

Call Admission Control for Preemptive and Partially Blocking Service Integration Schemes in ATM Networks

Ernst Nordström

Department of Computer Systems, Information Technology, Uppsala University,

Box 325, S-751 05 Uppsala, Sweden

Tel: +46 18 471 31 11, Fax: +46 18 55 02 25

Email: ernstn@docs.uu.se

This paper evaluates a Markov decision approach to single-link Call Admission Control for CBR/VBR and ABR/UBR services. Two different schemes that support integration of narrow-band ABR/UBR and wide-band CBR/VBR services are evaluated: the standard preemptive scheme and the modified partial blocking scheme. The structure of the Markov decision policy shows an "intelligent blocking" feature, which implements bandwidth reservation for wide-band calls. The numerical results show that the Markov decision method yields higher long-term reward than the complete sharing method when the ability to create sufficient capacity for wide-band calls through partial blocking/preemption is limited. The results also show that the modified partial blocking scheme, which allows total preemption, gives the highest average reward rate.

1. INTRODUCTION

Call Admission Control (CAC) in Asynchronous Transfer Mode (ATM) networks should support an efficient integration of the Variable Bit Rate (VBR), Constant Bit Rate (CBR), Available Bit Rate (ABR) and Unspecified Bit Rate (UBR) service classes. One of the main design issues is how to share the capacity between guaranteed services (CBR and VBR) and best effort services (ABR and UBR). The design must utilize that fact that best effort calls have the ability to reduce their bandwidth in case of congestion. Two methods that meet this constraint are the standard preemptive scheme and the standard partial blocking scheme.

In the standard preemptive scheme, best effort calls are preempted when guaranteed service calls arrive to a busy link. In this paper, the best effort calls that are chosen for preemption are

selected at random. When calls depart from the link such that sufficient free capacity becomes available, a preempted best effort call enters service again. The preemptive scheme was analyzed in [3] in the case when all calls enter a queue before service. It was found that the scheme is capable of improving the link utilization at the expense of fairness. The common FIFO policy was shown to maintain fairness at some expense of link utilization.

In the standard partial blocking scheme [1, 2], the best effort services adapt their bandwidth requirement to the available capacity such that the bandwidth - holding time product remains constant. Each best effort call can specify a minimal accepted service ratio, $r_{min} \in (0,1]$ (along with the bandwidth requirement, b) which is used in the call negotiation process. A best effort call is accepted only if the available bandwidth b_a fulfills the criteria: $r_{min}b \leq b_a \leq b$. Throughout the lifetime of a call, the instantaneous service rate $r(t)$, defined as $b_a(t)/b$, may fluctuate according to the current load and available capacity on the link. The standard partial blocking scheme was analyzed in [1, 2] where it was found that the scheme gives low blocking probability and efficient link utilization for best effort calls.

The standard preemptive and partial blocking scheme was evaluated in [9] using optimal call admission control policies derived from Markov decision theory [10]. The two methods were shown to yield high average reward rates for different mixes of narrow-band and wide-band traffic. Several alternative methods to the Markov decision approach have been proposed in the literature, e.g. class limitation, trunk reservation and dynamic trunk reservation. The comparison presented in [6] indicates that for many cases, the trunk reservation and dynamic trunk reservation policies can provide fair, bandwidth efficient solutions, having performance close to the optimal Markov decision policy.

This paper evaluates the efficiency of Markov decision based call admission control policies for the standard preemptive scheme and a modified version of the partial blocking scheme. The modified partial blocking scheme is controlled by two different minimal service ratios. The first ratio, $r_{min,dim} \in (0,1]$, controls the access of best effort calls and limits the number of accepted best effort calls. The second ratio, $r_{min,user} \in [0,1]$, controls the access of guaranteed service calls. Using two minimal service ratios it is possible to both limit the time spent in the system for best effort calls and to allow a zero instantaneous service ratio. Note that the preempt-

tion occurring with $r_{min,user}=0$ is fair since all best effort calls will have their bandwidth reduced to zero upon preemption, which is not the case with the standard preemptive scheme.

Markov decision theory provides a computationally efficient technique to find the optimal CAC policy in terms of long-term reward. The Markov decision policy maps states to admission decisions (actions), i.e. to accept or reject a new call. The Markov decision approach evaluates the long-term reward of each action in each state, and chooses the action which maximizes the reward. The evaluation is based on a Markov model of the decision task, which comprises the state transition probabilities and the expected reward delivered at each state transition. The decision task model is parameterized by the call arrival and departure rates, which are assumed to be measured on line.

The Markov decision technique has been applied to the link access control problem [7] and the network routing problem [4] assuming that blocked calls are lost. The technique has also been applied to link allocation [8] and routing problems [5] in the context of blockable narrow-band and queueable wide-band call traffic.

This paper is organized as follows. In the next section, the CAC problem is introduced. Section 3 presents a Markov decision model for the CAC task for the standard preemptive scheme and for the modified partial blocking scheme. Section 4 describes the policy iteration technique of Markov decision theory in which the value determination problem is handled by solving a sparse linear equation system. Section 5 presents the numerical results. Finally, section 6 concludes the paper.

2. THE CAC PROBLEM

In the CAC problem, a link with capacity C [units/s] is offered calls from K traffic classes of CBR¹ and ABR calls. Calls belonging to class $j \in J = \{1, 2, \dots, K\}$ have the same bandwidth requirements and similar arrival and holding time dynamics. For ease of presentation, we consider $K=2$ traffic classes throughout the rest of this paper. The two classes consists of a narrow-band ABR class and a wide-band CBR class, indexed by 1 and 2, respectively.

We assume that class- j calls with peak bandwidth requirement b_j arrive according to a Poisson process with average rate λ_j [s^{-1}], and that the CBR call holding time is exponentially distributed

1. VBR calls can be modelled the same way adopting the notion of effective bandwidth.

with average $1/\mu_2$ [s]. The ABR call holding time for the preemptive scheme and the partial blocking schemes is exponentially distributed with average $1/\mu_1$ in the case when the call experiences no preemption and no partial blocking, respectively. If the ABR calls are partially blocked, the call holding time can be calculated by techniques from Markov driven workload processes, see [2].

The task is to find a CAC policy π that maps *request states* $(j,x) \in J \times X$ to *admission actions* $a \in A$, $\pi: J \times X \rightarrow A$, such that the long-term reward is maximized. The set A contains the possible admission actions, $\{ACCEPT, REJECT\}$. The set X contains all feasible *system states*. For the preemptive scheme it is given by:

$$X_1 = \left\{ (n_1, n_2, p) : p = 0, n_j \geq 0, \sum_{j \in J} n_j b_j \leq C \right\} \cup \left\{ (n_1, n_2, p) : p \in \{1, 2, \dots, p_{max}\}, n_j \geq 0, \sum_{j \in J} n_j b_j = C \right\}, \quad (1)$$

where n_j is the number of class- j calls accepted on the link, and p is the number of preempted ABR calls, which can take on the values $p \in P = \{0, 1, \dots, p_{max}\}$. For later use, we also introduce the set of feasible *link states* for the preemptive scheme:

$$N = \left\{ (n_1, n_2) : n_j \geq 0, \sum_{j \in J} n_j b_j \leq C \right\}. \quad (2)$$

For the partial blocking scheme, the set of feasible system states to enter when admitting best effort calls is given by:

$$X_{2,dim} = \left((n_1, n_2) : n_j \geq 0, n_1 b_1 r_{min,dim} + n_2 b_2 \leq C \right) \quad (3)$$

The set of feasible system states to enter when admitting guaranteed service calls is given by:

$$X_{2,user} = \left\{ (n_1, n_2) : 0 \leq n_1 \leq \lfloor C / (b_1 r_{min,user}) \rfloor, 0 \leq n_2 \leq \lfloor C / b_2 \rfloor, n_1 b_1 r_{min,user} + n_2 b_2 \leq C \right\} \quad (4)$$

where $r_{min,dim} \in (0,1]$ is a minimal accepted service ratio used for dimensioning purposes, i.e. to control the number of ABR calls in the system, and $r_{min,user} \in [0,1]$ is the minimal service ratio acceptable for the user when admitting guaranteed service calls. Note that $r_{min,dim} \geq r_{min,user}$.

3. A MARKOV DECISION MODEL FOR CAC

This section presents a Markov decision model for CAC for the standard preemptive scheme and the modified partial blocking scheme. The Markov decision model specifies a Markov chain which is controlled by actions in each state. The actions result in state transitions and reward delivery to the system. The control objective is to find the actions that maximize the average reward accumulated over time. In the current application, the Markov chain evolves in continuous time, and we therefore face a semi-Markov decision problem (SMDP).

The SMDP state x corresponds to the system state in the previous section, i.e. $x=(n_1,n_2,p)$ for the preemptive scheme, and $x=(n_1,n_2)$ for the partial blocking scheme. The SMDP action a is represented by a vector $a=(a_1,a_2)$, corresponding to admission decisions for presumptive call requests. The action space for both the preemptive and the partial blocking scheme becomes:

$$A = \{(a_1,a_2) : a_j \in \{0,1\}, j \in J\}. \quad (5)$$

where $a_j=0$ denotes call rejection and $a_j=1$ denotes call acceptance. The permissible action space in state x is a state-dependent subset of A . For the preemptive scheme, the permissible action space becomes:

$$A_1(x) = \{(a_1,a_2) \in A : a_1 = 0 \text{ if } n + \delta_1 \notin N, \\ a_2 = 0 \text{ if } n + \delta_2 - \Delta(n_1,n_2)\delta_1 \notin N \text{ or } p + \Delta(n_1,n_2) \notin P\} \quad (6)$$

where $n=(n_1,n_2)$, δ_j denotes a vector with zeros except for a one at position j , and

$$\Delta(n_1,n_2) = \theta \left[\frac{1}{b_1} \left[\sum_{j \in J} n_j b_j - C + b_2 \right] \right] \quad (7)$$

where $\theta(s)=0$ if $s \leq 0$ and $\theta(s)=\lfloor s \rfloor$ if $s > 0$. The quantity $\Delta(n_1, n_2)$ denotes the number of ABR calls that should be preempted in link state (n_1, n_2) in order to reserve capacity for a new CBR call. For the partial blocking scheme, the permissible action space becomes:

$$A_2(x) = \{(a_1, a_2) \in A: a_1 = 0 \text{ if } n + \delta_1 \notin X_{2,dim}, a_2 = 0 \text{ if } n + \delta_2 \notin X_{2,user}\} \quad (8)$$

The Markov chain is characterized by state transition probabilities $p_{xy}(a)$ which expresses the probability that the next state is y , given that action a is taken in state x . For the preemptive scheme, the state transition probabilities for $j \in J$ become:

$$p_{xy}(a) = \begin{cases} \lambda_j a_j \tau(x, a), & n_y = n_x + \delta_j \in N \\ & p_y = p_x = 0, \\ \lambda_2 a_2 \tau(x, a), & n_y = n_x + \delta_2 - \Delta(n_1, n_2) \delta_1 \in N, n_x + \delta_2 \notin N, \\ & p_y = p_x + \Delta(n_1, n_2) \in P, \\ n_{xj} \mu_j \tau(x, a), & n_y = n_x - \delta_j + \min(b_j/b_1, p_x) \delta_1 \in N, \\ & p_y = \max(p_x - b_j/b_1, 0) \in P, \\ n_{xj} \mu_j \tau(x, a), & n_y = n_x - \delta_j \in N, \\ & p_y = p_x = 0, \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

where the quantity $\tau(x, a)$ denotes the average sojourn time in state x :

$$\tau(x, a) = \left[\sum_{j \in J} n_{xj} \mu_j + a_j \lambda_j \right]^{-1} \quad (10)$$

The first term in the state transition probability expression above gives the state transition probability for a CBR or ABR call arrival to a link with some free capacity without any preempt-

tion of ABR calls. The second term gives the state transition probability for a CBR call arrival to a link with sufficient free capacity after preemption of ABR calls. The third term gives the state transition probability for CBR or ABR call departures when the preemption queue is non-empty. The fourth term gives the state transition probability for CBR or ABR call departures when the preemption queue is empty.

For the partial blocking scheme, the state transition probabilities become:

$$p_{xy}(a) = \begin{cases} \lambda_1 a_1 \tau(x, a), & n_y = n_x + \delta_1 \in X_{2,dim} \\ \lambda_2 a_2 \tau(x, a), & n_y = n_x + \delta_2 \in X_{2,user} \\ n_{x1} \mu_1 r(x) \tau(x, a), & n_y = n_x - \delta_1 \in X_{2,user} \\ n_{x2} \mu_2 \tau(x, a), & n_y = n_x - \delta_2 \in X_{2,user} \\ 0, & \text{otherwise} \end{cases} \quad (11)$$

where $r(x)$ denotes the instantaneous service ratio in state x :

$$r(x) = \begin{cases} 1, & \sum_{j \in J} n_j b_j \leq C, \\ [C - n_{x2} b_2] / (n_{x1} b_1), & \sum_{j \in J} n_j b_j > C, \end{cases} \quad (12)$$

The average sojourn time in state x is given by:

$$\tau(x, a) = \left[n_{x1} \mu_1 r(x) + n_{x2} \mu_2 + \sum_{j \in J} a_j \lambda_j \right]^{-1} \quad (13)$$

The expected accumulated reward in state x is given by $R(x, a) = q(x) \tau(x, a)$. For the preemptive scheme the reward accumulation rate is given by $q(x) = \sum_{j \in J} r_j n_{xj} \mu_j$. For the partial blocking scheme the reward accumulation rate is given by $q(x) = r_1 n_{x1} \mu_1 r(x) + r_2 n_{x2} \mu_2$. The quantity r_j , which specifies the reward for carrying a type- j call, can be written $r_j = r_j' b_j / \mu_j$, where r_j' denotes the normalized reward parameter. In this paper, we let the normalized reward parameter depend on the pricing model used for call charging.

4. MARKOV DECISION COMPUTATIONS

This section describes a method for solving the CAC task, formulated as a SMDP. The method of choice is *policy iteration*, which is one of the computational techniques within Markov decision theory to determine an optimal policy.

The admission to the link is controlled by the so-called gain function, $g_j(x, \pi)$. This function simply measures the increase in long-term reward due to acceptance of a class j call in state x under policy π . Calls are accepted if the gain function is positive and rejected otherwise. The gain function can be expressed in terms of the relative value function, $v(x, \pi)$, as $g_j(x, \pi) = v(x + \delta_j, \pi) - v(x, \pi)$. The difference $v(x, \pi) - v(y, \pi)$ can be interpreted as the expected difference in accumulated reward over an infinite interval starting in state x instead of in state y under policy π . The relative value function is computed by the policy iteration algorithm.

The policy iteration algorithm computes a series of improved policies in an iterative manner. The computation of an improved policy π_{k+1} from the current policy π_k involves three steps:

- task identification
- value determination
- policy improvement

The first step involves determining the Markov decision model, i.e. the state transition probabilities and the expected rewards. These quantities are parameterized by link call arrival rates λ_j and call departure rates μ_j , see section 3. The arrival/departure rates are obtained from measurements to make the Markov decision model adaptive to actual traffic characteristics. The measurement period corresponds to the policy improvement period. The measurement period should be of sufficient duration for the system to attain statistical equilibrium.

The second step involves computing the relative value function for the current policy. The value determination step consists of solving the set of linear equations:

$$\begin{aligned} v(x, \pi) &= R(x, a) - g(\pi)\tau(x, a) + \sum_{y \in X} p_{xy}(a)v(y, \pi) \quad ; x \in X \\ v(x_r, \pi) &= 0 \end{aligned} \tag{14}$$

where x_r is an arbitrary chosen reference state (e.g. the empty state) and $g(\pi)$ denotes the average reward rate. The solution involving all the $v(x, \pi)$ and $g(\pi)$ can be obtained by any standard routine for sparse linear systems.

The third step is the actual policy improvement. This step consists of finding the action that maximizes the relative value in each state:

$$\max_{a \in A(x)} \left\{ R(x, a) - g(\pi)\tau(x, a) + \sum_{y \in X} p_{xy}(a)v(y, \pi) \right\} \quad ; x \in X \quad (15)$$

Policy iteration can be proved to converge to an optimal policy in a finite number of iterations in the case of finite state and action space [10].

The proposed method can be summarized as follows. Choose an initial admission policy π and a relative value function $v(x, \pi)$. During a finite period, allocate calls according to the gain function associated with the chosen relative value function. At the same time, measure traffic statistics (call arrival rates and call departure rates) in order to determine the Markov decision task for the current policy. Evaluate the applied policy in the context of the current Markov decision task, by solving a sparse linear equation system, and improve the policy. Apply the new policy during the next period, measure the traffic statistics and repeat the policy evaluation and the policy improvement step and so forth.

5. NUMERICAL RESULTS

This section evaluates the performance of two CAC methods for the preemptive scheme and the partial blocking scheme: the Markov decision (MD) method and the complete sharing (CS) method. Performance measures of interest are the average reward rate and the average time an ABR call spends in the system (the call holding time). For the preemptive scheme, the preemption probability is also evaluated.

The results are based on simulations for a single link with capacity $C=48$ [units/s], which is offered different mixes of ABR (class 1) and CBR (class 2) traffic. The bandwidth requirements are $b_1=1, b_2=6$ [units/s], and the mean call holding times $1/\mu_1=1/\mu_2=1$ [s], assuming that the ABR calls experiences no preemption and no partial blocking.

The arrival rates λ_1 and λ_2 were varied so that the average offered traffic equalled the link capacity:

$$\frac{b_1\lambda_1}{C\mu_1} + \frac{b_2\lambda_2}{C\mu_2} = 1.0 \quad (16)$$

A step size of 0.2 in the arrival rate ratio λ_1/λ_2 has been used when plotting all the figures. Moreover, the curves presented in the figures are obtained after averaging over 30 simulation runs and 95% confidence intervals, computed assuming normally distributed values, are also shown for each curve.

Figure 1 and 2 shows the average reward rate for the preemptive scheme for different arrival rate ratios, different maximal sizes of the preemption queue, and different normalized reward parameters for the ABR class.




Figure 1: Average reward rate for different arrival rate ratios for the preemptive scheme with $r_1'=0.05$. Case 1 has $p_{max}=24$ and case 2 has $p_{max}=6$.

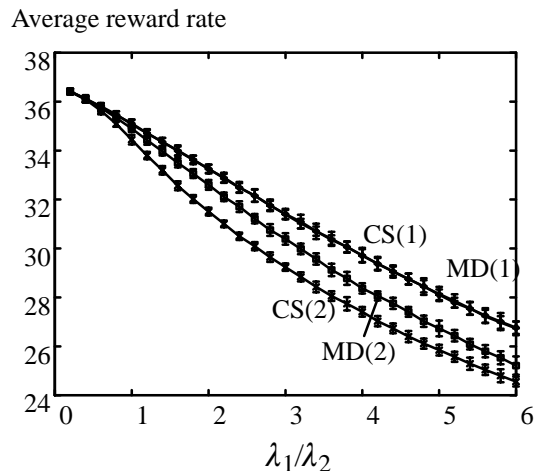


Figure 2: Average reward rate for different arrival rate ratios for the preemptive scheme with $r_1'=0.20$. Case 1 has $p_{max}=24$ and case 2 has $p_{max}=6$.

When the maximal size of the preemption queue is large ($p_{max}=24$), the average reward rate of the MD and CS method are similar. When the maximal queue size is small ($p_{max}=6$), the MD method gives a larger average reward rate compared to the CS method. The reason is that for small maximal queue sizes the MD method implements so called “intelligent blocking“ in individual states. By rejecting narrow-band call requests, typically when the free capacity equals the size of a wide-band call, bandwidth is reserved for the wide-band class, which increases the long-term reward.

Figure 3 and 4 shows the average time an ABR call spends in the system in the preemptive scheme with the MD method for different maximal sizes of the preemption queue. Three different curves are shown in each figure. The lower curve shows the average system time for calls that are not preempted. The middle curve shows the average system time taking all calls (preempted

and not preempted) into account. The upper curve shows the average system time for calls that are preempted. The lower curve is below 1 since short calls are more likely not to be preempted.

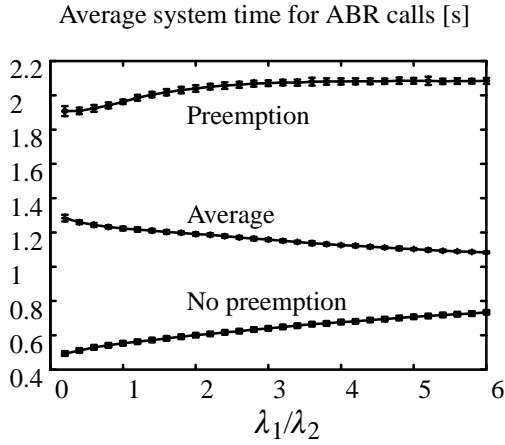


Figure 3: Average system time for ABR calls for different arrival rate ratios for the preemptive/MD scheme with $p_{max}=24$ and

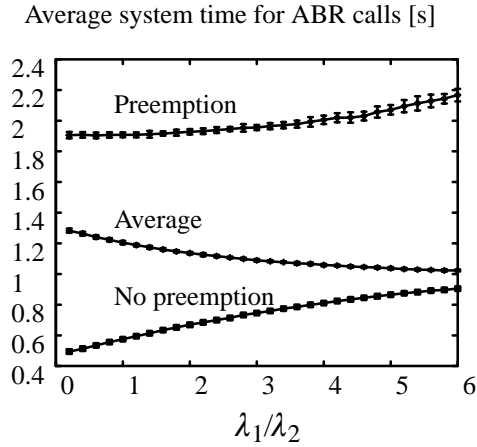


Figure 4: Average system time for ABR calls for different arrival rate ratios for the preemptive/MD scheme with $p_{max}=6$ and $r_1'=0.05$.




Figure 5: Preemption probability for ABR calls for different arrival rate ratios for the preemptive/MD scheme with $p_{max}=24$ and $r_1'=0.05$.

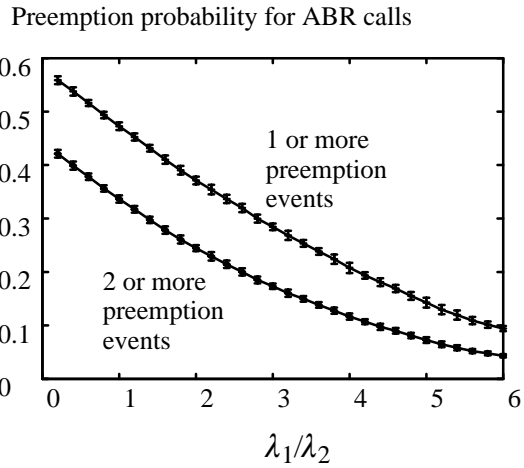


Figure 6: Preemption probability for ABR calls for different arrival rate ratios for the preemptive/MD scheme with $p_{max}=6$ and $r_1'=0.05$.

Figure 5 and 6 show the preemption probability for the preemptive scheme with the MD method for different arrival rate ratios and different maximal queue sizes. Two different curves are shown in each figure. The upper curve shows the probability of preemption occurring 1 or more times during the lifetime of an ABR call. The lower curve shows the probability of preemption occurring 2 or more times.

Figure 7 and 8 shows the average reward rate for the partial blocking scheme for different arrival rate ratios, different values of $r_{min,dim}$ and $r_{min,user}$ and different values of the normalized reward parameter for the ABR class. When $r_{min,user}=0$, i.e. when total preemption is allowed, there is no performance difference between the MD and CS method. When $r_{min,user}=0.5$, the intelligent blocking feature of the MD method results in a higher average reward rate. The narrow-band ABR class is blocked in all link states when the normalized reward parameter for the ABR class is low ($r_1'=0.05$). When the value is higher ($r_1'=0.20$) the MD policy blocks ABR calls in all link states when $\lambda_1/\lambda_2 < 2$, and in individual link states when $\lambda_1/\lambda_2 > 2$.

When the narrow-band ABR class is completely blocked, we face a severe fairness problem. However, the complete blocking can be avoided by increasing the normalized reward parameter r_1' for the ABR class. Of course, we can not expect the average reward rate to be as high as when the narrow-band ABR class is completely blocked since the blocking probability for the wide-band CBR class will increase. Nevertheless, changing the normalized reward parameters is a simple way to control the distribution of blocking probabilities among different call classes [4].

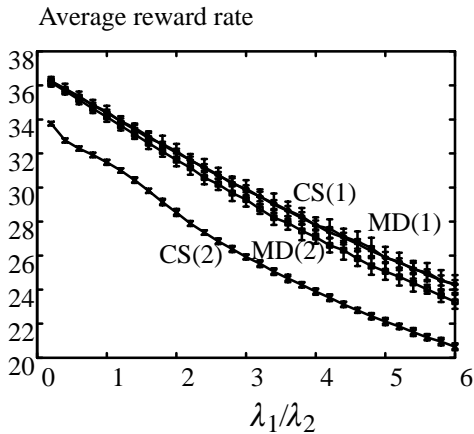


Figure 7: Average reward rate for different arrival rate ratios for the partial blocking scheme with price factor $r_1'=0.05$. Case 1 has $r_{min,dim}=0.5$ and $r_{min,user}=0$. Case 2 has $r_{min,dim}=0.5$ and $r_{min,user}=0.5$.

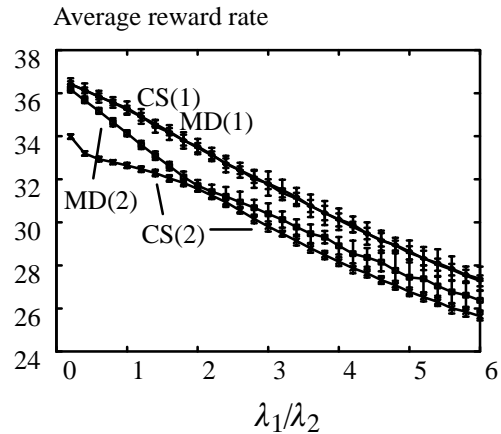


Figure 8: Average reward rate for different arrival rate ratios for the partial blocking scheme with price factor $r_1'=0.20$. Case 1 has $r_{min,dim}=0.5$ and $r_{min,user}=0$. Case 2 has $r_{min,dim}=0.5$ and $r_{min,user}=0.5$.

Figure 9 and 10 shows the average system time for ABR calls in the partial blocking scheme with the MD method for different arrival rate ratios and different values of $r_{min,dim}$ and $r_{min,user}$. No curve is shown for the case when ABR calls are blocked in each link state. In figure 9, the case 1 curve for the Markov decision method has larger confidence intervals than the case 1 curve

for the complete sharing method.

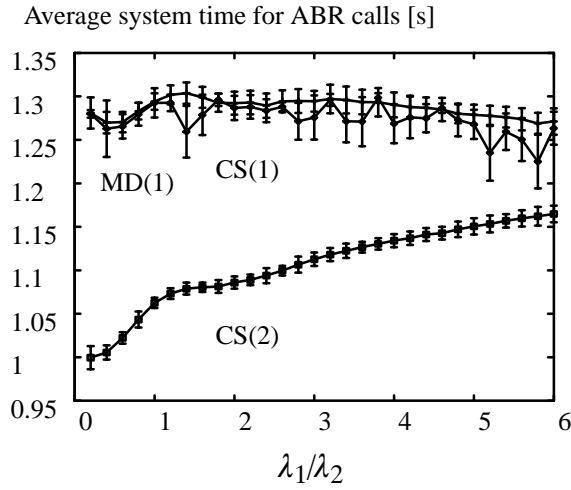


Figure 9: Average system time for ABR calls for different arrival rate ratios for the partial blocking scheme with $r_1'=0.05$.
Case 1 has $r_{min,dim}=0.5$ and $r_{min,user}=0$.
Case 2 has $r_{min,dim}=0.5$ and $r_{min,user}=0.5$.

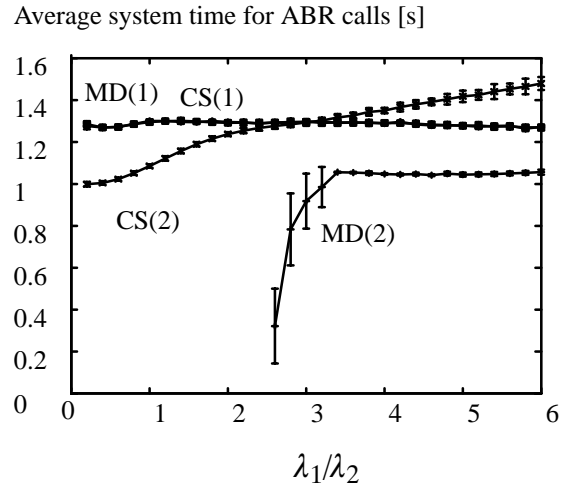


Figure 10: Average system time for ABR calls for different arrival rate ratios for the partial blocking scheme with $r_1'=0.20$.
Case 1 has $r_{min,dim}=0.5$ and $r_{min,user}=0$.
Case 2 has $r_{min,dim}=0.5$ and $r_{min,user}=0.5$.

For comparison, figure 11 and 12 shows the average reward rate for different realizations of the preemptive and the partial blocking scheme with CAC based on the MD method. The method with highest average reward rate is obviously partial blocking with $r_{min,user}=0$, i.e. when total preemption is allowed. Confidence intervals are not shown in order to improve the readability of the figures.




Figure 11: Average reward rate comparison between the preemptive/MD and partial blocking/MD scheme for different arrival rate ratios with $r_1'=0.05$.

PB case 1 has $r_{min,dim}=0.5$ and $r_{min,user}=0$.
 PB case 2 has $r_{min,dim}=0.5$ and $r_{min,user}=0.5$.
 PRE case 1 has $p_{max}=24$.
 PRE case 2 has $p_{max}=6$.

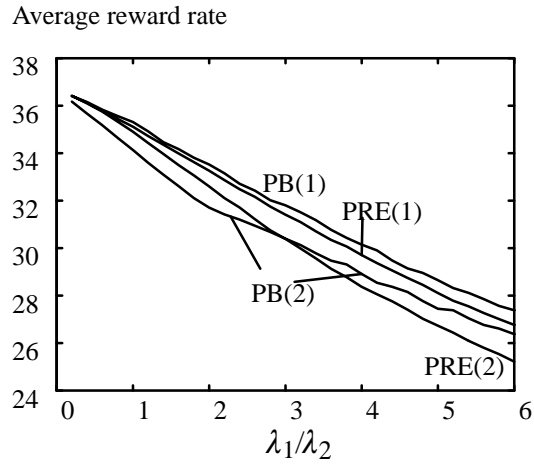


Figure 12: Average reward rate comparison between the preemptive/MD and partial blocking/MD scheme for different arrival rate ratios with $r_1'=0.20$.

PB case 1 has $r_{min,dim}=0.5$ and $r_{min,user}=0$.
 PB case 2 has $r_{min,dim}=0.5$ and $r_{min,user}=0.5$.
 PRE case 1 has $p_{max}=24$.
 PRE case 2 has $p_{max}=6$.

The results presented in the figures were obtained after 10 adaptation epoches with the policy iteration method. Each adaptation period contained 100 000 simulated call events. The performance values in the figures are based on measurements of 400 000 call events after policy convergence.

6. CONCLUSION

This paper has evaluated the efficiency of Call Admission Control (CAC) based on Markov decision theory for two schemes that supports integration of guaranteed services and best effort services: the standard preemptive scheme and the modified partial blocking scheme. The Markov decision technique can be used to compute CAC policies that are optimal in terms of long-term reward. The optimality is achieved by intelligent blocking of narrow-band ABR/UBR calls, either completely, or at link states where typically the free capacity equals the size of a wide-band call.

The presented numerical results show that the Markov decision method yields higher long-term reward than the complete sharing method when the ability to create sufficient capacity for wide-band CBR calls through partial blocking/preemption is limited. The results also show that

the modified partial blocking scheme, which allows total preemption ($r_{min,user}=0$), gives the highest average reward rate.

ACKNOWLEDGEMENTS

The author would like to thank Jakob Carlström, Søren Blaabjerg and Gábor Fodor for stimulating discussions. This work was financially supported by NUTEK, the Swedish National Board for Industrial and Technical Development.

REFERENCES

1. S. Blaabjerg and G. Fodor, A Generalization of the Multirate Circuit Switched Loss Model to Model ABR Services in ATM Networks, in Proc. of the IEEE International Conference on Communications, Singapore (1996).
2. S. Blaabjerg, G. Fodor and A. Andersen, A Partially Blocking-Queueing System with CBR/VBR and ABR/UBR Arrival Streams”, in Proc. of 5th International Conference on Telecommunications Systems, Nashville, USA (1997).
3. B. Kraimeche and M. Schwartz, Bandwidth Allocation Strategies in Wide-Band Integrated Networks, IEEE Journal on Selected areas in Commun., vol. SAC-4, no. 6, pp. 869-878 (1986).
4. Z. Dziong and L. Mason, Call Admission and Routing in Multi-Service Loss Networks, IEEE Trans. on Commun., vol. 42, no. 2, pp. 2011-2022 (1994).
5. Z. Dziong, K. Liao and L. Mason, Flow Control Models for Multi-Service Networks with Delayed Call Set Up, in proc. of INFOCOM'90, pp. 39-46, San Francisco, USA, (1990).
6. Z. Dziong and L. Mason, Fair-Efficient Call Admission Control Policies For Broadband Networks - A Game Theoretic Framework, *IEEE/ACM Trans. on Networking.*, vol. 4, no. 1, pp. 123-136, (1996).
7. K. Ross and D. Tsang, Optimal Circuit Access Policies in an ISDN Environment: A Markov Decision Approach, IEEE Trans. on Commun., vol. 37, no. 9, pp. 934-939, (1989).
8. E. Nordström, Near-Optimal Link Allocation of Blockable Narrow-Band and Queueable Wide-Band Call Traffic in ATM Networks, in Proc. of 15th International Congress on Telecommunications, ITC'15, Washington D.C., USA (1997).
9. E. Nordström, S. Blaabjerg and G. Fodor, Admission Control of CBR/VBR and ABR/UBR Call Arrival Streams: A Markov Decision Approach, in proc. of IEEE ATM'97 Workshop, Lisboa, Portugal (1997).
10. H. Tijms, Stochastic Modeling and Analysis - a Computational Approach, Wiley (1986).