Near-Optimal Link Allocation of Blockable Narrow-Band and Queueable Wide-Band Call Traffic in ATM Networks

Ernst Nordström

Department of Computer Systems, Uppsala University, Box 325, S–751 05 Uppsala, Sweden

This paper presents an adaptive scheme for a sub-function in ATM network routing, called link allocation, under blockable narrow-band and queueable wide-band call traffic assumptions. The scheme adapts the link allocation policy to the offered Poisson call traffic such that the long-term reward is maximized. It decomposes the link allocation task into a set of link admission control tasks, formulated as semi-Markov Decision Problems. The link admission control policies are adapted by the policy iteration algorithm. The long-term reward (throughput), fairness and average waiting time in the queue of the scheme are evaluated numerically.

1. INTRODUCTION

The link allocation problem arises in ATM network routing when a trunk group consists of several parallel physical links. The objective of the link allocation function is to select links for new calls such that the reward over time is maximized and the network availability is maintained. In order to maximize the reward over time, narrow-band calls should usually not be accepted on links which have a free capacity that equals the size of a wide-band call. The rejection of narrow-band calls reserves bandwidth for future wide-band calls on the link, which increase the reward over time. In general, the reward loss due to rejection of narrow-band calls should be traded off against the reward gain of accepting a future wide-band call.

Fairness is a typical network availability constraint that can be obtained by access control and/or call queueing. Partial sharing (class limitation) and trunk reservation are examples of link access control schemes that can provide fairness, in contrast to the complete sharing scheme which gives excessive blocking for wide-band calls [1]. Another alternative is queueing of wide-band calls when the link is busy. File transfer and medical imaging are examples of wide-band services that can tolerate moderate call set up delay. In [2], Serres and Mason derive performance measures for a single link system with blockable narrow-band and queueable wide-band call traffic. In their work, narrow-band calls enter the link whenever there is sufficient capacity, independently of the number of wide-band calls in the queue. There is also a cut-off parameter r_0 that specifies the maximum number of wide-band calls on the link. An arriving wide-band call is delayed in an infinite queue if r_0 wide-band calls are active on the link or if there is insufficient capacity to carry a wide-band call.

In this paper, we consider a multi-link system with blockable narrow-band calls and queueable wide-band calls. Each link has an associated finite queue for wide-band calls. The link allocation policy is implemented by a set of independent link admission control policies which are state-dependent. This means that a link can block a narrow-band call although there is sufficient capacity on the link. A further difference with [2] is that no cut-off parameter is used for wide-band calls.

In a previous paper, we have analyzed a link allocation scheme with blockable call traffic using the Markov decision approach [3].

An optimal link admission control policy can be formulated using Markov decision theory [4,6,7]. The Markov decision policy maps states to admission decisions (actions), i.e. accept or reject. At each state transition, a state-dependent reward is delivered to the system. The Markov decision approach evaluates the long-term reward obtained when starting in different states with different actions. This evaluation is used to find the optimal policy. The decision task model comprises the state transition probabilities and the expected reward delivered at each state transition.

Dziong et al. have applied the Markov decision approach to the network routing problem, also assuming blockable narrow-band and queueable wide-band call traffic [5]. Our approaches differs mainly in the definition of the state-dependent reward. Dziong et al. let the reward depend on the current queue size and on the number of active narrow-band calls. We let the reward depend on the number of active narrow-band calls *and* on the number of active wide-band calls, but not on the queue size. Hence, we consider only the problem of efficient resource utilization, and not the problem of restricting the queueing delay.

Our numerical results show that the fairness between the call classes depends on the maximal queue size. The results also show that the Markov decision method is more efficient in terms of long-term reward than the method proposed by Serres and Mason.

The paper is organized as follows. In the next section, the link allocation problem and its decomposition are introduced. Section 3 presents a Markov decision model for the link admission control tasks. Section 4 describes the policy iteration technique of Markov decision theory. Section 5 presents the numerical results. Section 6 concludes the paper and points out directions for future work.

2. THE LINK ALLOCATION PROBLEM AND ITS DECOMPOSITION

In the target link allocation problem, a group of M links with capacity C_i [units/s], $i \in I = \{1,...,M\}$, is offered calls from a narrow-band and a wide-band call class. Calls belonging to a class $j \in J = \{1,2\}$ have the same bandwidth requirements b_j [units/s], and similar arrival and holding time dynamics. Class index 1 and 2 corresponds to the narrow-band and wide-band call class, respectively. We assume that class-j calls arrive according to a Poisson process with rate λ_j [s^{-1}], and that the call holding time is exponentially distributed with mean $1/\mu_j$ [s]. In this work, the parameter b_j is given by the peak ATM cell transmission rate, since deterministic cell multiplexing is assumed. Moreover, we assume a uniform call charging policy, which means that the long-term reward is proportional to the cell throughput at the call level.

Blocked narrow-band calls are lost while blocked wide-band calls are either lost or delayed in a finite queue associated with one of the links. The length of each queue can take on the values $q_i \in Q_i = \{0,1,...,q_{max}\}$. An individual queue is served when there is sufficient free capacity on the link to accept a wide-band call.

The task is to find a link allocation policy π that maps *request states* $(j,x) \in J \times X$ to *allocation actions* $a \in A$, $\pi: J \times X \rightarrow A$, such that the long-term reward is maximized. The set A contains the possible allocation actions, $\{ACCEPT(1), ..., ACCEPT(M), REJECT\}$, where $ACCEPT(i)$ corresponds to accepting the call on link $i \in I$. The set X contains all feasible system states, and is given by the Cartesian product of sets of feasible link sub system states, X_i,

$$X_i = \{(n_{i1}, n_{i2}, q_i) : q_i = 0, n_{i1}, n_{i2} \geq 0, n_{i1}b_1 + n_{i2}b_2 \leq C_i\} \cup$$

$$\{(n_{i1}, n_{i2}, q_i) : q_i \in \{1, 2, ..., q_{max}\}, n_{i1}, n_{i2} \geq 0, C_i - b_2 < n_{i1}b_1 + n_{i2}b_2 \leq C_i\}, \qquad (1)$$

where n_{ij} is the number of class-j calls accepted on link i. For later use, we also introduce the set of feasible link states when the current queue size is zero:

$$N_i = \left\{(n_{i1}, n_{i2}) : n_{i1}, n_{i2} \geq 0,\ n_{i1}b_1 + n_{i2}b_2 \leq C_i\right\}, \tag{2}$$

The size of the system state space X increases rapidly with number of links in the group. To obtain a feasible computational complexity we therefore decompose the link allocation task into a set of independent link admission control tasks, and formulate these as semi-Markov decision problems.

3. A MARKOV DECISION MODEL FOR A SINGLE LINK ADMISSON CONTROL

This section presents a Markov decision model for a single link admission control task. 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 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). For ease of presentation, the link index sub script is dropped below. For example, we write n_x instead of n_{xi}.

The SMDP state x corresponds to a link sub system state in the previous section, i.e. $x=(n_{x1},n_{x2},q_x)\in X$. The SMDP action a is represented by a vector $a=(a_1, a_2)$, corresponding to admission decisions for presumptive call requests. Thus, the action space becomes

$$A = \{(a_1,a_2) : a_j \in \{0,1\},\ j=1, 2\}, \tag{3}$$

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:

$$A(x) = \{(a_1,a_2)\in A\text{: } a_1=0 \text{ if } n_{x1}b_1+n_{x2}b_2=C \text{ and } a_2=0 \text{ if } q=q_{max}\}. \tag{4}$$

The Markov chain is characterized by state transition probabilities between state pairs (x,y). The state transition probability from state x to state y under action $a=(a_1,a_2)$ in state x is:

$$p_{xy}(a) = \begin{cases} \lambda_{ij}a_j\tau(x,a), & n_y=n_x+\delta_j, \quad q_y=q_x, \quad n_x+\delta_j \in N, \\ \lambda_{i2}\tau(x,a), & n_y=n_x, \quad q_y=q_x+1\in Q, \quad n_x+\delta_2 \notin N, \\ n_{xj}\mu_j\tau(x,a), & n_y = n_x - \delta_j, \quad q_y = q_x, \quad n_x - \delta_j + \delta_2 \notin N, \\ n_{xj}\mu_j\tau(x,a), & n_y=n_x-\delta_j+\delta_2, \quad q_y=q_x-1\in Q, \quad n_x-\delta_j+\delta_2\in N, \\ 0 & \text{otherwise} \end{cases} \tag{5}$$

where n_x and n_y denotes the link states (n_{x1},n_{x2}) and (n_{y1},n_{y2}) respectively, δ_j denotes a vector with zeros except for a one at position j, and λ_{ij} denotes the arrival rate of class-j calls to link i, and is defined by a load sharing assumption [4]:

$$\lambda_{ij} = \lambda_j \frac{\bar{\lambda}_{ij}}{\sum_{k \in I} \bar{\lambda}_{kj}},\ i \in I,\ j \in J, \tag{6}$$

where $\bar{\lambda}_{ij}$ denotes the measured rate of accepted class-j calls on link i. The quantity $\tau(x,a)$ denotes the average time until the next call event in state x: $\tau(x,a)=\{\sum_{j \in J} [n_{xj}\mu_j + \lambda_{ij} a_j]\}^{-1}$.

The first term in the state transition probability expression above gives the state transition probability for a class-j call arrival to a link with some free capacity. The second term gives the state transition probability for a wide-band call arriving to a busy link with a non-full wide-band call queue. The third term gives the state transition probability for a class-j call departure which does *not* result in sufficient free capacity to accept a wide-band call from the queue. The fourth term gives the state transition probability for a class-j call departure which allows a wide-band call to be accepted from a non-empty wide-band call queue.

The expected accumulated reward in state x is given by $R(x,a)=q(x)\tau(x,a)$, where the reward accumulation rate is given by $q(x)=\Sigma_{j \in J} r_j n_{xj} \mu_j$. The quantity r_j is the absolute reward of carrying a type-j call. In order to maximize the overall call level throughput, r_j should be equal to the product of class-j's bandwidth requirement and its mean holding time, i.e. $r_j = b_j\, 1/\mu_j$.

In order to solve the value determination step of the Markov decision task, the continuous-time SMDP model must first be transformed into a discrete-time MDP model [7]:

$$\tilde{R}(x,a) = \frac{\tau}{\tau(x,a)} R(x,a) \qquad x \in X \text{ and } a \in A(x),$$

$$\tilde{p}_{xy}(a) = \begin{cases} \dfrac{\tau}{\tau(x,a)} p_{xy}(a), & y \neq x,\ x \in X \text{ and } a \in A(x), \\ \dfrac{\tau}{\tau(x,a)} p_{xy}(a) + [1 - \dfrac{\tau}{\tau(x,a)}], & y = x,\ x \in X \text{ and } a \in A(x), \end{cases} \tag{7}$$

where τ is the size of the discrete time step, chosen such that $0 < \tau \leq \min_{x,a} \tau(x,a)$. For example, $\tau = \{\sum_{j \in J} [N_{ij} \mu_j + \lambda_{ij}]\}^{-1}$, where $N_{ij} = C_i / b_j$ denotes the maximum number of class-j calls carried by link i.

4. ADAPTIVE POLICY ITERATION

This section describes a method for solving the link admission control tasks, formulated as semi-Markov decision problems. The method of choice is *policy iteration*, which is one of the computational techniques within Markov decision theory to determine an optimal policy. Another approach is *reinforcement learning*, which can be used to do model-free Markov decision optimization [3].

A fundamental quantity of Markov decision theory is the *evaluation function.* The evaluation function is defined for each state in the state space and measures the accumulated reward received during an infinite time interval, starting in the given state. The evaluation function is used as a tool to find the optimal policy.

The policy iteration approach 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_{ij} and call departure rates μ_j. The arrival/departure rates are obtained from measurements to make 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 evaluation function for the current policy. This is efficiently done by the method of successive approximations. This method relies on a basic equation of Markov decision theory. The equation states that the reward received within *n* decision epochs starting in given state *x*, should equal the expected immediate reward received after the first decision epoch, plus the expected accumulated reward within *n*–1 decision epochs starting from the neighbor states {*y*}:

$$V_n(x, \pi_k) = \tilde{R}(x, a) + \sum_{y \in X} \tilde{p}_{xy}(a)\, V_{n-1}(y, \pi_k) \qquad ; x \in X. \tag{8}$$

A full description of the value determination algorithm, called the method of successive approximations, can be found in the appendix. The method of successive approximations can be proved to converge to the correct evaluation function in a finite number of steps, provided that the state and action space are finite [7].

The third step is the actual policy improvement. The new action in each state (the new policy π_{k+1}) is determined by searching for the action that maximizes the sum of the immediate reward and the expected evaluation of the neighbor states:

$$max_{a \in A(x)} \{\tilde{R}(x, a) + \sum_{y \in X} \tilde{p}_{xy}(a)\, V(y, \pi_k)\} \qquad ; x \in X. \tag{9}$$

Where $V(y,\pi_k)$ denotes the evaluation function obtained from the value determination step. Since the search involves the evaluation function of the current policy π_k and not of the new policy π_{k+1} (which is unknown) we are not sure to find the optimal action. However, the method can be proved to converge to an optimal policy in a finite number of iterations in the case of finite state and action space [7].

During a policy improvement period, new calls are allocated according to two functions obtained from the Markov decision computations: the admission gain function and the queue gain function. The gain functions measure the increase in long-term reward induced by the control action. The admission gain function for a class-*j* call offered to a link is given by the difference $V(n_x+\delta_j,q_x,\pi_k)–V(n_x,q_x,\pi_k)$. The corresponding queue gain is defined as $V(n_x,q_x+1,\pi_k)–V(n_x,q_x,\pi_k)$. The admission gain function is computed for each link that has some free capacity. The queue gain function is computed for wide-band call arrivals which find the link busy. The link which has the maximal positive admission gain or queue gain is selected to carry or queue the new call. The call is rejected if the maximal admission gain or queue gain is negative.

The proposed method can be summarized as follows: Choose initial link admission control policies and evaluation functions for all the links. During a finite period, allocate calls according to the admission/queue gain functions associated with the chosen evaluation function. At the same time, measure traffic statistics (call arrival rates, link call acceptance rates and call departure rates) in order to determine the Markov decision task for the current policies. Evaluate the applied policies in the context of the current Markov decision tasks, using the method of successive approximations, and improve the policies for all the links. Apply the new policies 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 the proposed link allocation scheme. The performance measures of interest are the call level throughput, the per-class blocking probabilities, and the average waiting time in the queue.

The results are based on simulations for a link group with M=3 links with capacities C_i=C=24 [units/s], which is offered different mixes of narrow-band and wide-band call 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]. The curves presented in all figures are obtained after averaging over 30 simulation runs. 95% confidence intervals are also showed for each curve.

The first figure shows the call level throughput for different mixes of narrow-band and wide-band call traffic. The arrival rates λ_1 and λ_2 [s^{-1}] were varied so that:

$$\frac{b_1\lambda_1}{MC\mu} + \frac{b_2\lambda_2}{MC\mu} = 1.0 \qquad (10)$$

The maximal value of the arrival rate ratio λ_1/λ_2 is 6, which corresponds to a 50% load fraction for the narrow-band class. A step size of 0.2 in the arrival rate ratio has been used when plotting figures 1 and 2. As can be seen in the figure 1, the throughput is improved when increasing the maximal queue size. However, as the maximal queue size increases, the relative improvement is reduced.

The second figure shows the per class blocking probabilities for a maximal queue size of 1, for the same traffic mixes as in figure 1. When $\lambda_1/\lambda_2 = 0.2$, the narrow-band class is usually completely blocked. The unfairness is due to the fact that the Markov decision approach attempts to maximize the throughput. However, the fairness of the traffic classes can be controlled by varying the absolute reward r_j of carrying a class-j call [4]. When the relative reward of carrying a given class is increased, its blocking probability will decrease.

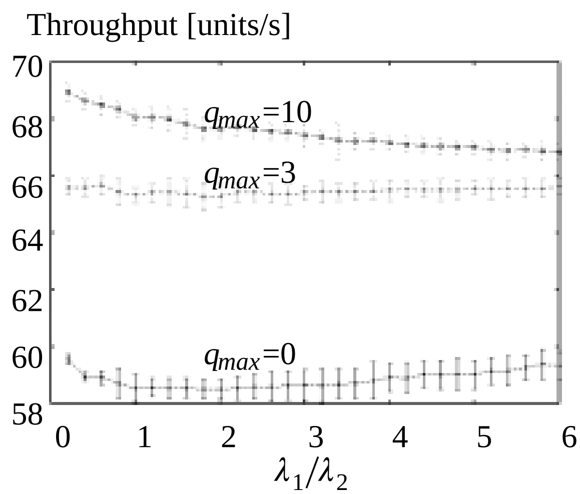

Figure 1. Call level throughput versus arrival rate ratio for different maximal queue sizes.

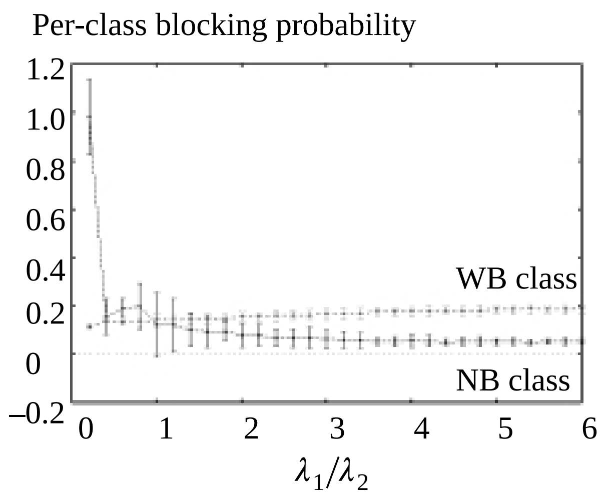

Figure 2. Per-class blocking probability versus arrival rate ratio for q_{max}=1.

A further investigation of the link admission control policies obtained for figure 2 shows that for arrival rate ratios in the interval $0.2 < \lambda_1/\lambda_2 \leq 1.0$, two links will usually block narrow-band calls completely. Similarly, in the interval $1.0 < \lambda_1/\lambda_2 \leq 2.0$, one link will usually block narrow-band calls completely and in the interval $2.0 < \lambda_1/\lambda_2 \leq 6$, all three links will typically carry narrow-band calls. Furthermore, so called "intelligent blocking" of the narrow-band class is usually performed at one or several links in the interval $0.2 < \lambda_1/\lambda_2 \leq 2.0$. The intelligent blocking typically occurs when the link has a free capacity that equals the size of a wide-band call. By rejecting the narrow-band call request, bandwidth is reserved for the wide-band class, which increases the long-term reward. However, if many narrow-band calls are accepted on the link, at least one of them is likely to depart before the next wide-band call arrival. Hence, narrow-band calls can be accepted, although the free capacity equals the size of a wide-band call.

An examination of the queue gain functions $V(n_x,q_x+1,\pi_k)-V(n_x,q_x,\pi_k)$ reveals that the queue gain values are higher when many of the active calls are from the wide-band class. This is due to the fact that when many wide-band calls are present on the link, there is a relatively low waiting time for service of the queue.

The third figure shows the per class blocking probabilities for different values of the maximal queue size, assuming the traffic mix $\lambda_1/\lambda_2=6$. The blocking probability for the wide-band class drops fairly quickly with the maximal queue size, and is approximately equal to the narrow-band blocking probability for a maximal queue size of 2. Obviously, the fairness between the traffic classes can be controlled by varying the maximal queue size.

The fourth figure shows the average waiting time in the queue for different values of the maximal queue size, again assuming the traffic mix $\lambda_1/\lambda_2=6$. The average waiting time is computed over all the queues in the link group. Apparently, the waiting time increases almost linearly with the maximal queue size. This figure is useful when determining the maximal allowed queue size for a given upper limit of the waiting time.

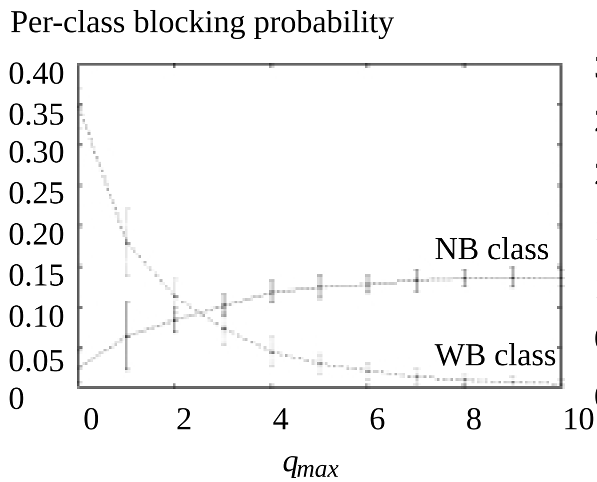

Figure 3. Per-class blocking probability for different maximal queue sizes.

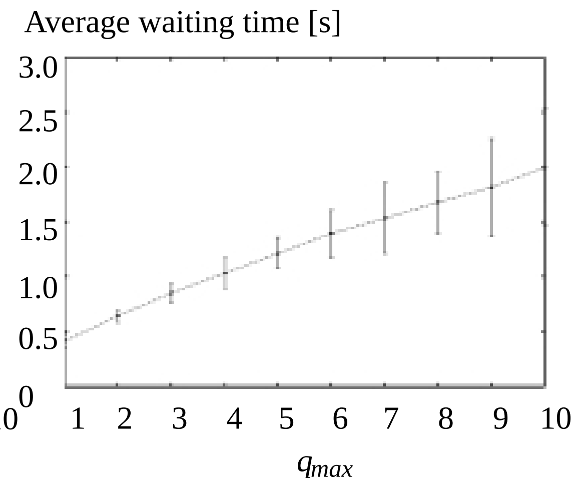

Figure 4. Average waiting time in queue for different maximal queue sizes.

The fifth figure compares the throughput in the single link case for the Markov decision method and the simple method proposed by Serres and Mason [2]. The latter method is modified so that it implements a finite queue. The reason is that the original method [2] requires the offered wide-band traffic load to be relatively low so that the queue size does not grow to infinity. In figure 5, the arrival rate ratio has been varied so that the utilization becomes 100% as in figure 1 and 2. As can be seen in the figure, the Markov decision method yields a higher throughput in the interval $0<\lambda_1/\lambda_2 \leq 4.0$, i.e. when intelligent blocking is relatively important. The cut-off parameter r_0 of the simple method for the maximal number of active wide-band calls were set to 3.

The sixth figure shows the per-class blocking probabilities for the Markov decision method and the simple method described above. The simple method has much smother blocking curves than the Markov decision method. The explanation is again the intelligent blocking feature of the Markov decision approach.

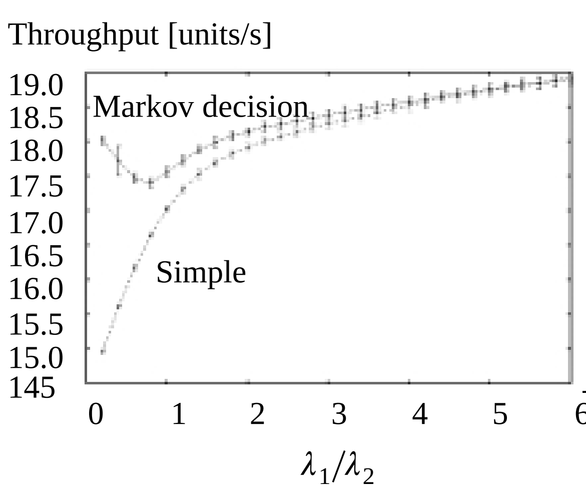

Figure 5: Call level throughput versus arrival rate ratio for q_{max}=1.

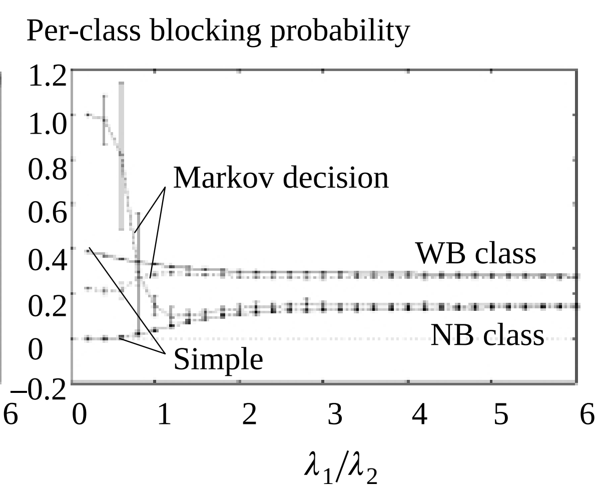

Figure 6. Per-class blocking probability versus arrival rate ratio for q_{max}=1.

The results presented in the figures where obtained after 5 adaptation epochs with the adaptive policy iteration method. Each adaptation period contained 1000 simulated call events. The parameters of the value determination algorithm were set as follows: τ=1/60, and $\varepsilon=0.005$. The performance values in the figures are based on measurements of 300 000 call events after policy convergence.

6. CONCLUSION

This paper has proposed a link allocation scheme for blockable narrow-band and queueable wide-band call traffic in ATM networks. The queueing increases not only the the fairness between the traffic classes, but also the long-term reward. The choice of maximal queue size depends on the fairness objective and the allowed average waiting time in the queue. A comparison with the simple scheme [2] in the single link case shows that the Markov decision method yields a higher long-term reward for certain traffic mixes.

The link allocation task is decomposed into a set of independent link admission control tasks, which are formulated as semi-Markov decision problems. Link call arrival rates and call departure rates, obtained from on-line measurements, parameterize the Markov decision model of a link admission control task. The policy iteration algorithm of Markov decision theory is used to find link admission control policies with optimal long-term reward. During the policy improvement period, new calls are allocated according to the admission gain function and the queue gain function, which are computed by the policy iteration algorithm.

In the future work, we will consider routing of blockable narrow-band and queueable wide-band call traffic. Then we also will consider the trade off between efficient resource utilization and queueing delay.

APPENDIX: Value Determination Algorithm

Step 1 (initialization). Choose initial values $V_0(x,\pi_k)$, $x \in X$. Let n :=1.

If $k=0$ then

$$0 \le V_0(x,\pi_k) \le max_{a \in A(x)} \tilde{R}(x,a),$$

else

Let $V_0(x,\pi_k)=V_n(x,\pi_{k-1})$.

Step 2 (recursion step). Compute the n-step value function:

$$V_n(x,\pi_k) = \tilde{R}(x,a) + \sum_{y \in X} \tilde{p}_{xy}(a)\, V_{n-1}(y,\pi_k) \qquad ;x \in X. \tag{A1}$$

Substitution for the MDP data transformation yields:

$$V_n(x,\pi_k) = R(x,a)\,\tau/\tau(x,a) + V_{n-1}(x,\pi_k) + \\ + \tau/\tau(x,a) \sum_{y \in X} p_{xy}(a)\,\{V_{n-1}(y,\pi_k) - V_{n-1}(x,\pi_k)\} \quad ;x \in X. \tag{A2}$$

Finally, inserting the expressions for the state transition probabilities we get:

$$V_n(n_x, q_x, \pi_k) = q(x)\tau + V_{n-1}(n_x, q_x, \pi_k)$$

$$+ \sum_{j \in J} \left(\lambda_{ij} a_j \tau \{ V_{n-1}(n_x + \delta_j, q_x, \pi_k) - V_{n-1}(n_x, q_x, \pi_k) \} \right)_{\{n_x + \delta_j \in N\}}$$

$$+ \left(\lambda_{i2} \tau \{ V_{n-1}(n_x, q_x + 1, \pi_k) - V_{n-1}(n_x, q_x, \pi_k) \} \right)_{\{n_x + \delta_2 \notin N,\ q_x + 1 \in Q\}}$$

$$+ \sum_{j \in J} \left(n_{xj_j} \mu_j \tau \{ V_{n-1}(n_x - \delta_j, q_x, \pi_k) - V_{n-1}(n_x, q_x, \pi_k) \} \right)_{\{n_x - \delta_j + \delta_2 \notin N\}}$$

$$+ \sum_{j \in J} \left(n_{xj_j} \mu_j \tau \{ V_{n-1}(n_x - \delta_j + \delta_2, q_x - 1, \pi_k) - V_{n-1}(n_x, q_x, \pi_k) \} \right)_{\{n_x - \delta_j + \delta_2 \in N,\ q_x - 1 \in Q\}}$$

$$x \in X \qquad \text{(A3)}$$

The sub script at the right parentheses in the formula above gives the condition for inclusion of the parenthesized term in the overall sum.

Step 3 (convergence test). Compute the bounds

$$m_n = min_{\, x \in X} \, \Delta V_n(x,\pi_k), \text{ and}$$

$$M_n = max_{\, x \in X} \, \Delta V_n(x,\pi_k),$$

where $\Delta V_n(x,\pi_k) = V_n(x,\pi_k) - V_{n-1}(x,\pi_k)$.

Stop the iteration if $0 \le M_n - m_n \le \varepsilon m_n$, where ε is a predetermined error.
Otherwise, let n:=n+1 and go to step 2.

ACKNOWLEDGEMENTS

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

REFERENCES

1. COST 242, Methods for the Performance Evaluation and Design of Broadband Multi-service Networks, Final report, part 3 (1996).
2. Y. Serres and L. Mason, A Multiserver Queue with Narrow- and Wide-Band Customers and Wide-Band Restricted Access, IEEE Trans. on Commun., vol. 36, no. 6 (1988).

3. E. Nordström and J. Carlström, A Reinforcement Learning Scheme for Adaptive Link Allocation in ATM Networks”, in proc. Int. Workshop on Appl. of Neural Networks to Telecommun., IWANNT’95, pp. 88-95, Stockholm, Sweden (1995).
4. Z. Dziong and L. Mason, Call Admission and Routing in Multi-Service Loss Networks, IEEE Trans. on Commun., vol. 42, no. 2 (1994).
5. Z. Dziong, K. Liao and L. Mason, Flow Control Methods in Multi-Service Networks with Delayed Call Set Up, in proc. of INFOCOM’90, pp. 39–46, San Francisco, USA, (1990).
6. 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 (1989).
7. H. Tijms, Stochastic Modeling and Analysis - a Computational Approach, Wiley (1986).