

.....
ATM Forum Document Number: ATM_Forum/97-0178R1.
.....

Title: ATM Switch Performance Testing Experiences
.....

Abstract: We experimented with the latency, throughput, fairness, and frame loss rate metrics. The results of these measurements are helpful in refining the baseline text. This revised version includes corrected and new measurements for throughput and latency.
.....

Source:

Gojko Babic, Arjan Durresi, Raj Jain, Justin Dolske, Shabbir Shahpurwala
The Ohio State University
Department of CIS Columbus, OH 43210-1277
Phone: 614-292-3989, Fax: 614-292-2911, Email: Jain@ACM.Org

The presentation of this contribution at the ATM Forum is sponsored by NASA.
.....

Date: April 1997
.....

Distribution: ATM Forum Technical Working Group Members (AF-TEST, AF-TM)
.....

Notice:

This contribution has been prepared to assist the ATM Forum. It is offered to the Forum as a basis for discussion and is not a binding proposal on the part of any of the contributing organizations. The statements are subject to change in form and content after further study. Specifically, the contributors reserve the right to add to, amend or modify the statements contained herein.
.....

An earlier version of this contribution was presented at the February 1997 meeting of the ATM Forum. This revised version includes new and corrected information in Table 2.2, 4.1, 4.2, and 4.3. The text associated with these tables has also been revised.

Introduction

We measured throughput and frame latency of a commercial switch using a commercial ATM monitor. The purpose of this contribution is to highlight our experiences in the following areas:

- **Frame Statistics from Cell Statistics:** Contemporary ATM monitors do not measure any frame level statistics. Therefore, it is necessary to derive frame level statistics based on cell statistics.
- **ATM Monitor Overhead:** ATM monitors have finite accuracy. It is necessary to take this into account when computing the frame level performance.
- **Background Traffic:** The baseline document does not yet contain information on background traffic. The tests presented here will help make progress in that direction.
- **Cell Transfer Delay:** Cell level latency has a high variance and, therefore, the average cell transfer latency is statistically not very meaningful.
- **Loss-less, Peak and Full-Load Throughput:** The baseline defines three types of throughputs. However, we found that lossless throughput is the only one that is meaningful. Others can be inferred.
- **Test Configurations:** The baseline defines 4 test configurations. Practical considerations help us identify better configurations that measure switch performance with less equipment.

1. Computing MIMO Frame Latency from CTD

Most current ATM monitors measure cell transfer delay (CTD), which is defined as the time between “Last bit in to first bit out” (LIFO) for the cell. At the December 1996 meeting, LIFO was rejected as the frame level metric. There were several reasons for this, including the fact that LIFO will be negative for most frames, since frames are not contiguous and most switches will be able to output first cell of a frame much before receiving the last cell of the frame. The accepted definition of latency is MIMO frame latency. In this section, we first define MIMO latency and then show how it can be obtained from current ATM monitors.

MIMO latency (Message-In Message-Out) is a general definition of the latency that applies to an ATM switch or a group of ATM switches and it is defined as follows:

$$\text{MIMO latency} = \min \{ \text{LILO latency}, \text{FILO latency} - \text{NFOT} \}$$

where:

- LILO latency = Time between the last-bit entry and the last-bit exit
- FILO latency = Time between the first-bit entry and the last-bit exit
- NFOT = Nominal Frame Output Time = Frame Input Time \times Input Rate/Output Rate
- Frame Input Time = Time between the first-bit entry and the last-bit entry

The following is an equivalent definition for MIMO Latency:

$$\text{MIMO latency} \begin{cases} \text{LILO latency} & \text{if input rate} \leq \text{output rate} \\ \text{FILO latency} - \text{NFOT} & \text{if input rate} \geq \text{output rate} \end{cases}$$

In cases where the input link rate is equal to the output link rate:

$$\text{MIMO latency} = \text{LILO latency} = \text{FILO latency} - \text{NFOT}$$

Note that for contiguous frames on input:

$$\begin{aligned} \text{Frame Input Time} &= \text{Frame Size} / \text{Input rate} \\ \Rightarrow \text{NFOT} &= \text{Frame Size} / \text{Output rate} \end{aligned}$$

An explanation of MIMO latency and its justification is presented in the ATM Performance Testing baseline document [1].

In our performance measuring experiments, we used a commercial ATM monitor as a traffic generator as well as a traffic analyzer. This monitor and, as far as we are aware, all other similar systems can provide measurement data on delays and inter-arrival times only at the cell level. Considering that the definition of MIMO latency requests bit level data, provided here is an analysis which results in adjustments to the above expression, so that data at the cell level can be used to calculate MIMO latency.

First, some observations about ATM monitors:

- The cell transfer delay is defined as the amount of time it takes for a cell to begin leaving the generator and to finish arriving at the analyzer, i.e. the time between the first bit in and the last bit out. Most commercial ATM monitors measure this delay with a finite granularity. Our ATM monitor has a resolution of 0.5 μsec .

We obtained the average cell transfer delay of 3.33 μsec for the case of a closed loop on the ATM monitor with a 10-meter fiber-optic cable (155 Mbps OC-3c). The measured delay is about 15% (0.4 μsec) larger than the theoretical value of the cell transmit time over a 155 Mbps link, plus the propagation delay for a 10 meter link. This discrepancy can be attributed to delays internal to the ATM monitor and its resolution of 0.5 μsec . Similar results are obtained when an UTP-5 closed loop connector was used on another 155 Mbps port instead of a fiber optic cable.

- The cell inter-arrival time is defined as the time between arrival of the last bit of the first cell and the last bit of the second cell. The resolution is 0.5 μsec .

We found that inter-arrival times measured by our ATM monitor are very accurate. For example, when we generated traffic at its maximum rate over a 155 Mbps closed

loop, the average cell inter-arrival time reported by the ATM monitor was 2.83 μsec , which is exactly the time needed to transmit one cell at that rate. This implies that all cells were received (and sent) back to back at the maximum transmit rate. One reason for this is that only one port is involved in the traffic analysis. (In the case of CTD, the clock generated from one port has to be subtracted from the clock at the receiving port.)

Now we analyze two cases for MIMO latency calculation

1.1 Case 1 Calculation: $\text{Input rate} \leq \text{Output rate}$

In cases when the input link rate is less than or equal to the output link rate:

$$\text{MIMO latency} = \text{LILO latency.}$$

From Figure 1.1, it can be observed that:

$$\text{LILO Latency} = \text{Last cell transfer delay} - \text{Last cell input transmit time}$$

where:

- Cell input transmit time = Time to transmit one cell into the input link.

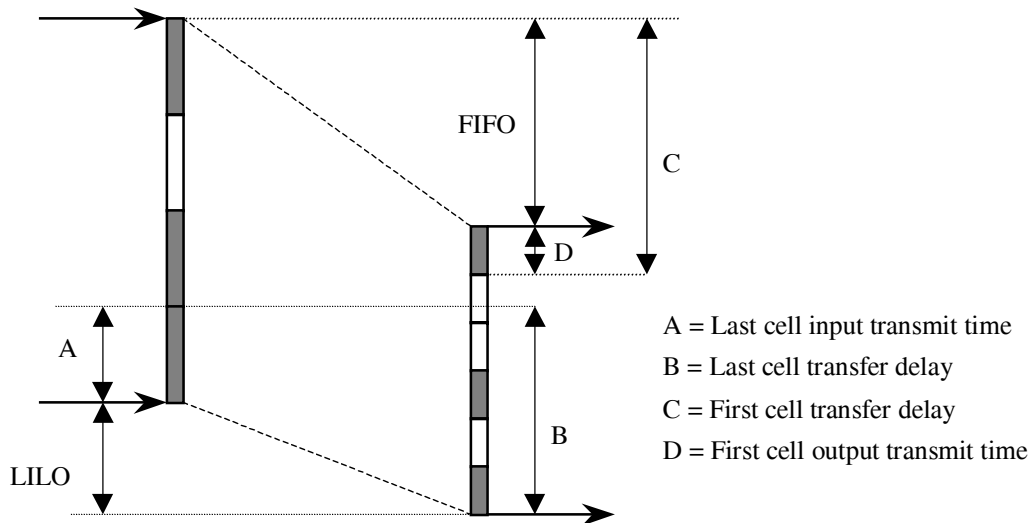


Figure 1.1: Input Rate < Output Rate

To account for the overhead in the monitor, the following adjustment is made in the previous expression:

$$\text{LILO Latency} = \text{Last cell transfer delay} - (\text{Last cell input transmit time} + \text{Monitor overhead}) \quad (1)$$

In conclusion, to calculate MIMO latency when the input link rate is less than or equal to the output link rate, it is sufficient to measure only the last cell delay.

It can be also observed, from Figure 1.1, that:

$$\text{FIFO Latency} = \text{First cell transfer delay} - (\text{First cell output transmit time} + \text{Monitor overhead}) \quad (2)$$

where:

- FIFO latency = Time between the first-bit entry and the first-bit exit
- Cell output transmit time = Time to transmit one cell into the output link.

This expression is included because it is needed later in this document.

1.2 Case 2 Calculation: Input rate \geq Output rate

In cases when the input link rate is greater than or equal to the output link rate:

$$\text{MIMO latency} = \text{FILO latency} - \text{NFOT} \quad (3a)$$

NFOT can be calculated, given the cell pattern of the frame on input and rates of input and output links, while FILO latency has to be measured.

From Figure 1.2, it can be observed:

$$\text{FILO latency} = \text{FIFO latency} + \text{FOLO time} \quad (3b)$$

where:

- FOLO time = time between the first bit out and the last bit out.




Figure 1.2: Input Rate > Output Rate

Also:

$$\text{FIFO latency} = \text{First cell transfer delay} - \text{First cell output transmit time} + \text{Monitor overhead} \quad (3c)$$

$$\text{FOLO time} = \text{First cell-Last cell inter-arrival time} + \text{Last cell output transmit time} \quad (3d)$$

where :

- Cell output transmit time = Time to transmit one cell into the output link.

Note that because the measurements of cell inter-arrival times are very accurate, we do not need any corrections in the FOLO expression due to any Monitor overhead.

In conclusion, to calculate MIMO latency when the input link rate is greater than or equal to the output link rate, it is necessary to measure the first cell transfer delay and inter-arrival time between the first cell and the last cell of a frame.

2. MIMO latency measurement tests without background traffic

2.1 Configuration

The test configuration for MIMO latency measurements without background traffic is shown in Figure 2.1. This configuration includes one ATM monitor and one ATM switch

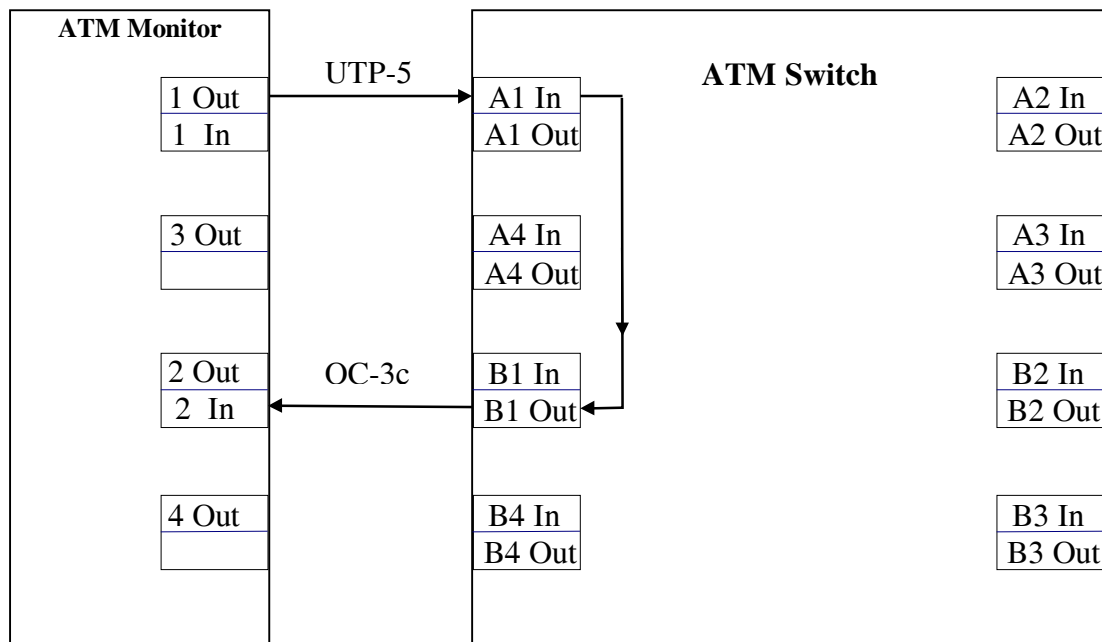


Figure 2.1: Test configuration for measurements of MIMO latency without background traffic

with a 155 Mbps UTP-5 link between the monitor port 1 and the switch port A1 and a 155 Mbps OC-3c link between the monitor port 2 and the switch port B1. The switch has two network modules A and B with four ports on each module. The ports are numbered A1, ..., A4 and B1, ..., B4. A permanent virtual path connection (VPC) or a permanent virtual channel connection (VCC) is established between the monitor ports 1 and 2 through the switch ports A1 and B1. That VPC (or VCC) is used for transmission of frames whose latency is measured, and it is referred as the foreground VPC (or the foreground VCC). Figure 2.1 also indicates the traffic flow direction.

2.2 Methodology, Measurement Results and Analysis

Note that when the input link rate is equal to the output link rate, as it is in our test configuration, MIMO latency calculations can be done according to either Case 1 calculation or Case 2 calculation. Here, we first use Case 1 calculation, for which it is sufficient to measure only the last cell delay. At the end of this section, we compare results obtained from Case 1 calculation and from Case 2 calculation.

Measurements of MIMO latency are performed slightly differently than as given in ATMF document [1]. For each test run, first, a sequence of equally spaced 192 cell frames is being sent over the foreground VPC (or the foreground VCC), at a rate of 4.63 frames/sec, i.e. with an inter-frame time (the time between beginnings of two successive frames) of 0.216 sec. After the flow has been established, we record the average transfer delays of the last cells in the next 1,000 consecutive frames. In different test runs, besides the average delay of 192nd (last) cell of all frame (1,000 samples), we also record a different subset of the following:

- Average delay of 1st cell of all frames (1,000 samples)
- Average delay of 2nd cell of all frames (1,000 samples)
- Average delay of 97th cell of all frames (1,000 samples)
- Average delay of 191st cell of all frames (1,000 samples)
- Average delay of 2nd through 191st cells of all frames (190,000 samples)
- Average delay of 3rd through 190th cells of all frames (188,000 samples)
- Average delay of 3rd through 96th cells of all frames (94,000 samples)
- Average delay of 98th through 190th cells of all frames (93,000 samples)

Table 2.1 presents measurements (average transfer delays in μsec) from 7 test runs. The first five runs use a VPC, the last two runs use a VCC. It can be observed that there is no significant difference in average transfer delays (and consequently in MIMO latency) for the switch performing either path or circuit switching.

Table 2.1

Run #	1st cell	2nd cell	2-191 cells	3-96 cells	97th cell	3-190 cells	98-190 cells	191st cell	192nd cell
1	19.02		20.52						20.77
2	19.06		20.54						20.78
3	19.04	19.21				20.53		20.79	20.77
4	19.07	19.21		20.31	20.75		20.76		20.78
5	19.07	19.19		20.32	20.73		20.76		20.78
6	19.14		20.58						20.83
7	19.13		20.58						20.81

Table 2.1 above clearly indicates that differences in cell transfer delays are a function of the cell's order inside a frame. Easily, it can be observed that cells at the beginning of the frame have lesser transfer delays than those towards the end of the frames.

Here is the MIMO latency calculation for test run #4 (or #5) using expression (1):

$$\text{MIMO latency} = 192\text{nd Cell transfer delay} - 3.33\mu\text{sec} = 20.78 - 3.33 = 17.45 \mu\text{sec}$$

Note that for the same test run, the FIFO latency using expression (2) is:

$$\text{FIFO latency} = \text{First Cell transfer delay} - 3.33 \mu\text{sec} = 19.07 - 3.33 = 15.74 \mu\text{sec}$$

The above indicates that FIFO latency is about 10% lesser than MIMO latency.

Table 2.2 presents measurement data for two randomly chosen frames (from run #7 in Table 2.1) and results from the Case 1 and Case 2 MIMO latency calculations. The first three columns show the first cell transfer delay, the last cell transfer delay and inter-arrival time between the first cell and the last cell for those two frames. The next column (labeled "MIMO latency [1]") shows results from the Case 1 calculation. The next three columns include intermediate results from the Case 2 calculation, i.e. FIFO latency, FOLO time and FILO latency). The last column (labeled "MIMO latency [2]") shows results from the Case 2 calculation. All data are in μsec .

In the Case 2 calculation, we need to calculate NFOT. In our tests, input frames are contiguous, and in these cases NFOT can be calculated as follows:

$$\begin{aligned} \text{NFOT} &= \text{Frame size} / \text{Output rate} = 192 \text{ cells} / 353,207.55 \text{ cells/sec} \\ &= 543.59 \mu\text{sec} \end{aligned}$$

Table 2.2

1st cell delay	last cell delay	1st cell-last cell inter-arrival time	MIMO latency [1]	FIFO latency	FOLO time	FILO latency	MIMO latency [2]
21.5	21.5	541.0	18.2	18.2	543.83	562.03	18.44
18.5	21.0	543.5	17.7	15.2	546.33	561.53	17.94

The purpose of Table 2.2 is to illustrate that both expression (1) and expressions (3a-d) forms of MIMO latency calculation provide the same values in the cases when the input link rate is equal to the output link rate.

From Table 2.2, it can be observed that although measured data (the first cell delay, the last cell delay and inter-arrival time) are different for two considered frames, the calculated values for MIMO latency are nearly identical and within the ATM monitor resolution time of 0.5 μ sec.

3. MIMO latency with background traffic

The ATMF document [1] states that details of measurements with background traffic are for further study. The results presented in this section can be used to provide future text for the document.

3.1 Configuration

The test configuration for measuring MIMO latency with background traffic is given in Figure 3.1. The configuration includes one ATM monitor and one ATM switch. There are two 155 Mbps UTP-5 links between monitor ports 1 and 3 and switch ports A1 and A4, respectively. There are also two 155 Mbps OC-3c links between monitor ports 2 and 4, and switch ports B1 and B4, respectively. In addition, we made external loopbacks on 155 Mbps OC-3c switch ports A2 and A3 (through a 10-meter fiber optic cable) and on 155 Mbps UTP-5 switch ports B2 and B3 (through connectors)

Several (permanent) virtual path connections are established as indicated in Figure 3.1. The foreground VPC is established between monitor ports 1 and 2 through switch ports A1 and B1. Traffic and frames transferred over the foreground VPC are referred as foreground traffic and measured frames, respectively.

Background traffic is generated from three monitor ports (ports 2, 3, and 4) through 6 VPCs (indicated with dotted lines). Each VPC starts and ends at the monitor so that the traffic generation can be controlled and the receiving traffic can be analyzed. In all tests, all background VPCs are loaded equally. This results in equal load on all switch ports (except for the two ports used for foreground traffic).

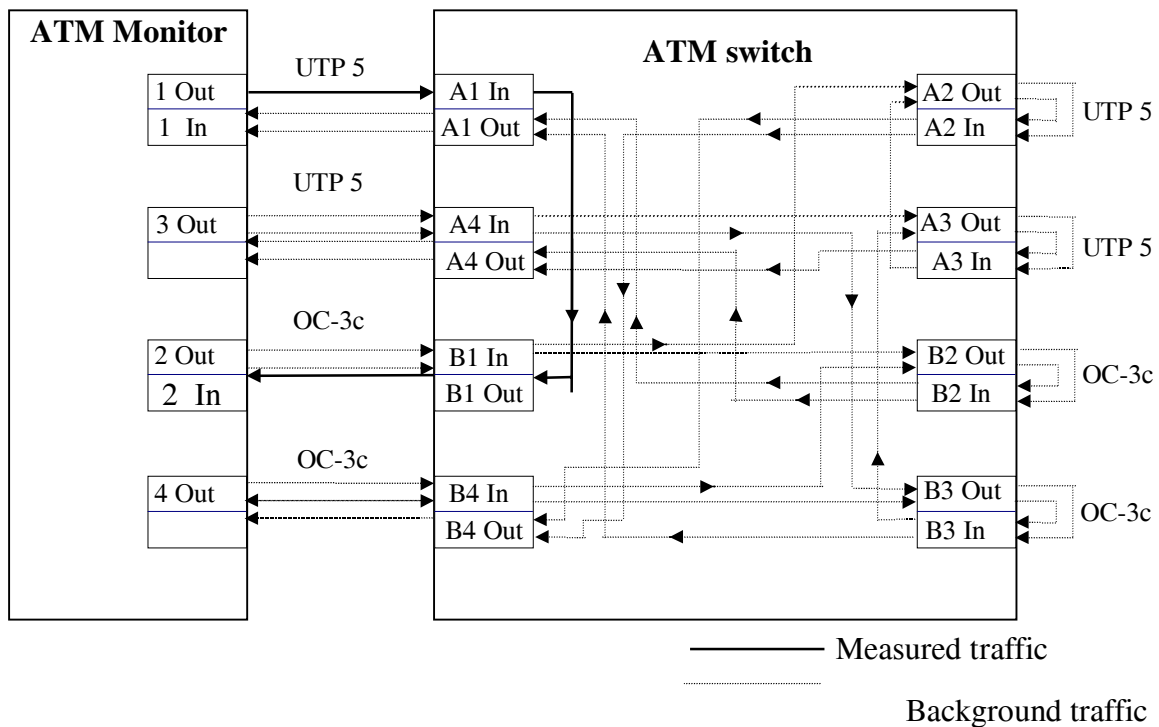


Figure 3.1:
Test configuration for measurements of MIMO latency with background traffic

The link between the monitor port 1 and switch port A1 is used to transfer foreground traffic in one direction and background traffic in the another direction. The link between the monitor port 2 and switch port B1 is used to transfer foreground traffic in one direction and background traffic in the another direction. Note that the foreground traffic does not share any generator or analyzer port with any other VCs in the same direction. This avoids possible distortions in the foreground traffic.

Using loop-back connections and background traffic VPC's as shown in Figure 3.1, we are able to load the seven ports of the switch at 100%. The maximum background load offered to the switch in our measurements equals $7 \times 149.76 \text{ Mbps} = 1.048 \text{ Gbps}$. We called this load the maximum background load (MBL) for the given switch. For an n-port switch, the MBL is (n-1) times the port capacity.

3.2 Methodology

Measurements for the MIMO latency with background traffic are performed in the following steps:

- a) Background traffic flow for the given load (a percentage of MBL) is started and allowed to stabilize.

- b) One or more of the following test runs are then performed. For each test run, a sequence of equally spaced frames of the given length is first sent over the foreground VPC at a very low rate. In our tests, this rate was set at the minimum rate allowed by the monitor. The rate was 4.63 frames/sec, i.e. with an inter-frame time (the time between beginnings of two successive frame) of 0.216 sec. After the flow has been established, we record the average transfer delays of the last cells of the next 1,000 frames. Note that when the input link rate is equal to the output link rate, as is the case in our configuration, we can calculate MIMO latency according to either Case 1 or Case 2. We have chosen to use the Case 1 approach for which it is sufficient to measure only the last cell delay. In each test run, we measure and record:
- the average transfer delay of the last cell of all frames (1,000 samples),
 - the average transfer delay of the first cells of all frames (1,000 samples), and
 - the average transfer delays of all cells between the first and the last cells of all frames (190,000 samples).
- c) The background load is increased, and steps a) and b) are repeated. We stop when background load reaches 100% of MBL (or very close to it).

We chose two different frame lengths for the foreground VPC: 192 cells and 1000 cells. The VPC uses UBR class of service. The following background traffic types were used:

- UBR traffic with burst size = 2,004 cells, i.e. 2,004 cell frames are sent at given rate,
- UBR traffic with burst size = 384 cells, i.e. 384 cell frames are sent at given rate,
- CBR traffic.

Note that in the first two cases, both foreground and background traffics are of the same priority. In the third case, the background traffic has priority over the foreground traffic.

3.3 Measurement Results

Results of our measurements are presented in Tables 3.1 through 3.4. The first column in each table indicates a background load as a percentage of the MBL. The next three columns in the tables include the average cell transfer delays (in μsec) for the cells of measured frames as indicated. The fifth column includes MIMO frame latency calculated according to expression (1), and the sixth column includes FIFO latency calculated according to expression (2). The last column indicates the percentage difference between MIMO and FIFO latencies.

Table 3.1 presents results for UBR background traffic with burst size of 2,004 cells. The length of measured frames equals 1,000 cells.

Table 3.1

Load %	1st cell	2nd -999th cells	Last cell	MIMO	FIFO	Dif %
25	19.2	20.96	21.08	17.75	15.87	11.85
40	19.24	21.13	21.32	17.99	15.91	13.07
50	19.28	21.20	21.36	18.03	15.95	13.04
60	19.38	21.54	21.95	18.62	16.05	16.01
65	19.4	21.38	21.6	18.27	16.07	13.69
70	19.52	21.43	21.65	18.32	16.19	13.16
75	19.47	28.81	36.31	32.98	16.14	194.34
80*	19.32	57.54	94.17	90.84	15.99	468.11
90*	19.4	98	168.08	164.75	16.07	925.2
97*	19.49	122.42	207.54	204.21	16.16	1163.7

* Losses in foreground traffic observed

Table 3.2 presents results for UBR background traffic with burst size of 384 cells. The length of measured frames equals 192 cells

Table 3.2

Load %	1st cell	2nd -191nd cells	Last cell	MIMO	FIFO	DIF%
50	18.57	20.28	20.55	17.22	15.24	12.99
50	18.98	20.32	20.64	17.31	15.65	10.61
70	18.76	20.54	20.85	17.52	15.43	13.55
70	18.67	20.47	20.86	17.53	15.34	14.28
80	18.73	20.57	20.94	17.61	15.40	14.35
80	18.82	20.64	21.00	17.67	15.49	14.07
90	18.78	20.74	21.06	17.73	15.45	14.76
90	18.97	20.73	21.00	17.67	15.64	12.98
95	18.88	20.85	21.20	17.87	15.55	14.92
95	18.94	20.85	21.18	17.85	15.61	14.35

Table 3.3 presents results for UBR background traffic with burst size of 1,000 cells. The length of measured frames equals 192 cells.

Table 3.3

Load %	1st cell	2nd -191nd cells	Last cell	MIMO	FIFO	DIF%
10	19.11	20.58	20.81	17.48	15.78	10.77
10	19.11	20.58	20.83	17.50	15.78	10.90
20	19.1	20.65	20.90	17.57	15.77	11.41
20	19.6	20.50	20.73	17.40	16.27	6.95
25	19.04	20.70	20.98	17.65	15.71	12.35
25	19.01	20.67	20.92	17.59	15.68	12.18
50	19.3	20.84	21.10	17.77	15.97	11.27
50	19.31	20.85	21.10	17.77	15.98	11.20
60	19.11	20.63	20.91	17.58	15.78	11.41
60	19.09	20.99	21.51	18.18	15.76	15.36
60	19.06	20.72	20.96	17.63	15.73	12.08
60	19.36	20.71	20.99	17.66	16.03	10.17
60	19.29	20.67	20.94	17.61	15.96	10.34
70	19.18	22.93	25.41	22.08	15.85	39.31
70	19.18	21.75	23.05	19.72	15.85	24.42
75	19.17	24.21	28.00	24.67	15.84	55.74
75	19.23	24.58	28.71	25.38	15.90	59.62
75	19.56	28.11	33.70	30.37	16.23	87.12
75	19.56	21.52	22.41	19.08	16.23	17.56
75	19.52	24.91	29.33	26.00	16.19	60.59
75	19.46	21.33	22.14	18.81	16.13	16.62
75	19.45	22.27	23.99	20.66	16.12	28.16
75	19.52	26.88	33.31	29.98	16.19	85.18
75	19.47	24.46	28.43	25.10	16.14	55.51
80	19.26	27.64	34.84	31.51	15.93	97.80
80	19.27	26.82	33.20	29.87	15.94	87.39
80	19.49	22.24	23.85	20.52	16.16	26.98
80	19.54	28.68	36.68	33.35	16.21	105.74
80	19.53	21.02	21.40	18.07	16.20	11.54
80	19.56	31.35	42.67	39.34	16.23	142.39
80	19.54	24.11	27.73	24.40	16.21	50.52
80	19.49	25.56	30.60	27.27	16.16	68.75
90	19.4	33.26	45.97	42.64	16.07	165.34
90	19.43	33.26	46.02	42.69	16.10	165.16
90	19.64	30.49	40.44	37.11	16.31	127.53
90	19.63	30.76	40.96	37.63	16.30	130.86
95	19.39	32.19	43.91	40.58	16.06	152.68
95	19.35	31.79	43.11	39.78	16.02	148.31
95	19.69	32.00	43.36	40.03	16.36	144.68
95	19.67	33.82	47.05	43.72	16.34	167.56
97.5	19.41	35.58	50.59	47.26	16.08	193.91
97.5	19.41	34.92	49.41	46.08	16.08	186.57

Table 3.4 presents results for CBR background traffic. The length of the measured frames equals 192 cells.

Table 3.4

Load %	1st cell	2nd -191nd cells	Last cell	MIMO	FIFO	DIF%
50	19.16	21.77	22.24	18.91	15.83	19.46
50	19.14	21.78	22.24	18.91	15.81	19.61
70	19.44	22.72	23.29	19.96	16.11	23.90
70	19.48	22.73	23.29	19.96	16.15	23.59
80	19.62	23.12	23.64	20.31	16.29	24.68
80	19.7	23.17	23.70	20.37	16.37	24.43
90*	20.15	23.82	24.31	20.98	16.82	24.73
90*	20.43	23.91	24.42	21.09	17.10	23.33

* Losses in foreground traffic observed

3.4 Analysis

A number of observations and conclusions can be made based on these measurements:

- FIFO latency has little or no sensitivity to the length of measured frames or the background traffic load. For example, FIFO latency without background traffic differs only about 3%-4% from FIFO latency with UBR background traffic of even 97% of MBL. With CBR background traffic, similar behavior is observed for loads up to 90% of MBL. All of these apply for all lengths of measured frames. Thus, FIFO latency does not really measure frame latency. It measures only the first cells latency, which is the minimum of all cells of a frame.
- MIMO latency results are quite different from FIFO latency results. For cases with measured frames of large length (1,000 cells) and UBR background traffic with large sized bursts (2,004 cells), MIMO latency increases 100% and more for loads of 70% of MBL or higher. (Table 3-1)

For cases with measured frames of short length (192 cells) and UBR background traffic with medium size bursts (1,000 cells), a significant increase in MIMO latency is observed at loads of 70% of the MBL or higher, but not as much as in the previous case. (Table 3-2)

For cases with measured frames of short length (192 cells) and UBR background traffic with small sized bursts (384 cells), MIMO latency does not change significantly when the background traffic load is increased. (Table 3-3)

For cases with measured frames of shorter length (192 cells) and CBR background traffic, there is no significant change in MIMO latency when the background traffic load is increased. (Table 3-4)

- From Tables 3.1 - 3.4, the average transfer delay of various cells in a frame show an interesting trend. The first cell has the lowest delay while the last cell has the highest. This can be explained by the fact that as the successive cells of a frame arrive, they have to wait in the switch queue for service. While the average cell transfer delay is one of standard ATM metrics, the CTD varies widely for various cells of a frame. Average CTD is therefore not very meaningful statistically.
- Tables 3.1 and 3.4 indicate that under certain background loads, there are losses in the foreground traffic. Although we calculated and presented finite values for MIMO latency in such cases, it should be noted that MIMO latency for lost frames is infinite.

4. Throughput measurement

4.1 Configuration

In throughput measurements, we use an n-to-1 configuration as given in the baseline document [1], i.e. the case with n traffic sources generating frames through input links to one output link, as shown in Figure 4.1. However, since our monitor has only 4 ports, we are able to perform tests only with the 4-to-1 configuration. We also perform tests with 2-to-1 and 3-to-1 configurations, but the results are similar to those reported here for the 4-to-1 case.




Figure 4.1: 4-to-1 configuration

The 4-to-1 configuration for throughput measurements is given in Figure 4.2. The configuration includes one ATM monitor and one ATM switch with two 155 Mbps UTP-5 links and two 155 Mbps OC-3c links. Four permanent virtual path connections (VPC) are established between the monitor ports. Note that the link between the monitor port 2 and the switch port B1 is used in one direction as the output link and in the another direction as one of the input links.

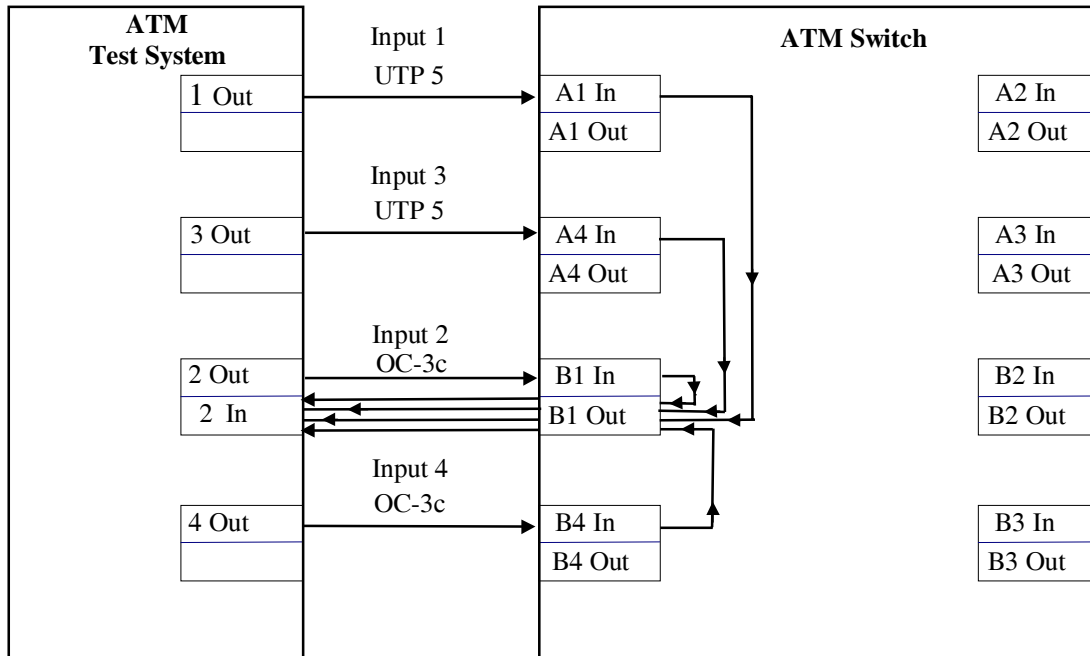


Figure 4.2: Test configuration for throughput measurements

4.2. Methodology, Measurement Results and Analysis

Four traffic sources generate, over VPC's, fixed length frames (106 cells) at identical rates with equally spaced frames. All frames are generated in simulated AAL 5 format.

A frame in simulated AAL 5 format is transmitted as 106 back-to back cells, with PT field in the ATM header set to 0 in the first 105 cells, while set to 1 in the last cell. Since we are interested not only in frame losses but also in cell losses to compare with, each cell payload includes a 16 bit cell sequence number and a 10 bit CRC field. With such cells, undetected cell loss is unlikely. We could not include the CPCS-PDU trailer, but as mentioned above, we have other (and better) means to detect corrupted frames

Changing the frame rate varies the input load. Each test run lasts 180 sec. Our measurements show that as long as the total input load is less than the output link rate, no loss of frames (or cells) is observed. For example, no loss is detected even when the load on each input line is 24.94% of its rate resulting in a total load of 99.76% ($=4 \times 24.94$) of the output link rate. The same behavior is observed regardless of the early packet discard (EPD) feature being turned off or on.

If the total input rate is even slightly higher than the output link rate, the frames are lost at a high rate.

Table 4.1 presents measurement results for the case when the total load is 100.32% (= 4x25.08%) of the output link rate. Measured results include cell loss ratio and frame loss ratio.

Table 4.1: Total offered load = 100.32% of output link rate

Metric	Input 1	Input 2	Input 3	Input 4	Mean
Cell Loss Ratio	0.00333	0.00381	0.00387	0.00278	0.00345
Frame Loss Ratio	0.288	0.283	0.204	0.283	0.265

Table 4.2 presents the same results when the total offered load to the output link is 120% (= 4x30%) of its rate.

Table 4.2: Total offered load = 120% of output link rate

Metric	Input 1	Input 2	Input 3	Input 4	Mean
Cell Loss Ratio	0.177	0.187	0.157	0.146	0.167
Frame Loss Ratio	0.817	0.784	0.736	0.820	0.789

Table 4.3 presents the same results when the total offered load to the output link is 400% (= 4x100%) of its rate.

Table 4.3: Total offered load = 400% of output link rate

Metric	Input 1	Input 2	Input 3	Input 4	Mean
Cell Loss Ratio	0.742	0.743	0.743	0.744	0.743
Frame Loss Ratio	1.0	1.0	1.0	1.0	1.0

From Table 4.1, it is observed that even with loads just slightly over the output link rate, the cell loss ratio is small but the frame loss ratio is high. The frame loss ratio is two orders of magnitude larger than the cell loss ratio. Note that frame loss rate varies between four traffic sources (within the range 20%-29%) resulting in some unfairness.

From Table 4.2, it is seen that with an offered load of 20% over the output link rate, the frame loss ratio is considerable, and 73% to 82 % of input frames are lost.

From Table 4.3, it is observed that with an offered load 300 % over the output link rate (full load per each input), all input frames are lost.

Although the manufacturer of the ATM switch we tested claims that Early Packet Discard is implemented, our tests did not show any improvements in frame loss rates with EPD on.

In conclusion, for the n-to-1 configurations, the lossless throughput for the switch under test is 155 Mb/s (ie equal to the output link rate). Obviously, in this case the lossless

throughput equals the peak throughput. Also, from the results presented in Table 4.3, we have found that for this particular ATM switch, the full load throughput for the n-to-1 configuration does not make sense, because even with EPD turned on practically all the frames are lost.

5. Summary

- It is possible to compute MIMO frame latency using current ATM monitors that give only cell level statistics.
- The cell transfer delays of various cells in a frame are widely different. The first cell of a frame has much lower latency than later cells. Therefore the average cell transfer delay is not statistically meaningful.
- The frame transfer delay depends upon the background traffic. The key parameters of the background traffic are its frame size, load level, and priority. A simple UBR traffic pattern with a few different frame sizes can provide a useful background load at the same priority as the measured traffic or CBR traffic can be used as a higher priority background load.
- There is an excessive loss of measured frames when the background traffic is close to full load even though the background traffic does not share ports with the measured traffic. This configuration and other similar to it need to be added to delay measurements.
- Peak throughput is equal to lossless throughput. If we find the same pattern on many switches, one of the two metrics could be removed.
- Variance in throughput measurements is negligible, and so we may remove the requirement for specifying the standard error of throughput.
- For the n-to-1 configuration, the frame level throughput vs. input load graph is a straight-line until the throughput reaches the output capacity. It then drops suddenly to zero. Thus the full load throughput is zero in n-to-1 configurations. We may, therefore, reduce the number of test configurations and/or remove the full load throughput metric.
- Throughput for different VCs is identical as long as there is no loss. Thus, fairness of throughput is not a useful metric.
- The frame loss rate for different VCs in a n-to-1 configuration is not identical. Therefore, fairness of frame loss rate is a useful metric to add.

We are continuing further experiments before suggesting specific changes to the baseline text.

References

- [1] ATM Forum Performance Testing Specification, BTD-TEST-TM-PERF.00.01 (96-0810R4): January 24, 1997.

All of our other related ATM Forum contributions and papers can be obtained on our web page: <http://www.cis.ohio-state.edu/~jain/>