Agilent Technologies

ATM Quality of Service

What It Is, How It Is Acheived, How To Test It

White Paper

By Gregan Crawford
Product Marketing Manager

Agilent Technologies
Network Systems Test Division

Agilent Technologies
Innovating the HP Way

Contents

Introduction

One attribute that makes ATM stand out is its ability to deliver a guaranteed “quality of service” (QoS), end to end over an entire network. Work is continuing in the standards bodies to develop methods for delivering QoS via other technologies but so far ATM is the only flexible technology that really achieves this goal. This white paper discusses the subject of ATM quality of service – what it is, why guarantees are possible, and how to test ATM QoS.

No attempt is made to give an exhaustive coverage of the subject – those interested should dive into the ITU-T’s recommendations (I.356, I.357, I.371, etc) and/or version 4.1 of the ATM Forum’s Traffic Management specification (af-tm-0056.000). Instead, this paper aims to explain in straight forward terms what ATM QoS is all about.

What is ATM QoS

For ATM, the degradation of quality of service is measured in terms of the cell loss ratio (CLR – cells lost with respect to cells sent), cell delay variation (CDV) and cell transfer delay (CTD) impairments introduced by the network. The impact of these impairments on a particular application depends on the nature of that application. For example, a lost cell will usually destroy a data packet, requiring that a higher layer process resend it, while cell delay variation may not impact it at all. Conversely, if a constant bit rate service (e.g. a DS1 or E1 structure) is being carried across an ATM network using “circuit emulation service” (CES), excessive cell delay variation introduced by the network may delay a cell so long that the carried samples belonging to the service arrive too late to be useful. In this case, the late cell might as well have been lost.

Other ATM impairments that can occur include “errored cells” and “misinserted cells”. Errored cells contain one or more bit errors and are measured as the cell error ratio (CER) – that is, cells errored with respect to cells sent. Misinserted cells are cells originally belonging to another virtual channel which have erroneously arrived in the wrong virtual channel. Misinserted cells are measured as the cell misinsertion rate (CMR) – that is, the number of cells misinserted with respect to time. This can happen, for example, due to undetected bit errors in the cell header, although it is relatively rare in practice.

The major advantage ATM has over some other technologies is that it is “connection orientated”. This means that the path over which cells travel is fixed at call set-up time (whether by provisioning or by signaling). Packet orientated technologies usually cannot guarantee that a particular path is always followed, although “multi-protocol label switching” (MPLS) methods are beginning to solve this.

So why is connection orientation so important? Because it is much easier to manage the traffic over such a connection. The answer lies also in the agreement that the user and network make, known as the “traffic contract”.

Traffic Contract and Traffic Parameters

In order that the network can provide adequate resources to allow it to guarantee the user’s QoS requirements, it needs to be told by the user the detailed behavior of the traffic that will be sent over the virtual channel. The traffic contract thus established at call set-up time provides benefits and responsibilities to each party:

User: “I need a guaranteed QoS in terms of CLR and CDV values. In return I promise not to exceed the usage parameter values (bandwidth, burstiness, etc.) I have requested.”

Network: “I promise to deliver the agreed QoS provided you do not exceed the agreed usage parameters. I will police your usage and, if you break the contract, I will take action including discarding your cells in order to protect other users of the network.”

or

> "I do not have the network resources you have requested / I cannot guarantee the QoS you require, so I cannot accept your call" (equivalent to "equipment busy").

The user traffic is always specified in terms of its PCR (peak cell rate) and, if it is bursty in nature, by its SCR (sustainable cell rate, i.e. the long term average rate) and the MBS (maximum burst size, i.e. the number of cells in the maximum sized burst). Note that the PCR of traffic with a constant cell rate (derived from emulation of a constant bit rate service) is the same as that cell rate – this often gives rise to confusion since a constant bit rate service is, in principal, not peaky!

The above parameters could describe the shape of traffic which is carrying data packets. For example, when a packet is being sent, cells carrying that packet may cross the UNI at a constant rate (the PCR). Then, when the packet has been sent, there is silence on that virtual channel until the next packet needs to be sent. If, for example, each packet in a multiple packet stream arrived at a rate of 15 cells per second, the packets were 30 cells long (i.e. they take two seconds to cross the UNI), and they occurred once every 6 seconds, we would have PCR=15, MBS=30, SCR=5. Note that the PCR may not be the same as the link rate (i.e. maximum cell rate).

It is clear from this example that, if we were really lucky regarding timing, we could interleave three similar packet streams (each described by the same traffic parameters) on three different virtual channels on the same ATM link without exceeding the overall PCR. This would be possible because of the extra detail provided by the SCR and MBS parameters. Had we not had these parameters, we would have had to treat each virtual channel as if it were carrying a constant bit rate (i.e. constant cell rate) service at the PCR and this would have wasted link and network bandwidth.

In reality, cell multiplexing between different virtual channels on the same link usually causes cells during a burst to be displaced from their ideal position – they can become bunched, though the average rate during the burst might remain the same. Instantaneously, the PCR may be exceeded (rather contradicts the term "peak", doesn't it?). The network has to tolerate this, provided the excess rate does not occur for too long and is accompanied by a correspondingly cell rate which is less than the PCR. An additional parameter is thus required, the cell delay variation tolerance (CDVT), which is usually measured in microseconds.

To summarize:

- PCR: peak cell rate (= average cell rate for a CBR service)

- SCR: sustainable cell rate – the long term average cell rate of a bursty traffic profile

- MBS: maximum burst size – the maximum number of cells in a burst

- CDVT: cell delay variation tolerance – the amount of cell bunching to be tolerated (that is, the cell rate above the PCR, typically up to the line rate).

Policing

The network does not trust the user to keep to the contract unsupervised, so a policing mechanism is implemented on the network side of the user-network interface (UNI), normally in the input port of the first ATM switch or cross-connect at the edge of the ATM network. If the network did not have this policing mechanism, individual users could accidentally or maliciously send excessive traffic causing congestion and consequent damage (cell loss and cell delay) to the traffic of other users in unforeseeable ways.

Where permanent virtual circuits (PVCs) are concerned, the policing mechanism is supplied with the usage parameters via the network management system. For switched virtual circuits (SVCs), the contract is negotiated and parameters passed via UNI signaling messages and the call set-up procedure to the policing mechanism.

Note that the policing mechanism treats each virtual circuit on an ATM link individually and may police many thousands of virtual circuits at an interface. This sounds like an overwhelming job but it is worth remembering that, regardless of the number of virtual circuits, only one cell crosses the interface at any given time and that the more virtual circuits there are, the lower the average bandwidth of each. Most policing mechanisms use context switching methods based upon the VPI-VCI of the current arriving cell to realize a “virtual” policer.

Leaky Buckets

The policing mechanism implements one or more instantiations of the “generic cell rate algorithm” (GCRA) which behaves in a similar way to a leaky bucket. Consider the leaky bucket diagram in Figure 1. If a drip of water were to enter the bucket every time a cell crossed the UNI (arrived at the policer), and if water were to drip from the hole in the bucket at a constant rate equivalent to the peak cell rate (PCR), then the bucket will not fill up, provided the drips caused by the arriving cells do not, on average, occur faster than the drips leaking from the bucket. Clearly, a burst of drips at a rate above the PCR would cause the bucket to fill, but if that burst is matched by a period when the drips occur at a lower rate than the PCR, the bucket would tend to empty. If the burst of drips caused by the arrival of cells occurs at a rate above the PCR for too long, the bucket will overflow. On the other hand, if the bucket empties completely (because cells stop arriving), the drips from the hole in the bucket cease, of course. The depth of the bucket is equivalent to the cell delay variation tolerance (CDVT).

The simple circuit on the right in Figure 1 is equivalent to the leaky bucket. The up-down counter is incremented by the arrival of a cell and decremented by the tick from a clock running at the peak cell rate. The counter is not allowed to underflow, ticks from the clock being stopped when the counter is at zero. A comparator compares the counter value to the CDVT value stored in a register; if the counter exceeds the value in this register, the comparator signals that the cell which caused this condition is “non-conforming”, a situation equivalent to the bucket in our analogy overflowing.

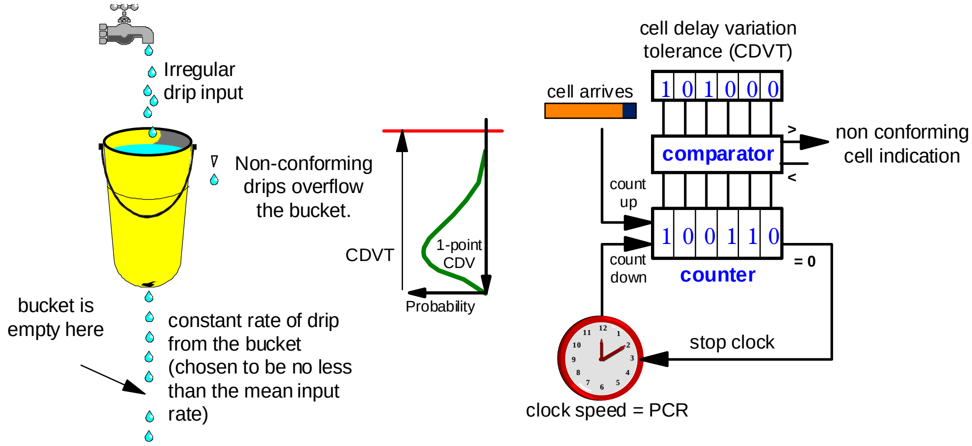
Irregular drip input
Non-conforming drips overflow the bucket.
bucket is empty here
constant rate of drip from the bucket (chosen to be no less than the mean input rate)
CDVT
1-point CDV
Probability
cell delay variation tolerance (CDVT)
cell arrives
1 0 1 0 0 0
comparator
>
<
non conforming cell indication
count up
count down
1 0 0 1 1 0
counter
= 0
stop clock
clock speed = PCR

Figure 1: Single leaky bucket.

This single leaky bucket would be the minimum required policing mechanism that would have to be implemented (effectively for each virtual connection) at the UNI. Cells which are declared to be non-conforming to the PCR/CDVT leaky bucket (GCRA) are *always* discarded at the interface. A single leaky bucket implementation is adequate for policing constant bit rate traffic, and could be used, even, for bursty traffic. However, it would be more bandwidth efficient in this case to implement a dual leaky bucket mechanism (as shown in Figure 2) to take advantage of the SCR and MBS traffic parameters discussed above.

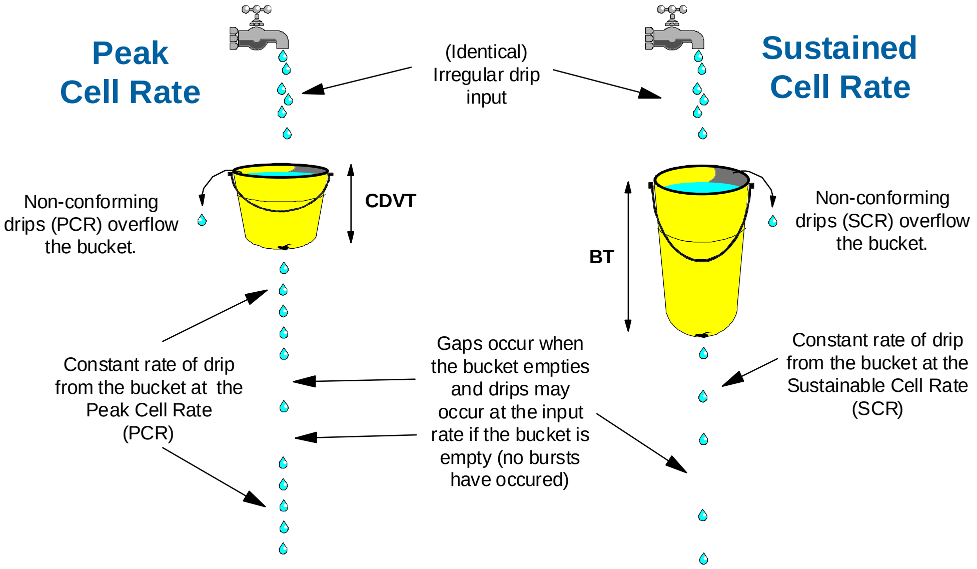
Peak Cell Rate
(Identical) Irregular drip input
Sustained Cell Rate
Non-conforming drips (PCR) overflow the bucket.
CDVT
BT
Non-conforming drips (SCR) overflow the bucket.
Constant rate of drip from the bucket at the Peak Cell Rate (PCR)
Gaps occur when the bucket empties and drips may occur at the input rate if the bucket is empty (no bursts have occured)
Constant rate of drip from the bucket at the Sustainable Cell Rate (SCR)

Figure 2: Dual leaky bucket.

In this case, the second leaky bucket receives identical drips to the first bucket. However, this time the hole is smaller so that the drips occur at the lower sustainable cell rate (SCR). The depth of the second bucket is equivalent to the “burst tolerance” (BT, also called the “intrinsic burst tolerance” (IBT) by the ITU-T) plus the CDVT. Figure 3 may help in understanding what is going on.

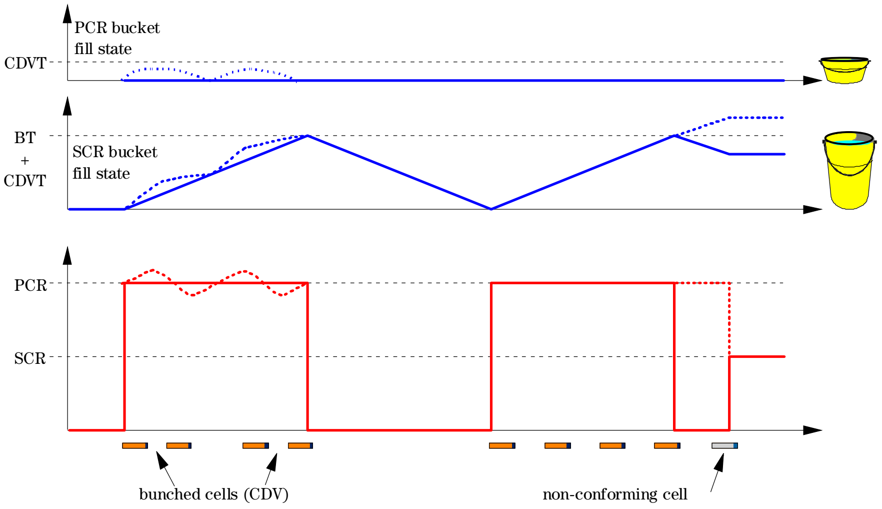

Figure 3: Dual leaky bucket fill patterns.

As it is the maximum burst size (MBS) that is normally provided in signaling messages, the burst tolerance value has to be derived from the PCR, SCR and MBS, according to the formula, where τ_{SCR} is the burst tolerance:

$$\tau_{SCR} = \left\lceil (\mathrm{MBS} - 1)\,(T_{SCR} - T_{PCR}) \right\rceil \text{ seconds}$$

where $\lceil x \rceil$ stands for the first value above x out of the “generic list” of values and where T_{SCR} and T_{PCR} are the time intervals 1/SCR and 1/PCR respectively. The generic list is a set of values that can be used for either the (intrinsic) burst tolerance τ_{SCR} or the cell delay variation tolerance τ_{PCR} according to the following formula (*SCR* can be substituted for *PCR*):

$$\tau_{PCR} = 2^{e_{PCR} - 32} \cdot 2^{9} \cdot \left(1 + \frac{w_{PCR} \cdot 2^{5}}{2^{10}} \right) \text{seconds}$$

$$0 \leq e_{PCR} \leq 31$$

$$0 \leq w_{PCR} \leq 31$$

where e_{PCR} and w_{PCR} are integer values.

Just as the values for CDVT and BT are standardized, so too are the values for PCR and SCR. Another generic list of values can be calculated for cell rates between 1 cell/s to over 4 Gcells/s using the following formula:

$$\Lambda_{PCR} = 2^{m_{PCR}} \cdot \left(1 + \frac{k_{PCR}}{512} \right) \text{ cells per second}$$

$$0 \leq m_{PCR} \leq 31$$

$$0 \leq k_{PCR} \leq 511$$

where Λ_{PCR} is the peak cell rate value and m_{PCR} and k_{PCR} are integer values.

The above formulas makes it clear that only certain discrete values can be used for the cell delay variation and burst tolerances and for the peak and sustainable cell rates. This standardizes policing.

Conformance to the Traffic Contract

We have seen how policing can be done but what, in practice are the options for policing? The ATM Forum and the ITU-T have defined sets of options for policing for different kinds of traffic and for the different capabilities of different networks.

It is important to distinguish between the shape of traffic offered to the network and what the network is prepared to tolerate; these may be essentially the same or they may be different. For example, if the user provides bursty traffic which has “perfect” bursts (equally spaced cells at the PCR during the burst then no cells between bursts), the contract is easy to specify and police. Furthermore, the user can get the maximum out of the contract as there is a good match between the parameters supplied and the shape of the traffic. Variable traffic need not have the above “ideal” shape, however, but traffic parameters still have to be supplied even if the traffic is continuously variable. This can give rise to difficulty in specifying ideal parameters and will often result in inefficient use of the network, as more resources may have to be reserved to ensure a sufficiently good QoS.

Constant Bit Rate Service

The simplest kind of traffic is what the ATM Forum calls “constant bit rate” (CBR) traffic and what the ITU-T calls “deterministic bit rate” (DBR) traffic. As CBR/DBR traffic results in a constant cell rate, only one GCRA or leaky bucket process is needed to police it based upon the PCR and the CDVT. All cells in the virtual connection are policed equivalently, so no distinction is made for the cell loss priority. Consequently, the following is used for policing this service:

$\text{CBR.1} = \text{DBR} = \text{GCRA}(1/\text{PCR}_{0+1}, \text{CDVT}_{0+1})$

Note 1: conventionally “0” is used to indicate cells of low loss priority, i.e. high priority and “1” is used to indicate cells of high loss priority, i.e. low priority and “0+1” is used to indicate cells of both priorities, i.e. all cells in the virtual connection.

Note 2: conventionally the GCRA expression takes parameters of time; 1/PCR is the minimum “emission interval”, the minimum interval between cells occurring at the peak cell rate.

The action that the policng mechanism takes if non-conformance to the above GCRA definition occurs is to delete the cell that caused this. That is, this cell is said to be “non-conforming”. Any cell can potentially be non-conforming, only the past history of cell behavior determines which. Once a cell has been declared non-conforming and consequently deleted, the GCRA behaves for the next cell as if the deleted cell had not existed which means that the next cell is almost bound to conform. If measurements of cell loss were to be made across the network, policing action of this kind would be indistinguishable from cells lost due to congestion in the network. However, because the user traffic violated the traffic contract, the guarantees made by the network as part of the contract do not have to be honored by the network. This gives rise to situations where the user blames the network for QoS when the user is, in fact, to blame.

Other Bit Rate Services

Bursty traffic can be treated by the network in several ways. If the user negotiates a “variable bit rate” (VBR) contract with the network, the network guarantees resources for the level of QoS requested.

Alternatively, the user could choose a less expensive option by requesting the ATM Forum’s “unspecified bit rate” (UBR) service in which the network makes no guarantee about the successful delivery of cells with this service but makes its “best effort”. UBR has, therefore, a lower priority than CBR and VBR services. Note that there is no such thing as “UBR traffic”, there is only a “UBR service” – the traffic would look the same whether the network was supplying a UBR or a VBR service. A UBR service would be attractive in a lightly loaded network because it would be tariffed at a lower rate. As the network became busier, the delivered QoS available from a UBR service would get worse until, perhaps, it would become unusable. The user would then be faced with transferring to a VBR service contract which would guarantee QoS at some extra cost.

A fourth type of service is “available bit rate” (ABR). This type of service makes guarantees about the minimum cell rate (MCR) and it also delivers a level of quality of service in terms of cell loss but not in terms of CDV. It operates by making use of spare capacity in the network but it differs from UBR because of its QoS guarantee. It can make this QoS guarantee because it has the ability to control the source traffic through an ATM layer flow control feedback process using special “resource management” (RM) cells. ABR is not discussed in any more detail here – it is a subject all its own!

Recently the ATM Forum completed version 4.1 of the Traffic Management specification (af-tm-0121-000) which introduced a fifth service type, the “Guaranteed Frame Rate” (GFR) service type. The purpose of the GFR service type is to support non-real-time applications which are frame based; in this instance, a frame is an AAL-5 CPCS (common part convergence sub-layer) PDU. AAL-5 CPCS PDUs are delineated at the ATM layer by identifying the last cell of the PDU (UUI bit = 1 in the PTI). If a cell is discarded within an AAL-5 PDU, the frame is destroyed, so the goal of GFR is to be aware of frames and to discard complete frames as part of the policing process. This is really bringing “early frame discard” into the traffic contract procedures. GFR guarantees a minimum cell rate (MCR), where MCR may be zero. A maximum frame size (MFR) is also specified, along with the MBS (maximum burst size). In the other service types we came across “conforming” and “non-conforming” cells; with GFR we talk about “conforming” and “non-conforming” frames, where policing action deletes complete frames. This paper will not go further into GFR – the reader is directed to the ATM Forum’s specification mentioned above.

Variable Bit Rate Service

The ATM Forum has defined “variable bit rate” service – the ITU-T has the same service which it calls “statistical bit rate” service. Three contract modes are defined for VBR/SBR to suite the capability of the network. In all cases, two GCRAs (leaky buckets) are used and in all cases the first GCRA polices the PCR/CDVT for all cells (high and low priority) of the virtual channel, discarding any cells which do not conform. The differences lie in how the policing mechanism handles the sustainable cell rate (SCR).

The first and simplest contract is called “VBR.1” by The ATM Forum – the ITU-T calls this “SBR1”. In this case, the second GCRA works on all cells of the virtual channel and discards any which do not conform. Tagging (changing cell loss priority from “0” to “1”) is not allowed:

VBR.1 = SBR1 = GCRA(1/PCR_{0+1}, CDVT_{0+1}), GCRA(1/SCR_{0+1}, BT_{0+1}), no tagging

The second option is called VBR.2 (SBR2). In this case, the second GCRA polices high priority cells only (these cells must, of course, have passed the peak cell rate GCRA). Low priority cells are not considered. Tagging, again, is not an option:

VBR.2 = SBR2 = GCRA(1/PCR_{0+1}, CDVT_{0+1}), GCRA(1/SCR_{0}, BT_{0}), no tagging

The third option is like the second option except that tagging is allowed. What this means is that, instead of discarding non-conforming cells, the policing mechanism changes the priority of the non-conforming (high priority) cells to low priority (CLPà1). These cells then become candidates for deletion at congested nodes in the network. Most of these cells probably, on balance, make it through the network successfully:

VBR.3 = SBR3 = GCRA(1/PCR$_{0+1}$, CDVT$_{0+1}$), GCRA(1/SCR$_{0+1}$, BT$_{0+1}$), tagging supported

Clearly this option is the most sophisticated but not all networks can provide it.

Unspecified Bit Rate Service
This service is specified only by the ATM Forum – the ITU-T does not specify it. Two variants exist (UBR.1 and UBR.2) and they differ only upon whether or not tagging is supported. Note that where tagging is supported, unlike the VBR services, the network can tag any cell, whether or not it is conforming. The ATM Forum specification is vague about which circumstances tagging would be done.

UBR.1 = GCRA(1/PCR$_{0+1}$, CDVT$_{0+1}$), no tagging
UBR.2 = GCRA(1/PCR$_{0+1}$, CDVT$_{0+1}$), tagging supported

Policing for the available bit rate (ABR) service is not discussed here.

To summarize:

Conformance Definition		PCR flow	SCR flow	Tagging option active	CLR on
ATMF	ITU-T				
CBR.1	DBR	0 + 1	n/s	n/a	0 + 1
VBR.1	SBR1	0 + 1	0 + 1	n/a	0 + 1
VBR.2	SBR2	0 + 1	0	No	0
VBR.3	SBR3	0 + 1	0	Yes	0
UBR.1	-	0 + 1	n/s	n/a	U
UBR.2	-	0 + 1	n/s	Yes	U

Network QoS Guarantees

Up until now we have seen how the network protects itself against violations of the contract. However, the network has its side of the bargain to deliver – QoS. Each of the above service types makes its own QoS guarantees (or none).

Before we discuss QoS guarantees, it is necessary to get into more detail about the ATM Forum's VBR traffic type. The ATM Forum distinguishes between "non-real time" VBR (nrt-VBR) and "real-time" VBR (rt-VBR). The former is used when delay insensitive data is being serviced and the latter is used for delay sensitive traffic (e.g. packetized video). The policing parameters discussed above are independent of this distinction, which is why discussion of it has been delayed until now. Remember that policing exists to protect the network and has nothing to do with the QoS guarantees that the network may be asked to make.

The CBR/DBR service will require guarantees from the network regarding cell loss, CDV, and maximum CTD. For the same reasons, so will rt-VBR. On the other hand, nrt-VBR/SBR will not require guarantees regarding CDV but will require cell loss to be minimized. With UBR service, the network makes no guarantees about either cell loss or CDV. With ABR, cell loss is controlled but there is no way that CDV guarantees could be met because of the flow control mechanisms. The following table shows the traffic parameters and the associated QoS parameters, indicating which have to be specified for the different categories of service:

	ATM Layer Service Category				
Attribute	**CBR**	**rt-VBR**	**nrt-VBR**	**UBR**	**ABR**
Traffic Parameters:					
PCR and CDVT	specified			specified	specified
SCR, MBS, CDVT	n/a	specified		n/a	
MCR	n/a			n/a	specified
QoS Parameters:					
Peak-to-peak CDV	specified		unspecified		
MaxCTD	specified		unspecified		
CLR	specified			unspecified	may be specfied

Where QoS parameters are specified, worse case values are standardized according to QoS class.

How is ATM QoS Tested?

There are several different aspects to testing what has been discussed above. A user experiencing problems may suspect that the network is not providing the QoS promised. In this case, test equipment can provide the means for establishing the truth. Testing QoS from the user perspective would involve providing test traffic (using test cells, for example) at one side of the network and measuring the QoS in terms of cell loss, cell misinsertion, cell errors, cell delay variation and cell transfer delay at the other side. The network guarantees QoS in terms of cell loss and, in certain cases, CDV. We have discussed how cell loss can be caused not only by congestion in the network but also by the action of policing at the ingress of the network, if the user traffic is violating the traffic contract. It is important to be able to distinguish this so as to establish whether it is the user traffic or the network which is to blame for the poorer than expected QoS.

Another scenario might be that the user or service provider may wish to monitor the conformance of "live" traffic entering the network on a particular virtual channel. The service provider may be able to access information from the policing mechanism via the management information base (MIB) using the integrated local management interface (ILMI) and the network management system (the user may not have access to this). Test equipment that can measure conformance to the traffic contract can therefore be useful in identifying problems ("finger pointing"). Test equipment can also be used to establish how close to the contract user traffic gets – there is no point in paying for resources that are not used since it is often the case that a service provider will charge for reserved capacity rather than actual usage. Reserved capacity for one user can mean the inability of the service provider to offer service to another user.

Desirable capabilities of test equipment for testing QoS can thus be summarized:

- It should be capable of generating test cells for providing measurement of QoS parameters.
- It should be capable of shaping the test traffic in terms of the traffic contract parameters.
- It should be capable of measuring conformance, i.e. it should be able to mimic the policing mechanism and indicate what action the network policer will be taking.

QoS Measurements Using Standardized Test Cells

The ITU-T, in liaison with the ATM Forum, has created Recommendation O.191 which provides the definition for a standardized test cell. This recommendation also specifies the methodology for using the test cell to deliver reliable and repeatable results using test equipment from different vendors when measuring permanent virtual connections. The advantage of this for the user of the test equipment is that different testers should be able to work together to deliver reliable results. Each test equipment vendor has the opportunity to differentiate their product to provide value added benefits – for example, there is a provision in the standard to allow for proprietary extensions to the test cell capability. Figure 4 shows the O.191 test cell:

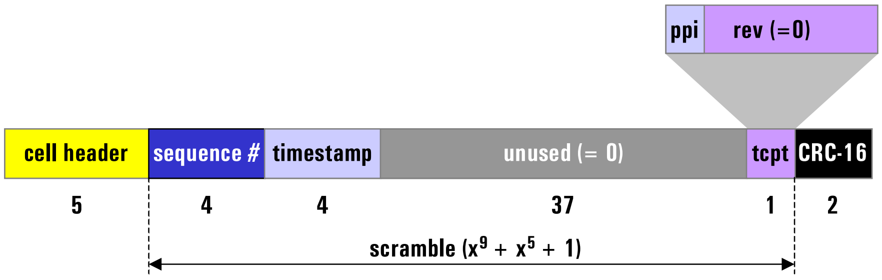

tcpt: test cell payload type
ppi: proprietary payload indicator

Figure 4: O.191 Test Cell

The test cell comprises:

- A standard ATM cell header (any VPI-VCI user value allowed).
- A 32 bit sequence number field, used in the measurement of i.e. VCI > 31 cell loss and cell misinsertion.
- A 32 bit timestamp field, used for end-to-end absolute delay measurements, including 2-point CDV (the timestamp is inserted by the test equipment as the test cell is transmitted and compared with the time of arrival in the test equipment at the other side of the system under test).
- An unused field of 37 bytes, which could be used for proprietary purposes.
- A test cell payload type field, seven bits of which are used as a revision indicator and the most significant bit used as a proprietary payload indicator (for the unused field).
- A 16 bit error detection code field, which uses the same CRC-16 as is used in frame relay HDLC frames to detect cell errors.
- Scrambling applied over most of the payload to ensure that all payload bits are exercised in order to detect stuck bits in memory-based switching fabric. Note that the scrambler is freshly "seeded" for each test cell by the first 9 bits of the cell payload. These bits are part of the sequence number and, in order that the scrambler seed changes every subsequent test cell to maximize randomness,the order of the sequence number field bytes is reversed (the order of bits within each byte is not reversed, though).

With the use of the associated "out-of-service cell transfer outcome measurement algorithm", O.191 plus its "Addendum 1" delivers standardized results for cell loss ratio, cell misinsertion rate, cell error ratio, mean cell transfer delay, 1-point and 2-point cell delay variation, and severely errored cell block ratio.

Delay measurements can present a problem since round trip delay may disguise uni-directional asymmetry. Synchronizing a pair of remote testers in time as well as clock stability is not easy and may require use of GPS (Global Positioning System) techniques. However, in ATM systems cell delay variation is generally more important than the absolute delay, and techniques exist to measure uni-directional 1-point and 2-point CDV. Note: 1-point CDV is delay variation measured at one point and assumes nominally constant bit rate traffic; 2-point CDV measurements normally require the use of a timestamp embedded within a test cell, as with O.191, and observing the change in the difference between arrival and departure times for a sequence of test cells. This does not require GPS techniques. Work has continued in the ITU-T to enhance 0.191 to take account of GPS synchronizing techniques.

Testing QoS and Usage Conformance

Testing quality of service in the presence of policing and testing usage conformance are really two sides of the same coin. In the first case, test traffic (using test cells) is generated such that it conforms to the traffic contract – perhaps only just – and, by processing the received test cells, measurements are made at the far side of the network to determine the QoS. This is an intrusive "out-of-service" type of test which will consume network bandwidth. It would be performed by either the user, who suspects that the network is not delivering QoS to the traffic contract, or by the service provider, who wishes to check for quality control purposes. In both cases, test traffic should be injected before the policing function. The user has no choice in this, of course, and the same is generally true for the service provider, as the policing is usually done directly on the line card of the switch. Figure 5 shows the general arrangement:

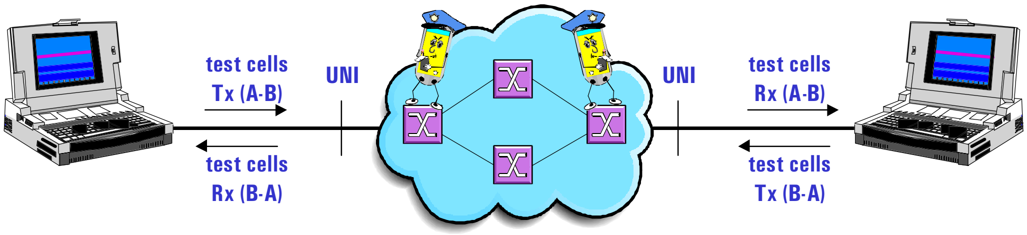

Figure 5: QoS Measurement Setup

Direct (passive) monitoring of user traffic to measure conformance or to determine what usage parameters are appropriate for the traffic contract can also be done using test equipment. Of course, counts of non-conforming cells would normally be collected by the network equipment MIB but this information may not be easily accessed by the user side (the ILMI should allow such access, of course). Another scenario may be more useful – the ability to vary the traffic parameters (PCR, SCR, MBS and CDVT) to see if those originally chosen are appropriate. If, for example, the PCR has had to be chosen to be dramatically higher than the SCR, this is a clear indication that adding a traffic shaper, which would “smooth” the traffic entering the policed network, could significantly reduce the PCR required (at the expense of some added delay and delay variation) and, in many cases, and reduce the cost of the service from the provider. Figure 6 shows the general arrangement:

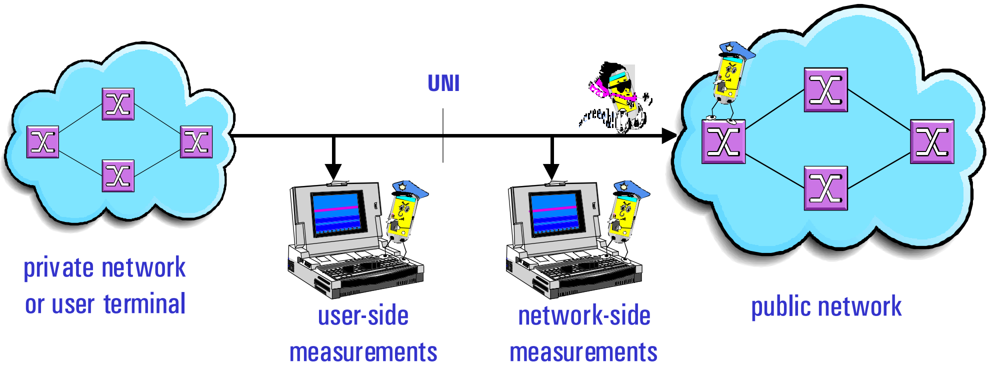

Figure 6: Testing Policing

How Can Agilent Help?

Agilent has a number of ATM testers for different applications (R&D testing, core network testing, manufacturing testing. etc.). One particular tester which can perform all of the tests discussed in this paper is the Internet Advisor ATM, a portable protocol analyzer designed to aid the installation, troubleshooting, and maintenance in the network access area. The Internet Advisor follows the ITU-T O.191 standard for performing QoS measurements of cell loss, cell misinsertion, cell errors, cell delay and delay variation, and severely errored cell blocks. It can shape the test cell traffic using traffic contract usage parameters in order to test network QoS. It can measure the impact of specific policing configurations on live user traffic using the discrete values of usage parameters from the “generic list”. All this is in addition to the other comprehensive protocol analysis capability of this versatile and inexpensive tester for ATM and other technologies. For more details, contact your Agilent representative.

Summary

This paper has discussed issues surrounding quality of service within an ATM environment. We have seen how ATM is able to deliver guaranteed quality of service, when required to do so, and how the network is able to protect itself from abuse. Most service types have been discussed, although discussion of ABR has been restricted. Much time has been spent on the traffic contract and what it means for both the user and the network. Finally, testing the different aspects of QoS, including policing, has been covered.

References

af-tm-0121.000	ATM Forum Traffic Management Specification v 4.1 (03/99)
ITU-T I.356	B-ISDN ATM Layer Cell Transfer Performance (10/96)
ITU-T I.357	B-ISDN Semi-Permanent Connection Availability (08/96)
ITU-T I.371	Traffic Control and Congestion Control in B-ISDN (08/96)
ITU-T O.191	Equipment to Assess ATM Layer Cell Transfer Performance (10/97)

Acronyms

ABR	Available Bit Rate
BT	Burst Tolerance
CBR	Constant Bit Rate (ATM Forum equivalent of DBR)
CDV	Cell Delay Variation
CDVT	Cell Delay Variation Tolerance
CER	Cell Error Ratio
CES	Channel Emulation Service
CLR	Cell Loss Ratio
CMR	Cell Misinsertion Rate
CPCS	Common Part Convergence Sub-layer (of AAL-5)
DBR	Deterministic Bit Rate (ITU-T equivalent of CBR)
GCRA	Generic Cell Rate Algorithm
GFR	Guaranteed Frame Rate
GPS	Global Positioning System
HDLC	High-level Data Link Control
IBT	Intrinsic Burst Tolerance (ITU-T equivalent of BT)
ILMI	Integrated (formally “Interim”) Local Management Interface
ITU-T	International Telecommunications Union – Telecommunications Sector
MBS	Maximum Burst Size
MCR	Minimum Cell Rate
MFS	Maximum Frame Size
MIB	Management Information Base
PCR	Peak Cell Rate
PPI	Proprietary Payload Indicator
PVC	Permanent Virtual Circuit/Connection
QoS	Quality of Service
SCR	Sustainable Bit Rate
SBR	Statistical Bit Rate (ITU-T equivalent of VBR)
SVC	Switched Virtual Circuit/Connection
TCPT	Test Cell Payload Type
UBR	Unspecified Bit Rate
UNI	User-Network Interface
VBR	Variable Bit Rate (ATM Forum equivalent of SBR)
VCI	Virtual Channel Identifier
VPI	Virtual Path Identifier

Contact us with FaxBack

By Returning This FaxBack Page, with the following required information, you can facilitate your initial contact to speak with a Customer Care Representative

Fax to: 1-303-662-2038

OR...

E-mail to: csp_telesales@agilent.com

Visit our web site

www.agilent.com/comms/onenetworks

First Name____________________

Last Name ____________________

Company Name ____________________

Job Title ____________________

Business Mailing Address ____________________

____

City ____________________

State/Province ____________________

Country____________________

Zip Code ____________________

E-Mail Address ____________________

____

Phone Number ____________________
(incl. area/country code)

Fax Number____________________

Do you have a budget set for this application?

□ Yes
□ No
□ In process

What is your time frame to implement this product?

□ 30 days
□ 90 days
□ 180 days
□ Other – (please define)____________________

____

____

Product(s) of Interest

□ ***The Agilent Advisor*** – Integrated, High-Performance Troubleshooting for:
____ Advisor LAN
____ Advisor WAN
____ Advisor ATM

□ ***The LAN Analyzer*** – Scaleable Ethernet and Token Ring Test Solutions

□ ***Telegra Fax Test*** – Fax Protocol and Low Generation Analysis

□ ***Telegra Voice Quality Tester*** – Detailed Voice Analysis for Clarity, Echo and Delay using PSQM and PAMS

□ ***Telegra Voice and Fax over IP*** – Protocol Analysis

□ ***FASTest*** – Automated Service Verification for PSTN and IP Networks

What is the main problem you need to solve on your network?

____

____

____

Agilent Technologies
Innovating the HP Way

Notes

Notes

Notes

Connect with us!
http://www.agilent.com/comms/onenetworks

This Product is Y2K Compliant

About The Author

Gregan Crawford is currently a product marketing manager with Agilent Technologies's Network Systems Test Division where he has responsibilities for the strategic direction of broadband test products for service providers and ATM Forum activities.

Following a period of research at Edinburgh University's Artificial Intelligence department, Gregan joined Hewlett-Packard's Queensferry Telecommunications Division in 1976 as an engineer in the research and development department. From 1987 to 1992 he managed Hewlett-Packard's participation in the Parasol consortium of the European Communities' RACE program (RACE: Research in Advanced Communications for Europe) when he first started working in the field of ATM.

At the start of 1993 Gregan became the principal representative for HP's Communications Measurement Division at the ATM Forum. In late '93 he was invited to become the first chairman of the newly formed Test Working Group of the Worldwide Technical Committee of the ATM Forum and he continued in this rôle until quite recently.

Gregan speaks regularly at conferences and seminars and last summer completed a seminar tour of the US and Canada during which he delivered papers on high speed ATM and network management OAM functions. He also participated in the ATM Year 97 conference in San Jose during which he led a session on ATM testing and presented at the ATM Year 98 (Europe) conference in London on the subject of ATM quality of service.

By internet, phone or fax, get assistance with all your Test and Measurement needs.

Online assistance:
http://www.agilent.com/find/assist

United States:
(Tel) 1 800 452 4844

Canada:
(Tel) 1 877 894 4414
(Fax) (905) 206 4120

Europe:
(Tel) (31 20) 547 2323
(Fax) (31 20) 547 2390

Japan:
(Tel) (81) 426 56 7832
(Fax) (81) 426 56 7840

Latin America:
(Tel) (305) 269 7500
(Fax) (305) 269 7599

Australia:
(Tel) 1-800 629 485
(Fax) (61 3) 9272 0749

New Zealand:
(Tel) 0 800 738 378
(Fax) (64 4) 495 8950

Asia Pacific:
(Tel) (852) 3197 7777
(Fax) (852) 2506 9284

Product specifications and descriptions in this document subject to change without notice.

Printed in U.S.A. 10/00

5968-8556E

Agilent Technologies
Innovating the HP Way