

PERFORMANCE EVALUATION OF THE CONFORMANCE DEFINITION FOR THE ABR SERVICE IN ATM NETWORKS

L. Cerdà¹, B. Van Houdt², O. Casals¹ and C. Blondia² *

¹*Polytechnic University of Catalonia* ²*University of Antwerp*
Computer Architecture Dept. *Dept. Math and Computer Science*
c/ Jordi Girona, 1-3, Modulo C6, *Universiteitsplein, 1*
E-08034 Barcelona, Spain *B-2610 Antwerp, Belgium*
{llorenc,olga}@ac.upc.es {vanhoudt,blondia}@uia.ua.ac.be

Abstract The standardization bodies have defined the Dynamic Generic Cell Rate Algorithm (DGCRA) as the Conformance Definition for ABR. This algorithm may be implemented in the Usage Parameter Control (UPC) for policing. In this paper we describe an equivalent queuing model of the DGRCA and we solve this model using a matrix analytic approach. The analytical results are validated by simulation.

1. INTRODUCTION

The ABR Service has been defined to efficiently multiplex sources that can adapt their transmission rate to the congestion state of the network. In this scheme, a source maintains the Allowed Cell Rate (ACR) parameter that defines the maximum rate at which cells may be scheduled for transmission. This parameter is controlled by special control cells called Resource Management Cells (RM-Cells). RM-Cells are transmitted embedded in the Data-Cell flow from the Source End System (SES) to the Destination End System (DES). The DES “turns around” the RM-Cells, which are sent back to the SES along the same path carrying congestion information. Depending on the congestion information received in the RM-Cell, the SES increases or decreases the ACR. A detailed description of the ABR Service can be found in ATM Forum, 1996.

*The first and third authors of this work were supported by the Ministry of Education of Spain under grant TIC96-2042-CE. The second and fourth authors were supported by the *Vlaams Actieprogramma Informatietechnologie* under project ITA/950214/INTEC.

The Conformance Definition for ABR is the Dynamic Generic Cell Rate Algorithm (DGCRA). It is based on the GCRA, which has been defined as the conformance definition for the PCR/SCR of CBR and VBR service categories. The decision of cell conformance in the DGCRA is made by measuring the inter-cell arrival time of a connection and checking whether it deviates from the inverse of the expected rate less than a tolerance called Cell Delay Variation Tolerance (CDVT). The CDVT is negotiated at the connection set up and is an upper bound of the unavoidable CDV introduced by the ATM layer functions and multiplexing stages up to the measuring point. A network operator may use a Usage Parameter Control (UPC) which considers as non-conforming the cells with a CDV higher than the CDVT. Non-conforming cells may be marked or discarded. Consequently, a correct dimensioning of the CDVT is needed to guarantee a low non-conforming cell probability at the UPC.

In order to compute the expected rate, the UPC keeps track of the rate changes conveyed by the RM-Cells. These rate changes are applied in the forward direction, after the round trip delay between the UPC and the source. Since this round trip delay is variable, the standard specifies that two time constants referred to as τ_2 and τ_3 have to be given with the CDVT at the connection set up. These are respectively an upper and a lower bound of the round trip delay between the UPC and the source. If these delay bounds are not set correctly, the UPC may not compute the expected rate properly, causing cell rejection.

The paper is organized as follows. We first give a definition of the DGCRA in Section 2. In Section 3 we describe an equivalent queuing model for the DGCRA. Section 4 presents a detailed description of the analytical model used in the evaluation. This model is solved using a matrix analytical approach as shown in Section 7. Section 8 gives numerical results obtained from the analytical model. The results are validated by simulation. Finally, in Section 9 some concluding remarks are formulated.

2. SPECIFICATION OF THE DGCRA

The DGCRA has been defined by the ATM Forum as the conformance definition for an ABR connection. At arrival instant of cell n , conformance is decided by measuring the CDV value $y_n = c_n - a_n$, where a_n is the arrival epoch and c_n is the theoretical arrival time. Cell n is non-conforming if y_n is greater than τ_1 and is conforming otherwise. The parameter τ_1 is the CDVT for the ABR connection. This notation is used in the standard ATM Forum, 1996 and we will indistinctly use CDVT and τ_1 in the rest of the paper. The theoretical arrival time is computed at each cell arrival with the following algorithm:

After the initializations $LVST_0 = a_0, I_0^{old} = I_0$, the set of theoretical arrival times $\{c_n\}_{n>0}$ is computed at the arrival epochs a_n as:

$$\begin{aligned} c_n &= LVST_{n-1} + \min(I_{n-1}^{old}, I_n) \\ LVST_n &= \begin{cases} \max(c_k, a_k) & \text{if } y_k \leq \tau_1 \quad (\text{cell conforming}) \\ LVST_{n-1} & \text{if } \tau_1 < y_n \quad (\text{cell non conforming}) \end{cases} \\ I_n^{old} &= \begin{cases} I_n & \text{if } y_n \leq \tau_1 \quad (\text{cell conforming}) \\ I_{n-1}^{old} & \text{if } \tau_1 < y_n \quad (\text{cell non conforming}) \end{cases} \end{aligned} \quad (1)$$

$LVST_{n-1}$ stands for the Last Virtual Scheduled Time at cell n arrival. The theoretical arrival time c_n is given by $LVST_{n-1}$ plus an increment $\min(I_{n-1}^{old}, I_n)$ equal to the inverse of the expected emission time that should be used by the source between cell $n-1$ and n . I_n is the inverse of the last rate change to be received by the source before cell n emission. This is computed by the DGCRA based on the feedback conveyed by the backward RM-cell flow up to cell n arrival time a_n . The algorithm takes $\min(I_{n-1}^{old}, I_n)$ because the first cell received after a new increase I_n is scheduled may be received at this increase or at the previous increment I_{n-1}^{old} . By taking the minimum the algorithm stays on the safe side.

The computation of the sequence I_n is not an easy task because a change of rate conveyed by a backward RM-cell received at the measuring point at a given time may be applied to the forward cell flow after a delay equal to the round trip delay between the measuring point and the source. To cope with this problem two time constants τ_2 and τ_3 have been introduced which are respectively an upper bound and a lower bound of this round trip delay. Furthermore, two algorithms ‘‘A’’ and ‘‘B’’ have been defined to determine I_n (see ATM Forum, 1996 for details).

3. QUEUING MODEL

The GCRA can be modeled as a single server queue with a workload (waiting time) limited to the CDVT (see e.g. Roberts et al., 1996). This queuing model can be extended to the DGCRA conformance definition given by equations (1). In the model we make the following assumptions: (i) the same rate changes followed by the source are policed by the DGCRA, (ii) the DGCRA is able to properly schedule the increments I_n corresponding to these rate changes, (iii) the source schedules cell transmissions at a new rate at cell emission times. With these assumptions, when cell n arrives at the DGCRA the increment I_n is equal to the source emission interval between cell n and cell $n+1$. This will be our definition of I_n in the rest of the paper.

The behavior of the queuing system which we use to model the DGCRA is shown in the time diagram of Figure 1. In this queuing model the cells correspond to customers. The service time is the increment applied by the

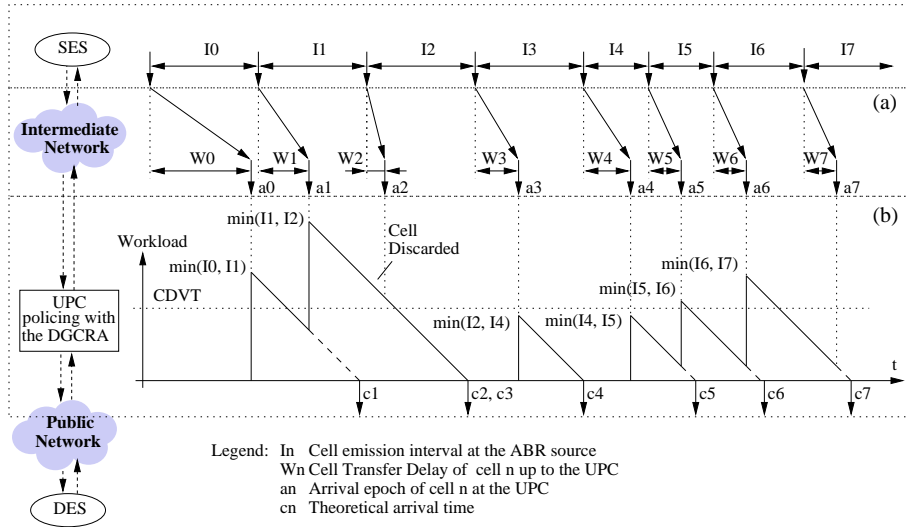


Figure 1 CDV introduced by an intermediate network (a) and the workload at the equivalent queuing model of the DGCRA (b).

DGCRA and the theoretical arrival time of cell n is the departure time of the previous accepted cell (c_n in the figure). The CDV value y_n of Equations (1) is given by the workload of the queue when $y_n > 0$. Therefore, when a cell arrival finds a workload higher than the CDVT, a non-conforming condition is given.

Note that the increment added at the cell n arrival epoch a_n in (1), is the service time to be added at the cell arrival epoch a_{n-1} in the equivalent queuing model. In other words, the workload increment added at the arrival epoch a_n in the queuing model is the increment that would be added by the DGCRA at the cell arrival epoch a_{n+1} , and is therefore given by $\min(I_n^{old}, I_{n+1})$ if cell $n+1$ is accepted, and 0 otherwise. For sake of simplicity, in Figure 1 we have approximated this by adding $\min(I_n^{old}, I_{n+1})$ to the workload if the cell n is accepted and 0 if cell n is not accepted.

4. MODEL USED IN THE EVALUATION

In order to evaluate the DGCRA we consider an ABR source multiplexed with a VBR source into a single switch. The VBR source has full priority over the ABR source. This multiplexing stage models the jitter introduced in the intermediate network shown in Figure 1.

The assumptions made in the equivalent queue of the DGCRA described in Section 3 apply to our model. Note that these assumptions imply that the

ABR source generates as much traffic as possible without violating the allowed cell rate (ACR). We also assume that the ACR is not fixed by the intermediate network but by the switches located after the UPC (such that the UPC can keep track of these rate changes). This is a foreseeable situation, at least as long as the VBR and the ABR sources do not cause a heavy congestion condition at the switch in the intermediate network. As a likely consequence of these assumptions, we consider the VBR and the ABR rate changes to be independent.

Finally, we assume a propagation delay equal to zero between the UPC and the ABR source. This is a plausible assumption since it is a fixed delay which should have no influence on the policing function. We also consider that a backward RM cell experiences no delay between the UPC and the ABR source. With these assumptions the only component of the round trip delay between the ABR source and the UPC considered in our model is the delay introduced by the switch in the intermediate network.

In the following we give a detailed description of the analytical model considered for each device.

4.1 THE ABR SOURCE BEHAVIOR

The ABR traffic that is multiplexed together with the VBR background traffic is described here. A set of N cell rates r_s , $s = 1, \dots, N$ is associated with the ABR source. The ABR source always transmits at one of these rates for which we have: $MCR \leq r_1 \leq \dots \leq r_{N-1} \leq r_N \leq PCR$

The parameters MCR and PCR represent the minimum and peak cell rate of the ABR connection. When the ABR source is transmitting at rate r_i , $i = 1, \dots, N$, a cell is forwarded respectively every $1/r_i$, $i = 1, \dots, N$, slots. In our analytical model we need $1/r_i$, $i = 1, \dots, N$ to be an integer value.

The source behavior can be modeled by associating $1/r_i$ states to each possible rate r_i . We denote these states by (i, j) , $i = 1, \dots, N$; $j = 1, \dots, 1/r_i$. We say that the ABR source is in state (i, j) , $j = 1, \dots, 1/r_i$, when it is transmitting at rate r_i . As long as the cell rate remains the same, e.g. r_i , the states $(i, 1)$ to $(i, 1/r_i)$ are traversed periodically. i.e. at each slot a transition occurs from (i, j) to $(i, j+1)$, $j = 1, \dots, 1/r_i - 1$ (note that in case of $1/r_i = 1$ there is only one of these states, the state $(i, 1)$, and the ABR source remains in it until a transition to a different rate state occurs). A cell is only transmitted in state $(i, 1/r_i)$ and no cell is generated in states $(i, 1)$ to $(i, 1/r_i - 1)$.

The ABR source considered here schedules cell transmissions at a new rate at cell emission times (the $(i, 1/r_i)$, $i = 1, \dots, N$ states in our model). Therefore, when a rate change occurs, e.g. from r_i to $r_{i'}$, there is a transition from state $(i, 1/r_i)$ to state $(i', 1)$. We denote by $P_{ABR}(i, i')$, $i, i' = 1, \dots, N$ the probability that such state transition occurs and assume that they are Markovian.

4.2 THE BACKGROUND TRAFFIC

We consider as background traffic a VBR source which will be multiplexed together with the ABR source at the switch. This VBR source is modeled by a Markov Chain with M states, each with an associated rate. We say that the source is in state k when it is transmitting at rate v_k , $k = 1, \dots, M$. These rates obey the following relation: $v_1 < v_2 < \dots < v_M$

A cell is generated by the VBR source with probability v_k while being in state k , $k = 1, \dots, M$. The VBR source can change its state at the end of each slot. We define $P_{\text{VBR}}(k, k')$ as the probability that a transition occurs from state k to state k' , ($k, k' = 1, \dots, M$). Notice that these transitions are independent of the ABR rate changes.

4.3 THE SWITCH

The switch which precedes the UPC device is used to model the jitter caused by the intermediate network on the ABR traffic. The input of this switch consists of an ABR and a VBR traffic stream generated by the traffic sources described above. As the VBR traffic has full priority over the ABR source and the VBR source never generates more than one cell in a slot, the switch only needs a buffer to store delayed ABR cells. Delayed ABR cells are forwarded by the switch towards the UPC device when there is no VBR cell arrival. If a VBR cell arrives, this cell is forwarded and the ABR cells have to wait.

4.4 THE UPC DEVICE

In our model the current state of the UPC device is characterized by its workload U as described in Section 3. Recall from Section 3 that the workload U is increased by $\min(I_n^{\text{old}}, I_{n+1})$ upon the arrival of cell n , if the cell is accepted, and is decremented by one at the end of each slot. For sake of simplicity we have taken I_n instead of $\min(I_n^{\text{old}}, I_{n+1})$ in the analytical model. The validity of this approximation is checked by simulation.

As described in Section 2, the expected rate at the interface is computed taking into account the upper and lower bounds of the round trip delay between the UPC and the source, τ_2 and τ_3 respectively. In order to see the influence of these delay bounds we have considered two scenarios. In the first one we model a UPC which immediately applies a rate change when scheduled. This is equivalent to setting $\tau_2 = \tau_3$. Clearly, if there are delayed cells at the switch buffer when a rate decrease occurs, these will be likely considered as non conforming (because an increment higher than the emission interval used by the source will be applied at the UPC).

In the second scenario we consider a UPC with τ_2 properly set to an upper bound of the round trip delay between the UPC and the source. Such a UPC

guarantees that no higher increment than the emission interval used by a “well behaving” source is applied at the UPC. We refer as “well behaving” a source that follows the rate changes conveyed by the backward RM-Cells. Note that this is the kind of source we use in our model. In the following the analytical model we use for these two scenarios is described.

5. SCENARIO WITH $\tau_2 = \tau_3$

In this case we consider a UPC which does not apply a time tolerance to the scheduled rate changes. Since in our model we consider a propagation delay equal to zero, we have $\tau_2 = \tau_3 = 0$. In our model this is equivalent to using the inverse of the rate associated with the ABR state at the time that a cell arrives at the UPC. Clearly this is not necessarily the rate of the ABR source when this cell was generated, as the rate of the ABR source might have changed if the cell was delayed in the switch.

We recall that as long as the ABR cell rate remains the same, e.g. r_i , it traverses the states $(i, 1), \dots, (i, 1/r_i)$ emitting one cell and possibly changing the rate in state $(i, 1/r_i)$. Therefore, we use $1/r_i$, $i = 1, \dots, N$ as the increment I_n associated with the states (i, j) , $j = 1, \dots, 1/r_i - 1$. The increment associated with the state $(i, 1/r_i)$ depends on whether or not a rate change occurs. If no rate change occurs the increment $1/r_i$ is used, otherwise we use $1/r_{i'}$ where $r_{i'}$ represents the new rate.

Note that we are making the following approximation. In a real situation with $\tau_2 = \tau_3 = 0$ the UPC would apply I_n immediately after the backward RM-Cell conveying the new rate traversed the UPC. In our model I_n is applied when the ABR source effectively performs the rate change, which happens at the cell emission epoch. This approximation is confronted with simulation results.

6. SCENARIO WITH τ_2 PROPERLY TUNED

In this case the UPC postpones the scheduled rate decreases until after a delay bound $\tau_2 > \tau_3$. Since in our model we consider a propagation delay equal to zero, this implies that during the first τ_2 slots after the scheduling of a rate reduction (from r_i to $r_{i'}$), the UPC will continue to use the smaller increment $1/r_i$.

We approximate this scenario by flushing the switch buffer each time that a rate reduction occurs. To assess the increments at the UPC we use the same rules as in the previous scenario. Note that by doing this we guarantee that no higher increment than the emission interval used by the source will be applied to a cell arriving at the UPC.

To be able to solve the analytical model, the flushing is only performed with probability $1 - \alpha$, where α is small, e.g. $\alpha < 10^{-12}$ (see Appendix A.2).

7. PERFORMANCE ANALYSIS

The system is observed at the end of each time slot. A Markov Chain is obtained by looking at the stochastic vector:

$$(Q, U, (i, j), k) \quad (2)$$

Where Q represents the queue length of the buffer inside the switch, U equals the remaining workload at the UPC, (i, j) , $i = 1, \dots, N$; $j = 1, \dots, 1/r_i$ is the state of the ABR source and k , $k = 1, \dots, M$ the state of the VBR source.

Denote by $P(S, S')$ the one slot transition probability from state $S = (Q, U, (i, j), k)$ to state $S' = (Q', U', (i', j'), k')$. By ordering the states $(Q, U, (i, j), k)$ lexicographically the probabilities $P(S, S')$ define a stochastic transition probability matrix \mathbf{P} with the block structure $\mathbf{P} = (\mathbf{Q}_{m,n})$. The sub-matrices $\mathbf{Q}_{m,n}$ govern the state $(U, (i, j), k)$ transitions when a queue length change from m to n occurs. Therefore, $\mathbf{Q}_{m,n}$ are square matrices of order equal to $U_{max} + 1$ times the number of ABR states times the number of VBR states, where U_{max} is the maximum workload. Notice that $U_{max} = \tau_1 + 1/r_{min}$, where r_{min} is the minimum of the cell rates considered for the ABR source. In the Appendix we describe how to derive the matrix \mathbf{P} and how to find the stationary probabilities in each scenario.

Having calculated the stationary probability vector of the process (2), we denote its components as $\pi(Q, U, (i, j), k)$, i.e. the probability that we are in state $(Q, U, (i, j), k)$. The rejection probability can then be found as:

$$\frac{\sum_{Q \geq 1} \sum_{\substack{U > \tau_1 + 1 \\ i, j, k}} \pi(Q, U, (i, j), k) (1 - v_k) + \sum_{\substack{U > \tau_1 + 1 \\ i, k}} \pi(0, U, (i, \frac{1}{r_i}), k) (1 - v_k)}{\sum_{Q \geq 1} \sum_{U, i, j, k} \pi(Q, U, (i, j), k) (1 - v_k) + \sum_{U, i, k} \pi(0, U, (i, \frac{1}{r_i}), k) (1 - v_k)}$$

8. NUMERICAL RESULTS

In this section we show the numerical results obtained with the analytical model described in Section 7. In order to validate the analytical model, all the results shown in this section have been verified by simulation.

We have used two different rates for the ABR source. In order to assess $P_{ABR}(i, i')$, $i, i' = 1, 2$ we have assumed that a Markovian process governs the rate changes. This process alternatively changes between two states, namely E_i , $i = 1, 2$. In this model a change into a certain state E_i represents a backward RM-Cell arrival conveying a new rate equal to r_i . Therefore, when the ABR source is in state $(i, 1/r_i)$, $i = 1, 2$, a change to state $(i', 1)$ occurs if $E_{i'}$, $i' = 1, 2$ is the current state. We have taken the sojourn time in each state E_i , $i = 1, 2$, to be the same and equal to p slots. With these assumptions we

have:

$$P_{\text{ABR}}(i, i') = \sum_{n=0}^{1/r_i} \binom{1/r_i}{n} 1/p^n (1 - 1/p)^{1/r_i - n} \cdot 1[(i = i' \text{ and } n \text{ is even}) \text{ or } (i \neq i' \text{ and } n \text{ is odd})], \quad i, i' = 1, 2$$

where we use the indicator function $1[\text{condition}]$ equal to 1 if *condition* is true and 0 otherwise. Note that as long as the inverse of the ABR rates is small compared to p , the mean time between two consecutive rate changes of the ABR source is equal to p . Thus, we will refer to $1/p$ as the ABR rate change frequency.

For the VBR source only one state is possible, and thus one rate v_1 . Obviously $P_{\text{VBR}}(k, k') = 1$, $k, k' = 1$. Note that the switch load ρ in this model is approximately given by: $\rho = v_1 + (r_1 + r_2)/2$

In the following we first investigate the validity of the analytical results by comparing them with the simulation results. Then the figures are analyzed, in order to derive some engineering rules.

8.1 VALIDATION


In the simulation we have used the DGCRA given by Equations 1. Moreover, the approximations made in the analytical model to make it tractable have been removed.

Figures 2.A and 2.B correspond to the scenario where τ_2 is properly tuned. The figures show a good agreement between the analytical and simulation results. Only when there is a heavy load and the ABR rate change frequency is high, the model yields an underestimated rejection probability.

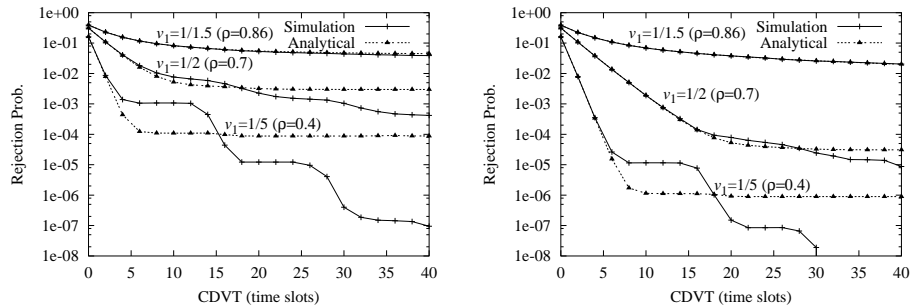
This can be explained by the following reasoning. If the queue length flushed when there is a rate reduction, e.g. from r_i to $r_{i'}$, is small compared to the cells emitted by the source while the rate was r_i , this approximation will clearly have a small influence on the rejection probability. Therefore, the approximation is worse for high loads and small sojourn times in the states with higher rates.

Figures 3.A and 3.B show the analytical and simulation results for the scenario where $\tau_2 = \tau_3$. In these figures we can see that the analytical model gives a good approximation for high loads, but it becomes worse when the load decreases. To explain these results we have to take into account the two main reasons that may lead to a cell rejection in this scenario: (i) the jitter of the ABR cell stream and (ii) the usage of an increment higher than the source cell emission interval.

The analytical model is able to capture very well the cells rejected due to condition (i). When the load is high, the condition (i) is predominant and thus the analytical model yields a good approximation. For lower loads, condition



2.A ABR rate change frequency = 1/200 2.B ABR rate change frequency = 1/20000
 Figure 2 Rejec. Prob. in the τ_2 properly tuned scenario. ABR rates: $r_1 = 1/3$, $r_2 = 1/15$.



3.A ABR rate change frequency = 1/200 3.B ABR rate change frequency = 1/20000
 Figure 3 Rejec. Prob. in the $\tau_2 = \tau_3$ scenario. ABR rates: $r_1 = 1/3$, $r_2 = 1/15$.




Figure 4 Rejec. Prob. when there are no ABR rate changes. ABR rate: $r_1 = 1/5$.

(i) is only predominant when the CDVT is close to zero. When the CDVT increases, cell rejection is mainly due to (ii). When this happens the analytical model shows that the rejection probability remains nearly constant.

When condition (ii) is predominant, Figures 3.A and 3.B show that the analytical model is initially more optimistic (it yields lower rejection probability), and then the curve obtained by simulation falls down showing a step behavior. The optimistic results are explained because the rate changes are applied later in the analytical model than in the simulation (remember that rate changes occur

at the cell emission epochs in the analytical model and at the backward RM-Cell arrivals in the simulation). The steps are caused because $\min(I_n^{old}, I_{n+1})$ is used in the simulation while I_n is used in the analytical model. This makes that each time a rate increase occurs in the simulator, e.g. from r_i to $r_{i'}$, the workload is increased by $1/r_{i'}$ instead of by the emission interval $1/r_i$, which causes an average reduction of $1/r_i - 1/r_{i'}$. Consequently, each time the CDVT added to these reductions compensates the wrong increments mentioned above, there is a reduction on the rejection probability.

The former differences between the simulator (which models the real DGCRA) and the analytical model could be removed. However, this would considerably increase the complexity of the analytical model. For dimensioning purposes the scenario with τ_2 properly tuned would be used, and in this scenario these differences do not have any influence.

Finally, Figure 4 shows the analytical and simulation results when there are no ABR rate changes. Without rate changes there are no differences between the two scenarios considered in this paper. Moreover, in this case there are no differences between the simulation and analytical models. This is confirmed by Figure 4 which shows a perfect agreement of the analytical and simulation results.

8.2 ANALYSIS

The Figures 2-4 show the influence of the following items on the rejection probability:

- Tuning of the τ_2 delay bound,
- jitter on the ABR stream (the higher the VBR rate, the higher the jitter),
- difference between the ABR rates,
- frequency of the ABR rate changes.

In Figure 2 we can see that when τ_2 is properly tuned, the rejection probability decreases exponentially when increasing the CDVT for low to moderated loads. In this case rejection probabilities of 10^{-9} can be reached with CDVT in the order of tens. Moreover, the frequency of the ABR rate changes has little influence on the results.

Figure 3 shows that when τ_2 is not properly set, the rejection probability has a major degradation. Moreover, the rejection probability does not decrease exponentially with the CDVT and is much more sensitive to the frequency of the ABR rate changes.

Figure 4 shows the rejection probability when the ABR source rate does not change. Note that the average load of the ABR source is maintained constant in all the figures (approximately equal to 0.2). By comparing Figure 4 with

the other figures we can see that the higher the differences between the ABR rates (while maintaining the same load) the higher the rejection probability. The figures also show that this effect increases very rapidly with an increasing overall load.

9. CONCLUSIONS

In this paper we have investigated the parameter dimensioning in the conformance definition for the ABR Service, the Dynamic Generic Cell Rate Algorithm (DGCRA). We have proposed an equivalent queuing model of the DGCRA and we have solved this model using a matrix analytic approach. The analytical results have been validated by simulation.

In the DGCRA three parameters are negotiated at the connection setup: the Cell Delay Variation Tolerance (CDVT), and the upper and lower bounds of the round trip delay between the UPC and the source, τ_2 and τ_3 respectively.

In the model we have considered the jitter introduced by a VBR source on an ABR cell stream sharing a common multiplexing stage. The model shows the influence of the following parameters on the rejection probability at the UPC: (i) Tuning of the τ_2 delay bound, (ii) jitter on the ABR stream (the higher the VBR rate, the higher the jitter), (iii) the difference between the ABR rates, (iv) the frequency of the ABR rate changes. These are investigated in two scenarios which show the influence of τ_2 :

1. Scenario with τ_2 properly tuned:

- For loads low to moderate, the rejection probability decreases exponentially when increasing the CDVT. In this case, rejection probabilities of 10^{-9} can be reached with CDVT in the order of tens.
- The higher the differences between the ABR rates (while maintaining the same load) the higher the rejection probability. This effect increases very rapidly with increasing overall loads.
- The frequency of the ABR rate changes has a minor influence on the rejection probability.

2. Scenario with $\tau_2 = \tau_3$ (the UPC does not apply a time tolerance to the scheduled rate changes):

- Compared with the former scenario, results show a major degradation of the rejection probability. This probability does not decrease exponentially when increasing the CDVT, but decreases at a slower rate.
- Frequency and amplitude of the ABR rate changes have a remarkable influence on the rejection probability.

Appendix: Solution of the Transition Probability Matrix

A.1 SCENARIO WITH $\tau_2 = \tau_3$

Since the ABR source can only emit one cell at each slot, $P(S, S')$ elements with queue length increments or decrements higher than one (i.e. $|Q' - Q| > 1$) are zero. Moreover, the transitions are exactly the same for all values with $Q \geq 1$. Therefore, we define: $\mathbf{B}_0 = \mathbf{Q}_{0,0}$, $\mathbf{A}_0 = \mathbf{Q}_{n+1,n}$, $n > 0$, $\mathbf{A}_1 = \mathbf{Q}_{n,n}$, $n > 0$ and $\mathbf{A}_2 = \mathbf{Q}_{n,n+1}$, $n \geq 0$. This yields the following structure for the matrix \mathbf{P} :

$$\mathbf{P} = \begin{pmatrix} \mathbf{B}_0 & \mathbf{A}_2 & \mathbf{0} & \mathbf{0} & \dots \\ \mathbf{A}_0 & \mathbf{A}_1 & \mathbf{A}_2 & \mathbf{0} & \dots \\ \mathbf{0} & \mathbf{A}_0 & \mathbf{A}_1 & \mathbf{A}_2 & \dots \\ \dots & \dots & \dots & \dots & \ddots \end{pmatrix} \quad (\text{A.1})$$

The following partitioned solution (π_0, π_1, \dots) of the stationary probabilities exists for this type of processes (see Neuts, 1981) (π_i , $i = 0, 1, \dots$ are vectors of length equal to the order of the blocks of the matrix \mathbf{P}):

$$\begin{aligned} \pi_k &= \pi_0 \mathbf{R}^k, & k \geq 1 \\ \pi_0 &= \pi_0 [\mathbf{B}_0 + \mathbf{R}\mathbf{A}_0] \\ \pi_0 (\mathbf{I} - \mathbf{R})^{-1} \mathbf{e} &= 1 \end{aligned} \quad (\text{A.2})$$

where \mathbf{R} has the same order as \mathbf{A}_i . \mathbf{I} is the unity matrix with the same order and \mathbf{e} is an all ones vector with corresponding length. To find \mathbf{R} we use the logarithmic reduction algorithm of Latouche and Ramaswami (see Latouche and Ramaswami, 1993).

Now we show how \mathbf{A}_2 may be derived, similar reasoning can be applied to derive \mathbf{A}_0 , \mathbf{A}_1 and \mathbf{B}_0 . In this case the queue length at the switch is increased by one, i.e. $Q' = Q + 1$. This can only happen if a cell is emitted by the ABR source (and thus $j = 1/r_i$) and a cell is emitted by the VBR source. Remember from Section 4.1 that the ABR source changes to state $j' = 1$ after a cell emission. Since there is no cell arrival at the UPC the workload is decreased by 1, and thus $U' = \max(U - 1, 0)$. Therefore, we have:

$$\begin{aligned} \mathbf{A}_2(S, S') &= P_{\text{ABR}}(i, i') \cdot P_{\text{VBR}}(k, k') \cdot v_k \cdot \\ &1[j = 1/r_i \text{ and } j' = 1 \text{ and } U' = \max(U - 1, 0)] \end{aligned}$$

where $\mathbf{A}(S, S')$ denotes the element (S, S') of the matrix \mathbf{A} .

A.2 SCENARIO WITH τ_2 PROPERLY TUNED

We approximate this scenario by flushing the switch buffer each time that a rate reduction occurs, i.e. $r_i > r_{i'}$ (see Section 6). For the reasons explained

below, the flushing is only performed with probability $1 - \alpha$. A transition from $Q \geq 0$ to $Q' = 0$ occurs when the queue is flushed, therefore, we obtain the following structure for the transition probability matrix $\mathbf{P}^{(p)}$ (by the superscript $^{(p)}$ we shall distinguish the matrices derived in this scenario from those used in Appendix A.1):

$$\mathbf{P}^{(p)} = \begin{pmatrix} \mathbf{B}_0^{(p)} & \mathbf{A}_2^{(p)} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \dots \\ \mathbf{A}_0^{(p)} + \mathbf{A}_3^{(p)} & \mathbf{A}_1^{(p)} & \mathbf{A}_2^{(p)} & \mathbf{0} & \mathbf{0} & \dots \\ \mathbf{A}_3^{(p)} & \mathbf{A}_0^{(p)} & \mathbf{A}_1^{(p)} & \mathbf{A}_2^{(p)} & \mathbf{0} & \dots \\ \mathbf{A}_3^{(p)} & \mathbf{0} & \mathbf{A}_0^{(p)} & \mathbf{A}_1^{(p)} & \mathbf{A}_2^{(p)} & \dots \\ \vdots & \vdots & \vdots & \ddots & \ddots & \ddots \end{pmatrix} \quad (\text{A.3})$$

The submatrices are given by: $\mathbf{A}_0^{(p)} = \mathbf{A}_0$, $\mathbf{A}_1^{(p)} = \mathbf{A}'_1$, $\mathbf{A}_2^{(p)} = \mathbf{A}'_2$, $\mathbf{A}_3^{(p)} = \mathbf{A}''_1 + \mathbf{A}''_2$, $\mathbf{B}_0^{(p)} = \mathbf{B}_0 + \mathbf{A}''_2$. where the matrices \mathbf{A}_0 and \mathbf{B}_0 are obtained as in Appendix A.1. The matrices \mathbf{A}'_1 and \mathbf{A}'_2 are respectively obtained replacing $P_{\text{ABR}}(i, i')$ by $P_{\text{ABR}}(i, i') \cdot (1[r_i \leq r_{i'}] + \alpha \cdot 1[r_i > r_{i'}])$ in the relations given for \mathbf{A}_1 and \mathbf{A}_2 in Appendix A.1. Similarly, the matrices \mathbf{A}''_1 and \mathbf{A}''_2 are obtained replacing $P_{\text{ABR}}(i, i')$ by $P_{\text{ABR}}(i, i') \cdot (1 - \alpha) \cdot 1[r_i > r_{i'}]$ in the relations of \mathbf{A}_1 and \mathbf{A}_2 .

To solve the former matrix $\mathbf{P}^{(p)}$, the matrix $\mathbf{A}_0^{(p)} + \mathbf{A}_1^{(p)} + \mathbf{A}_2^{(p)}$ must be irreducible (see Neuts, 1978). By choosing $\alpha > 0$, this condition is fulfilled. $\mathbf{P}^{(p)}$ is then solved using equations (A.2) except that the second equation has to be replaced by $\pi_0 = \pi_0 [\mathbf{B}_0^{(p)} + \mathbf{R} \mathbf{A}_0^{(p)} + ((\mathbf{I} - \mathbf{R})^{-1} - \mathbf{I}) \mathbf{A}_3^{(p)}]$. The matrix \mathbf{R} still obeys a similar equation as before and therefore can be found using the L-R algorithm (see Latouche and Ramaswami, 1993).

References

- ATM Forum (1996). “ATM Forum Traffic Management Specification Version 4.0”. ATM Forum Technical Committee Traffic Management Working Group.
- Latouche, G. and Ramaswami, V. (1993). “A logarithmic reduction algorithm for Quasi-Birth-Death processes”. *Journal of Applied Prob.*, 30:650–674.
- Neuts, M. (1978). “Markov Chains with Applications in Queueing Theory, which have a Matrix-Geometric Invariant Probability Vector”. *Journal of Applied Prob.*, 10:185–212.
- Neuts, M. (1981). “*Matrix-Geometric Solutions in Stochastic Models*”. The John Hopkins University Press, Baltimore.
- Roberts, J., Mocchi, U., and Virtamo, J., editors (1996). “*Broadband Network Teletraffic - Final Report of Action COST 242*”. Springer Verlag.