



The ATM Forum
Technical Committee

**~~Introduction to~~ ATM Forum
Performance ~~Benchmarking~~
Testing Specifications**

~~May~~ 07/17/96 10:21 AM ~~June, 1996~~

96-0810R1

Introduction to ATM Forum Performance Benchmarking Testing Specifications

Version 1.0

May, June 1996

(C) 1996 The ATM Forum. All Rights Reserved. No part of this publication may be reproduced in any form or by any means.

The information in this publication is believed to be accurate at its publication date. Such information is subject to change without notice and the ATM Forum is not responsible for any errors. The ATM Forum does not assume any responsibility to update or correct any information in this publication. Notwithstanding anything to the contrary, neither The ATM Forum nor the publisher make any representation or warranty, expressed or implied, concerning the completeness, accuracy, or applicability of any information contained in this publication. No liability of any kind shall be assumed by The ATM Forum or the publisher as a result of reliance upon any information contained in this publication.

The receipt or any use of this document or its contents does not in any way create by implication or otherwise:

- Any express or implied license or right to or under any ATM Forum member company's patent, copyright, trademark or trade secret rights which are or may be associated with the ideas, techniques, concepts or expressions contained herein; nor
- Any warranty or representation that any ATM Forum member companies will announce any product(s) and/or service(s) related thereto, or if such announcements are made, that such announced product(s) and/or service(s) embody any or all of the ideas, technologies, or concepts contained herein; nor
- Any form of relationship between any ATM Forum member companies and the recipient or user of this document.

Implementation or use of specific ATM recommendations and/or specifications or recommendations of the ATM Forum or any committee of the ATM Forum will be voluntary, and no company shall agree or be obliged to implement them by virtue of participation in the ATM Forum.

The ATM Forum is a non-profit international organization accelerating industry cooperation on ATM technology. The ATM Forum does not, expressly or otherwise, endorse or promote any specific products or services.

Table of Contents

1. INTRODUCTION	1
1.1 SCOPE	1
1.2 GOALS OF PERFORMANCE TESTING	2
1.3 NON-GOALS OF PERFORMANCE TESTING	3
1.5 TERMINOLOGY	4
1.6 ABBREVIATIONS	4
2. CLASSES OF APPLICATIONS	4
2.1 PERFORMANCE TESTING ABOVE THE ATM LAYER	5
2.2 PERFORMANCE TESTING AT THE ATM LAYER	5
3. PERFORMANCE METRICS	7
3.1 THROUGHPUT	7
3.1.1 Definitions	7
3.1.2 Units	8
3.1.3 Statistical Variations	8
3.1.4 Traffic Pattern	9
3.1.5 Background Traffic	10
3.2 FRAME LATENCY	10
3.2.1 Definition	10
3.2.2 Units	11
3.2.3 Statistical Variations	11
3.2.4 Traffic Pattern	11
3.2.5 Background Traffic	12
3.3 APPLICATION GOODPUT	12
3.4 REPORTING RESULTS	13
3.5 DEFAULT PARAMETER VALUES	14
APPENDIX A: MIMO LATENCY	15
1. INTRODUCTION	11
1.1 SCOPE	11
1.2 GOALS OF PERFORMANCE TESTING	21
1.3 NON-GOALS OF PERFORMANCE TESTING	32
1.4 STRUCTURE OF DOCUMENT	43
1.5 TERMINOLOGY	43
1.6 ABBREVIATIONS	43
2. OBJECTIVES OF PERFORMANCE TESTING	44
3. CLASSES OF APPLICATIONS	45
3.1 PERFORMANCE TESTING ABOVE THE ATM LAYER	56
3.2 PERFORMANCE TESTING AT THE ATM LAYER	56
4. PERFORMANCE METRICS	78
4.1 THROUGHPUT	78
4.1.1 Definitions	78
4.1.2 Units	89
4.1.3 Statistical Variations	89

4.1.4 Traffic Pattern	910
4.1.5 Background Traffic	1012
4.2 FRAME LATENCY	1012
4.2.1 Definition	1012
4.2.2 Units	1113
4.2.3 Statistical Variations	1113
4.2.4 Traffic Pattern	1113
4.2.5 Background Traffic	1214
4.3 APPLICATION GOODPUT	1214
4.4 REPORTING RESULTS	1315
4.5 DEFAULT PARAMETER VALUES	1416
APPENDIX A: MIMO LATENCY	1517

1. Introduction

Performance testing in ATM deals with the measurement of the level of quality of a SUT or a IUT under well-known conditions. The level of quality can be expressed in the form of metrics such as latency, end-to-end delay, effective throughput. Performance testing can be carried at the end-user application level (e.g., ftp, nfs) or at or above the ATM layers (e.g., cell switching, signaling, etc.). Performance testing also describes in details the procedures for testing the IUTs in the form of test suites. These procedures are intended to test the SUT or IUT and should not assume or imply any specific implementation or architecture of these systems.

This document ~~contains a proposal for a baseline document for performance testing.~~ It highlights the objectives of performance testing and suggests an approach for the development of the test suites.

~~Performance testing has been referred to as "performance benchmarking" in other contributions presented at the ATM Forum. For consistency with reference [1], it is recommended to use the terminology "performance testing" instead.~~

1.1 Scope

Asynchronous Transfer Mode, as an enabling technology for the integration of services, is gaining an increasing interest and popularity. ATM networks are being progressively deployed and in most cases a smooth migration to ATM is prescribed. This means that most of the existing applications can still operate over ATM via service emulation or service interworking along with the proper adaptation of data formats. At the same time, several new applications are being developed to take full advantage of the capabilities of the ATM technology through an Application Protocol Interface (API).

While ATM provides an elegant solution to the integration of services and allows for high levels of scalability, the performance of a given application may vary substantially with the IUT or the SUT utilized. The variation in the performance is due to the complexity of the dynamic interaction between the different layers. For example, an application running with TCP/IP stacks will yield different levels of performance depending on the interaction of the TCP window flow control mechanism and the ATM network congestion control mechanism used. Hence, the following points and recommendations are made. First, ATM adopters need guidelines on the measurement of the performance of user applications under different SUTs. Second, some functions above the ATM layer, e.g., adaptation, signaling, constitute applications (i.e. IUTs) and as such should be considered for performance testing. Also, it is essential that these layers be implemented in compliance with the ATM Forum specifications. Third, performance testing can be executed at the ATM layer in relation to the QoS provided by the different service categories. Finally, because of the extensive list of available applications, it is preferable to group applications in generic classes. Each class of applications requires different testing environment such as

metrics, test suites and traffic test patterns. It is noted that the same application, e.g., ftp, can yield different performance results depending on the underlying layers used (TCP/IP to ATM versus TCP/IP to MAC layer to ATM). Thus performance results should be compared based on the utilization of the same protocol stack.

Performance testing benchmarking is related to user perceived performance of ATM technology. For the success of ATM technology, it is important that the performance of existing and new applications be better than that of other competing networking technologies. In other words, goodness of ATM will not be measured by cell level performance but by frame-level performance and performance perceived at higher layers.

Most of the quality of Service (QoS) metrics, such as cell transfer delay (CTD), cell delay variation (CDV), cell loss ratio (CLR), and so on, may or may not be reflected directly in the performance perceived by the user. For example, while comparing two switches if one gives a CLR of 0.1% and a frame loss ratio of 0.1% while the other gives a CLR 1% but a frame loss of 0.05%, the second switch will be considered superior by many users,

ATM Forum and ITU have standardized the definitions of QoS metrics. We need to do the same for higher level performance metrics. Without a standard definition, each vendor will use their own definition of common metrics such as throughput and latency resulting in a confusion in the market place. Avoiding such a confusion will help buyers eventually leading to better sales resulting in the success of the ATM technology.

The initial work at the ATM Forum will be restricted to the native ATM layer and the adaptation layer. Any work on the performance of the higher layers will be deferred.

1.2 Goals of Performance **Benehmarking Testing**

The goal of this effort is to enhance the marketability of ATM technology and equipment. Any additional criteria that helps in achieving that goal can be added later to this list.

- a. The ATM Forum shall define metrics that will help compare various ATM equipment in terms of performance.
- b. The metrics shall be such that they are independent of switch or NIC architecture.
 - (D) The same metrics shall apply to all architectures.
- c. The metrics can be used to help predict the performance of an application or to design a network configuration to meet specific performance objectives.
- d. The ATM Forum will develop a precise methodology for measuring these metrics.

~~(i)~~ The methodology will include a set of configurations and traffic patterns that will allow vendors as well as users to conduct their own measurements.

-e. The benchmarking testing shall cover all classes of service including CBR, VBRRT, VBRNRT, ABR, and UBR.

-f. The metrics and methodology for different service classes may be different.

-g. The testing shall cover as many protocol stacks and ATM services as possible.

(i) As an example, benchmarks for verifying the performance of services such as IP, Frame Relay and SMDS over ATM may be included.

h. The ~~t~~esting shall include metrics to measure performance of network management, connection setup, and normal data transfer.

I. The following objectives are set for ATM performance testing:

(i) Definition of criteria to be used to distinguish classes of applications.

(ii) Definition of classes of applications, at or above the ATM Layer, for which performance metrics are to be provided.

(iii) Identification of the functions at or above the ATM Layer which influence the perceived performance of a given class of applications. Example of such functions include traffic shaping, quality of service, adaptation, etc. These functions need to be measured in order to assess the performance of the applications within that class.

(iv) Definition of common performance metrics for the assessment of the performance of all applications within a class. The metrics should reflect the effect of the functions identified in (iii).

(v) Provision of detailed test cases for the measurement of the defined performance metrics.

1.3 Non-Goals of Performance Testing

a. The ATM Forum is not responsible for conducting any measurements.

-b. The ATM Forum will not certify measurements.

-c. The ATM Forum will not set thresholds such that equipment performing below those thresholds are called "unsatisfactory."

-d. The ATM Forum shall not establish any requirement that dictates a cost versus performance ratio.

e. The following areas are excluded from the scope of ATM performance testing:

- (i) Applications whose performance cannot be assessed by common implementation independent metrics. In this case the performance is tightly related to the implementation. An example of such applications is network management which performance behavior depends on whether it is a centralized or a distributed implementation.
- (ii) Performance metrics which depend on the type of implementation or architecture of the SUT or the IUT.
- (iii) Test configurations and methodologies which assume or imply a specific implementation or architecture of the SUT or the IUT.
- (iv) Evaluation or assessment of results obtained by companies or other bodies.
- (v) Certification of conducted measurements or of bodies conducting the measurements.

1.4 Structure of Document

This document is structured as follows:

- ~~-Section 1 constitutes the introduction. It describes the scope, goals and non-goals of the ATM Forum. The structure of the document and the terminology are listed.~~
- ~~-Section 2 provides the objective of Performance Testing.~~
- ~~-Section 3 explains the classes of applications. Performance testing is divided into measurement above the ATM Layer which includes AAL5 layer and the higher layers and at the ATM Layer.~~
- ~~-Section 4 describes the Living List. Performance Metrics~~
- ~~-Section 5 gives the conclusion.~~

1.5 Terminology

The following definitions are used in this document. ~~Some of the definitions are taken wholly or partly from [1].~~

*Implementation Under Test (IUT): The part of the system that is to be tested.

*Metric: a variable or a function that can be measured or evaluated and which reflects quantitatively the response or the behavior of an IUT or an SUT.

*System Under Test (SUT): The system in which the IUT resides.

*Test Case: A series of test steps needed to put an IUT into a given state to observe and describe its behavior.

*Test Suite: A complete set of test cases, possibly combined into nested test groups, that is necessary to perform testing for an IUT or a protocol within an IUT.

1.6 Abbreviations

ISO	International Organization for Standardization
IUT	Implementation Under Test
NP	Network Performance
NPC	Network Parameter Control
PDU	Protocol Data Unit
PVC	Permanent Virtual Circuit
QoS	Quality of Service
SUT	System Under Test
SWG	Sub Working Group
SVC	Switched Virtual Circuit

2. Objectives of Performance Testing

Asynchronous Transfer Mode, as an enabling technology for the integration of services, is gaining an increasing interest and popularity. ATM networks are being progressively deployed and in most cases a smooth migration to ATM is prescribed. This means that most of the existing applications can still operate over ATM via service emulation or service interworking along with the proper adaptation of data formats. At the same time, several new applications are being developed to take full advantage of the capabilities of the ATM technology through an Application Protocol Interface (API).

While ATM provides an elegant solution to the integration of services and allows for high levels of scalability, the performance of a given application may vary substantially with the IUT or the SUT utilized. The variation in the performance is due to the complexity of the dynamic interaction between the different layers. For example, an application running with TCP/IP stacks will yield different levels of performance depending on the interaction of the TCP window flow control mechanism and the ATM network congestion control mechanism used. Hence, the following points and recommendations are made. First, ATM adopters need guidelines on the

~~measurement of the performance of user applications under different SUTs. Second, some functions above the ATM layer, e.g., adaptation, signaling, constitute applications (i.e. IUTs) and as such should be considered for performance testing. Also, it is essential that these layers be implemented in compliance with the ATM Forum specifications. Third, performance testing can be executed at the ATM layer in relation to the QoS provided by the different service categories. Finally, because of the extensive list of available applications it is preferable to group applications in generic classes. Each class of applications requires different testing environment such as metrics, test suites and traffic test patterns. It is noted that the same application, e.g., ftp, can yield different performance results depending on the underlying layers used (TCP/IP to ATM versus TCP/IP to MAC layer to ATM). Thus performance results should be compared based on the utilization of the same protocol stack.~~

~~The following objectives are set for ATM performance testing:~~

- ~~(i) Definition of criteria to be used to distinguish classes of applications. Example of criteria are given in Section 4.~~
- ~~(ii) Definition of classes of applications, at or above the ATM Layer, for which performance metrics are to be provided.~~
- ~~(iii) Identification of the functions at or above the ATM Layer which influence the perceived performance of a given class of applications. Example of such functions include traffic shaping, quality of service, adaptation, etc. These functions need to be measured in order to assess the performance of the applications within that class.~~
- ~~(iv) Definition of common performance metrics for the assessment of the performance of all applications within a class. The metrics should reflect the effect of the functions identified in (iii).~~
- ~~(v) Provision of detailed test cases for the measurement of the defined performance metrics.~~

~~The following areas are excluded from the scope of ATM performance testing:~~

- ~~(i) Applications which whose performance cannot be assessed by common implementation independent metrics. In this case the performance is tightly related to the —implementation. An example of such applications is network management which performance behavior depends on whether it is a centralized or a distributed implementation.~~
- ~~(ii) Performance metrics which depend on the type of implementation or architecture of the SUT or the IUT.~~
- ~~(iii) Test configurations and methodologies which assume or imply a specific implementation or architecture of the SUT or the IUT.~~
- ~~(iv) Evaluation or assessment of results obtained by companies or other bodies.~~

~~(v) Certification of conducted measurements or of bodies conducting the measurements.~~

2.3. Classes of Applications

Developing a test suite for each existing and new application can prove to be a difficult task. Instead, applications should be grouped into categories or classes. Applications in a given class have similar performance requirements and can be characterized by common performance metrics. This way, the defined performance metrics and test suites will be valid for a range of applications. Classes of applications can be defined based on one or a combination of criteria. The following criteria can be used in the definition of the classes:

- (i) Time or delay requirements: real-time versus non real-time applications.
- (ii) Distance requirements: LAN versus WAN applications.
- (iii) Media type: voice, video, data, or multimedia application.
- (iv) Quality level: for example desktop video versus broadcast quality video.
- (v) ATM service category used: some applications have stringent performance requirements and can only run over a given service category. Others can run on several service categories. An ATM service category relates application aspects to network functionalities. ~~Examples of user applications for the defined service categories in UNI 4.0 are given in [2].~~
- (vi) Others to be determined.

2.3.1 Performance Testing Above the ATM Layer

Performance metrics can be measured at the user application layer, and sometimes at the transport layer and the network layer, and can give an accurate assessment of the perceived performance. Since it is difficult to cover all the existing applications and all the possible combinations of applications and underlying protocol stacks, it is desirable to classify the applications into classes ~~as described in section 4.~~ Performance metrics and performance test suites can be provided for each class of applications.

The perceived performance of a user application running over an ATM network is dependent on many parameters. It can vary substantially by changing an underlying protocol stack, the ATM service category it uses, the congestion control mechanism used in the ATM network, etc. Furthermore, there is no direct and unique relationship between the ATM Layer Quality of Service (QoS) parameters and the perceived application performance. For example, in an ATM network implementing a packet level discard congestion mechanism, applications using TCP as

the transport protocol may see their effective throughput improved while the measured cell loss ratio may be relatively high. In practice, it is difficult to carry measurements in all the layers that span the region between the ATM Layer and the user application layer given the inaccessibility of testing points. More effort needs to be invested to define the performance at these layers. These layers include adaptation, signaling, etc

2.3.2 Performance Testing at the ATM Layer

The notion of application at the ATM Layer is related to the service categories provided by the ATM service architecture. The Traffic Management Specification, version 4.0, specifies five service categories [2]: CBR, rt-VBR, nrt-VBR, UBR, and ABR. Each service category defines a relation of the traffic characteristics and the Quality of Service (QoS) requirements to network behavior. A set of ATM QoS parameters is defined in [2]. There is an assessment criteria of the QoS associated with each of these parameters. These are summarized in Table 26.1.

QoS PERFORMANCE PARAMETER	QoS ASSESSMENT CRITERIA
Cell Error Ratio	Accuracy
Severely-Errored Cell Block Ratio	Accuracy
Cell Misinsertion Rate	Accuracy
Cell Loss Rate	Dependability
Cell Transfer Delay	Speed
Cell Delay Variation	Speed

Table 26.1: ATM Transfer Performance Parameters.

A few methods for the measurement of the QoS parameters are defined in [2]. However, detailed test cases and procedures, –as well as test configurations are needed for both in-service and out-of-service measurement of QoS parameters. An example of test configuration for the out-of-service measurement of QoS parameters is given in [1].

Performance testing at the ATM Layer covers the following categories:

- (i) In-service and out-of-service measurement of the QoS performance parameters for all five service categories (or application classes in the context of performance testing): CBR, rt-VBR, nrt-VBR, UBR, and ABR. The test configurations assume a non-overloaded SUT.
- (ii) Performance of the SUT under overload conditions. In this case, the efficiency of the congestion avoidance and congestion control mechanisms of the SUT are tested.

In order to provide common performance metrics that are applicable to a wide range of SUT's and that can be uniquely interpreted, the following requirements must be satisfied:

- (i) Reference load models for the five service categories CBR, rt-VBR, nrt-VBR, UBR, and ABR, are required. Reference load models are to be defined by the Traffic Management Working Group.
- (ii) Test cases and configurations must not assume or imply any specific implementation or architecture of the SUT.

3.4. 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.

3.4.1 THROUGHPUT

3.4.1.1 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.

3.4.1.2 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.

3.4.1.3 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 ith run, The mean and standard errors of the measurement should be computed as follows:

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

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

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

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

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

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

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

3.4.1.4 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.

3.4. 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.

3.4.1.5 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 testing will be done without any background traffic.

3.4.2 FRAME LATENCY

3.4.2.1 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 = Min{First-bit in to last-bit out latency - nominal frame output time, 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.

3.4.2.2 Units

The latency should be specified in micro-seconds.

3.4.2.3 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.

3.4.2.4 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 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.

3.4.2.5 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.

3.4.3 Application Goodput

Application-goodput captures the notion of what an application sees as useful data transmission in the long term. Application-goodput is the ratio of packets(frames) received to packets(frames) transmitted over a measurement interval.

The application-goodput (AG) is defined as:

$$\text{AG} = \frac{\text{Frames Received in Measurement Interval}}{\text{Frames Transmitted in Measurement Interval}}$$

where Measurement Interval is defined as the time interval from when a frame was successfully received to when the frame sequence number has advanced by n.

Note that traditionally goodput is measured in bits per sec. However, we are interested in a non-dimensional metric and are primarily interested in characterizing the useful work derived from the expended effort rather than the actual rate of transmission. While the application-goodput is intended to be used in a single-hop mode, it does have meaningful end-to-end semantics over multiple hops.

Notes:

1. This metric is useful when measured at the peak load which is characterized by varying the number of transmitted frames must be varied over a useful range from 2000 frames per second (fps) through 10000 fps at a nominal frame size of 64 bytes. Frame sizes are also varied through 64 bytes, 1518 bytes, and 9188 bytes to represent small, medium, and large frames respectively. Note that the frame sizes specified do not account for the overhead of accomodating the desired frame transmission rates over the ATM medium.
2. Choose the measurement interval to be large enough to accommodate the transmission of the largest packet (frame) over the connection and small enough to track short-term excursions of the average goodput.
3. It is important to not include network management frames and/or keep alive frames in the count of received frames.
4. There should be no changes of frame handling buffers during the measurement.
6. The results are to be reported as a table for the three different frame sizes.

3.44.4 REPORTING RESULTS

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

Table 3.1: Tabular format for reporting performance testing benchmarking results

	Throughput							
Traffic	Lossless		Peak		Full-load		Latency	
Pattern	Mean	Std Err	Mean	Std Err	Mean	Std Err	Mean	Std Err

<u>n-to-n</u> Straight								
<u>n-to-n</u> Cross								
<u>n-to-1</u>								
<u>1-to-n</u>								

3.4.5 DEFAULT PARAMETER VALUES

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

Table 2: List of Parameters and their default values

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

5. Status and Conclusion

This document has introduced performance testing, discussed the role of ATM Forum's goals and non-goals. Also the testing has been confined to the ATM layer and the AAL-5 layer for the present, though higher layer testing will be addressed later. Performance metrics for throughput and latency has been provided in detail. The message-in message-out (MIMO) latency is a general definition of latency that applies to a switch or a group of switches and a detailed analysis is provided in the appendix. Application-goodput captures the notion of what an application sees as useful data transmission in the long term and the definition is provided.

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 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:

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

$$\text{LIFO} = \text{FIFO} - \text{Nominal frame output time}$$

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

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

	<u>frame, zero-delay switch</u>				
<u>1d</u>	<u>Input rate = output rate, non-contiguous frame, nonzero-delay switch</u>	<u>×</u>	<u>√</u>	<u>√</u>	<u>√</u>
<u>2a</u>	<u>Input rate > output rate, contiguous frame, zero-delay switch</u>	<u>√</u>	<u>×</u>	<u>√</u>	<u>√</u>
<u>2b</u>	<u>Input rate > output rate, contiguous frame, nonzero-delay switch</u>	<u>√</u>	<u>×</u>	<u>√</u>	<u>√</u>
<u>2c</u>	<u>Input rate > output rate, non-contiguous frame, zero-delay switch</u>	<u>This case is not possible</u>			
<u>2d</u>	<u>Input rate > output rate, non-contiguous frame, nonzero-delay switch</u>	<u>×</u>	<u>×</u>	<u>√</u>	<u>√</u>
<u>3a</u>	<u>Input rate < output rate, contiguous frame, zero-delay switch</u>	<u>×</u>	<u>√</u>	<u>×</u>	<u>√</u>
<u>3b</u>	<u>Input rate < output rate, contiguous frame, nonzero-delay switch</u>	<u>×</u>	<u>√</u>	<u>×</u>	<u>√</u>
<u>3c</u>	<u>Input rate < output rate, non-contiguous frame, zero-delay switch</u>	<u>×</u>	<u>√</u>	<u>×</u>	<u>√</u>
<u>3d</u>	<u>Input rate < output rate, non-contiguous frame, nonzero-delay switch</u>	<u>×</u>	<u>√</u>	<u>×</u>	<u>√</u>

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:

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

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:

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

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.

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

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

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.

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

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

MIMO is also equal to LIFO 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.

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

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 < FILO-NFOT,

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

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.