
ATM Forum Document Number: ATM_Forum/97-0178

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.

Source:

Gojko Babic, Arjan Duresi, Raj Jain

The Ohio State University

Department of CIS Columbus, OH 43210-1277

Contact 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: February 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.

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

1. **Frame Statistics from Cell Statistics:** Current monitors do not measure any frame level statistics. Therefore, it is necessary to derive frame level statistics based on cell statistics.
2. **Monitor Overhead:** Monitors have finite accuracy. It is necessary to take this into account when computing the frame level performance.
3. **Background Traffic:** The baseline document does not yet contain information on background traffic. The tests presented here will help make progress in that direction.
4. **Cell Transfer Delay:** Cell level latency has a high variance and, therefore, the average cell transfer latency is statistically not very meaningful.
5. **Lossless, 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.

6. 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 monitors measure cell transfer delay (CTD), which is defined as the time between “Last bit in to first bit out” (LIFO) for the cell. In the December 1996 meeting, LIFO was rejected as the frame level metrics. There were several reasons for it, including the fact that LIFO for most frames will be negative 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 and then show how it can be obtained from current 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 = FIT \times Input Rate/Output Rate
- FIT = Frame Input Time = Time between the first-bit entry and the last-bit entry

Note that for contiguous frames on input:

$$\text{Frame Input Time} = \text{Frame Size} / \text{Input rate}$$

and :

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

The following is an equivalent definition for MIMO Latency:

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

Note that for input rate = output rate:

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

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 test system as a traffic generator as well as a traffic analyzer. This system and, as far as we are aware, all other similar systems can provide data on delays and inter-arrival times only at the cell level. Considering that the definition of MIMO latency requests bit level data, here we provide analysis which results into adjustments to the above expression, so data at the cell level can be used to calculate MIMO latency. First some observations about commercial monitors (based on a sample of one):

1. 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 last bit in and the first bit out. Most commercial monitors measure this delay with a finite granularity. Our monitor has a resolution of 0.5 μ sec.

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

2. 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 μ sec.

We found that inter-arrival times measured by our monitor were very accurate. For example, when we generated traffic at its maximum rate over 155 Mbps closed loop, the average cell inter-arrival time reported by the monitor was 2.83 μ sec, 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 generation. (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. MIMO Latency Calculation: Case 1. (Input rate \leq Output rate)

We have MIMO latency = LILO latency in cases when Input link rate \leq Output link rate.

From Figure 1.1 we have:

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

where a cell input transmit time is the time to transmit one cell into the input link.




Figure 1.1: Case 1: Input Rate \leq Output Rate

To account for the overhead in the monitor, we have to make the following adjustment in the previous expression:

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

In conclusion, when input link rate is less than or equal to the output link rate, it is sufficient to measure only the last cell delay and:

$$\text{MIMO Latency} = \text{LILO Latency} = \text{Last cell transfer delay} - (\text{Last cell input transmit time} + \text{Test system overhead})$$

From Figure 1.1, we also have:

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

where a cell output transmit time is the time to transmit one cell into the output link. This expression is included because it is needed later in this document.

1.2. MIMO Latency Calculation: Case 2. (Input rate \geq Output rate)

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

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

NFOT can be calculated while the FILO latency has to be measured.

From Figure 1.2 we have:

$$\text{FILO latency} = \text{FIFO latency} + \text{FOLO time}$$

where

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

Also, we have:

FIFO latency = First cell transfer delay -
(First cell output transmit time + Test system overhead)

FOLO time = First cell-Last cell inter-arrival time + Last cell output transmit time

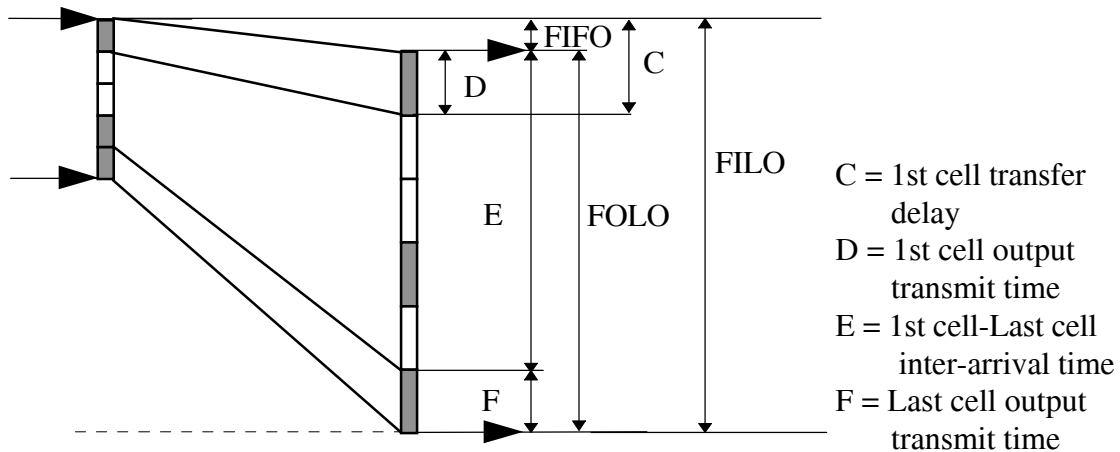


Figure 1.2 : Case 2 Input Rate \geq Output Rate

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 test system overhead.

In conclusion, to calculate MIMO latency when 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. The configuration includes one ATM test system (monitor) and one ATM switch with an 155 Mbps UTP-5 link between monitor port 1 and switch port A1 and an 155 Mbps OC-3c link between monitor port A2 and the switch port B1. The switch has two cards A and B with four ports on each card. The ports are numbered A1, ..., A4 and B1, ..., B4. A permanent virtual path connection (VPC) is established between the monitor ports 1 and 2 through switch ports A1 and B1. That VPC is used for

transmission of frames whose latency is measured, and it is referred as the measuring VPC. Figure 2.1 also indicates the traffic flow direction.

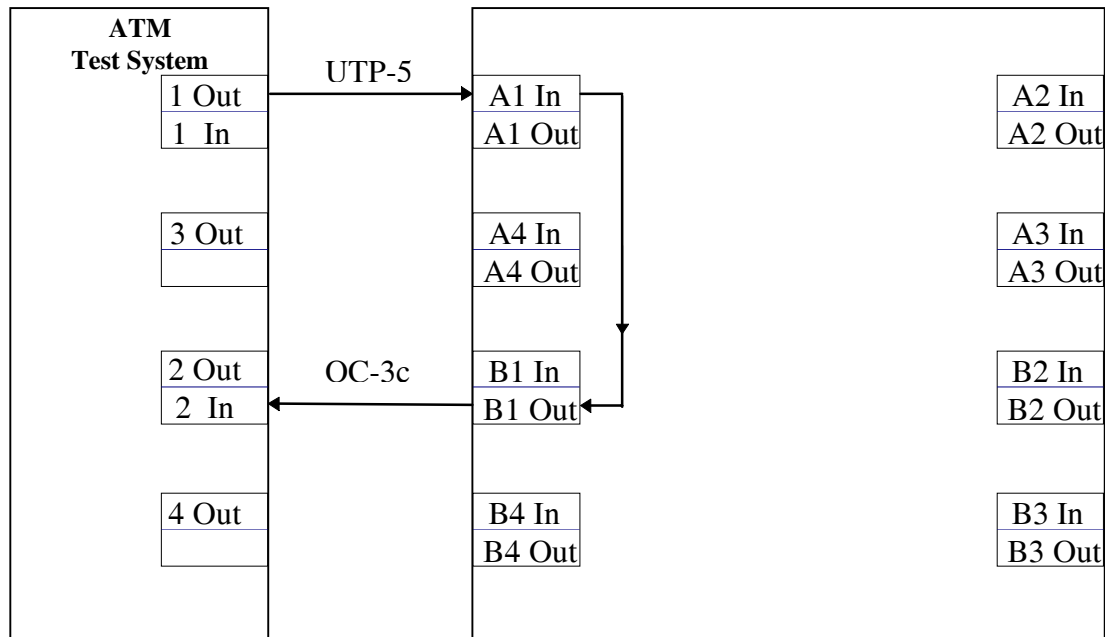


Figure 2.1.

Test configuration for measurements of MIMO latency without background traffic

2.2. Methodology, Measurement Results and Analysis

Note that when Input link rate = Output link rate, as it is in our configuration, we can calculate MIMO latency according to either Case 1 or Case 2. We compared results obtained from those two approaches, and found excellent agreements. We have chosen here to present MIMO calculation using Case 1 approach for which it is sufficient to measure only the last cell delay. Also, we found no significant difference in MIMO latency for the switch performing either path or circuit switching, and here we present results when measured frames are transferred over VPC.

Measurements of MIMO latency are performed slightly differently than given in ATMF document [1]. For each test run, first, a sequence of equally spaced 192 cell frames is being sent over the measuring VPC, at a rate of 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 last cells of 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 different subset of the followings:

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

Table 2.1. presents results (average transfer delays in μsec) from 5 test runs.

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

The table above clearly indicates that differences in cell transfer delays are function of the cell 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 MIMO latency calculation for test runs #4 and #5 using the expression (1):
MIMO latency = 192nd Cell transfer delay - $3.33\mu\text{sec}$ = $20.78 - 3.33 = 17.45 \mu\text{sec}$

Note that, for the same test runs from the expression (2), we have:

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

Above indicates that FIFO latency is about 10% off from MIMO latency.

3. MIMO latency with background traffic

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 MIMO latency measurements with background traffic is given in Figure 3.1. The configuration includes one ATM test system and one ATM switch with a 155 Mbps UTP-5 link between the monitor port 1 and the switch port A1 and an 155 Mbps OC-3c link between the monitor port A2 and switch port B1. In addition, we made external loopbacks on 155 Mbps OC-3c multi mode switch ports A2 and A3 (through 10 meter fiber optic cable) and on 155 Mbps UTP-5 switch ports B2 and B3 (through connectors)

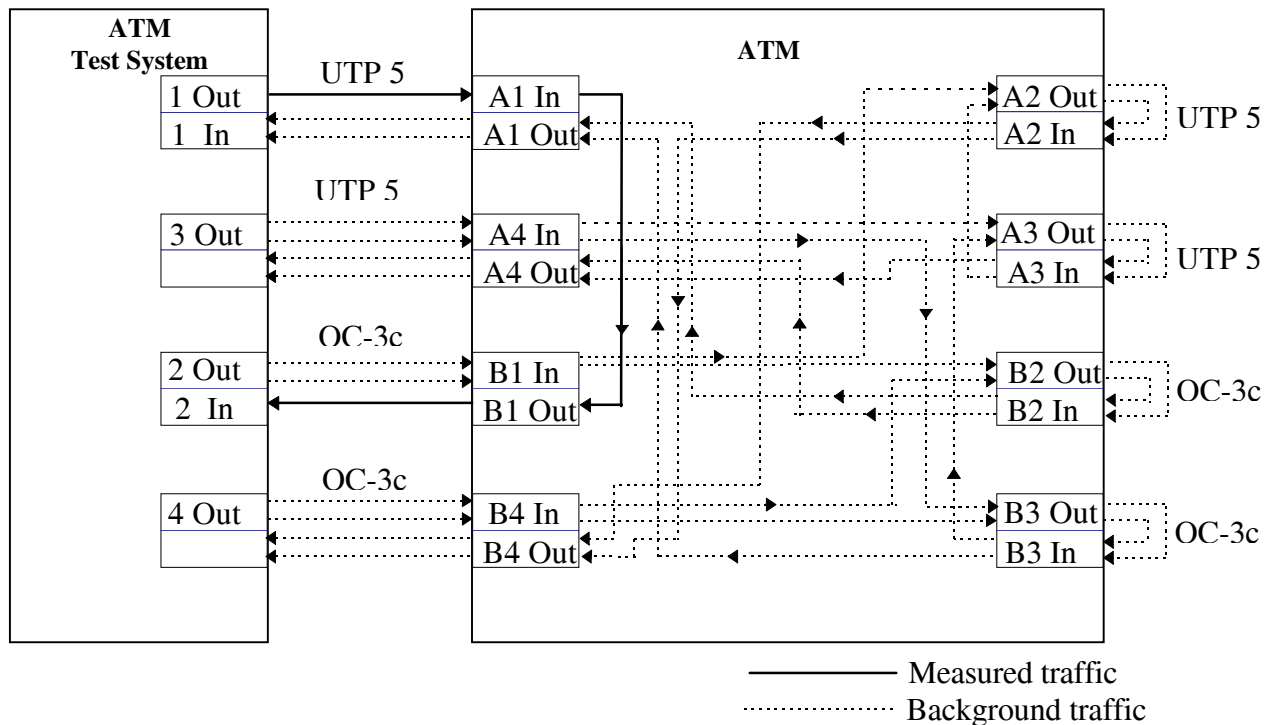


Figure 3.1:

Test configuration for measurements of MIMO latency with background traffic

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

Background traffic is generated from three monitor ports (ports 2, 3, and 4) through 6 VPC's (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 the two ports used for measured traffic). The link between the monitor port 1 and switch port A1 is used to transfer measured 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 measured traffic in one direction and background traffic in the another direction. Note that the measured traffic does not share any generator or analyzer logic with any

other VCs in the same direction. This avoids distortions in the measured traffic that could be caused otherwise.

With loop-back connections and background traffic VPC's as given 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) Then one or more of the following test runs are performed. For each test run, first, a sequence of equally spaced frames of the given length is sent over the measuring 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 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 Case 1 approach for which it is sufficient to measure only the last cell delay. In each test run, besides the average delay of the last cell of all frames (1,000 samples), we also record the average transfer delays of the first cells of all frames (1,000 samples), so we can calculate FIFO latency, and the average transfer delays of all cells between the first and the last cells of all frames (190,000 samples).
- c) Background load is increased and steps a) and b) are repeated. We stop at 100% of MBL (or very close to it).

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

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

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

3.3. Measurement Results

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

Table 3.2. presents results for UBR background traffic with burst size of 1,000 cells. The length of measured frames equals 192 cells.

Table 3.2.

Load %	1st cell	2nd -999th 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 3.3. presents results for UBR background traffic with burst size of 384 cells. The length of measured frames equals 192 cells

Table 3.3.

Load %	1st cell	2nd -999th 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 3.4. presents results for CBR background traffic. The length of the measured frames equals 192 cells.

Table 3.4.

Load %	1st cell	2nd -999th 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 measured traffic observed

3.4. Analysis

A number of observations and conclusions can be made based on measurement results.

- 1) 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. For 90% or higher loads, we observe losses in the measured traffic. All of these apply for all lengths of measured frames. 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.

- 2) MIMO latency results are quite different from FIFO latency results. For cases with measured frames of larger length (1,000 cells) and UBR background traffic with longer size bursts (2,004 cells), MIMO latency increases 100% and more for loads of 70% or higher of MBL.

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

For cases with measured frames of shorter length (192 cells) and UBR background traffic with shorter size bursts (384 cells), MIMO latency does not change significantly with increase of background traffic load.

For cases with measured frames of shorter length (192 cells) and CBR background traffic, there is no significant change in MIMO latency with increase of background traffic load.

- 3) From Tables 3.1 - 3.4, it is interesting to compare the average transfer delay of various cells of a frame. The first cell has the lowest delay while the last cell has the highest. As 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 statistically not very meaningful.

4. Throughput measurement.

4.1 Configuration

In throughput measurements, we used an n-to-1 configuration 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 were able to perform only the 4-to-1 tests. We also performed 2-to-1 and 3-to-1 tests, but the results were similar to those reported here for the 4-to-1 case.




Figure 4.1: 4-to-1 traffic pattern

The test configuration for throughput measurements, corresponding to the 4-to-1 traffic pattern is given in Figure 4.2. The configuration includes one ATM test system and one ATM switch with two 155 Mbps UTP-5 links and two 155 Mbps OC-3c multimode fiber 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.

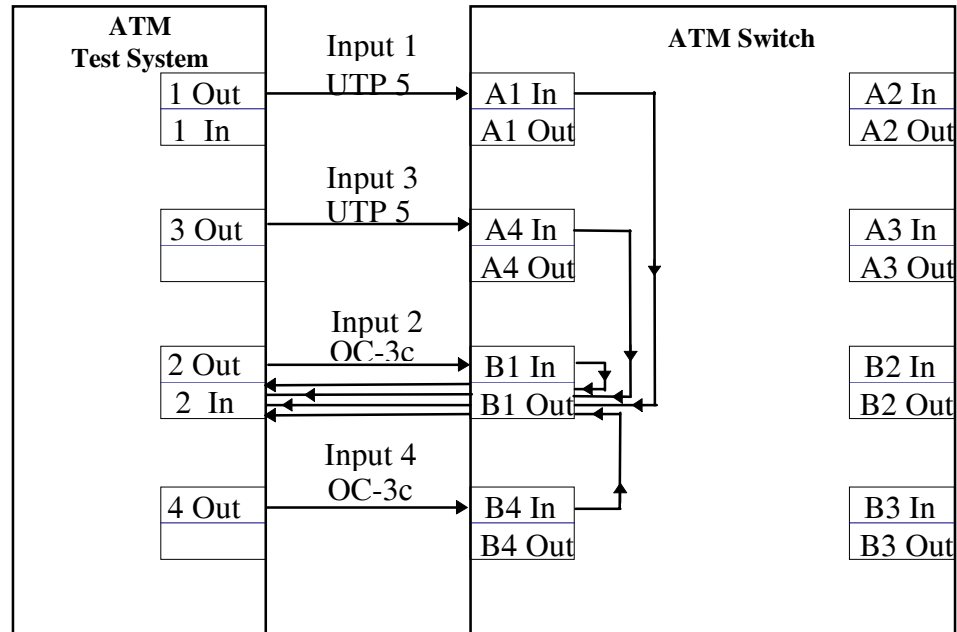


Figure 4.2. Test configuration for throughput measurements.

4.2 Methodology, Measurement Results and Analysis

Four traffic sources generate, over corresponding VPC's, frames at identical rates with equally spaced frames. All frames are generated in AAL 3/4 CPCS-PDU format, and each PDU is segmented in 106 cells. We had to use AAL 3/4 because our monitor allows only one AAL5 VCs on any one port.

The input load is varied by changing the frame rate. 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 examples no loss is detected even when load in 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.

When 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, frame loss ratio and cell mis-insertion rate in cells/sec. Incidentally, cell mis-insertion is defined as delivering cells that do not belong to the VC. This may happen, for example, when there are errors in the cell header.

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

Metric	Input 1	Input 2	Input 3	Input 4
Cell Loss Ratio	0.0036	0.0022	0.0033	0.0026
Frame Loss Ratio	0.2620	0.2050	0.2890	0.2260
Cell Mis-ins. Rate [cell/sec]	0.0000	0.0000	0.0000	0.0000

The 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
Cell Loss Ratio	0.0637	0.0520	0.0630	0.0771
Frame Loss Ratio	0.7340	0.7350	0.6310	0.8760
Cell Mis-ins. Rate[cell/sec]	146	117	137	181

The 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
Cell Loss Ratio	0.3530	0.3500	0.3500	0.3500
Frame Loss Ratio	1.0000	1.0000	0.9890	0.9890
Cell Mis-ins. Rate[cell/sec]	1800	1790	1790	1820

From Table 4.1, it can be observed that 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 order of magnitude larger than the cell loss ratio. It is noted that frame loss rate also varies considerably between four traffic sources (within the range 20%-29%) resulting in some unfairness.

From Table 4.2, it can be observed that with offered load of 20% over output link rate, the frame loss ratio is considerable. Approximately, 63% to 87 % of input frames are lost. The mis-inserted cell rate is also high.

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

In conclusion, for the n-to-1 traffic pattern the lossless throughput for the switch under test is 155 Mbps or 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 can conclude that the full load throughput for this traffic pattern does not make sense, because in this case, practically all the frames are lost.

5. Summary

1. It is possible to compute MIMO frame latency with current ATM monitors that give only cell level statistics.
2. The cell transfer delays of various cells of a frame are widely different. First cell of a frame has much lower latency than later cells. Therefore cell transfer delay is not statistically meaningful.
3. 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 may provide a useful background load at the same priority as the measured traffic. While CBR can be used as a higher priority background load.
4. 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 the port with the measured traffic. This configuration and other similar to it are needed to be added for delay measurements.
5. Peak load is equal to lossless throughput. If we find the same pattern on many switches, then it may be wise to remove one of the two metrics.
6. Variance in the throughput is negligible and so we may remove the requirement for specifying standard error of throughput.
7. For 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, 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.
8. Throughput for various VCs is identical as long as there is no loss. Thus, fairness of throughput is not a useful metric.
9. The frame loss rate for different VCs in a n-to-1 configuration is different. 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/>