

CONTRIBUTION TO T1 STANDARDS PROJECT

 STANDARDS PROJECT: Specification and Allocation of ISDN Performance (T1Q1-10)

TITLE: Frame Delay Through ATM Switches: MIMO Latency

AUTHOR: Gojko Babic, Raj Jain, Arjan Durresi

SOURCE: Ohio State University

CONTACT: Raj Jain

Raj Jain is now at
Washington University in Saint Louis
Jain@cse.wustl.edu
<http://www.cse.wustl.edu/~jain/>

tion Science
 , OH 43210-1277
 911 (Fax), jain@cis.ohio-state.edu

DATE: October 26, 1998

DISTRIBUTION: Working Group T1A1.3

ABSTRACT: This contribution addresses the problem of measuring frame latency in ATM switches. The frames consisting of several ATM cells may arrive with numerous gaps between cells. It is important that the gaps present in the input stream be not counted towards the switch's contribution to the frame delay. The proposed solution called "MIMO" (Message-In Message-Out) latency improves upon FILO (First-In Last-Out) latency commonly used for continuous frame technologies such as frame relay.

Briefly, the MIMO latency is defined as the difference between FILO latency through the switch and that through an ideal switch. The definition and the discussion also apply to any network of switches as well.

NOTICE

This document has been prepared to assist the Standards Committee T1 - Telecommunications. It is offered to the Committee as a basis for discussion and is not a binding proposal on Ohio State University. The requirements presented in this document are subject to change in form and numerical value after more study. Ohio State University specifically reserves the right to add to, or amend, the statements contained herein.

1. Problem Statement

The performance of ATM equipment and the quality of services have been defined in terms of cell-level metrics such as cell transfer delay (CTD), cell delay variation (CDV) and cell loss ratio (CLR). However, cell-level metrics do not very often reflect the performance as experienced (or desired) by end users. For example, a video user sending 30 frames/sec would like frames to be completely delivered every 33 ms and it does not matter whether the cells belonging to a frame arrive back-to-back or regularly spaced. Thus, it is the frame delay and its variation that matters, not CTD and CDV.

A frame is defined here as the ATM Adaptation Layer (AAL) protocol data unit (PDU). One problem in measuring the frame delay in ATM networks is that when seen inside the network, the frames may be discontinuous with numerous gaps between the cells as well as cells of other frames. Note that the monitoring equipment, if placed inside the host, will be affected by the performance of the host and may not accurately reflect the performance of the switch. Thus, the test probes of the monitoring equipment should be placed at the entrance and the exit of the system to be measured, as in Figure 1.




Figure 1. Measurement Point

Although we use the term "switch" throughout this contribution, the discussion applies equally well to any network element (including switches, routers, multiplexers, inverse-multiplexers, wires, etc) or a network as whole.

The delay of switch at the cell level is generally measured by FILO (first-bit in to the last-bit out) latency as indicated in Figure 2. Other alternatives such as FIFO (first-bit in to the first-bit out), LILO (last-bit in to the last-bit out), and LIFO (last-bit in to the first-bit out) latencies can be easily obtained from the figure. Most ITU documents measure cell level delay using FILO metric. Therefore, we will use FILO for our discussion.




Figure 2. FILO latency at cell level

One way to measure switch delay at the frame level is to measure the delay between the first bit in and the last bit out events for the frame. This is so called FILO latency introduced by the switch at the frame level. For example, consider a frame consisting of two cells as shown in Figure 3. Let us assume that the input link rate is identical to output link rate and that the cell input or output time is 1 ms. The two cells arrive with a gap of 1 ms. The switch A introduces a delay of 5 ms to each cell. As a result, the FILO latency (interval between the first bit in and the last bit out events of the frame) is 8 ms.




Figure 3. FILO latency for the switch A that delays each cell by 5 ms

Generally, the measured performance of a system depends upon the system as well as the workload. Some metrics are highly workload dependent while others are less dependent. A metric, which depends more on the system and less on the workload, is generally preferred particularly if the users are interested in comparing the systems and not the workloads. It turns out that the FILO frame latency as defined above has the undesirable property that it depends heavily on the workload. For example, see Figure 4. Here the two cells of the frame arrive with a gap of 5 ms, the switch B delays each cell by 1 ms. The FILO frame latency is 8 ms, which is the same as in Figure 3 for switch A. Clearly, switch B is better, but the FILO latency does not reflect that fact.




Figure 4. FILO latency for the switch B that delays each cell by 1 ms

To show the problem in its extreme case, consider the situation in Figure 5, where the two cells of the frame arrive two days apart. Switch B delays each cell by 1 ms. But the FILO

frame latency is 2 days plus 3 ms. It mostly reflects the arrival gap and is nowhere close to the actual delay introduced by the switch.




Figure 5. FILO latency for the switch B that delays each cell by 1 ms

In this contribution, we propose a new metric called MIMO (Message In Message Out) latency that measures the true contribution of the switch to the frame latency and is not affected by the arrival patterns (gaps) of the cells constituting the frame. We introduce the concept of an ideal switch that does the best possible processing of its frames. MIMO latency is calculated for any given arrival pattern as the FILO frame latency for the pattern through the ideal switch ($FILO_0$) subtracted from the measured FILO frame latency of the switch under test gives, i.e.:

$$\text{MIMO latency} = \text{FILO latency} - FILO_0 \quad (1)$$

The concept of ideal switch is defined more concretely later in this contribution. For the examples discussed so far a wire of zero length can serve as an ideal switch. With this wire, the bits depart as soon as they arrive. In the example shown in Figure 3, the FILO frame latency through the ideal switch, $FILO_0$, is 3 ms and so MIMO latency for the switch A is 8-3 or 5 ms. Similarly, for the example shown in Figure 4, $FILO_0$ is 7 ms and MIMO latency for the switch B is 1 ms. Finally, for the example shown in Figure 5, $FILO_0$ is 2 days plus 2 ms and so MIMO latency for the switch B is again 1 ms. Notice that in each case, MIMO latency reflects the switch behavior and is not affected by the arrival pattern.

In Section 2 of this contribution we present a more rigorous definition of the MIMO latency. The ideal switch is defined in Section 3. Section 4 presents some of our measurement tests of MIMO latency.

2. MIMO Latency Definition

As discussed above, MIMO latency is defined as:

$$\text{MIMO latency} = \text{FILO latency} - FILO_0$$

$FIL0_0$ for a given frame is equal to the FILO latency of that frame passing through *an ideal switch*. An ideal switch is defined as a switch that handles incoming frames in such way that they are transmitted on the output link without any unnecessary time consumption, i.e. the best any switch can do. By definition, MIMO latency for an ideal switch is zero. Hence, an ideal switch can also be called a zero-delay switch.

The procedure for $FIL0_0$ calculation is as follows:

- Initially $FIL0_0 = 0$ and time t is measured from the arrival of the first bit of the first cell.
- For each cell with its first bit arriving at time t , update $FIL0_0$ as follows:

$$FIL0_0 = \max\{t, FIL0_0\} + \max\{CIT, COT\}$$

where:

CIT = cell input time = 424 bits / Input Link Rate in bps

COT = cell output time = 424 bits / Output Link Rate in bps

Note that MIMO latency, as a switch delay metric, accounts only for delays caused by node processing, such as switching, routing and queuing delays, and not by transmission delays introduced by communication links. MIMO latency is not limited to ATM switches and it applies to all types of communication devices, including multiplexers, store-and-forward or cut-through bridges, routers, repeaters, wires, or any combination of these.

3. Cell and Frame Latency through an Ideal Switch

The concept of an ideal switch is explored in depth in this section. Figure 6 illustrates how an ideal switch would handle a cell. The switch behavior depends upon the relationship between the input and output link rates. In the case when the input link rate is equal to the output link rate, as presented in Figure 6a, an ideal switch transmits each bit as soon as it arrives. Thus, each bit of the cell experiences zero latency in an ideal switch.





Figure 6a. Cell Processing of an Ideal Switch for Input Rate = Output Rate

Figure 6b illustrates the case when the input link rate is higher than the output link rate. In this case, outputting (transmitting) a bit takes longer than inputting it. The ideal switch can transmit only the first bit as soon as it is received. The other bits of the cell can not be transmitted immediately as they arrive, because the transmission of all previously received bits has not yet finished. Bits at the end of the cell wait longer than bits at the beginning. Thus, an ideal switch in this situation should be intelligent to do appropriate buffering of incoming bits.






Figure 6b. Cell Processing of an Ideal Switch for Input Rate > Output Rate

Figure 6c illustrates the case when the input link rate is lower than the output link rate. An ideal switch does not start transmission of the first bit immediately after it is received, but after an appropriate delay. Bits at the beginning of the cell are delayed more than bits at the end, with larger delays for slower output link rates. Only the last bit of a cell has no delay and it is transmitted immediately upon its arrival. Thus, an ideal switch should be intelligent to avoid under-runs by appropriately delaying the transmission of incoming bits.




Figure 6c. Cell Processing of an Ideal Switch for Input Rate < Output Rate

It should be easily realized that the illustrations in Figures 6 apply not only to cells, but also to contiguous frames. Note that none of the usual latencies (FILO, LILO, FIFO, or LIFO) has a zero value in all three cases, as it should be for delays of a cell (frame) passing through an ideal switch.

The rest of this section considers how an ideal switch handles discontinuous frames in an ATM environment.

Figures 7 present two possible cases of a frame passing through an ideal switch with the input link rate higher than the output link rate. Figure 7a illustrates the case when cells of a frame do not have to wait. The given frame includes two cells and the input link rate is 4 times the output link rate. The two cells start arriving at time $t = 0$ and $t = 5$, respectively. An ideal switch will start transmitting the first cell at time $t = 0$ and finish at time $t = 4$. The second cell can be transmitted without waiting and the transmission is finished at $t = 9$. This is how long an ideal switch will take to transmit this frame. Hence, FILO latency of an ideal switch for this frame is 9. This is $FILLO_0$ for the given input pattern, and the same value is obtained using the procedure defined in Section 2.




Figure 7a. No-Cell-Waiting Operation of an Ideal Switch for Input Rate > Output Rate

Figure 7b shows the another possible case of a frame passing through an ideal switch with an input link rate higher than the output link rate when cells of a frame have to wait. As in Figure 7a, the given frame has two cells and the input link rate is 4 times the output link rate. However, the frame has a different gap pattern. The second cell arrives at time $t = 2$ and thus has to wait. An ideal switch will start transmitting the first cell at time $t = 0$ and finish at time $t = 4$. The second cell transmission starts at $t = 4$ and it is finished at $t = 8$. Hence, FILO latency of an ideal switch for this frame is 8, i.e. $FILO_0 = 8$.




Figure 7b. Cell-Waiting Operation of an Ideal Switch for Input Rate > Output Rate

Thus, Figures 7 illustrate possibilities that an incoming cell can be transmitted immediately without waiting and that an incoming cell has to wait for previously received cells of the same frame to be transmitted.

In general, for a given discontinuous frame when the input link rate is higher than the output link rate, it is possible that some cells have to wait on previously received cells of the same frame, while some cells can be transmitted without waiting. Also, notice that ideal switch on output decreases the size of each gap from input, with some gaps being completely removed.

Figure 8 illustrates the only possible case of a frame passing through an ideal switch with an input rate lower than the output rate. Again, the frame includes two cells but the output link rate is now four times the input link rate. The two cells arrive at time $t = 0$ and $t = 5$, respectively. An ideal switch will start transmitting the first cell at time $t = 3$ (not at

$t = 0$, in order to avoid an underrun), and finish at time $t = 4$. The second cell transmission starts at $t = 8$ and finishes at $t = 9$. This is how long an ideal switch will take to transmit this frame. Hence, the FILO latency of an ideal switch for this frame is 9, i.e. $FILO_0 = 9$.




Figure 8. Operations of an Ideal Switch for Input Rate < Output Rate

Figure 9 illustrates the only possible case of a frame passing through an ideal switch with an input rate equal to the output rate. Again, the frame includes two cells. The two cells arrive at time $t = 0$ and $t = 5$, respectively. An ideal switch will start transmitting the first cell at time $t = 0$ and finish at time $t = 1$. The second cell transmission starts at $t = 5$ and finishes at $t = 6$. This is how long an ideal switch will take to transmit this frame. Hence, the FILO latency of an ideal switch for this frame is 6, i.e. $FILO_0 = 6$.




Figure 9. Operations of an Ideal Switch for Input Rate = Output Rate

Note that in the cases when the input rate is less than or equal to the output rate, a cell never has to wait for completion of transmissions of previously received cells. The $FILO_0$ in such cases is equal to the frame input time (first-bit-in to the last bit in) and MIMO latency becomes equal to the delay of the last bit of the last cell, i.e. LILO latency. Thus, when input link rate \leq output link rate, we have:

$$\text{MIMO latency} = \text{LILO latency} \quad (2)$$

4. Measurement Experiences

In this section we describe several measurements performed in our performance laboratory using a commercial available 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 at the cell level.

Here are some observations about ATM monitors:

- The cell transfer delay (CTD) is measured, as defined in many current standards, as FILO latency. Also, ATM monitors measure FILO cell delay with a finite granularity. Our ATM monitor has a resolution of 0.5 μ s. We obtained the average cell transfer delay of 3.33 μ s 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 μ s) 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 μ s. 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 between any two cells is defined as the time between arrival of the last bit of the first cell and the last bit of the second cell. The specified resolution was 0.5 μ s. However, 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 μ s, 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.)

The following two relations, which can be easily derived, are used later in this section for MIMO latency calculation:

$$\text{FILO latency} = \text{First cell to last cell inter-arrival time at the output} \\ + \text{First cell transfer delay}$$

(3)

$$\text{LIFO latency} = \text{Last cell transfer delay} - \text{Cell input time}$$

(4)

4.1. Tests with Input Rate equal to Output Rate

The test configuration for MIMO latency measurements is shown in Figure 10. The configuration includes one ATM monitor and one ATM switch with 155 Mbps UTP-5 link between the monitor port 1 and the switch port A1 and 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. A permanent virtual channel connection (VCC) is

established between the monitor ports 1 and 2 through the switch ports A1 and B1. That VCC is used for transmission of frames whose latency is measured. Figure 10 also indicates the traffic flow direction.




Figure 10. Test configuration for measurements of MIMO latency

For each test run, a sequence of equally spaced 192 cell frames (cells of each frame are generated back-to-back) was sent over the VCC at a rate of 4.63 frames/s. After the flow had been established, we recorded 20.78 μ s as the average transfer delays of the last cells in the next 1,000 consecutive frames.

Since in this configuration:

- $CIT = 424[\text{bits}] / \text{Input Link Rate} = 424[\text{bits}] / 149.76 [\text{Mbps}] = 2.83 \mu\text{s}$

the average MIMO latency calculation using expressions (2) and (4) is given as:

$$\text{MIMO latency} = \text{Last cell transfer delay} - \text{CIT} = 20.78 - 2.83\mu\text{s} = 17.95 \mu\text{s}$$

Table 1 presents measurement data for two randomly chosen frames and calculated MIMO latency. When the input link rate is equal to the output link rate, there are two expressions (equation 1 and Equation 2) for MIMO latency calculation. As shown in the fourth and sixth columns of Table 1, both expressions provide the same values (within the 0.5 μ s resolution).

Table 1. (all times in μ s)

Last cell CTD	1 st cell CTD	1 st cell to last cell inter-arrival time	MIMO latency (2)	FILO latency (3)	MIMO latency (1)
21.5	21.5	541.0	18.67	562.5	18.91
21.0	18.5	543.5	18.17	562.0	18.41

Note that if the frames of a cell arrive back to back at the output, $FILO_0$ can be calculated simply as follows:

$$FILO_0 = \text{Frame size} / \text{Output rate} = 192 \text{ cells} / 353,207.55 \text{ cells/s} = 543.59 \mu\text{s}$$

4.2. Tests with Input Rate Higher Than Output Rate

The test configuration for the MIMO latency measurements for the case with the input link rate higher than the output link rate, shown in Figure 11. It uses a 155 Mbps UTP-5 link between the monitor port 1 and the switch port A1 and a 25 Mbps link between the monitor port 2 and the switch port D1. Figure 11 also indicates the traffic flow direction.




Figure 11. Test configuration for measurements of MIMO latency

In this configurations:

- CIT = 2.83 μ s
- COT = 424[bits] / Output Link Rate = 424[bits] / 25.6 [Mbps] = 16.56 μ s

We performed all our tests with 32-cell frames. One of the measurements used contiguous frames, i.e. cells of the test frame were transmitted back-to-back. In the rest of the tests, we introduce identical gaps (unassigned cells or cells of other frames) between cells of the test frame.

Table 2 presents measurement results for eight test runs, from which MIMO latency is calculated. The first test uses a contiguous test frame on input. All other tests use discontinuous frames on input, with gaps between cells of the test frame, as indicated in the second column. Our tests do not show any significant difference if gaps include unassigned cells or cells of other frames, which leave the switch through output links other than the one used by the test frames.

Table 2: (All times are in μ s)

Test No.	Frame Pattern	1 st cell CTD	1 st cell to last cell inter-arrival time	FILO ₀	FILO latency (3)	MIMO Latency (1)
----------	---------------	--------------------------	--	-------------------	------------------	------------------

1	No gap	36.8	526.5	530.0	563.3	33.3
2	1-cell gaps	35.8	526.0	530.0	561.8	31.8
3	2-cell gaps	36.8	526.0	530.0	562.8	32.8
4	3-cell gaps	34.8	526.5	530.0	561.3	31.3
5	4-cell gaps	40.8	519.5	530.0	560.3	30.3
6	5-cell gaps	36.8	526.5	542.9	562.8	19.9
7	6-cell gaps	36.8	616.0	630.6	652.8	22.2
8	7-cell gaps	35.3	705.0	718.4	740.3	21.9

The third and fourth columns present measurement results for the first cell delay and inter-arrival time between the first and the last cells. The fifth column includes calculated values for $FILO_0$, as explained in Section 2, given a frame pattern on input. Here is how we calculate those values. For the first five tests, it can be found that each cell entering an ideal switch has to wait for transmission of the previously received cell to finish. Thus, on output we should have back-to-back cells, i.e. a contiguous frame. Therefore, we can calculate $FILO_0$ for 32-cell frames in all those cases as:

$$FILO_0 = 32 \cdot COT = 32 \cdot 16.56 = 530 \text{ ns}$$

In the last three tests, the gaps on input are large enough that no cells have to wait on a previously received cell. In the case with 5-cell gaps, the first bit of the 32nd (last) cell arrives at an ideal switch at time t , where

$$t = (CIT + 5\text{-cell gap}) \times 31 = 6 \text{ CIT} \times 31 = 526.4 \text{ } \mu\text{s}$$

and then

$$FILO_0 = t + COT = 526.4 + 16.5 = 542.9$$

Similarly in the cases with 6-cell gaps and 7-cell gaps, $FILO_0$ is calculated as 630.6 μs and 718.4 μs , respectively.

The sixth column shows FILO latency calculated, according to the expression (3) as the sum of terms in the third and the fourth column. In the last column, according to the expression (1), MIMO latency values are obtained subtracting terms in the fifth column from terms in sixth column.

Note that the switch latency is higher in the first 5 tests due to cell queueing. In the last three tests, the gap between the cells is large and there is no queueing. MIMO latency clearly reflects this effect.

4.4. Tests with Input Link Rate Lower Than Output Link Rate

We also performed tests using the configuration in Figure 11, but with the traffic flow in the opposite direction as indicated in the figure. Thus, this is the configuration with the input link rate lower than the output link rate. In this case, we have:

- CIT = cell input time = 16.56 μ s
- COT = cell output time = 2.83 μ s

We performed tests with 32-cell frames, with random idle periods between cells. Table 3 includes measurement data from two tests for which MIMO latency is also calculated. Since the input link rate is lower than the output link rate, both the expression (1) and the expression (2) can be used to calculate MIMO latency.

Table 3. (All times are in μ s)

Last cell CTD	MIMO latency (2)	1 st cell CTD	1 st cell to last cell inter-arrival time	FILO ₀	FILO Latency	MIMO latency (1)
32.0	15.44	31.0	535.0	550.0	566.0	16.0
32.5	15.94	33.0	1067.5	1082.6	1100.5	17.9

The results in Table 3 show clearly that MIMO latency reflects the switch behavior and is not affected by the arrival pattern. On the other hand, it is shown that FILO latency is strongly affected by the arrival pattern. It can be observed that good agreement of MIMO latency values can be obtained using any of the two expressions for its calculation.

References:

[BAB97a] G. Babic, A. Durresi, R. Jain, J. Dolske, S. Shahpurwala, "ATM Switch Performance Testing Experience," ATM_Forum/97-0178R1, April 1997, <http://www.cis.ohio-state.edu/~jain/atmf/a-0178r1.htm>

[BAB97b] G. Babic, A. Durresi, J. Dolske, R. Jain, "Measurement Experience with the Revised MIMO Latency Definition," ATM_Forum/97-0859, September 1997, <http://www.cis.ohio-state.edu/~jain/atmf/a97-0859.htm>

[BRA91] S. Bradner, "Benchmarking Terminology for Network Interconnection Devices," RFC 1242, July 1991

[CCI92] CCITT Recommendation X.135, "Speed of Service (Delay and Throughput) Performance Values for Public Data Networks when Providing International Packet Switched Service," 1992

[ITU95] ITU-T Recommendation I.356, "B-ISDN ATM Layer Specification," ITU-Study Group 13, Geneva, 1995