ABSTRACT

The Test Working Group of the ATM Forum is developing a specification for performance testing of ATM switches and networks. The emphasis is on the user-perceived frame-level performance. This article explains what is different about this new effort and describes its status.

Performance Testing Effort at the ATM Forum: An Overview

Raj Jain and Gojko Babic, The Ohio State University

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

*A*synchronous transfer mode (ATM) technology is now being deployed in operational networks. Most of the specifications required for operation have been developed. This includes signaling (user-network interface: UNI 4.0), routing (private network-network interface: PNNI 1.0), traffic management (TM 4.0), and numerous physical-layer, network management, and testing specifications. As the technology moves from laboratories to the field, users have a need to benchmark and compare various ATM switches and other devices. *The Devil's DP Dictionary* defines performance benchmarking as follows:

Benchmark *v. trans.* To subject (a system) to a series of tests in order to obtain prearranged results not available on competitive systems [1].

This definition is not very far from the truth. In the absence of a performance testing standard, each vendor is free to use whatever metric it chooses and to measure that in an arbitrary manner. This can lead to confusion among buyers and users of the technology and can hurt the technology. Therefore, it is important to develop a set of standard performance metrics and precisely define their measurement and reporting procedures.

In October 1995, we therefore made a proposal for starting the performance testing specification in the ATM Forum [2]. The forum members have been enthusiastic, and considerable progress has been made since then. At every bimonthly meeting of the Forum, joint meetings of the Traffic Management and Test Working Groups have been held to discuss performance testing. A baseline document [3] is being prepared. Although much work remains to be done, the goal of this article is to provide advance insight into this effort for ATM switching and monitoring product designers and other readers.[1]

CELL-LEVEL VS. FRAME-LEVEL METRICS

*O*ne of the key distinguishing features of this new effort is its emphasis on frame-level metrics. In the past, the performance of ATM equipment as well as the quality of service were defined in terms of cell-level metrics. Cell loss ratio (CLR), cell delay variation (CDV), and cell transfer delay (CTD) are examples of cell-level metrics. Unfortunately, cell-level metrics do not reflect the performance as experienced (or desired) by end users. Most users have frames to send and, for the same CLR, user-perceived performance can be very different depending on whether the cells dropped belong to a few or many frames. The user is more interested in frame loss ratio. Here, the term "frame" refers to an ATM adaptation layer (AAL) or higher-layer protocol data unit (PDU).

A similar argument can be made against other cell-level metrics. For example, a video user sending 30 frames/s would like to receive complete frames every 33 ms. It does not matter whether the cells belonging to a frame arrive together or regularly spaced. Thus, it is the frame delay variation that matters, not the CDV.

Frame-level metrics are also helpful in allowing ATM technology to be compared to non-ATM technology. For example, given a traffic pattern, a user could compare the performance of several network design alternatives; some may be ATM-based, others non-ATM-based.

Based on these arguments, the ATM Forum Test Working Group members decided to start the work on defining frame-level metrics.

GOALS OF THE ATM FORUM WORK

*T*he objective of the ATM Forum work on performance testing is to enhance the marketability of ATM technology and equipment. The Forum will define metrics to help compare various ATM equipment in terms of performance.

The metrics should be independent of switch architecture. For example, "percentage of frames cut through without delay" applies only to switches with the cut-through feature and is not meaningfully applicable to other (store-and-forward) switches. Such architecture-dependent metrics will not be defined.

The Forum plans to develop precise methodologies for measuring the metrics so that anyone can measure and produce the same result. The methodologies will include specific configurations, traffic patterns, and measurement procedures.

NON-GOALS OF THE ATM FORUM WORK

*T*he ATM Forum does not intend to perform any measurements itself. Any vendor, user, or independent laboratory can use the methodologies and metrics developed by the Forum. The Forum does not intend to certify any particular

---

[1] *The views expressed here are solely those of the authors and are not meant to reflect the official position of the ATM Forum.*

measurements or laboratories. Also, the Forum does not intend to set any thresholds of required performance. The frame loss ratio or frame delay variation that is acceptable is left to the user and the supplier. Generally, there is a trade-off between cost and acceptable performance. Users may accept equipment that is slow if it is cheap, while they may expect faster performance from expensive equipment. Different vendors will try to provide different cost-performance trade-off points, and such differentiation is generally good for a technology.

■ Figure 1. *Peak, lossless, and full-load throughput.*

Foreground VCs can be permanent or switched, virtual path or virtual channel connections, established between ports on the same network module on the switch or between ports on different network modules, or between ports on different switching fabrics.

A system with *n* ports is tested for the following connection configurations.

n-to-n Straight — Input from one port exits to another port. This represents almost no path interference among VCs. There are *n* VCs.

METRICS

*M*ost of the metrics discussed here apply to a single switch as well as a network of switches. Therefore, we use "system under test" or just "system" to refer to the device(s) being tested. A partial list of the metrics includes throughput, frame latency, throughput fairness, frame loss ratio, maximum frame burst size, call establishment latency, and application goodput. A brief overview of these metrics follows.

THROUGHPUT

*T*hree different frame-level throughput metrics are defined. *Lossless throughput* is the maximum rate at which none of the offered frames is dropped by the system. *Peak throughput* is the maximum rate at which the system operates regardless of frames dropped. The maximum rate can actually occur when the loss is nonzero. *Full-load throughput* is the rate at which the system operates when the input links are loaded at 100 percent of their capacity.

A model graph of throughput vs. input rate is shown in Fig. 1. Level *x* defines the lossless throughput, level *y* defines the peak throughput and level *z* defines the 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. The peak throughput is the maximum throughput that can be achieved in spite of losses. The full-load throughput is the throughput of the system at 100 percent 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.

Throughput is expressed in effective bits per second, counting only bits from AAL payloads, excluding the overhead introduced by the ATM technology and transmission systems. This is preferred over specifying it in frames per second or cells per second. Frames per second requires specifying the frame size. The throughput values in frames per second at various frame sizes cannot be compared without first being converted into bits per second. Cells per second is not a good unit for frame-level performance because the cells are not seen by the user.

Before starting measurements, a number of virtual channel connections (VCCs), or virtual path connections (VPCs), called *foreground* virtual connections (VCs), are established through the system. Foreground VCs are used to transfer only the traffic whose performance is being measured. That traffic is referred as the foreground traffic. Foreground traffic is specified by the type of foreground VC, connection configuration, service class, arrival patterns, frame length, and input rate.
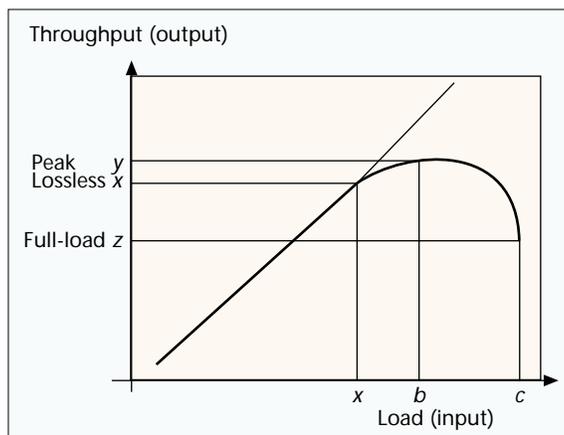
n-to-(n *– 1) Full Cross* — Input from each port is divided equally to exit on each of other (*n* – 1) ports. This represents intense competition for the switching fabric by VCs. There are *n* x (*n* – 1) VCs.

*n-to-*m *Partial Cross* — Input from each port is divided equally to exit on other *m* ports (1 < *m* < *n* – 1). This represents partial competition for the switching fabric by VCs. There are *n* x *m* VCs. Note that *n*-to-*n* straight and *n*-to-*n* full cross are special cases of *n*-to-*m* partial cross with *m* = 1 and *m* = *n* – 1, respectively.

k-to-1 — Input from *k* (1 < *k* < *n*) ports is destined to one output port. This stresses the output port logic. There are *k* VCs.

1-to-(n *– 1)* — Input from one port is multicast to all other output ports. This tests the multicast performance of the switch. There is only one VC.

Different connection configurations are illustrated in Fig. 2, where each configuration includes one ATM switch with four ports, with their input components shown on the left and their output components shown on the right.

The following service classes, arrival patterns and frame lengths for foreground traffic are used for testing:
• Unspecified bite rate (UBR) service class: foreground traffic consists of equally spaced frames of fixed length (uniform bit rate). Measurements are performed at AAL payload sizes of 64 bytes, 1518 bytes, 9188 bytes, and 64 kbytes. Variable-length frames and other arrival patterns (e.g., self-similar) are under study.
• Available bit rate (ABR) service class is under study.

Higher-priority traffic like variable bit rate (VBR) or constant bit rate (CBR) can act as background traffic. Details of background traffic characteristics have not yet been defined.

The input rate of foreground traffic is expressed in effective bits per second, counting only bits from AAL payloads, excluding the overhead introduced by the ATM technology and transmission systems.

It is obvious that testing larger systems (e.g., switches with larger number of ports) could require very extensive (and expensive) measurement equipment. Hence, we introduce scalable test configurations for throughput measurements that require only one ATM monitor with one generator/analyzer pair. Figure 3 presents a sample test configuration for an ATM switch with eight ports in an 8-to-2 partial cross-connection configuration. The configuration emulates 16 foreground VCs.

There is one link between the ATM monitor and the switch. The other seven ports have external loopbacks. A loopback on

the given port causes the frames transmitted over the output of the port to be received by the input of the same port.

The test configuration in Fig. 3 assumes two network modules in the switch, with switch ports P0–P3 in one network module and switch ports P4–P7 in the another network module. In this case, foreground VCs are always established from a port in one network module to a port in another network module.

This connection configuration could be more demanding on the system than the cases where each VC uses ports in the same network module. An even more demanding case could be when foreground VCs use different fabrics of a multifabric switch.

Similar approaches can be used for *n*-to-*n* straight, *n*-to-*n* full cross and other types of *n*-to-*m* partial cross-connection configurations, as well as for larger switches.

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

■ Table 1. *Quantile values for commonly used confidence levels.*

FRAME LATENCY

MIMO latency (message-in message-out) is a general definition of the latency that applies to an ATM switch or a group of ATM switches and is defined as follows:

MIMO latency = min{LILO latency, FILO latency – NFOT}

where:
- LILO (last-in last-out) latency = time between the last-bit entry and the last-bit exit
- FILO (first-in last-out) latency = time between the first-bit entry and the last-bit exit
- NFOT (nominal frame output time) = frame latency through a zero-delay switch

An explanation of MIMO latency and its justification are presented in [3].

The MIMO is a general definition that applies even when the frames are discontinuous at the input and/or output or the input and output rates are different. To measure MIMO latency for a given frame, the times of the following three events should be recorded:

1 The first bit of the frame enters the system.
2 The last bit of the frame enters into the system.
3 The last bit of the frame exits from the system.

The time between events 1 and 3 is FILO latency, and the time between events 2 and 3 is LILO latency. Also, given the frame size and (input and output) link rates, NFOT can be calculated. Then, substituting LILO latency, FILO latency, and NFOT in the MIMO latency formula would give the frame latency of the system.

Contemporary ATM monitors provide measurement data only at the cell level (e.g., CTD and cell interarrival time). This data is sufficient to calculate MIMO frame latency as follows.

If the input link rate is less than or equal to the output link rate,

MIMO latency = Last cell's transfer delay – (Last cell's input transmit time + Monitor overhead)

where:
- The cell input transmit time is the time to transmit one cell into the input link. It can easily be calculated.
- The monitor overhead is the overhead introduced by the ATM monitor when measuring

CTD and is usually nonzero. It can be calculated as the difference between the measured CTD for the case of a closed loop on the ATM monitor and the theoretical value for the cell transmit time plus any propagation delay.

Thus, to calculate MIMO latency when the input link rate is less than or equal to the output link rate, it is sufficient to measure the transfer delay of the last cell of a frame.
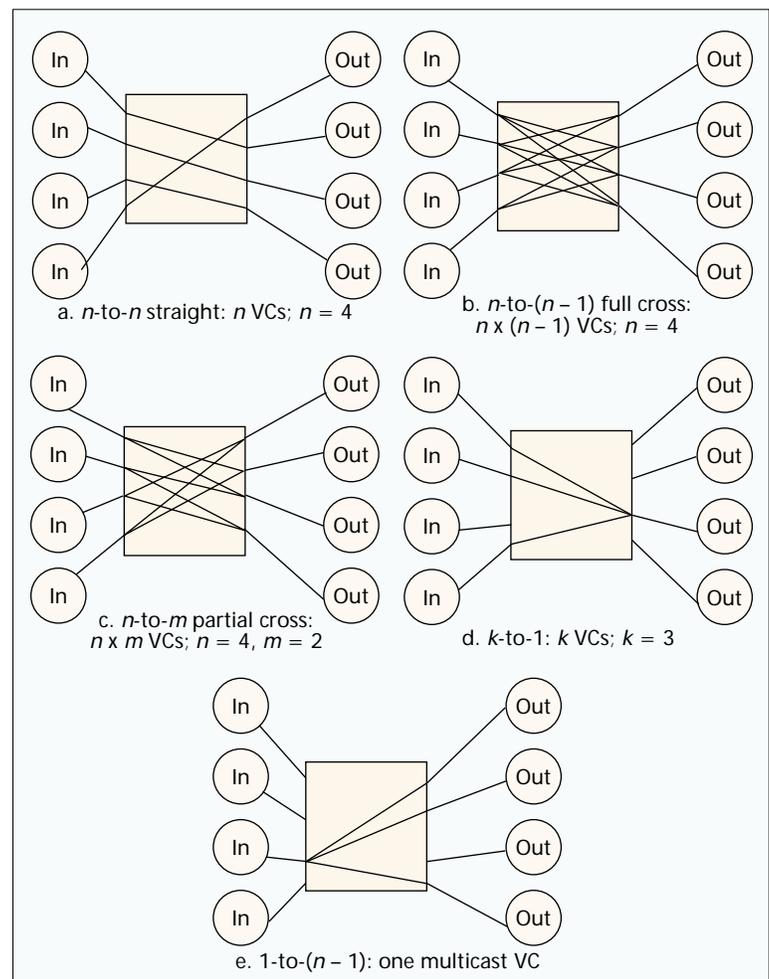
If the input link rate is greater than or equal to the output link rate:

MIMO latency = FIFO latency + FOLO time – NFOT

where:
- FIFO latency = First cell's transfer delay – (First cell's output transmit time + Monitor overhead)
- FOLO time = First cell to last cell interarrival time + Last cell's output transmit time
- The cell output transmit time is the time to transmit one cell into the output link. Again, it can easily be calculated.

Thus, to calculate MIMO latency when the input link rate is greater than or equal to the output link rate, it is necessary

■ Figure 2. *Connection configurations for throughput measurements.*

a. *n*-to-*n* straight: *n* VCs; *n* = 4

b. *n*-to-(*n* – 1) full cross: *n* x (*n* – 1) VCs; *n* = 4

c. *n*-to-*m* partial cross: *n* x *m* VCs; *n* = 4, *m* = 2

d. *k*-to-1: *k* VCs; *k* = 3

e. 1-to-(*n* – 1): one multicast VC

to measure the first cell's transfer delay and the interarrival time between the first and last cells of a frame.

For MIMO latency measurements, it is first necessary to establish one VCC or VPC used only by foreground traffic (the foreground VC), and a number of VCCs or VPCs used only by background traffic (background VCs). Then the background traffic is generated. When the flow of the background traffic has been established, the foreground traffic is generated. After the steady-state flow of foreground traffic has been reached, required times and/or delays needed for MIMO latency calculation are recorded for p consecutive frames, while the flow of background traffic continues uninterrupted. Here, p is a parameter.

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

Mean MIMO latency $= (\Sigma M_i)/p$

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

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

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

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

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

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

The foreground VC can be permanent or switched, VPC or VCC, established between ports on the same network module, between ports on different network modules, or between ports on different fabrics.

For the UBR service class, the foreground traffic consists of equally spaced frames of fixed length. Measurements are performed at AAL payload sizes of 64 bytes, 1518 bytes, 9188 bytes and 64 kbytes. Variable-length frames and other arrival
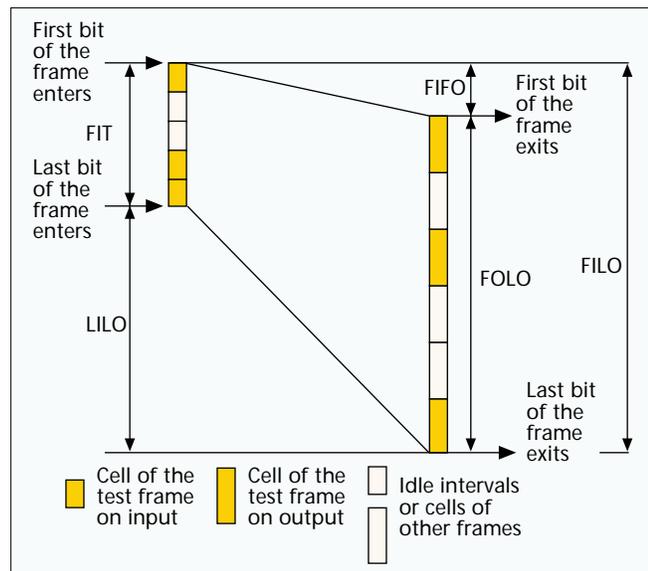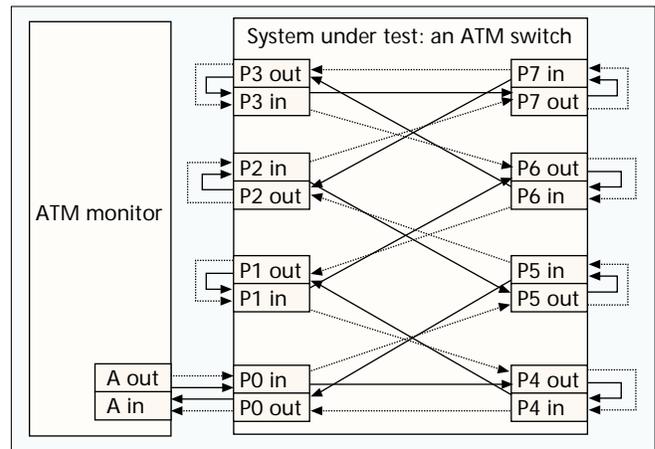
■ **Figure 3**. *A scalable test configuration for throughput measurements using only one generator/analyzer pair with an 8-port switch and an 8-to-2 partial cross-connection configuration.*

patterns (e.g., self-similar) are under study. The ABR service class is also under study.

The input rate of foreground traffic is expressed in effective bits per second, counting only bits from AAL payloads, excluding the overhead introduced by the ATM technology and transmission systems.

The first measurement run is performed at the lowest possible foreground input rate (for the given test equipment). For later runs, the foreground load is increased up to the point where losses occur or up to the full foreground load (FFL). FFL is equal to the lesser of input or output link rate used by the foreground VC.

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

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

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

A system with $w (= n + 1)$ ports is measured for the following background traffic connection configurations:
• n-to-n straight, with n VCs
• n-to-$(n - 1)$ full cross, with n x $(n - 1)$ VCs
• n-to-m partial cross, $1 < m < n - 1$, with n x m VCs
• 1-to-$(n - 1)$, with one multicast VC
These configurations are the same as those shown in Fig. 2.

The following service classes, arrival patterns (if applicable), and frame lengths (if applicable) are used for the background traffic:
• UBR service class — Traffic consists of equally spaced frames of fixed length. Measurements are performed at AAL payload size of 64 bytes, 1518 bytes, 9188 bytes, and 64 kbytes. This is a case of bursty background traffic of priority equal to or lower than that of the foreground

■ **Figure 4**. *Latency metrics.*
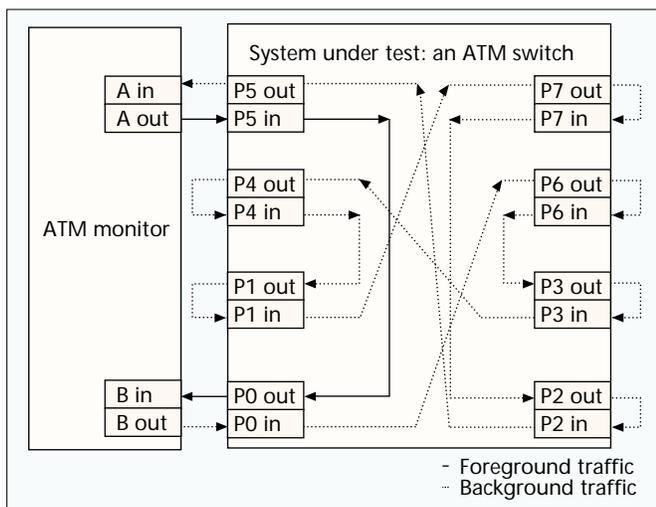
■ Figure 5. *A scalable test configuration for measurement of MIMO latency using only two generator/analyzer pairs with 8-port switch and a 7-to-7 straight connection configuration for the background traffic.*

traffic. Variable-length frames and other arrival patterns (e.g., self-similar) are under study.

• CBR service class — Traffic consists of a contiguous stream of cells at a given rate. This is a case of nonbursty background traffic of priority higher than that of the foreground traffic.

• VBR and ABR service classes — Under study.

Scalable test configurations for MIMO latency measurements require only one ATM monitor with two generator/analyzer pairs. Figure 5 presents the test configuration with an ATM switch with eight ports ($w = 8$). There are two links between the ATM monitor and the switch, and they are used in one direction by the background traffic and in the other by the foreground traffic, as indicated. The other six ($w – 2$) ports of the switch are used only by the background traffic, and they have external loopbacks.

Figure 5 shows a 7-to-7 straight connection configuration for the background traffic. The n-to-$(n – 1)$ full cross and n-to-m partial cross-connection configurations can also be similarly implemented.

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

It should be noted that in these test configurations, if all link rates are not identical, it is not possible to generate background traffic (without losses) equal to MBL. The maximum background traffic input rate in such cases equals ($n – 1$) x lowest link rate. Only if all link rates are identical, it is possible to obtain MBL level without losses in the background traffic.

THROUGHPUT FAIRNESS

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, the ideal allocation is defined by the max-min allocation[2] 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.

---

[2] *Other policies can be used, but must be specified.*

If the actual measured throughputs of n VCs sharing a system (a single switch or a network of switches) are found to be $\{T_1, T_2, ..., T_n\}$, where the optimal max-min throughputs should be $\{\hat{T}_1, \hat{T}_2, ..., \hat{T}_n\}$, then the fairness of the system under test is quantified by the "fairness index" computed as follows [4]:

$$\text{Fairness index} = (\Sigma x_i)^2/(n \times (\Sigma x_i^2))$$

where, $x_i = T_i/\hat{T}_i$ is the relative allocation to the ith VC.

This fairness index has the following desirable properties:

• It is dimensionless. The units used to measure the throughput (bits, cells, or frames per second) do not affect its value.

• It is a normalized measure that ranges between 0 and 1. The maximum fairness is 100 percent, the minimum 0 percent. This makes it intuitive to interpret and present.

• If all x_is are equal, the allocation is fair and the fairness index is one.

• If $n – k$ of n x_is 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.

Throughput fairness is quantified by the fairness index for each of the throughput experiments in which there are either multiple VCs or multiple input output ports. Thus, it applies to all three throughput measures (lossless, peak, and full-load), all connection configurations and all traffic patterns. No additional experiments are required for throughput fairness. The detailed results obtained for the throughput tests are analyzed to compute the fairness.

The throughput tests are run several times for a specified duration. The fairness is computed for each individual run. Let F_i be the fairness index for the ith run; then the mean fairness is computed as follows:

$$\text{Mean fairness} = (\Sigma F_i)/\text{Number of repetitions}$$

Note that the fairness index is not limited to throughput. It can be applied to other metrics, such as latency. However, extreme unfairness in latency is expected to appear as unfairness in throughput and vice versa.

FRAME LOSS RATIO

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

Frame loss ratio = (input frame count – output frame count)/input frame count

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

Peak Throughput Frame Loss Ratio — The frame loss ratio at the input load corresponding to the peak throughput.

Full-Load Throughput Frame Loss Ratio — The frame loss ratio at the input load corresponding to the full-load throughput.

These metrics are related to the throughput as follows:

Frame loss ratio = (input rate – throughput)/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. The throughput experiments are repeated a specified number of times. If FLR_i is the frame loss ratio for the ith run,

$$FLR_i = (\text{input rate}_i - \text{throughput}_i)/\text{input rate}_i$$

Since frame loss ratio is a "ratio," its average cannot be computed via straight summation [4]. The average frame loss ratio (FLR) for multiple runs is computed as follows:

$$FLR = (\Sigma \text{ input rate}_i - \Sigma \text{ throughput}_i)/\Sigma \text{ input rate}_i$$

MAXIMUM FRAME BURST SIZE

*M*aximum frame burst size (MFBS) is the maximum number of frames that a source end system can send at the peak rate through a system without incurring any loss.

MFBS measures the data buffering capability of the system 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, MFBS provides a useful indication.

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

MFBS is expressed in octets of AAL payload. This is preferred over number of frames or cells because the former requires specifying the frame size and the latter is not very meaningful for a frame-level metric. Also, the 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, 536, 1518, and 9188 bytes may be used.

The number of frames sent in the burst is increased successively until a loss is observed. The maximum number of frames that can be sent without loss are reported as MFBS.

CALL ESTABLISHMENT LATENCY

*F*or short-duration VCs, call establishment latency is an important part of 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 (Fig. 6). 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 cells.
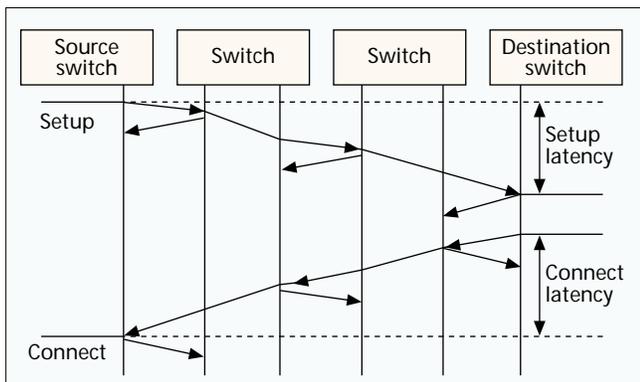
■ **Figure 6.** *Call establishment latency.*

Unlike X.25, frame relay, and integrated services digital network (ISDN) networks, the messages in ATM networks are not contiguous; therefore, the MIMO latency metric is used. Thus,

Call establishment latency = MIMO latency for setup message + MIMO latency for connect message

The call establishment latency as defined above applies to any network of switches. In practice, it has been found that the latency depends on 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 are indicated.

APPLICATION GOODPUT

*A*pplication 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 = frames received in the measurement interval/frames transmitted in the measurement interval

Traditionally, goodput is measured in bits per second. However, we are interested in a nondimensional metric, and primarily 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

- This metric is useful when measured at the peak load. The number of transmitted frames is varied over a useful range from 2000 frames/s (fps) through 10,000 fps at a nominal frame size of 64 bytes. Frame sizes are also varied through 64, 1518, and 9188 bytes to represent small, medium, and large frames, respectively. Note that the frame sizes specified do not account for the overhead of accommodating the desired frame transmission rates over the ATM medium.
- The measurement interval should be chosen to be large enough to accommodate the transmission of the largest packet (frame) over the connection and small enough to track short-term variations of the average goodput.
- It is important not to include network management frames and/or keep-alive frames in the received frames count.
- There should be no changes of handling buffers during the measurement interval.
- The results are to be reported as a table for the three different frame sizes.

THE OSU ATM BENCHMARKING LABORATORY

*A*ny vendor or user can run the benchmarks and tests developed by the ATM Forum. Still, there is a need for an independent measuring organization that can conduct the tests and publish results on a regular basis. The Ohio State University ATM benchmarking laboratory will play this role. The role of this laboratory for ATM testing will be similar to that at Harvard for router and local area network (LAN) switch testing. We have been awarded funding by the National Science Foundation and the State of Ohio for this laboratory.

SUMMARY

*T*he ATM Forum Test and Traffic Management groups are jointly working on defining a set of standard performance metrics and tests. The key distinguishing feature of this work is that it considers the user-perceived performance and therefore uses frame-level metrics rather than the cell-level metrics of the past.

In this article, we provided a brief overview of several metrics that are being defined. The metrics, and their definitions and tests, are currently being refined.

REFERENCES[3]

[1] S. Kelly-Bootle, *The Devil's DP Dictionary*, New York: McGraw-Hill.
[2] R. Jain, "Performance Benchmarking BOF," ATM Forum 95-1347, Oct. 1995, ftp://netlab.ohio-state.edu/pub/jain/atmf/af_perf2.ps.
[3] R. Jain, A. Durresi, and G. Babic, "ATM Forum Performance Testing Specification — Baseline Text," ATM Forum BTD-TEST-TM-PERF.00.01 (96-0810R4), Dec. 1996, ftp://netlab.ohio-state.edu/pub/ jain/atmf/btd-test-tm-perf.00.01.ps.
[4] R. Jain, *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation, and Modeling*, New York: Wiley-Interscience, Apr. 1991.

ADDITIONAL READING

[1] S. Bradner, "Benchmarking Methodology for Network Interconnect Devices," RFC 1944, May 17, 1996.
[2] S. Bradner, "Benchmarking Terminology for Network Interconnection Devices," RFC 1242, July 1991.
[3] R. Jain and B. Nagendra, "Performance Benchmarking of ATM Switches," ATM Forum 95-1662, Dec. 1995, ftp://netlab.ohio-state.edu/ pub/jain/atmf/af_prf22.ps.
[4] R. Jain, B. Nagendra, and G. Babic, "Scope for ATM Forum's Performance Benchmarking Work Item," ATM Forum 96-0180, Feb. 1996, ftp://netlab.ohio-state.edu/pub/jain/atmf/af_prf32.ps.
[5] R. Jain, B. Nagendra, and G. Babic, "General Considerations for Frame-Level Performance Measurement of ATM Switches," ATM Forum 96-0519, Apr. 1996, ftp://netlab.ohio-state.edu/pub/jain/atmf/af_prf42.ps.
[6] R. Jain, B. Nagendra, and G. Babic, "Considerations for Frame Level Throughput and Latency Measurements of ATM Switches," ATM Forum 96-0520, Apr. 1996, ftp://netlab.ohio-state.edu/pub/jain/atmf/ af_prf52.ps.
[7] G. Babic, A. Durresi, and R. Jain, "ATM Switch Performance Testing Experiences, Feb. 1997," ATM Forum/97-0178, ftp://netlab.ohio-state.edu/pub/jain/atmf/atm97-0178.ps.
[8] C. D. Krivda, "Analyzing ATM Adapter Performance: The Real-World Meaning of Benchmarks," Efficient Networks white paper, http://www.efficient.com/dox/EM.html.
[9] R. Nair and D. Hunt, "Performance Metrics for Traffic Management," ATM Forum 95-1680, Dec. 1995.
[10] R. Nair and D. Hunt, "On the Scope of the Performance Benchmarking BOF," ATM Forum 96-0255, Feb. 1996.
[11] R. Nair, "Performance Benchmarking Methodology for ABR Applications," ATM Forum 96-0580, Apr. 1996.
[12] R. Nair, "Performance Metrics for ATM Layer Data Services," ATM Forum/96-0902, June 1996.
[13] M. Rowe, "Wealth of ATM Testers — Answers Most Needs," *Test and Measurement World*, Sept. 1995.
[14] Wakid *et al.*, "Architectures for BISDN Networks: A Performance Study," NIST report, http://www.hpcc.gov/blue94/ section.4.7.html.
[15] LanQuest Inc., "ATM vs Ethernet Performance Benchmark Comparison," May 1996, http://www. lanquest.com/reports/pool/rptadpt.htm.
[16] R. Jain, D. Chiu, and W. Hawe, "A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Computer Systems," DEC Tech. Rep. DEC- TR-301, Sept. 1984, ftp://netlab.ohio-state.edu/pub/jain/papers/fairness.ps

BIOGRAPHIES

RAJ JAIN [F] is a professor of computer and information sciences at Ohio State University in Columbus, Ohio. Prior to joining the University in April 1994, he was a senior consulting engineer at Digital Equipment Corporation, Littleton, Massachusetts, where he was involved in design and analysis of many computer systems and networks, including VAX clusters, Ethernet, DECnet, OSI, FDDI, and ATM networks. Currently he is very active in the Traffic Management Working Group of ATM Forum and has influenced its direction considerably. He is also the editor of the ATM Forum's Performance Testing Specification. He received a B.E. in electrical engineering from A.P.S. University, Rewa, India, in 1972, an M.E. in automation from the Indian Institute of Science, Bangalore, 1974, and a Ph.D. in applied math (computer science) from Harvard University, Cambridge, Massachusetts, in 1978. He is the author of *The Art of Computer Systems Performance Analysis* (Wiley, 1991) and *The FDDI Handbook: High-Speed Networking with Fiber and Other Media* (Addison Wesley, 1994).

GOJKO BABIC [SM] received his graduate degree in electrical engineering from the University of Sarajevo in 1972, an M.S. degree in computer science from Florida Institute of Technology in 1975, and a Ph.D. in computer science from Ohio State University in 1978. From 1983 to 1990, he was a leader and principal investigator of the program Energonet, the largest and most important research and development project in computer networking in the former Yugoslavia. Its main results were a public X.25 network for Bosnia-Herzegovina and an X.25 network in China for the Ministry of Public Security. From 1988 to 1992, he was director of the Computer Networking Department at IRIS-Computer, a division of Energoinvest Corporation, and an associate professor at the University of Sarajevo. He is currently a senior lecturer at Ohio State University. He has been an ACM member since 1975.

---

[3] *All our past ATM forum contributions and presentations are available on-line at http://www.cis. ohio-state.edu/~ jain/*