

# Testing Policing in ATM Networks

Policing is one of the key mechanisms used in ATM (Asynchronous Transfer Mode) networks to avoid network congestion. The HP E4223A policing and traffic characterization test application has been developed to test policing implementations in ATM switches before the switches are deployed for commercial service.

by **Mohammad Makarechian and Nicholas J. Malcolm**

---

The Asynchronous Transfer Mode (ATM) is a network technology that can satisfy the quality-of-service requirements of many different types of traffic. The ability of ATM to handle many different types of traffic and its ability to operate at high bandwidths position it to be one of the core technologies behind future broadband wide area networks and the Internet. To provide quality-of-service guarantees, ATM relies crucially upon avoiding network congestion. Congestion can result in unacceptably large cell loss or delays. Cells that are lost may have to be retransmitted, which can result in increased congestion. Cells with excessive delays can cause higher-layer protocol timers (e.g., TCP/IP timers) to expire, which will result in even more cells being retransmitted. For video traffic, excessive delays or cell delay variation can result in underflow in the video decoder buffers. This can cause jagged movements or screen freezes when playing back the video.

Policing is one of the key mechanisms used by ATM to avoid network congestion. Policing is responsible for monitoring the amount of traffic sent by a connection. If a connection is sending more than the agreed-upon amount of traffic, then policing can discard traffic from the offending connection. By preventing too much traffic from entering the network, policing helps to avoid network congestion. This ensures that existing connections in the network will continue to receive their required quality of service.

Policing occurs at the *user-network interface* (UNI), where user traffic first enters a public network, and at the *broadband ISDN intercarrier interface* (B-ICI), where traffic crosses from one public network to another. Policing is known as *usage parameter control* (UPC) at the UNI and *network parameter control* (NPC) at the B-ICI.

Given the importance of policing to ATM, it is essential that policing be well-tested. Policing must be tested both by network equipment manufacturers when developing switches, and by network providers when commissioning switches. The HP E4223A policing and traffic characterization test application has been developed to test policing implementations in ATM switches. This product allows users to generate policing test traffic and to measure how the traffic is affected by policing. In this way the HP E4223A can thoroughly test policing in ATM switches before the switches are deployed for commercial service. The HP E4223A can also analyze the traffic originating from a traffic source to determine whether the source is sending too much traffic into the network.

## How Policing Works

For ATM to meet its quality-of-service commitments, it is essential to reduce or eliminate network congestion. Congestion can result in unacceptably poor network performance. ATM attempts to avoid congestion by managing network resources (e.g., transmission links, buffer space inside switches) in such a way that congestion will not occur. A connection will only be established if there are enough network resources to provide an acceptable quality of service to the new connection without disrupting the service provided to existing connections. Once a connection has been established, usage parameter control (UPC) is responsible for policing the connection traffic when it enters the network and ensuring that the traffic does not exceed the agreed-upon traffic rate.

Depending on the requirements of the traffic source, ATM provides a variety of service categories, as shown in Fig. 1.<sup>1</sup> For example, an application such as digital voice may be suited for *constant bit rate* (CBR) service, while compressed video may be suited for *real-time variable bit rate* (rt-VBR) service. The service category used by a connection is chosen at connection setup time. For each service category, several traffic parameters are given to the network to describe the type of traffic that will be sent by the connection. For *switched virtual connections* (SVCs), the traffic parameters are given to the network during the call setup or renegotiation phase of signaling. For *permanent virtual connections* (PVCs) the traffic parameters can be specified manually at subscription time. The traffic parameters are used by the network to police the traffic on the connection and to determine how many network resources must be reserved to support the connection.

Cells in an ATM network can be given a high priority or a low priority. High-priority cells have the *cell loss priority* (CLP) bit in their headers set to 0, while low-priority cells have a CLP of 1. Low-priority cells are more likely to be discarded if the network becomes congested. At the minimum, the traffic parameters declared to the network at connection setup time include the *peak cell rate* (PCR) and the *cell delay variation tolerance* (CDVT) for CLP = 0+1 cells, that is, for all cells in the connection, regardless of priority.

Service Category *	Characteristics	Example Applications	Traffic Parameters
Constant Bit Rate (CBR)	<ul style="list-style-type: none"> <li>Tightly bounded cell delay variation</li> <li>Static amount of bandwidth available throughout a connection's lifetime</li> <li>Low cell loss ratio</li> </ul>	Video/audio on demand, video conferencing, digital telephony	<ul style="list-style-type: none"> <li>PCR, CDVT on CLP= 0+ 1 cells</li> </ul>
Real-Time Variable Bit Rate (rt-VBR)	<ul style="list-style-type: none"> <li>Supports bursty traffic</li> <li>Tightly bounded cell delay variation</li> <li>Low cell loss ratio</li> </ul>	Compressed video, distributed classroom	<ul style="list-style-type: none"> <li>PCR, CDVT and SCR, MBS on CLP= 0+ 1 cells</li> <li>PCR, CDVT on CLP= 0+ 1 cells, SCR, MBS on CLP= 0 cells</li> <li>PCR, CDVT on CLP= 0+ 1 cells, SCR, MBS on CLP= 0 cells, tagging applicable</li> </ul>
Non-Real-Time Variable Bit Rate (nrt-VBR)	<ul style="list-style-type: none"> <li>Supports bursty traffic</li> <li>No cell delay variation bounds</li> <li>Low cell loss ratio</li> </ul>	Airline reservations, banking transactions	<ul style="list-style-type: none"> <li>PCR, CDVT on CLP= 0+ 1 cells, SCR, MBS on CLP= 0 cells, tagging applicable</li> </ul>
Unspecified Bit Rate (UBR)	<ul style="list-style-type: none"> <li>No cell delay variation bounds</li> <li>No cell loss ratio bounds</li> <li>"Best effort" service</li> </ul>	File transfer, e-mail	<ul style="list-style-type: none"> <li>PCR, CDVT on CLP= 0+ 1 cells</li> <li>PCR, CDVT on CLP= 0+ 1 cells, tagging applicable</li> </ul>

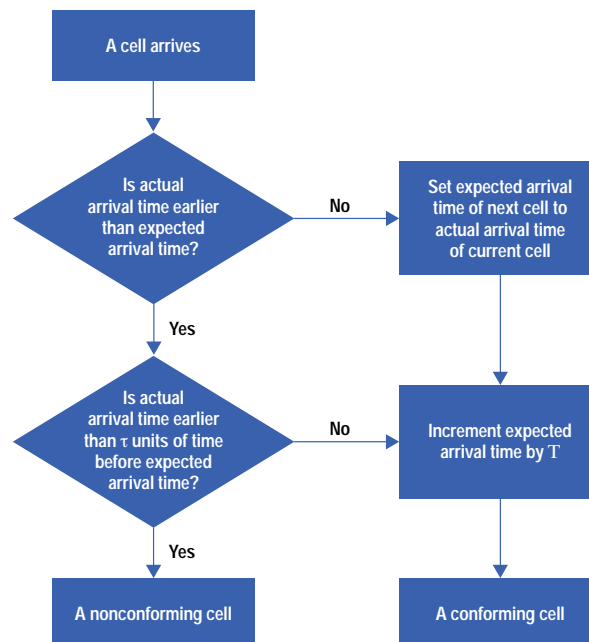
\* The ATM Forum also defines a service category called available bit rate (ABR). Policing of ABR connections is not discussed in this article.

**Fig. 1. ATM service categories.**

The PCR is the maximum rate at which the source may generate traffic. The CDVT indicates how many back-to-back cells there may be at the user-network interface. Together, the PCR and the CDVT give the network an idea of when to expect the next arrival of a cell given that one has just arrived. In addition, for variable bit rate service categories, a traffic source can also specify a *sustainable cell rate* (SCR) and a *maximum burst size* (MBS). The SCR gives an upper bound on the conforming cell rate of a VBR connection. The MBS gives the maximum burst size for a VBR connection, assuming that the cells in the burst arrive at the PCR. Specifying the SCR and MBS allows the network to allocate resources such as buffer space more efficiently because the network has more knowledge about the type of traffic that will be generated.

The UPC function is responsible for ensuring that traffic on a connection does not exceed the agreed-upon rate. The policing function in a switch first validates the VPI/VCI (*virtual path identifier/virtual channel identifier*) of arriving cells, and then determines whether or not the cells are conforming to the agreed-upon PCR or SCR. Whether or not a cell is conforming is determined by an algorithm called the *generic cell rate algorithm* (GCRA),\* popularly known as the "leaky bucket" algorithm (Fig. 2). The GCRA has two parameters, denoted T and  $\tau$ . The first parameter, T, is the emission interval and can be regarded as the expected interarrival time of conforming cells. The second parameter,  $\tau$ , is the cell delay variation tolerance (CDVT) and determines how many back-to-back cells are allowed. The GCRA maintains a variable called the *theoretical arrival time* (TAT), which gives the expected arrival time of the next cell. Cells arriving more than  $\tau$  units of time before the TAT are considered to be nonconforming. Nonconforming cells can be tagged (given a lower priority) or

\* The GCRA is a reference algorithm used to define conformance. An actual UPC implementation may use the GCRA or another algorithm provided that the quality-of-service objectives for connections are met.



**Fig. 2. Generic cell rate (leaky bucket) algorithm.**

discarded by the switch. Cells arriving  $\tau$  units of time before the TAT or later are considered to be conforming. For each conforming cell, the TAT is updated to give the expected arrival time of the next cell in the connection.

For an example of how the GCRA can be used to police a CBR connection, suppose that a connection has a peak cell rate of  $PCR = 8000$  cells/s and a CDVT of  $0 \mu s$ . The GCRA emission interval is calculated as  $T = 1/PCR = 125 \mu s$ . The GCRA CDVT is  $\tau = 0 \mu s$ . Table I shows the resultant conforming and nonconforming cells assuming that a cell on the connection arrives every  $120 \mu s$ .

**Table I**  
Conforming and Nonconforming Cells  
for  $PCR = 8000$  cells/s,  $\tau = 0 \mu s$

Cell #	$t_{arrival}$ ( $\mu s$ )	TAT ( $\mu s$ )	Conforming?
1	0	0	yes
2	120	125	no
3	240	125	yes
4	360	365	no
5	480	365	yes
6	600	605	no
7	720	605	yes
8	840	845	no
9	960	845	yes
10	1080	1085	no

To see how the CDVT can be increased to allow more back-to-back cells, suppose that the CDVT used by the GCRA is increased to  $\tau = CDVT = 11 \mu s$ . The resultant pattern of conforming and nonconforming cells is shown in Table II.

**Table II**  
Conforming and Nonconforming Cells  
for  $PCR = 8000$  cells/s,  $\tau = 11 \mu s$

Cell #	$t_{arrival}$ ( $\mu s$ )	TAT ( $\mu s$ )	Conforming?
1	0	0	yes
2	120	125	yes
3	240	250	yes
4	360	375	no
5	480	375	yes
6	600	605	yes
7	720	730	yes
8	840	855	no
9	960	855	yes
10	1080	1085	yes

As mentioned earlier, some service categories also have sustainable cell rate (SCR) and maximum burst size (MBS) parameters. In this case, one GCRA is used to police the peak cell rate and another GCRA is used to police the sustainable cell rate. The GCRA tolerance used with the SCR GCRA ( $\tau_{scr}$ ) is derived from the PCR, SCR, MBS, and CDVT parameters,<sup>1</sup> and permits a burst of MBS cells at the peak cell rate. The PCR and SCR GCRA's form a dual leaky bucket algorithm and operate in lockstep fashion. A cell is only considered to be conforming if it conforms to both GCRA's. If tagging is allowed, then high-priority  $CLP = 0$  cells can be tagged (given a lower priority) if they do not conform to the SCR GCRA.

In addition to its role in the UPC function, the GCRA can also be used inside a traffic source to ensure that the outgoing cell flow conforms to a particular cell rate. This is referred to as traffic shaping. When nonconformance is detected during shaping, the offending cells are delayed until their transmission will be conforming. In this way, traffic can be guaranteed conforming before it enters the network.

### The HP BSTS Policing Application

The HP E4223A policing and traffic characterization test application is designed to test UPC implementations in network equipment and to analyze the characteristics of traffic on a connection. The HP E4223A is part of the HP Broadband Series Test System (BSTS).<sup>2</sup> The HP Broadband Series Test System contains a number of VXibus modules that allow testing of broadband networks over a variety of physical interfaces. For brevity, the HP E4223A will be denoted the HP BSTS policing application during the remainder of this article.

The HP BSTS policing application works with the HP E4209 cell protocol processor, a VXibus module forming part of the HP BSTS. The cell protocol processor in conjunction with a line interface module can transmit and receive ATM cells for testing purposes. Received ATM cells can be stored in a capture RAM for later analysis. The HP BSTS policing application consists of embedded software running on the cell protocol processor module and a software component running on HP 9000 Series 400 or 700 workstations.

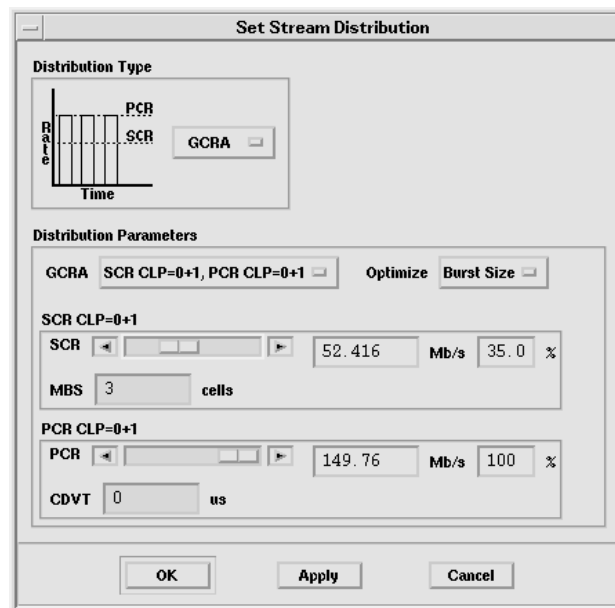
The HP BSTS policing application provides the following functions:

- Generates traffic conforming to a single or dual leaky bucket algorithm (GCRA).
- Generates UPC test cells, which are test cells designed for testing policing.
- Makes a number of policing-related measurements on captured ATM cells, such as the number of nonconforming cells and the number of cells that were lost or tagged.
- Makes general performance measurements on captured ATM cells, such as cell delay, interarrival time, and one-point cell delay variation.

## Traffic Generation

When generating traffic to test policing, it is important to test the limits of the GCRA being used for policing. This means that it is important to generate traffic that has the maximum cell rate and burst size but is still conforming. Because policing is configured in a switch using the parameters of a GCRA, it is convenient to use the parameters of a GCRA when specifying traffic to test policing. Consequently, the HP BSTS policing application provides a GCRA distribution, which allows traffic to be generated using the parameters of a GCRA (Fig. 3). The GCRA combinations supported are:

- PCR CLP = 0+1 (single leaky bucket)
- SCR CLP = 0+1 and PCR CLP = 0+1 (dual leaky bucket).



**Fig. 3.** Specifying the generic cell rate algorithm distribution with the HP E4223A policing and traffic characterization test application.

The GCRA distribution can be optimized to generate traffic based on either the cell rate or the burst size. This allows independent testing of how policing implementations handle cell rates and burst sizes.

Traffic that optimizes the burst size consists of repeated bursts of the maximum possible conforming burst size. The burst is at the line rate for the single leaky bucket. For the dual leaky bucket, the burst consists of MBS cells at the peak cell rate. The spacing between the bursts is the minimum necessary to maintain conforming traffic. For example, suppose a dual leaky bucket is chosen with SCR = 35%, MBS = 3 cells, PCR = 100%, and CDVT = 0 ms. The generated traffic will consist of repeated bursts of three cells at 100% of the line rate, with a gap of six cells between any two bursts. Because cell transmission is quantized, the generated load is 33.3%, which is less than the SCR. Traffic that optimizes the burst size has very precisely controlled burst sizes, but the rate of the generated traffic may be less than requested.


Traffic that optimizes the cell rate consists of traffic generated at the maximum conforming rate. The traffic will consist of an initial burst of the maximum conforming burst size, followed by cells at the PCR (for a single leaky bucket) or at the SCR (for a dual leaky bucket). For example, with a dual leaky bucket with SCR = 35%, MBS = 3 cells, PCR = 100%, and CDVT = 0 ms, the generated traffic consists of an initial burst of three cells, followed by conforming traffic at the SCR. When

optimizing the cell rate, the generated traffic rate can be much closer to the requested rate than traffic that optimizes the burst size.

## Policing Measurements


When generating traffic to test policing, the number of tagged or discarded cells must be measured. It is difficult for test equipment to make these measurements with regular user traffic. This is because the test equipment usually does not know how many user cells were transmitted or the original priority of the user cells. For this reason, test cells are often used to measure policing performance. The HP BSTS policing application can transmit sequences of UPC test cells, each cell having the format shown in Fig. 4. The UPC test cells are specifically designed for testing policing with the HP BSTS policing application. The payload of each UPC test cell within a sequence contains the following information:

- $SN_{0+1}$ . The number of previous test cells in the sequence.
- RES. Reserved bytes, set to zero.
- $SL_{0+1}$ . The total number of test cells in the sequence. The HP BSTS policing application can repeatedly transmit sequences of 512 or 1024 test cells.
- $SL_0$ . The total number of high-priority test cells in the sequence.
- OCLP. The priority of the test cell when it is first transmitted. The low-order bit is set to 0 for high-priority cells and 1 for low-priority cells. The remaining bits in this field are set to zero.
- $SN_0$ . The number of previous high-priority test cells in the sequence.
- VN. The version number of the test cell format. The current version number is 0.
- CRC-16. A cyclic redundancy check error code to provide protection and validation of the encoded payload information. The CRC-16 code is computed using the polynomial  $x^{16} + x^{12} + x^5 + 1$ .



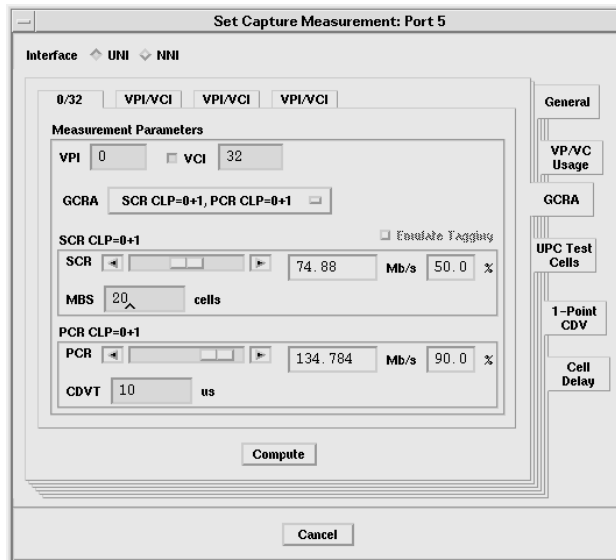
**Fig. 4.** Payload of HP BSTS policing application UPC (usage parameter control) test cell.

The information contained in the payload of UPC test cells allows a number of measurements to be made on captured ATM cells (Fig. 5). These measurements include the number of lost or tagged cells, which are measurements directly relevant to testing policing.



**Fig. 5.** UPC test cell measurements.

In addition to making measurements with UPC test cells, the HP BSTS policing application can measure the conformance of traffic on a connection (Fig. 6). Conformance is measured by saving ATM cells in the capture RAM and then measuring the number of captured nonconforming cells. These measurements can be used to test the number of nonconforming cells detected by a switch or to test whether the traffic on a connection is conforming.



**Fig. 6.** GCRA (generic cell rate algorithm) conformance measurements.

### Testing Policing in a Switch

Policing in ATM switches must work correctly if ATM is to realize its potential for providing guaranteed quality of service for many different types of traffic. This requires that policing be thoroughly tested, both during switch development and during switch deployment. There are basically two aspects of policing to be tested: conforming cells should not be tagged or discarded, and nonconforming cells should be tagged or discarded to protect the quality of service provided to other connections. To test the above aspects of policing, the number of lost cells, the number of tagged cells, and the number of lost high-priority cells must all be measured (Table III).

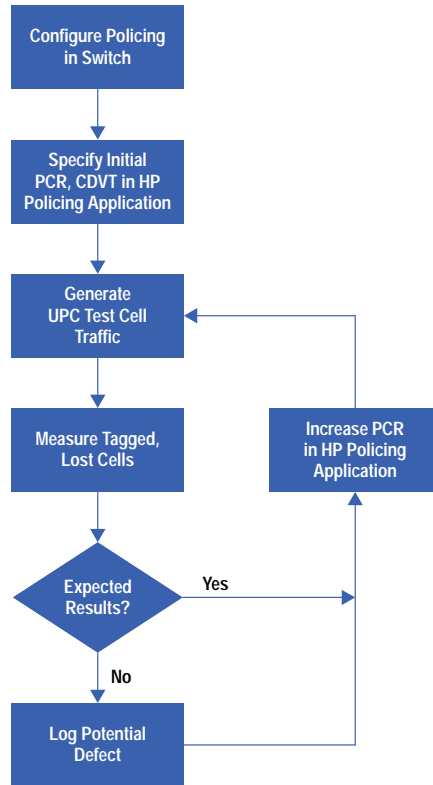
**Table III**  
Parameters to Measure when Testing Policing

Parameter	Description
Lost cells	Number of cells discarded or lost in the switch. This parameter is used to check that policing is not discarding too many cells.
Tagged cells	Number of cells tagged (changed from high to low priority) by the switch. This parameter is used to check that policing only changes the priority of a cell when necessary.
Lost high-priority cells	Number of high-priority cells discarded or lost in the switch. This parameter is used to check that policing is not discarding too many high-priority cells.

When testing policing in a switch with the HP BSTS policing application, the approach is to transmit a well-understood stream of test cells into the switch, capture the cells after they have traversed the switch, and then calculate how many cells were tagged or discarded. This approach is well-suited for stimulus-response type testing to test the capabilities of the switch systematically.

To simplify testing, the overall philosophy when testing policing in a switch is to test one GCRA parameter at a time. This means keeping the cell rate constant while varying the burst size, or keeping the burst size constant while varying the cell rate. Fig. 7 shows how to test the cell rate of a single leaky bucket. The switch is first configured with the leaky bucket parameters to be tested—in this case, the PCR and CDVT for a single leaky bucket. The PCR and CDVT used to generate test traffic are then entered, with the PCR used to generate the traffic being lower than the PCR in the switch. The testing then iterates between measuring the number of tagged or lost cells and increasing the PCR. If the number of tagged or lost cells differs from what is expected, a potential defect is logged.

**Example: Testing a single leaky bucket in a switch.** The approach in Fig. 7 was followed to test a single leaky bucket (GCRA) in a switch with PCR = 8 Mbits/s and CDVT = 60 ms on a 155-Mbit/s SONET port. The HP BSTS policing application was used to generate traffic consisting of a repeating sequence of 1024 UPC test cells conforming to a GCRA with PCR = 4% (5.9 Mbits/s) and CDVT = 60 ms. The traffic was sent through the switch and back into the HP BSTS, where it was placed in the cell protocol processor capture RAM. The ratio of lost cells was calculated based on the captured UPC test cells. The PCR was then incremented by 0.5% for the next iteration of the test. The expected cell loss ratio was 0 if the generated cell rate was



**Fig. 7. Testing the PCR (peak cell rate) of a single leaky bucket.**

less than the policing cell rate, otherwise the expected cell loss ratio was the proportion of generated cell rate greater than the policing cell rate, that is  $(PCR_{\text{traffic}} - PCR_{\text{police}})/PCR_{\text{traffic}}$ . Table IV shows the test results.

**Table IV  
Test Results for a Single GCRA in a Switch**


PCR for Generating Traffic	Cell Loss Ratio	Expected Result?
4.0%	0.00	yes
4.5%	0.00	yes
5.0%	0.00	yes
5.5%	0.02	yes
6.0%	0.10	yes
6.5%	0.17	yes
7.0%	0.23	yes

### Testing Traffic Conformance

Although policing in network switches will protect the network from traffic sources that send too much traffic, it is also important for traffic sources themselves to generate conforming traffic if possible. If a source generates nonconforming traffic, then the nonconforming cells will be discarded by the network and may have to be retransmitted by the source. This can significantly degrade the network performance experienced by the traffic source. The HP BSTS policing application can be used to check whether a source is generating conforming traffic.

**Example: Testing the conformance of MPEG-2 video traffic.** This example demonstrates how to use the HP BSTS policing application to measure the conformance of a traffic source. As shown in Fig. 8, a laser disk player was connected to a commercial MPEG-2 encoder with a 45-Mbit/s DS3 ATM output. The encoder was set up to generate MPEG-2 video over ATM at 4 Mbits/s. The user's guide for the encoder states that a CDVT of 100 ms should be sufficient to compensate for the effects of adapting MPEG-2 packets to ATM cells.

The ATM output of the MPEG-2 encoder was first sent directly to the HP BSTS, where the MPEG-2 traffic was placed in the cell protocol processor capture RAM. To verify that the MPEG-2 traffic was being captured correctly, the HP E4226B MPEG-2 protocol viewer test software was used to play back the captured video segment. The HP BSTS policing application was then used to measure the number of nonconforming cells with a PCR CLP = 0+1 GCRA. The GCRA parameters were



**Fig. 8.** Testing the conformance of video traffic.

chosen conservatively to be PCR = 4.07 Mbits/s and CDVT = 200 ms. The HP BSTS policing application measurements showed that the MPEG-2 encoder was not well-behaved, with approximately 25% of the cells being nonconforming.

To see the effect of the nonconforming cells on the video traffic, the output of the MPEG-2 encoder was then directed to an ATM switch before being routed to the HP BSTS. The switch was configured to police the MPEG-2 traffic with PCR = 4.07 Mbits/s and CDVT = 200 ms. Like the HP BSTS, the switch detected approximately 25% of the cells as being nonconforming. These nonconforming cells were discarded by the switch. The remaining cells passed through the switch and were captured in the cell protocol processor capture RAM. However, because of the large number of ATM cells that were discarded by the switch, it was not possible to play back even one video frame. This example clearly demonstrates the importance of generating conforming ATM traffic and shows how the HP BSTS policing application can be used to test the conformance of a traffic source.

## Conclusion

Policing network traffic at the UNI or B-ICI is crucial to maintaining quality-of-service guarantees in ATM-based networks. The ability to support the quality-of-service requirements of many different types of traffic is one of the distinguishing features of ATM. This feature means that ATM is well-suited to providing the backbone network for future broadband wide area networks and the Internet. The HP BSTS policing application enables switch vendors and service providers to test policing and helps ensure the successful deployment of ATM.

## Acknowledgments

The authors would like to acknowledge the contributions of many individuals who participated in the development and deployment of the HP BSTS policing application, including Brian Smith (project manager), Lawrence Croft (product design), Judith Watson (learning products), Mark Leonard (usability engineer), Drew Paterson, Jack Lam, and Scott Reynolds (QA testing), and Reto Brader (product marketing).

---



---

## References

1. *ATM Forum Traffic Management Specification Version 4.0*, The ATM Forum Technical Committee, March 1996.
  2. The URL for the HP BSTS is <http://www.hp.com/go/bsts>.
- 
- 

- ▶ [Go to Subarticle 11a](#)
- ▶ [Go to Next Article](#)
- ▶ [Go to Journal Home Page](#)