ELSEVIER

Computer Networks and ISDN Systems 27 (1995) 521–535

COMPUTER NETWORKS and ISDN SYSTEMS

Comparative performance of interleaved and non-interleaved pipelining in ATM terminal adapters

Anurag Kumar [a], Robert G. Cole [b,*]

[a] *Department of Electrical Communication Engineering, Indian Institute of Science, Bangalore-560 012, India*
[b] *AT&T Bell Laboratories, Holmdel, New Jersey 07733, USA*

Accepted 10 November 1993

Abstract

In this paper we compare the end-to-end delay performance of two service disciplines that an Asynchronous Transfer Mode (ATM) multiplexor can use to multiplex pipelined synchronous native protocol frames arriving over low speed access lines onto higher speed ATM trunks. The ATM multiplexor must convert native protocol frames on the access lines (typically hundreds of bytes in length) into ATM cells on the trunks (of fixed length 53 bytes) at the network access and the corresponding ATM (de)multiplexor must reassemble these multiple ATM cells into a bit continuous synchronous native protocol frame at the network egress.

In one service discipline, the multiplexor (hereafter referred to as a terminal adaptor, or TA) pipelines one synchronous frame at a time onto the high speed trunk, waiting, if it needs to, for the successive ATM cells from this one synchronous frame to accumulate. This scheme introduces a larger queueing delay at the ingress TA but no playout delay is required at the egress TA. In the other scheme the TA interleaves the ATM cells from the various synchronous frames arriving concurrently over the low speed lines. An interleaving scheme has a smaller queueing delay, but introduces a large variability in the network insertion times of the ATM cells corresponding to a native protocol frame. Hence the egress TA must enforce a delay before playing out the original bit-synchronous native frame over a low speed line, so as to increase the probability that the ATM cells are available when required to properly reassemble the native synchronous frame into a bit continuous stream.

Using an analytical model for the first scheme and a simulation model for the second scheme we analyze the end-to-end delay of a synchronous frame, and compare the mean end-to-end delay for various trunk speeds (TA-to-network switch) to "access" line speed ratios. We conclude that, for the cases studied in this paper, if this ratio is greater than or equal to 4 then, over the range of useful loading, interleaved pipelining typically yields smaller mean delays than noninterleaved pipelining. If the ratio is less than 4 then interleaving typically yields larger mean delays for native protocols due to increased playout delays at the egress TA.

Keywords: ATM multiplexor; M/D/1 queue with gradual input

* Corresponding author.

SSDI 0169-7552(93)E0101-J

1. Introduction

International standards bodies, e.g., CCITT, define *terminal adaption* [3] as the function of converting end user, non-standard, native protocols, e.g., High-Level Data Link Control (HDLC) protocol, to standards based network access protocols, e.g., Broadband-Integrated Services Digital Network (B-ISDN) standards. More generally, terminal adaption functions are performed at network access points, which must wrap user's native protocols inside wide-area-network (WAN) protocols for transport.

An example of this is shown in Fig. 1, where the WAN is based on the Asynchronous Transfer Mode (ATM) technology, the user native protocol is frame relay (FR), and the terminal adaptor (TA) is located at the network interface point. On access, the native protocol frames are wrapped, or enveloped, in the B-ISDN frame format, which is ATM at layer 2, and the ATM cells from multiple access ports on the TA are multiplexed onto network trunks. The TA function which performs this wrapping appears to provide the end system a virtual private line (VPL). In other words, the native protocol is not terminated but is simply wrapped and the native frames, including control frames, pass transparently through the network. The network relies on the end system's capability to recover from errors.

The TAs on access into the network break up the large native frames (e.g., a 552 byte Internet

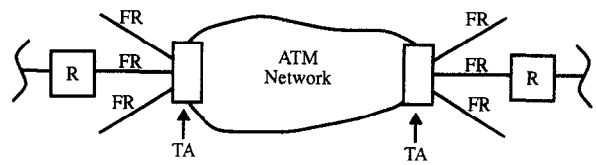

Fig. 1. Terminal adaption functions as an integral part of non-ATM protocol, e.g. frame relay (FR), wrapping for transport over an ATM network. Here the TA is multiplexing FR traffic from multiple customer premises routers (R).

Protocol (IP) packet encapsulated in a frame relay) into smaller (e.g., 53 byte) ATM cells (see Fig. 2). As soon as enough bytes to constitute an ATM cell have been accumulated from a native frame (indicated by the dashed vertical lines intersecting the horizontal line labeled "ingress line" in Fig. 2), the access TA can wrap them into an ATM cell and ship this packet over the designated virtual circuit (defined as streaming mode service by CCITT Recommendation I.363 [4]). When these ATM cells arrive at the (egress) TA at the other side of the virtual circuit, this TA has a choice in playing out the native frame over the egress line. It can choose to accumulate the full native frame or it can choose to begin transmission of the native frame upon receipt of the first ATM cell. In the former choice, the network edge acts in a store-and-forward fashion (indicated by the line labeled "egress line without pipelining" in Fig. 2), while in the latter the network performs egress pipelining (indicated by the line labeled "egress line with pipelining" in

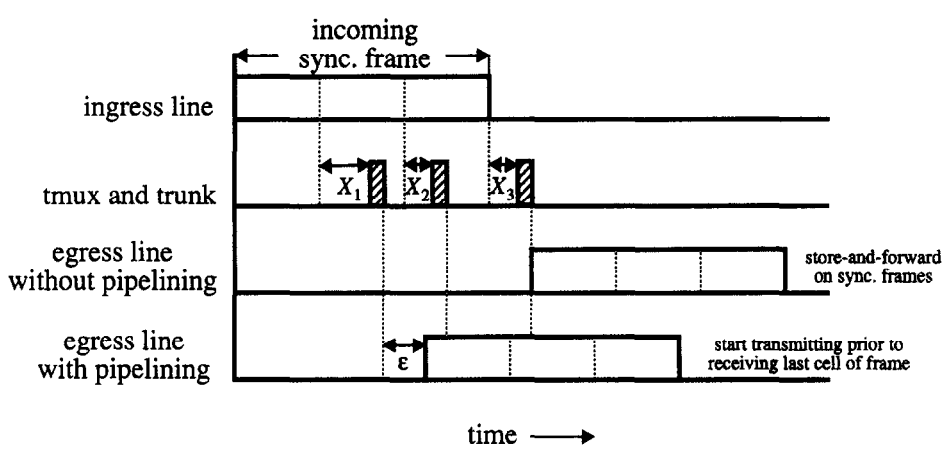

Fig. 2. ATM wrapping and unwrapping of a synchronous frame with and without pipelining.

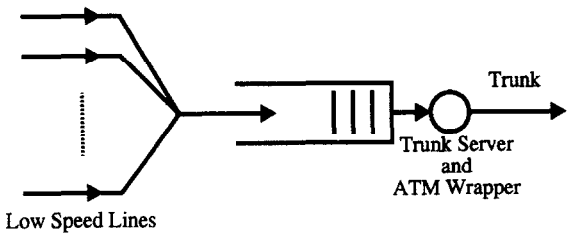

Fig. 3. Schematic of multiplexer with noninterleaved operation.

Fig. 2). Egress pipelining can dramatically improve the end-to-end delay perceived by the user, see Refs. [5] and [9].

Egress pipelining reduces the transport time, but the risk of *synchronization* loss between the egress TA and the customer premise equipment now exists for bit-synchronous protocols. Due to variations in the network transport times of the ATM cells (the X_is in Fig. 2), the possibility exists that the egress TA will have exhausted transmission of the "already arrived bits" of a native protocol frame *prior* to the arrival of the next ATM cell of the native frame. Thus, an apparent gap appears in the native frame on the egress line initiating an error recovery procedure due to the synchronization loss. To reduce the likelihood of such a synchronization loss, the egress TA can hold up the first ATM cell for a short period of time before initiating transmission of the native frame on the egress line [7]. This allows the trailing ATM cells to catch up. This imposed time delay is referred to as a *playout delay* (denoted by ϵ in Fig. 2 and defined in Section 2).

The playout delay can be chosen so as to meet a target probability of synchronization loss. The playout delay required to achieve a desired probability of synchronization loss (P_ϵ), depends on the multiplexing strategy used by the TA, on the virtual circuit characteristics and on the network load. An adaptive playout scheme dynamically adjusts ϵ, based on real-time measurements on the virtual circuit, so as to maintain the desired synchronization loss probability (see, for example, Ref. [10]). Hereafter, we shall assume that the value of the playout, ϵ, is such that it ensures a probability of synchronization loss of 10^{-3}.

Assuming that the ingress TA pipelines the synchronous protocol frames over the outgoing trunk, there are two ways in which the multiplexor can multiplex ATM cells from the various incoming synchronous native frames.

(1) *Non-interleaved multiplexing* (FIFO trunk server, see Fig. 3): Frames arriving over the low speed lines are queued in their order of arrival in a commonly shared buffer. The trunk is assigned to the frame at the head of the queue; note that this maybe a partial frame and the remainder of it may still be arriving over the low speed line. As the requisite number of bytes from the frame become available, ATM cells are formed (i.e., the header, CRC, etc., are attached) and put onto the trunk. Observe that in this multiplexing scheme, ATM cells from different synchronous frames are not interleaved. Thus if the frame at the head of the FIFO server has not yet fully arrived then the trunk will stay idle while waiting for the successive ATM cells to accumulate.

(2) *Interleaved multiplexing* (Round Robin or FCFS trunk server, see Fig. 4): As the frames arrive over the low speed lines they are wrapped into ATM cells. These ATM cells are queued up, and then served by the trunk server. There are

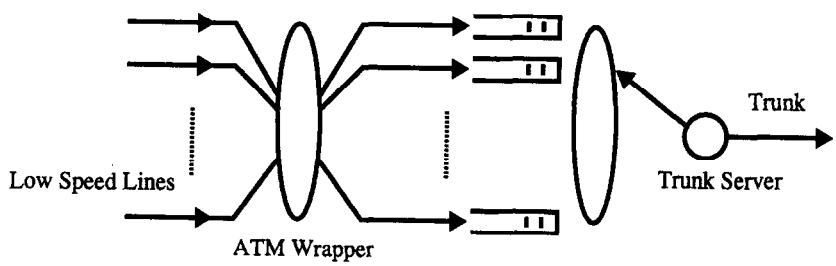

Fig. 4. Schematic of multiplexor with interleaved operation and round robin service.

several alternatives for queueing and serving the ATM cells. The ATM cells from each line may be queued in their order of arrival in a per line buffer, and then served in a round robin fashion by the trunk server. This is depicted in Fig. 4. Observe that in such multiplexing schemes, ATM cells from different synchronous frames will be interleaved. Thus the trunk need not stay idle waiting for the successive ATM cells from a single frame, but can start serving frames from other lines. With this service discipline, however, there is a probability that when ATM cells from a given synchronous frame arrive at the TA at the other end of the virtual circuit, they may be separated in time more than when they arrived over the low speed line. A playout delay as described previously is therefore required before the TA at the other end can start playing out the unwrapped frames over the outgoing synchronous line.

An alternative interleaved multiplexing scheme is implemented with a FCFS trunk server. Here the synchronous frames are wrapped as they arrive from the low speed access line. The ATM cells from each line are queued in a common buffer (as opposed to per line buffering as depicted in Fig. 4 above). The trunk server serves the ATM cells in a FCFS fashion. A playout delay is also required for this FCFS interleaved multiplexing scheme. However, the value of the playout will be different for the two interleaving multiplexors.

In this paper we are principally interested in the performance differences between the TA multiplexing schemes. We analyze the end-to-end delay of a synchronous frame, defined as the time between the first bit of the frame entering a TA at one end and the last bit of the frame being delivered by the TA at the other end. We limit ourselves to the simplest ATM virtual circuit topology of one trunk connecting the two TAs (see Fig. 5).

If the virtual circuit is composed of several trunks and network switches then the end-to-end delay in either multiplexing scheme will be increased by: (i) the queueing and insertion delays at the tandem switches and trunks; and (ii) a build-out delay to allow for ATM cell delay variability caused by queueing in the network. If there is a separation between the various trunk speeds, then the queuing and playout delays will be dominated by the slowest trunk. Therefore, we ignore these additional network queueing delays in the analysis and their effects on the required playout delay and focus our attention on the simplified virtual circuit model in Fig. 5. See Refs. [5,9,12] for discussion and illustration of these network performance affects.

The rest of this paper is organized as follows. In Section 2 we present an analytical model for the non-interleaved scheme and describe a simulation model for the interleaved schemes. In Section 3 we present some results obtained from these models, and in Section 4 we summarize our conclusions. The Appendix contains the details of the analysis of the non-interleaved scheme.

2. Performance analysis of the two multiplexing schemes

After setting down some notation, in Section 2.1 we present an analytical model for the non-interleaving scheme and an expression for mean end-to-end delay obtained from its analysis. Analysis of the interleaving schemes, particularly the calculation of the playout delay, is much more

Fig. 5. ATM virtual circuit topology.

difficult. If we denote by X_i the multiplexor queuing delay of the ith ATM cell of a synchronous frame ($1 \leqslant i \leqslant N$), then it is easily observed from Fig. 2 that for no synchronization loss on this frame, $\epsilon \geqslant X_i - X_1$, for $2 \leqslant i \leqslant N$. But, of course, $\epsilon \geqslant 0$, and hence the playout delay is given by $\epsilon = \max_{1 \leqslant i \leqslant N}(X_i - X_1)$. For a synchronization loss probability, P_ϵ, of 10^{-3}, ϵ should be chosen to be the 99.9th percentile of the stationary distribution of $\max_{1 \leqslant i \leqslant N}(X_i - X_1)$. The difficultly arises from the fact that the X_is are not independent. We have, therefore, resorted to simulation models to obtain the mean end-to-end delays for the interleaving schemes. We discuss these simulation models in Section 2.2.

We first set down some notation:

s_l = line speed
s_t = trunk speed
r = s_t/s_l
b_l = mean time taken to transmit a synchronous frame over a low speed line
b_t = mean time taken to transmit a synchronous frame over the trunk (without any ATM headers)
N = number of ATM cells in each synchronous frame
h_t = time taken to transmit a ATM header over the trunk
ϕ_t = $(Nh_t)/b_t$ = fractional trunk transmission overhead due to ATM wrapping
W_{int} = end-to-end delay of a native synchronous frame with interleaving
W_{nonint} = end-to-end delay of a native synchronous frame with non-interleaving
λ = aggregate arrival rate of synchronous frames to the multiplexor
ρ_t = $\lambda(b_t = Nh_t)$ = trunk occupancy in one direction
n_l = number of access lines into the multiplexor

Here, the end-to-end delay is defined as the time between the arrival of the first bit of the synchronous frame at the access TA to the transmission of the last bit from the egress TA.

2.1. Analysis of non-interleaving multiplexor

To obtain the mean end-to-end delay of a synchronous frame with the non-interleaving

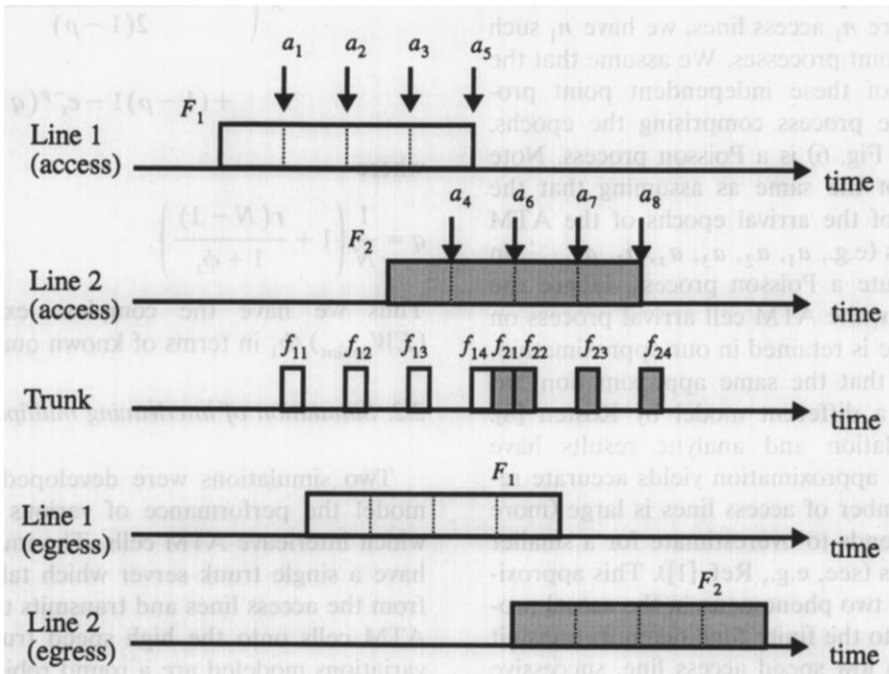

Fig. 6. Noninterleaved multiplexing of synchronous frames (dotted lines indicate cell boundaries).

scheme we proceed as follows. Since the trunk is allocated to a frame only after the arrival of a number of bits sufficient to comprise a ATM cell, we shall consider the arrival epoch of the first ATM cell boundary as the arrival epoch (e.g., the epochs a_1, a_4 in Fig. 6) of the synchronous frame to our multiplexing delay model (described in the following).

This will yield (see Fig. 6)

$$EW_{\text{nonint}} = \frac{b_l}{N} + EW^{(1)}_{\text{nonint}} + \frac{b_t}{N} + h_t + b_l,$$

where

$W^{(1)}_{\text{nonint}}$ = waiting time of the first ATM cell of a synchronous frame, after its arrival to the FIFO queue.

Normalizing this mean end-to-end delay to b_l, we get

$$\frac{EW_{\text{nonint}}}{b_l} = 1 + \frac{1}{N}\left[1 + \frac{1}{r}(1+\phi_t)\right] + \frac{EW^{(1)}_{\text{nonint}}}{b_l}.$$

We obtain $EW^{(1)}_{\text{nonint}}$ analytically. Consider the point process comprising the epochs of arrival of the cell boundary of the first ATM cell in the successive native protocol frames on an access line. If there are n_l access lines, we have n_l such independent point processes. We assume that the superposition of these independent point processes (i.e., the process comprising the epochs, a_1, a_4, etc., in Fig. 6) is a Poisson process. Note that this is not the same as assuming that the superposition of the arrival epochs of the ATM cell boundaries (e.g., a_1, a_2, a_3, a_4, a_5, a_6, ..., in Fig. 6) constitute a Poisson process. Hence the bursty nature of the ATM cell arrival process on each access line is retained in our approximation. We point out that the same approximation has been used in a different model by Kosten [8]. Further, simulation and analytic results have shown that the approximation yields accurate results if the number of access lines is large (more than 20) and tends to overestimate for a smaller number of lines (see, e.g., Ref. [1]). This approximation ignores two phenomena in the actual process: (i) owing to the finite time taken to transmit a frame over a low speed access line, successive frames on a line arrive at least b_l time units after each other; and (ii) since there is a finite number of lines, if all lines are transmitting frames, no more frames can arrive; this is the finite source effect. The Poisson assumption makes analysis of the model tractable. We also assume that the frame lengths are deterministic, that each synchronous frame can be broken up into exactly N ATM cells, and that protocol processing time is negligible.

With the above assumptions, $W^{(1)}_{\text{nonint}}$ becomes the waiting time in a variation of the well known M/D/1 FIFO model [6]. *The variation is that a customer does not arrive to the queue all at once, but gradually.* If the customer that is arriving is at the head of the queue then the server is allocated to the customer until it finishes arriving and is fully served. The waiting time (i.e., time between the customer beginning to arrive and beginning to get served) can be analyzed via level-crossing analysis (see Ref. [2]). The analysis is detailed in the Appendix and yields

$$\frac{EW^{(1)}_{\text{nonint}}}{b_l} = -\frac{(1+\phi_t)}{\rho_t r} \times \left(\frac{\rho_t^2(2q-3) + 2\rho_t(1-q)}{2(1-\rho)} + (1-\rho)1 - e_t^{-\rho}(q-1) \right),$$

where

$$q = \frac{1}{N}\left(1 + \frac{r(N-1)}{1+\phi_t}\right).$$

Thus we have the complete expression for $(EW_{\text{nonint}})/b_l$ in terms of known quantities.

2.2. Simulation of interleaving multiplexors

Two simulations were developed in order to model the performance of various multiplexors which interleave ATM cells. The multiplexors all have a single trunk server which takes inputs n_l from the access lines and transmits the individual ATM cells onto the high speed trunk. The two variations modeled are a round robin server providing limited service (serving a single ATM cell

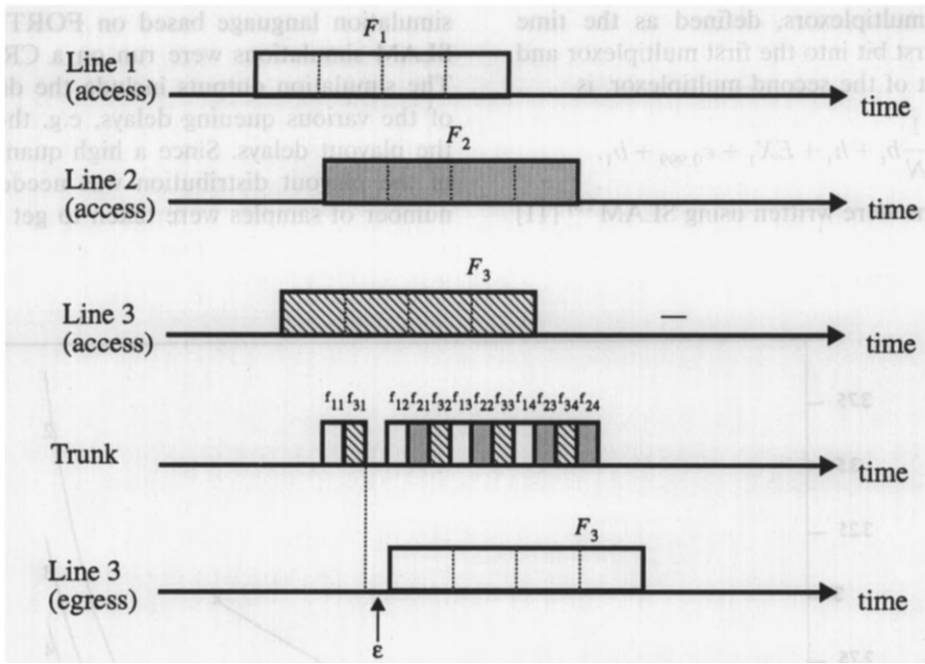

Fig. 7. Interleaved multiplexing of synchronous frames (dotted lines indicate ATM cell boundaries).

per visit to each access line queue, see Fig. 7) and a FCFS server (no per-access-line queues). Both multiplexors accept input from n_1 identical access lines with speed s_1 and multiplex these onto a single output trunk with speed $s_t \geq s_1$. The assumed arrival process on each access line is shown in Fig. 8. Here the line activity is modeled as an alternating sequence of idle and active periods. Idle periods are distributed exponentially with mean λ_1^{-1} during which no inputs (i.e. ATM cells) arrive. Active periods are deterministic in length, during which N identical ATM cells arrive, equally spaced in time. The fixed duration of the active periods is b_1.

The key simulation outputs are EX_1, the mean queuing delay of the first ATM cell of any synchronous frame, and $\epsilon_{0.999}$, the 99.9th percentile of the stationary distribution of $\max_{1 \leq i \leq N}(X_i - X_1)$ (i.e. the playout distribution). Given EX_1 and $\epsilon_{0.999}$, then the mean delay (see Fig. 7) through

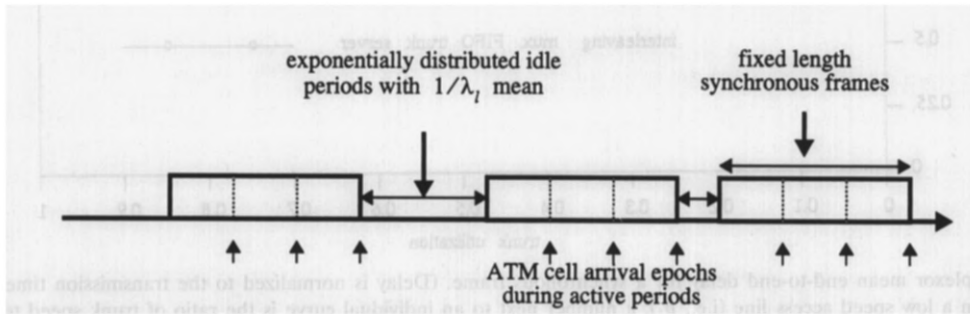

Fig. 8. Arrival process on the access lines into the statistical multiplexor for the simulation results.

back-to-back multiplexors, defined as the time between the first bit into the first multiplexor and the last bit out of the second multiplexor, is

$$EW_{\text{int}} = \frac{b_1}{N} + \frac{1}{N} b_t + h_t + EX_1 + \epsilon_{0.999} + b_1.$$

The simulations were written using SLAM™ [11] simulation language based on FORTRAN. The SLAM simulations were run on a CRAY XMP. The simulation outputs include the distributions of the various queuing delays, e.g. the X_is, and the playout delays. Since a high quantile (0.999) of the playout distribution was needed, a large number of samples were taken to get reasonably

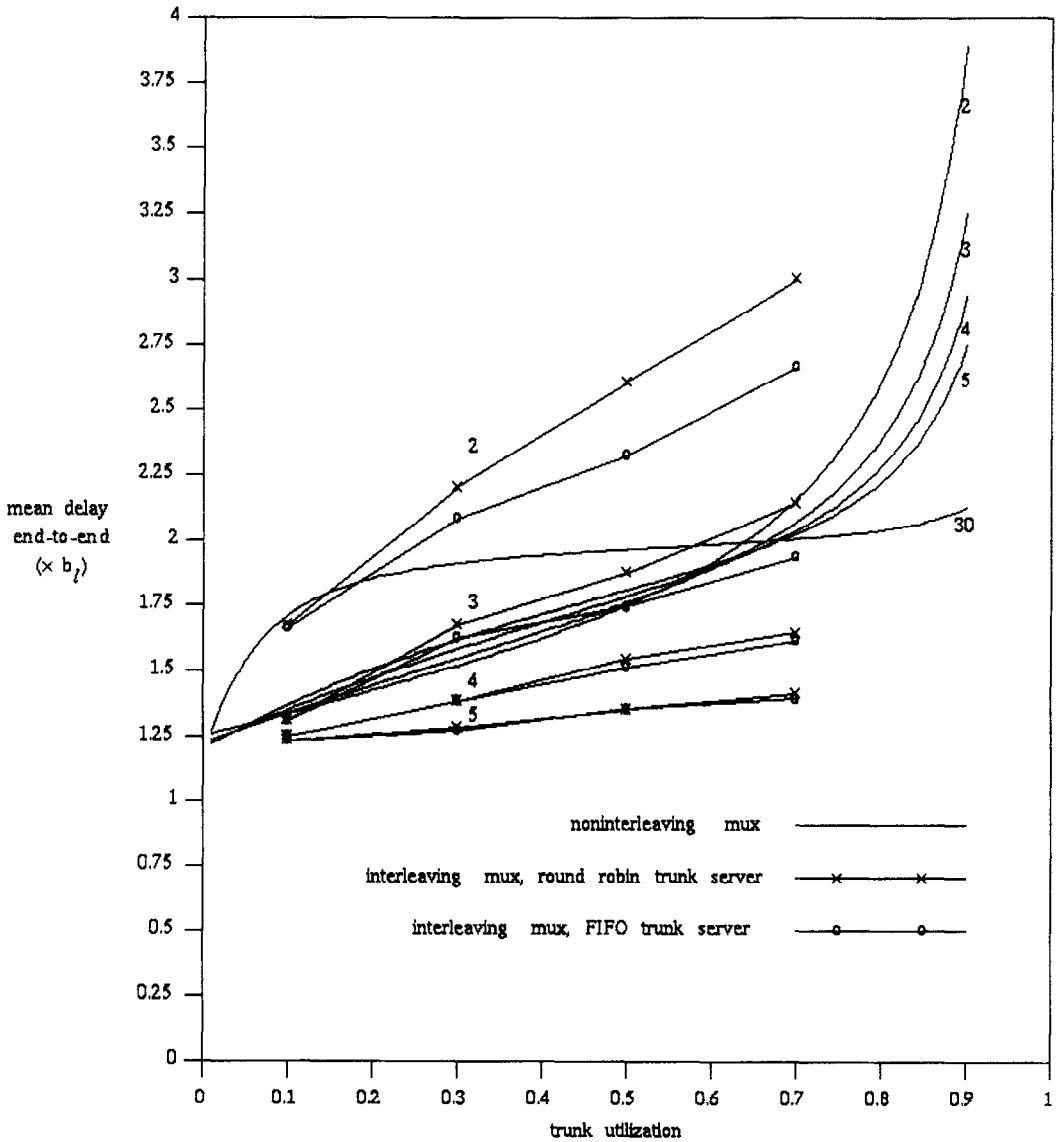

Fig. 9. Six line multiplexor mean end-to-end delay for a synchronous frame. (Delay is normalized to the transmission time of a synchronous frame on a low speed access line (i.e., b_1); a number next to an individual curve is the ratio of trunk speed to line speed for that curve.)

narrow confidence intervals for $\epsilon_{0.999}$. In order to get intervals that lie within $(\epsilon_{0.9985}, \epsilon_{0.9995})$, we need 10^5 samples. Hence, each simulation was run 10_5 until synchronous frames had passed through the system. Multiple runs of each configuration were performed and the averages taken. Each run is characterized by the multiplexor type, n_l, r, N, b_l, and ρ_t. Assuming that the header length, h_t, is zero, the trunk occupancy is given by

$$\rho_t = \frac{n_l}{r}\left(\frac{b_l/\lambda_l^{-1}}{1 + b_l/\lambda_l^{-1}}\right),$$

where λ_l^{-1} is the mean duration of the access line idle periods (see Fig. 8). The results of the simulations are discussed in the next section.

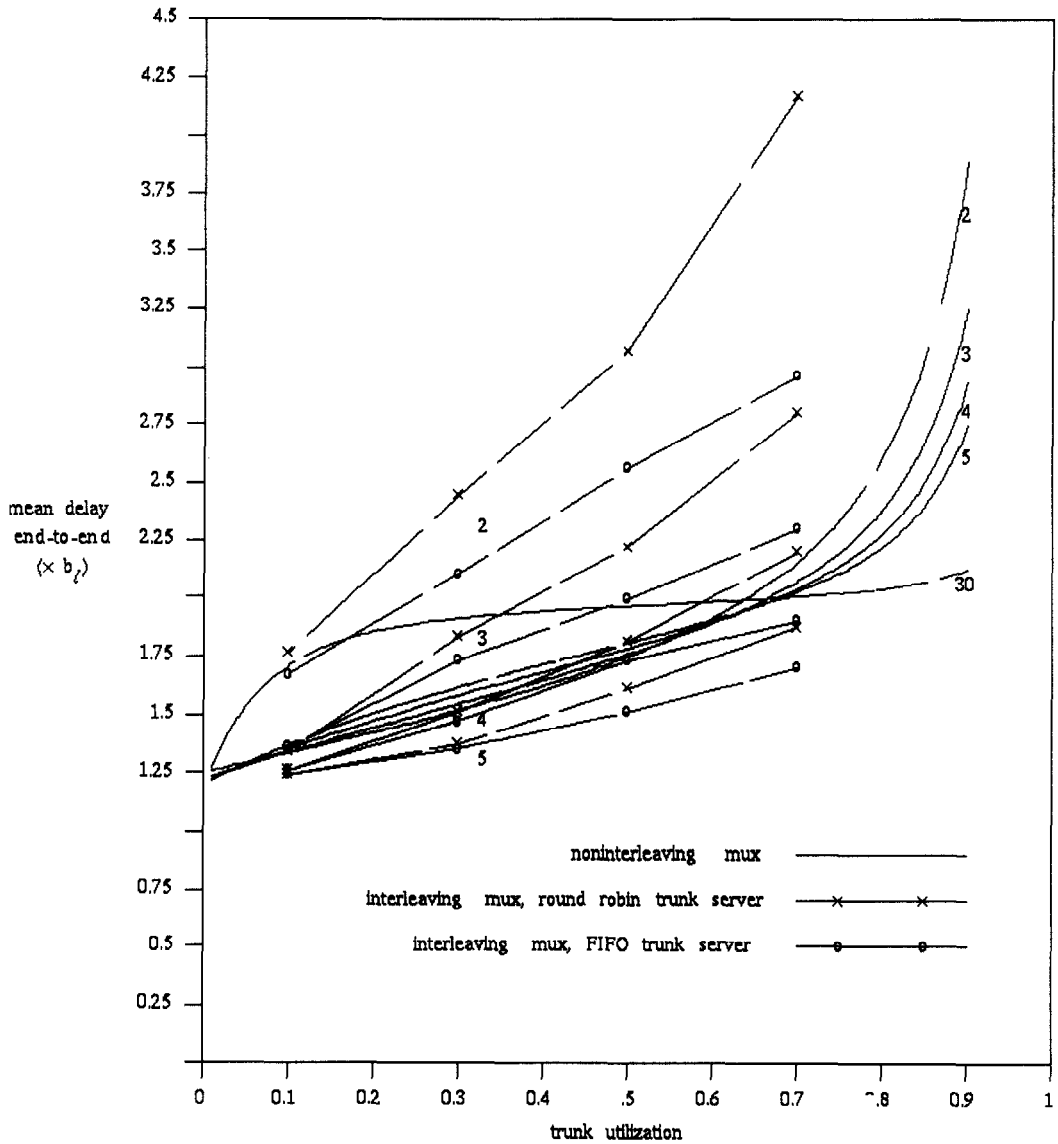

Fig. 10. Ten line multiplexor mean end-to-end delay for a synchronous frame. (Delay is normalized to the transmission time of a synchronous frame on a low speed access line (i.e. b_l); a number next to an individual curve is the ratio of trunk speed to line speed for that curve.)

3. Results

The analytical results for the mean end-to-end delay as a function of the trunk utilization for the non-interleaved case are presented in Figs. 9 and 10. Delay curves for $r = 2$, 3, 4, 5, and 30 are presented (note that $r = 30$ corresponds to, for example, $s_l = 1.544$ Mbps and $s_t = 45$ Mbps). These curves show an unusual shape, as highlighted in the $r = 30$ case, but also found in the other cases as well. At low trunk utilizations, the non-interleaving system demonstrates a behavior similar to an M/D/1 system with service time dependent on the slower access line speed. At high trunk utilizations, this case demonstrates behavior similar to an M/D/1 system but with a service time dependent on the faster trunk speed. At intermediate loads a cross-over occurs resulting in a flattening of the load–delay curve which is more pronounced as r increases.

The simulation results are presented in Tables 1 through 4. Evident in all the simulation results is the fact that $\epsilon_{0.999}$ is extremely sensitive to trunk utilization, increasing in some cases by as much as an order of magnitude as the load increases from 0.1 to 0.7. When $\epsilon_{0.999}$ becomes comparable to the time required to accumulate the entire synchronous frame, then pipelining is no longer a beneficial feature. Therefore, a measure of the pipelining efficiency is $\alpha_\epsilon \equiv (N/(N-1))\epsilon_{0.999}/b_l$. When α_ϵ is less than unity, then pipelining improves performance, otherwise it is of no benefit. The effectiveness of pipelining in improving performance diminishes as r decreases, and in general shows little if any benefits for $r < 4$ except for low trunk utilizations. However, the cutoff value of $r = 4$ is not sharp and depends on many factors, including utilization, number of access lines, multiplexing schemes, etc.

Increasing the number of access lines affects the playout values. Tables 1 and 2 present results for six line multiplexors with round robin and FCFS servers, respectively. Tables 3 and 4 present the comparable ten line multiplexor results.

Table 1
Simulation results on the six line round robin multiplexor

r	b_l	ρ_t	EX_1	$\epsilon_{0.999}$	EW_{int}/b_l	α_ϵ
1	107	0.1	18.6	178	3.00	2.0
		0.3	21.7	278	3.97	3.1
		0.5	29.6	373	4.93	4.2
		0.7	54.8	448	5.87	5.0
2	213	0.1	18.6	89	1.67	0.5
		0.3	21.1	200	2.20	1.1
		0.5	25.9	280	2.60	1.6
		0.7	40.3	350	3.00	2.0
3	107	0.1	6.4	10	1.31	0.1
		0.3	7.1	48	1.67	0.5
		0.5	8.3	69	1.87	0.8
		0.7	11.3	96	2.14	1.0
4	427	0.1	18.6	18	1.25	0.1
		0.3	20.6	72	1.38	0.2
		0.5	23.5	137	1.54	0.4
		0.7	29.1	174	1.64	0.5
5	183	0.1	6.4	6	1.23	0.0
		0.3	7.0	14	1.28	0.1
		0.5	7.9	26	1.35	0.2
		0.7	9.3	36	1.41	0.2
6	640	0.1	18.6	18	1.22	0.0
		0.3	20.6	36	1.26	0.1
		0.5	23.1	51	1.28	0.1
		0.7	26.5	62	1.30	0.1

Table 2
Simulation results on the six line FIFO multiplexor

r	b_l	ρ_t	EX_1	$\epsilon_{0.999}$	EW_{int}/b_l	α_ϵ
1	107	0.1	19.0	144	2.69	1.6
		0.3	24.9	197	3.24	2.2
		0.5	39.8	249	3.87	2.8
		0.7	83.3	283	4.59	3.2
2	213	0.1	18.6	87	1.66	0.5
		0.3	21.6	173	2.08	1.0
		0.5	29.8	216	2.32	1.2
		0.7	55.5	262	2.66	1.5
3	110	0.1	6.4	9	1.31	0.1
		0.3	7.1	43	1.62	0.5
		0.5	8.7	54	1.74	0.6
		0.7	13.5	70	1.93	0.8
4	427	0.1	18.6	18	1.25	0.1
		0.3	20.6	71	1.38	0.2
		0.5	23.7	122	1.51	0.3
		0.7	30.8	160	1.61	0.5
5	183	0.1	6.4	6	1.23	0.0
		0.3	7.0	12	1.27	0.1
		0.5	7.9	25	1.35	0.2
		0.7	9.3	31	1.39	0.2
6	640	0.1	18.6	18	1.22	0.0
		0.3	20.5	32	1.25	0.1
		0.5	23.1	42	1.27	0.1
		0.7	26.4	47	1.28	0.1

Table 3
Simulation results on the ten line round robin multiplexor

r	b_1	ρ_t	EX_1	$\epsilon_{0.999}$	EW_{int}/b_1	α_ϵ
2	213	0.1	18.7	107	1.76	0.6
		0.3	21.5	250	2.44	1.4
		0.5	27.0	379	3.07	2.1
		0.7	46.4	593	4.17	3.3
3	110	0.1	6.4	13	1.34	0.1
		0.3	7.2	66	1.83	0.7
		0.5	8.8	107	2.22	1.2
		0.7	13.9	166	2.80	1.8
4	427	0.1	18.7	20	1.26	0.1
		0.3	21.0	127	1.51	0.4
		0.5	25.4	249	1.81	0.7
		0.7	36.2	405	2.20	1.1
5	183	0.1	6.4	7	1.24	0.0
		0.3	7.2	30	1.37	0.2
		0.5	8.5	73	1.61	0.5
		0.7	11.5	117	1.87	0.8

For light loads the number of access lines have little effect on the value of the playout values. For the higher loads, however, increasing the number of lines from six to ten in many cases caused the playout values to double. Increasing the number of access lines allows for the possibility that more synchronous frames will be simultaneously interleaved, therefore increasing $\epsilon_{0.999}$.

The playout value even shows significant sensitivity to the service discipline, i.e. round robin versus FCFS. By examining possible frame arrival patterns on access lines to the ATM multiplexor,

Table 4
Simulation results on the ten line FIFO multiplexor

r	b_1	ρ_t	EX_1	$\epsilon_{0.999}$	EW_{int}/b_1	α_ϵ
2	213	0.1	18.7	89	1.67	0.5
		0.5	33.5	264	2.56	1.5
		0.7	69.3	312	2.96	1.8
3	110	0.1	6.4	15	1.36	0.2
		0.3	7.3	54	1.73	0.6
		0.5	9.8	81	1.99	0.9
		0.7	19.1	106	2.30	1.2
4	427	0.1	18.7	20	1.26	0.1
		0.3	21.1	109	1.47	0.3
		0.5	26.6	216	1.73	0.6
		0.7	44.0	270	1.90	0.8
5	183	0.1	6.4	7	1.24	0.1
		0.3	7.2	26	1.35	0.2
		0.5	8.6	55	1.51	0.4

it becomes clear that $\epsilon_{0.999}$ (round robin) $\geqslant$ $\epsilon_{0.999}$(FCFS), since the configurations of active and idle access lines that contribute to the non-zero part of the playout distribution for FCFS also contribute for round robin. However, by examining frame arrival patterns which contribute to the non-zero part of the playout distribution for the round robin scheme, some of these configurations of active and idle access lines do not (also) contribute to the non-zero part of the playout distribution for the FCFS scheme. For example, when $r = 2$ and three access lines are simultaneously active, then a playout will be required. For the FCFS scheme the playout at egress must compensate for the fact that between each ATM cell of a given native frame two ATM cells will arrive from the other active lines. However, for the round robin scheme, more than two ATM cells may be interleaved between the ATM cells of a given native frame in the event that a fourth access line becomes active before the transmission on the trunk of the earlier frames has completed. This forces the playout of the round robin scheme to be larger than the FCFS scheme to account for this additional degree of interleaving. However, in no instance will the converse be true, that the FCFS scheme will induce a greater degree of interleaving than round robin. Therefore, we conclude that $\epsilon_{0.999}$(round robin) $\geqslant \epsilon_{0.999}$ (FCFS) as illustrated in the simulation results.

Finally, in Figs. 9 and 10 the results for non-interleaving and interleaving are compared graphically for the six and ten line multiplexors, respectively. The performance of the non-interleaving multiplexor is less sensitive to r than the performance of the corresponding interleaving multiplexor. For large r interleaving is preferred and for small r non-interleaving is preferred as anticipated. The cross-over value is around $r = 4$. We recall that the approximate analysis for the non-interleaving case yields an overestimate of the mean delay. Hence we can say that, for the cases studied, when ATM cell-interleaving should be employed. From Figs. 9 and 10, when $r \geqslant 4$ little difference between round robin and FCFS is observed. The decision concerning the service discipline should then be based upon other con-

siderations, e.g. fairness in overload or complexity.

4. Conclusions

We have presented two alternate strategies for ATM cell wrapping and multiplexing of synchronous traffic for transport over a virtual circuit in a B-ISDN ATM-based network. The first strategy, called non-interleaving, transmits ATM cells from one synchronous frame consecutively onto the trunk. The alternative strategy, called interleaving, interleaves ATM cells from multiple synchronous frames onto the trunk. When the customer protocol is bit or byte-synchronous, comparison of analytical results for the non-interleaving case and simulation results for the interleaving case show that, for the configurations studied in this paper, interleaving minimized end-to-end delays when r, the ratio of the trunk to access line speed, was greater than or equal to approximately 4. Otherwise, non-interleaving can significantly reduce mean transport delays.

Appendix A – Derivation of $EW^{(1)}_{\text{nonint}}$

For notational simplicity, we shall first analyze an idealized version of the model; a simple interpretation of the parameters of the idealized model will yield $EW^{(1)}_{\text{nonint}}$. As discussed in Section 2.1, the idealized model can be thought of as an M/D/1 model in which a customer does not arrive to the queue all at once, but gradually. The server is allocated to customers in a FCFS fashion; once allocated to a customer, the server is deallocated only after the customer has fully arrived and has been fully served. With reference to the multiplexer problem, the customers correspond to the arriving frames, and the server to the trunk. We idealize the situation by assuming that the frames are continuous (i.e., not discretized into bits), and the ATM cell length is infinitesimally small. This implies that the trunk can start transmitting a frame as soon as it is allocated to the frame.

Let

$\lambda =$ *arrival rate of the frames*

$\sigma =$ *service time of frame over the trunk*

$\alpha =$ *time taken for frame to arrive over a low speed line*

We naturally assume that $\sigma < \alpha$. Let $\{W_t,\ t \geqslant 0\}$ denote the virtual waiting time process for the above queue, i.e., if a new frame were to start arriving at time t then its service will begin time units later. Now observe that if service of a frame begins more than σ time units before it finishes arriving over the low speed access line then the trunk will be able to "keep up" with the frame, and its transmission over the trunk will complete at the same instant that it finishes arriving over a line. Such a situation will occur if at the arrival epoch of a frame $\alpha - W_t \geqslant \sigma$; in this case there will be a jump in W_t of an amount $\alpha - W_t$. If at an arrival epoch $\alpha - W_t < \sigma$, then the trunk can-

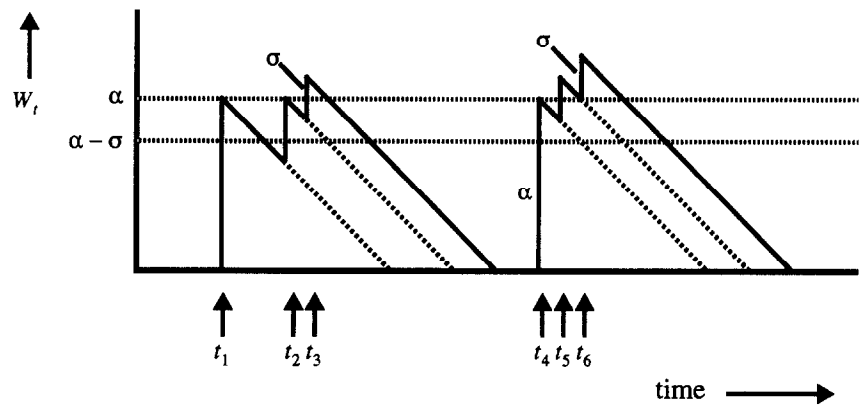

Fig. A.1. Sample paths for the non-interleaved multiplexor case.

not finish serving the frame before it finishes arriving. In this case there will be a jump of σ in W_t.

More formally, denoting by $t_1, t_2, \ldots, t_n, \ldots$ the arrival epochs of the frames, W_t will have right continuous sample paths and will decrease at a unit rate between jumps. The work following an arrival epoch is given by

$$W_{t_n^+} = \begin{cases} \alpha & \text{if } (W_{t^-} \leqslant \alpha - \sigma) \\ W_t^- + \sigma & \text{if } (W_{t^-} > \alpha - \sigma) \end{cases}$$

and the jumps at the arrival epochs will be given by

$$W_{t_n^+} = W_{t_n^-} + \max\left[\alpha - W_{t_n^-}, \sigma\right] = \max\left[\alpha, W_{t_n^-} + \sigma\right]. \quad \text{(A.1)}$$

This is illustrated in Fig. A.1.

Thus $\{W_t, t \geqslant 0\}$ is a Markov process. For large values of W_t (in fact for $W_t > \alpha - \sigma$), the process behaves exactly like the virtual waiting time process of an M/D/1 queue with arrival rate λ and service time σ. Hence W_t is ergodic if $\lambda\sigma < 1$. Let $w(x)$, $x \geqslant 0$ denote the density of the continuous part of the stationary distribution of $\{W_t, t \geqslant 0\}$, and w_0 denote the stationary probability that $W_t = 0$. Equating the rate of down-crossing of a level *at* x with the rate of up-crossing, we get the following equations for $w(\cdot)$, and w_0

$$w(x) = \lambda\left\{w_0 + \int_0^x w(u)\,du\right\}, \quad 0 < x \leqslant \alpha, \quad \text{(A.2)}$$

$$w(x) = \lambda \int_{(x-\sigma)}^{x} w(u)\,du, \qquad x > \alpha. \quad \text{(A.3)}$$

We now proceed to solve these integral equations to obtain the Laplace Transform $\tilde{w}(s)$ of $w(\cdot)$.

It immediately follows that

$$w(x) = \lambda w_0 e^{\lambda x}, \quad 0 < x \leqslant \alpha.$$

Now

$$\tilde{w}(s) = \int_0^\infty w(x) e^{-sx}\,dx = \int_0^\alpha \lambda w_0 e^{\lambda x} e^{-sx}\,dx + \int_0^\infty w(y+\alpha) e^{-s(y+\alpha)}\,dy.$$

Defining $\nu(x) = w(x+\alpha)$, $x \geqslant 0$ we have

$$\tilde{w}(s) = \frac{\lambda w_0}{(\lambda - s)}\{e^{(\lambda - s)\alpha} - 1\} + e^{-s\alpha}\tilde{\nu}(s), \quad \text{(A.4)}$$

where $\tilde{\nu}(s)$ is the Laplace Transform of $\nu(x)$. We next proceed to find $\tilde{\nu}(s)$. It follows from Eq. (A.1) that

$$w'(x) = \lambda\left[w(x) - w(x-\sigma)\right], \quad x > \alpha,$$

$$\nu'(y) = \lambda\left[\nu(y) - 1_{\{0 \leqslant y \leqslant \sigma\}}\lambda w_0 e^{\lambda(y+\alpha-\sigma)} - 1_{\{y > \sigma\}}\nu(y-\sigma)\right],$$

$$s\tilde{\nu}(s) - \nu(0) = \lambda\left[\tilde{\nu}(s) - \int_0^\sigma \lambda w_0 e^{\lambda(y+\alpha-\sigma)} e^{-ys}\,dy - e^{-s\sigma}\tilde{\nu}(s)\right],$$

$$\tilde{\nu}(s)\{s - \lambda + \lambda e^{-s\sigma}\} = \nu(0) - \lambda^2 w_0 e^{\lambda(\alpha-\sigma)} \int_0^\sigma e^{(\lambda - s)}\,dy.$$

Now from Eq. (A.3) it follows that

$$\nu(0) = \lambda \int_{(\alpha-\sigma)}^{\alpha} \lambda w_0 e^{\lambda u}\,du = \lambda w_0 e^{\lambda\alpha}(1 - e^{-\lambda\sigma}).$$

Substituting and simplifying we get

$$\tilde{\nu}(s) = \lambda w_0 e^{\lambda\alpha}\left[\frac{-1}{(\lambda - s)} + \frac{se^{-\lambda\sigma}}{(\lambda - s)(s - \lambda + \lambda e^{-\sigma s})}\right].$$

Substituting this expression for $\tilde{\nu}(s)$ into Eq. (A.4), and simplifying, we find

$$\tilde{w}(s) = \frac{\lambda w_0}{\lambda - s}\left(\frac{se^{\lambda(\alpha-\sigma)}e^{-\alpha s}}{(s - \lambda + \lambda e^{-\sigma s})} - 1\right). \quad \text{(A.5)}$$

It easily follows that

$$\lim_{s \to 0} \tilde{w}(s) = w_0\left(\frac{e^{\lambda(\alpha-\sigma)}}{1 - \lambda\sigma} - 1\right),$$

and since $w_0 + \lim_{s \to 0}\tilde{w}(s) = 1$, we obtain

$$w_0 = (1 - \lambda\sigma)e^{-\lambda(\alpha-\sigma)}. \quad \text{(A.6)}$$

As expected, $w_0 < (1 - \lambda\sigma)$ owing to the fact that trunk capacity is sometimes wasted while a frame

is arriving. Denote by W the random variable whose distribution is the stationary distribution of $\{W_t,\ t \geqslant 0\}$. Then

$$EW = -\lim_{s \to 0} \tilde{w}'(s).$$

A tedious differentiation of $\tilde{w}(s)$, followed by use of L'Hospital's rule to evaluate the limit yields (this calculation was also checked using MACSYMA!)

$$EW = -\frac{1}{\lambda}\left\{\frac{\lambda^2\sigma^2\left(2\frac{\alpha}{\sigma} - 3\right) + 2\lambda\sigma\left(1 - \frac{\alpha}{\sigma}\right)}{2(1-\lambda\sigma)}\right\} - \left\{\frac{1-\lambda\sigma}{\lambda}\right\}\{1 - e^{-\lambda(\alpha-\sigma)}\},$$

or, defining $\rho = \lambda\sigma$,

$$EW = -\frac{\sigma}{\rho}\left[\left\{\frac{\rho^2\left(2\frac{\alpha}{\sigma} - 3\right) + 2\rho\left(1 - \frac{\alpha}{\sigma}\right)}{2(1-\rho)}\right\} + (1-\rho)\{1 - e^{-\rho(\alpha/\sigma - 1)}\}\right].$$

This completes the mean delay analysis for the idealized model. The original problem was to obtain the mean delay of the first ATM cell after its arrival to the FIFO queue, i.e., $EW^{(1)}_{\text{nonint}}$. It is easily argued that the jumps in $W_t^{(1)}$ will be given by

$$W^{(1)}_{t_n^+} = W^{(1)}_{t_n^-} + \max\left[\left(\frac{(N-1)}{N}b_1 + \frac{b_t}{N} + h_t - W^{(1)}_{t_n^-}\right), b_t + Nh_t\right].$$

Thus $EW^{(1)}_{\text{nonint}}$ is obtained by replacing α with $\{[(N-1)/N]b_1 + b_t/N + h_t\}$ and σ with $b_t + Nh_t$ in the above formula for EW. After some simplification this yields the formula displayed in Section 2.1.

References

[1] D. Anick, D. Mitra and M.M. Sondhi, Stochastic theory of data handling systems with multiple sources, *Bell Systems Technical J.* **61** (8) (1982).

[2] P.H. Brill and J.M. Posner, Level crossing in point processes applied to queues: single server case, *Operations Res.* **25** (4) (1977) 662–674.

[3] CCITT Blue Book, 1980.

[4] CCITT Recommendation I.363, 1988.

[5] R.G. Cole, Performance models of binary synchronous communications multipoint circuits over virtual private line frame relay networks, *AT&T Technical J.* **67** (5) (1988) 41–56.

[6] R.B. Cooper, *Introduction to Queueing Theory* (North-Holland, 2nd Edition, New York, NY, 1981).

[7] P.M. Gopal, J.W. Wong and J.C. Majithia, Analysis of playout strategies for voice transmission using packet switching techniques, *Performance Evaluation* **4** (1984) 11–18.

[8] L. Kosten, Stochastic theory of a multi-entry buffer (1), *Delft Progress Rep.* **1** (1988) 10–18.

[9] A. Kumar, Performance of SNA/SDLC over virtual circuits in a data network (trademark of IBM), *AT&T Technical J.* **67** (5) (1988) 27–40.

[10] W.A. Montgomery, Techniques for packet voice synchronization, *IEEE J. Selected Areas Comm.* **SAC-1** (6) (1983) 1022–1028.

[11] A. Alan B. Pritsker, *Introduction to Simulation and SLAM II* (A Halsted Press Book, John Wiley and Sons, New York, 1986).

[12] D.D. Sheng, Virtual private-line performance and customer cost impacts, *AT&T Technical J.* **67** (6) (1988) 33–62.

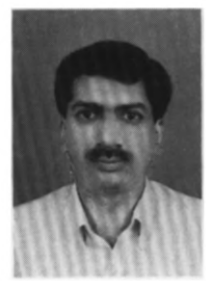

Anurag Kumar obtained his Bachelor's Degree in Electrical Engineering from the Indian Institute of Technology at Kanpur in 1977, and his PhD from Cornell University in 1981. He was a Member of Technical Staff at AT&T Bell Laboratories, Holmdel, for over 6 years. During this period he worked on the performance analysis of computer systems, communication networks, and manufacturing systems. Since 1988 he has been with the Indian Institute of Science (IISc), Bangalore, in the Department of Electrical Communication Engineering, where he is now Associate Professor. He is also the Coordinator at IISc of the Education and Research Network Project, which has set up a country-wide computer network for academic and research institutions, and conducts R&D and training in the area of communication networks. His own research interests are in the area of modelling, analysis, control and optimization problems arising in communication networks and distributed systems.

Robert G. Cole obtained his Professional Bachelor's Degree in Chemistry from the University of Miami in Coral Cables, Florida in 1977, and his PhD in Theoretical Chemistry from Iowa State University in 1982. He held postdoctoral positions in Theoretical Chemistry at Yale, Boston and Brown Universities from 1983 to 1985. He joined AT&T Bell Laboratories, Holmdel, in 1985 as a Member of Technical Staff. In 1990 he became a Distinguished Member of Technical Staff. During this period he worked on the performance analysis of computer networks and data networking products. His current interests are in the area of modelling architectures supporting service guarantees in high speed communications networks.