Optimal Admission Control for High Speed Networks: A Dynamic Programming Approach

Tania Jiménez, IEEE Student Member
CeSiMo (U.L.A., Venezuela)
and INRIA, 2004, route des Lucioles B.P.93
Sophia Antipolis - FRANCE,
Email: Tania.Jimenez@sophia.inria.fr

Abstract

We consider in this paper the problem of Call Admission Control of guaranteed performance (GP) connections (such as the CBR and VBR traffic classes in ATM) in the presence of best effort (BE) connections that use the bandwidth left over by the guaranteed performance connections. We assume that the BE sessions do not require a minimum cell rate (MCR) and are thus not subject to call admission control. By slightly increasing rejection rate of GP sessions one may decrease dramatically the delay of BE sessions. We formulate the admission problem as a Markov decision problem and obtain the optimal policy. In particular, we show that it is of a switching curve type. We then compare numerically the performance of the optimal policy to threshold policies as well as to the policy which ignores the BE traffic (and accepts GP sessions as long as there is available bandwidth for them). We show that threshold policies are good approximations for the overall optimal policy.

1 Introduction

Optimal admission control for high speed networks with traffic that requires guaranteed Quality of Service (QoS) is still an open and challenging problem. Network traffic can be classified in two types: traffic from applications which need a guaranteed QoS and traffic from more flexible applications that can handle large variations of bandwidth and delays. These two types of traffic will be noted GP (Guaranteed Performance) and BE (Best Effort) respectively.

In ATM networks, the Connection Admission Control (CAC) function is used for managing the network's resources and decide if a new connection can be accepted or must be rejected. Because GP and BE traffic share the same resource, the CAC policy affects the performance of both types of traffic. Yet most approaches used to design CACs focused only on GP traffic. We consider also BE traffic in the admission control decision and we show that an optimal policy that takes into account the two types of traffic performs better than a greedy policy that always accepts GP traffic sessions ignoring the performance of BE sessions.

In [3] a similar problem has been studied under more general statistical assumptions. It was shown numerically (yet not proved) that a limitting control problem that is obtained in the heavy traffic regime has optimal policies given by (decreasing) switching curves. In this paper we restrict to exponential assumptions on the session duration. Under this assumption we show that this structure holds not only for the limitting heavy traffic regime but for any loading of the system. (Since the results in [3] are only sensitive to the first moment of the processes and not to the whole distribution, our results here imply that the switching curve structure of optimal policies is in fact optimal in the heavy traffic limit under much weaker assumptions on the distribution of the session durations.)

We used an event-based dynamic programming approach. This type of approach have been successfully used in other Call Admission Control problems which did not involve BE traffic, see [4, 5]. We show that our model has some properties which implies that the optimal policy has a switching curve form that is easily computed.

We compare the optimal policy with others simpler policies, numerically. We show that threshold policies are good aproximations to the optimal policy.

The paper is organized as follows. The next section 2 presents the model. Section 3 defines increasingness, convexity and supermodularity properties and presents the proof that the optimal cost in our model has such properties. This is shown to imply that the optimal policies are described via monotone decreasing switching curves. In section 4 we compare the optimal policies obtained by dynamic programming and other simpler policies. We finish with some conclusions in section 5.

2 The model

We consider the above mentioned two types of traffic arriving to a single shared resource. The shared resource typically represents a link or a router in a network.[1]

In this model, the variables are the number of traffic sessions of each type, let x_0 denote the number of sessions of BE traffic and x_1 the GP traffic ones.

We assume that the number of GP sessions is always bounded by a constant $x_1 \leq B$, whereas we do not assume a bound for the number of BE sessions. We also assume that all BE traffic sessions share equally the bandwidth unused by the GP traffic.

Traffic sessions arrival are Poisson with parameters λ_0 and

[1] We do not consider here the case of several resources. The latter case has been successfully addressed in the past by decomposing the network into a set of independent resources (links) assumed to have independent traffic and reward process [6].

λ_1 and each GP connection remains for a period exponentially distributed with parameter μ_1.
Using the standard uniformization technique [7], (in which we sample the continuous time controlled Markov chain with a Poisson process of rate $\nu \geq \lambda_0 + \lambda_1 + \bar{\mu}$; $\bar{\mu} \geq B\mu_1$) the problem is reduced to a discrete time MDP, with

- **State space** $Y = \{\mathbb{N} \times [0, ..., B]\}$,
- **Action space** $W = \{0, 1\}$, (where 0 corresponds to rejection and 1 to acceptance), where 0 is the only action available at state (x_0, B).
- **Immediate cost** ρ composed of the sum of a convex increasing holding cost $C(x_1)$ for BE sessions, and a constant cost c for rejection of a GP session. We assume that f grows at most polynomially. Let ρ_n be the immediate cost at the nth transition.
- **Policy** A policy u is a sequence $(u_1, u_2, ...)$ where u_n is the decision rule at time n. u_n is allowed to depend on the entire history (i.e. all previous states and actions as well as the current state at time n); furthermore it is allowed to be a random decision rule. More precisely, u_n is a function from the history till time n to the set of probabilities on W. We shall be in particular interested by the class of *pure stationary* policies, i.e. policies that choose an action (without randomizing) only as a function of the current state.

An initial state $x = (x_0, x_1)$ together with a policy u define a probability measure P_x^u over the trajectories of states and actions. We denote by E_x^u the corresponding expectation.

Remark Our choice of costs represents the fact that congestion is felt differently by GP and BE sessions. Indeed, the impact of congestion on GP sessions is the rejection of calls, whereas the impact of congestion on BE sessions is in long delays. We are thus in particular interested in holding costs of the form $C(x_0) = C \cdot x_0$. The average holding cost is then proportional to the average queue length, which is in turn proportional to the *average delay* of best effort sessions (due to Little's Law). An alternative cost function that has been considered in the literature corresponds to the total revenue for the network [2]; although it concerns also BE traffic, it does not include their delay.

We are interested in minimizing the following *infinite horizon expected discounted cost*: $C_\alpha(x, u) = \sum_{n=1}^{\infty} \alpha^{n-1} E_x^u \rho_n$

To solve the infinite horizon problem, we define first the *finite horizon* cost: $C_\alpha^m(x, u) = v \sum_{n=1}^{m} \alpha^{n-1} E_x^u \rho_n$

Define the followings events:

- A BE traffic session arrival : $\mathcal{A}_0(x) = (x_0 + 1, x_1)$,
- A GP traffic session arrival : $\mathcal{A}_1(x) = (x_0, x_1 + 1)$,
- A session departure of BE traffic: $\mathcal{D}_0(x) = (\max(x_0 - 1, 0), x_1)$,
- A session departure of GP traffic: $\mathcal{D}_1(x) = (x_0, \max(x_1 - 1, 0))$.

Define the value function of the infinite horizon control problem as $C_\alpha(x) = \inf_u C_\alpha(x, u)$. We also define the value function of the finite horizon problem as $V_n = \inf_u C_\alpha^n(x, u)$. We shall use the following dynamic programming results [7]:

Lemma 1 *1. The optimal infinite horizon cost is the unique solution (among those that are polynomially bounded) of*

$$\phi(x) = T_\alpha \phi(x) := C(x) + \alpha S\phi(x) \tag{1}$$

where

$$S\phi(x) = \nu^{-1}\Big((\bar{\mu} - x_1\mu_1)\phi(\mathcal{D}_0 x) + x_1\mu_1\phi(\mathcal{D}_1 x) + \lambda_0\phi(\mathcal{A}_0 x) + \lambda_1 \min\big(\phi(x) + c, \phi(\mathcal{A}_1 x)\big) + (\nu - \bar{\mu} - \lambda_0 - \lambda_1)\phi(x)\Big)$$

2. Consider the pure stationary policy u that chooses at state x action 0 if and only if the first term achieves the minimum in the above minimization. Then u is optimal.
3. The optimal cost for a horizon n is given recursively by:

$$V_{n+1}(x) = T_\alpha V_n(x), \quad V_0 = 0. \tag{2}$$

4. Value iteration: the optimal value for the infinite horizon problem is given as $\lim_{n\to\infty} T_\alpha^n(0)$.

Definition A map $g : Y \to W$ is called *monotone decreasing switching curve* if g is decreasing. This means that
– for each fixed x_0 there is a threshold $l_0(x_0)$ such that $g(x_0, x_1) = 0$ if and only if $x_1 \geq l_0(x_0)$.
– for each fixed x_1 there is a threshold $l_1(x_1)$ such that $g(x_0, x_1) = 0$ if and only if $x_0 \geq l_1(x_1)$.

It follows from step 2 that a sufficient condition for the policy u to have a monotone decreasing switching curve structure is that $\phi(\mathcal{A}_1 x) - \phi(x)$ be monotone increasing.
The approach used in this paper for deriving monotonicity results, and, in particular, the structure of the optimal policy is based on [8].

3 The Structure of the Optimal Policy

Let us rewrite (2):

$$V_{n+1}(x) = C(x) + \frac{\alpha}{\nu}\Big(\bar{\mu}V_n(\mathcal{D}_0 x) + x_1\mu_1\big(V_n(\mathcal{D}_1 x) - V_n(\mathcal{D}_0 x)\big) + \lambda_0 V_n(\mathcal{A}_0 x) + \lambda_1 \min\big(V_n(x) + c, V_n(\mathcal{A}_1 x)\big) + (\nu - \bar{\mu} - \lambda_0 - \lambda_1)V_n(x)\Big) \tag{3}$$

Define the following operators:

$$T_{cost} f(x) = C(x) + \alpha f(x)$$
$$T_{arr_0} f(x) = f(x + e_0)$$
$$T_{ac} f(x) = min\big(c + f(x), f(x + e_1)\big)$$
$$T_{dep} f(x) = \begin{cases} x_1 f(x - e_1)^+ + (B - x_1) f(x - e_0)^+ & \text{if } x_1 < B \\ B f(x - e_1) & \text{otherwise} \end{cases}$$
$$T_{unif}(f_1(x), f_2(x), f_3(x), f_4(x)) = (\lambda_0/\nu)f_1(x) + (\lambda_1/\nu)f_2(x) + (\mu_1/\nu)f_3(x) + (\upsilon/\nu)f_4(x)$$

where: $\upsilon = \nu - \bar{\mu} - \lambda_1 - \lambda_0$, e_i denotes the vector with 1 in the ith component and zero in all the others components and $x^+ = max(x, 0)$ componentwise.

The value iteration for computing V_n can be rewritten as:

$$V_{n+1}(x) = T_{cost}\Bigg(T_{unif}\Big(T_{arr_0}\big(V_n(x)\big), T_{ac}\big(V_n(x)\big), T_{dep}\big(V_n(x)\big), V_n(x)\Big)\Bigg) \tag{4}$$

The following properties are defined for functions of the state if the corresponding inequalities hold:

Increasingness. $Inc(i) : f(x) \le f(x + e_i)$, for $i = 0, 1$

Supermodularity.
$Super : f(x + e_0) + f(x + e_1) \le f(x) + f(x + e_0 + e_1)$,
We note that $C(x)$ satisfies properties $Inc(i), Super$.
We first present a few Lemmas that establish some properties of the dynamic programming operators. This will enable us to prove the optimality of switching curve policies.

Lemma 2 *The following relations hold for* $i = 0, 1$*:*
1. $T_{cost} : Inc(i) \Longrightarrow Inc(i)$, *2.* $T_{unif} : Inc(i) \Longrightarrow Inc(i)$, *3.* $T_{arr_0} : Inc(i) \Longrightarrow Inc(i)$, *4.* $T_{ac} : Inc(i) \Longrightarrow Inc(i)$, *5.* $T_{dep} : Inc(i) \Longrightarrow Inc(i)$.

Proof: The two first operators are based on the operators with the same name in section (3) of [8] and as exposed there, the results follow directly as these properties are closed under convex combinations. For T_{arr_0}, named $T_{A(i)}$ in [8], the result is straightforward using the variable change $x = y + e_0$ in the inequalities. The proof for the operator T_{ac} can be found in [8] section (3), lemma 3.1. It remains to establish the proof for the operator T_{dep}. We thus assume that $f(x)$ is increasing in x_0 and x_1 and we show that this implies that so is T_{dep}.

We first check increasingness in x_0: $x_1 f(x - e_1) + (B - x_1)f(x - e_0) \le x_1 f(x - e_1) + (B - x_1)f(x)$. This is equivalent to $(B - x_1)f(y) \le (B - x_1)f(y + e_0)$ which clearly holds. Hence the monotonicity in x_0.
Next we check the monotonicity in x_1:

$$\begin{aligned} &x_1 f(x - e_1) + (B - x_1)f(x - e_0) \le \\ &\quad (x_1 + 1)f(x) + (B - x_1 - 1)f(x - e_0 + e_1) \end{aligned}$$

This is obtained as the sum of the three inequalities: $x_1 f(x - e_1) \le x_1 f(x)$, $f(x - e_0) \le f(x)$, and $(B - x_1 - 1)f(x - e_0) \le (B - x_1 - 1)f(x - e_0 + e_1)$. All these inequalities hold since $f(x)$ is increasing. This establishes the monotonicity in x_1 as well. ■

Definition A function f of the state is said to satisfy the *Partial Superconvexity (PSC) property* if the following inequality holds for all x such that $x + e_1 \le B$.

$$f(x + e_0) + f(x + e_0 + e_1) \le f(x + e_1) + f(x + 2e_0) \quad (5)$$

Clearly $C(x)$ satisfies the PSC property.

Lemma 3 *The following relations hold:*
– $T_{cost} : PSC \Longrightarrow PSC; Super \Longrightarrow Super$.
– $T_{unif} : PSC \Longrightarrow PSC; Super \Longrightarrow Super$.
– $T_{arr_0} : PSC \Longrightarrow PSC; Super \Longrightarrow Super$
– $T_{ac} : PSC \Longrightarrow PSC; Super \Longrightarrow Super$.
– $T_{dep} : PSC \Longrightarrow PSC; Super, PSC \Longrightarrow Super$.

Proof: (5) is part of the property named superconvexity (SuperC(i,j)) in section (4) in [8]. It is established in this paper that the operators $T_{cost}, T_{unif}, T_{arr_0}, T_{ac}$ defined above satisfy this inequality if $f(x)$ does.
We will prove that the operator T_{dep} satisfies (5) on the assumption that $f(x)$ satisfies it. Inequality (5) for T_{dep} has the following form in the case that $0 < x_1 < B$, $x_0 > 0$:

$$\begin{aligned} &x_1 f(x + e_0 - e_1) + (B - x_1)f(x) + (x_1 + 1)f(x + e_0) \\ &\quad +(B - x_1 - 1)f(x + e_1) \le (x_1 + 1)f(x) \\ &\quad +(B - x_1 - 1)f(x - e_0 + e_1) + x_1 f(x + 2e_0 - e_1) \\ &\quad +(B - x_1)f(x + e_0). \end{aligned}$$

This inequality can be seen to be the sum of the inequalities

$$\begin{aligned} &(B - x_1 - 1)\big(f(y + e_0) + f(y + e_0 + e_1)\big) \le \\ &\quad (B - x_1 - 1)\big(f(y + e_1) + f(y + 2e_0)\big) \\ \text{and} \quad &x_1\big(f(y + e_0) + f(y + e_0 + e_1)\big) \le \\ &\quad x_1\big(f(y + 2e_0) + f(y + e_1)\big) \end{aligned}$$

where $y = x - e_1$. Both equations above hold since f satisfies (5). Hence the departure operator satisfies (5) for $0 < x_1 < B$, $x_0 > 0$. It remains to check the boundary cases:
1.- $x_1 = 0$. Inequality (5) for T_{dep} has the following form:

$$\begin{aligned} &Bf(x) + f(x + e_0) + (B - 1)f(x + e_1) \le \\ &\quad f(x) + (B - 1)f(x + e_1 - e_0) + Bf(x + e_0) \end{aligned}$$

which is equivalent to $(B - 1)\big(f(x) + f(x + e_1)\big) \le (B - 1)\big(f(x - e_0 + e_1) + f(x + e_0)\big)$. The inequality above indeed holds since f satisfies (5).
2.- $x_0 = 0$. (5) for T_{dep} has the following form:

$$\begin{aligned} &x_1 f(x + e_0 - e_1) + (B - x_1)f(x) + \\ &\quad (x_1 + 1)f(x + e_0) + (B - x_1 - 1)f(x + e_1) \le \\ &\quad (x_1 + 1)f(x) + (B - x_1 - 1)f(x - e_0 + e_1) \\ &\quad +x_1 f(x + 2e_0 - e_1) + (B - x_1)f(x + e_0) \end{aligned}$$

It can be obtained as the sum of the following inequalities:

$$\begin{aligned} &(B - x_1 - 1)\big(f(x) + f(x + e_1)\big) \le \\ &\quad (B - x_1 - 1)\big(f(x - e_0 + e_1) + f(x + e_0)\big) \\ \text{and} \quad &x_1\big(f(x + e_0 - e_1) + f(x + e_0)\big) \le \\ &\quad x_1\big(f(x) + f(x + 2e_0 - e_1)\big) \end{aligned}$$

These two inequalities indeed hold since f satisfies (5).
3.- $x_1 = B - 1$. Inequality (5) for T_{dep} has the form:

$$\begin{aligned} &(B - 1)f(x + e_0 - e_1) + f(x) + Bf(x + e_0) \le \\ &\quad Bf(x) + (B - 1)f(x + 2e_0 - e_1) + f(x + e_0) \end{aligned}$$

It can be rewritten as:

$$\begin{aligned} &(B - 1)\big(f(x + e_0 - e_1) + f(x + e_0)\big) \le \\ &\quad (B - 1)\big(f(x) + f(x + 2e_0 - e_1)\big) \end{aligned}$$

which again follows since f satisfies inequality (5). We conclude that all the operators satisfy (5) if $f(x)$ does, which establishes the relation $T_{dep} : PSC \Longrightarrow PSC$.

To prove supermodularity we assume that $f(x)$ is supermodular and satisfies inequality (5). The supermodluarity for $T_{dep}f$ has the following form away from the boundaries:

$$\begin{aligned} &x_1 f(x + e_0 - e_1) + (B - x_1)f(x) + \\ &\quad (x_1 + 1)f(x) + (B - x_1 - 1)f(x + e_1 - e_0) \\ \le \quad &x_1 f(x - e_1) + (B - x_1)f(x - e_0) + \\ &\quad (x_1 + 1)f(x + e_0) + (B - x_1 - 1)f(x + e_1) \end{aligned}$$

The inequality above is the sum of the following inequalities:

$$(B-x_1)\big(f(x)+f(x+e_1-e_0)\big) \leq (B-x_1)\big(f(x-e_0)+f(x+e_1)\big),$$

$$x_1\big(f(x+e_0-e_1)+f(x)\big) \leq x_1\big(f(x-e_1)+f(x+e_0)\big),$$

$$f(x)-f(x-e_0+e_1) \leq f(x+e_0)-f(x+e_1)$$

The first and second inequalities hold by supermodularity (using the changes $z=x-e_0$ and $z=x-e_1$ resp.), the third equation holds because $f(x)$ satisfies inequality (5).
Next, we show that supermodularity holds also on the boundaries. Consider four cases: 1) $x_1=0$, and $x_0 \neq 0$; 2) $x_0=0$, and $x_1 \neq 0$; 3) $x_0=0$ and $x_1=0$; and 4) $x_1=B-1$.
1.- $x_1=0$, and $x_0 \neq 0$. The supermodularity for T_{dep} takes the form

$$\begin{aligned} &Bf(x)+f(x)+(B-1)f(x-e_0+e_1) \\ \leq \quad & Bf(x-e_0)+f(x+e_0)+(B-1)f(x+e_1) \end{aligned}$$

This is the sum of $(B-1)\big(f(x+e_1-e_0)+f(x)\big) \leq (B-1)\big(f(x-e_0)+f(x+e_1)\big)$ and $2f(x) \leq f(x-e_0)+f(x+e_0)$. The first equation holds by supermodularity and the second by convexity in direction e_0 of f (note that supermodularity plus equation (5) give convexity in direction e_0).
2.- $x_0=0$, and $x_1 \neq 0$

$$\begin{aligned} &x_1 f(x+e_0-e_1)+(B-x_1)f(x)+(x_1+1)f(x) \\ &+(B-x_1-1)f(x+e_1) \leq \\ &x_1 f(x-e_1)+(B-x_1)f(x) \\ &+(x_1+1)f(x+e_0)+(B-x_1-1)f(x+e_1) \end{aligned}$$

This is the sum of $x_1\big(f(x+e_0-e_1)+f(x)\big) \leq x_1\big(f(x-e_1)+f(x+e_0)\big)$ and $f(x) \leq f(x+e_0)$. The first equation holds by supermodularity and the second by increasigness in x_0.
3.- $x_0=0$ and $x_1=0$

$$\begin{aligned} &Bf(x)+f(x)+(B-1)f(x+e_1) \leq \\ &Bf(x)+f(x+e_0)+(B-1)f(x+e_1) \end{aligned}$$

which reduces to $f(x) \leq f(x+e_0)$. This holds true by incressingness in x_0.
4.- $x_1=B-1$.

$$\begin{aligned} &(B-1)f(x+e_0-e_1)+Bf(x)+f(x) \\ \leq \quad & (B-1)f(x-e_1)+Bf(x+e_0)+f(x-e_0) \end{aligned}$$

This is the sum of $Bf(x+e_0-e_1)+Bf(x) \leq Bf(x-e_1)+Bf(x+e_0)$ and $f(x)+f(x-e_1) \leq f(x-e_0)+f(x+e_0-e_1)$. The first inequality holds by supermodularity and the second due to inequality (5). Note that if $x_0=0$ the second inequality holds by increassingness in x_0. ∎

Using the previous Lemmas we now obtain our main result on the structure of the optimal policy:

Theorem 1 *The optimal admission policy u^* (from step (2) of Lemma 1) is of monotone decreasing switching curve type.*

Proof: We use the value iteration procedure from Lemma 1 (4) as well as an inductive argument. We show that V_n satisfies $Inc(i)$, $i=0,1$, as well as PSC and $Super$. Clearly $V_0(x)=0$ has these properties. We thus make the inductive assumption that V_n satisfies them. It follows from the previous Lemmas that V_{n+1} also satisfies them. It is easily seen that these properties carry on to the limit as $n \to \infty$. Hence, by Lemma 1 (4), the value function $C_\alpha(x)$ satisfies these properties. By definition of u^*, these properties imply its desired structure. ∎

4 Numerical Results

In this section we present some numerical examples and we compare between the performance of the following policies: (i) The optimal switching curve policy, (ii) The greedy policy which always accepts GP sessions if there is available bandwidth for them, and (iii) Threshold policies.
A threshold policy is defined through a threshold l (that does not depend on x_1) that accepts a GP session if the number x_0 of BE sessions is not greater than l.
We shall consider two threshold policies: the *average* threshold policy and the *maximal* threshold policy. These two policies are both constructed from the optimal switching curve policy. The switching curve policy can be viewed as a threshold policy in which the threshold l_1 depends on the state x_1: we accept a GP session if and only if $x_0 \leq l_1(x_1)$.
The average threshold policy is taken by choosing a threshold l_{av} which is the threshold of the switching curve policies averaged over all values of x_1.
The maximal threshold policy is defined through the maximum of all thresholds that appear in the switching curve, i.e. through a threshold $l_{\max} := \max_{x_1} l_1(x_1)$.
The dynamic programming was coded in C language and all the examples were run on a Ultra 5 SUN Station.

A regime of good utilisation of the resources We first consider a situation in which the sum of input rates are less than the total available capacity. More precisely, we chose input rates of $\lambda_0=0.2$ and $\lambda_1=0.2$, the number of resources is $B=10$ and the capacity per resource is $\mu=0.045$. Thus the total available capacity is $\bar{\mu}=0.45$. This can be classified as a regime of good utilisation of the resources: $\hat{\rho}=8/9$ ($\hat{\rho}$ is defined as the ratio between the sum of input rates and the output rate). The discount factor is $\alpha=0.998$.

The following Figures illustrate the performance of all policies: the average delay of BE sessions (2) and the rejection rate of the GP sessions (1). The rejection probabilities are obtained from the rejection rates by dividing them by the input rate λ_1.

We see that the delays of BE are considerably larger under the greedy policy than under the three other policies: the optimal and the threshold ones. The delay performance for BE traffic is almost the same under the latter policies, with slightly smaller delay for the optimal policy. We see from Fig. 1 that the price that we pay in terms of rejection rates for this improvement in delay is very small.

Note that, since we use discounted costs, the averages of the

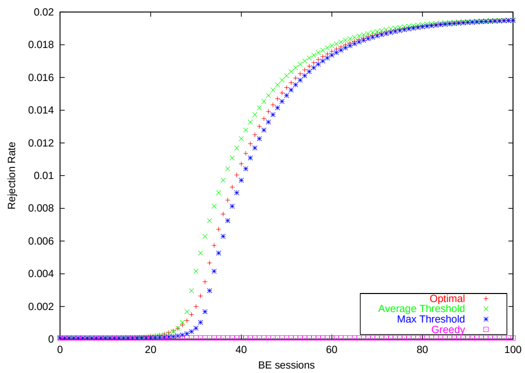

Figure 1: Rejection Rates for Different Policies

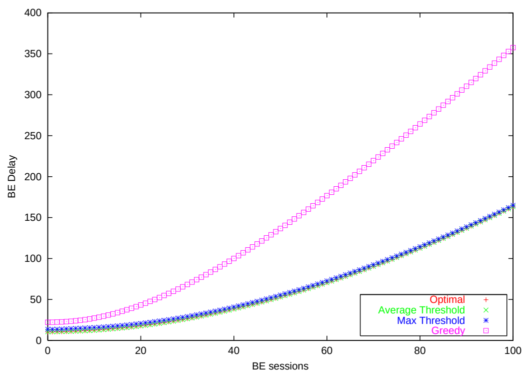

Figure 2: BE Delay for Different Policies

delays and rejection rates are taken in the Abel sense (i.e. they are discounted) rather than in the Cesaro sense, which explains the dependence on the initial state.

Heavy traffic regime We chose a heavy traffic regime, i.e. a situation in which the sum of input rates is close to the total available capacity: $\lambda_0 = 0.2$ and $\lambda_1 = 0.2$, the number of resources is $B = 10$ and the capacity per resource is $\mu = 0.039$. Thus the total available capacity is $\bar{\mu} = 0.39$. Here the input rate is slightly larger than the output rate ($\hat{\rho} > 1$).

Fig. 3 presents the optimal switching curve policies as well as their approximations by threshold policies, for a rejection cost of 150 units per GP session, and for three different linear BE cost funtions $C(x_0) = cx_0$ where we took $c = 5, c = 10$ and $c = 15$. The horizontal (respectively, vertical) axis represents the number of GP (respectively, BE) sessions. The corresponding curve have the following interpretation: above the curve we reject an arriving GP session, whereas below or at the curve we accept an arriving GP session. Thus for each value x_1 of the horizontal axis (number of GP sessions), we accept as long as the number x_0 of BE sessions is not above the value of the curve at x_1. For example, for the cost $c = 10$ we will accept a new GP session when we are in the state $x_1 = 2$ if and only if $x_0 \leq 4$. The discount factor is $\alpha = 0.98$.

The optimal policies are seen to be almost linear for all three cost parameters. Thus instead of memorizing a whole switching curve it it suffices to store their value at $x_1 = 0$ and the slope. A simple calculation is then performed

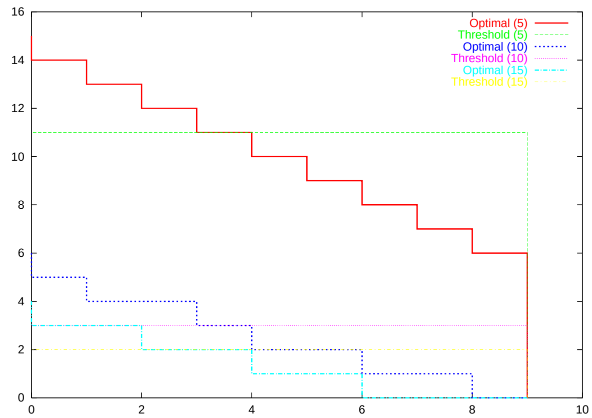

Figure 3: Optimal and Threshold Policies

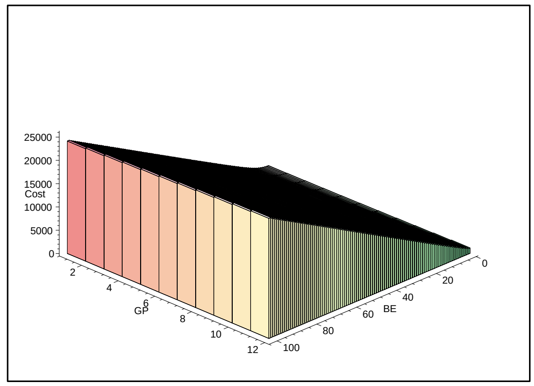

Figure 4: Cost of the optimal policy

upon arrival of a GP session in order to decide upon acceptance/rejection.

In Figure 4 we depict the cost corresponding to the optimal policy for $c = 5$ and rejection cost of 150. We see that the optimal cost is indeed convex in direction x_0 (as already mentioned, supermodularity plus the equation (5) imply convexity in direction x_0).

Figure 5 illustrates the benefit of use of the optimal policy rather than the greedy policy. The difference between the costs of the two policies is depicted. (The average threshold is 11, as seen from Fig. 3.) We see some discontinuity in this difference around the threshold value 11.

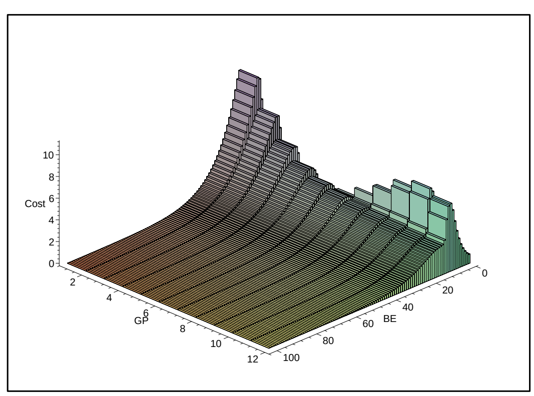

Figure 5: Average Threshold vs Optimal Policies

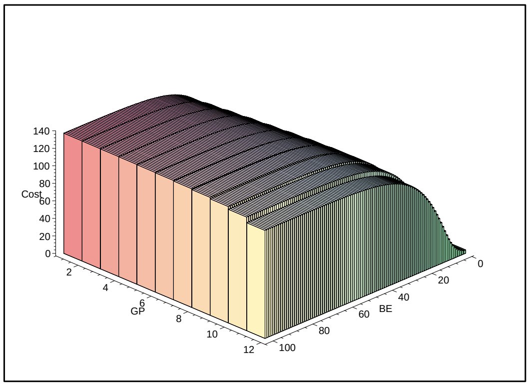

Figure 6: Greedy vs Optimal Policies

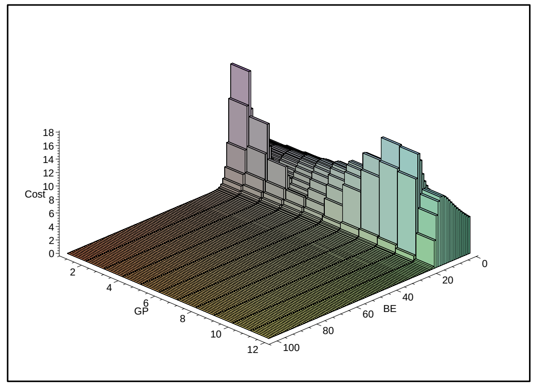

Figure 7: Average Threshold vs Optimal Policies

Figure 6 shows the difference in the cost between the optimal policy and the greedy policy. The difference between the performances corresponding to these two policies is much larger (more than ten times) than the difference in performance between the optimal and the average threshold. We conclude that the average threshold is a good approximation (in term of overall performance) for the optimal policy.

Note that the maximum threshold policy is closer to the greedy policy than the average threshold policy, in the sense that its acceptance region contains the one of the average threshold policies (and is itself contained in the acceptance region of the greedy policy). The difference in performances between the two threshold policies is very small, hence in practice one may choose either thresholds (both perform better than the greedy policy and only slightly worse than the optimal policy).

The overloaded regime Here we consider the case of input rates which are much larger than the available service rate ($\hat{\rho} >> 1$). In this regime the sum of input rates is considerably larger than the network capacity. It is of course an undesirable operational mode of the network that results in high congestion (which translates in large delays of BE sessions and large rejection rate of GP sessions). Yet we believe that even a network that is dimensioned so as to avoid this regime, it still may occur during some peak hours.
The parameters used in the example are: The total number of resources (B=10), the discount factor ($\alpha = 0.98$), the arrival rates ($\lambda_0 = 0.4$ and $\lambda_1 = 0.3$), the service rate or total capacity $\bar{\mu} = 0.2$, and $\nu = 1$. The cost values and the rejection cost vary according to the differents examples and are noted in each of them. We present in the last figure the difference in performance between average threshold and optimal policies.

5 Conclusions

We have studied in this paper the problem of admission control of GP (guaranteed performance) sessions, taking into account the impact on the performance of both GP as well as BE (best effort) sessions. Using dynamic programming techniques we were able to compute the optimal admission policy. We studied its structure, as well as the the properties of the optimal cost. We then performed a numerical study of several policies and compared their performance.

The numerical study shows that the optimal policies can be well approximated by threshold policies in terms of their performance. The performance of both threshold as well as the optimal policies are considerably better than the performance of the greedy policy which does not take into consideration the delay of BE sessions, and accepts GP sessions whenever there is available bandwidth for them.
The fact that threshold policies perform considerably better than the greedy one was already discovered in [9]. Our analysis allows to justify the use of the threshold policies, as they were shown to perform very well in comparison to the optimal policies.

References

[1] The ATM Forum Technical Committee, *af-tm-0056*, Traffic Management Specification, Version 4.0, Apr 1996.

[2] J. Carlström E. Nordström, "Call admission control and routing for integrated CBR/VBR and ABR services: a Markov decision approach," *ATM'99 workshop*, pp. 24–27, May 1996, Kochi, Japan.

[3] E. Altman and H. J. Kushner, "Admission control for combined guaranteed performance and best effort sessions under heavy traffic," in *IEEE CDC*, 1998.

[4] T. Jiménez E. Altman and G. Koole, "On optimal call admission control," in *IEEE CDC*, 1998.

[5] K. Ross and D. Tsang, "The stochastic knapsack problem," *IEEE Trans. on Communications*, **37**, 1989.

[6] E. Nordström, *Markov Decision Problems in ATM traffic control*, Ph.D. thesis, Dept. Computer Systems, Uppsala University, 1998.

[7] S.A. Lippman, "On dynamic programming with unbounded rewards," *Management Science*, **21**, pp. 1225–1233, 1975.

[8] G. Koole, "Structural results for the control of queueing systems using event-based dynamic programming," *Queueing Systems*, vol. 30, pp. 323–339, 1998.

[9] E. Altman, D. Artiges, and K. Traore, "On the integration of best-effort and guaranteed performance services," *European Trans. on Telecommunications*, Feb 1999.