $\mathcal{H}_i^*$ (defined in (26)) for each video client $i \in \mathcal{N}$. Note that for each $i \in \mathcal{N}$, $\mathcal{H}_i^*$ is a set obtained by projecting $\mathcal{H}^*$ on a lower dimensional space (by considering only video client i's components and 'dropping' the components $(\mathbf{v},\boldsymbol{\sigma},\boldsymbol{\rho})$). Hence, the following result along with Theorem 4 (which relates evolution of NOVA parameters to the differential inclusion) help us to establish the desired convergence property for NOVA parameters. The proof of this result requires several intermediate results (using

Lemma (c), optimality properties related to RNOVA$(\mathbf{b}, c)$, QNOVA$_i(\boldsymbol{\theta}_i, f_i)$, OPTSTAT etc.) and extensions of ideas in [17], [22] etc.

Theorem 3. *(a) For $\widehat{\boldsymbol{\Theta}} = \left(\widehat{\mathbf{m}}, \widehat{\mathbf{v}}, \widehat{\mathbf{b}}, \widehat{\mathbf{d}}, \widehat{\boldsymbol{\sigma}}, \widehat{\boldsymbol{\rho}}\right) \in \mathcal{H}$, and some $(\boldsymbol{\rho}^\pi, \mathbf{b}^\pi, \mathbf{d}^\pi) \in \mathcal{X}^\pi$, let*

$$L\left(\widehat{\boldsymbol{\Theta}}\right) := -\sum_{i\in\mathcal{N}} \left(1+\overline{\beta}_i\right) l_i \left(\widehat{m}_i - \eta_i \widehat{v}_i\right)$$
$$+ \sum_{i\in\mathcal{N}} \left(1+\overline{\beta}_i\right) \left(l_i d_i^\pi \left(\frac{p_i^d \widehat{\sigma}_i}{\overline{p}_i} - 1\right) + \int_{\underline{d}}^{\widehat{d}_i} \left(h_i^D(e) - d_i^\pi\right) de \right)$$
$$+ \sum_{i\in\mathcal{N}} \left(l_i b_i^\pi \widehat{\sigma}_i - \tau_{slot} b_i^\pi \widehat{\rho}_i\right) + \sum_{i\in\mathcal{N}} \sigma_i^\pi \int_{\underline{b}}^{\widehat{b}_i} \left(h_i^B(e) - b_i^\pi\right) de$$
$$+ \sum_{i\in\mathcal{N}} \left(1+\overline{\beta}_i\right) l_i \left(\widehat{m}_i - m_i^\pi\right)^2 + \chi d\left(\widehat{\boldsymbol{\Theta}}, \widetilde{\mathcal{H}}\right),$$

where χ_0 is the positive constant from Lemma 5, and χ_2 is an appropriately chosen (large) positive constant. If $\widehat{\boldsymbol{\Theta}}(0) \in \mathcal{H}$, then for almost all t

$$\frac{dL\left(\widehat{\boldsymbol{\Theta}}(t)\right)}{dt} \begin{cases} \leq 0, \ \forall\ \widehat{\boldsymbol{\Theta}}(t) \in \mathcal{H}, \\ < 0, \ \forall\ \widehat{\boldsymbol{\Theta}}(t) \notin \mathcal{H}^*. \end{cases}$$

(b) If $\widehat{\boldsymbol{\Theta}}(0) \in \mathcal{H}$, then

$$\lim_{t\to\infty} d_{6N}\left(\widehat{\boldsymbol{\Theta}}(t), \mathcal{H}^*\right) = 0.$$

C. NOVA parameters also converge to the optimal parameter set, and proving Theorem 1

The main focus of this subsection is Theorem 4 which relates NOVA to the auxiliary differential inclusion (31)-(36), and obtains the desired convergence result for NOVA by using the convergence result in Theorem 3 for the differential inclusion. Our approach here relies on viewing the update equations ((11)-(13) and (27)-(29)) of NOVA as an asynchronous stochastic approximation update equation (see Chapter 12 of [20] for a detailed discussion on asynchronous stochastic approximation) to relate NOVA to the differential inclusion using tools from the theory of stochastic approximation. After obtaining the convergence result for NOVA in Theorem 4, we conclude this section with the proof of Theorem 1.

Next, we define two auxiliary variables $b_{R,i,k}$ and b_{Q,i,s_i+1}. At the beginning of slot k, let $b_{R,i,k} = b_{i,k}$ for each $i \in \mathcal{N}$ and thus the variable stores the value of $b_{i,k}$ used while deciding allocation for k-th slot. In slot k, if any video client $i \in \mathcal{N}$ finishes download of s_i th segment, let $b_{Q,i,s_i+1} = b_{i,k+1}$, and thus the variable stores the value of $b_{i,k}$ used while deciding the quality for video client i's (s_i+1)-th segment. In this following, we use the superscript ϵ on NOVA parameters $(m_{i,s}^\epsilon)_{i\in\mathcal{N}}$, $(v_{i,s}^\epsilon)_{i\in\mathcal{N}}$, $(b_{Q,i,s}^\epsilon)_{i\in\mathcal{N}}$, $(b_{R,i,k}^\epsilon)_{i\in\mathcal{N}}$, $(b_{i,k}^\epsilon)_{i\in\mathcal{N}}$, $(d_{i,s}^\epsilon)_{i\in\mathcal{N}}$, $(\sigma_{i,s}^\epsilon)_{i\in\mathcal{N}}$ and $(\rho_{i,k}^\epsilon)_{i\in\mathcal{N}}$ to emphasize their dependence on ϵ (see NOVA updates in (10)-(13) to see the dependence). We refer to the update of NOVA parameters $(m_{i,s_i}, b_{i,k}, d_{i,s_i})$ in (11)-(13) carried out after the selection of segment quality for video client i (following a segment download) as a Q$_i$-update. Let $\delta\tau_{Q,i,s}^\epsilon$ denote the time (in seconds) between the sth and $(s+1)$th Q$_i$-updates. Let $\tau_{Q,i,s}^\epsilon = \epsilon \sum_{j=0}^{s-1} \delta\tau_{Q,i,j}^\epsilon$. denote ϵ times the cumulative time for the first s Q$_i$-updates.

Next, we define time interpolated processes $\left(\widehat{\mathbf{m}}^\epsilon(t), \widehat{\mathbf{v}}^\epsilon(t), \widehat{\mathbf{b}}^\epsilon(t), \widehat{\mathbf{d}}^\epsilon(t), \widehat{\boldsymbol{\sigma}}^\epsilon(t), \widehat{\boldsymbol{\rho}}^\epsilon(t)\right)$ associated with NOVA's parameters. For each $i \in \mathcal{N}$ and for $t \in \left[\tau_{Q,i,s}^\epsilon, \tau_{Q,i,s+1}^\epsilon\right)$, let $\widehat{m}_i^\epsilon(t) = m_{i,s}^\epsilon$, $\widehat{v}_i^\epsilon(t) = v_{i,s}^\epsilon$, $\widehat{b}_{Q,i}^\epsilon(t) = b_{Q,i,s}^\epsilon$, $\widehat{d}_i^\epsilon(t) = d_{i,s}^\epsilon$ and $\widehat{\sigma}_i^\epsilon(t) = \sigma_{i,s}^\epsilon$. Also, for $t \in [k\tau_{slot}\epsilon, (k+1)\tau_{slot}\epsilon)$, let $\widehat{b}_{R,i}^\epsilon(t) = b_{R,i,k}^\epsilon$ and $\widehat{\rho}_i^\epsilon(t) = \rho_{i,k}^\epsilon$. For each t, let

$$\begin{aligned} \widehat{\boldsymbol{\Theta}}_Q^\epsilon(t) &= \left(\widehat{\mathbf{m}}^\epsilon(t), \widehat{\mathbf{v}}^\epsilon(t), \widehat{\mathbf{b}}_Q^\epsilon(t), \widehat{\mathbf{d}}^\epsilon(t), \widehat{\boldsymbol{\sigma}}^\epsilon(t), \widehat{\boldsymbol{\rho}}^\epsilon(t)\right), \\ \widehat{\boldsymbol{\Theta}}_R^\epsilon(t) &= \left(\widehat{\mathbf{m}}^\epsilon(t), \widehat{\mathbf{v}}^\epsilon(t), \widehat{\mathbf{b}}_R^\epsilon(t), \widehat{\mathbf{d}}^\epsilon(t), \widehat{\boldsymbol{\sigma}}^\epsilon(t), \widehat{\boldsymbol{\rho}}^\epsilon(t)\right), \end{aligned}$$

Note that $\widehat{\boldsymbol{\Theta}}_Q^\epsilon(.)$ and $\widehat{\boldsymbol{\Theta}}_R^\epsilon(.)$ are different only for components $2N+1$ to $3N$. The next result states that for small enough ϵ, the time interpolated versions of NOVA parameters $\widehat{\boldsymbol{\Theta}}_Q^\epsilon(.)$ and $\widehat{\boldsymbol{\Theta}}_R^\epsilon(.)$ stay close to the set $\mathcal{H}^*$ (defined in (38)) most of the time over long time windows. This result is an *extension* of Theorem 3.4 in Chapter 12 of [20]. The proof relies on relating $\widehat{\boldsymbol{\Theta}}_Q^\epsilon(.)$ and $\widehat{\boldsymbol{\Theta}}_R^\epsilon(.)$ associated with NOVA to the auxiliary differential inclusion (31)-(36) (by viewing the update equations (11)-(13) of NOVA as an asynchronous stochastic approximation update equation), and using Theorem 3 which states that the differential inclusion converges to the set $\mathcal{H}^*$.

Theorem 4. *Let $\widehat{\boldsymbol{\Theta}}_Q^\epsilon(0) = \widehat{\boldsymbol{\Theta}}^\epsilon(0) \in \mathcal{H}$. Then, the fraction of time in the time interval $[0, T]$ that $\widehat{\boldsymbol{\Theta}}_Q^\epsilon(.)$ and $\widehat{\boldsymbol{\Theta}}_R^\epsilon(.)$ spend in a small neighborhood of $\mathcal{H}^*$ converges to one in probability as $\epsilon \to 0$ and $T \to \infty$.*

We have the following corollary of Theorem 4 which says that for small enough ϵ and after running NOVA for long enough, video client i's NOVA parameter stays close to $\mathcal{H}_i^*$(defined in (26)) most of the time with high probability.

Corollary 1. *Let $\widehat{\boldsymbol{\Theta}}^\epsilon(0) \in \mathcal{H}$ and $S_\epsilon = \frac{S}{\epsilon}$. Then for each $i \in \mathcal{N}$, the following holds: for any $\delta > 0$, the fraction of segment indices for which $(\boldsymbol{\theta}_{i,s})_{1\leq s\leq S_\epsilon}$ is in a δ-neighborhood of $\mathcal{H}_i^*$ converges to one in probability as $\epsilon \to 0$ and $S \to \infty$.*

We have now obtained all the intermediate results required to prove Theorem 1 which is given below.

Proof of Theorem 1: A detailed proof of part (a) of Theorem 1 can be found in [18] and it primarily relies on the fact that $b_{i,k}$ and $d_{i,s}$ are bounded (from Lemma 1).

Next, we prove part (b). Using Corollary 1 (which says that $(\boldsymbol{\theta}_{i,s})_{1\leq s\leq S_\epsilon}$ essentially converges to $\mathcal{H}_i^*$) and Lemma 2 (a) (which says that $q_i^*(\boldsymbol{\theta}_i, f_i)$ is a continuous function of $\boldsymbol{\theta}_i$), we can conclude that for $\boldsymbol{\theta}_i^\pi \in \mathcal{H}_i^*$

$$\lim_{S\to\infty} \lim_{\epsilon\to 0} \Big(\phi_{S_\epsilon}\left(\left(\left(q_i^*\left(\boldsymbol{\theta}_{i,s}, f_{i,s}\right)\right)_{i\in\mathcal{N}} \right)_{1\leq s\leq S_\epsilon} \right) - \phi_{S_\epsilon}\left(\left(\left(q_i^*\left(\boldsymbol{\theta}_i^\pi, f_{i,s}\right)\right)_{i\in\mathcal{N}} \right)_{1\leq s\leq S_\epsilon} \right) \Big)$$

goes to zero in probability. Now, part (b) of Theorem 1 follows from the above observation and Theorem 2 which states that for each $i \in \mathcal{N}$ and for almost all sample paths

$$\lim_{S\to\infty} \left(\phi_S\left(\left(\left(q_i^*\left(\boldsymbol{\theta}_i^\pi, f_{i,s}\right)\right)_{i\in\mathcal{N}} \right)_{1\leq s\leq S} \right) - \phi_S^{opt} \right) = 0. \qquad \blacksquare$$

VI. Extensions

In [18], NOVA has been extended in several important directions and they are discussed briefly next. [18] considers a more general framework allowing more flexibility in imposing QoE fairness (as pointed out in Section II), general QoE models (i.e., generalizations of (3) that allow more flexibility in variability penalty), and more general allocation constraints (described in terms of finite number of convex functions allowing the modeling of the network resources available in the form of sub-resources like sub-bands).

Analysis of NOVA's optimality for certain important special settings is included in [18], and these settings include networks with legacy resource allocation policies, with just a single standalone video client and with other traffic (e.g., data traffic).

[18] also analyzes the performance of NOVA in networks with discrete network resources (i.e., when the set of feasible resource allocations in a slot is discrete), with user dynamics, and with several practical video client implementation considerations such as finiteness of the number of representations, impact of choice of ϵ, $(h_i^B(.))_{i\in\mathcal{N}}$ and $(h_i^D(.))_{i\in\mathcal{N}}$, reduction of startup delay and frequency of rebuffering, playback buffer limits, playback pauses, ads etc.

VII. Simulations

In this section, we evaluate NOVA using Matlab simulations to compare the performance of a wireless network operating under NOVA vs one using Proportionally Fair (PF) network resource allocation (see [19]) and quality adaptation based on Rate Matching (RM). We discuss PF and RM below. We restrict the discussion to the key features of the setting used for simulations, and finer details can be found in [18] and [23].

We consider a wireless network with $\tau_{slot} = 10$ msecs, and with allocation constraints of the form $c_k(\mathbf{r}_k) = \sum_{i\in\mathcal{N}} \frac{r_{i,k}}{p_{i,k}} - 1$ in each slot k, where $p_{i,k}$ denotes the peak resource allocation for video client i in slot k, i.e., if we only allocate resources to video client i in slot k, then $r_{i,k} = p_{i,k}$ is the maximum resource allocation to the video client. We used traces for peak resource allocation based on data for an HSDPA system[3] and we used randomly scaled versions of these traces to model heterogeneous channels for video clients.

Under PF (see [19]), an optimal solution to

$$\max_{\mathbf{r}} \left\{ \sum_{i\in\mathcal{N}} \frac{r_i}{\rho_{i,k}} : \ c_k(\mathbf{r}) \le 0, \ r_i \ge r_{i,\min} \ \forall\, i \in \mathcal{N} \right\}, \quad (39)$$

is the network resource allocation in slot k. Here the parameters $(\rho_{i,k})_{i\in\mathcal{N}}$ track the mean allocation to the video clients.

In our simulations, we consider video clients downloading *different parts* of three open source movies Oceania, Route 66 and Valkaama where the segments are of duration 1 second each and have 5-6 different representations. We obtained[4] proxy subjective VQA metric for the representations based on the corresponding value of MSSSIM-Y metric ([24]). To account for finiteness of available representations, we modify the optimization problem $\text{QNOVA}_i(\boldsymbol{\theta}_i, f_i)$, used for quality adaptation in NOVA by imposing an additional restriction that the quality for segment s of video client i is picked from the finite set of quality choices available for the segment.

In quality adaptation based on RM (Rate Matching), a video client tries to 'match' the effective compression rate of the selected representation to (current estimate of) mean resource allocation in bits per second, and further modifies the selection to respond to the state of the playback buffer by switching to aggressive and cautious modes (see [18] for details). This is basic feature in many compression rate adaptation algorithms, for instance, see [25] where (following their terminology) we see that 'requested bitrate' (i.e., size of the representation) stays close to the 'average throughput' (i.e., $\rho_{i,k}$ in our setting) in Microsoft Smooth Streaming player and Netflix player.

For our simulations of NOVA, we let $\epsilon = 0.05$, $r_{i,min} = 0.001$ bits, $\eta_i = 0.05$, $\overline{\beta}_i = 0$ and $p_i^d = 0.01$ dollars per bit for each $i \in \mathcal{N}$. While evaluating the rebuffering time in the simulation results, we allow for a startup delay of 3 secs. For each $i \in \mathcal{N}$, we chose $h_i^D(d_i) = 10d_i$ and $h_i^B(b_i) = 0.005\left(\frac{b_i}{0.05} + \max\left(\frac{b_i-20}{0.05}, 0\right)^2\right)$, $m_{i,0} = 25$, $b_{i,0} = \frac{40}{0.05}$ and $d_{i,0} = 1$ (these choices are discussed in more detail in [18]).

Each point in the plots discussed below is obtained by running the associated algorithm 50 times where each simulation is run until all the video clients have downloaded a video of duration at least 10 minutes. Each point corresponds to a fixed number of video clients N taking values in $\{12,\ 15,\ 18,\ 21,\ 24,\ 27,\ 30,\ 33\}$. We refer to the combination of PF resource allocation and RM quality adaptation as PF-RM. We also study the performance of PF-QNOVA which uses PF resource allocation and quality adaptation in NOVA. NOVA, PF-QNOVA and PF-RM correspond to setting with no price constraints, and their modifications with price constraint of 3 dollars per bit are referred to as NOVA(3), PF-QNOVA(3) and PF-RM(3) respectively. NOVA(3) and PF-QNOVA(3) implementations use a more stringent/conservative price constraint of 0.95×3.

In Fig. 2 (a), we compare the QoE of the video clients under different algorithms, where we measure QoE using the metric QoE_1 which is the average across simulation runs of

$$\frac{1}{N}\sum_{i\in\mathcal{N}} \left(m_i^{600}(q_i) - \sqrt{\text{Var}_i^{600}(q_i)} \right),$$

where $m_i^{600}(q_i) - \sqrt{\text{Var}_i^{600}(q_i)}$ is the metric proposed in [4] with the scaling constant for $\sqrt{\text{Var}_i^{600}(q_i)}$ set to unity (and $m_i^{600}(q_i)$ and $\text{Var}_i^{600}(q_i)$ are defined in Section II). On comparing QoE_1 using Fig. 2 (a), we see that NOVA performs much better than PF-RM and PF-QNOVA, and in fact provides 'network capacity gains' of about 60% over PF-RM, i.e., given a requirement on average QoE_1, we can support about 60% more video clients by using NOVA than that under PF-RM. For instance, if we consider the horizontal dashed line in Fig. 2 (a) that corresponds to an average QoE_1 requirement of about 43, we see that PF-RM can only support 20 video clients while meeting this requirement whereas NOVA can support almost 33 video clients. Under price constraint (of

[3]This data was provided by a service provider. See [18] for more details on the generation of these sequences.

[4]See [18] for details including plots depicting diversity of the QR tradeoffs.

3 dollars per second) also, we see that NOVA(3) provides network capacity gains of about 60% over PF-RM(3). The gain from the adaptation component of NOVA is also visible in Fig. 2 (a), where we see that PF-QNOVA provides network capacity gains of about 25% over PF-RM respectively.

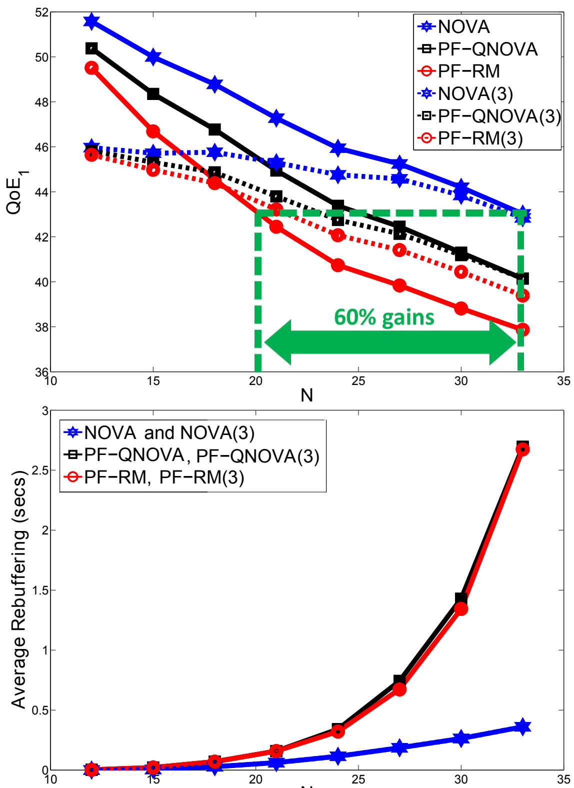

Fig. 2. (a) Top figure: QoE_1 gains from NOVA; (b) Bottom figure: Reduction in rebuffering under NOVA

The results in Fig. 2 (b) depict the significant reduction in the amount of time spent rebuffering under NOVA and NOVA(3). Using Fig. 2, we see that NOVA outperforms PF-RM in both the metric QoE_1 and the amount of time spent rebuffering which cover some of the most important factors affecting video clients' QoE (see the discussion in Section I).

Our simulations results also showed capacity gains of about 50% with respect to another metric QoE_2 obtained by replacing $\mathrm{Var}_i^{600}(q_i)$ in QoE_1 with $\mathrm{MSD}_i^{600}(q_i) := \frac{1}{600}\sum_{s=1}^{600}(q_{i,s+1} - q_{i,s})^2$ which penalizes short term variability. Further, the results also showed that NOVA even has a slightly higher mean quality (*in addition to lower variability in quality*) in all but lightly loaded networks.

More details (e.g., fairness gains under NOVA) of the results for the above setting is given in [18] . We carried out extensive simulations validating the performance of NOVA in other setting too, and these results can also be found in [18].

VIII. Conclusions and Future Directions

We developed a simple online algorithm NOVA for optimizing video delivery, well suited for today's networks supporting DASH-based video clients. Interesting future directions include exploration of the potential of *learning* user preferences, and developing 'NOVA-like' algorithms for networks with contention based medium access by modulating the back-off timers using information about parameters like $b_{i,k}$.

References

[1] CISCO, "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2010-2015," Feb. 2011.

[2] MPEG-DASH. mpeg.chiariglione.org/standards/mpeg-dash.

[3] K. Seshadrinathan and A. C. Bovik, "Automatic prediction of perceptual quality of multimedia signals–a survey," *International Journal of Multimedia Tools and Applications*, Jan. 2011.

[4] C. Yim and A. C. Bovik, "Evaluation of temporal variation of video quality in packet loss networks," *Signal Processing: Image Communication*, Jan. 2011.

[5] T. Kim and M. Ammar, "Optimal quality adaptation for scalable encoded video," *IEEE Journal on Selected Areas in Communications*, Dec. 2005.

[6] P. Seeling, M. Reisslein, and B. Kulapala, "Network performance evaluation using frame size and quality traces of single-layer and two-layer video: A tutorial," *IEEE Communications Surveys Tutorials*, 2004.

[7] R. Mok, E. Chan, and R. Chang, "Measuring the quality of experience of HTTP video streaming," *IFIP/IEEE International Symposium on Integrated Network Management*, 2011.

[8] F. Fu and M. van der Schaar, "A systematic framework for dynamically optimizing multi-user wireless video transmission," *IEEE Journal on Selected Areas in Communications*, Apr. 2010.

[9] J. Huang, Z. Li, M. Chiang, and A. Katsaggelos, "Joint source adaptation and resource allocation for multi-user wireless video streaming," *IEEE Transactions on Circuits and Systems for Video Technology*, May 2008.

[10] H. Hu, X. Zhu, Y. Wang, R. Pan, J. Zhu, and F. Bonomi, "Qoe-based multi-stream scalable video adaptation over wireless networks with proxy," in *ICC*, 2012.

[11] D. Bethanabhotla, G. Caire, and M. J. Neely, "Joint transmission scheduling and congestion control for adaptive video streaming in small-cell networks," *arXiv:1304.8083 [cs.NI]*, 2013.

[12] V. Joseph and G. de Veciana, "Jointly Optimizing Multi-user Rate Adaptation for Video Transport over Wireless Systems: Mean-Fairness-Variability Tradeoffs," *Technical Report*, Jul. 2011. www.ece.utexas.edu/~gustavo/VariabilityAwareVideoRateAdapt.pdf.

[13] V. Joseph, S. Borst, and M. Reiman, "Optimal rate allocation for adaptive wireless video streaming in networks with user dynamics," *IEEE INFOCOM*, 2014.

[14] B. Blaszczyszyn, M. Jovanovic, and M. Karray, "Quality of real-time streaming in wireless cellular networks- stochastic modeling and analysis," *arXiv:1304.5034 [cs.NI]*, 2013.

[15] S. Khan, Y. Peng, E. Steinbach, M. Sgroi, and W. Kellerer, "Application-driven cross-layer optimization for video streaming over wireless networks," *IEEE Communications Magazine*, Jan. 2006.

[16] S. Shakkottai and R. Srikant, "Network optimization and control," *Foundations and Trends in Networking*, Jan. 2007.

[17] V. Joseph, G. de Veciana, and A. Arapostathis, "Resource Allocation: Realizing Mean-Variability-Fairness Tradeoffs," *Submitted to IEEE Transactions on Automatic Control*, 2013.

[18] V. Joseph and G. de Veciana, "NOVA: QoE-driven Optimization of DASH-based Video Delivery in Networks," *Full version available on arXiv.org, http://wncg.org/publications/dl.php?file=VJGdV13.pdf*, 2013.

[19] H. Kushner and P. Whiting, "Convergence of proportional-fair sharing algorithms under general conditions," *IEEE Transactions on Wireless Communications*, Jul. 2004.

[20] H. Kushner and G. G. Yin, *Stochastic Approximation and Recursive Algorithms and Applications*. Springer, 2003.

[21] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2009.

[22] A. L. Stolyar, "Greedy primal-dual algorithm for dynamic resource allocation in complex networks," *Queueing Systems*, Nov. 2006.

[23] V. Joseph, "Mean-Variability-Fairness Tradeoffs in Resource Allocation with Applications to Video Delivery," *Ph.D. dissertation*, Aug. 2013. ECE Department, UT Austin.

[24] Z. Wang, E. Simoncelli, and A. Bovik, "Multiscale structural similarity for image quality assessment," *Asilomar Conference on Signals, Systems and Computers*, Nov. 2003.

[25] S. Akhshabi, S. Narayanaswamy, A. C. Begen, and C. Dovrolis, "An experimental evaluation of rate-adaptive video players over HTTP," *Signal Processing: Image Communication*, 2012.