
ATM Forum Document Number: ATM_Forum/96-0811R1

Title: Frame-level throughput and latency metrics - proposed text.

Abstract:

In the April meeting we presented definitions for throughput and latency, which were generally agreed. This contribution proposes exact text for inclusion in the performance benchmarking document.

Source:

Raj Jain, Gojko Babic, Bhavana Nagendra
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: June 1996

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

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.

In AF-TEST April 96 meeting in Alaska, we presented definitions for two throughput metrics and one latency metrics [1]. There was general agreement about the metrics. Also, It was suggested that we add one more throughput metrics. Based on this discussion, it is proposed that the following be considered for addition to the performance benchmarking document:

Performance Metrics:

In the following description System Under Test (SUT) refers to an ATM switch. However, the definitions and measurement procedures are general and may be used for other devices or a network consisting of multiple switches as well.

1. THROUGHPUT

1.1 Definitions:

There are three frame-level throughput metrics that are of interest to a user.

- i. **Lossless throughput** - It is the maximum rate at which none of the offered frames are dropped by the SUT.
- ii. **Peak throughput** - It is the maximum rate regardless of frames dropped at which the SUT operates. The maximum rate can actually occur when the loss is not zero.
- iii. **Full-load throughput** - Its the rate at which SUT operates when the input links are loaded at 100% of their capacity.

A model graph of throughput vs input rate is shown in Figure 1. Level X defines the loss-less throughput, level Y defines the peak throughput and level Z defines the full-load throughput.




Figure 1: Peak, lossless and full-load throughput

The lossless throughput is the highest load at which the count of the output frames equals the count of the input frames.

Peak throughput is the maximum throughput that can be achieved in spite of the losses. Full-load throughput is the throughput of the system at 100% load on input links. Note that the peak throughput may equal the lossless throughput in some cases.

Only frames that are received completely without errors are included in frame-level throughput computation. Partial frames and frames with CRC errors are not included.

1.2 Units:

Throughput should be expressed in bits/sec. This is preferred over specifying it in frames/sec or cells/sec. Frames/sec requires specifying the frame size. The throughput values in frames/sec at various frame sizes cannot be compared without first being converted into bits/sec. Cells/sec is not a good unit for frame-level performance since the cells aren't seen by the user.

1.3 Statistical Variations:

The tests should be run NRT times for TRT seconds each. Here NRT and TRT are parameters. These and other such parameters and their default values are listed later in Table 2.

If T_i is the throughput in i th run, The mean and standard errors of the measurement should be computed as follows:

$$\text{Mean throughput} = (\sum T_i)/n$$

$$\text{Standard deviation of throughput} = (\sum (T_i - \text{Mean throughput})^2)/(n-1)$$

$$\text{Standard error} = \text{Standard deviation of throughput}/\sqrt{n}$$

Given mean and standard errors, the users can compute an $100(1-\alpha)$ -percent confidence interval as follows:

$$100(1-\alpha)\text{-percent confidence interval} = (\text{mean} - z \times \text{std error}, \text{mean} + z \times \text{std error})$$

Here, z is the $(1-\alpha/2)$ -quantile of the unit normal variate. For commonly used confidence levels, the quantile values are as follows:

Confidence	α	Quantile
90%	0.1	1.615
99%	0.01	2.346
99.9%	0.001	3.291

1.4 Traffic Pattern:

The input traffic will consist of frames of length FSA bytes each. Before starting the throughput measurements, all required VCs will be set up (for an n -port SUT) in one of the following four configurations:

1. **n-to-n straight:** All frames input from port i exit to port $i+1$ modulo n . This represents almost no path interference among the VCs. Total n VCs.

2. **n-to-n cross:** Input from port each port is divided equally to exit on each of the n output ports. Total n^2 VCs.

3. **n-to-1**: Input from all ports is destined to one output port. Total n VCs.

4. **1-to-n**: Input from a port is multicast to all output ports. Total 1 VC.

The frames will be delivered to the layer under test equally spaced at a given input rate. The rate at which the cells reach SUT may vary depending upon the service used. For example, for ABR traffic, the allowed cell rate may be less than the link rate in some configurations.

At each value of the input rate to the layer under test, the total number of frames sent to SUT and received from SUT are recorded. The input rate is computed based on the time from the first bit of first frame enters the SUT to the last bit of the last frame enters the SUT. The throughput (output rate) is computed based on the time from the first bit of the first frame exits the SUT to the last bit of the last frame exits SUT.

If the input frame count and the output frame count are the same then the input rate is increased and the test is conducted again. The lossless throughput is the highest throughput at which the count of the output frames equals the count of the input frames. If the input rate is increased even further, although some frames will be lost, the throughput may increase till it reaches the peak throughput value after which the further increase in input rate will result in a decrease in the throughput. The input rate is increased further till 100% load is reached and the full-load throughput is recorded.

1.5 Background Traffic:

The tests can be conducted under two conditions - with background traffic and without background traffic.

Higher priority traffic like VBR can act as background traffic for the experiment. Further details of measurements with background traffic (multiple service classes simultaneously) are to be specified. Until then all benchmarking will be done without any background traffic.

2. FRAME LATENCY:

2.1 Definition:

The frame latency for a system under test is measured using a "Message-in Message-out (MIMO)" definition. Succinctly, MIMO latency is defined as follows:

MIMO Latency = $\text{Min}\{\text{First-bit in to last-bit out latency} - \text{nominal frame output time}, \text{last-bit in to last-bit out latency}\}$

An explanation of MIMO latency and its justification is presented in Appendix A.

To measure MIMO latency, a sequence of equally spaced frames are sent at a particular rate. After the flow has been established, one of the frames in the flow is marked and the time of the following four events is recorded for the marked frame while the flow continues unperturbed:

1. First-bit of the frame enters into the SUT
2. Last-bit of the frame enters into the SUT
3. First-bit of the frame exits from the SUT
4. Last-bit of the frame exits from the SUT

The time between the first-bit entry and the last bit exit (events 1 and 4 above) is called first-bit in to last-bit out (FILO) latency. The time between the last-bit entry to the last-bit exit (events 2 and 4 above) is called last-bit in to last-bit out (LILO) latency. Given the frame size and the nominal output link rate, the nominal frame output time is computed as follows:

$$\text{Nominal frame output time} = \text{Frame size} / \text{Nominal output link rate}$$

Substituting the FILO latency, LILO latency, and Nominal frame output time in the MIMO latency formula gives the frame level latency of the SUT.

2.2 Units:

The latency should be specified in micro-seconds.

2.3 Statistical Variations:

NML samples of the latency are obtained by sending NML marked frames at $TTL/(NML + 1)$ intervals for a total test duration of TTL seconds. Here, NML and TTL are parameters. Their default values are specified in Table 2. The mean and standard errors computed (in a manner similar to that explained in Section 1.3 for Throughput) from these samples are reported as the test results.

2.4 Traffic Pattern:

The input traffic will consist of frames of length FSA bytes. Here, FSA is a parameter. Its default value is specified in Table 2.

Before starting the throughput measurements, all required VCs will be set up (for an n-port SUT) in one of the following configurations:

1. **n-to-n straight:** All frames input from port i exit to port $i+1$ modulo n. This represents almost no path interference among the VCs.
2. **n-to-n cross:** Input from port each port is divided equally to exit on each of the n output ports.
3. **n-to-1 :** Input from all ports is destined to one output port.
4. **1-to-n:** Input from a port is multicast to all output ports. Total 1 VC.

The frames will be delivered to the layer under test equally spaced at a given input rate. For latency measurement, the input rate will be set at the input rate corresponding to the lossless throughput. This avoids the problem of lost marked cells and missing samples.

2.5 Background Traffic:

The tests can be conducted under two conditions - with background traffic and without background traffic.

Higher priority traffic like VBR can act as background traffic for the experiment. Further details of measurements with background traffic (multiple service classes simultaneously) are to be specified. Initially all tests will be conducted without the background traffic.

3. REPORTING RESULTS

The throughput and latency results will be reported in a tabular format as follows:

Table 1: Tabular format for reporting performance benchmarking results

	Throughput							
Traffic	Lossless		Peak		Full-load		Latency	
Pattern	Mean	Std Err	Mean	Std Err	Mean	Std Err	Mean	Std Err
n-to-n Straight								
n-to-n Cross								
n-to-1								
1-to-n								

4. DEFAULT PARAMETER VALUES

The default values of the parameters used in performance benchmarking are listed in Table 2.

Table 2: List of Parameters and their default values

Parameter	Meaning	Default
NRT	Number of repetitions of throughput experiments	30
TRT	Time of each repetition of throughput experiment	60 sec
FSA	Frame Size for AAL performance experiments	9188 Bytes
NML	Number of marked frames sent in latency experiments	30
TTL	Total time of latency experiments	31 sec

APPENDIX A: MIMO LATENCY

The message-in message-out (MIMO) latency is a general definition of latency that applies to a switch or a group of switches when the frames equal to output link rate.

For a single bit, the latency is generally defined as the time from bit in to bit out.




Figure A.1: Latency for single-bit frames

For a multi-bit frame, there are several possible latency definitions. First, consider the case of contiguous frames. All bits of the frames are delivered contiguously without any gap between them. In this case, latency can be defined in one of the following four ways:

1. First bit in to first bit out (FIFO)
2. Last bit in to last bit out (LILO)
3. First bit in to last bit out (FILO)
4. Last bit in to first bit out (LIFO)




Figure A.2: Latency for multibit Frames

If the input link and the output links are of the same speed and the frames are contiguous, the FIFO and LILO latencies are identical. FILO and LIFO latencies can be computed from FIFO (or LILO) given the frame time:

$$\begin{aligned} \text{FILO} &= \text{FIFO} + \text{Nominal frame output time} \\ \text{LIFO} &= \text{FIFO} - \text{Nominal frame output time} \end{aligned}$$

It is clear that FIFO (or LILO) is a preferred metrics in this case since it may be independent of the frame time while FILO and LIFO would be different for each frame size.

Unfortunately, none of the above four metrics apply to an ATM network (or a switch) since the frames are not always delivered contiguously. There may be idle time between cells of a frame. Also, the input and output link may be of different speeds.

In the following we consider twelve different cases. For each case, we compare four possible metrics (FIFO, LILO, FILO-nominal frame output time, and MIMO) and show that MIMO is the correct metrics in all cases while other metrics apply to some cases but give wrong answers in others.

The twelve cases and the applicability of the four metrics is shown in Table A.1

Table A.1: Applicability of various latency definitions

No.	Case	FIFO	LILO	FILO-NFOT	MIMO
1a	Input rate = output rate, contiguous frame, zero-delay switch	√	√	√	√
1b	Input rate = output rate, contiguous frame, nonzero-delay switch	√	√	√	√
1c	Input rate = output rate, non-contiguous frame, zero-delay switch	This case is not possible			
1d	Input rate = output rate, non-contiguous frame, nonzero-delay switch	×	√	√	√
2a	Input rate > output rate, contiguous frame, zero-delay switch	√	×	√	√
2b	Input rate > output rate, contiguous frame, nonzero-delay switch	√	×	√	√
2c	Input rate > output rate, non-contiguous frame, zero-delay switch	This case is not possible			
2d	Input rate > output rate, non-contiguous frame, nonzero-delay switch	×	×	√	√
3a	Input rate < output rate, contiguous frame, zero-delay switch	×	√	×	√
3b	Input rate < output rate, contiguous frame, nonzero-delay switch	×	√	×	√
3c	Input rate < output rate, non-contiguous frame, zero-delay switch	×	√	×	√
3d	Input rate < output rate, non-contiguous frame, nonzero-delay switch	×	√	×	√

NFOT = Nominal frame output time = Frame size/output link rate

√ => The metric gives a valid result

× => The metric gives an invalid result

CASE 1a: Input Rate = Output Rate, Contiguous Frame, Zero-Delay Switch

One way to verify the validity of a latency definition is to apply it to a single input single output zero delay switch (basically a very short wire). In this case, the bits appear on the output as soon as they enter on the input. All four metrics give a delay of zero and therefore valid.




Figure A.1a: Input Rate = Output Rate, Contiguous Frame, Zero-Delay Switch

Notice that FILO and LIFO will give a non-zero delay equal to frame time. Since we are interested in only switch delay and know that the switch delay in this case is zero, FILO and LIFO are not good switch delay metrics and will not be considered any further.

The nominal frame output time (NFOT) is computed as the frame size divided by the output link rate. It indicates how long it will take to output the frame at the link speed. FILO - NFOT indicates switch's contribution to the latency and is therefore a candidate for further discussion.

CASE 1b: Input Rate = Output Rate, Contiguous frame, non-zero delay switch

Figure A.1b shows the flow in this case.




Figure A.1b: Input=Output, contiguous frame, nonzero-delay

In this case, the total delay FILO can be divided into two parts: switch latency and frame time:

$$\text{FILO} = \text{Switch latency} + \text{Nominal frame output time}$$

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

$$\text{LILO} = \text{FIFO} = \text{FILO} - \text{NFOT}$$

$$\text{MIMO} = \text{Min}\{\text{FILO} - \text{NFOT}, \text{LILO}\} = \text{LILO} = \text{FILO} - \text{NFOT} = \text{FIFO}$$

All four metrics again give identical and meaningful result.

CASE 1c: Input Rate = Output Rate, Non-contiguous frame, Zero-delay Switch

On a zero-delay switch, the bits will appear on the output as soon as they enter the input. Since the input frame is continuous, the output frame will also be contiguous and therefore this case is not possible.

CASE 1d: Input Rate = Output Rate, Non-contiguous frame, Nonzero-Delay Switch

This case is shown in Figure A.1d. There are several gaps between the cells of the frame at the output. FIFO latency does not reflect performance degradation caused by gaps that appear after the first cell. It is, therefore, not a good switch latency metrics..




Figure A.1d: Input rate=output rate, non-contiguous frame, nonzero-delay switch

FILO, LILO, and MIMO are related as follows:

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

Either one of these three metrics can be used as switch latency.

CASE 2a: Input Rate > Output Rate, Contiguous frame, Zero-delay Switch

In this case, the switch consists of a single-input single-output memory buffer. The frame flow is shown in Figure A.2a.




Figure A.2a: Input Rate > Output Rate, Contiguous frame, Zero-delay Switch

For this case, FIFO, FILO, and MIMO are related as follows:

$$\text{LILO} > \text{FIFO} = \text{FILO} - \text{NFOT} = \text{min}\{\text{FILO-NFOT}, \text{LILO}\} = \text{MIMO} = 0$$

In this case, FIFO, FILO-NFOT, and MIMO give the correct (zero) latency. LILO will produce a non-zero result. LILO is affected by the output link speed and does not correctly represent the switch latency.

CASE 2b: Input Rate > Output Rate, Contiguous frame, Nonzero-delay Switch

The frame flow is shown in Figure A.2b.




Figure A.2b: Input Rate > Output Rate, Contiguous frame, Nonzero-delay Switch

Note that the following relationship among various metrics still holds as in case 2a:

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

Thus, LILO gives incorrect answer. It is affected by the output link speed. While the other three metrics give the correct answer.

CASE 2c: Input Rate > Output Rate, Non-contiguous frame, Zero-delay Switch

A zero-delay switch will not introduce any gaps. Thus, this case is not possible.

CASE 2d: Input Rate > Output Rate, Non-contiguous frame, Nonzero-Delay Switch

In this case, (see Figure A.2d)




Figure A.2d: Input Rate > Output Rate, Non-contiguous frame, Nonzero-Delay Switch

In this case, FIFO does not reflect the degradation caused by the gaps and is therefore, not a correct measure of switch latency. It can be made arbitrarily small by delivering the first cell fast but later introducing large gaps.

LILO is affected by the output link speed. It can be made arbitrarily large by decreasing the output rate (and not changing the switch otherwise).

Thus, FILO-NFOT and MIMO are the only two metrics that can be considered valid in this case. Both give the same result:

$$\text{LILO} > \text{FILO} - \text{NFOT} = \text{Min}\{\text{FILO-NFOT}, \text{LILO}\} = \text{MIMO}$$

CASE 3a: Input Rate < Output Rate, Contiguous frame, Zero-delay Switch

This case is shown in Figure A.3a.




Figure A.3a: Input Rate < Output Rate, Contiguous frame, Zero-delay Switch

Contiguous frames are possible only if the transmission of the first bit is scheduled such that there will not be any buffer underflow until the last frame. Thus, the FIFO delay depends upon the frame time. It is non-zero and is incorrect. FILO-NFOT is similarly incorrect.

$$\text{FILO-NFOT} = \text{FIFO} > 0$$

$$\text{LILO} = \text{min}\{\text{FILO-NFOT}, \text{LILO}\} = \text{MIMO} = 0$$

Both LILO and MIMO give the correct result of zero.

CASE 3b: Input Rate < Output Rate, Contiguous frame, Nonzero-delay Switch

This case is shown in Figure A.3b.




Figure A.3b: Input Rate < Output Rate, Contiguous frame, Nonzero-delay Switch

As in Case 3a, FIFO latency depends upon the output speed. It can be made arbitrarily large by increasing the output link rate (and not changing the switch otherwise). FIFO is not a good indicator of switch latency.

FILO-NFOT is equal to FIFO latency and is also incorrect.

LILO is the only metric that can be argued to be the correct measure of latency.

$$\text{LILO} \text{ is less than FILO-NFOT. Therefore, } \text{LILO} = \text{Min}\{\text{FILO-NFOT}, \text{LILO}\} = \text{MIMO}$$

MIMO is also equal to LILO and is therefore a correct measure.

CASE 3c: Input Rate < Output Rate, Non-contiguous frame, Zero-delay Switch

This case is shown in Figure A.3c.




Figure A.3c: Input Rate < Output Rate, non-contiguous frame, zero-delay Switch

Even though the frame is non-contiguous. The cells are contiguous. To maintain frame contiguity, the departure of the first bit of each cell has to be scheduled such that there is no underflow during the first cell time. FIFO latency, therefore, depends upon the output link speed and is not a correct measure of switch latency.

FILO-NFOT is non-zero and, therefore, incorrect.

$$LILO = \min\{\text{FILO-NFOT}, LILO\} = \text{MIMO} = 0$$

Both LILO and MIMO give the correct result of zero.

CASE 3d: Input Rate < Output Rate, Non-contiguous frame, Nonzero-Delay Switch




Figure A.3d: Input Rate < Output Rate, Non-contiguous frame, Nonzero-Delay Switch

In this case, FIFO can be made small by sending the first cell fast and then introducing large time gaps in the output. FIFO is, therefore, not a valid switch latency metric in this case.

FILO - NFOT > FIFO is similarly incorrect.

LILO is the only metric that can be argued to be correct in this case.

Since $LILO < \text{FILO-NFOT}$,

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

MIMO is also a correct measure.

Once again looking at Table A.1, we find that MIMO is the only metric that applies to all input and output link rates and contiguous and non-contiguous frames.

MOTION:

Adopt the text of this contribution for inclusion in the performance benchmarking draft baseline document.

REFERENCES:

[1] R. Jain, G. Babic, and B. Nagendra, "Considerations for Frame-level Throughput and Latency Measurements of ATM switches," ATM_Forum 96-0520, April 1996.

All our papers and past ATM Forum contributions are available on-line:
<http://www.cis.ohio-state.edu/~jain/>