

.....
ATM Forum Document Number: ATM_Forum/97-0611
.....

Title: Modifications to the latency sections of Performance Testing Baseline Text
.....

Abstract: Improved text for measurement procedures, foreground and background traffic, and scalable configurations for the latency section of the baseline text is presented.
.....

Source:
Gojko Babic, Arjan Durresi, Raj Jain, Justin Dolske
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: July 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.
.....

This contribution is a resubmission of a part of our contribution 97-0426 submitted in April. In this, we explain the proposed changes to the latency section of the baseline text. As instructed in the last meeting, this contribution has three parts. In the first part, we describe the changes. The second part contains the proposed text and the third part shows the changes from the current baseline.

A. Explanations of Changes in “3.2. Frame Latency”

- a. 3.2.1. – No changes
- b. 3.2.2. – No changes
- c. 3.2.3. – Instead of defining time interval for measurement, number of frames defined. Also, this paragraph now includes expressions for the mean and the standard deviation (previously this paragraph referenced to 3.1.3, but expressions for the mean and the standard deviation have been removed from 3.1.3).
- d. 3.2.4 – Improved description of measurement procedures; in the old 3.2.4, this topic was just vaguely mentioned.
- e. 3.2.5. – Description of foreground traffic characteristics; in the old 3.2.4, this topic was just vaguely mentioned.
- f. 3.2.6. – Description of background traffic characteristics including:
 - types of VCs,
 - connection configurations (completely rewritten and improved from old Paragraph 3.1.4.), with addition of n-to-m partial cross and n-to-(n-1) multicast configurations, while n-to-1 configuration is removed.
 - service classes,
 - arrival pattern,
 - frame rate,
 - input rate.
- g. 3.2.7. – Guidelines for scaleable test configurations (new).
- h. 3.2.8. – Improved requirements for reporting of measurement results.

B. Revised Text

3.2. Frame Latency

3.2.1. 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.

3.2.2. Units

The latency should be specified in μsec .

3.2.3. Statistical Variations

For the given foreground traffic and background traffic, the required times and/or delays, needed for MIMO latency calculation, are recorded for p frames, according to the procedures described in 3.2.4. Here p is a parameter and its default (and the minimal value) is 100.

Let M_i be the MIMO latency of the i th frame. Note that MIMO latency is considered to be infinite for lost or corrupted frames. The mean and standard errors of the measurement are computed as follows:

$$\text{Mean MIMO latency} = (\sum M_i) / p$$

Standard deviation of MIMO latency = $(\sum(M_i - \text{mean MIMO latency})^2) / (p-1)$

Standard error = standard deviation of MIMO latency / $p^{1/2}$

Given the mean and the standard error, the users can compute a 100(1- α)-percent confidence interval as follows:

100(1- α)-percent confidence interval = (mean - $z \times$ standard error, mean + $z \times$ standard 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

The value of p can be chosen differently from its default value to obtain the desired confidence level.

3.2.4. Measurement Procedures

For MIMO latency measurements, it is first necessary to establish one VCC (or VPC) used only by foreground traffic, and a number of VCCs or VPCs used only by background traffic. Then, the background traffic is generated. Characteristics of background traffic are described in section 3.2.6. When flow of the background traffic has been established, the foreground traffic is generated. Characteristics of foreground traffic are specified in section 3.2.5. After the steady state flow of foreground traffic has been reached the required times and/or delays needed for MIMO latency calculation are recorded for p consecutive frames from the foreground traffic, while the flow of background traffic continue uninterrupted. The entire procedure is referred to as one measurement run.

3.2.5. Foreground traffic

MIMO latency depends upon several characteristics of foreground traffic. These include the type of foreground VCC, service class, arrival patterns, frame length, and input rate.

The foreground VCC can be a permanent or switched, virtual path or virtual channel connection established between ports on the same network module of the switch, or between ports on different network modules, or between ports on different switching fabrics.

For the UBR service class, the foreground traffic consists of equally spaced frames of fixed length. Measurements are performed on AAL payload sizes of 64 B, 1518 B, 9188 B and 64 kB. Variable length frames and other arrival patterns (e.g. self-similar) are under study. ABR service class is also under study.

Input rate of foreground traffic is expressed in the effective bits/sec, counting only bits from AAL payload excluding the overhead introduced by the ATM technology and transmission systems.

The first measurement run is performed at the lowest possible foreground input rate (for the given test equipment). For later measurement runs, the foreground load is increased up to the point when losses in the traffic occur or up to the full foreground load (FFL). FFL is equal to the lesser of the input and the output link rates used by the foreground VCC. Suggested input rates for the foreground traffic are: 0.5, 0.75, 0.875, 0.9375, 0.9687, ..., i.e. $1 - 2^{-k}$, $k = 1, 2, 3, 4, 5, \dots$, of FFL.

3.2.6. Background Traffic

Background traffic characteristics that affect frame latency are the type of background VCCs, connection configuration, service class, arrival patterns (if applicable), frame length (if applicable) and input rate.

Like the foreground VCC, background VCCs can be permanent or switched, virtual path or channel connections, established between ports on the same network module on the switch, or between ports on different network modules, or between ports on different switching fabrics. To avoid interference on the traffic generator/analyzer equipment, background VCCs are established in such way that they do not use the input link or the output link of the foreground VCC in the same direction.

For a SUT with w ports, the background traffic can use $(w-2)$ ports, not used by the foreground traffic, for both input and output. The port with the input link of the foreground traffic can be used as an output port for the background traffic. Similarly, the output port of the foreground traffic can be used as an input port for the background traffic. Overall, background traffic can use an equivalent of $n=w-1$ ports. The maximum background load (MBL) is defined as the sum of rates of all links, except the one used as the input link for the foreground traffic, in the maximum possible switch configuration.

A SUT with $w (=n+1)$ ports is measured for the following background traffic connection configurations:

- n -to- n straight, with n background VCCs, (Figure 3.2.a);
- n -to- $(n-1)$ full cross, with $n \times (n-1)$ background VCCs. (Figure 3.2.b);
- n -to- m partial cross, $1 \leq m \leq n-1$, with $n \times m$ background VCCs. (Figure 3.2.c);
- 1-to- $(n-1)$ multicast, with one (multicast) background VCC. (Figure 3.2.e);
- n -to- $(n-1)$ multicast, with n (multicast) background VCC. (Figure 3.2.d);

Use of the 1-to-(n-1)multicast and n-to-(n-1) multicast connection configurations for the background traffic is under study.

The following service classes, arrival patterns (if applicable) and frame lengths (if applicable) are used for the background traffic:

- UBR service class: Traffic consists of equally spaced frames of fixed length. Measurements are performed at AAL payload size of 64 B, 1518 B, 9188 B and 64 kB. This is a case of bursty background traffic with priority equal to or lower than that of the foreground traffic. Variable length frames and other arrival patterns (e.g. self-similar) are for further study.
- CBR service class: Traffic consists of a contiguous stream of cells at a given rate. This is a case of non-bursty background traffic with priority higher than that of the foreground traffic.
- VBR and ABR service classes are under study.

Input rate of the background traffic is expressed in the effective bits/sec, counting only bits from frames excluding the overhead introduced by the ATM technology and transmission systems.

In the cases of n-to-n straight, n-to-(n-1) full cross and n-to-m partial cross connection configurations, measurement are performed at input rates of 0, 0.5, 0.75, 0.875, 0.9375, 0.9687, ... ($1 - 2^{-k}$, $k = 0, 1, 2, 3, 4, 5, \dots$) of MBL. The required traffic load is obtained by loading each input link by the same fraction of its input rate. In this way, the input rate of background traffic can also be expressed as a fraction (percentage) of input link rates.

3.2.7. Guidelines For Scaleable Test Configurations

Scaleable test configurations for MIMO latency measurements require only one ATM test system with two generator/analyzer pairs. Figure 3.6 presents the test configuration with an ATM switch with eight ports ($w=8$). There are two links between the ATM monitor and the switch, and they are used in one direction by the background traffic and in the another direction by the foreground traffic, as indicated. The other six ($w-2$) ports of the switch are used only by the background traffic and they have external loopbacks. A loopback on a given port causes the frames transmitted over the output of the port to be received by the input of the same port.

Figure 3.6 shows a 7-to-7 straight connection configuration for the background traffic. The n-to-(n-1) full cross configuration and the n-to-m partial cross configurations can also be similarly implemented.

The test configuration shown assumes two network modules in the switch with ports P0-P3 in one network module and ports P4-P7 in the another network module. Here, the foreground VCC and background VCCs are established between ports in different network modules.

It should be noted that in the proposed test configurations, because of loopbacks, only permanent VCCs or VPCs can be established.

It should also be realized that in test configurations, if all link rates are not identical, it is not possible to generate background traffic (without losses) equal to MBL. The maximum background traffic input rate in those cases equals $(n-1) \times$ lowest link rate. Only in the case where all link rates are identical is it possible to obtain MBL level without losses in background traffic.

If the link rates are different, it is possible to obtain MBL in the n-to-n straight case, but background traffic will have losses. In this case, the foreground traffic should use the lowest rate port in the switch as the input, while the highest rate port in the switch should be used as the output. The background traffic enters the SUT through the highest rate port and passes successively through ports of decreasing speeds. At the end, the background traffic exits the switch through the lowest rate port.

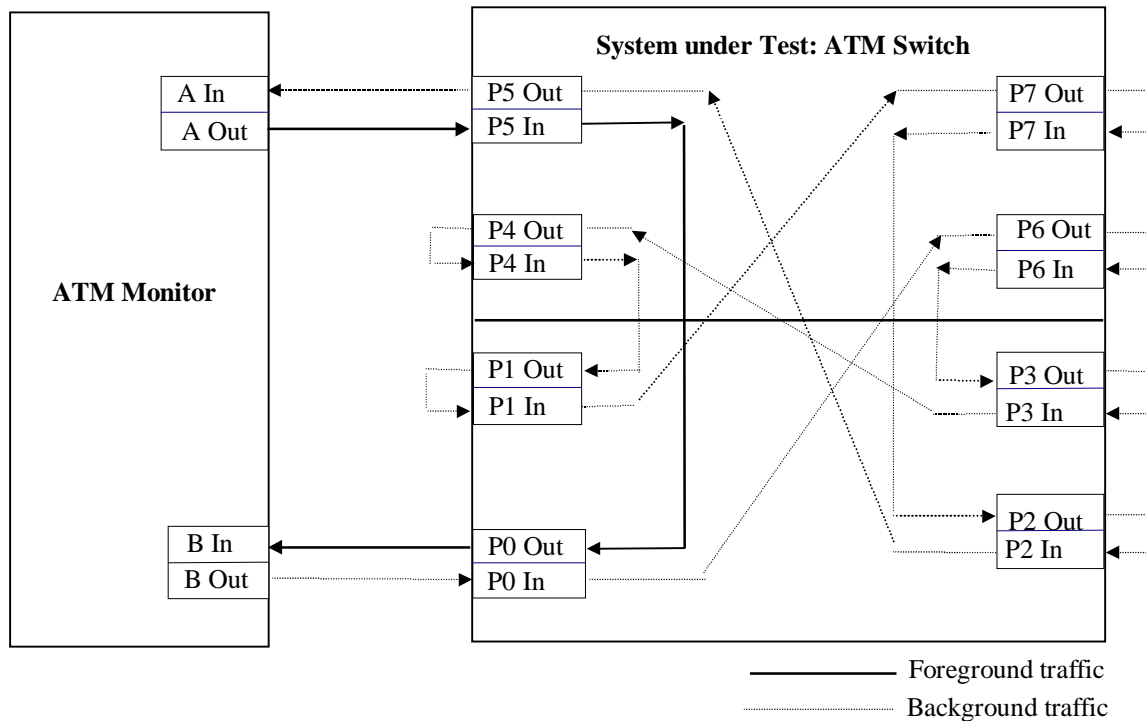


Figure 3.7: A scalable test configuration for measurements of MIMO latency using only two generator analyzer pairs with 8-port switch and 7-to-7 straight configuration for background traffic

3.2.8. Reporting results

Reported results should include detailed description of the SUT, such as the number of ports, rate of each port, number of ports per network module, number of network

modules, number of network modules per fabric, number of fabrics, the software version and any other relevant information.

Values of the mean and the standard error of MIMO latency are reported along with values of foreground and background traffic characteristics for each measurement run.

The list of foreground and background traffic characteristics and their possible values are now provided:

Foreground traffic:

- type of foreground VCC: permanent virtual path connection, switched virtual path connection, **permanent virtual channel connection**, switch virtual channel connection;
- foreground VCC established: between ports inside a network module, **between ports on different network modules**, between ports on different switching fabrics;
- service class: **UBR**, ABR;
- arrival patterns: **equally spaced frames**, self-similar, random;
- frame length: 64 B, **1518 B**, **9188 B** or 64 kB, variable;
- full foreground load (FFL);
- input rate: **the lowest rate possible for the given test equipment**, and 0.5, **0.75**, 0.875, 0.9375, 0.9687, ..., (i.e., $1 - 2^{-k}$, $k = 1, 2, 3, 4, 5, \dots$) of FFL.

Background traffic:

- type of background VCC's: permanent virtual path connections, switched virtual path connections, **permanent virtual channel connections**, switch virtual channel connections;
- foreground VCCs established: between ports inside a network module, **between ports on different network modules**, between ports on different switching fabrics, some combination of previous cases;
- connection configuration: n-to-n straight, n-to-(n-1) full cross, **n-to-m partial cross with $m = 2, 3, 4, \dots, n-1$** , 1-to-(n-1) multicast, n-to-(n-1) multicast;
- service class: **UBR**, **CBR**, ABR, VBR;
- arrival patterns (when applicable): **equally spaced frames**, self-similar, random;
- frame length (when applicable): 64 B, 1518 B, **9188 B**, 64 kB, variable;
- maximum background load (MBL);
- input rate: **0**, 0.5, 0.75, **0.875**, 0.9375, 0.9687, ... (i.e., $1 - 2^{-k}$, $k = 0, 1, 2, 3, 4, 5, \dots$) of MBL.

Values in bold indicate traffic characteristics for which measurement tests must be performed and for which MIMO latency values must be reported.

C. Differences Between the Revised and Old Text

~~3.2 FRAME LATENCY~~**3.2. Frame Latency**

~~3.2.1~~**3.2.1. 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:

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.

~~3.2.2~~**3.2.2. Units**

The latency should be specified in ~~micro-seconds~~ μsec .

~~3.2.3~~**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 3.2. The mean and standard errors computed (in a manner similar to that explained in Section 3.1 for Throughput) from these samples are reported as the test results.~~

~~3.2.4~~**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 3.2.~~

~~Before starting the throughput measurements, all required VCs will be set up (for an n-port SUT) in one of the following configurations (see Figure 3.2):~~

- ~~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.

3.2.5 ~~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 for further study. Initially all tests will be conducted without the background traffic.

3.2.6 ~~Guidelines For Using This Metric~~

To be specified.

3.2.3. Statistical Variations

For the given foreground traffic and background traffic, the required times and/or delays, needed for MIMO latency calculation, are recorded for p frames, according to the procedures described in 3.2.4. Here p is a parameter and its default (and the minimal value) is 100.

Let M_i be the MIMO latency of the i th frame. Note that MIMO latency is considered to be infinite for lost or corrupted frames. The mean and standard errors of the measurement are computed as follows:

$$\text{Mean MIMO latency} = (\sum M_i) / p$$

$$\text{Standard deviation of MIMO latency} = (\sum(M_i - \text{mean MIMO latency})^2) / (p-1)$$

$$\text{Standard error} = \text{standard deviation of MIMO latency} / p^{1/2}$$

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

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

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

<u>Confidence</u>	<u>α</u>	<u>Quantile</u>
<u>90%</u>	<u>0.1</u>	<u>1.615</u>
<u>99%</u>	<u>0.01</u>	<u>2.346</u>
<u>99.9%</u>	<u>0.001</u>	<u>3.291</u>

The value of p can be chosen differently from its default value to obtain the desired confidence level.

3.2.4. Measurement Procedures

For MIMO latency measurements, it is first necessary to establish one VCC (or VPC) used only by foreground traffic, and a number of VCCs or VPCs used only by background traffic. Then, the background traffic is generated. Characteristics of background traffic are described in section 3.2.6. When flow of the background traffic has been established, the foreground traffic is generated. Characteristics of foreground traffic are specified in section 3.2.5. After the steady state flow of foreground traffic has been reached the required times and/or delays needed for MIMO latency calculation are recorded for p consecutive frames from the foreground traffic, while the flow of background traffic continue uninterrupted. The entire procedure is referred to as one measurement run.

3.2.5. Foreground traffic

MIMO latency depends upon several characteristics of foreground traffic. These include the type of foreground VCC, service class, arrival patterns, frame length, and input rate.

The foreground VCC can be a permanent or switched, virtual path or virtual channel connection established between ports on the same network module of the switch, or between ports on different network modules, or between ports on different switching fabrics.

For the UBR service class, the foreground traffic consists of equally spaced frames of fixed length. Measurements are performed on AAL payload sizes of 64 B, 1518 B, 9188 B and 64 kB. Variable length frames and other arrival patterns (e.g. self-similar) are under study. ABR service class is also under study.

Input rate of foreground traffic is expressed in the effective bits/sec, counting only bits from AAL payload excluding the overhead introduced by the ATM technology and transmission systems.

The first measurement run is performed at the lowest possible foreground input rate (for the given test equipment). For later measurement runs, the foreground load is increased up to the point when losses in the traffic occur or up to the full foreground load (FFL). FFL is equal to the lesser of the input and the output link rates used by the foreground VCC. Suggested input rates for the foreground traffic are: 0.5, 0.75, 0.875, 0.9375, 0.9687, ..., i.e. $1 - 2^{-k}$, $k = 1, 2, 3, 4, 5, \dots$, of FFL.

3.2.6. Background Traffic

Background traffic characteristics that affect frame latency are the type of background VCCs, connection configuration, service class, arrival patterns (if applicable), frame length (if applicable) and input rate.

Like the foreground VCC, background VCCs can be permanent or switched, virtual path or channel connections, established between ports on the same network module on the switch, or between ports on different network modules, or between ports on different switching fabrics. To avoid interference on the traffic generator/analyzer equipment, background VCCs are established in such way that they do not use the input link or the output link of the foreground VCC in the same direction.

For a SUT with w ports, the background traffic can use $(w-2)$ ports, not used by the foreground traffic, for both input and output. The port with the input link of the foreground traffic can be used as an output port for the background traffic. Similarly, the output port of the foreground traffic can be used as an input port for the background traffic. Overall, background traffic can use an equivalent of $n=w-1$ ports. The maximum background load (MBL) is defined as the sum of rates of all links, except the one used as the input link for the foreground traffic, in the maximum possible switch configuration.

A SUT with $w (=n+1)$ ports is measured for the following background traffic connection configurations:

- n-to-n straight, with n background VCCs, (Figure 3.2.a);
- n-to-($n-1$) full cross, with $n \times (n-1)$ background VCCs. (Figure 3.2.b);
- n-to- m partial cross, $1 \leq m \leq n-1$, with $n \times m$ background VCCs. (Figure 3.2.c);
- 1-to-($n-1$) multicast, with one (multicast) background VCC. (Figure 3.2.e);
- n-to-($n-1$) multicast, with n (multicast) background VCC. (Figure 3.2.d);

Use of the 1-to-($n-1$) multicast and n-to-($n-1$) multicast connection configurations for the background traffic is under study.

The following service classes, arrival patterns (if applicable) and frame lengths (if applicable) are used for the background traffic:

- UBR service class: Traffic consists of equally spaced frames of fixed length. Measurements are performed at AAL payload size of 64 B, 1518 B, 9188 B and 64 kB. This is a case of bursty background traffic with priority equal to or lower than that of the foreground traffic. Variable length frames and other arrival patterns (e.g. self-similar) are for further study.
- CBR service class: Traffic consists of a contiguous stream of cells at a given rate. This is a case of non-bursty background traffic with priority higher than that of the foreground traffic.
- VBR and ABR service classes are under study.

Input rate of the background traffic is expressed in the effective bits/sec, counting only bits from frames excluding the overhead introduced by the ATM technology and transmission systems.

In the cases of n-to-n straight, n-to-(n-1) full cross and n-to-m partial cross connection configurations, measurement are performed at input rates of 0, 0.5, 0.75, 0.875, 0.9375, 0.9687, ... $(1 - 2^{-k})$, $k = 0, 1, 2, 3, 4, 5, \dots$) of MBL. The required traffic load is obtained by loading each input link by the same fraction of its input rate. In this way, the input rate of background traffic can also be expressed as a fraction (percentage) of input link rates.

3.2.7. Guidelines For Scalable Test Configurations

Scalable test configurations for MIMO latency measurements require only one ATM test system with two generator/analyzer pairs. Figure 3.6 presents the test configuration with an ATM switch with eight ports ($w=8$). There are two links between the ATM monitor and the switch, and they are used in one direction by the background traffic and in the another direction by the foreground traffic, as indicated. The other six ($w-2$) ports of the switch are used only by the background traffic and they have external loopbacks. A loopback on a given port causes the frames transmitted over the output of the port to be received by the input of the same port.

Figure 3.6 shows a 7-to-7 straight connection configuration for the background traffic. The n-to-(n-1) full cross configuration and the n-to-m partial cross configurations can also be similarly implemented.

The test configuration shown assumes two network modules in the switch with ports P0-P3 in one network module and ports P4-P7 in the another network module. Here, the foreground VCC and background VCCs are established between ports in different network modules.

It should be noted that in the proposed test configurations, because of loopbacks, only permanent VCCs or VPCs can be established.

It should also be realized that in test configurations, if all link rates are not identical, it is not possible to generate background traffic (without losses) equal to MBL. The maximum background traffic input rate in those cases equals $(n-1) \times$ lowest link rate. Only in the case where all link rates are identical is it possible to obtain MBL level without losses in background traffic.

If the link rates are different, it is possible to obtain MBL in the n-to-n straight case, but background traffic will have losses. In this case, the foreground traffic should use the lowest rate port in the switch as the input, while the highest rate port in the switch should be used as the output. The background traffic enters the SUT through the highest rate port and passes successively through ports of decreasing speeds. At the end, the background traffic exits the switch through the lowest rate port.




Figure 3.7: A scalable test configuration for measurements of MIMO latency using only two generator analyzer pairs with 8-port switch and 7-to-7 straight configuration for background traffic

3.2.8. Reporting results

Reported results should include detailed description of the SUT, such as the number of ports, rate of each port, number of ports per network module, number of network modules, number of network modules per fabric, number of fabrics, the software version and any other relevant information.

Values of the mean and the standard error of MIMO latency are reported along with values of foreground and background traffic characteristics for each measurement run.

The list of foreground and background traffic characteristics and their possible values are now provided:

Foreground traffic:

- type of foreground VCC: permanent virtual path connection, switched virtual path connection, **permanent virtual channel connection**, switch virtual channel connection;
- foreground VCC established: between ports inside a network module, **between ports on different network modules**, between ports on different switching fabrics;
- service class: **UBR**, ABR;
- arrival patterns: **equally spaced frames**, self-similar, random;
- frame length: 64 B, **1518 B**, **9188 B** or 64 kB, variable;
- full foreground load (FFL);
- input rate: **the lowest rate possible for the given test equipment**, and 0.5, **0.75**, 0.875, 0.9375, 0.9687, ..., (i.e., $1 - 2^{-k}$, $k = 1, 2, 3, 4, 5, \dots$) of FFL.

Background traffic:

- type of background VCC's: permanent virtual path connections, switched virtual path connections, **permanent virtual channel connections**, switch virtual channel connections;
- foreground VCCs established: between ports inside a network module, **between ports on different network modules**, between ports on different switching fabrics, some combination of previous cases;
- connection configuration: n-to-n straight, n-to-(n-1) full cross, **n-to-m partial cross with $m = 2, 3, 4, \dots, n-1$** , 1-to-(n-1) multicast, n-to-(n-1) multicast;
- service class: **UBR**, **CBR**, ABR, VBR;
- arrival patterns (when applicable): **equally spaced frames**, self-similar, random;
- frame length (when applicable): 64 B, 1518 B, **9188 B**, 64 kB, variable;
- maximum background load (MBL);
- input rate: **0**, 0.5, 0.75, **0.875**, 0.9375, 0.9687, ... (i.e., $1 - 2^{-k}$, $k = 0, 1, 2, 3, 4, 5, \dots$) of MBL.

Values in bold indicate traffic characteristics for which measurement tests must be performed and for which MIMO latency values must be reported.