

.....

ATM Forum Document Number: ATM_Forum/97-0614.

.....

Title: Modifications to the throughput section of Performance Testing Baseline Text

.....

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

.....

Source:

Gojko Babic, Arjan Duresi, Raj Jain, Justin Dolske
The Ohio State University

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

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 throughput 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 “3.1. Throughput”

- a. 3.1.1. – Minor editorial changes
- b. 3.1.2. – A small change clarifying units; A term effective bits/sec used, instead of just bits/sec.
- c. 3.1.3. – Expressions for the mean and the standard deviation removed. Our (and some other) studies indicate that there are no significant variations in throughput. The old paragraph 3.1.3 with minor changes is moved in the section “3.2. Frame Latency”
- d. 3.1.4 – Improved description of measurement procedures; in the old 3.1.4, this topic was only vaguely specified.
- e. 3.1.5. – This paragraph includes and expands topics from the old 3.1.4. It now provides the complete list and description of foreground traffic characteristics including:
 - types of VCs,
 - connection configurations (previously referred as traffic patterns),
 - service classes,
 - arrival pattern,
 - frame rate,
 - input rate.
 Changes in connection configurations:
 - n-to-(n-1) full cross replaces n-to-n cross
 - n-to-m partial cross introduced
 - k-to-1 replaces n-to-1
 - 1-to-(n-1) multicast replaces 1-to-n straight
 - n-to-(n-1) multicast replaces n-to-n multicast
- f. 3.1.6. – This is old 3.1.5., without its first sentence. That sentence is elaborated in the new 3.1.4.
- g. 3.1.7. –Guidelines for scaleable test configurations (new).
- h. 3.1.8. – Improved requirements for reporting of measurement results

B. Revised Text

3.1. Throughput

3.1.1. Definitions

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

- **Loss-less throughput** - It is the maximum rate at which none of the offered frames is dropped by the SUT.
- **Peak throughput** - It is the maximum rate at which the SUT operates regardless of frames dropped. The maximum rate can actually occur when the loss is not zero.
- **Full-load throughput** - It is the rate at which the 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 3.1. Level X defines the loss-less throughput, level Y defines the peak throughput and level Z defines the full-load throughput.




Figure 3.1: Peak, loss-less and full-load throughput

The loss-less throughput is the highest load at which the count of the output frames equals the count of the input frames. The peak throughput is the maximum throughput that can be achieved in spite of the losses. The full-load throughput is the throughput of the system at 100% load on input links. Note that the peak throughput may equal the loss-less 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.

3.1.2. Units

Throughput should be expressed in the effective bits/sec, counting only bits from frames excluding the overhead introduced by the ATM technology and transmission systems.

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.

3.1.3. Statistical Variations

There is no need for obtaining more than one sample for any of the three frame-level throughput metrics. Consequently, there is no need for calculation of the means and/or standard deviations of throughputs.

3.1.4. Measurement Procedures

Before starting measurements, a number of VCCs (or VPCs), henceforth referred to as “foreground VCCs”, are established through the SUT. Foreground VCCs are used to transfer only the traffic whose performance is measured. That traffic is referred to as the foreground traffic. Characteristics of foreground traffic are specified in 3.1.5.

The tests can be conducted under two conditions:

- without background traffic;
- with background traffic;

Procedure without background traffic

The procedure to measure throughput in this case includes a number of test runs. A test run starts with the traffic being sent at a given input rate over the foreground VCCs with early packet discard disabled (if this feature is available in the SUT and can be turned off). The average cell transfer delay is constantly monitored. A test run ends and the foreground traffic is stopped when the average cell transfer delay has not significantly changed (not more than 5%) during a period of at least 5 minutes.

During the test run period, the total number of frames sent to the SUT and the total number of frames received from the SUT are recorded. The throughput (output rate) is computed based on the duration of a test run and the number of received frames.

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 loss-less throughput is the highest throughput at which the count of the output frames equals the count of the input frames.

The input rate is then increased even further (with early packet discard enabled, if available). Although some frames will be lost, the throughput may increase till it reaches the peak throughput value. After this point, any further increase in the input rate will result in a decrease in the throughput.

The input rate is finally increased to 100% of the link input rates and the full-load throughput is recorded.

Procedure with background traffic

Measurements of throughput with background traffic are under study.

3.1.5. Foreground Traffic

Foreground traffic is specified by the type of foreground VCCs, connection configuration, service class, arrival patterns, frame length and input rate.

Foreground VCCs can be permanent or switched, virtual path or virtual 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.


A system with n ports can be tested for the following connection configurations:

- n -to- n straight,
- n -to- $(n-1)$ full cross,
- n -to- m partial cross, $1 \leq m \leq n-1$,
- k -to-1, $1 < k < n$,
- 1-to- $(n-1)$ multicast,
- n -to- $(n-1)$ multicast.


Different connection configurations are illustrated in Figure 3.2, where each configuration includes one ATM switch with four ports, with their input components shown on the left and their output components shown the right.

In the case of n -to- n straight, input from one port exits to another port. This represents almost no path interference among the foreground VCCs. There are n foreground VCCs. See Figure 3.2a.


In the case of n -to- $(n-1)$ full cross, input from each port is divided equally to exit on each of the other $(n-1)$ ports. This represents intense competition for the switching fabric by the foreground VCCs. There are $n \times (n-1)$ foreground VCCs. See Figure 3.2b.




a. n-to-n straight: n VCCs; $n=4$




b. n-to-(n-1) full cross: $n \times (n-1)$ VCCs; $n=4$




c. n-to-m partial cross: $n \times m$ VCCs; $n=4$, $m=2$



d. k-to-1: k VCCs; $k=3$



e. 1-to-(n-1) multicast: one (multicast) VCC



f. n-to-(n-1) multicast: n (multicast) VCCs; $n=4$

Figure 3.2: Connection configurations for foreground traffic

In the case of n -to- m partial cross, input from each port is divided equally to exit on the other m ports ($1 \leq m \leq n-1$). This represents partial competition for the switching fabrics by the foreground VCCs. There are $n \times m$ foreground VCCs as shown in Figure 3.2c. Note that n -to- n straight and n -to- $(n-1)$ full cross are special cases of n -to- m partial cross with $m=1$ and $m=n-1$, respectively.

In the case of k -to-1, input from k ($1 < k < n$) ports is destined to one output port. This stresses the output port logic. There are k foreground VCCs as shown in Figure 3.2d.

In the case of 1-to- $(n-1)$ multicast, all foreground frames input on the one designated port are multicast to all other $(n-1)$ ports. This tests single multicast performance of the switch. There is only one (multicast) foreground VCC as shown in Figure 3.2e.

Use of the 1-to- $(n-1)$ multicast connection configuration for the foreground traffic is under study.

In the case of n -to- $(n-1)$ multicast, input from each port is multicast to all other $(n-1)$ ports. This tests multiple multicast performance of the switch.. There are n (multicast) foreground VCCs. See Figure 3.2f.

Use of the n -to- $(n-1)$ multicast connection configuration for the foreground traffic is under study.

The following service classes, arrival patterns and frame lengths for foreground traffic are used for testing:

- 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. Variable length frames and other arrival patterns (e.g. self-similar) are under study.
- ABR and VBR service classes are under study.

The required input rate of foreground traffic is obtained by loading each link by the same fraction of its input rate. In this way, the input rate of foreground traffic can also be referred to as a fraction (percentage) of input link rates. The maximum foreground load (MFL) is defined as the sum of rates of all links in the maximum possible switch configuration. Input rate of the foreground traffic is expressed in the effective bits/sec, counting only bits from frames, excluding the overhead introduced by the ATM technology and transmission systems.

3.1.6. Background Traffic

Higher priority traffic (like VBR or CBR) can act as background traffic for experiments. Further details of measurements with background traffic using multiple service classes simultaneously are under study. Until then, all testing will be done without any background traffic.

3.1.7. Guidelines For Scaleable Test Configurations

It is obvious that testing larger systems, e.g., switches with larger number of ports, could require very extensive (and expensive) measurement equipment. Hence, we introduce scaleable test configurations for throughput measurements that require only one ATM monitor with one generator/analyzer pair. Figure 3.3 presents a simple test configuration for an ATM switch with eight ports in a 8-to-8 straight connection configuration. Figure 3.4 presents a test configuration with the same switch in an 8-to-2 partial cross connection configuration. The former configuration emulates 8 foreground VCCs, while the later emulates 16 foreground VCCs.

In both test configurations, there is one link between the ATM monitor and the switch. The other seven ports 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.

The test configurations in Figure 3.3 and Figure 3.4 assume two network modules in the switch, with switch ports P0-P3 in one network module and switch ports P4-P7 in the another network module. Foreground VCCs are always established from a port in one network module to a port in the another network module. These connection configurations could be more demanding on the SUT than the cases where each VCC uses ports in the same network module. An even more demanding case could be when foreground VCCs use different fabrics of a multi-fabric switch.

Approaches similar to those in Figure 3.3 and Figure 3.4 can be used for n -to- $(n-1)$ full cross and other types of n -to- m partial cross connection configurations, as well as for larger switches. Guidelines to set up scaleable test configurations for the k -to-1 connection configuration are under study.

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 the test configurations with loopbacks, if all link rates are not identical, it is not possible to generate foreground traffic equal to the MFL. The maximum foreground traffic load for a n -port switch in those cases equals $n \times$ lowest link rate. Only in the case when all link rates are identical is it possible to obtain MFL level. If all link rates are not identical, and the MFL level needs to be reached, it is necessary to have more than one analyzer/generator pair.




Figure 3.3: A scalable test configuration for throughput measurements using only one generator/analyzer pair with 8-port switch and a 8-to-8 straight connection configuration

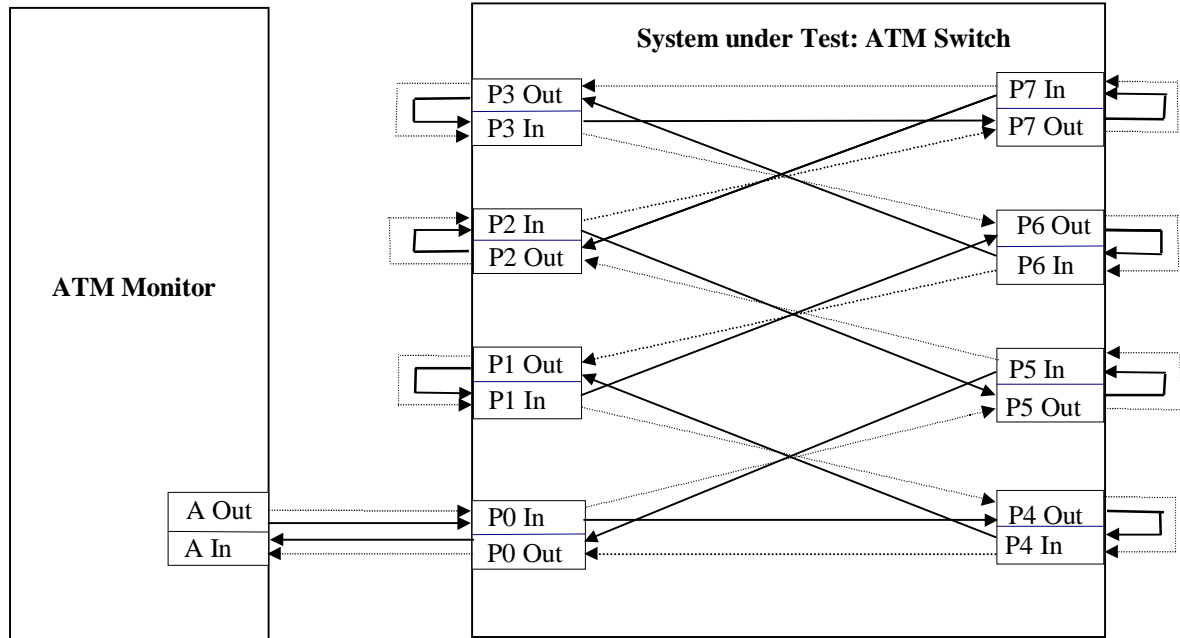


Figure 3.4: A scalable test configuration for throughput measurements using only one generator/analyzer pair with 8-port switch and a 8-to-2 partial cross connection configuration

3.1.8. Reporting results

Results should include a 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, maximum foreground load (MFL), software version, and any other relevant information.

Values for the loss-less throughput, the peak throughput with corresponding input load, and the full-load throughput with corresponding input load (if different from MFL) are reported along with foreground (and background, if any) traffic characteristics.

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

- type of foreground VCCs: 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 fabrics, some combination of previous cases;
- connection configuration: n-to-n straight, n-to-(n-1) full cross, **n-to-m partial cross** with $m = 2, 3, 4, \dots, n-1$, **k-to-1** with $k=2, 3, 4, 5, 6, \dots$, 1-to-(n-1) multicast, n-to-(n-1) multicast;
- service class: **UBR**, ABR, VBR;
- arrival patterns: **equally spaced frames**, self-similar, random;
- frame length: 64 B, **1518 B**, **9188 B** or 64 kB, variable;

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

C. Differences Between the Revised and Old Text

3.1 THROUGHPUT


3.1. Throughput

~~3.1.1~~ 3.1.1. Definitions

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

- ~~i-~~ **Lossless** ~~Loss-less~~ **throughput** - It is the maximum rate at which none of the offered frames is dropped by the SUT.
- ~~ii-~~ **Peak throughput** - It is the maximum rate at which the SUT operates 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** - It is the rate at which the 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 3.1. Level X defines the loss-less throughput, level Y defines the peak throughput and level Z defines the full-load throughput.






Figure 3.1: Peak, loss-less and full-load throughput

The loss-less throughput is the highest load at which the count of the output frames equals the count of the input frames. **Peak**The **peak** throughput is the maximum throughput that can be achieved in spite of the losses. **Full-load**The **full-load** throughput is the throughput of the system at 100% load on input links. Note that the peak throughput may equal the loss-less 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.

3.1.2 3.1.2. Units

Throughput should be expressed in the effective bits/sec, counting only bits from frames excluding the overhead introduced by the ATM technology and transmission systems. 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.

3.1.3 Statistical Variations

~~The tests should be run NRT times for TRT seconds each. Here NRT (number of repetitions for throughput tests) and TRT (time per repetition for throughput tests) are parameters. These and other such parameters and their default values are listed later in Table 3.2.~~

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

~~Mean throughput = $(\sum T_i)/n$~~

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

Standard error = 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

3.1.4 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 (see Figure 3.2):




(a) n-to-n Straight: n VCs, i to $i+1 \bmod n$



(b) n-to-n Cross: n^2 VCs



(c) n-to-1: n VCs



(d) 1-to-n Straight: 1 VC

Figure 3.2: Configurations for throughput measurements

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 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.

3. **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 it reaches 100% of the link rate. The full load throughput is then recorded.

3.1.5 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 for further study. Until then all testing will be done without any background traffic.

3.1.6 Guidelines For Using This Metric

To be specified.

3.1.3. Statistical Variations

There is no need for obtaining more than one sample for any of the three frame-level throughput metrics. Consequently, there is no need for calculation of the means and/or standard deviations of throughputs.

3.1.4. Measurement Procedures

Before starting measurements, a number of VCCs (or VPCs), henceforth referred to as “foreground VCCs”, are established through the SUT. Foreground VCCs are used to transfer only the traffic whose performance is measured. That traffic is referred as the foreground traffic. Characteristics of foreground traffic are specified in 3.1.5.

The tests can be conducted under two conditions:

- without background traffic;
- with background traffic;

Procedure without background traffic

The procedure to measure throughput in this case includes a number of test runs. A test run starts with the traffic being sent at a given input rate over the foreground VCCs with early packet discard disabled (if this feature is available in the SUT and can be turned off). The average cell transfer delay is constantly monitored. A test run ends and the foreground traffic is stopped when the average cell transfer delay has not significantly changed (not more than 5%) during a period of at least 5 minutes.

During the test run period, the total number of frames sent to the SUT and the total number of frames received from the SUT are recorded. The throughput (output rate) is computed based on the duration of a test run and the number of received frames.

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 loss-less throughput is the highest throughput at which the count of the output frames equals the count of the input frames.

The input rate is then increased even further (with early packet discard enabled, if available). Although some frames will be lost, the throughput may increase till it reaches the peak throughput value. After this point, any further increase in the input rate will result in a decrease in the throughput.

The input rate is finally increased to 100% of the link input rates and the full-load throughput is recorded.

Procedure with background traffic

Measurements of throughput with background traffic are under study.

3.1.5. Foreground Traffic

Foreground traffic is specified by the type of foreground VCCs, connection configuration, service class, arrival patterns, frame length and input rate.

Foreground VCCs can be permanent or switched, virtual path or virtual 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.


A system with n ports can be tested for the following connection configurations:

- n-to-n straight,
- n-to-(n-1) full cross,
- n-to-m partial cross, $1 \leq m \leq n-1$,
- k-to-1, $1 < k < n$,
- 1-to-(n-1) multicast,
- n-to-(n-1) multicast.


Different connection configurations are illustrated in Figure 3.2, where each configuration includes one ATM switch with four ports, with their input components shown on the left and their output components shown on the right.

In the case of n-to-n straight, input from one port exits to another port. This represents almost no path interference among the foreground VCCs. There are n foreground VCCs. See Figure 3.2a.


In the case of n-to-(n-1) full cross, input from each port is divided equally to exit on each of the other (n-1) ports. This represents intense competition for the switching fabric by the foreground VCCs. There are $n \times (n-1)$ foreground VCCs. See Figure 3.2b.




a. n-to-n straight: n VCCs; $n=4$




b. n-to-(n-1) full cross: $n \times (n-1)$ VCCs; $n=4$



c. n-to-m partial cross: $n \times m$ VCCs; $n=4$, $m=2$



d. k-to-1: k VCCs; $k=3$



e. 1-to-(n-1) multicast: one (multicast) VCC f. n-to-(n-1) multicast: n (multicast) VCCs, $n=4$

Figure 3.2: Connection configurations for foreground traffic

In the case of n -to- m partial cross, input from each port is divided equally to exit on the other m ports ($1 \leq m \leq n-1$). This represents partial competition for the switching fabrics by the foreground VCCs. There are $n \times m$ foreground VCCs as shown in Figure 3.2c. Note that n -to- n straight and n -to- $(n-1)$ full cross are special cases of n -to- m partial cross with $m=1$ and $m=n-1$, respectively.

In the case of k -to-1, input from k ($1 < k < n$) ports is destined to one output port. This stresses the output port logic. There are k foreground VCCs as shown in Figure 3.2d.

In the case of 1-to- $(n-1)$ multicast, all foreground frames input on the one designated port are multicast to all other $(n-1)$ ports. This tests single multicast performance of the switch. There is only one (multicast) foreground VCC as shown in Figure 3.2e.

Use of the 1-to- $(n-1)$ multicast connection configuration for the foreground traffic is under study.

In the case of n -to- $(n-1)$ multicast, input from each port is multicast to all other $(n-1)$ ports. This tests multiple multicast performance of the switch. There are n (multicast) foreground VCCs. See Figure 3.2f.

Use of the n -to- $(n-1)$ multicast connection configuration for the foreground traffic is under study.

The following service classes, arrival patterns and frame lengths for foreground traffic are used for testing:

- 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. Variable length frames and other arrival patterns (e.g. self-similar) are under study.
- ABR and VBR service classes are under study.

The required input rate of foreground traffic is obtained by loading each link by the same fraction of its input rate. In this way, the input rate of foreground traffic can also be referred to as a fraction (percentage) of input link rates. The maximum foreground load (MFL) is defined as the sum of rates of all links in the maximum possible switch configuration. Input rate of the foreground traffic is expressed in the effective bits/sec, counting only bits from frames, excluding the overhead introduced by the ATM technology and transmission systems.

3.1.6. Background Traffic

Higher priority traffic (like VBR or CBR) can act as background traffic for experiments. Further details of measurements with background traffic using multiple service classes simultaneously are under study. Until then, all testing will be done without any background traffic.

3.1.7. Guidelines For Scalable Test Configurations

It is obvious that testing larger systems, e.g., switches with larger number of ports, could require very extensive (and expensive) measurement equipment. Hence, we introduce scalable test configurations for throughput measurements that require only one ATM monitor with one generator/analyzer pair. Figure 3.3 presents a simple test configuration for an ATM switch with eight ports in a 8-to-8 straight connection configuration. Figure 3.4 presents a test configuration with the same switch in an 8-to-2 partial cross connection configuration. The former configuration emulates 8 foreground VCCs, while the later emulates 16 foreground VCCs.

In both test configurations, there is one link between the ATM monitor and the switch. The other seven ports 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.

The test configurations in Figure 3.3 and Figure 3.4 assume two network modules in the switch, with switch ports P0-P3 in one network module and switch ports P4-P7 in the another network module. Foreground VCCs are always established from a port in one network module to a port in the another network module. These connection configurations could be more demanding on the SUT than the cases where each VCC uses ports in the same network module. An even more demanding case could be when foreground VCCs use different fabrics of a multi-fabric switch.

Approaches similar to those in Figure 3.3 and Figure 3.4 can be used for n-to-(n-1) full cross and other types of n-to-m partial cross connection configurations, as well as for larger switches. Guidelines to set up scalable test configurations for the k-to-1 connection configuration are under study.

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 the test configurations with loopbacks, if all link rates are not identical, it is not possible to generate foreground traffic equal to the MFL. The maximum foreground traffic load for a n-port switch in those cases equals $n \times$ lowest link rate. Only in the case when all link rates are identical is it possible to obtain MFL level. If all link rates are not identical, and the MFL level needs to be reached, it is necessary to have more than one analyzer/generator pair.

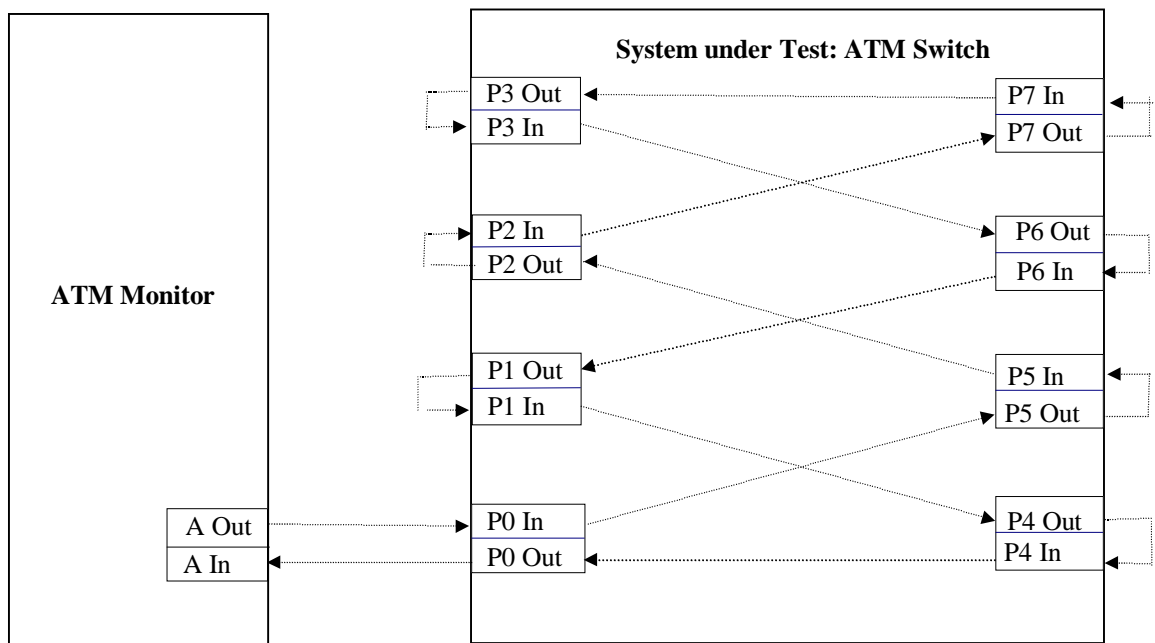


Figure 3.3: A scalable test configuration for throughput measurements using only one generator/analyzer pair with 8-port switch and a 8-to-8 straight connection configuration




Figure 3.4: A scaleable test configuration for throughput measurements using only one generator/analyzer pair with 8-port switch and a 8-to-2 partial cross connection configuration

3.1.8. Reporting results

Results should include a 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, maximum foreground load (MFL), software version, and any other relevant information.

Values for the loss-less throughput, the peak throughput with corresponding input load, and the full-load throughput with corresponding input load (if different from MFL) are reported along with foreground (and background, if any) traffic characteristics.

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

- type of foreground VCCs: 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 fabrics, some combination of previous cases;
- connection configuration: n-to-n straight, n-to-(n-1) full cross, n-to-m partial cross with $m = 2, 3, 4, \dots, n-1$, k-to-1 with $k=2, 3, 4, 5, 6, \dots$, 1-to-(n-1) multicast, n-to-(n-1) multicast;
- service class: UBR, ABR, VBR;
- arrival patterns: equally spaced frames, self-similar, random;

- frame length: 64 B, **1518 B**, **9188 B** or 64 kB, variable;

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