

# **Adaptive Polling Schemes for an ATM Bus with Bursty Arrivals**

**Johan M. Karlsson \***

Communication Systems Department  
Lund University  
Box 118, S-22100  
Lund, Sweden

**Harry G. Perros**

Computer Science Department and  
Center for Communications and Signal Processing  
North Carolina State University  
Raleigh, NC 27695

**Ioannis Viniotis**

Department of Electrical and Computer Engineering and  
Center for Communications and Signal Processing  
North Carolina State University  
Raleigh, NC 27695

---

\*Work done while visiting the Center for Communications and Signal Processing, North Carolina State University.

## **Abstract**

We consider a high speed ATM bus, which can serve a small number of stations ( $\leq 16$ ). The bus is slotted and each slot is long enough to contain one ATM cell. Transmission on the bus is controlled by an arbitrator. Each station sends its status information to the arbitrator at regular intervals. Based on this information the arbitrator makes polling decisions. Due to transmission and processing times the arbitrator makes decisions based on status information which is a few bus slots old. In this paper, we propose adaptive polling schemes assuming bursty arrivals. The schemes take into account both the cell loss probability and the mean waiting times. The advantage of these new schemes over ordinary polling is also demonstrated.

**Key words:** Adaptive Polling, ATM, bus architecture, bursty arrivals.

# 1 Introduction

The Asynchronous Transfer Mode (ATM) appears to be the most promising solution for broadband ISDN. ATM provides the means for transporting different types of highly bursty traffic such as voice, video and bulk files. It is based on the principle of packet switching. In ATM, a packet has a fixed length of 53 bytes, and it is known as a *cell*. Various switch architectures have been proposed for an ATM network. For a review of these architectures see [1].

In this paper, we consider a high speed bus for ATM traffic. The bus can be used as a local switch for a small number of input ports. Transmission on the bus is controlled by an arbitrator. The arbitrator receives status information from each of the stations, and then based upon this information, it makes polling decisions.

The problem of polling has been extensively studied in the literature, see [2] and [3]. There are three well known polling schemes, namely: *limited*, *gated* and *exhaustive*. In the limited polling scheme, a polled station transmits a predefined maximum number of packets. In the gated scheme a polled station transmits all the packets that were in its queue when it was polled. Finally, in the exhaustive scheme, a polled station transmits all the packets that were in the queue when it was polled, plus any other packets that might arrive during their transmission time. Several variants of the above basic polling schemes include probabilistic polling, and periodic polling, see [4].

In this paper we investigate various *adaptive* polling schemes through simulation. To the best of our knowledge, adaptive polling has not as yet been investigated in the literature.

The paper is organized as follows. In the following section we describe the structure of the bus, and we discuss how bursty arrivals are modelled. In section 3, we discuss the adaptive polling schemes examined in this paper. The results are given in section 4, and the conclusions in section 5.

## 2 Model Description

For some applications a need for a switch between different user groups would be necessary. Each group of users is connected to the switch through a station, as shown in Figure 1. In our approach this switch is accomplished by a high speed bus. Physically this bus is rather short ( $\sim 1$  yard) and can

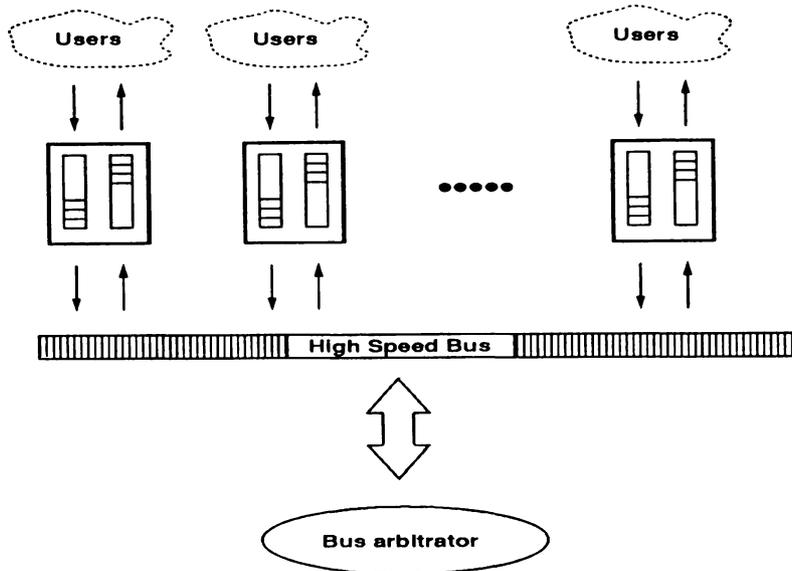


Figure 1: Model of the high speed bus under study.

fit in a rack. Up to 16 stations can be attached to the bus. Due to these restrictions, the propagation delay is very small. Each station serves a pool of users, which can be high speed users, a gateway to a small network, a LAN, a super computer, etc. Each station consists of an input and an output buffer, a counter which keeps track of the number of cells in the input buffer, and a scanning point for communication to and from the arbitrator. A generic model of a station is shown in Figure 2. When the input buffer becomes full, the arriving cells from the users get lost. The same holds true for the output buffer. During the time it is full, arriving cells from other stations get lost.

## 2.1 Bus structure

The bus is slotted, and each slot is long enough to contain one cell. The speed of the bus is assumed to be  $\sim 2$  Gbps. The speed of the incoming and outgoing lines of each station is assumed to be one third of the speed of the bus ( $\sim 600$  Mbps). Transmission on the bus is controlled by an arbitrator. In particular, each station sends information about its status to the bus

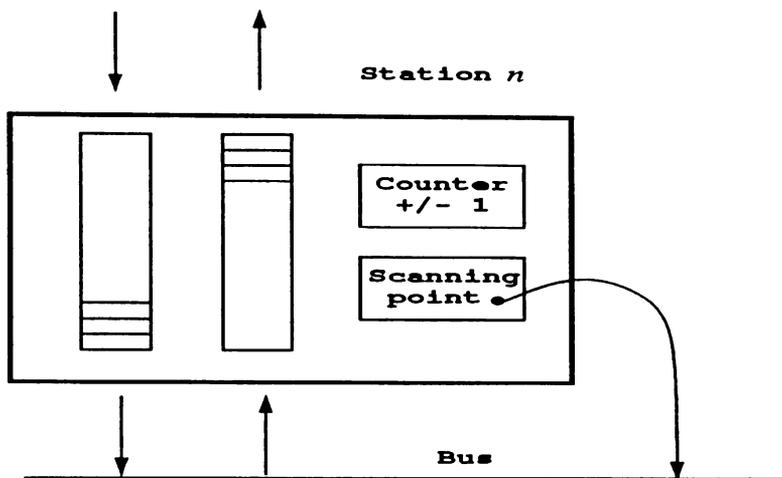


Figure 2: A generic model of a station connected to the high speed bus.

arbitrator at regular intervals. The arbitrator gathers the information from all the stations and then decides in which order the stations will be served. A small delay between the time the stations transmit their status information and the decision of the arbitrator occurs. This delay (hereafter referred to as *decision delay*) consists of the following time components:

- transmission time of status information to the arbitrator;
- processing time at the arbitrator in order to decide the polling sequence;
- transmission time of the polling sequence to the stations.

In view of this delay, the arbitrator makes decisions based on information which is a few bus slots old. For instance, if it takes  $\nu$  bus slots to receive, process and transmit back to the stations the polling sequence, then all stations send information every  $\nu$ th slot, and the order in which the stations

will be served is decided for  $\nu$  consecutive slots at each decision time. This information is transmitted by using some of the capacity of the bus, or it is transmitted on a separate bus. For the bus speed given above, the decision delay would be two bus slots.

## 2.2 The Arrival Process

The arrival process from a user group to a station is assumed to be an Interrupted Bernoulli Process (IBP). That is, for a geometrically distributed period (active state) arrivals occur according to a Bernoulli process. This period is followed by another period (idle state) which is also geometrically distributed, during which no arrivals occur. An IBP is governed by a two state Markov chain, in which the time unit is equal to a slot of the incoming link to a station. Given that we are in the active state, the process will remain in this state with probability  $p$  or it will change to the idle state with probability  $1 - p$ . If the process is in the idle state, it will remain in that state with probability  $q$ , or it will change to the active state with probability  $1 - q$ . When in active state, a slot contains a cell with probability  $\alpha$ . Each station is characterized by its own  $p$ ,  $q$  and  $\alpha$ . It can be shown that the generating function,  $G(z)$ , of the probability distribution of the interarrival time,  $\xi_A$ , is:

$$\begin{aligned} G(z) &= E(z^{\xi_A}) \\ &= \frac{z\alpha(p + z(1 - p - q))}{(1 - \alpha)(p + q - 1)z^2 - (q + p(1 - \alpha))z + 1} \end{aligned} \quad (1)$$

From this generating function we can obtain the mean interarrival time  $E(\xi_A)$  and the second moment of the interarrival time  $E(\xi_A^2)$ ,

$$E(\xi_A) = \frac{(2 - p - q)}{\alpha(1 - q)} \quad (2)$$

$$\begin{aligned} E(\xi_A^2) & \quad (3) \\ &= \frac{8 - 4\alpha + (5\alpha - 8)(p + q) + (2 - \alpha)q^2 + 2(1 - \alpha)p^2 + (4 - 3\alpha)pq}{\alpha^2(1 - q)^2} \end{aligned}$$

From the above equations, we can also obtain the squared coefficient of variation of the interarrival time,  $C^2$ ,

$$\begin{aligned} C^2 &= \frac{\text{Var}(\xi_A)}{[E(\xi_A)]^2} \\ &= 1 + \alpha \left( \frac{(1-p)(q+p)}{(2-p-q)^2} - 1 \right). \end{aligned} \quad (4)$$

The probability  $\rho$ , that any slot is busy (i.e. it carries a cell) is equal to the mean number of cells transmitted during the active period (which is equal to the mean length of the active period) over the mean length of the silence and active period. This is equal to  $1/E(\xi_A)$ , or

$$\rho = \frac{\alpha(1-q)}{2-p-q}. \quad (5)$$

We note that  $\rho$  is also known as the load or the mean bandwidth. In our analysis of the system we will assume that  $\alpha = 1$  which gives us,

$$\rho = \frac{1-q}{2-p-q} \quad \text{and} \quad C^2 = \frac{(p+q)(1-p)}{(2-(p+q))^2}. \quad (6)$$

Now, by varying  $p$  and  $q$  we can alter the values of the mean bandwidth  $\rho$ , and of  $C^2$ , which can be seen as a measure of burstiness. In particular, from (6) we have

$p, q$	$\rho$	$C^2$
$p \rightarrow 0, q \rightarrow 0$	$\rho \rightarrow 0.5$	$C^2 \rightarrow 0$
$p \rightarrow 0, q \rightarrow 1$	$\rho \rightarrow 0$	$C^2 \rightarrow 1$
$p \rightarrow 1, q \rightarrow 0$	$\rho \rightarrow 1$	$C^2 \rightarrow 0$
$p \rightarrow 1, q \rightarrow 1$	$\rho \rightarrow 0.5$	$C^2 \rightarrow \infty$

We note that in this study, we have concentrated on the interaction between the input buffers of the stations, the bus, and the bus arbitrator. The output buffers of the stations were not considered. Also, the user groups were modelled indirectly by the arrival process to each station that was described above. Below, we shall refer to the input buffer of a station as simply the “*station*”.

### 3 Adaptive Polling Schemes

The objective in this paper is to develop an adaptive polling scheme that equalizes the cell loss probability and the mean waiting time at each station. We mainly focus on equalizing the mean waiting time at each station, seeing that due to the bandwidth allocation algorithm the cell loss probability at each station is expected to be less than  $10^{-10}$ . That is, we assume that the traffic streams offered to each station have already been accepted by a bandwidth allocation algorithm, which guarantees a certain cell loss probability. In view of this, it does not make sense to develop an adaptive polling scheme that equalizes or minimizes cell loss probabilities, even if it means that one station may lose a few cells more than another one. Below, we examine three different policies for equalizing the mean waiting time. These policies guarantee a degree of fairness among the stations. They can be easily modified to provide for priorities among the stations<sup>1</sup>. They are easily implementable since they only depend on knowledge of the arrival rate  $\lambda_i$  (assumed known from the bandwidth allocation algorithm), and not on any other arrival/service statistics. We also consider a fourth polling policy which equalizes the mean waiting time in each station, but which in certain cases switches to equalizing the cell loss probabilities. We introduce it for the purpose of providing better control of the transient behaviour. For a detailed description and a theoretical analysis of adaptive rules (in a CPU scheduling model), the reader is referred to [5, 6].

These policies are as follows:

- *Policy 1:* This is the simplest policy, where at time  $t$  each station sends its current queue size  $n_i(t)$  to the bus arbitrator. The arbitrator makes its decision based on the queue sizes. In view of this, we shall refer to this policy as the *queue size (QS)* policy.
- *Policy 2:* Due to the negotiation at set-up time, the arbitrator knows the mean bandwidth  $\lambda_i$  of the traffic arriving at each station  $i$ . (In real life the arbitrator may only have an estimate of the mean bandwidth.

---

<sup>1</sup>Instead of equalizing the mean waiting times, we could equalize  $E(\text{waiting time})/g_i$ , for some given constant  $g_i$ . Then, the higher the value of  $g_i$ , the lower the priority of station  $i$ . In this case, the policy would calculate the ratios  $\frac{n_i(t)}{\lambda_i g_i}$ .

It uses this estimate and updates it from measurements.) This information is used together with the queue size  $n_i(t)$  transmitted by each station  $i$ , in order to make a decision. In particular, for station  $i$ , the arbitrator calculates the ratio  $\frac{n_i(t)}{\lambda_i}$ . From Little's law, this can be seen as an estimate of the momentary waiting time at station  $i$  at time  $t$ . Because of this, we shall refer to this policy as the *momentary waiting time (MWT)* policy.

*Policy 3:* The above policy uses instant values for the number of cells in each station. This policy is similar to Policy 2 (MWT), only we use a cumulative value for the number of cells in a station. In particular, the following cumulative value for each station  $i$  is used,

$$\sum_{t=t_{decision}-t_0}^{t_{decision}} \frac{n_i(t)}{\lambda_i}, \quad i = 1, 2, \dots, N$$

where  $t_{decision}$  is the time at which the polling decision sequence for the next decision cycle is made and  $t_0$  is a value between 0 and  $t_{decision}$ . These two values determine a window for the above summation. The window reflects a compromise between two desirable characteristics of the policy. More, specifically, we want the policy to be responsive to transient fluctuations of delay (small window). On the other hand we also want the policy to equalize average delays (large window). In this paper, the value of  $t_0$  is set equal to  $t_{decision}$ . We shall refer to this policy as the *cumulative waiting time (CWT)* policy.

Policy 3 appears more complex than policy 2. Observe, however, that the above cumulative sum can be implemented recursively and thus it requires only one addition per decision.

*Policy 4:* This scheme is a combination of the QS and CWT policies. The arbitrator normally operates under the CWT policy. However, if one or more stations lose more than a predefined number of consecutive cells, the arbitrator switches to the QS policy. The motivation behind this policy is that the arbitrator will equalize mean waiting times using the CWT policy. However, if one or more stations start losing packets, it will switch to the QS policy in order to equalize cell loss probabilities

Each station maintains a flag. When a station  $i$  loses more than  $K_i$  consecutive cells, the station raises its flag. The station lowers its flag when the number of cells in its queue drops below  $N_i - K_i$ , where  $N_i$  is the stations maximum capacity. When the arbitrator sees that a station has raised its flag, it switches to the QS policy for all stations. When the station lowers its flag, the arbitrator switches back to the CWT policy for all stations. During this time, it is possible that other stations may raise their flag. In this case, the arbitrator does not switch back to the CWT policy until all flags have been lowered.

During the time that one or more flags are up and the arbitrator operates under the QS policy, the arbitrator continues to update the  $\sum_t \frac{n_i(t)}{\lambda_i}$  values for all stations. Because this policy alternates between the CWT and QS policies, we shall refer to it as the *CWT-QS* policy.

We note that when two or more stations have the same priority (for instance, they have the same queue size under the QS policy,) one of these stations is selected randomly. In addition to using one of the above four policies, the arbitrator makes use of one of the three classical polling rules, i.e. limited, gated, or exhaustive. Each of these four policies, therefore, can be seen as providing a rule for deciding which station will transmit next. How many cells the station will transmit, depends on the polling rule that the arbitrator uses. In order to clarify this point further, let us assume that the arbitrator operates under the CWT policy and under the limited polling rule whereby each station is allowed to transmit only one cell. Furthermore, the decision delay is 3 bus slots. That is, the arbitrator has to decide on the polling sequence for the next three slots under information which is three slots old. Using the CWT policy, the arbitrator decides which station will transmit next. Let  $j$  be this station. The number of cells in the  $j$ th station is reduced by one, and its cumulative waiting time  $\sum_t \frac{n_j(t)}{\lambda_j}$  is recalculated assuming no new arrivals. The arbitrator now applies the CWT policy once more to decide which station will transmit next using the new value for the  $j$ th station, and the same values for all the other stations. This process repeats once more, before the polling sequence has been completely decided.

Let us now assume that in the above example the arbitrator uses the gated polling rule. As above, let  $j$  be the station that will transmit first according

to the CWT policy. If there are more than three cells in its buffer, say 5 cells, station  $j$  will transmit for the next three successive slots. The 2 cells left in the buffer are served during the next decision cycle. If station  $j$  has less than three cells, say 2 cells, it will transmit for the next two slots. The number of cells in station  $j$  is reduced by 2, and its value used by the CWT policy is updated assuming no new arrivals. Then, using the CWT policy the arbitrator decides which station will transmit during the third slot. The new value for station  $j$  and the same values for the other stations are used for this decision.

Finally, let us consider the case where exhaustive polling is used under the CWT policy. Again, let station  $j$  be the station that will transmit first, according to the CWT policy. Let us assume that the station contains 4 cells. Then the first three cells will be transmitted during the next decision cycle, call it cycle  $\tau$ . The remaining cell plus any new arrivals will be transmitted during decision cycle  $\tau + 1$ . Due to the assumed difference between the bus speed and the incoming lines, at most one cell may arrive during a decision cycle. Let us assume that station  $j$  reports 2 cells at the end of decision cycle  $\tau$ . Then, during decision cycle  $\tau + 1$ , station  $j$  will transmit 2 cells. The number of cells in station  $j$  is reduced by 2, and its value for the CWT policy is recalculated assuming no new arrivals. That is, if a cell arrives at the station during the transmission time of these two cells, the cell will not be transmitted until the same station is chosen again by the CWT policy. The arbitrator then applies the CWT policy in order to decide which station will transmit next in the  $\tau + 1$  decision cycle.

## 4 Results

The performance of each of these four policies under the limited, gated and exhaustive polling rule, was evaluated using simulation techniques. Five stations were simulated, each with a buffer of 10. The speed of the bus was assumed to be three times the speed of the incoming line into each station. The decision delay was assumed to be two bus slots. To get a more objective view of the performance of the polling policies presented above, different traffic loads have been considered. All results in Figures 3-5 were obtained for symmetrical loads, while the results in Figures 6-16 were obtained for the three different asymmetrical traffic data sets A,B and C shown in Table 1

Station	Traffic Data Set					
	A		B		C	
	$\rho$	$C^2$	$\rho$	$C^2$	$\rho$	$C^2$
1	0.6	10	0.7	10	0.4	5
2	0.5	10	0.5	10	0.6	5
3	0.4	10	0.3	10	0.6	5
4	0.5	10	0.5	10	0.6	5
5	0.6	10	0.7	10	0.4	5

Table 1:  $\rho$  and  $C^2$  for the different traffic data sets A, B and C.

Observe that the total load on the bus (sum of the  $\rho$  values of all stations) is assumed to be less than the capacity of the bus. Since the bus slot is one third of the arrival slot, the total input load should be less than 3. For example, in the symmetric case reported in Figures 3-5, the value of  $\rho$  of each station cannot be more than 0.6.

We first note that in the well known limited, gated and exhaustive polling, bus cycles are lost while empty stations are being polled. Under the same polling rules, a significant improvement in the cell loss probability and the mean waiting time can be achieved by using the fact that the arbitrator has knowledge of the status of each station. (This knowledge is, of course, a few slots old.) Using this knowledge, the arbitrator will never poll an empty station, i.e. a station which at the beginning of the decision cycle was empty. The cell loss probability and the mean waiting time are given in Figures 3 and 4 respectively, for various values of  $\rho$ , for limited, gated and exhaustive polling assuming  $C^2 = 10$ . For comparison purposes, we also give the results for the time division multiplexing (TDM) rule. Under this rule, the arbitrator visits each station cyclically, and only one cell per station is transmitted. If a station is empty, the bus slot is wasted. The difference in performance is quite substantial. In Figure 5, we plot the mean waiting time for the exhaustive polling rule as a function of  $\rho$ , and for three different values of the decision delay, assuming that  $C^2 = 100$ . We note that as the decision delay increases, the mean waiting time increases as well. This is mainly due to the fact that when the decision delay gets longer, the cells in the stations have to wait longer to obtain service. Finally, the cell loss probabilities and the

mean waiting times per station for each of the three polling rules are shown in Figures 6 and 7 respectively, assuming a decision delay of 2 slots. We note that the performance with respect to cell loss probabilities and mean waiting times, would be affected by the variation in the load ( $\rho$ ) and burstiness ( $C^2$ ) of the arrival process. In particular, the mean waiting time increases as  $\rho$  and/or  $C^2$  increase. We now proceed to discuss the behaviour of the adaptive polling policies described in the previous section.

In Figures 8 and 9, the QS and the CWT policies are compared, assuming exhaustive polling. We note that under the QS policy, the cell loss probabilities are equalized, and under the CWT policy, the mean waiting times are equalized. In Figure 10, the QS and the MWT policies are compared for the gated polling rule. Since the MWT policy favours stations with lower mean arrival rates  $\lambda_i$  these stations achieve a smaller mean waiting time than the other stations with higher arrival rates  $\lambda_i$  which have a longer mean waiting time. This is shown in Figure 10. Similar conclusions are drawn for the MWT policy and the CWT policy, in Figure 11.

In Figure 12, the mean waiting time per station is given for the CWT and CWT-QS policies, assuming exhaustive polling. The CWT policy is an appropriate policy for equalizing the mean waiting times. We note that the CWT-QS policy performs quite well with respect to equalizing the mean waiting times. In spite of the fact that the CWT-QS policy is not using the optimal rules for equalizing the mean waiting times, these times remain, almost unaffected as it could be seen in Figure 12. In Figure 13, the cell loss probability is given for the same traffic data set to show the impact of the CWT-QS policy on the cell loss probability. Exhaustive polling is assumed. As it can be seen, the CWT-QS policy levels out the cell loss probability in contrast to the CWT policy.

In Figure 14 and 15 the performance of the QS policy is compared to the CWT-QS policy for traffic data set B which is more asymmetrical than traffic data set A or C. Exhaustive polling is assumed. The CWT-QS policy equalizes the mean waiting times at the expense of higher cell loss probabilities.

From the above results, we conclude that the adaptive policy CWT-QS is the best policy to use in the high speed bus. It achieves a compromise between

the QS policy that equalizes cell loss rate and the CWT policy that equalizes mean waiting times.

## 4.1 Burst size

The probability of loss is higher during long bursts of arrivals rather than during short bursts of arrivals. This is shown in Figure 16, where the cell loss probability for station 1 has been plotted as a function of burst size; the CWT-QS policy under the gated policy rule was used. Burst sizes have been classified into 10 groups, so that group  $x$  means bursts of size  $[10(x - 1), 10(x - 1) + 9]$ . The last group also includes burst sizes greater than 200. This is compared to the overall calculated mean cell loss probability, which is shown by the dotted line in the graph.

In the IBP arrival process considered in this paper, the bursty period is identical to the busy period of the process. Therefore, the expected burst size is equal to  $\frac{1}{1-\rho}$ . This can be expressed as follows in terms of  $C^2$  and  $\rho$ :

$$E(\text{burst size}) = \frac{C^2 + (1 - \rho)}{2(1 - \rho)^2}. \quad (7)$$

We note that for the results given in Figure 16, the expected burst size is equal to 7.78. Therefore, if we treat all burst sizes to be equal to the expected burst size, we overestimate the mean cell loss probability.

## 5 Conclusion

In this study, we have presented a set of new adaptive polling policies. The adaptiveness in our policies is due to the fact that we have some knowledge of the status (number of cells in the input buffer) of each station, and therefore we are able to make adaptive changes to the polling decisions. The polling policies focus on equalizing the mean waiting time (delay) in the buffers, before a cell gets transmitted on the bus. In ATM bus environments, users go through a bandwidth allocation algorithm which guarantees that cell loss probabilities are very small (i.e.  $10^{-9}$  or  $10^{-10}$ ).

A number of policies has been investigated, mainly in terms of waiting times

Simulations have shown that the policies work very well at equalizing mean waiting times, even when they operate with limited, outdated state knowledge. No prior knowledge of user statistics is assumed (except for arrival rates, which are known to the bandwidth allocation algorithm, anyway). The price of having slightly higher cell loss probabilities is, of course, paid. An improved, slightly more complex rule, which still equalizes delays and exhibits lower cell loss probabilities is also analyzed. This rule pays attention to the station(s) which are in bursty mode and lose a lot of cells. The results show considerably lower waiting times as well as cell loss probabilities.

## Acknowledgements

We would like to thank Dan Stevenson, Microelectronics Center of North Carolina, for describing this problem to us. Also, we would like to thank the two anonymous referees for their suggestions.

## References

- [1] F.Tobagi, "Fast Packet Switch Architectures for Broadband Integrated Services Networks", *Proc. IEEE*, Vol 78 (1990) 1133-1167.
- [2] D. Grillo, "Polling Mechanism Models in Communication Systems - Some Application Examples", in: Takagi, Ed., *Stochastic Analysis of Computer and Communication Systems* (North Holland, Amsterdam, 1990), 659-698.
- [3] H. Takagi, "Queueing Analysis of Polling Models: An Update", in: Takagi, Ed., *Stochastic Analysis of Computer and Communication Systems* (North Holland, Amsterdam, 1990), 267-318.
- [4] O. J. Boxma, H. Levy and J. A. Westrate, "Optimization of Polling Systems", in: King, Mitrani, Pooley, *PERFORMANCE'90* (North Holland, Amsterdam, 1990), 349-361.
- [5] P. P. Bhattacharrya, L. Georgiadis, P. Tsoucas, I. Viniotis, "Ergodicity and Optimality of an Adaptive Multiobjective Algorithm", submitted to *Mathematics of Operations Research*. 1990.

- [6] L. Georgiadis, C. Nikolaou, I. Viniotis, “*A Family of Job Scheduling Strategies for Average Response Time Constraints*”, Proceedings of the 27th Annual Allerton Conference on Control, Communications and Computers, Urbana, Illinois, 1989, 810-819.

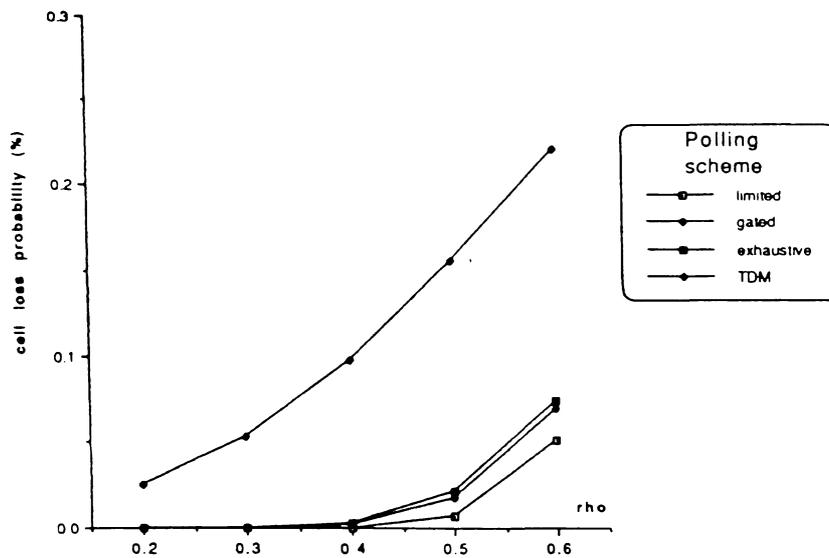


Fig. 3. Cell loss probability vs rho for different polling schemes

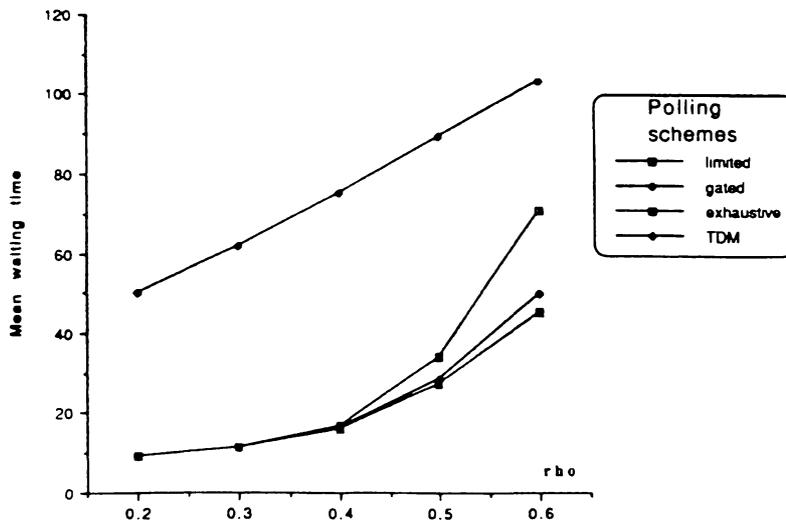


Fig. 4. Mean waiting time vs rho for different polling schemes

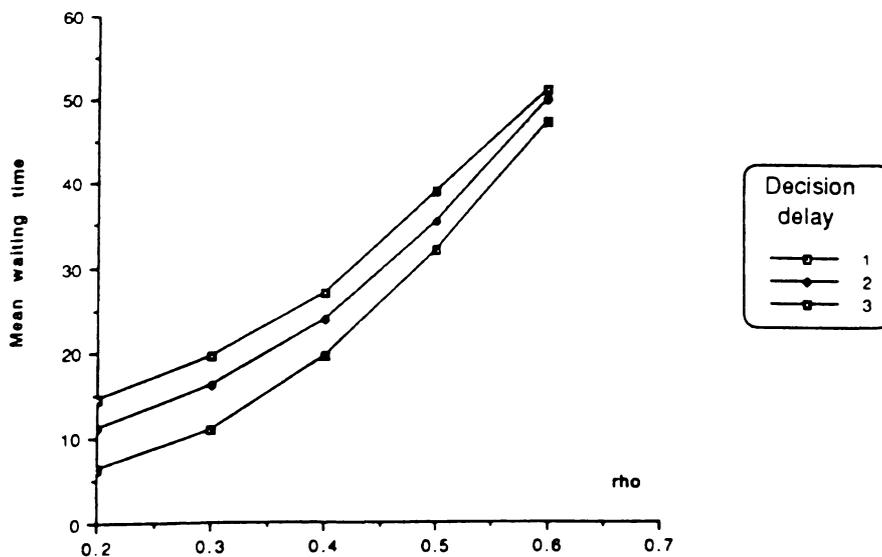


Fig. 5. Mean waiting time vs rho for different decision delays

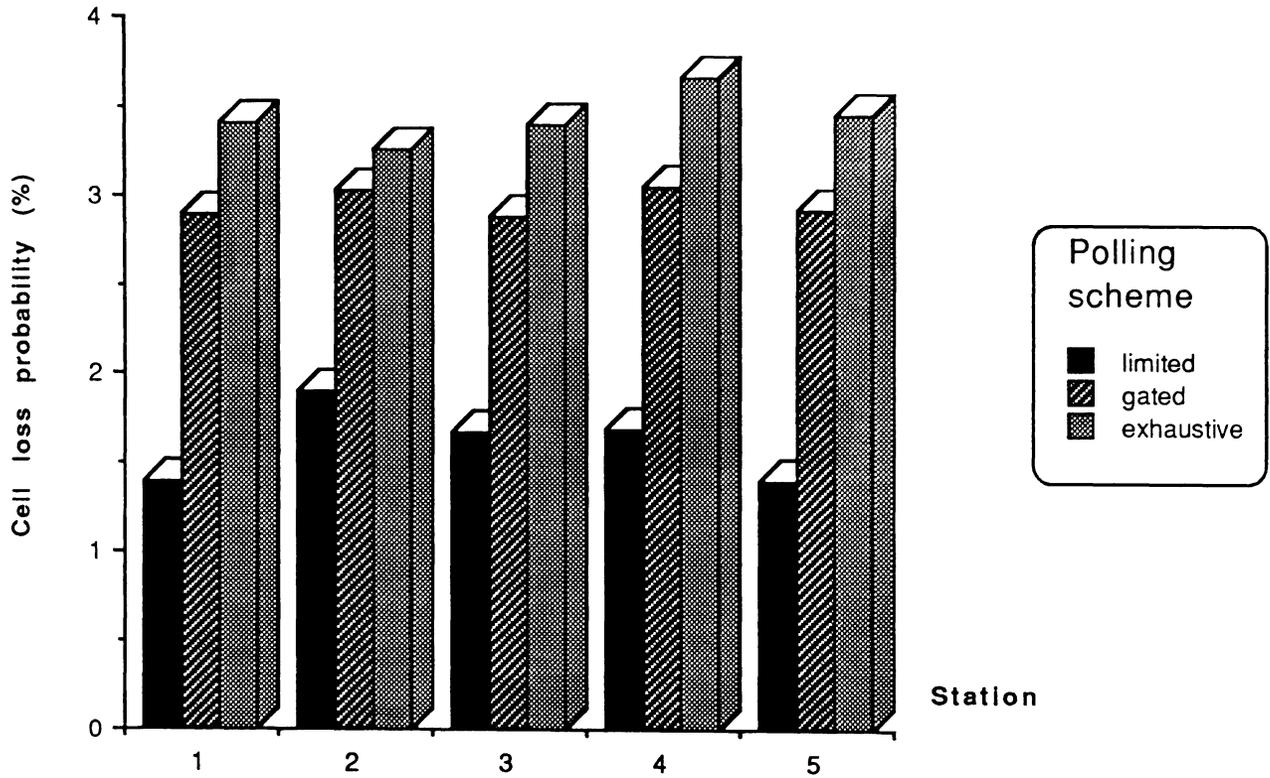


Fig 6. Cell loss probability for the stations for different polling schemes; traffic data set A.

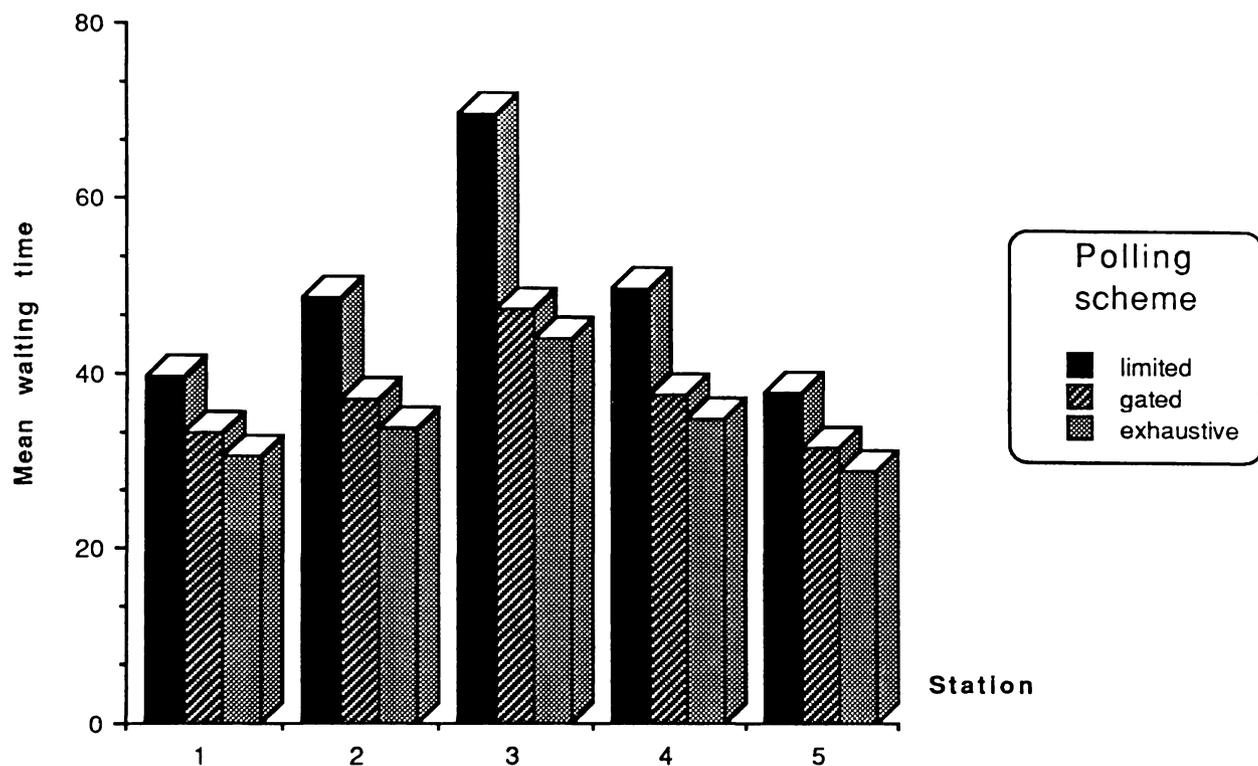


Fig 7. Mean waiting time for the stations for different polling schemes; traffic data set A.

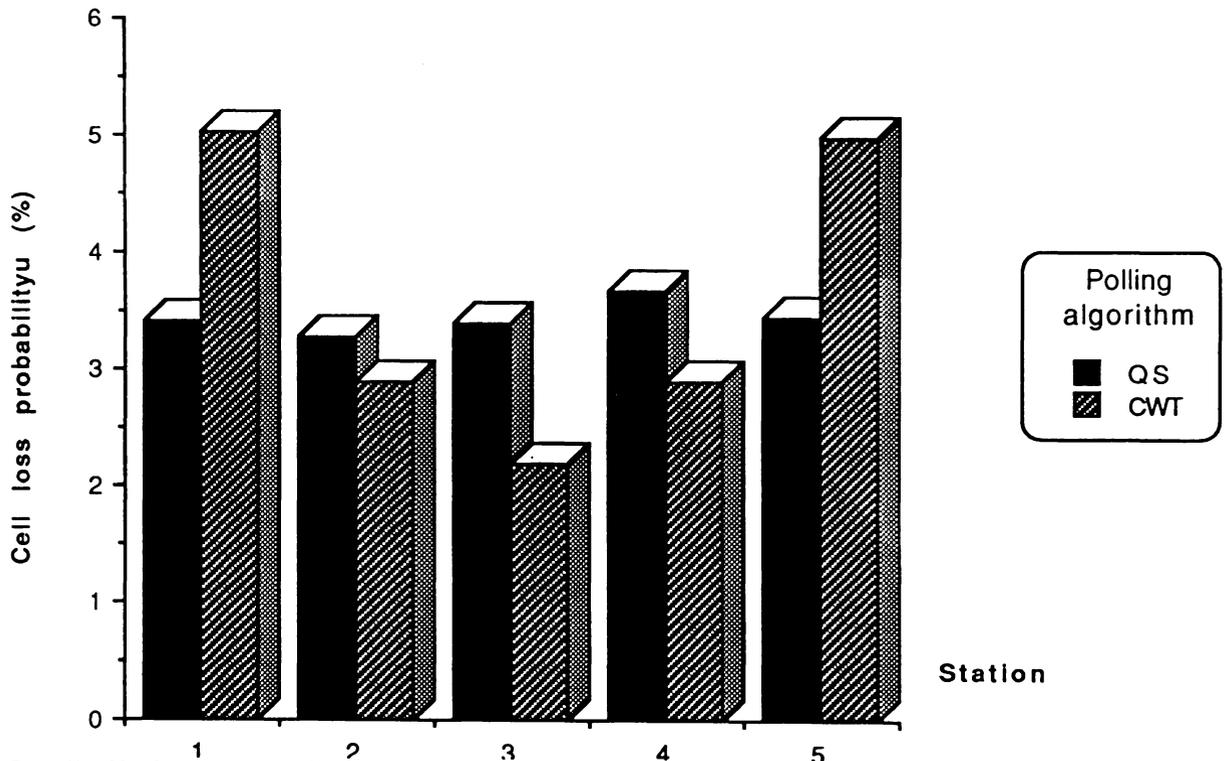


Fig 8. Cell loss probability at each station for the QS and CWT policies; traffic data set A.

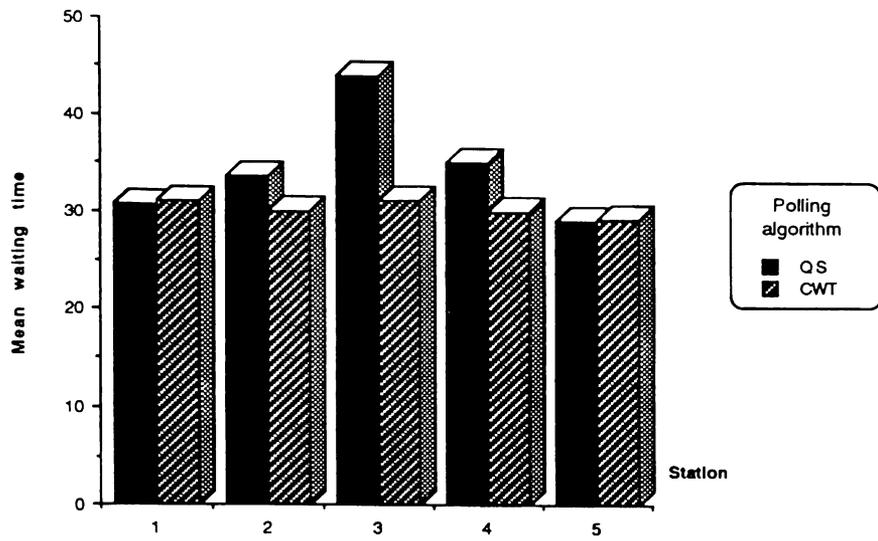


Fig 9. Mean waiting time at each stations for the QS and CWT policies; traffic data set A.

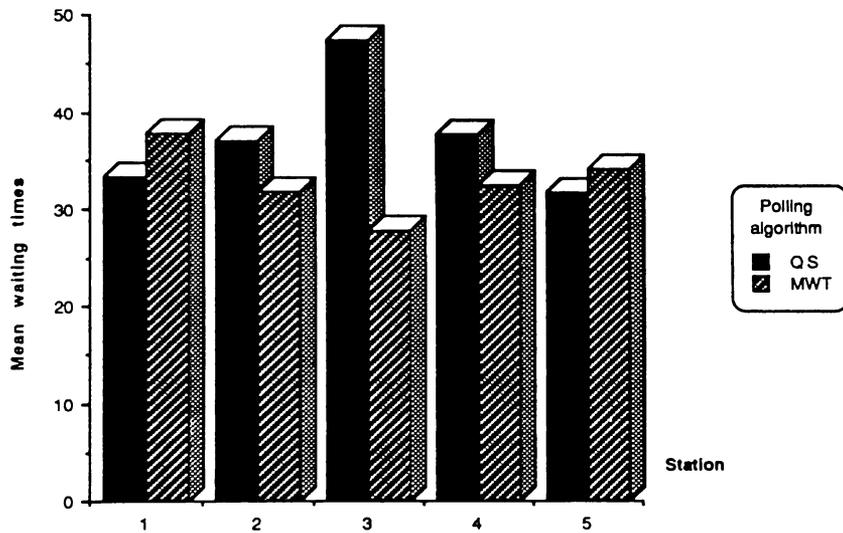


Fig 10. Mean waiting time at each station for the QS and MWT policies; traffic data set B.

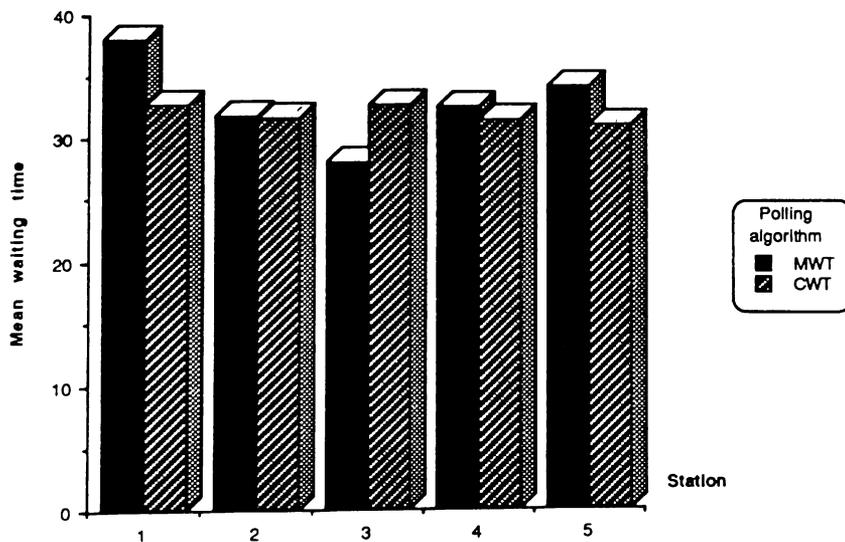


Fig 11. Mean waiting time at each station for the MWT and CWT policies; traffic data set B.

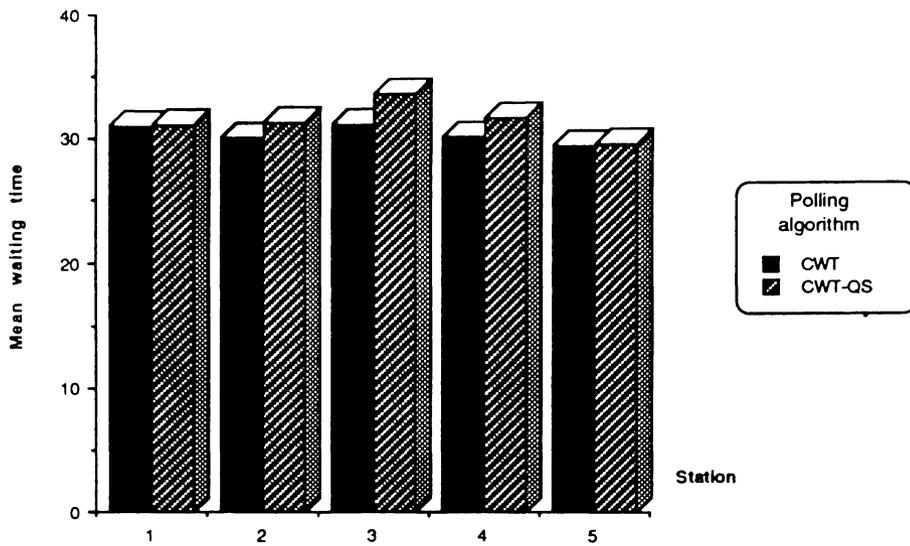


Fig 12. Mean waiting time at each station for the CWT and CWT-QS policies; traffic data set A.

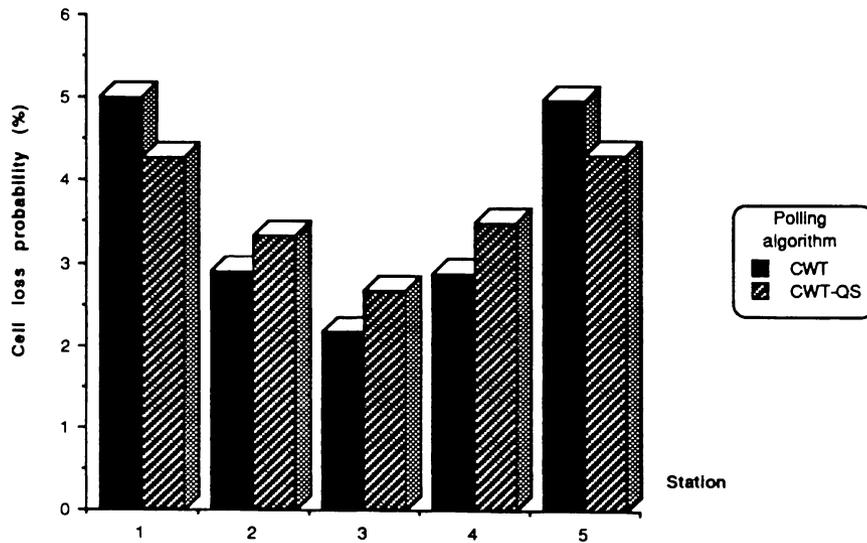


Fig 13. Cell loss probability at each station for the CWT and CWT-QS policies; traffic data set A.

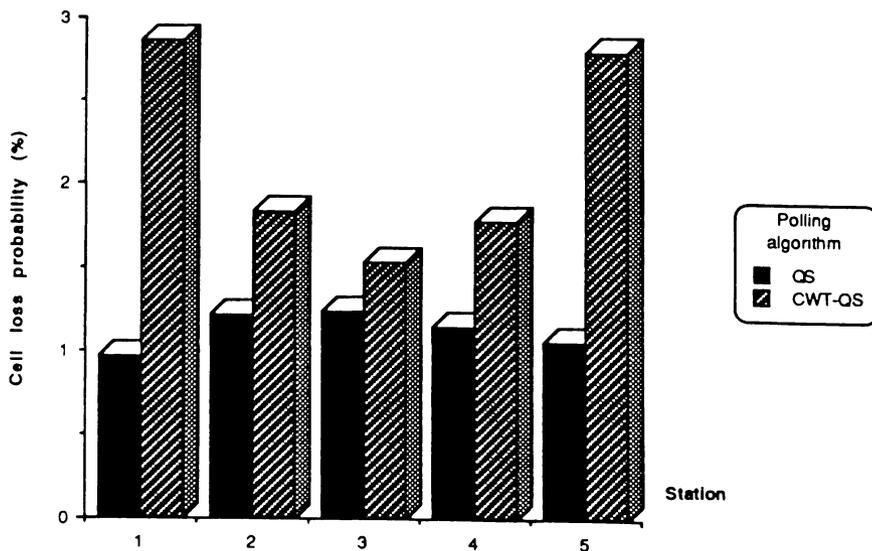


Fig 14. Cell loss probability at each station for the QS and CWT-QS policies; traffic data set B.

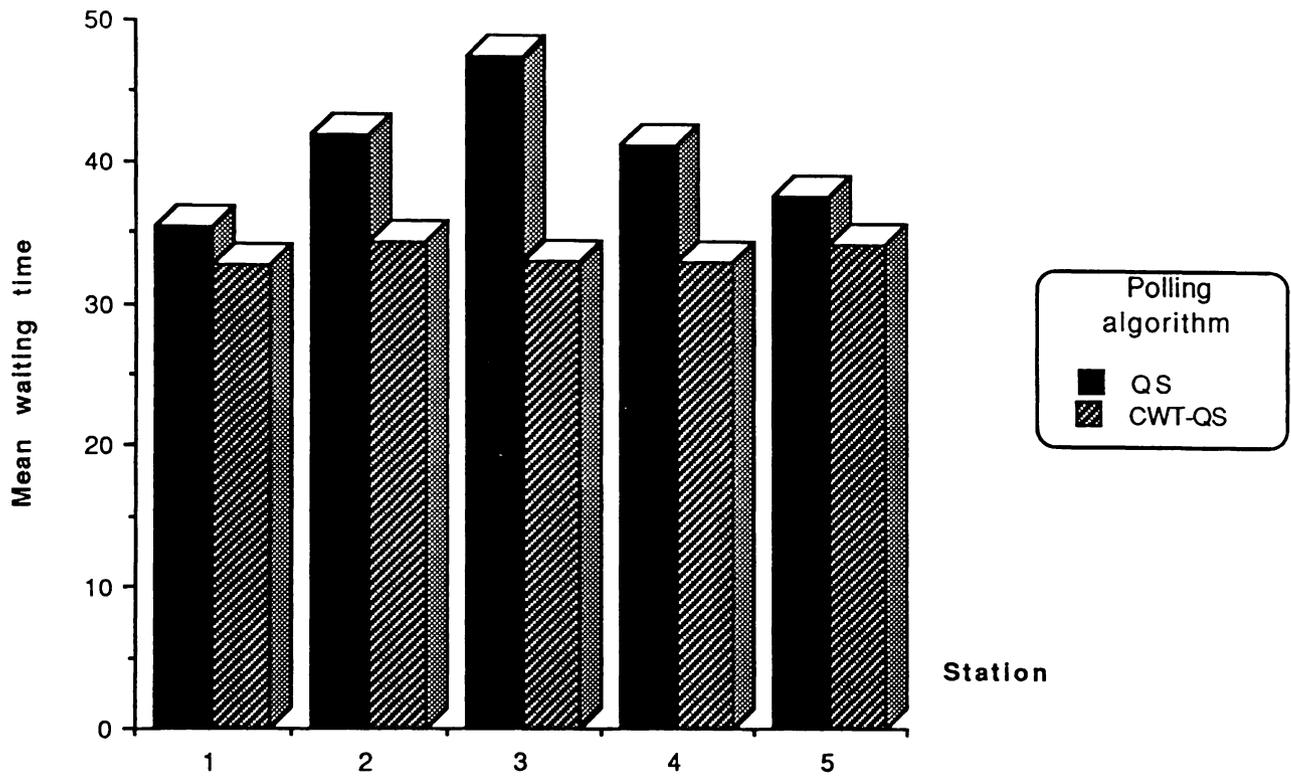


Fig 15. Mean waiting time at each station for the QS and CWT-QS policies traffic data set B.

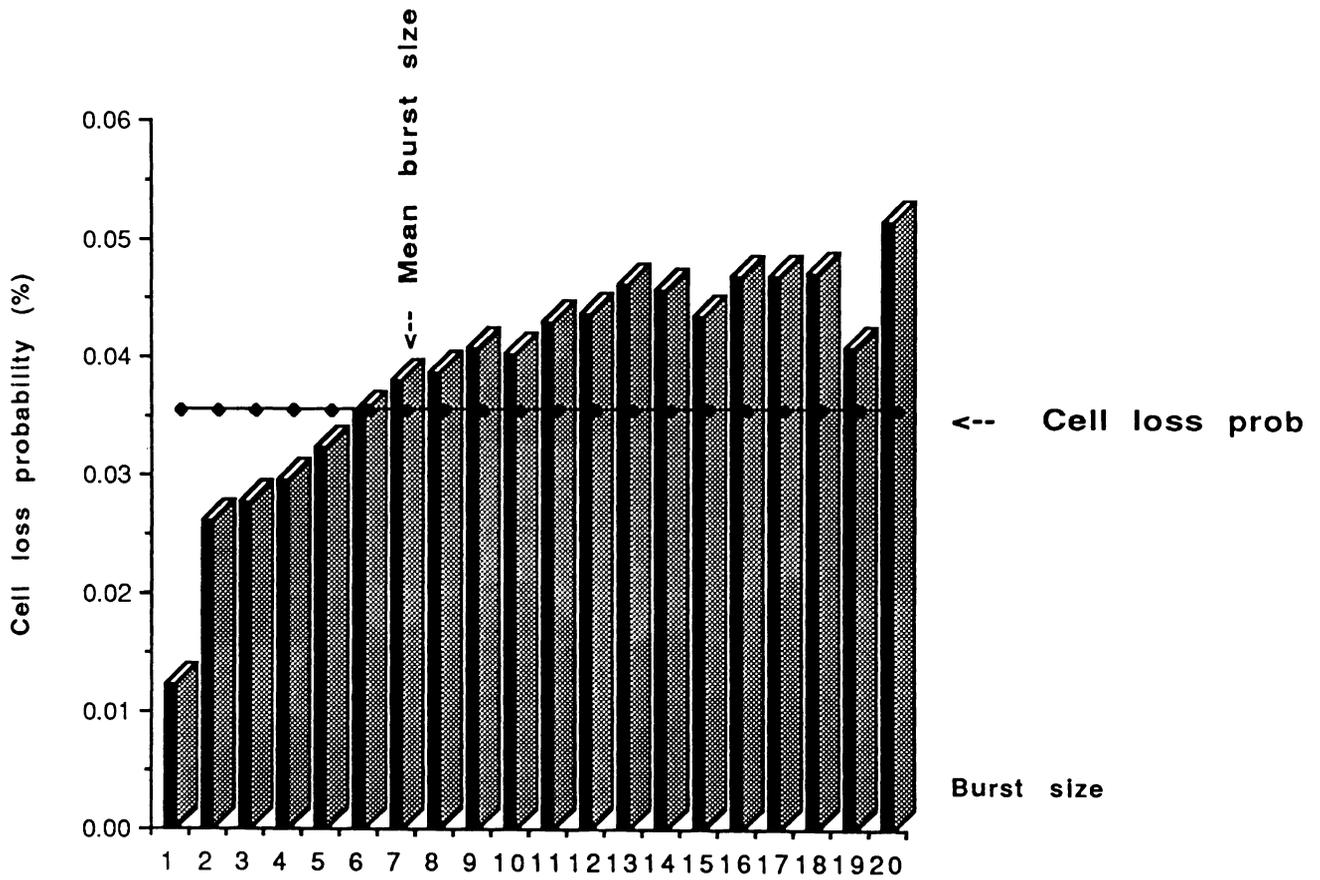


Fig 16. Cell loss probability for different burst sizes for station 1; traffic data set C.