
ATM Forum Document Number: BTD-TEST-TM-PERF.00.00 (96-0810R3)

Title: ATM Forum Performance Testing Specification - Baseline Text

Abstract: This baseline document includes all text related to performance testing that has been agreed so far by the ATM Forum Testing Working Group.

Source:

Raj Jain
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: December 1996

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.



The ATM Forum
Technical Committee

**ATM Forum Performance
Testing Specification**

11/18/96 8:41 PM

BTD-TEST-TM-PERF.00.00 (96-0810R3)

ATM Forum Performance Testing Specifications

Version 1.0, December 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 APPLICATION	4
2.1 PERFORMANCE TESTING ABOVE THE ATM LAYER.....	5
2.2 PERFORMANCE TESTING AT THE ATM LAYER.....	5
3. PERFORMANCE METRICS	6
3.1 THROUGHPUT.....	7
3.1.1 <i>Definitions</i>	7
3.1.2 <i>Units</i>	8
3.1.3 <i>Statistical Variations</i>	8
3.1.4 <i>Traffic Pattern</i>	9
3.1.5 <i>Background Traffic</i>	10
3.1.6 <i>Guidelines For Using This Metric</i>	10
3.2 FRAME LATENCY	10
3.2.1 <i>Definition</i>	10
3.2.2 <i>Units</i>	11
3.2.3 <i>Statistical Variations</i>	11
3.2.4 <i>Traffic Pattern</i>	11
3.2.5 <i>Background Traffic</i>	12
3.2.6 <i>Guidelines For Using This Metric</i>	12
3.3. THROUGHPUT FAIRNESS.....	12
3.3.1 <i>Definition</i>	12
3.3.2 <i>Load Level and Traffic Pattern</i>	13
3.3.3 <i>Statistical Variation</i>	13
3.3.4 <i>Background Traffic</i>	14
3.3.5 <i>Reporting Results</i>	14
3.3.6 <i>Guidelines For Using This Metric</i>	14
3.4. FRAME LOSS RATIO	14
3.4.1 <i>Definition</i>	14
3.4.2 <i>Unit</i>	15
3.4.3 <i>Traffic Patterns</i>	15
3.4.4 <i>Statistical Variation</i>	15
3.4.5 <i>Reporting Results</i>	15
3.4.6 <i>Guidelines For Using This Metric</i>	15
3.5. MAXIMUM FRAME BURST SIZE (MFBS).....	15
3.5.1 <i>Definition</i>	15
3.5.2 <i>Units</i>	16
3.5.3 <i>Statistical Variations</i>	16
3.5.4 <i>Traffic Patterns</i>	16
3.5.5 <i>Guidelines For Using This Metric</i>	16
3.6. CALL ESTABLISHMENT LATENCY	16
3.6.1 <i>Definition</i>	16
3.6.2 <i>Units</i>	17
3.6.3 <i>Configurations</i>	18

BTD-TEST-TM-PERF.00.00 (96-0810R3) ATM Forum Performance Testing Specification

3.6.4 *Statistical Variations*..... 18
3.6.5 *Guidelines For Using This Metric* 18
3.7 APPLICATION GOODPUT 18
 3.7.1 *Guidelines For Using This Metric* 19
3.8 REPORTING RESULTS 19
3.9 DEFAULT PARAMETER VALUES 20
APPENDIX A: MIMO LATENCY..... **21**

1. Introduction

Performance testing in ATM deals with the measurement of the level of quality of a system under test (SUT) or an interface under test (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), 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 do not assume or imply any specific implementation or architecture of these systems.

This document highlights the objectives of performance testing and suggests an approach for the development of the test suites.

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 over different systems. 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 is related to user perceived performance of ATM technology. In other words, goodness of ATM will be measured not only by cell level performance but also 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 ATM layer 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 is being deferred for further study.

1.2 Goals of Performance 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.
 - (i) 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 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, measurements for verifying the performance of services such as IP, Frame Relay and SMDS over ATM may be included.

h. The testing 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 will 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, whose 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.

1.5 Terminology

The following definitions are used in this document.

*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
SVC	Switched Virtual Circuit
WG	Working Group

2. Classes of Application

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 application 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.
- (vi) Others to be determined.

2.1 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. 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 out 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

2.2 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: 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. There is an assessment criteria of the QoS associated with each of these parameters. These are summarized below.

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

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. 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.

3.1 THROUGHPUT

3.1.1 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 is 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 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, 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.

3.1.2 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.

3.1.3 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:

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

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

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

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


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

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


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

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.

3.2 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 micro-seconds.

3.2.3 Statistical Variations

NML samples of the latency are obtained by sending NML marked frames at $\text{TTL}/(\text{NML} + 1)$ 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.

3.3. THROUGHPUT FAIRNESS

3.3.1 Definition

Given n contenders for the resources, throughput fairness indicates how far the actual individual allocations are from the ideal allocations. In the most general case of a network, ideal allocation is defined by max-min allocation to various contending virtual circuits. For the simplest case of n VCs sharing a link with a total throughput T , the throughput of each VC should be T/n .

If the actual measured throughputs of n VCs sharing a system (a single switch or a network of switches) are found to be $\{T_1, T_2, \dots, T_n\}$, where the optimal max-min throughputs¹ should be

¹ Other policies could be used but must be specified.

$\{\bar{T}_1, \bar{T}_2, \dots, \bar{T}_n\}$, then the fairness of the system under test is quantified by the "fairness index" computed as follows:

$$\text{Fairness Index} = \frac{(\sum x_i)^2}{n \sum x_i^2}$$

Where, $x_i = T_i / \bar{T}_i$ is the relative allocation to i th VC.

This Fairness Index has the following desirable properties:

1. It is dimensionless. The units used to measure the throughput (bits/sec, cells/sec, frames/sec) do not affect its value.
2. It is a normalized measure that ranges between zero and one. The maximum fairness is 100% and the minimum 0%. This makes it intuitive to interpret and present.
3. If all x_i 's are equal, the allocation is fair and the fairness index is one.
4. If $n-k$ of n x_i 's are zero, while the remaining k x_i 's are equal and non-zero, the fairness index is k/n . Thus, a system which allocates all its capacity to 80% of VCs has a fairness index of 0.8 and so on.

3.3.2 Load Level and Traffic Pattern

Throughput fairness is quantified via the fairness index for each of the throughput experiments in which there are either multiple VCs or multiple input or output ports. Thus, it applies to all three throughput measures (lossless, peak, and full-load) and all four traffic patterns (n-to-n straight, n-to-n cross, n-to-1, and 1-to-n) described in Section 3.1.4.

Note that in the case of n-to-n cross, there are n^2 VCs and, therefore, n^2 should be substituted in place of n in the fairness index.

In the case of a 1-to-n pattern, there is only one VC and all input is expected to be multicast to n output ports. The fairness will measure the equality of throughput to the output ports.

No additional experiments are required for throughput fairness. The detailed results obtained for the throughput tests are analyzed to compute the fairness.

3.3.3 Statistical Variation

The throughput tests are run NRT times for TRT seconds each. Recall that NRT and TRT are parameters. The fairness is computed for each individual run. Let F_i be the fairness for the i th run, then the mean fairness is computed as follows:

$$\text{Mean Fairness} = \text{sum}(F_i) / \text{NRT}$$

3.3.4 Background Traffic

The throughput tests are conducted with and without background traffic. Higher priority VBR traffic can act as background traffic. Further details for measurements with background traffic (multiple service classes simultaneously) are for further study. Until then all performance testing will be done without any background traffic.

3.3.5 Reporting Results

The fairness index values are reported for each of the throughput experiments in the tabular format specified in Table 3.1.

Note that fairness index is not limited to throughput. It can be applied to other metrics, such as latency. However, extreme unfairness in latency is expected show up as unfairness in throughput and vice versa. Therefore, it is not required to quantify fairness of latency.

3.3.6 Guidelines For Using This Metric

To be specified.

3.4. FRAME LOSS RATIO

3.4.1 Definition

Frame loss ratio is defined as the fraction of frames that are not forwarded by a system under test (SUT) due to lack of resources. Partially delivered frames are considered lost.

$$\text{Frame loss ratio} = (\text{Input frame count} - \text{output frame count}) / (\text{input frame count})$$

There are two frame loss ratio metrics that are of interest to a user.

- i. **Peak-throughput frame loss ratio** - It is the frame loss ratio at a frame load for the peak throughput.
- ii. **Full-load frame loss ratio** - It is the frame loss ratio at a frame load for the full-load throughput.

These metrics are related to the throughput:

$$\text{Frame Loss Ratio} = (\text{Input Rate} - \text{Throughput}) / \text{Input Rate}$$

Thus, no additional experiments are required for frame loss ratios. These can be derived from tests performed for throughput measurements provided the input rates are recorded.

3.4.2 Unit

The frame loss ratio is expressed as a fraction of input frames.

3.4.3 Traffic Patterns

FLRs are measured for each of the four traffic patterns (n-to-n straight, n-to-n cross, n-to-1, and 1-to-n) specified for throughput measurements in Section 3.1.4. All frames are of the same size.

3.4.4 Statistical Variation

The throughput experiments are repeated NRT times for TRT seconds each. Here, NRT and TRT are parameters. If FLR_i is the frame loss ratio for the i th run:

$$\text{Frame Loss Ratio } FLR_i = (\text{Input Rate}_i - \text{Throughput}_i) / \text{Input Rate}_i$$

Since frame loss ratio is a "ratio," its average cannot be computed via straight summation. The average frame loss ratio for NRT runs is computed as follows:

$$\text{Average Frame Loss Ratio } FLR = [\sum \text{Input Rate}_i - \sum \text{Throughput}_i] / \sum \text{Input Rate}_i$$

The average is reported as the FLR for the experiment.

3.4.5 Reporting Results

FLR values are reported for peak throughput and full-load throughput experiments in the tabular format specified in Table 3.1.

3.4.6 Guidelines For Using This Metric

To be specified.

3.5. MAXIMUM FRAME BURST SIZE (MFBS)

3.5.1 Definition

Maximum Frame Burst Size (MFBS) is the maximum number of frames that source end systems can send at the peak rate through a system under test without incurring any loss.

MFBS measures the data buffering capability of the SUT and its ability to handle back-to-back frames.

Many applications and transport layer protocol drivers often present a burst of frames to AAL for transmission. For such applications, Maximum Frame Burst Size provides a useful indication.

This metric is particularly relevant to UBR service category since the UBR sources are always allowed to send a burst at peak rate. ABR sources may be throttled down to a lower rate if a switch runs out of buffer.

3.5.2 Units

MFBS should be expressed in octets of AAL payload field. This is preferred over number of frames or cells. The former requires specifying the frame size and the latter is not very meaningful for a frame-level metric. Also, number of cells has to be converted to octets for use by AAL users.

It may be useful to indicate the frame size for which MFBS has been measured. If MFBS is found to be highly variable with frame size, a number of common AAL payload field sizes such as 64 octets, 536 octets, 1518 octets, and 9188 octets may be used (exact sizes are for further study).

3.5.3 Statistical Variations

The number of frames sent in the burst is increased successively until a loss is observed on any VC. The maximum number of frames that can be sent without loss are reported as MFBS. The tests should be repeated NRT times. The average of NRT repetitions is reported as the MFBS for the system under test.

3.5.4 Traffic Patterns

The MFBS is measured for n-to-1 traffic pattern specified in Section 3.1.4. Optionally, it can be measured for other traffic patterns also. The value obtained for n-to-1 pattern is expected to be smaller than that for other patterns.

3.5.5 Guidelines For Using This Metric

To be specified.

3.6. CALL ESTABLISHMENT LATENCY

3.6.1 Definition

For short duration VCs, call establishment latency is an important part of the user perceived performance. Informally, the time between submission of a call setup request to a network and the receipt of the connect message from the network is defined as the call establishment latency. The time lost at the destination while the destination was deciding whether to accept the call is not under network control and is, therefore, not included in call setup latency (See Figure 3.1).

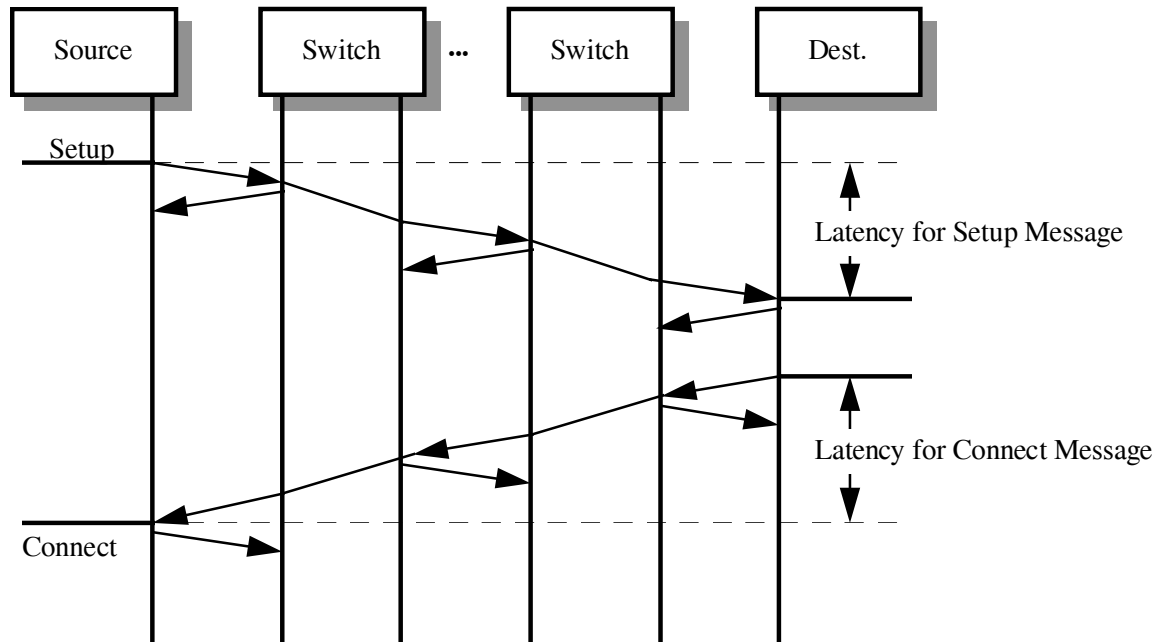


Figure 3.1: Call establishment

Thus, the sum of the latency experienced by the setup message and the resulting connect message is the call setup latency.

The main problem in measuring these latencies is that both these messages span multiple cells with intervening idle/unassigned cells. Unlike X.25, frame relay, and ISDN networks, the messages in ATM networks are not contiguous. Therefore, the MIMO latency metric defined in Section 3.2 is used². Thus,

$$\text{Call Establishment Latency} = \text{MIMO Latency for SETUP message} + \text{MIMO latency for the corresponding Connect message}$$

Recall that the MIMO latency for a frame is defined as the minimum of last-bit-in-to-last-bit-out (LILO) and the difference of first-bit-in-to-last-bit-out (FILO) and normal frame output time (NFOT).

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

3.6.2 Units

Call establishment latency is measured in units of time.

² Applies only if cells of setup and connect messages are contiguous at the input port.

3.6.3 Configurations

The call establishment latency as defined above applies to any network of switches. In practice, it has been found that the latency depends upon the number of switches and the number of PNNI group hierarchies traversed by the call. It is expected that measurements will be conducted on multiple switches connected in a variety of ways. In all cases, the number of switches and number of PNNI group hierarchies traversed should be indicated.

The simplest configuration is that of a single switch connecting both the source and the destination end systems. Further configurations are for further study.

3.6.4 Statistical Variations

The latency measurement is repeated NRT times. Each time a different node pair is selected randomly as the source and destination end system. The average and standard error of NRT such measurements is reported. For a single n-port switch it is expected that all n ports are equally probable candidates to be source and destination end-system.

3.6.5 Guidelines For Using This Metric

To be specified.

3.7 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:

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

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 64bytes. 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 frametransmission 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 not to 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.
5. The results are to be reported as a table for the three different frame sizes.

3.7.1 Guidelines For Using This Metric

To be specified.

3.8 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 results

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

3.9 DEFAULT PARAMETER VALUES

The default values of the parameters used in performance testing are listed in Table 3.2.

Table 3.2: List of Parameters and their default values

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

APPENDIX A: MIMO LATENCY

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

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




Figure A.1: Latency for single-bit frames

For a multi-bit frame, there are several possible 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:

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

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

It is clear that FIFO (or LILO) is a preferred metric 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 applies 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 metric in all cases while other metrics apply to some cases but give wrong answers in others.

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

Table A.1: Applicability of various latency definitions

No.	Case	FIFO	LILO	FILO-NFOT	MIMO
1a	Input rate = output rate, contiguous frame, zero-delay switch	√	√	√	√
1b	Input rate = output rate, contiguous frame,	√	√	√	√

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

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

√ => The metric gives a valid result

× => The metric gives an invalid result

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

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




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

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

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

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

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




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

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

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

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

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

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

All four metrics again give identical and meaningful result.

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

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

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

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




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

FILO, LILO, and MIMO are related as follows:

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

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

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

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




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

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

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

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

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

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




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

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

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

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

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

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

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

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




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

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

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

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

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

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

This case is shown in Figure A.3a.




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

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

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

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

Both LILO and MIMO give the correct result of zero.

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

This case is shown in Figure A.3b.




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

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

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

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

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

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

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

This case is shown in Figure A.3c.




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

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

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

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

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,

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

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.