On CBR Service

M. Grossglauser and S. Keshav
AT&T Bell Labs, Murray Hill, NJ 07974, USA
{grossgla, keshav}@research.att.com

Abstract

We investigate the performance of CBR traffic in the context of large-scale networks, where many connections and switches coexist and interact. We develop a framework for simulating such networks, decoupling the influence of breadth and depth. Our results are briefly as follows: we found that a Poisson stream is a good approximation to a superposition of many CBR streams with differing phases and bandwidths. Delays incurred by a reference stream with cross traffic composed of many CBR streams with different bandwidths and phases do not exceed a few cell times even under heavy load, which means that buildout buffers of 10 to 20 cells seem to be sufficient after traversing 20 switches. CBR traffic can be efficiently served by the First Come First Served (FCFS) scheduling discipline, which has the least implementation cost. Surprisingly, the Round Robin (RR) and Weighted Round Robin (WRR) disciplines perform worse than FCFS, despite their greater implementation complexity.

We also compare an analytical approximation method, based on the Multiclass Parametric Decomposition Method, with the simulation results and found it to be suitable to estimate end-to-end delays for the FCFS discipline.

1 Introduction

Constant Bit Rate (CBR) traffic is interesting for several reasons. Due to its simplicity, it will probably be the first service class offered by the BISDN. It is backwards compatible with circuit-switched networks, which means that a large user base already exists for this type of service: indeed, current video and audio codecs mostly produce CBR traffic. Former leased-line users wishing to switch to Virtual Private Networking (VPN) will use CBR service as well, emulating mostly T1 and T3 lines. In recent work, we have shown that a Renegotiated CBR service is well-suited for multiple-time-scale traffic [1]. Finally, CBR is easier to describe, handle and administer for both the user and the network than service types with more degrees of freedom, such as Variable Bit Rate (VBR).

However, there still are open issues that have to be addressed to build a large-scale network offering CBR service. For example, it is not clear yet what the effect of bunching or clustering is on a CBR stream as it travels through many switches, and what buildout buffer sizes at the endpoint are necessary to restore initial cell equispacing. Also, it is not clear what scheduling discipline among the obvious candidates, First Come First Served (FCFS), Round Robin (RR) and Weighted Round Robin (WRR) is best suited, especially when implementation complexity is taken into account. Finally, good analytical approximations for end-to-end delays and buildout buffer sizes for provisioning and call admission are still not available.

We address all these issues in this paper. The goal is to come up with a set of engineering guidelines that will support the deployment of CBR service. In Section 2 we give a survey of related literature. Section 3 discusses the metrics used, and Section 4 presents results from simulation and analysis. Section 5 concludes the paper.

2 Related Work

In this section we review past work in the quantitative analysis of CBR traffic. Roberts and Virtamo [2] and Dron, Ramamurthy and Sengupta [3] derive the queue size distribution for superposed CBR streams with identical periods ($nD/D/1$). Roberts and Virtamo [2] also present upper and lower bounds on this distribution for different periods ($\sum D_i/D/1$). However, they do not discuss the properties of the output process of the queue. Thus, their work does not extend to tandems of queues.

Matragi, Bisdikian and Sohraby [4, 5] derive analytical approximations for jitter for the single and the multihop case. However, their notion of jitter is based on the difference of the interdeparture process at subsequent switches and does not allow us to derive the end-to-end delay histogram, which is necessary to estimate buildout buffer sizes. Gruenenfelder [6] observes that the end-to-end delay of a reference connection going through multiple queues where it is multiplexed with cross traffic depends largely on the autocovariance of the latter. DeSimone [7] studies a network of three queues in tandem with one cross traffic stream and uses squared coefficients of variation as a measure of burstiness. He uses simulation and a simple analytic model to compare cell-level and packet-level FCFS and Round Robin (RR). Verma [8] has done multihop simulations, but his cross traffic is composed of CBR streams of equal bandwidths (among other cross traffic models). He observes bunching by measuring the change of the minimum inter-cell spacing after a number of hops. He does not discuss, however, the implications for end-to-end delay and the impact of correlated cross traffic.

Whitt [9] discusses how squared coefficients of variation can be applied to approximate open queueing networks, and also presents results for the case of multiclass networks [10]. Golestani [11] has shown by an example that for FCFS, bursts can be formed simply by superposing CBR streams.

To our knowledge, three aspects of our work are new. First, we study large-scale networks: some of our experiments involve as many as 600 cross traffic connections interacting with a reference connection. Second, we consider not only scheduling performance, but also attempt to consider the *implementation cost* of scheduling disciplines in relation with their performance. Finally, we give analytical expressions for the end-to-end delay of a reference connection with interference by cross traffic streams, and compare analytical results with detailed simulations.

3 Metrics

In this section, we present the metrics used to assess delay, clustering (or bunching) and implementation complexity.

3.1 Delay metrics

The 99%-percentile delay measures the delay necessary for reconstruction of the initial equal cell spacing through buffering (while allowing for 1% of cell loss). It also determines the amount of buffering needed at the end points, which translates directly into a hardware cost for memory.

3.2 Clustering metrics

We measure clustering of a cell stream with the *index of dispersion for intervals (IDI)* of its interarrival process, which provides us with the variability of cell arrivals on different time scales. The IDI for an interarrival process $\{X_i\}$ is defined as follows [12, 13]:

$$J_i(n) = \frac{n \cdot \mathrm{var}[\sum_{k=1}^{n} X_{i+k}]}{\mathrm{E}^2[\sum_{k=1}^{n} X_{i+k}]} \qquad (1)$$

If the process is assumed to be wide-sense stationary, so that the coefficient of variation is constant in the observed interval, then $J_i(n) = J(n)$ for all i. The function $J(n)$ describes the variation of arrivals over different time frames. Note that for a Poisson process, $J(n) = 1$ for $n = 1, 2, \ldots$, and for a CBR stream, $J(n) = 0$. Also note that $J(1)$ is the squared coefficient of variation c^2 of the stream. Furthermore, the IDI of point processes with positive correlation coefficients monotonically increase in n [12]. If $J(n)$ increases up to a certain $n = n_0$ and then remains constant for $n > n_0$, then we have correlation in the interarrival process up to a lag of n_0. Thus, the IDI is a powerful metric to measure burstiness over different time scales and to detect correlation in arrival processes.

3.3 Complexity metrics

As hardware implementation cost in VLSI directly depends on the required chip area, which in turn is dominated by *memory*, the cost of an ATM switch depends mostly on its state memory size. Thus, it makes sense to consider the amount of state a scheduling discipline requires. As the memory necessary to hold the cells itself is independent of how the cells are scheduled, we do not take it into account.

We assume that pointers consist of $n_p = \lceil \log_2(P) \rceil$ bits, where P is the number of words in the memory that can be addressed. Queues need a head and a tail pointer; each queue, therefore, accounts for $2n_p$ bits of state. Let N be the number of connections and Q the maximum size of a service quantum, with $n_q = \lceil \log_2(Q) \rceil$. Then, the state required for the scheduling disciplines we will study are given below.

Scheduling discipline	amount of state [bits]
FCFS	$2n_p$
RR	$N(3n_p + 1) + 2n_p$
WRR	$N(3n_p + n_q + 1) + 2n_p + n_q$

4 Results

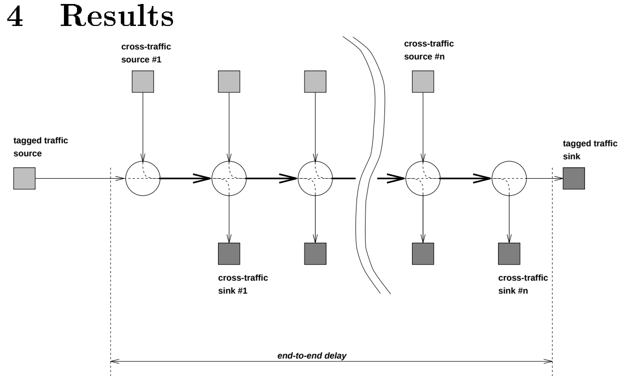

Figure 1: The *multihop* topology used to simulate the interactions with other streams and the resulting impact on the characteristics of a reference connection.

In our experiments, we focus on a single traffic stream (the *tagged connection*) and study how its characteristics are altered as it interacts with *cross traffic*. In most studies, the interpacket spacing T is the same for all streams. The superposition, then, has periodicity T, and maximum cell delay is T. Burstiness re-

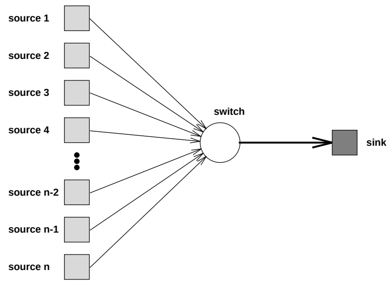

Figure 2: The topology used to study the characteristics of a CBR superposition process.

sults purely from the phase relationships between the constituent streams. If the phases are arbitrary and independent, then ensemble averages for the delay experienced by such a CBR stream can be derived [14, pages 117-119].

It is quite likely that in future ATM networks users will be allowed to choose the bandwidth they desire, with a fairly small granularity. Thus, bandwidths and phases are likely to be arbitrary, and we believe that this situation, though not well studied, is likely to be the dominant situation in the future.

We will first study the superposition of CBR streams to study the impact of *breadth* (Fig. 2), and then the impact of *depth* (Fig. 1) through the multihop topology. Then we extend the multihop experiment by replacing the CBR cross traffic with superpositions of CBR streams. This will allow us to draw conclusions about large networks, where each connection typically travels through many hops, and we will be able to compare the performance of two scheduling disciplines. We will compare this performance based on the 99%-percentile end-to-end delay and the mean end-to-end delay.

In all experiments, we compare the FCFS with the RR scheduling discipline. The main goal is to find out if FCFS performs well enough under the circumstances described above so that the additional implementation complexity of a discipline providing for explicit fairness between connections, such as RR, is not required.

Some simulation parameters that were common to all of the experiments below are given in the following table:

Simulation parameters:			
propagation delay	switching delay	cell size	buffer size
0	0	1000 bit	∞

4.1 Multiplexed CBR

We wanted a mix of high and low bandwidth streams in our study of multiplexing (Fig. 2). To achieve this, the total bandwidth of the incoming link is partitioned in the following way: We call *target bandwidth* the average bandwidth of the superposition stream (which will later be used as cross traffic). The first stream's bandwidth is a uniform random variable on the interval from 0 to half of the target bandwidth. The second stream's bandwidth is chosen in the same way between 0 and the remaining bandwidth, i.e. the target bandwidth minus the first stream's bandwidth, and so on. The last stream gets the remaining bandwidth such that all streams sum up to the target bandwidth. Our results are averaged over several bandwidth partitions.

Simulation parameters:	
Incoming line speed	1 MB/s
Outgoing line speed	100 Mb/s
Target bandwidth	0.8 Mb/s

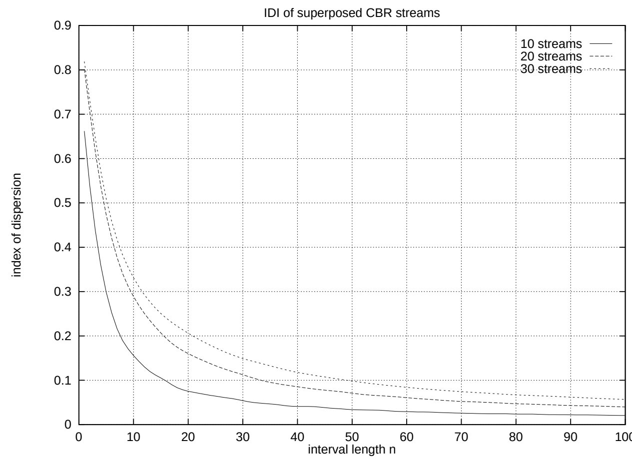

Figure 3: The IDI for the superposition of a number of CBR streams with arbitrary bandwidths.

The fairly large IDI at a time frame of a couple of cells (for comparison, a Poisson process has an IDI of 1) indicates considerable variability on the outgoing link. As the observed interval increases (n large), the arrival variability decreases (which means that the arrival process does not have a positive correlation function). Furthermore, we observe that the IDI function increases over the entire range of n as more connections are multiplexed together.

It is important to note that this phenomenon has nothing to do with the scheduling discipline, and any work-conserving discipline would show the same results. The nonzero index of dispersion of the interarrival process on the outgoing link is a pure result of the *superposition* of the incoming CBR streams.

These results suggest that a superposition of CBR connections carried by one link, for example in a backbone network, cannot be replaced by an "equivalent" CBR stream with a bandwidth equal to the sum of the bandwidths of the constituent streams when assessing scheduling performance. Furthermore, the order of magnitude of the IDI function for small n suggests the use of a Poisson process as an approximation of the cell arrival for the study of scheduling performance when delays in the order of a few cells are of importance.

4.2 Multihop CBR with CBR cross traffic

In the following experiment, we look at a CBR stream going through 21 switches, sharing each link with another CBR stream (Fig. 1). The link utilization was set to a fairly high value ($\rho = 0.8$). The tagged stream's bandwidth was then varied over a range of values (and the cross traffic accordingly, to keep total utilization at ρ). We were interested in determining how the CBR stream changes as it moves through the network.

Simulation parameters:	
Line speed	1 Mb/s
Traffic intensity	0.9
Tagged traffic intensity ρ	$0.10 \ldots 0.80$
Cross traffic intensity ρ_T	$0.80 \ldots 0.10$
End-to-end transmission delay	22ms

The transmission delay is composed as follows: 20 cell service times on the links connecting switches, one cell service time on the tagged stream's access link, and one cell service time on the link going into the sink.

The cell delay pattern is highly dependent on the phases of the two streams. For example, if the two streams have the same bandwidth allocation, then three situations can occur, depending on the phase difference between the two streams: (1) no cells are delayed; (2) every tagged cell is delayed; (3) every cross traffic cell is delayed. As we do not wish to assume anything about the phase relationship between the connections, we would like to measure ensemble averages over all possible phases. We approximate this in our simulation by introducing arbitrary phase changes in intervals of length $\gg T$. Another reason for the interest in ensemble averages is the possibility of clock drifts between the user and the network, which would destroy this periodic delay patterns.

As the end-to-end delay increases almost linearly with the hopcount (Fig. 4) means that the bunching of tagged cells seems not to be important. If bunching did occur and accumulate as the tagged connection goes through more and more hops, the delay would not be a linear function in n. We will make this same important observation later with Poisson cross traffic in Section 4.4.

We can see that the two scheduling disciplines have comparable performance over the entire range of bandwidth partitions (Fig. (5)). Interestingly, the end-to-end delay of the tagged connection decreases as its share of the link bandwidth increases. This is because, on average, a cell belonging to the tagged stream encounters a nonempty queue at a switch with a (ensemble) probability ρ_{CT}, the link utilization of the cross traffic. As the cross traffic bandwidth decreases, the tagged stream has a higher chance of finding queues empty and experiences less delay, even as its own bandwidth increases.

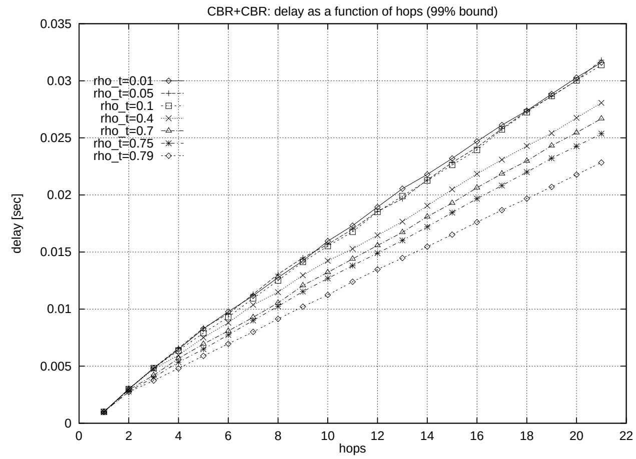

Figure 4: The 99-percentile delay as a function of the number of hops the cell goes through (FCFS).

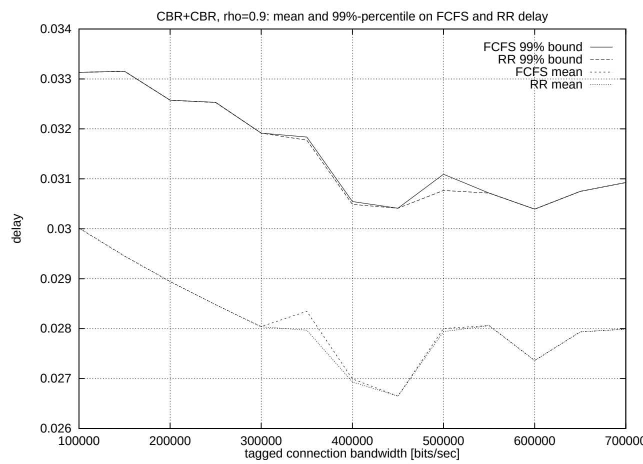

Figure 5: The end-to-end delays (mean and 99%-percentile) for FCFS and RR scheduling for different bandwidth partitions. The total link utilization ρ was kept constant at 0.9. The values shown are ensemble averages over phase combinations.

4.3 CBR with aggregated CBR cross traffic

We now go a step further by replacing the cross traffic with a superposition of CBR streams. The bandwidths are chosen as described in Section 4.1 above. We repeated the experiment with cross traffic consisting of 10, 20 and 30 superposed CBR streams. The simulation parameters are as follows [1].

Simulation parameters:	
Trunk line speed	1 Mb/s
Cross traffic access line speed	100 Mb/s
Traffic intensity ρ	0.9
Tagged traffic intensity ρ_T	0.1
Cross traffic intensity ρ_{CT}	0.8
End-to-end transmission delay	22ms

Note that we intentionally did not use the same trace (representing the same bandwidth partition) for each cross traffic, but a different one at each switch.

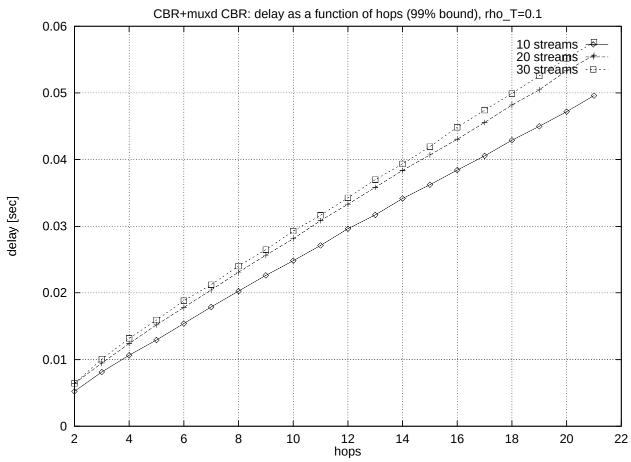

Figure 6: The 99-percentile delay as a function of the number of hops (FCFS).

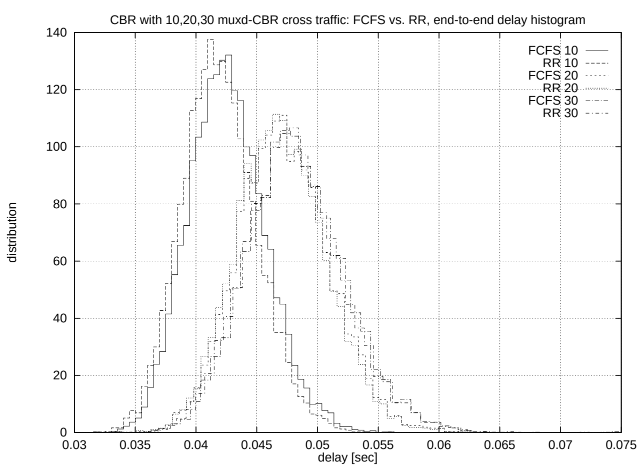

Figure 7: The end-to-end delay histogram with cross traffic consisting of 10, 20 and 30 CBR streams

Again, the delay is almost linear in the number of hops (Fig. 6), and bunching does not adversely affect the delay experienced by tagged cells.

The difference between the two scheduling disciplines is very small (Fig.7). Note that in this experiment, and unlike in the Poisson cross traffic experiment to be discussed in Section 4.4, the cross traffic is composed of a mix of "thick" and "thin" streams. The property of the RR discipline to allocate bandwidth fairly between all streams can therefore result in lower delays for the tagged stream if the cross traffic contains thick streams, as seen in Fig. 7. In the Poisson experiments, we will assume that infinitely thin streams compose the cross traffic.

As a rule of thumb we can observe that the mean delay in both cases is approximately one cell service time per switch, which is about the delay that would be seen on average if the cross traffic was purely CBR (the ensemble average over phases) at the same link utilization and a "thin" tagged stream. The 99%-delay-percentile is about 2 cell service times per switch, which is tolerable in a high speed cell-based network. For example, in the context of ATM with its 53 bytes cells, two cell service times correspond to $2.74\mu s$ on a 155 Mb/s link.

Both the 99%-percentile and the histogram width of the end-to-end delay increase as the number of multiplexed CBR streams increases. In the next experiment, we try to find an upper bound on the delay by using an conservative approximation for the cross traffic, based on the observations in Section 4.1.

For an analytic discussion of this problem (limited to one switch), the reader is referred to [14, pages 122-129].

4.4 CBR with Poisson cross traffic

In this section, we replace the cross traffic stream in experiment 4.3 with Poisson arrival processes, based on the results of the multiplexing experiment. This experiment must be viewed as a limiting case, when the number of streams constituting the cross traffic tends to infinity [15]. The results found in this section will be *conservative* and have also been used for the analysis of an *M+D/D/1* queue in [14, page 129]. Switches implement one of the FCFS or RR disciplines.

Simulation parameters:	
Line speed	1 Mb/s
Traffic intensity ρ	0.9
Tagged traffic intensity ρ_T	$0.01 \ldots 0.1$
Cross traffic intensity ρ_{CT}	$0.89 \ldots 0.8$
End-to-end transmission delay	22 ms

Fig. 8 shows the 99-percentile delay for different tagged-stream bandwidths after 21 hops. FCFS and RR have comparable delay for small tagged stream intensity. As this intensity increases, the RR delay gets worse, while the FCFS delay stays approx-

[1] The *trunk lines* are the lines interconnecting switches; the *cross traffic access lines* are the ones connecting the cross traffic sources to the switches. The cross traffic line speed has been chosen higher in order to emulate multiple incoming lines of the same speed as the trunk lines

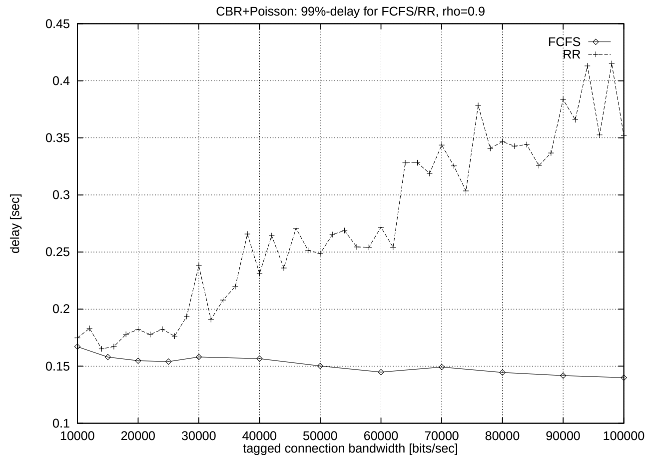

Figure 8: The 99%-percentile of the measured end-to-end delay after going through 21 switches for FCFS and RR.

imately constant. The largest delay is about 170ms for FCFS. Taking into account the 22 ms of end-to-end transmission delay means that FCFS has about $(170\text{ms}-22\text{ms})/20 \approx 7.4\text{ms}$ or 7.4 cell service times of queueing delay per switch. The size of the buildout buffer can also be estimated from Fig. 8: For example, when the tagged connection has 100kb/s bandwidth, we observe a histogram width of about 90ms for FCFS, which corresponds to a buildout buffer size of approximately $90\text{ms} \times 100\text{kb/s} \times (1/\text{cell size}) = 9$ cells. When the tagged connection has 10kb/s bandwidth, the histogram width is about 150ms (170ms-22ms), corresponding to 1.5 cells.

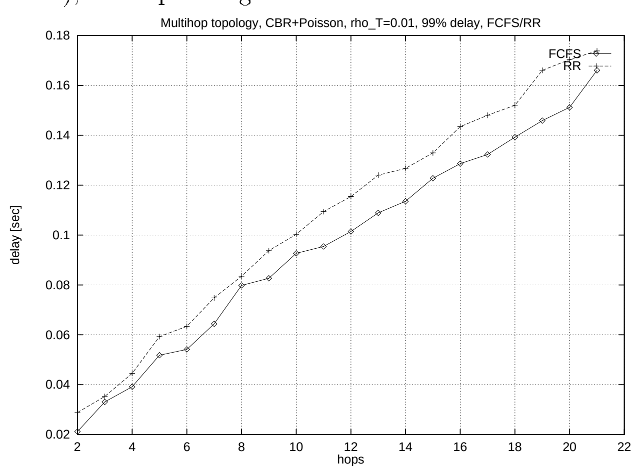

Figure 9: The 99%-delay-percentile as a function of the number of hops the tagged stream goes through, with $\rho_T = 0.01$.

From Fig. 9, we see that the delay is again approximately linear as a function of the hopcount, as observed previously. Bunching still has no notable effect on the delays.

We now state, without proof, a theorem that confirms that the delay of each tagged cell in the RR case in greater or equal to the delay in the FCFS case (the proof has been omitted due to lack of space). This theorem is based on the thin-cross-traffic assumption made earlier, i.e. that the per-connection queues of the cross traffic streams contain at most one cell at any time. The arrival and departure times of a cell i are called $a(i)$ and $d(i)$, respectively.

Theorem 1 *For given arrival process $\mathcal{A} = \{a(i) | i \in \mathcal{I}\}$ the queueing delay $d(i) - a(i)$ of any tagged cell i in the RR case is greater or equal to the delay experienced in the FCFS case, if the RR per-connection queues for the cross traffic never contain more than one cell.*

Intuitively, RR attempts to share the link equally between all connections. If the tagged connection has a higher bandwidth than the cross traffic streams, then this connection experiences a degradation in performance. Clearly, this is not desirable.

A way of explicitely giving each connection the desired share of bandwidth is by using Weighted Round Robin (WRR). Unfortunately, this discipline has the drawback of implicit bunching. WRR allocates a certain service quantum to each connection. When this connection's identifier is at the head of the service list, then up to `service_quantum` number of cells are served. As these cells leave the WRR scheduler back-to-back, although they might have arrived with some spacing between them, some bunching is introduced. This bunching is implicit to the WRR scheduling discipline.

4.5 An Analytical Approach for FCFS delays

In this section, we present an analytical approach that can be used to estimate the delay experienced by the tagged connection in the FCFS case with Poisson cross traffic.

Our model is based on work by Whitt that gives accurate results even if the deterministic traffic is not negligible and if the number of switches is small [9, 10]. Whitt's methods allow us to derive approximate relationships for the coefficients of variation of the tagged stream at each switch in the network, and to derive the expected wait time at each switch. Each (wide sense stationary) interarrival process $\{X_i\}$ can be associated with its *squared coefficient of variation (SCV)* c_a^2, which is defined as

$$c_a^2 = \frac{\mathrm{var}[X_i]}{\mathrm{E}^2[X_i]} \quad (2)$$

As the interdeparture process ($c^2_{d,T}$) of the tagged stream at switch i is at the same time the interarrival process of the tagged stream at switch $i+1$, we calculate the SCVs for the multihop topology recursively from left to right. The variables used for one segment are depicted in Fig. 10.

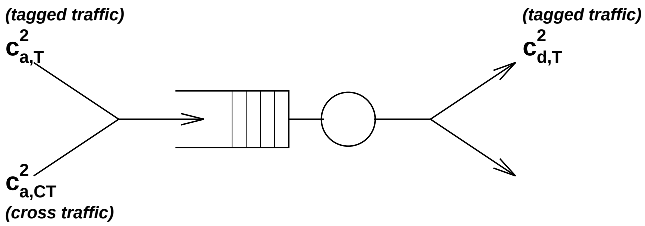

Figure 10: The model used to approximate one segment in the multihop topology.

Equation (11) in [10] (below) allows us to determine the departure SCV of the tagged cells. Note that the service time SCV c^2_s is zero in our case, as cells are of constant size. ρ is the total traffic intensity, and $\lambda_{T,CT}$ are the intensities of the tagged and the cross traffic streams, respectively. $p_1 = \lambda_T/(\lambda_T + \lambda_{CT})$ is the fraction of cells belonging to the tagged stream. The cross traffic is a Poisson process, and therefore has SCV $c^2_{a,CT} = 1$.

$$c^2_{d,T} = \rho^2 p_T c^2_s + (2-\rho)p_T(1-p_T)c^2_{a,CT} + [(1-p_T)^2 + (1-\rho^2)p^2_T]c^2_{a,T} \quad (3)$$

If we denote the arrival SCV at switch i with $c^2_{a,T}(i)$, then we obtain the recursion

$$c^2_{a,T}(i+1) = (2-\rho)p_T(1-p_T) + [(1-p_T)^2 + (1-\rho^2)p^2_T]c^2_{a,T}(i) \quad (4)$$

(with $c^2_{a,T}(1) = 0$), which allows us to determine the departure SCV at each switch.

Once the $c^2_{a,T}$ are known at each switch, we need to determine the SCV of the *superposed* interarrival process at each switch. We use (29), (30) and (33) in [9].

$$c^2_a = \omega\left(\frac{\lambda_T c^2_{a,T} + \lambda_{CT} c^2_{a,T}}{\lambda_T + \lambda_C T}\right) + 1 - \omega \quad (5)$$

$$\omega = \frac{1}{1 + 4(1-\rho)^2(v-1)} \quad (6)$$

$$v = 1 + 2\left(\frac{\lambda_T \lambda_{CT}}{\lambda^2_T + \lambda^2_{CT}}\right) \quad (7)$$

This allows us to determine the expected wait time at each switch, given by (44) and (45) in [9]. τ is the mean cell service time (1ms in our case).

$$\mathrm{E}[W] = \frac{\tau\rho(c^2_a + c^2_s)g}{2(1-\rho)} \quad (8)$$

where

$$g(\rho, c^2_a, c^2_s) = \begin{cases} \exp\left[-\frac{2(1-\rho)}{3\rho}\frac{(1-c^2_a)^2}{c^2_a + c^2_s}\right] & c^2_a < 1 \\ 1 & c^2_a \geq 1 \end{cases} \quad (9)$$

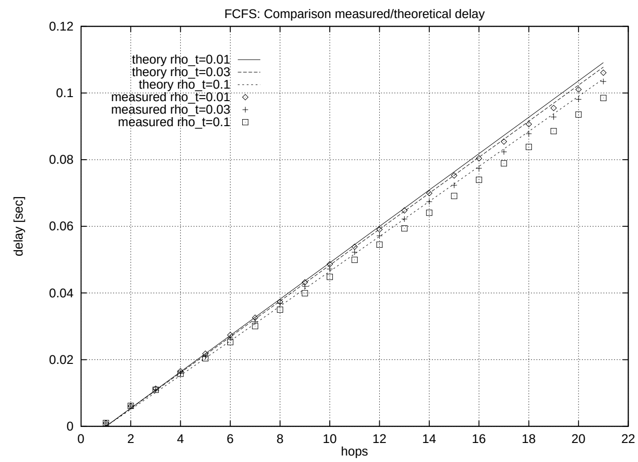

Figure 11: Comparison of analytic and measured delay as a function of number of hops for different tagged stream intensities ρ_T.

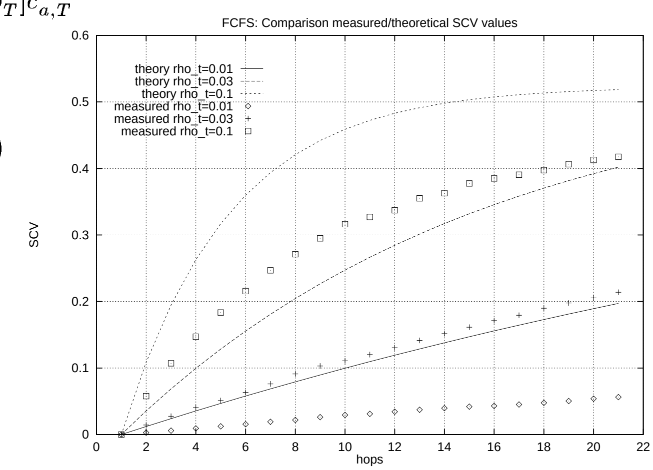

Figure 12: Comparison of analytic and measured SCV as a function of number of hops for different tagged stream intensities ρ_T.

Over the whole range of tagged traffic intensities, the mean delay is predicted very well by Whitt's method (Fig. 11). However, for small ρ_T, we observe considerable differences between the SCV values predicted by theory and the ones measured. The approximation tends to improve as ρ_T increases. We conclude that Whitt's Parametric Decomposition method, at least in our setting, gives excellent, slightly conservative, mean delay estimates.

5 Conclusion

We have carried out a detailed study of the behavior of CBR traffic in the presence of cross traffic, using available analytical techniques as well as exhaustive simulation. Earlier work had indicated [11] that multiplexing of such sources could lead to bunching - in our work, we have characterized the bunching using the IDI and 99% delay metrics. We have also compared simulation results with the results predicted by multiclass parametric decomposition techniques proposed by Whitt. Finally, we compare the behavior of the FCFS and RR disciplines with respect to bunching of CBR traffic.

Our main results are as follows. First, our simulations indicate that when CBR traffic is multiplexed with a small number of up to 30 other CBR streams, the difference in end-to-end delays for FCFS and RR are small (see Sections 4.2 and 4.3), and there is no appreciable bunching, even as the depth in the network increases to 20 hops (Fig. 4 and Fig. 6).

Second, as the number of heterogeneous CBR sources multiplexed at a single link (the *breadth*) increases, the output stream looks like a Poisson stream at smaller time scales (2-10 cell arrival times, Fig. 3). Given this observation, we model the aggregate cross traffic by a Poisson stream, which is admittedly a pessimistic modeling, in order to find conservative delay estimations.

Third, when a CBR source interacts with Poisson cross traffic, the scheduling discipline has an impact on the output stream (Fig. 8). When the scheduling discipline is FCFS, the output stream experiences some bunching and delays, which increases as the load or the depth in the network increases. We have found that with FCFS, very reasonable buildout buffer sizes are sufficient. As a rule of thumb, for a connection using one tenth of the trunk bandwidth, less than a cell of buffer space per hop is required to keep losses below 1%. The performance of RR is comparable to FCFS for low intensity streams, but deteriorates as a connection's intensity increases.

Fourth, in all three multihop experiments, the end-to-end delay grows linearly in function of the number of hops. This suggests that bunching of cells belonging to the tagged stream is not sufficient to have an impact on delays experienced by these cells (Fig. 9). This means that for CBR streams, no special measures have to be taken to avoid bunching, such as per-connection reregulation at intermediate switches. It is sufficient to restore cell-equispacing in a buildout buffer at the connection endpoint.

Fifth, we have found that the expected waiting delays for FCFS scheduling are well predicted by theory (Fig. 11), although the predicted squared coefficient of variation differs substantially from the observed values (Fig. 12).

Our observations about FCFS's superiority are reinforced by implementation complexity of the FCFS and RR scheduling disciplines, which is dominated by the amount of state that has to be maintained. To get an idea about the relative cost of these implementations, we present the cost given some typical values of the parameters. We assume that the output queue can serve 16K conversations at any time, so that $N = 16K$. Let the largest service quantum be 256 quanta, so that $n_q = 8$. We assume that the queue has 16 MB of memory, which is cell addressable, so that we can store approximately 39570 cells, so that $n_p = 16$ (otherwise, with byte addressable memory, $n_p = 24$). Thus, the state information required for FCFS is $2n_p$ or 32 bits (48 bits with byte addressing). The state information for WRR is 116 Kbytes with cell addressing and 165 Kbytes with byte addressing. The state information with RR is 100 Kbytes with cell addressing and 149 Kbytes with byte addressing. Thus, FCFS not only performs better than RR (and WRR), but requires much less state information.

To the best of our knowledge, this is the first comprehensive simulation study that confirms the hypothesis that the FCFS scheduling discipline is sufficient for CBR traffic even in a large-scale network. This has previously been claimed (for example [16]), but not been verified. Based on our work, we recommend that CBR networks be built with FCFS scheduling, with a few cells of buffering per switch, and a buildout buffer proportional to the number of hops in the path.

A limitation of this study is that the 99%-percentile on delay corresponds to a relatively large cell loss rate of 1% in the buildout buffer. Future ATM networks will have to offer considerably smaller error probabilities, e.g. 10^{-6}. In order to do experimental work with loss probabilities in this range, one has to resort to more advanced techniques, e.g. rare event simulation [17]. A second limitation is that we assume that incoming CBR traffic is smooth at the cell level. If this traffic is not smooth at this level, a regulator would have to be introduced at the network entrance to perform smoothing. We do not consider the smoothing delay at this regulator, but it is easily computed as the worst case buildup over the averaging interval of the CBR source.

6 Acknowledgments

Discussions with E. Biersack, S.P.Morgan and W. Whitt have been an important contribution to this work. We thank them for many helpful suggestions. W.Whitt has helped with applying the Multiclass Parametric Decomposition Method to this particular problem. C. Lund provided us with accurate random number sequences for calibrating our measurements. The technique of choosing connection bandwidths in Section 4 is due to C. Lund, S. Phillips and N. Reingold. Our thanks to them.

References

[1] M. Grossglauser and S. Keshav and D. Tse, "RCBR: A Simple and Efficient Service for Multiple Time-Scale Traffic," in *Proc. ACM SIGCOMM '95*, (Boston, Mass.), pp. 219–230, August 1995.

[2] J. W. Roberts and J. T. Virtamo, "The Superposition of Periodic Cell Arrival Streams in an ATM Multiplexer," *IEEE Transactions on Communications*, vol. 39, pp. 298–303, February 1991.

[3] L. G. Dron, G. Ramamurthy, and B. Sengupta, "Delay Analysis of Continuous Bit Rate Traffic Over an ATM Network," *IEEE Journal on Selected Areas in Communications*, vol. 9, pp. 402–407, April 1991.

[4] W. Matragi, C. Bisdikian, and K. Sohraby, "Jitter Calculus in ATM networks: Single Node Case," in *IEEE INFOCOM '94*, (Toronto, Canada), June 1994.

[5] W. Matragi, C. Bisdikian, and K. Sohraby, "Jitter Calculus in ATM networks: Multiple Node Case," in *IEEE INFOCOM '94*, (Toronto, Canada), June 1994.

[6] R. Gruenenfelder, "A Correlation Based End-to-End Cell Queueing Delay Characterization in an ATM Network," in *Proc. Thirteenth International Teletraffic Congress (ITC-13)*, vol. 15, (Copenhagen, Denmark), pp. 59–64, North-Holland Studies in Telecommunications, June 1991.

[7] A. DeSimone, "Generating Burstiness in Networks: A Simulation Study of Correlation Effects in Networks of Queues," in *ACM Computer Communication Review*, pp. 24–31, 1991.

[8] D. C. Verma, *Guaranteed Performance Communication in High Speed Networks.* PhD thesis, University of California, Berkeley, Berkeley, CA 94720, December 1991. Report No. UCB/CSD 91/663.

[9] W. Whitt, "The Queueing Network Analyzer," *The Bell System Technical Journal*, vol. 62, pp. 2779–2815, November 1983.

[10] W. Whitt, "Towards better multi-class parametric-decomposition approximations for open queueing networks," *Annals of Operations Research*, vol. 48, pp. 221–248, 1994.

[11] S. J. Golestani, "Congestion-Free Transmission of Real-Time Traffic in Packet Networks," in *Proc. IEEE INFOCOM '90*, (San Francisco, California), pp. 527–536, June 1990.

[12] R. Gusella, "Characterizing the Variability of Arrival Processes with Indices of Dispersion," *IEEE Journal on Selected Areas in Communications*, vol. 9, pp. 203–211, February 1991.

[13] K. Fendick and W. Whitt, "Measurements and Approximations to Describe the Offered Traffic and Predict the Average Workload in a Single-Server Queue," *Proceedings of the IEEE*, vol. 77, pp. 171–194, January 1989.

[14] J. Roberts, ed., *COST 224: Performance evaluation and design of multiservice networks.* Commission of the European Communities: information technologies and sciences, October 1991.

[15] E. Cinlar, "Superposition of Point Processes," in *Stochastic Point Processes: Statistical Analysis, Theory, and Applications* (P.A.W.Lewis, ed.), pp. 549–606, Wiley Interscience, 1972.

[16] S. P. Morgan, "personal communication," 1994.

[17] V. S. Frost and B. Melamed, "Traffic Modeling For Telecommunications Networks," *IEEE Communications Magazine*, pp. 70–81, March 1994.