Channel Sharing by Multi-class Rate Adaptive Streams: Performance Region and Optimization ⋆

Nikos Argiriou, Leonidas Georgiadis

Electrical and Computer Engineering Dept., Aristotle University of Thessaloniki, Thessaloniki, Greece.

Abstract

We consider the problem of channel sharing by rate adaptive streams belonging to various classes. The performance metric per class is the mean scaled bandwidth allocated to connections in the class possibly scaled by appropriate factors. We first provide a bandwidth adaptation policy that maximizes a linear combination of class performance metrics; then, we use this result to characterize the region where the class performance metrics lie under any bandwidth adaptation policy. Based on the results above we use stochastic approximation techniques to provide a policy that optimizes a combination of concave rewards associated with class performance metrics. Finally, we propose a modification of the optimal policy to account for the case where connection holding times are unknown, and we study its performance through simulations.

Key words: Bandwidth Sharing, Rate-adaptive Streams, Wireless Channel Sharing, Performance Region, Stochastic Approximation.

1 Introduction

We consider a communication channel whose bandwidth is shared by randomly arriving connections belonging to a number of classes. Connection bandwidth

⋆ Part of the results of this paper were presented at the HET-NETs 2005 Conference, UK.

Email addresses: narg@egnatia.ee.auth.gr (Nikos Argiriou), leonid@eng.auth.gr (Leonidas Georgiadis).

may be adapted within a given range, hence giving the opportunity for bandwidth management in order to achieve good reception quality with low blocking probability. Connection bandwidth adaptation can be achieved by various coding techniques such as layered coding [1], [2] and adaptation of compression parameters [3], [4], [5]. Depending on the technique, rate adaptation can take one of a number of discrete values, or it can take any value within a specific range. In particular, wavelet coding [5] is particularly well suited for continuous rate adaptation. Rate adaptation implies some variability in the perceived quality of the application. There is a relatively large class of applications that can tolerate this variability, such as entertainment video, low-cost information distribution such as video or speech news and videoconferencing.

Using rate adaptation applications can adapt their transmission rate to changing network conditions in order avoid congestion [6], [7], [8]. On the other hand, there are proposals where Variable Bit Rate (VBR) connections request Constant Bit Rate (CBR) service from the network depending on their current needs [9].

The ability of applications to adapt their transmission rate is particularly useful in shared-channel environments such as Hybrid Fiber Coax (HFC) networks and broadband wireless cellular networks. These channels are shared by a number of users. The advantage of channel sharing is that when the number of active users is small they can share all the available bandwidth and hence receive very good QoS. The disadvantage is that as the number of users that share the same bandwidth increases, if the system is left uncontrolled, the perceived quality of multimedia applications reduces significantly. However proper admission control, combined with application rate adaptation has the potential of guaranteeing acceptable quality of reception while achieving large system utilization. With this approach, when the number of active users is small, applications are admitted by the system with their maximum requested rate, while as the system load increases the application transmission rate is reduced, while still remaining within acceptable levels, so that more connections can be admitted. This process is facilitated by the existence of controllers (headend in HFC networks [10] and base stations in wireless cellular networks) that can convey feedback to the already running applications through the downstream channel (see Figure 1), in order to reduce their rate accordingly.

In recent years, several works addressed the problem of bandwidth adaptation management under various assumptions on channel characteristics and the bandwidth adaptation policies [11], [12], [13], [14], [15], [16], [17], [18]. Our approach follows [15], [17], where for single class systems, bandwidth management policies optimizing the average connection scaled bandwidth were presented. For the single class system we show in Section 3 that the policy presented in [15], [17], is optimal under quite general conditions. We then

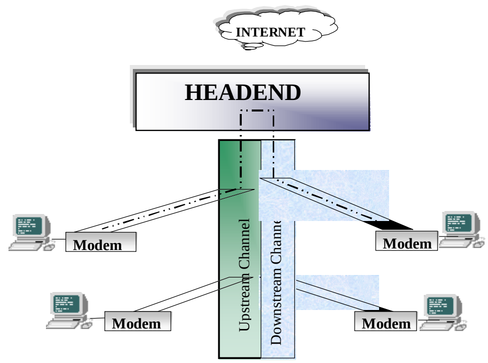

Fig. 1. CATV System

consider multiclass systems (Section 4). We provide a policy for maximizing a linear combination of class performance measures in Section 4.1; using this, we characterize the performance region of the system in Section 4.2. These results permit us to apply the general framework developed in [19], [20], to provide a bandwidth adaptation policy that optimizes a combination of concave rewards associated with class performance measures (Section 5). Based on the insight obtained by the optimal policy in Section 5 we propose in Section 6 a modified bandwidth adaptation policy that does not require the knowledge of connection holding times and study its performance relative to the optimal through simulations. Finally, in Section 7 we discuss generalizations of our results and directions for further research.

2 System Model and Notation

Throughout the paper, sets are denoted by calligraphic capital letters. The number of elements in a set $\mathcal{X}$ is denoted by the corresponding normal capital letter, i.e., $X = |\mathcal{X}|$. Vectors are denoted by boldface letters, e.g., $\mathbf{x} = \{x_i\}_{i=1}^N$, $\mathbf{x} \in \mathbb{R}^N$ ($\mathbb{R}$ is the set of real numbers). The Euclidean norm of a vector $\mathbf{x}$ is denoted by $\|\mathbf{x}\|$.

We consider a communication channel of bandwidth B bps. Extensions to the case where the bandwidth of the link may vary over time are discussed in Section 7. Connections arrive for transmission over the link. Let $\mathcal{A}(t)$ be the set of connections that arrived and have been accepted by the system up to time t. Let also $\mathcal{N}(t)$ be the set of connections that are present in the system at time t.

Connection i arrives at time a_i and departs at time $d_i > a_i$. The connection holding time $h_i = d_i - a_i > 0$, is assumed known, e.g., connections may correspond to transmission of prerecorded movies. Through appropriate compression techniques, transmission of connection i may take place at rates belonging to a set $\mathcal{B}_i$. We call these rates "bandwidth levels". We denote by $\underline{B}_i$ and $\overline{B}_i$ the minimum and maximum bandwidth levels in $\mathcal{B}_i$. Transmission rates of connection i may be adapted over time, but must belong to $\mathcal{B}_i$ for acceptable reception quality. Hence, if $b_i(t)$ is the bandwidth allocated to connection i at time t, $a_i \leq t < d_i$, it must hold for any t,

$$\sum_{i \in \mathcal{N}(t)} b_i(t) \leq B, \tag{1a}$$

$$b_i(t) \in \mathcal{B}_i, \ \ i \in \mathcal{N}(t). \tag{1b}$$

In this paper we mainly consider multiclass systems where arriving connections belong to one of the classes in a set $\mathcal{C}$. In this case we define the sets $\mathcal{A}_c(t)$ and $\mathcal{N}_c(t)$ in a manner analogous to the definition of $\mathcal{A}(t)$, $\mathcal{N}(t)$ and respectively. Hence we have, $\mathcal{A}(t) = \cup_{c\in\mathcal{C}}\mathcal{A}_c(t)$ and $\mathcal{N}(t) = \cup_{c\in\mathcal{C}}\mathcal{N}_c(t)$. We also denote by $\mathcal{C}(t)$ the set of classes for which $\mathcal{N}_c(t) \neq \varnothing$, i.e., for any class in $\mathcal{C}(t)$, say class c, there is a least one class c connection present in the system at time t.

In order to operate the system, two policies must be defined: the "Connection Admission Policy" and the "Bandwidth Adaptation Policy". The Connection Admission policy decides whether to accept or reject a newly arriving connection, while the Bandwidth Adaptation policy adjusts at any time t the bandwidth of the connections that are currently in the system.

The Connection Admission policy affects the blocking probability of the classes. Hence and important issue is how to design such a policy in order to satisfy fairness or preferential treatment criteria related to class blocking probability. There is an extensive literature on the design of such a policy when the connection bandwidth requirements are fixed, e.g., [21], [22], [23] and a large number of other references in [24] where a nice collection of several Connection Admission Policies and their analysis can be found in [24]. In our case connection bandwidths can be adapted. In order to keep the design of Connection Admission and Bandwidth Adaptation policies separate, and to keep blocking probabilities as low as possible (for given class blocking probability objectives), we consider connection admission policies which operate based on the minimum acceptable bandwidth levels ($\underline{B}_i$). More specifically, we adopt the following general Connection Admission Policy.

Acceptable Connection Admission Policy. Any policy π designed for connections with fixed bandwidth requirements may be employed. Whenever a new connection arrives to the system, the policy π admits or rejects the connection using as connection bandwidths the minimum acceptable connection bandwidth levels $\underline{B}_i$.

An example is the "Complete Sharing" policy which operates based on (1) as follows. Let connection j arrive to the system at time t. The connection is admitted by the system if and only if the channel capacity is not exceeded by giving to all connections their minimal bandwidth requirements, i.e.,

$$\sum_{i \in \mathcal{N}(t)} \underline{B}_i \leq B. \tag{2}$$

This type of Connection Admission policy have been proposed in [12], [15], [16], [17].

Another example is the "Threshold" policy. With this policy a parameter $H_c \leq B$ is associated with each class c. Whenever a class c connection j arrives in the system, the connection is admitted if and only if a) the channel capacity is not exceeded by giving to all connections their minimal bandwidth requirements, i.e., (2) holds, and b) the sum of acceptable minimal class c connection bandwidths does not exceed, i.e.,

$$\sum_{i \in \mathcal{N}_c(t)} \underline{B}_i \leq H_c.$$

When $\sum_{c \in \mathcal{C}} H_c \leq B$, the policy is called "Complete Partitioning" since in effect channel bandwidth is partitioned in a fixed manner to the classes.

A third policy example is proposed in [13], [14] in order to provide preferential treatment to connections that arrive to a cell from other cells (hand-off operation) compared to connections that are generated within the cell. A locally generated class c connection is admitted (possibly by adapting the bandwidths of already admitted connections) only if the number of class c connections currently in the system is below a given threshold t_c. Connections arriving to the cell through hand-off are always admitted (if possible by bandwidth adaptation).

Partitioning (not necessarily complete) policies, general Coordinate Convex policies and Trunk Reservation policies (see [24] for definitions) operating by considering as connection bandwidth the minimal acceptable connection bandwidth level are also acceptable policies within our study.

The key property of an acceptable Connection Admission policy q that will be useful in the development that follows, is the following

> **Key Property of an Acceptable Connection Admission policy** q. Under q, $\mathcal{N}_c(t)$, $\mathcal{A}_c(t)$ and $\mathcal{C}(t)$ are independent of the employed Bandwidth Adaptation Policy.

While the achievement of proper class blocking probabilities is the objective of the Connection Admission Policy, the objective of the Bandwidth Adaptation policy is to rearrange connection bandwidths so that certain performance optimization criteria related to channel bandwidth sharing are satisfied. In this paper we concentrate on the design of Bandwidth Adaptation policy.

There are various performance metrics related to bandwidth allocation that may be defined and the issue which of them or combination thereof is appropriate, is still an open research problem [25], [26], [27], [28]. Among the most relevant ones is the average bandwidth allocated to a connection throughout its holding time, or its scaled version called "scaled mean connection bandwidth" [17]. The (scaled) mean bandwidth allocated to connection i when Bandwidth Adaptation policyπ is employed, is defined as

$$\widehat{b}_i^\pi = \frac{\int_{a_i}^{d_i} b_i^\pi(t)dt}{h_i\overline{B}_i}, \tag{3}$$

that is, $\widehat{b}_i^\pi$ represents the quality of the average bandwidth received by the connection relative to the best possible, $\overline{B}_i$. For definiteness, the development that follows uses the scaled mean connection bandwidth as a metric of connection performance. However, we note that the scaling by $\overline{B}_i$ is not essential for the development. Scaling factors other than $\overline{B}_i$ may also be used is desired, including no scaling, i.e., on can set $\overline{B}_i = 1$.

For our purposes, it will be convenient to extend the definition of $\widehat{b}_i^\pi$for all times $t \geq 0$, as follows.

$$\widehat{b}_i^\pi(t) = \begin{cases} 0 & t < a_i \\ \frac{1}{h_i\overline{B}_i}\int_{a_i}^{t} b_i^\pi(s)ds & a_i \leq t < d_i \\ \widehat{b}_i^\pi & t \geq d_i \end{cases} \tag{4}$$

That is, the scaled mean bandwidth allocated to connection i is zero before the connection arrives, the actual value $\widehat{b}_i^\pi$ after the connection departs, and its "instantaneous value" $\int_{a_i}^{t} b_i^\pi(s)ds/\left(h_i\overline{B}_i\right)$ while the connection resides in the system. Let $V_i = h_i\overline{B}_i$ and set $b_i^\pi(t) = 0$ if either $t < a_i$, or $t \geq d_i$. We can then rewrite (4) in the following form which will be useful in the sequel.

$$\widehat{b}_i^\pi(t) = \frac{\int_0^t b_i^\pi(s)ds}{V_i}. \tag{5}$$

Since $\underline{B}_i \leq b_i^\pi(t) \leq \overline{B}_i$, it holds for $i \in \mathcal{A}(t)$, $\widehat{b}_i^\pi(t) \leq 1$. The performance of the system at time t is defined as the average of the performance metrics of all connections that have been admitted by the system up to time t, that is,

$$\widehat{B}^\pi(t) = \frac{\sum_{i\in\mathcal{A}(t)} \widehat{b}_i^\pi(t)}{A(t)}. \tag{6}$$

For a single-class system, we are interested in defining Bandwidth Adaptation Policies that maximize (6) either at any time t, or asymptotically as $t \to \infty$ (the precise meaning will be given in Section 4.1). For a multiclass system, the class performance metric $\widehat{B}_c^\pi(t)$ is defined analogously to (6) as

$$\widehat{B}_c^\pi(t) = \frac{\sum_{i \in \mathcal{A}_c(t)} \widehat{b}_i^\pi(t)}{A_c(t)}.$$

In this case, we will be interested in optimizing functions of class performance metrics $\widehat{B}_c^\pi(t)$, $c \in \mathcal{C}$.

In the rest of the paper we will use the following conventions regarding summations

$$\sum_{i \in \mathcal{X}} x_i = 0, \text{ if } \mathcal{X} = \varnothing, \ \sum_{i=a}^{b} x_i = 0, \text{ if } a > b.$$

3 Single Class System

In [15], [17], a policy π with maximum $\widehat{B}^\pi \triangleq \lim_{t \to \infty} \widehat{B}^\pi(t)$ has been proposed for the single class system assuming Poisson arrivals and i.i.d. connection holding times. We generalize this result in this section by showing that the same policy optimizes $\widehat{B}^\pi(t)$ at any time t, under arbitrary arrival patterns and connection holding times, and under any acceptable Connection Admission policy.

The following identity, obtained by interchanging summation and integration, is important for the subsequent development.

$$\sum_{i \in \mathcal{A}(t)} \widehat{b}_i^\pi(t) = \sum_{i \in \mathcal{A}(t)} \frac{\int_0^t b_i^\pi(s)ds}{V_i} = \int_0^t \sum_{i \in \mathcal{A}(t)} \frac{b_i^\pi(s)}{V_i} ds. = \int_0^t \sum_{i \in \mathcal{N}(s)} \frac{b_i^\pi(s)}{V_i} ds. \quad (7)$$

The last equality follows from the fact that by definition $b_i^\pi(s) = 0$ if a connection is not present in the system at time t. Note that since $\mathcal{A}(t)$ is independent of the employed Bandwidth Adaptation Policy, in order to maximize $\widehat{B}^\pi(t)$ it suffices to maximize $\sum_{i \in \mathcal{A}(t)} \widehat{b}_i(t)$, or according to (7), to maximize,

$$\int_0^t \sum_{i \in \mathcal{N}(s)} \frac{b_i^\pi(s)}{V_i} ds.$$

Observe next that a) by the definition of an acceptable Connection Admission policy the set $\mathcal{N}(s)$, is independent of the employed Bandwidth Adaptation policy, and b) The Bandwidth Adaptation policy, may assign bandwidths to connections in $\mathcal{N}(s)$ independently of the bandwidths assigned at other times. Hence in order to maximize $\int_0^t \sum_{i\in\mathcal{N}(s)} \frac{b_i^{\pi}(s)}{V_i} ds$, it suffices to maximize $\sum_{i\in\mathcal{N}(t)} \frac{b_i^{\pi}(s)}{V_i}$ at any time s, subject to the constraints imposed by the channel bandwidth, the Connection Admission policy and the acceptable bandwidth levels. The constraint for the channel bandwidth is,

$$\sum_{i\in\mathcal{N}(s)} b_i^{\pi}(s) \le B. \tag{8}$$

Different Connection Admission Policies impose different additional constraints on the choice of connection bandwidths. For example, a Complete Sharing policy does not impose any further constraints. A Threshold policy imposes the following additional constraints

$$\sum_{i\in\mathcal{N}_c(s)} b_i^{\pi}(s) \le H_c, \ \ c \in \mathcal{C}.$$

In general, the Connection Admission policy restricts the vector of possible bandwidth allocations $(b_i^{\pi}(s)))_{i\in\mathcal{N}(s)}$ to lie in a region $\mathcal{A}(s)$.

We summarize the discussion above in the following theorem.

Theorem 1 *Under arbitrary connection arrival pattern, connection holding times and connection bandwidth levels, the policy π^* that at any time $s \ge 0$ allocates to connection $i \in \mathcal{N}(s)$ bandwidth $b_i^{\pi^*}(s) = b_i^*$, where b_i^* is the solution to the following optimization problem,*

$$\max \left\{ \sum_{i\in\mathcal{N}(s)} \frac{b_i}{V_i} \right\}$$

$$\sum_{i\in\mathcal{N}(s)} b_i^{\pi}(s) \le B. \tag{9}$$

$$(b_i)_{i\in\mathcal{N}(s)} \in \mathcal{A}(s), \tag{10}$$

$$b_i \in \mathcal{B}_i, \ \ i \in \mathcal{N}(s), \tag{11}$$

maximizes $\widehat{B}^{\pi}(t)$ for all $t \ge 0$.

Conditions (9), (10) represent respectively the channel bandwidth and Admission Policy Constraints, while those in (11) represent the constraints on the connection bandwidth levels.

Note: The solution to the optimization problem of Theorem 1 depends on s only through the set $\mathcal{N}(s)$, $\mathcal{A}(s)$. So, as long as $\mathcal{N}(s)$ and $\mathcal{A}(s)$ remain the same the bandwidth allocated to connections by π^* need not change. Since a change of $\mathcal{N}(s)$ and $\mathcal{A}(s)$ occurs only at connection arrival and departure instances, the Bandwidth Adaptation policy may be invoked only at these instances.

4 Multiple Class System

In this section we deal with a multiclass system. We first consider the problem of optimizing a linear combination of class performance measures, under quite general conditions. Next, with additional statistical assumptions on the arrival rates and holding times, we derive the performance region of the system, i.e., the region in $\mathbb{R}^C$ where the vector of class performances takes values. Finally, we consider the case where concave (instead of linear) rewards are associated with each class and provide an optimal Bandwidth Adaptation Policy.

For definiteness and simplicity in the presentation, in the discussion that follows we assume that the Connection Admission policy is Complete Sharing, the additional conditions (10) will not be explicitly described . Unless otherwise specified, the results hold for any acceptable Connection Admission policy, by incorporating appropriate constraints of the form (10).

4.1 Linear Rewards

In this section we assume that connections belonging to class c are accepted by the system at a long-term rate λ_c. Specifically,

$$\lim_{t\to\infty} \frac{A(t)}{t} = \lambda_c,\ \ 0 < \lambda_c < \infty,\ \ c \in \mathcal{C}. \tag{12}$$

Since connections belong to multiple classes, it may be desirable to provide discriminatory service depending on the class to which a connection belongs. The simplest such class discriminatory service is to associate with class c a reward $r_c \geq 0$ per unit of received performance (i.e., reward per unit of average class scaled mean bandwidth) . Let $\mathbf{r} = \{r_c\}_{c\in\mathcal{C}}$. Then, the total system reward at time t under Bandwidth Adaptation policy π becomes,

$$\widehat{B}^\pi_{\mathbf{r}}(t) = \sum_{c\in\mathcal{C}} r_c \widehat{B}^\pi_c(t) = \sum_{c\in\mathcal{C}} r_c \frac{\int_0^t \sum_{i\in\mathcal{N}_c(s)} \frac{b_i^\pi(s)}{V_i} ds}{A_c(t)} = \frac{1}{t}\int_0^t \sum_{c\in\mathcal{C}(s)} \sum_{i\in\mathcal{N}_c(s)} \frac{t r_c b_i^\pi(s)}{A_c(t) V_i} ds. \tag{13}$$

It may be argued that $V_i = h_i \overline{B}_i$ already discriminates connections through the scaling factors $\overline{B}_i$. As mentioned in Section 2, in case it is desirable to eliminate this extra factor, one can set instead $\overline{B}_i = 1$ without affecting the rest of the results.

The next theorem provides a policy that optimizes the system performance measure $\widehat{B}^\pi_{\mathbf{r}}(t)$ asymptotically, as $t \to \infty$. Since we do not make any assumptions on the connection holding times, bandwidth levels and on the manner a Bandwidth Adaptation policy operates, the limit $\lim_{t\to\infty} \widehat{B}^\pi_{\mathbf{r}}(t)$ may not exist. Therefore, regarding the asymptotic measures of performance, we consider either $\limsup_{t\to\infty} \widehat{B}_{\mathbf{r}}^{\pi}(t)$ or $\liminf_{t\to\infty} \widehat{B}_{\mathbf{r}}^{\pi}(t)$.

Theorem 2 *Let the connection acceptance rate be λ_c, $c \in \mathcal{C}$, $0 < \lambda_c < \infty$, and the connection holding times and bandwidth levels arbitrary. Let π^* be the Bandwidth Adaptation policy that at any time t allocates to connection $i \in \mathcal{N}(t)$ bandwidth $b_i^{\pi^*}(t) = b_i^*$, where b_i^* is the solution to the following optimization problem.*

$$\max\left\{\sum_{c\in\mathcal{C}(t)}\sum_{i\in\mathcal{N}_c(t)} \frac{t}{A_c(t)}\frac{r_c}{V_i}b_i\right\} \tag{14a}$$

$$\sum_{c\in\mathcal{C}(t)}\sum_{i\in\mathcal{N}_c(t)} b_i \le B \tag{14b}$$

$$b_i \in \mathcal{B}_i,\ i \in \mathcal{N}_c(t),\ c \in \mathcal{C}(t). \tag{14c}$$

Then, it holds for any Bandwidth Adaptation policy π.

$$\limsup_{t\to\infty} \widehat{B}_{\mathbf{r}}^{\pi}(t) \le \limsup_{t\to\infty} \widehat{B}_{\mathbf{r}}^{\pi^*}(t),$$
$$\liminf_{t\to\infty} \widehat{B}_{\mathbf{r}}^{\pi}(t) \le \liminf_{t\to\infty} \widehat{B}_{\mathbf{r}}^{\pi^*}(t).$$

Proof. The proof is given in Appendix A. ■

Note: Theorem 2 shows asymptotic optimality of the proposed policy, while Theorem 1 states that the policy is optimal at any time t. The reason for this discrepancy is that the integrand in the last term of (13) contains the factor $t/A_c(t)$, which is in effect the estimate of the inverse of connection acceptance rate λ_c at time t. In contrast the integrand in (7) is only a function of the integration variable s. If λ_c is known apriori, and in (14a) we replace $t/A_c(t)$ with $1/\lambda_c$, the proposed policy remains asymptotically optimal (this observation will be useful in Section 4.2). In this case, it follows from the discussion in Section 3, that the resulting policy optimizes at any time t, the quantity

$$\frac{1}{t}\int_0^t \sum_{c\in\mathcal{C}(s)}\sum_{i\in\mathcal{N}_c(s)} \frac{r_c b_i^{\pi}(s)}{\lambda_c V_i} ds.$$

As with the single-class system, since $\mathcal{C}(t)$ and $\mathcal{N}_c(t)$ are changing only at connection arrival and departure instants, the bandwidth allocated to connections by policy π^* may change only at those instants.

In this section we derive the performance region of a multiclass system. For this, we will need to make further statistical assumptions on the system parameters. These assumptions are the following.

Assumptions

(1) The arrival processes $A_c(t)$, $c \in \mathcal{C}$, are independent and Poisson. The arrival rate of class c connections is λ_c^I, $0 < \lambda_c^I < \infty$. We set $\lambda^I = \sum_{c\in\mathcal{C}} \lambda_c^I$. Note that λ_c^I is the rate connections arrive to the system, while λ_c denotes the connection acceptance rate by the system.
(2) The connection holding times are i.i.d. per class and independent among classes.
(3) Each class c has associated bandwidth levels $\mathcal{B}_c$, that are common for all connections belonging to the class. By $\underline{B}_c$ and $\overline{B}_c$ we denote the minimum and maximum level in $\mathcal{B}_c$, respectively. Without loss of generality we assume that $\underline{B}_c \leq B$, $c \in \mathcal{C}$ (if $\underline{B}_c > B$, connections from class c are never admitted in the system and can therefore be excluded from further consideration).

Taking into account that the Connection Admission policy is Complete Sharing, we conclude that under any Bandwidth Adaptation policy, a class e connection arriving at time t is admitted by the system if,

$$\sum_{c\in\mathcal{C}(t)} \underline{B}_c N_c(t^-) + \underline{B}_e \leq B, \tag{15}$$

where $N_c(t^-)$ is the number of class c connections in the system before the decision to accept or reject the newly arriving connection is made. This, and Assumptions 1 and 2 imply through the Insensitivity Property [24, Theorem 5.3] that the stationary distribution $P_S(\mathbf{n})$ of the process $\mathbf{N}(t) = \{N_c(t)\}_{c\in\mathcal{C}}$ exists and that

$$\lim_{t\to\infty} \Pr\left(\mathbf{N}(t) = 0\right) = P_S(\mathbf{0}) > \mathbf{0}.$$

We now state some preliminary results that follow from the regenerative structure of system processes (see [29] for the definition and a nice introduction to regenerative processes). Assume for simplicity that $\mathbf{N}(0) = \mathbf{0}$. Let $T_0 = 0$ and T_k be the kth time that the system empties, i.e., $\mathbf{N}(T_k^-) > 0$ and $\mathbf{N}(T_k) = \mathbf{0}$. Due to the assumptions stated above, $\{T_k\}_{k=1}^{\infty}$ is an i.i.d. process. We call $[T_k, T_{k+1})$ the "kth regeneration period" of the system. Let U_k, $k = 1, 2, ...,$ be the length of time in $[T_k, T_{k+1})$ that the system is empty, i.e., $\mathbf{N}(t) = \mathbf{0}$, $T_k \leq t < T_k + U_k$. Since connection arrivals are Poisson and a connection arriving in an empty system is always admitted due to the fact that $\underline{B}_c \leq B$,

$c \in \mathcal{C}$ it holds, $\mathrm{E}[U_1] = 1/\lambda^I$. Using regenerative arguments, we have,

$$\begin{aligned}\lim_{t\to\infty} \Pr\left(\mathbf{N}(t) = 0\right) &= \frac{\mathrm{E}\left[U_1\right]}{\mathrm{E}\left[T_1\right]} \\ &= \frac{1}{\lambda^I \mathrm{E}\left[T_1\right]} = P_S(\mathbf{0}) > 0,\end{aligned}$$

hence $\mathrm{E}[T_1] < \infty$. Using again regenerative arguments it can be seen that, the long-term average rates of connections that are accepted in the system, λ_c, $c \in \mathcal{C}$, exist and

$$0 < \lim_{t\to\infty} \frac{A_c\left(t\right)}{t} = \lambda_c = \frac{\mathrm{E}\left[A_c(T_1)\right]}{\mathrm{E}\left[T_1\right]} \leq \lambda_c^I. \tag{16}$$

Given a vector $\mathbf{r} = \{r_c\}_{c\in C}$, $r_c \geq 0$, consider the policy $\pi^{\mathbf{r}}$ that operates as follows:

policy $\pi^{\mathbf{r}}$: At time t allocate to connection $i \in \mathcal{N}(t)$ bandwidth $b_i^{\pi^{\mathbf{r}}}(t) = b_i^*$, where b_i^* is the solution to the following optimization problem.

$$\max\left\{\sum_{c\in\mathcal{C}(t)} \sum_{i\in\mathcal{N}_c(t)} \frac{r_c}{\lambda_c} \frac{b_i}{V_i}\right\} \tag{17a}$$

$$\sum_{c\in\mathcal{C}(t)} \sum_{i\in\mathcal{N}_c(t)} b_i \leq B \tag{17b}$$

$$b_i \in \mathcal{B}_c,\ i \in \mathcal{N}_c(t),\ c \in \mathcal{C}(t). \tag{17c}$$

Note that (17) differs from (14) only in that $t/A_c(t)$ is replaced by $1/\lambda_c$.

Consider the auxiliary variables

$$\widetilde{b}_c^{\pi^{\mathbf{r}}}(t) = \sum_{i\in\mathcal{N}_c(t)} \frac{b_i^{\pi^{\mathbf{r}}}(t)}{V_i}.$$

Under the stated assumptions, and since policy $\pi^{\mathbf{r}}$ depends only on the current state of the network, we conclude that the process $\left\{\widetilde{b}_c^{\pi^{\mathbf{r}}}(t)\right\}_{c\in\mathcal{C}}$ is regenerative with respect to T_k . Hence the long term average of $\widetilde{b}_c^{\pi^{\mathbf{r}}}(t)$ exist,

$$\lim_{t\to\infty} \frac{1}{t} \int_0^t \sum_{i\in\mathcal{N}_c(s)} \frac{b_i^{\pi^{\mathbf{r}}}(s)}{V_i} ds = \frac{\mathrm{E}\left[\int_0^{T_1} \widetilde{b}_c^{\pi^{\mathbf{r}}}(s) ds\right]}{\mathrm{E}\left[T_1\right]} = \widetilde{B}_{\mathbf{r},c}\ .$$

From (16) and (eql7) we conclude that limit of class performance metric under $\pi^{\mathbf{r}}$ exists,

$$\lim_{t\to\infty} \widehat{B}_c^{\pi^{\mathbf{r}}}(t) = \lim_{t\to\infty} \frac{t}{A_c(t)} \frac{1}{t} \int_0^t \sum_{i\in\mathcal{N}_c(s)} \frac{b_i^{\pi^{\mathbf{r}}}(s)}{V_i} ds = \frac{\widetilde{B}_{\mathbf{r},c}}{\lambda_c} \triangleq \widehat{B}_{\mathbf{r},c},$$

and hence the limit of system performance metric under $\pi^{\mathbf{r}}$ also exists,

$$\lim_{t\to\infty} \widehat{B}_{\mathbf{r}}^{\pi^{\mathbf{r}}}(t) = \lim_{t\to\infty} \frac{1}{t}\int_0^t \sum_{c\in\mathcal{C}} \sum_{i\in\mathcal{N}_c(s)} \frac{t}{A_c(t)} r_c \frac{b_i^{\pi}(s)}{V_i} ds = \sum_{c\in\mathcal{C}} r_c \widehat{B}_{\mathbf{r},c} \triangleq f(\mathbf{r}).$$

According to the discussion in Section 4.1 (see note at the end of that section) we have that under any Bandwidth Adaptation policy π,

$$\limsup_{t\to\infty} \sum_{c\in\mathcal{C}} r_c \widehat{B}_c^{\pi}(t) \le f(\mathbf{r}). \tag{18}$$

Let $\mathcal{Q}$ be the convex hull of the set of points

$$\mathcal{B} = \left\{ \mathbf{b} : \ \mathbf{b} = \left\{ \widehat{B}_{\mathbf{r},c} \right\}_{c\in\mathcal{C}}, \mathbf{r} = \{r_c\}_{c\in\mathcal{C}}, r_c \ge 0, \ \sum_{c\in\mathcal{C}} r_c > 0 \ \right\},$$

i.e., all linear combinations of the form $\sum_{j=1}^{J} p_j \mathbf{b}_j$, where $\mathbf{b}_j \in \mathcal{B}$ and $\sum_{j=1}^{J} p_j = 1$ [1]. Let also $\mathcal{P}$ be the set of all points that are coordinate-wise smaller than some point in $\mathcal{Q}$, that is,

$$\mathcal{P} = \left\{ \mathbf{b} = \{b_c\}_{c\in\mathcal{C}} : b_c \le b_c^q, \text{ for some } \mathbf{b}^q = \{b_c^q\}_{c\in\mathcal{C}} \in \mathcal{Q} \right\}.$$

The following lemma states that the system performance lies asymptotically in the region $\mathcal{P}$. We denote $\widehat{\mathbf{B}}^{\pi}(t) = \left\{ \widehat{B}_c^{\pi}(t) \right\}_{c\in\mathcal{C}}$.

Lemma 3 *Under any Bandwidth Adaptation policy π,*

$$\lim_{t\to\infty} \left(\inf\left\{ \left\| \widehat{\mathbf{B}}^{\pi}(t) - \mathbf{b} \right\| : \mathbf{b} \in \mathcal{P} \right\} \right) = 0.$$

Proof. Assume the contrary, that is,

$$\limsup_{t\to\infty} \left(\inf\left\{ \left\| \widehat{\mathbf{B}}^{\pi}(t) - \mathbf{b} \right\| : \mathbf{b} \in \mathcal{P} \right\} \right) > 0.$$

This, and the fact that $\widehat{\mathbf{B}}^{\pi}(t)$ is bounded imply that there is a subsequence t_k such that

$$\lim_{k\to\infty} \widehat{\mathbf{B}}^{\pi}(t_k) = \widehat{\mathbf{B}}^{\pi}, \tag{19}$$

and

$$\inf\left\{ \left\| \widehat{\mathbf{B}}^{\pi} - \mathbf{b} \right\| : \mathbf{b} \in \mathcal{P} \right\} > 0,$$

i.e., $\widehat{\mathbf{B}}^{\pi} \notin \mathcal{P}$. Since $\mathcal{P}$ is compact, by the Strict Separation Theorem [30], there is a vector $\mathbf{r}$ that separates $\widehat{\mathbf{B}}^{\pi}$ and $\mathcal{P}$, that is,

$$\sum_{c\in\mathcal{C}} r_c \widehat{B}_c^{\pi} > \sum_{c\in\mathcal{C}} r_c b_c, \text{ for all } \mathbf{b} \in \mathcal{P}. \tag{20}$$

[1] As can be seen from (17) policy $\pi^{\mathbf{r}}$ is identical to policy $\pi^{\gamma\mathbf{r}}$ where $\gamma > 0$ is any constant. Hence the set $\mathcal{B}$ can be restricted to $\mathcal{B}' = \left\{ \mathbf{b} : \ \mathbf{b} = \left\{ \widehat{B}_{\mathbf{r},c} \right\}_{c\in\mathcal{C}}, \mathbf{r} = \{r_c\}_{c\in\mathcal{C}}, r_c \ge 0, \sum_{c\in\mathcal{C}} r_c = 1 \right\}$.

We may assume without loss of generality that $r_c \geq 0,\ c \in \mathcal{C}$. Indeed, let $r_e < 0$ for some $e \in \mathcal{C}$. Notice that by the definition of $\mathcal{P}$, if $\mathbf{b} \in \mathcal{P}$ then the vector $\mathbf{b}'$ where $b'_c = b_c$ if $c \neq e,\ b_e = 0$, also belongs to $\mathcal{P}$. Define also, $r'_c = r_c$ if $c \neq e,\ r'_e = 0$. By (20) we have for any $\mathbf{b} \in \mathcal{P}$,

$$\begin{aligned}\sum_{c\in\mathcal{C}} r'_c \widehat{B}^\pi_c &= \sum_{c\in\mathcal{C}-\{e\}} r_c \widehat{B}^\pi_c \\ &\geq \sum_{c\in\mathcal{C}} r_c \widehat{B}^\pi_c \quad \text{since } r_e \widehat{B}^\pi_e \leq 0 \\ &> \sum_{c\in\mathcal{C}} r_c b'_c \ \ \text{since } \mathbf{b}' \in \mathcal{P} \\ &= \sum_{c\in\mathcal{C}} r'_c b_c \text{ by definition or } \mathbf{r}'_c \text{ and } \mathbf{b}'.\end{aligned}$$

Therefore, for any e such that $r_e < 0$ we may set $r_e = 0$ and the resulting vector will still be separating and will have nonnegative coordinates. Assume therefore that the separating vector in (20) has nonnegative coordinates and consider the point $\mathbf{b} = \left\{\widehat{B}_{\mathbf{r},c}\right\}_{c\in\mathcal{C}}$ which belongs to $\mathcal{P}$. We then have from (19) and (20) that

$$\lim_{k\to\infty} \sum_{c\in\mathcal{C}} r_c \widehat{B}^\pi_c(t_k) > \sum_{c\in\mathcal{C}} r_c \widehat{B}_{\mathbf{r},c} = f(\mathbf{r}).$$

However, the last inequality contradicts (18). ■

Notice the difference between (18) and Lemma 3. While (18) states that for a *given* $\mathbf{r}$ the system performance of any Bandwidth Adaptation policy is worse than the system performance of the policy $\pi^{\mathbf{r}}$, Lemma 3 states that the *performance vector* (i.e., the vector of class performances) of any Bandwidth Adaptation policy is componentwise smaller than the performance vector of some linear combination of performance vectors of a set of policies $\pi^{\mathbf{r}_j}$, $j = 1, .., J$.

The region $\mathcal{P}$ is characterized implicitly through the vectors $\mathbf{r}$ with nonnegative coordinates. For the purposes of developing optimal policies in Section 5, this characterization will suffice.

The structure of $\mathcal{P}$ is in general very complicated. A significant simplification is obtained when the Connection Admission policy is Complete Sharing, if we assume that all connections in a given class have the same holding times and that class bandwidth levels can take any value between the minimum and maximum possible. In this case, we have that $V_i = V_c$ for all connections $i \in c$, and it turns out that $\mathcal{P}$ is a polymatroid [31], as the following lemma indicates. Although we do not make use of this structure in our derivations, it gives insight on the structure of the performance region and may be useful for other applications.

Lemma 4 . *If the Connection Admission policy is Complete Sharing, con-*

nections in class c have the same holding time h_c and the bandwidth levels for class c can take any value between $\underline{B}_c$ and $\overline{B}_c$, then the region $\mathcal{P}$ is a polymatroid, i.e., the space in $\mathbb{R}^C$ defined by the following inequalities,

$$\sum_{c\in\mathcal{S}} \alpha_c x_c \leq F(\mathcal{S}), \ \textit{for all } \mathcal{S} \subseteq \mathcal{C},$$
$$x_c \geq 0, \ \textit{for all } c \in \mathcal{C},$$

where $\alpha_c = \lambda_c h_c \overline{B}_c$ and the set function $F(\mathcal{S})$ is submodular, i.e., it satisfies,

$$F(\mathcal{S}) = \varnothing,$$
$$F(\mathcal{S}) \leq F(\mathcal{T}), \ \textit{if } \mathcal{S} \subseteq \mathcal{T},$$
$$F(\mathcal{S}) + F(\mathcal{T}) \geq F(\mathcal{S} \cup \mathcal{T}) + F(\mathcal{S} \cap \mathcal{T}).$$

Proof. The proof is given in Appendix B. ■

As can be seen from the proof, the function $F(\mathcal{S})$ represents the average bandwidth received by connections belonging to classes in the set $\mathcal{S}$, when the system gives as much bandwidth as possible to these connections before proceeding to allocate bandwidth to connections belonging to classes in the set $\mathcal{C} - \mathcal{S}$. In Figure 2 we show the region $\mathcal{P}$ when there are two classes in the system. The point A represents the class performance vector when the system gives as much bandwidth as possible to connections in class 1 before allocating bandwidth to connections in class 2. Point B represents the class performance vector when the roles of classes 1 and 2 are interchanged. The line segment AB represents the set $\mathcal{Q}$; any point in this segment is a linear combination of the point $A$and B.

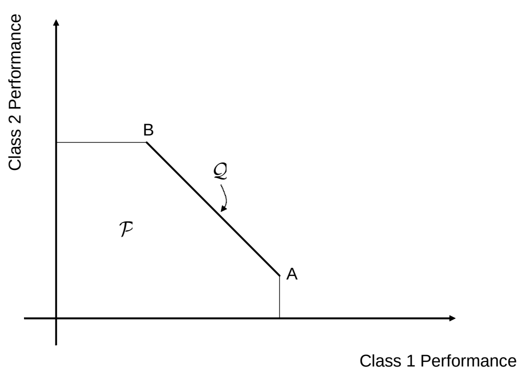

Fig. 2. The region of class performance vector for a two-class system.

While the structure of $\mathcal{P}$ is significantly simplified under the assumption of Lemma 4, it is still difficult to compute the values of $F(\mathcal{S})$ and this calculation demands knowledge of the statistical parameters of the system. However, this will not be an obstacle for the policies that will be presented in the next section.

5 An Adaptive Policy for Separable Convex Optimization

In this section we assume that with each class $c \in \mathcal{C}$ there is an associated reward function $f_c(b)$ so that, under Bandwidth Adaptation policy π, the reward of class c for receiving performance $\widehat{B}_c^{\pi}(t)$ by time t, is $f_c\left(\widehat{B}_c^{\pi}(t)\right)$. The function $f_c(b)$ is assumed to be concave, twice continuously differentiable. Then, the total system reward by time t is,

$$\sum_{c \in \mathcal{C}} f_c\left(\widehat{B}_c^{\pi}(t)\right). \tag{21}$$

Various choices of $f_c(x)$ provide various fairness criteria [32].. We describe below the most common ones.

- *Linear utilities.* In this case, $f_c(x) = r_c x$.This criterion is relevant when, for example, r_c represents the (monetary) reward received when class c has long-term performance metric x and the objective of the system is to maximize its reward.
- *Proportional Fairness.* In this case, $f_c(x) = \log(x)$. This allocation has several important properties discussed in [32]. Intuitively, since the $\log(x)$ function increases very slowly, this type of utilities express the fact that the satisfaction received by a given increase in bandwidth is higher if the already allocated bandwidth is small.
- *Max-min Fairness.* Intuitively here one attempts to maximize the bandwidth of the classes with the minimal allocated bandwidth, while splitting evenly whatever bandwidth remains to the rest of the classes. While this problem cannot be directly translated in the form (21) it can be well approximated by choosing

$$\phi_c(x) = c - g(x)^m$$

where c is a constant and g a differentiable, decreasing, convex and positive function.

We will define a Bandwidth Adaptation policy π^* that maximizes the long-term system reward, i.e. for any other policy π,

$$\limsup_{t \to \infty} \sum_{c \in \mathcal{C}} f_c\left(\widehat{B}_c^{\pi}(t)\right) \leq \lim_{t \to \infty} \sum_{c \in \mathcal{C}} f_c\left(\widehat{B}_c^{\pi^*}(t)\right).$$

In Section 4.1 we presented a simple optimal policy for the case of linear rewards, i.e., $f_c(b) = r_c b$ and in Section 4.2 we presented the performance region of the system under any Bandwidth Adaptation policy. We therefore have all the ingredients to apply the framework proposed in [20] in order to develop an optimal adaptive policy π^*. The framework in [20] is based on a methodology used in [19]. We present the policy π^*, describe the manner it which the various parameters used by the policy are updated and provide some intuition as to why is works. For technical details on the method the reader is referred to the above mentioned references.

The class c performance metric at time t,

$$\widehat{B}_c^{\pi}(t) = \frac{\sum_{i \in \mathcal{A}_c(t)} \widehat{b}_i^{\pi}(t)}{A_c(t)},$$

is the average of scaled bandwidths allocated to class c connections that arrived to the system up to time t. This metric can be easily measured and updated on-line (see below). The proposed policy below uses these measurements.

> **Optimal Policy π^*.** Let t_n be the successive times when a connection arrival or departure occurs. At time t_n, $n = 1, 2, ...$, set
>
> $$r_c(t_n) = \left. \frac{df_c(b)}{db} \right|_{b = \widehat{B}_c^{\pi^*}(t_n))}.$$
>
> At time $t \in (t_n, t_{n+1}]$, allocate to connection $i \in \mathcal{N}(t)$ bandwidth $b_i^{\pi^*}(t) = b_i^*$, where b_i^* is the solution to the following optimization problem.
>
> $$\max \left\{ \sum_{c \in \mathcal{C}(t)} \sum_{i \in \mathcal{N}_c(t)} r_c(t_n) \frac{t_n}{A_c(t_n)} \frac{b_i}{V_i} \right\} \tag{22a}$$
>
> $$\sum_{c \in \mathcal{C}(t)} \sum_{i \in \mathcal{N}_c(t)} b_i \leq B \tag{22b}$$
>
> $$b_i \in \mathcal{B}_c, \ i \in \mathcal{N}_c(t), \ c \in \mathcal{C}(t). \tag{22c}$$

It is important to note that no system statistics are required to be known for the policy to operate. All the required parameters are computed on-line. The metrics $\widehat{B}_c^{\pi^*}(t)$ can be updated at time t_n as follows. Then, since in the interval (t_n, t_{n+1}) no connection arrives or departs and no bandwidth updates occur, we have,

$$\mathcal{N}_c(t_{n+1}) = \begin{cases} \mathcal{N}_c(t_n) \cup \{j\} & \text{if a class } c \text{ connection } j \text{ is accepted at time } t_n \\ \mathcal{N}_c(t_n) - \{j\} & \text{if a class } c \text{ connection } j \text{ departs at time } t_n \\ \mathcal{N}_c(t_n) & \text{otherwise} \end{cases} \tag{23}$$

$$\mathcal{A}_c(t_{n+1}) = \begin{cases} \mathcal{A}_c(t_n) \cup \{j\} & \text{if a class } c \text{ connection } j \text{ is accepted at time } t_n \\ \mathcal{A}_c(t_n) & \text{otherwise} \end{cases}$$

and

$$
\begin{aligned}
\widehat{B}_c^{\pi^*}(t_{n+1}) &= \frac{\sum_{i\in\mathcal{A}_c(t_{n+1})} \widehat{b}_i^{\pi^*}(t_{n+1})}{A_c(t_{n+1})} \\
&= \frac{\sum_{i\in\mathcal{A}_c(t_n)} \widehat{b}_i^{\pi^*}(t_n) + \sum_{i\in\mathcal{N}_c(t_n)} \widehat{b}_i^{\pi^*}(t_n)\frac{(t_{n+1}-t_n)}{V_i}}{\mathcal{A}_c(t_{n+1})} \\
&= \frac{\mathcal{A}_c(t_n)}{\mathcal{A}_c(t_n+1)}\widehat{B}_c^{\pi^*}(t_n) + \frac{\sum_{i\in\mathcal{N}_c(t_n)} \widehat{b}_i^{\pi^*}(t_n)\frac{(t_{n+1}-t_n)}{V_i}}{\mathcal{A}_c(t_n+1)} \qquad (24)
\end{aligned}
$$

Hence at time t_{n+1}, knowing $\widehat{B}_c^{\pi^*}(t_n)$ and $\mathcal{A}_c(t_n)$ and the bandwidths $\widehat{b}_i^{\pi^*}(t_n)$ allocated to connections that are in the system at time t_n, one can obtain $\mathcal{A}_c(t_{n+1})$ and $\widehat{B}_c^{\pi^*}(t_{n+1})$ by using (23), (24)

Based on the computed parameters, the linear optimization problem in (22) needs to be solved. The complexity of this problem depends on the additional constraints imposed by the Connection Admission policy and of whether the bandwidth level sets $\mathcal{B}_c$, $c \in \mathcal{C}$ allow continuous or discrete variations of connection bandwidths. When the Connection Admission policy is Complete Sharing, the constraints are exactly those in (22). If moreover the connection bandwidths of all classes c can take any value in $[\underline{B}_c, \overline{B}_c]$ then the linear problem has a very simple and efficient solution, described in the Appendix as part of the proof of Lemma 4.

Policy π^* is essentially a stochastic version of the "conditional gradient algorithm" used in deterministic nonlinear optimization [30]. We outline next this method and discuss its relation to policy π^*.

Consider the optimization problem

$$
\begin{aligned}
&\max \ f(\mathbf{x}) && (25a)\\
&\mathbf{x}\in\mathcal{X}, && (25b)
\end{aligned}
$$

where $\mathbf{x} = (x_1, ..., x_N)$ is an $N-$dimensional vector, $f(\mathbf{x})$ a real function and $\mathcal{X}$ a compact $N-$dimensional set. The conditional gradient algorithm generates a sequence of points $\mathbf{x}_n$, $n = 1, 2, ...$ that under certain conditions converge to the optimal solution of (25). Point $\mathbf{x}_{n+1}$ is obtained from $\mathbf{x}_n$ as follows.

(1) The solution $\overline{\mathbf{x}}_n$ to the following linear optimization problem is obtained.

$$
\begin{aligned}
&\max \sum_{i=1}^{N} \frac{\partial f(\mathbf{x}_n)}{\partial x_{n,i}} x_i && (26a)\\
&\mathbf{x} \in \mathcal{X}. && (26b)
\end{aligned}
$$

(2) The new point $\mathbf{x}_{n+1}$ is obtained by the update

$$\mathbf{x}_{n+1} = a_n \mathbf{x}_n + (1 - a_n)\overline{\mathbf{x}}_n \tag{27}$$

where α_n are appropriately chosen parameters.

In the stochastic environment we are considering, the analogue of $\mathbf{x}_n$ is the vector $\widehat{\mathbf{B}}^{\pi^*}(t_n) = \left(\widehat{B}_c^{\pi^*}(t_n)\right)_{c \in \mathcal{C}}$. Since our optimization objective function is $f\left(\widehat{\mathbf{B}}^{\pi^*}(t_n)\right) = \sum_{c \in \mathcal{C}} f_c\left(\widehat{B}_c^{\pi}(t_n)\right)$, we have

$$\frac{\partial f\left(\widehat{\mathbf{B}}^{\pi^*}(t_n)\right)}{\partial \widehat{B}_c^{\pi^*}(t_n)} = \left.\frac{df_c(b)}{db}\right|_{b=\widehat{B}_c^{\pi^*}(t_n))}$$

According to Section 4.1, employing (22) provides a long-term solution to the optimization problem with linear rewards $\partial f\left(\widehat{\mathbf{B}}^{\pi^*}(t_n)\right)/\partial \widehat{B}_c^{\pi^*}(t_n),\ c \in \mathcal{C}$, which is the analogue of (26). Finally, the analogue of (27) is (24). Of course a number of technical issues need to be addressed in order to ensure that the deterministic algorithm can be adapted in a stochastic environment. These issues are addressed in [19], [20].

6 Unknown Holding Times

The development above assumed that the connection holding times are known. This assumption is valid when the connections represent prestored multimedia content, however there are other cases, e.g. videoconferencing, where connection holding times are not known apriori. However, based on the insight obtained, heuristics can be developed so that the policy operates also with unknown times, although we cannot claim optimality in this case. Below we present such a heuristic.

In a multiclass system it is reasonable to assume that the holding times of connections in a given class c are random variables concentrated around their mean value H_c. The basic idea of the heuristic is to employ a policy $\overline{\pi}$ that operates as π^* but by replacing the connection holding times with their expected value, as long as these connections reside in the system; hence for a connection i in class c we use $V_i = H_c \overline{B}_c$ (as before, $\overline{B}_c$ may be replaced by other scaling factors). Upon connection departure, the connection holding times are known and therefore their exact value can be used to update the calculation of $\widehat{B}_c^{\overline{\pi}}(t_n)$. More specifically, the update of $\widehat{B}_c^{\overline{\pi}}(t_n)$ is done as follows.

The system keeps track of the total bandwidth consumed by connection i,

$$\widetilde{b}_i^{\overline{\pi}} = \int_{a_i}^{d_i} b_i^{\overline{\pi}}(t)dt.$$

- If at time t_{n+1} no connection from class c departs, the update is the same as in (24).

$$\widehat{B}_c^{\overline{\pi}}(t_{n+1}) = \frac{\mathcal{A}_c(t_n)}{\mathcal{A}_c(t_n+1)}\widehat{B}_c^{\overline{\pi}}(t_n) + \frac{\sum_{i\in\mathcal{N}_c(t_n)} b_i^{\overline{\pi}}(t_n)\frac{(t_{n+1}-t_n)}{V_i}}{\mathcal{A}_c(t_n+1)} \tag{28}$$

- If at time t_{n+1} connection i from class c departs, then the update becomes,

$$\begin{aligned}\widehat{B}_c^{\overline{\pi}}(t_{n+1}) = {} & \frac{\mathcal{A}_c(t_n)}{\mathcal{A}_c(t_n+1)}\widehat{B}_c^{\overline{\pi}}(t_n) + \frac{\sum_{i\in\mathcal{N}_c(t_n)} b_i^{\overline{\pi}}(t_n)\frac{(t_{n+1}-t_n)}{V_i}}{\mathcal{A}_c(t_n+1)} \\ & + \frac{1}{\mathcal{A}_c(t_n+1)}\left(\frac{\widetilde{b}_i^{\overline{\pi}}}{h_i\overline{B}_c} - \frac{\widetilde{b}_i^{\overline{\pi}}}{H_c\overline{B}_c}\right)\end{aligned} \tag{29}$$

  The last term in (29) accounts for the correction done upon obtaining the knowledge of connection service time: while connection i resides in the system, its contribution to class performance is $\widetilde{b}_i/\left(H_c\overline{B}_c\right)$, which is replaced by the real contribution $\widetilde{b}_i/\left(h_c\overline{B}_c\right)$ upon connection departure.

Various modifications to the policy can be proposed. For example, assume that connection i resides in the system for time τ_i. If the statistics of the average connectional holding time $\mathrm{E}[h_i\,|\tau_i]$ is known, then this value may be used in the evaluation of V_i. However, this requires more statistical information and increases the complexity of the algorithm.

We next evaluate policy $\overline{\pi}$ by simulations.

[describe the simulation environment and what you plan to study]

7 Discussion and Suggestions for Further Work

We considered the problem of channel sharing by rate-adaptive multi-class streams. We presented policies maximizing a linear combination of average scaled connection bandwidths, under quite general conditions. Under additional but reasonable statistical assumptions, we described the performance region of the system and applied a general methodology to provide a Bandwidth Adaptation policy for maximizing a combination of convex class performance rewards.

In Section 2 we assumed that the channel bandwidth is fixed. If the channel bandwidth is a function of time as may be the case in wireless systems, then the results in Sections 3 and 4.1 still hold. Regarding the results in Sections 4.2 and 5, extensions are still possible by adding statistical assumptions on the channel bandwidth fluctuation. While these assumptions may not be severe, several technical questions need to be answered in order to completely justify the optimality of the proposed policy.

In this work, we concentrated on connection bandwidth as the basic performance measure. The results can be easily applied if as measure for the perceptual quality of reception is considered a function $r(b(t))$ of the bandwidth received by a connection at time t [12], [14]. The main difference is that now one has to solve nonlinear optimization problems at the decision instants.

Our final comment regards the adaptation capabilities of the proposed policy. Since the policy is based on the observation of past history (through the quantities $\widehat{B}_c^{\pi^*}(t)$), in order to adapt faster to changes in statistical parameters, it is important to introduce techniques that weigh recent history more. A standard way of doing this, is to work with the weighted means of the quantities involved. It can be shown that this way the adaptivity of the policy can be improved at the expense of small deviation from optimality, see e.g., the techniques used in [33].

There are several directions for further research on this subject. In this work we concentrated on average scaled connection bandwidth as the main performance metric. However, as commented in Section 2 there are other metrics that affect the end-user perception quality of the received stream. While the issue of which combination of these metrics is appropriate has not been resolved yet, it is generally accepted that the frequency of bandwidth updates and/or the size of the incremental updates does play a role. An interesting feature of the proposed policies is that bandwidth updates may take place only at connection arrival and departure times, hence update frequency is limited. However, depending on the application, even this may not be acceptable. Research towards the direction of taking into account multiple metrics has been presented in [12], [11], [14]. We believe that it is possible to extend the framework presented in this work to incorporate multiple metrics from system and end-user perspectives. Work is under way to address this problem. Another direction of research is extension of the policies presented here to a network environment. Work towards this direction has been done in [17], where a single class system was considered. The multiclass optimization, however is still an open issue.

References

[1] A. Eleftheriadis, D. Anastasiou, Optimal data partitioning of MPEG-2 coded video, in: First IEEE International Conference on Image Processing, 1994, pp. 273–277.

[2] P. Pancha, M. Zarki, Prioritized transmission of variable bit rate MPEG video, in: IEEE GLOBECOM, 1992, pp. 1135–1139.

[3] N. G. Duffield, K. K. Ramakrishnan, A. R. Reibman, SAVE: an algorithm for smoothed adaptive video over explicit rate networks, IEEE/ACM Transactions on Networking 6 (6) (1998) 717–728.

[4] J. Bolot, T. Turletti, A rate control mechanism for packet video in the internet, in: IEEE INFOCOM, 1994, pp. 1216–1223.

[5] D. Taubman, A. Zakhor, A common framework for rate and distortion based scaling of highly scalable compressed video, IEEE Transactions on Circuits and Systems for Video Technology 6 (4) (1996) 329–354.

[6] H. Kankia, P. P. Mishra, A. Reibmand, An adaptive congestion control scheme for real-time video transport, IEEE/ACM Transactions on Networking 3 (6) (1995) 671–682.

[7] T. V. Lakshman, P. Mishra, K. K. Ramakrishnan, Transporting compressed video over ATM networks with explicit rate feedback control, in: IEEE INFOCOM, Kobe, Japan, 1997.

[8] P. P. Mishra, Fair bandwidth sharing for video sources using distributed feedback control, in: IEEE GLOBECOM, Singapore, 1995.

[9] M. Grossglauser, S. Keshav, D. Tse, RCBR: A simple and efficient service for multiple time-scale traffic, in: ACM SIGCOMM'95, 1995.

[10] C. Adjih, N. Argiriou, M. Chaudier, E. Deberdt, F. Dumontet, L. Georgiadis, P. Jacquet, An architecture for IP quality of service provisioning over CATV networks, in: EMMSEC, Sweden, 1999.

[11] A. K. Talukdar, B. R. Badrinath, A. Acharya, Rate adaptation schemes in networks with mobile hosts, in: Proceedings of ACM/IEEE MobiCom '98, 1998, pp. 169–180.

[12] V. Bharghavan, K. Lee, S. Lu, S. Ha, D. Dwyer, The TIMELY adaptive resource management architecture, IEEE Personal Communications Magazine 5 (4) (1998) 20–31.

[13] T. Kwon, S. Kim, Y. Choi, N. Naghshineh, Threshold-type call admission control in Wireless/Mobile multimedia networks using prioritized adaptive framework, Electronics Letters 36 (9) (2000) 852–854.

[14] T. Kwon, Y. Choi, S. K. Das, Bandwidth adaptation algorithms for adaptive multimedia services in mobile cellular networks, Wireless Personal Communications 22 (2002) 337–357.

[15] S. Weber, G. de Veciana, Telecommunications Network Design, Academic Publisher, 2002, Ch. Asymptotic Analysis of Rate Adaptive Multimedia Streams.

[16] T. Kwon, Y. Choi, C. Bisdikian, M. Naghshineh, QoS provisioning in Wireless/Mobile multimedia networks using and adaptive framework, Wireless Networks 9 (2003) 51–59.

[17] S. Weber, G. Veciana, Rate adaptive multimedia streams: Optimization, admission control, and distributed algorithms, IEEE Transactions on Networking 1 (1) (2003) 1–16.

[18] N. Argiriou, L. Georgiadis, Channel sharing by rate-adaptive streaming applications, Performance Evaluation 55 (3-4) (2004) 211–229.

[19] P. P. Bhattacharya, L. Georgiadis, P. Tsoucas, Problems of adaptive optimization in multiclass M/GI/1 queues with bernoulli feedback, Mathematics of Operations Research 20 (1995) 356–380.

[20] V. Tsibonis, L. Georgiadis, An adaptive framework for addressing fairness issues in wireless networks, Computer Communications Journal 28 (10) (2005) 1167–1178.

[21] G. J. Foschini, B. Gopinath, Sharing memory optimally, IEEE Transactions on Communications 31 (3) (1983) 352–360.

[22] K. W. Ross, D. Tsang, The stochastic knapsack problem, IEEE Transactions on Communications 37 (1989) 934–939.

[23] A. Gavious, Z. Rosberg, A restricted complete sharing policy for a stochastic knapsack problem in a b-ISDN, IEEE Transactions on Communications 42 (1994) 2375–2379.

[24] K. W. Ross, Multiservice Loss Models for Broadband Telecommunication Networks, Springer, 1995.

[25] ITU, ITU-T recommendation P.861: Objective quality measurements of telephone band and (300-3400 hz) speech codecs, Tech. rep. (1998).

[26] ITU, ITU-500-R recommendation BT.500-8: Methodology for the subjective assessment of the quality of television pictures, Tech. rep. (1998).

[27] S. Winkler, A perceptual distortion metric for digital color video, in: Proceedings SPIE Human Vision and Electronic Imaging, Vol. 3644, 1999, pp. 175–184.

[28] A. R. Prasad, R. Esmailzadeh, S. Winkler, T. Ihara, B. Rohani, B. Pinguet, M. Capel, Perceptual quality measurement and control: Definition, application and performance, in: Proceedings 4th International Symposium on Wireless Personal Multimedia Communications, Aarborg, Denmark, 2001, pp. 547–552.

[29] R. W. Wolff, Stochastic Modeling and the Theory of Queues, Prentice Hall, 1989.

[30] D. R. Bertsekas, Nonlinear Programming, Athena Scientific, 1995.

[31] J. Edmonds, Submodular functions, matroids and certain polyhedra, in: Calgary International Conference of Combinatorial Structures and Applications, Calgary, Alta, 1969, pp. 69–87.

[32] J.-Y. Boudec, Rate adaptation, congestion control and fairness: A tutorial, Tech. rep.
URL `citeseer.ifi.unizh.ch/boudec00rate.html`

[33] P. P. Bhattacharya, L. Georgiadis, P. Tsoucas, I. Viniotis, Adaptive lexicographic optimization in multi-class M/GI/1 queues, Mathematics of Operations Research 18 (1993) 705–740.

APPENDIX

A Proof of Theorem 2

Proof. Since $\lim_{t\to\infty} \frac{A_c(t)}{t} = \lambda_c$ for all $c \in \mathcal{C}$, and $\lambda_c > 0$, $c \in \mathcal{C}$, it follows that for any $\varepsilon > 0$ there is a time t_ε such that

$$\lambda_c^{-1} - \varepsilon \leq \frac{t}{A_c(t)} \leq \lambda_c^{-1} + \varepsilon, \text{ for all } t \geq t_\varepsilon, \; c \in \mathcal{C}. \tag{A.1}$$

Let

$$\lambda_{\max} = \max_{c\in\mathcal{C}} \{\lambda_c\}, \; \lambda_{\min} = \min_{c\in\mathcal{C}} \{\lambda_c\}$$

In the following, without loss of generality we take $0 < \varepsilon < \lambda_{\max}^{-1}$. Define also,

$$r_{\max} = \max_{c\in\mathcal{C}} \{r_c\}, \; V_{\min}^{\varepsilon} = \min_{i\in A(t_\varepsilon)} \{V_i\}.$$

Clearly, we have $0 < \lambda_{\max}^{-1} \leq \lambda_{\min}^{-1} < \infty$, $r_{\max} < \infty$. Also notice that since $\lambda_c < \infty$, it holds $A_c(t_\varepsilon) < \infty$, $c \in \mathcal{C}$, and hence $V_{\min}^{\varepsilon} > 0$. Considering a time $t \geq t_\varepsilon$ and taking into account (13), (A.1), we have,

$$\begin{aligned}
\widehat{B}_{\mathbf{r}}^{\pi}(t) &= \frac{1}{t}\sum_{c\in\mathcal{C}} \frac{t}{A_c(t)} \int_0^{t_\varepsilon} \sum_{i\in\mathcal{N}_c(s)} r_c \frac{b_i^\pi(s)}{V_i} ds + \frac{1}{t}\sum_{c\in\mathcal{C}} \frac{t}{A_c(t)} \int_{t_\varepsilon}^{t} \sum_{i\in\mathcal{N}_c(s)} r_c \frac{b_i^\pi(s)}{V_i} ds \\
&\leq \frac{\lambda_{\min}^{-1} + \varepsilon}{t} \int_0^{t_\varepsilon} \sum_{c\in\mathcal{C}} \sum_{i\in\mathcal{N}_c(s)} r_c \frac{b_i^\pi(s)}{V_i} ds + \frac{1}{t}\varepsilon \int_{t_\varepsilon}^{t} \sum_{c\in\mathcal{C}} \sum_{i\in\mathcal{N}_c(s)} r_c \frac{b_i^\pi(s)}{V_i} ds. \\
&\qquad + \frac{1}{t} \int_{t_\varepsilon}^{t} \sum_{c\in\mathcal{C}} \sum_{i\in\mathcal{N}_c(s)} \frac{r_c}{\lambda_c} \frac{b_i^\pi(s)}{V_i} ds.
\end{aligned} \tag{A.2}$$

To proceed we need to provide estimates on certain of the quantities appearing in (A.2). First note that,

$$\int_0^{t_\varepsilon} \sum_{c\in\mathcal{C}} \sum_{i\in\mathcal{N}_c(s)} r_c \frac{b_i^\pi(s)}{V_i} ds \le \frac{r_{\max}}{V_{\min}^\varepsilon} \int_0^{t_\varepsilon} \sum_{c\in\mathcal{C}} \sum_{i\in\mathcal{N}_c(s)} b_i^\pi(s) ds = \frac{r_{\max}}{V_{\min}^\varepsilon} \int_0^{t_\varepsilon} \sum_{i\in\mathcal{N}(s)} b_i^\pi(s) ds$$
$$\le \frac{r_{\max}}{V_{\min}^\varepsilon} B t_\varepsilon \equiv M_\varepsilon, \tag{A.3}$$

where the last inequality follows from the fact that under any policy π, inequality (1a) must be satisfied. Next, note that since $\widehat{b}_i^\pi(t) \le 1$, it follows that $\widehat{B}_c^\pi(t) \le 1$, $c \in \mathcal{C}$. Taking also into account (13) and the left hand side of (A.1) we have for $t \ge t_\varepsilon$,

$$\left(\lambda_{\max}^{-1} - \varepsilon\right) \frac{1}{t} \int_{t_e}^{t} \sum_{c\in\mathcal{C}} \sum_{i\in\mathcal{N}_c(s)} r_c \frac{b_i^\pi(s)}{V_i} ds \le \frac{1}{t} \int_0^t \sum_{c\in\mathcal{C}} \sum_{i\in\mathcal{N}_c(s)} \frac{t}{A_c(t)} r_c \frac{b_i^\pi(s)}{V_i} ds$$
$$= \sum_{c\in\mathcal{C}} r_c \widehat{B}_c^\pi(t) \le \sum_{c\in\mathcal{C}} r_c. \tag{A.4}$$

From (13), (A.1), (A.3) and (A.4) we have,

$$\widehat{B}_{\mathbf{r}}^\pi(t) \le \frac{\lambda_{\min}^{-1} + \varepsilon}{t} M_\varepsilon + \frac{\varepsilon}{\lambda_{\max}^{-1} - \varepsilon} \sum_{c\in\mathcal{C}} r_c + \frac{1}{t} \int_{t_\varepsilon}^{t} \sum_{c\in\mathcal{C}} \sum_{i\in\mathcal{N}_c(s)} \frac{r_c}{\lambda_c} \frac{b_i^\pi(s)}{V_i} ds. \tag{A.5}$$

By the definition of policy π^* we have that for any time $s > 0$,

$$\sum_{c\in\mathcal{C}} \sum_{i\in\mathcal{N}_c(t)} r_c \frac{b_i^\pi(s)}{V_i} \frac{s}{A_c(s)} \le \sum_{c\in\mathcal{C}} \sum_{i\in\mathcal{N}_c(t)} r_c \frac{b_i^{\pi^*}(s)}{V_i} \frac{s}{A_c(s)}.$$

Taking into account (A.1) we conclude that for any $s > t_\varepsilon$ it holds,

$$\sum_{c\in\mathcal{C}} \sum_{i\in\mathcal{N}_c(s)} r_c \frac{b_i^\pi(s)}{V_i} \left(\lambda_c^{-1} - \varepsilon\right) \le \sum_{c\in\mathcal{C}} \sum_{i\in\mathcal{N}_c(s)} r_c \frac{b_i^{\pi^*}(s)}{V_i} \left(\lambda_c^{-1} + \varepsilon\right),$$

or

$$\sum_{c\in\mathcal{C}} \sum_{i\in\mathcal{N}_c(s)} \frac{r_c}{\lambda_c} \frac{b_i^\pi(s)}{V_i} - \varepsilon \sum_{c\in\mathcal{C}} \sum_{i\in\mathcal{N}_c(s)} r_c \frac{b_i^\pi(s)}{V_i} \le \sum_{c\in\mathcal{C}} \sum_{i\in\mathcal{N}_c(s)} \frac{r_c}{\lambda_c} \frac{b_i^{\pi^*}(s)}{V_i} + \varepsilon \sum_{c\in\mathcal{C}} \sum_{i\in\mathcal{N}_c(s)} r_c \frac{b_i^{\pi^*}(s)}{V_i}. \tag{A.6}$$

Integrating (A.6) and using (A.4) we have,

$$\frac{1}{t} \int_{t_\varepsilon}^{t} \sum_{c\in\mathcal{C}} \sum_{i\in\mathcal{N}_c(s)} \frac{r_c}{\lambda_c} \frac{b_i^\pi(s)}{V_i} ds \le \frac{1}{t} \int_{t_\varepsilon}^{t} \sum_{c\in\mathcal{C}} \sum_{i\in\mathcal{N}_c(s)} \frac{r_c}{\lambda_c} \frac{b_i^{\pi^*}(s)}{V_i} ds + \frac{2\varepsilon \sum_{c\in\mathcal{C}} r_c}{\lambda_{\max}^{-1} - \varepsilon} \tag{A.7}$$

From (A.5) and (A.7) we conclude,

$$
\begin{aligned}
\widehat{B}^{\pi}_{\mathbf{r}}(t) &\leq \frac{\lambda_{\min}^{-1}+\varepsilon}{t} + \frac{3\varepsilon\sum_{c\in\mathcal{C}} r_c}{\lambda_{\max}^{-1}-\varepsilon} + \frac{1}{t}\int_{t_\varepsilon}^{t}\sum_{c\in\mathcal{C}}\sum_{i\in\mathcal{N}_c(s)}\frac{r_c}{\lambda_c}\frac{b_i^{\pi^*}(s)}{V_i}ds \\
&\leq \frac{\lambda_{\min}^{-1}+\varepsilon}{t} + \frac{3\varepsilon\sum_{c\in\mathcal{C}} r_c}{\lambda_{\max}^{-1}-\varepsilon} + \frac{1}{t}\int_{t_\varepsilon}^{t}\sum_{c\in\mathcal{C}}\sum_{i\in\mathcal{N}_c(s)}\left(\frac{t}{A_c(t)}+\varepsilon\right) r_c\frac{b_i^{\pi^*}(s)}{V_i}ds \\
&\leq \frac{\lambda_{\min}^{-1}+\varepsilon}{t} + \frac{4\varepsilon\sum_{c\in\mathcal{C}} r_c}{\lambda_{\max}^{-1}-\varepsilon} + \frac{1}{t}\int_{0}^{t}\sum_{c\in\mathcal{C}}\sum_{i\in\mathcal{N}_c(s)}\frac{t}{A_c(t)} r_c\frac{b_i^{\pi^*}(s)}{V_i}ds \\
&= \frac{\lambda_{\min}^{-1}+\varepsilon}{t} + \frac{4\varepsilon\sum_{c\in\mathcal{C}} r_c}{\lambda_{\max}^{-1}-\varepsilon} + \widehat{B}^{\pi^*}_{\mathbf{r}}(t)
\end{aligned}
\tag{A.8}
$$

Finally, (A.8) implies,

$$
\limsup_{t\to\infty}\widehat{B}^{\pi}_{\mathbf{r}}(t) \leq \frac{4\varepsilon\sum_{c\in\mathcal{C}} r_c}{\lambda_{\max}^{-1}-\varepsilon} + \limsup_{t\to\infty}\widehat{B}^{\pi^*}_{\mathbf{r}}(t).
$$

Since ε is arbitrary, we conclude that $\limsup_{t\to\infty}\widehat{B}^{\pi}_{\mathbf{r}}(t) \leq \limsup_{t\to\infty}\widehat{B}^{\pi^*}_{\mathbf{r}}(t)$. The inequality for $\liminf$ follows by similar reasoning. ■

B Proof of Lemma 4

The region described in Lemma is known as a polymatroid. Polymatroids were introduced by Edmonds [31], where the following lemma - adapted to conform to our notation - was proved.

Edmonds Lemma. *Consider in $\mathbb{R}^C$ the polyhedron R defined by*

$$
\begin{aligned}
\sum_{i\in\mathcal{S}} a_c x_c \leq f(\mathcal{S}), &\text{ for all } \mathcal{S}\subseteq\mathcal{C} \\
x_c \geq 0, &\text{ for all } c\in\mathcal{C}.
\end{aligned}
$$

where $a_c \geq 0$, $c \in C$. If for any permutation $\sigma(i)$, $i = 1, ..., C$ of the class indices in C the vector obtained as a solution to

$$
\sum_{i=1}^{k} a_{\sigma(i)}x_{\sigma(i)} = f(\{\sigma(1), ..., \sigma(i)\}\,,\ k = 1, ...C,
$$

belongs to R, then the function f is submodular.

We can now proceed with the proof of the lemma.

Proof. Under the assumptions of the lemma, the optimization problem on which policy $\pi^{\mathbf{r}}$ is based becomes,

$$\max\left\{\sum_{c\in\mathcal{C}(t)}\frac{r_c}{\alpha_c}\sum_{i\in\mathcal{N}_c(t)} b_i\right\} \tag{B.1a}$$

$$\sum_{c\in\mathcal{C}(t)}\sum_{i\in\mathcal{N}_c(t)} b_i \le B, \tag{B.1b}$$

$$\underline{B}_c \le b_i \le \overline{B}_c,\ i\in\mathcal{N}_c(t),\ c\in\mathcal{C}(t). \tag{B.1c}$$

Set

$$B_c = \sum_{i\in\mathcal{N}_c(t)} b_i, \tag{B.2}$$

and consider the problem,

$$\max\left\{\sum_{c\in\mathcal{C}(t)}\frac{r_c}{\alpha_c}B_c\right\} \tag{B.3a}$$

$$\sum_{c\in\mathcal{C}(t)} B_c \le B \tag{B.3b}$$

$$\underline{b}_c N_c(t) \le B_c \le \overline{b}_c N_c(t),\ c\in\mathcal{C}(t). \tag{B.3c}$$

Any optimal solution to problem (B.1) provides through (B.2) a solution to problem (B.3) with the same value. Alternatively, any optimal solution to (B.3) provides a solution to (B.1) with the same value, by setting,

$$b_i = \frac{B_c}{N_c(t)},\ i\in\mathcal{N}_c(t). \tag{B.4}$$

Hence we may assume without loss of generality that policy $\pi^{\mathbf{r}}$ allocates bandwidth according to (B.3) and (B.4). The general solution to the optimization (17) has been discussed in [17]. For the specific form of (B.1), the solution consists of the following steps.

(1) Sort the coefficients $\frac{r_c}{\alpha_c}$ of all classes in $\mathcal{C}(t)$, in nonincreasing order. Let $\sigma(i)$ denote the class in $\mathcal{C}(t)$ whose order is i in this sorting, i.e.,

$$\frac{r_{\sigma(i)}}{\alpha_{\sigma(i)}} \ge \frac{r_{\sigma(i+1)}}{\alpha_{\sigma(i+1)}},\ i = 1, ..., C(t)-1.$$

(2) Select $\overline{n}$ as, $\overline{n} = \max\left\{n : \sum_{k=1}^{m}\overline{B}_{\sigma(i)}N_{\sigma(i)}(t) + \sum_{k=m+1}^{C(t)}\underline{B}_{\sigma(i)}N_{\sigma(i)}(t) \le B\right\}$

(3) If $\overline{n} = C(t)$, then set $B_c = \overline{B}_c N_c(t)$ for all $c\in\mathcal{C}(t)$. Else,

(4) Set

$$\begin{aligned}
B_{\sigma(i)} &= \overline{B}_{\sigma(i)}N_{\sigma(i)}(t), 1\le i\le \overline{n},\\
B_{\sigma(i)} &= \underline{B}_{\sigma(i)}N_{\sigma(i)}(t),\ \overline{n}+2\le i\le C(t),\\
B_{\sigma(\overline{n}+1)} &= B - \left(\sum_{k=1}^{\overline{n}}\overline{B}_{\sigma(i)}N_{\sigma(i)}(t) + \sum_{k=\overline{n}+2}^{C(t)}\underline{B}_{\sigma(i)}N_{\sigma(i)}(t)\right).
\end{aligned}$$

We observe from the above that the performance of policy $\pi^{\mathbf{r}}$ depends only on the order induced by vector $\mathbf{r}$ on the quantities $r_c/\alpha_c,\ c \in \mathcal{C}$. If any other vector $\mathbf{r}'$ induces the same order, then the resulting policy $\pi^{\mathbf{r}'}$ will have the same performance values $\widehat{\mathbf{B}}_{\mathbf{r}}$. Consider now a set $\mathcal{S} \subseteq \mathcal{C}$ and let $\Pi_{\mathcal{S}}$ be the class of policies $\pi^{\mathbf{r}}$ for which the induced order $\sigma_{\mathbf{r}}$ is such that the first S higher order classes is the set $\mathcal{S}$, that is, $\mathcal{S} = \{c : \ c = \sigma_{\mathbf{r}}(i),\ i = 1, 2, ...S\}$. For any policy in $\Pi_{\mathcal{S}}$, we see from the steps above, that the total bandwidth allocated to the classes in $\mathcal{S}$ at time t, is the same, $B_{\mathcal{S}}(t)$. Moreover, under any policy $\pi^{\mathbf{r}}$ not in $\Pi_{\mathcal{S}}$, the total bandwidth allocated to the classes in $\mathcal{S}$ at time $t, B_{\mathcal{S}}^{\pi^{\mathbf{r}}}(t)$, is at most $B_{\mathcal{S}}(t)$. Using regenerative arguments it can be seen that the steady-state average values of $B_{\mathcal{S}}(t)$ and $B_{\mathcal{S}}^{\pi^{\mathbf{r}}}(t)$ exist, i.e.,

$$\lim_{t\to\infty} \frac{1}{t} \int_0^t B_{\mathcal{S}}^{\pi^{\mathbf{r}}}(s)ds = \lim_{t\to\infty} \frac{1}{t} \int_0^t \sum_{c\in\mathcal{S}} \sum_{i\in\mathcal{N}_c(t)} b_i^{\pi^{\mathbf{r}}} dt = \sum_{c\in\mathcal{S}} \alpha_c \widehat{B}_c^{\pi^{\mathbf{r}}} \triangleq F(\mathcal{S}) \text{ for } \pi^{\mathbf{r}} \text{ in } \Pi_{\mathcal{S}}, \tag{B.5}$$

and

$$\lim_{t\to\infty} \frac{1}{t} \int_0^t B_{\mathcal{S}}^{\pi^{\mathbf{r}}}(s)ds = \sum_{c\in\mathcal{S}} \alpha_c \widehat{B}_c^{\pi^{\mathbf{r}}} \leq F(\mathcal{S}) \text{ for } \pi^{\mathbf{r}} \text{ not in } \Pi_{\mathcal{S}}. \tag{B.6}$$

From (B.5) and (B.6), using similar arguments to those used in the proof of Lemma 4 we conclude that for any $\mathbf{x} \in \mathcal{P}$ it holds,

$$\sum_{c\in\mathcal{S}} \alpha_c x_c \leq F(\mathcal{S}),\ \mathcal{S} \subseteq \mathcal{C}, \quad x_c \geq 0,\ c \in \mathcal{C}.$$

Finally, using regenerative arguments it can be seen that given any permutation order $\sigma_{\mathbf{r}}$ induced by a policy $\pi_{\mathbf{r}}$, it holds

$$\sum_{i=1}^{k} a_{\sigma_{\mathbf{r}(i)}} \widehat{B}_{\sigma_{\mathbf{r}(i)}}^{\pi^{\mathbf{r}}} = F\left(\{\sigma_{\mathbf{r}}(1), ...\sigma_{\mathbf{r}}(k)\}\right),\ k = 1, ..., C.$$

Since $\widehat{\mathbf{B}}^{\pi^{\mathbf{r}}} \in \mathcal{P}$, the proof is completed by resorting to Edmonds Lemma. ■