On ATM Support for Distributed Real-Time Applications

Chia Shen

June 10, 1996

http://www.merl.com/

Tutorial I

IEEE Real-Time Technology and Applications Symposium

Contents

This is NOT a tutorial on everything about ATM.

- ATM is much more than just traffic management and control.
- ATM Forum Technical Committee Working Groups:

  TM, SIG, SAA, MPOA, PNNI, PHY, B-ICI, RB, LANE, SEC, NM, TEST.

What is ATM?

Cell-based switching, and asynchronous multiplexing.
(Best of Packet Switching + TDM)

An ATM cell = 53 bytes. 5 bytes cell header, 48 bytes payload.

Why 48?
48 = [32(the European choice) + 64(the American choice)]/2

A typical ATM LAN consists of workstations/PC with ATM NICs connected to ATM switches.

For two applications to communicate via an ATM network, a virtual channel must be established.

History of ATM

In 1983, the first ideas on ATM and related techniques were published in 1983 by two research centers (CNET and AT&T Bell Labs).

In 1984, the Research Center of Alcatel Bell in Antwerp and other places started actual work on ATM.

In 1985, J. Turner published "Design of an integrated services packet network".

In 1987, then CCITT (now renamed ITU) selected ATM as the transfer mode for future BISDN.

History of ATM (Cont.)

Also in 1987, the early papers and experiments on multiplexing, integrating voice and data, fast cell switching, etc. in ATM were published.

In 1990, with all the early study, analysis and experiments, CCITT issued a series of Recommendations specifying the details of ATM for use in BISDN.

In 1991, ATM Forum was formed with four members. Today 750 members.

I. Why ATM?

A typical configuration of current industrial control systems

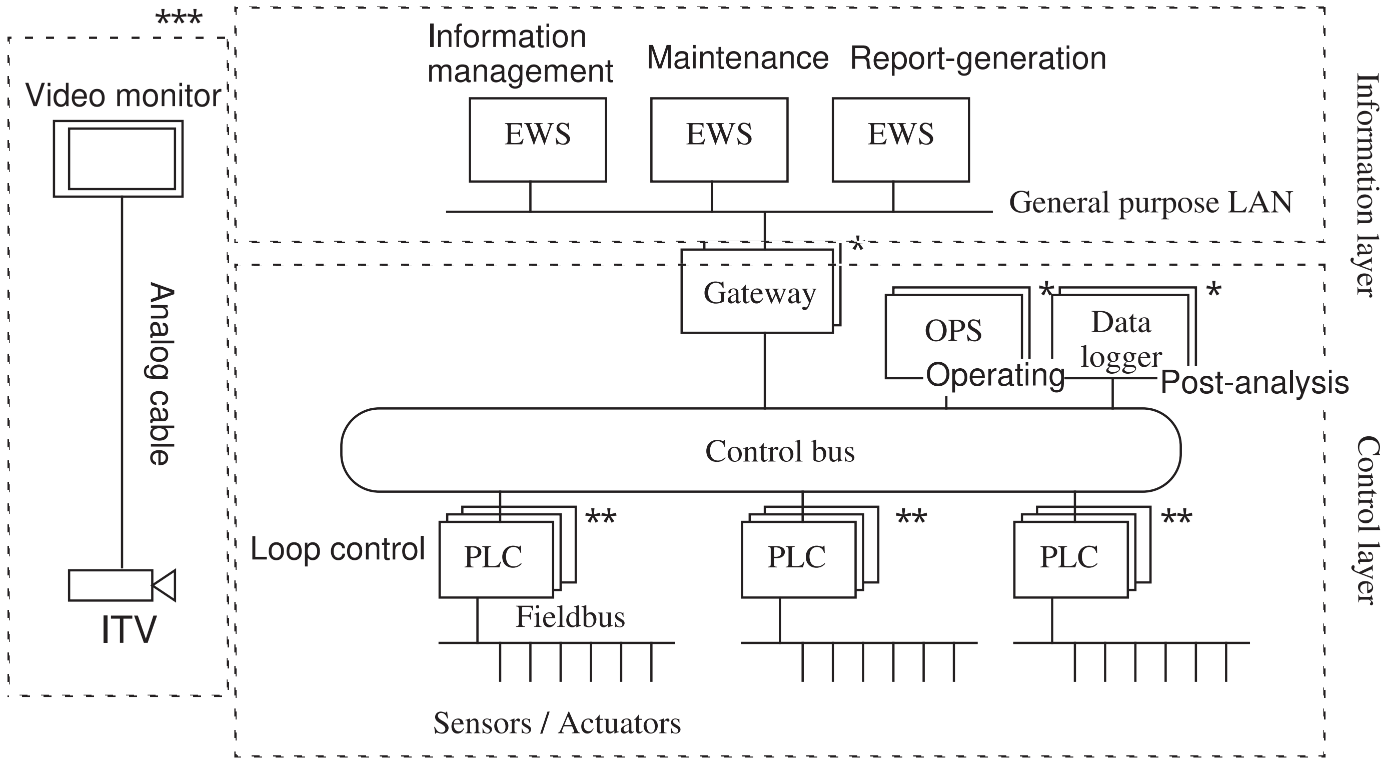

Problems: Scalability, Proprietary, Fixed QoS, Integration.

Difficult to Grow and Evolve

- Example — Static Cyclic Communication in current systems:
  - 4 Kbytes per 0.5msec = 64Mbps
  - ATM cell overhead (5 bytes/53 bytes) = 9.4%
  - A total bandwidth of $>$ 70 Mbps.

Applications do not need to see all 4 Kbytes per 0.5msec. But difficult to support real-time request-response in the layered, non-uniform network technology.

4 Key (promised) Advantages of ATM

(1) **High bandwidth** — 25Mbps, 155Mbps, 622Mbps...

(2) **Connection oriented** — Each virtual connection is guaranteed a level of QoS.

(3) **End-to-end QoS** (Quality of Service) — Resources are reserved along the entire path.

(4) **Scalability** — Flexible topology, distance and bandwidth.

No other types of network technology, e.g., Ethernet, FDDI, can give us all of these 4 key advantages.

How Can ATM LAN Change Future Distributed Industrial Control Applications?

(1) **Integrated** traffic (control data, video, audio, image, files).

- New functionality.

(2) Real-time built upon **open** network and system platform.

- Lower cost.

(3) **Virtual**, not physical isolation among connections.

- One network technology, Native ATM Services, plug and play.

(4) **Interactivity** possible, more demand driven.

- Client-server distributed real-time computing.

Vision: A New System Architecture with ATM

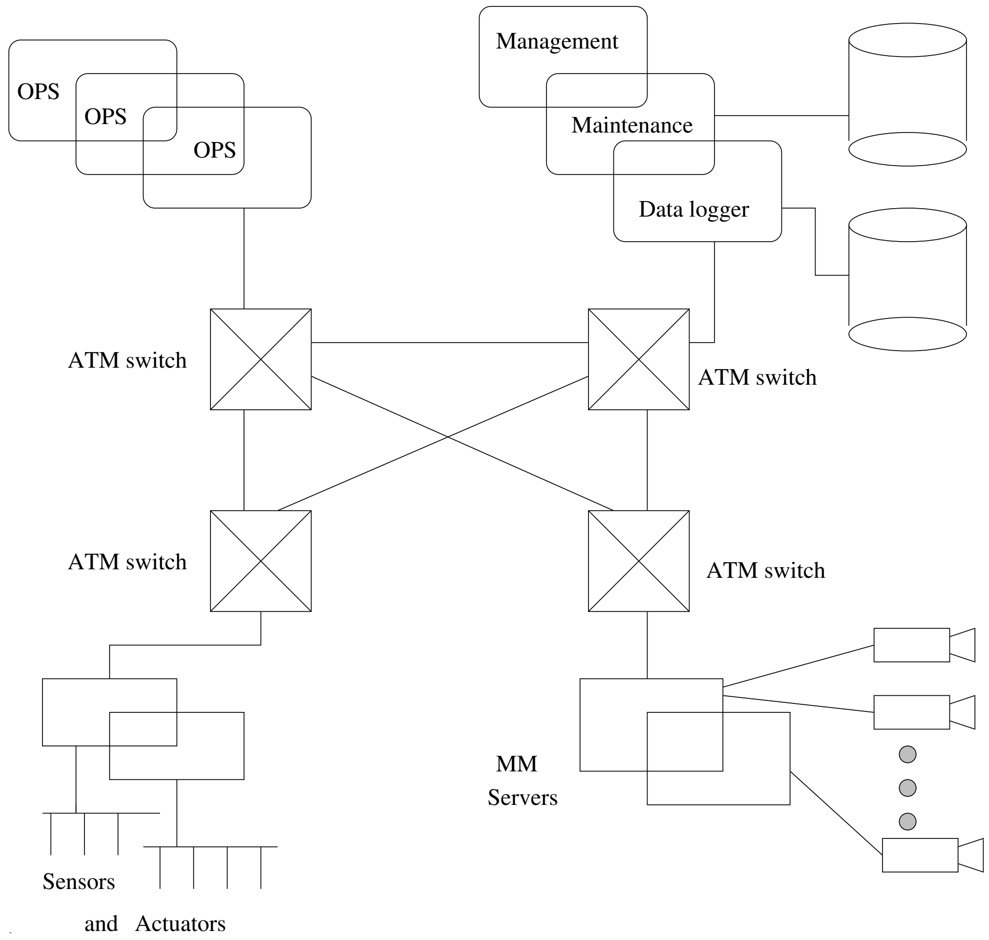

Scalable, Open Standard Based, Flexible QoS, Easy Integration.

How Can ATM Support Such A Vision?

We need to understand —

(1) The design space provided by ATM for real-time applications.

(2) How are data handled in an actual ATM network?

(3) What are the technical challenges?

(4) Limitations.

II. Design Space in ATM

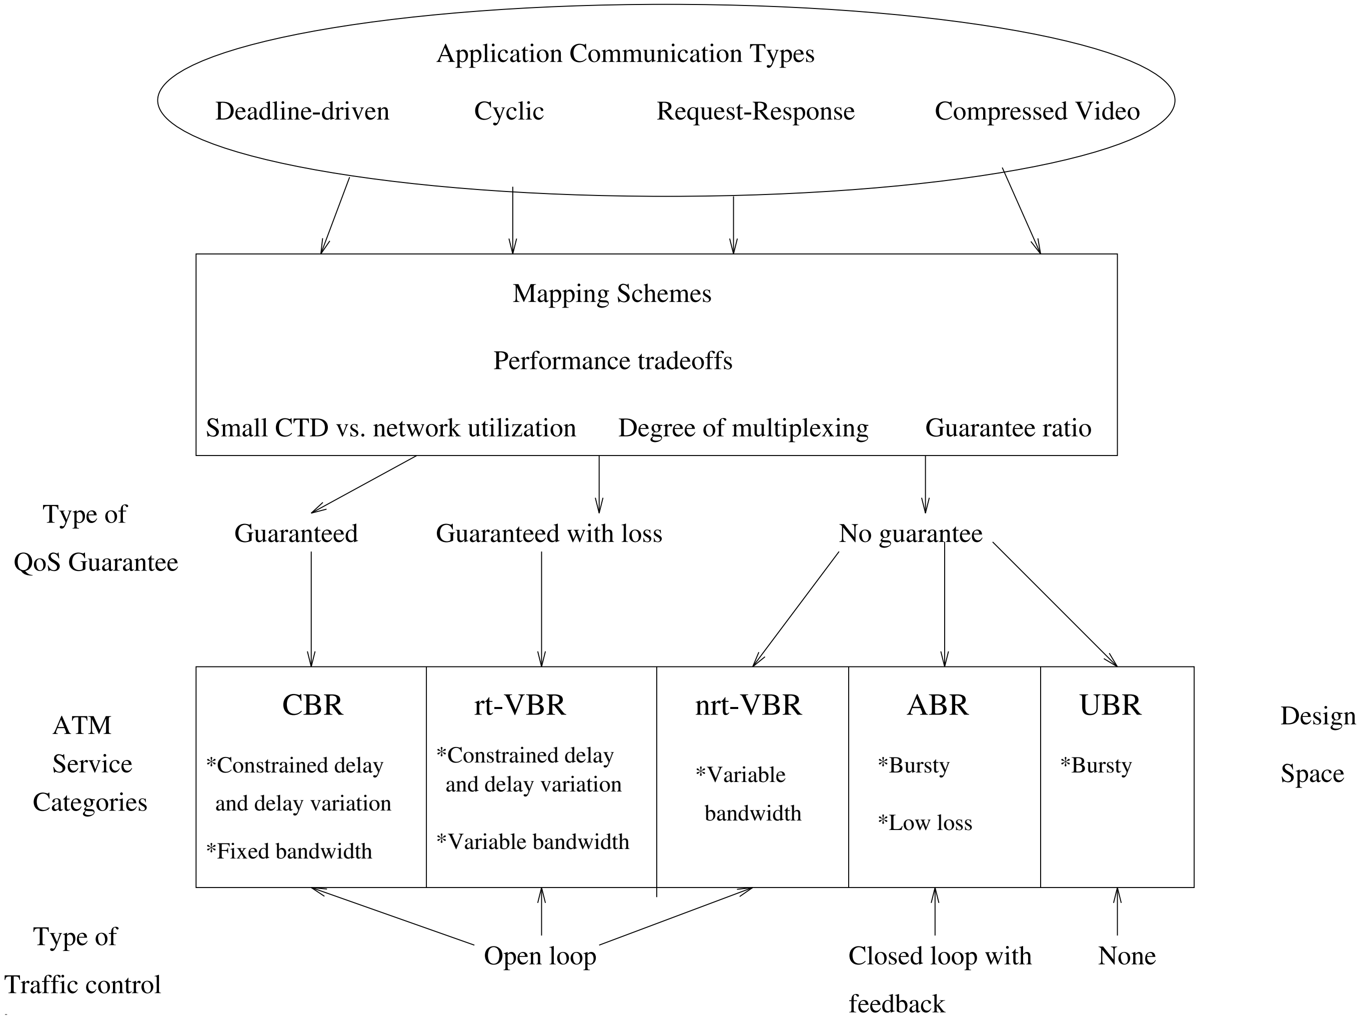

Open vs. Closed Loop Control

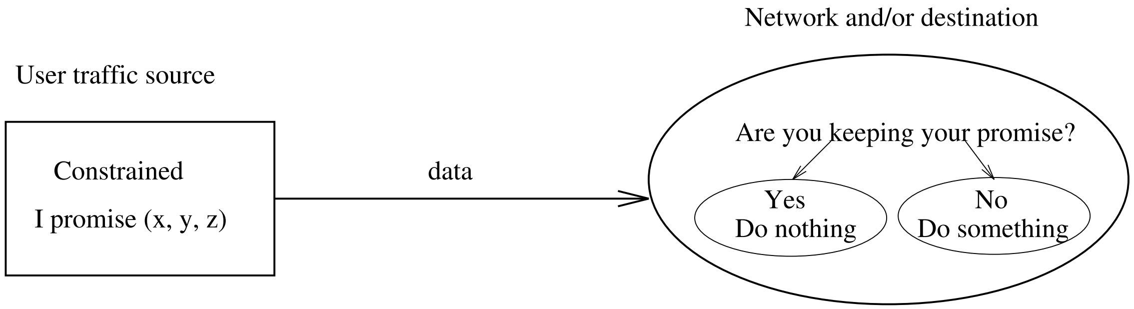

Open Loop Control

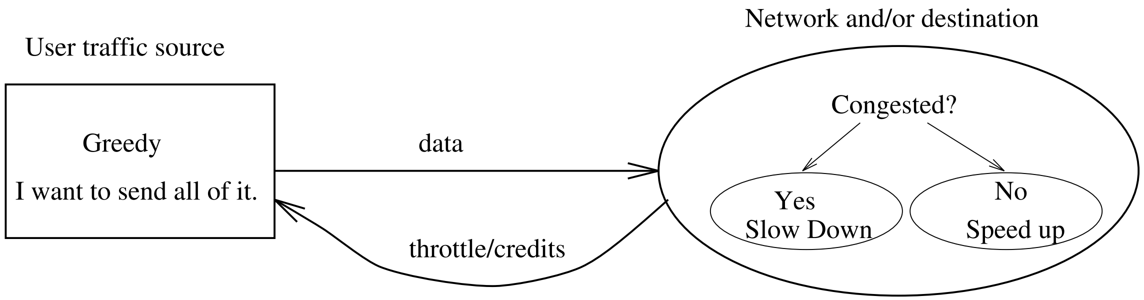

Closed Loop Control

Open vs. Closed Loop Control

Open loop control == Congestion avoidance via scheduling.

Open-loop control is particularly suited for real-time communication:

- Map real-time traffic onto one of the ATM service categories.
- The network guarantees delay bounds and delay variation bounds.

Open-loop traffic control is used for CBR, rt-VBR and nrt-VBR service categories in ATM.

CBR, rt-VBR, and nrt-VBR

(ATM Forum TM 4.0)

CBR — A static amount of bandwidth, and tightly constrained delay and delay variations (jitter) required.

rt-VBR — Bursty traffic. Tightly constrained delay variations (jitter) required. rt-VBR is for real-time applications which can benefit from statistical multiplexing — more efficient use of the network resources.

nrt-VBR — Bursty traffic. No delay bounds are associated with this service category.

CBR and rt-VBR are designed for real-time traffic.

ATM Open Loop Control for Real-Time QoS Guarantee

The key — The network commits to support the QoS for all connections that are compliant!

How is this commitment implemented?

Traffic Contract, Shaping, Policing and Scheduling

Traffic Contract and Shaping

Contract:

– The user: This is my traffic characteristic and my QoS requirement.

– The network: Connection admission control (CAC).

– Need traffic models, and conformance definitions.

Shaping:

Make sure my traffic passes the policer's conformance test.

Traffic Policing and Scheduling

Policing:

– Usage Parameter Control at UNI. Test for compliance.

– Need observable traffic streams — simple deterministic traffic models.

– Policing in real-time at very high data rate (e.g., 622 Mbps).

Scheduling:

– Whose turn? Which cell to send?

– Decisions depend on QoS parameters.

The Network vs. The User

Computer networks (LAN) are owned by the same organization.

ATM was originally designed by telecommunication

- Separation of the network provider and the user.

The network provider always wants to protect the network and its performance that the provider has carefully engineered.

This leads to the concept of traffic contract, user compliance, policing actions .... (We never heard of these in networks designed only for computer LAN, e.g., Ethernet, FDDI ...)

III. A Day in the (simplified) Life of an ATM Cell

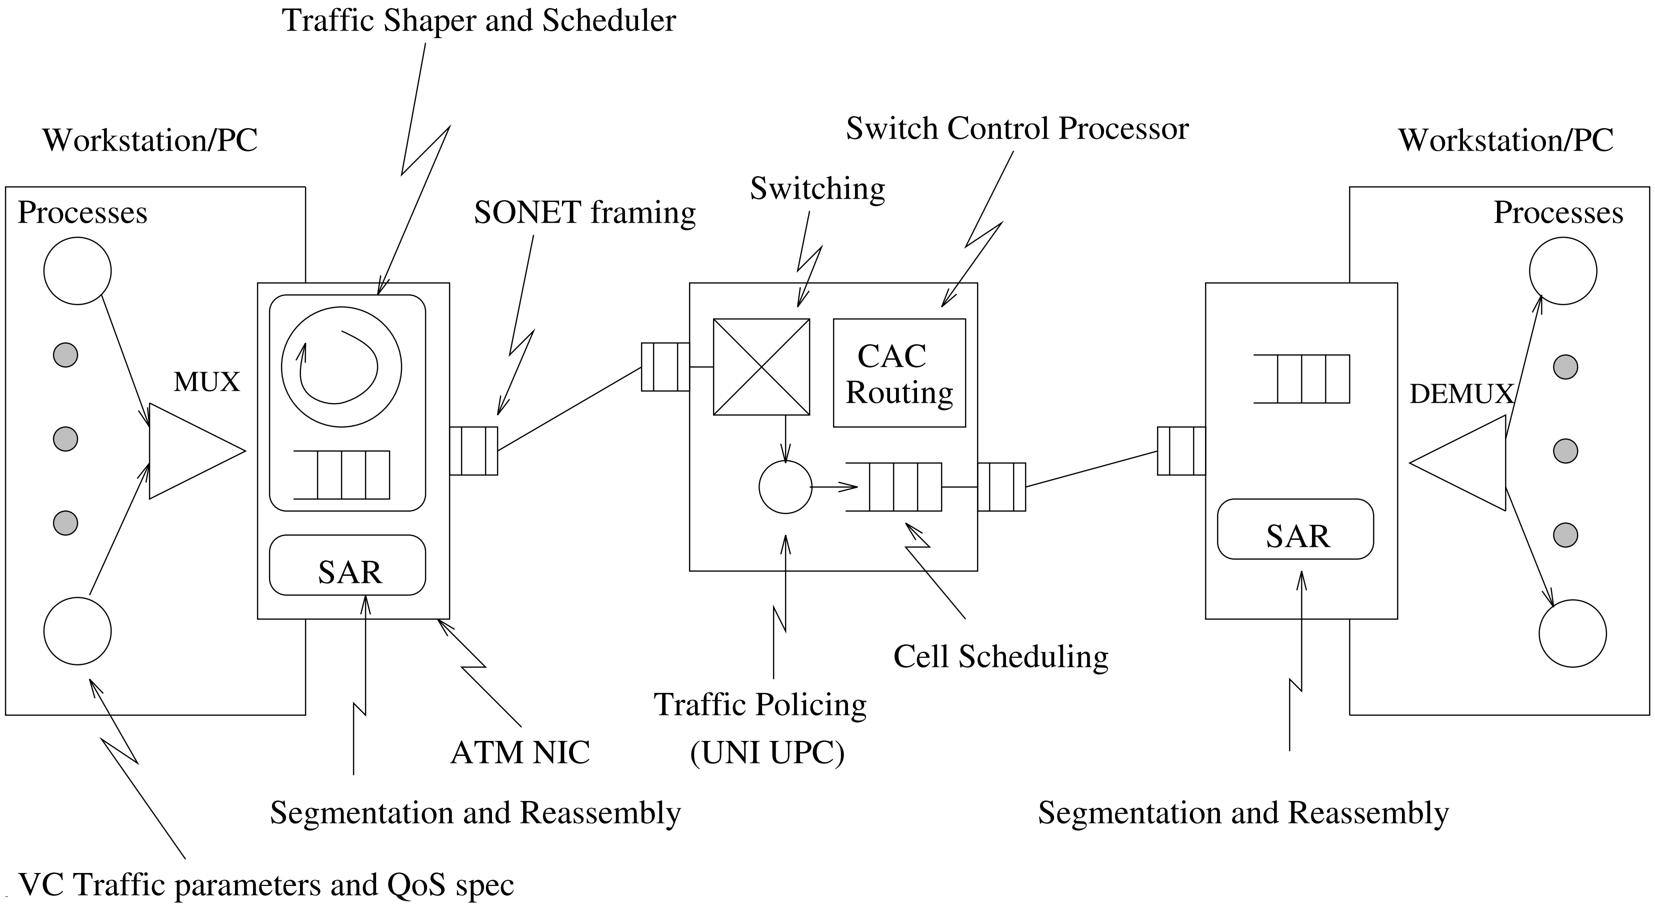

5 ATM Service Categories

Attribute	ATM Layer Service Category				
	CBR	rt-VBR	nrt-VBR	UBR	ABR
Traffic Parameters:					
PCR and $CDVT$	$\surd$	$\surd$	$\surd$	$\surd$	$\surd$
SCR, MBS and $CDVT$	n/a	$\surd$	$\surd$	n/a	n/a
MCR	n/a	n/a	n/a	n/a	$\surd$
QoS Parameters:					
peak-to-peak CDV	$\surd$	$\surd$	X	X	X
maxCTD	$\surd$	$\surd$	X	X	X
CLR	$\surd$	$\surd$	$\surd$	X	X

$\surd$ = specified

X = unspecified

n/a = not applicable

Definitions

of Traffic and QoS Parameters Used by CBR and VBR

Parameter	Description
PCR	Peak Cell Rate: The cell rate that a source can never exceed.
SCR	Sustainable Cell Rate: The long term average cell rate.
MBS	Maximum Burst Size.
CDV	End-to-end Cell Delay Variation induced by buffering and cell scheduling.
CTD	End-to-end Cell Transfer Delay from the exit (first bit out) of a source to the entry (last bit in) of a destination.

Traffic Descriptors vs QoS Parameters (ATM Forum TM 4.0)

Traffic Descriptor: PCR, SCR, MBS.

- A user uses these parameters to describe the source traffic.
- This is a deterministic traffic model.

QoS Parameters: maxCTD, peak-to-peak CDV, CLR.

- A user uses these parameters to specify the delay and loss requirements.

Cell Loss Ratio? I don't want to lose any cells!

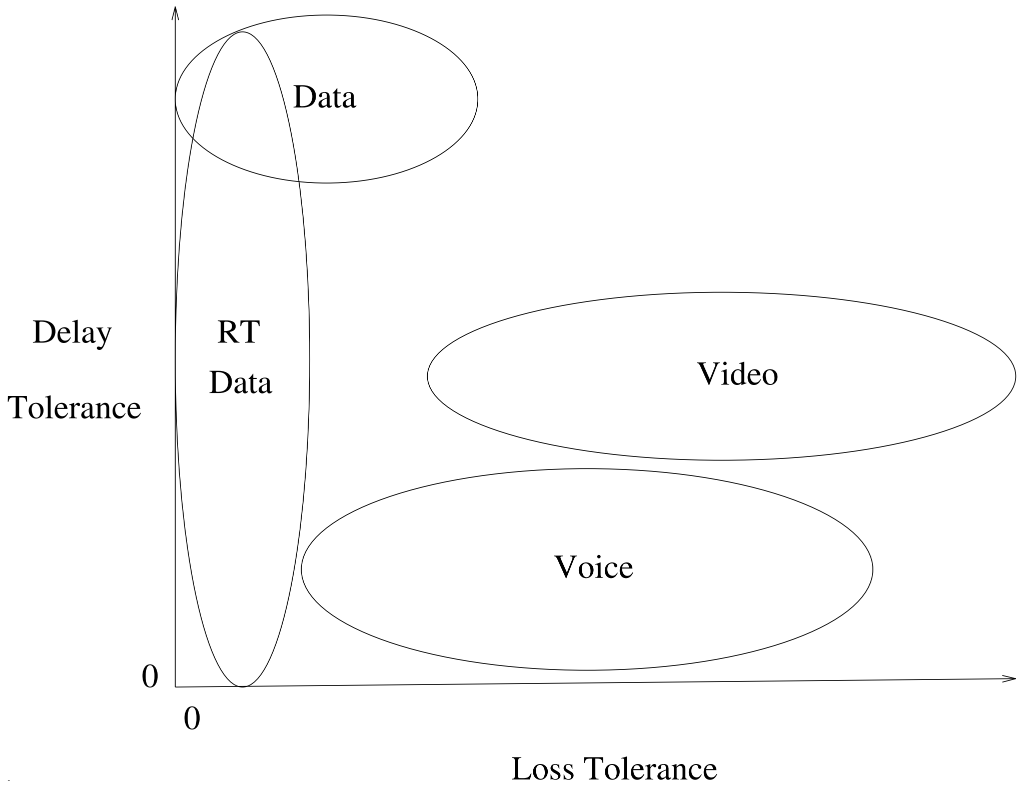

CLR is not very well defined. It is network specific.

The network offers a value of CLR over the lifetime of the connection.

QoS Parameter Accumulation

maxCTD and peak-to-peak CDV are **end-to-end** deley parameters.

MaxCTD and peak-to-Peak CDV are determined by

- each switch along the path accumulating its own values for these parameters during connection set up.

The worst case accumulation currently adopted by the ATM Forum Signaling group is

$$maxCTD = F + CDV$$

F = Fixed Delay due to propagation, fixed transmission and switching delay.

Definition of maxCTD and peak-to-peak CDV

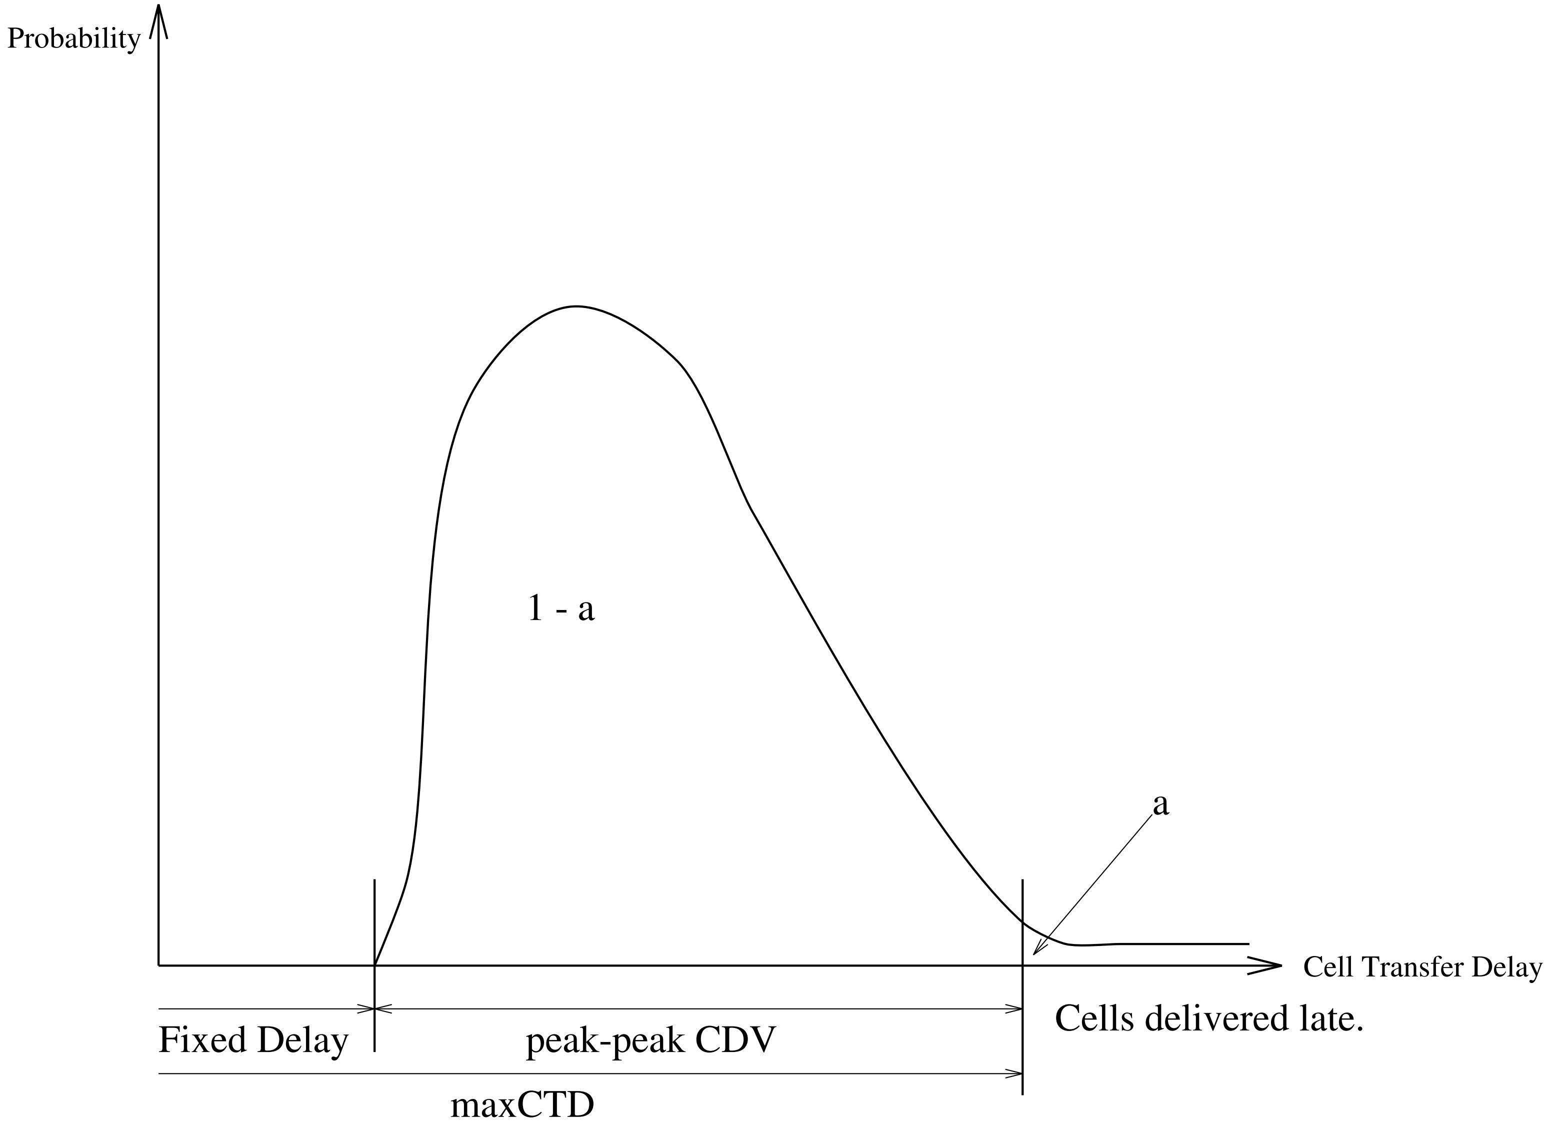

CBR and rt-VBR Traffic Specification

	CBR	rt-VBR
Traffic descriptor	PCR	PCR, SCR, MBS
QoS parameter	CDV, CTD	CDV, CTD
Conformance def	GCRA(T, CDVT)	GCRA(T, CDVT) GCRA(T_s, CDVT+τ)

$T = 1/\text{PCR}$, $T_s = 1/\text{SCR}$.

CDVT — Cell Delay Variation Tolerance. CDVT is NOT the CDV QoS parameter.

GCRA — Generic Cell Rate Algorithm (e.g., Continuous-state Leaky Bucket).

GCRA can be used (1) for cell scheduling, and (2) for UPC.

Understanding Traffic Contract and Conformance

Case 1: CBR

CBR is monitored at the UNI by one Leaky Bucket:

GCRA(1/PCR,CDVT).

Traffic of a CBR connection coming out of a NIC should almost be GCRA(1/PCR,0). It may suffer delay variation of up to CDVT inside the network.

Maximum number of conforming back-to-back cells =

$$\lfloor 1 + \frac{CDVT}{1/PCR - ct} \rfloor$$

ct = cell time at the link speed.

Understanding Traffic Contract and Conformance

Case 2: rt-VBR

For simplicity, assume CDVT = 0.

VBR is monitored at the UNI by two Leaky Buckets GCRA(1/PCR,0) and GCRA(1/SCR,τ) — An arriving cell must pass the tests by both leaky buckets to be conforming.

τ is the maximum earliness of a conforming cell.

$$\tau = (MBS - 1) * (1/SCR - 1/PCR)$$

Understanding Traffic Contract and Conformance

Case 2: rt-VBR (cont.)

Let N be the maximum number of cells generated by a conforming rt-VBR source during any time interval t, then $N \leq f(t)$ and

$$f(t) = MBS + \lfloor SCR * t \rfloor$$

$f(t)$ is called a traffic constraint function.

The minimum time interval I between maximum size conforming bursts of cells is

$$I = MBS * (1/SCR - 1/PCR)$$

What does the Leaky Bucket really mean?

(I've seen many misconceptions ...)

It does NOT mean a constant rate.

Cells do NOT leak out of the bucket!!

It does NOT mean a unique conforming traffic stream.

It is a cell conformance checking algorithm and used for traffic policing at the UNI.

A Continuous State Leaky Bucket for CBR

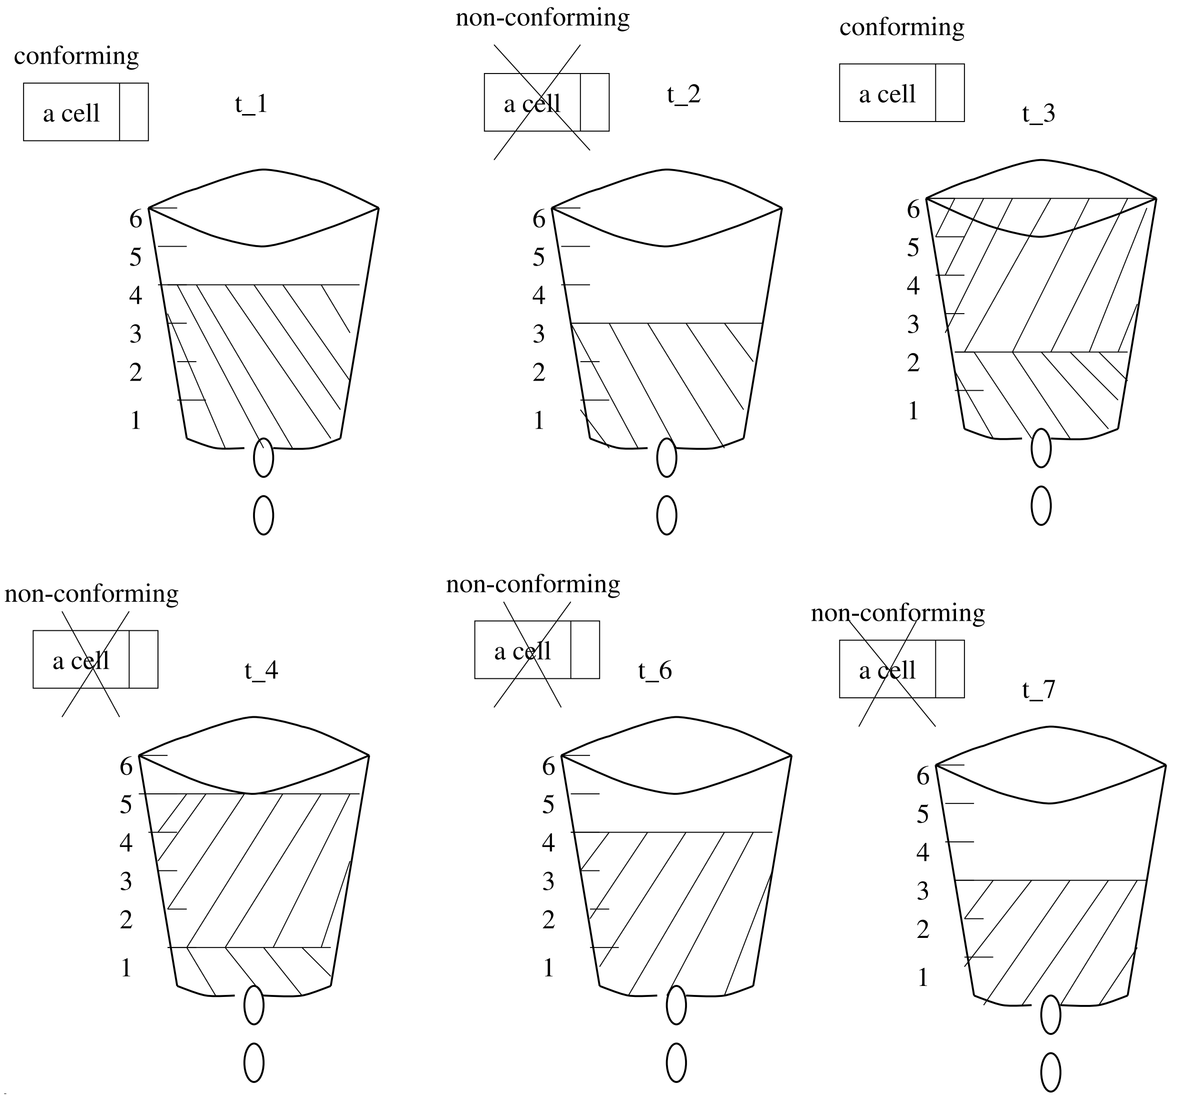

1/PCR = 4

bucket size = 6

drop size = 1

Every conforming cell
addes 4 units of water
into the bucket.

Observations of Leaky Buckets

CDVT does not give one a cell rate higher than PCR — the PCR is still the same in the long run.

A non-conforming cell has no effect on the state of the leaky bucket.

The policer is only concerned with cells that have arrived too early — *Late* cells will see empty buckets.

If a VBR source transmits a burst at a rate higher than SCR, then it must decrease its rate to less than SCR for some time before the next burst — to allow the SCR bucket to drain a bit.

Leaky Buckets

Can be combined in sequence for policing.

Leaky Buckets		
Bucket	Bucket Size	Drop Size
PCR	1/PCR + CDVT	1/PCR
SCR	τ + CDVT	1/SCR

What does the UNI do?

Quote from a recent paper: (**Misconception**)

"At the network entrance, the UNI regulates the incoming traffic specified by the (σ_i, ρ_i) model. The UNI regulates the maximum cell transmission rate into the network entrance below ρ_i by ...".

UNI does NOT buffer or smooth or rate control the traffic.

It only monitors and may mark cells.

The switch will switch a cell immediately!

Cells will only be buffered inside the network for scheduling/queueing in current ATM switches.

CDVT for PCR and SCR vs. CDV vs. MBS

CDVT allows 'back-to-back' (i.e., at link speed) cell clumping
— used by the policer.

CDVT is induced by
– the slotted nature of ATM,
– the physical layer overhead, e.g., SONET framing, and
– the ATM layer functions, e.g., cell multiplexing within the NIC.

CDVT is specified by the network
– can NOT be signalled/negotiated.

CDVT for PCR and SCR vs. CDV vs. MBS (cont.)

CDV is a user QoS parameter
– bounds the cell spacing deviation from the ideal inter-cell spacing. – E.g., The receiver/destination of a video stream needs this quantitative information to reassemble the frames.

CDV is induced by
– buffering and cell scheduling.

MBS is a user traffic descriptor
– allows MBS number of cells to be transmitted at the rate of PCR.

Jitter Is A 'Fact-of-Life' in A Network

Jitter == Variations in cell delay.

E.g., SONET framing at the physical layer:
– SONET is *synchronous* bit wise, ATM is *asynchronous* cell wise.

End-to-end delays:
– AAL: PDUs (Protocol Data Units). AAL5 PDU payload may be 1 to $2^{16} - 1 = 65,535$ bytes.
– NIC: SAR, scheduling.
– Physical layer: TC fifo, SONET framing.
– Switch: queueing, SONET framing, HEC.

Jitter Caused by SONET Framing

(OC3 155Mbps)

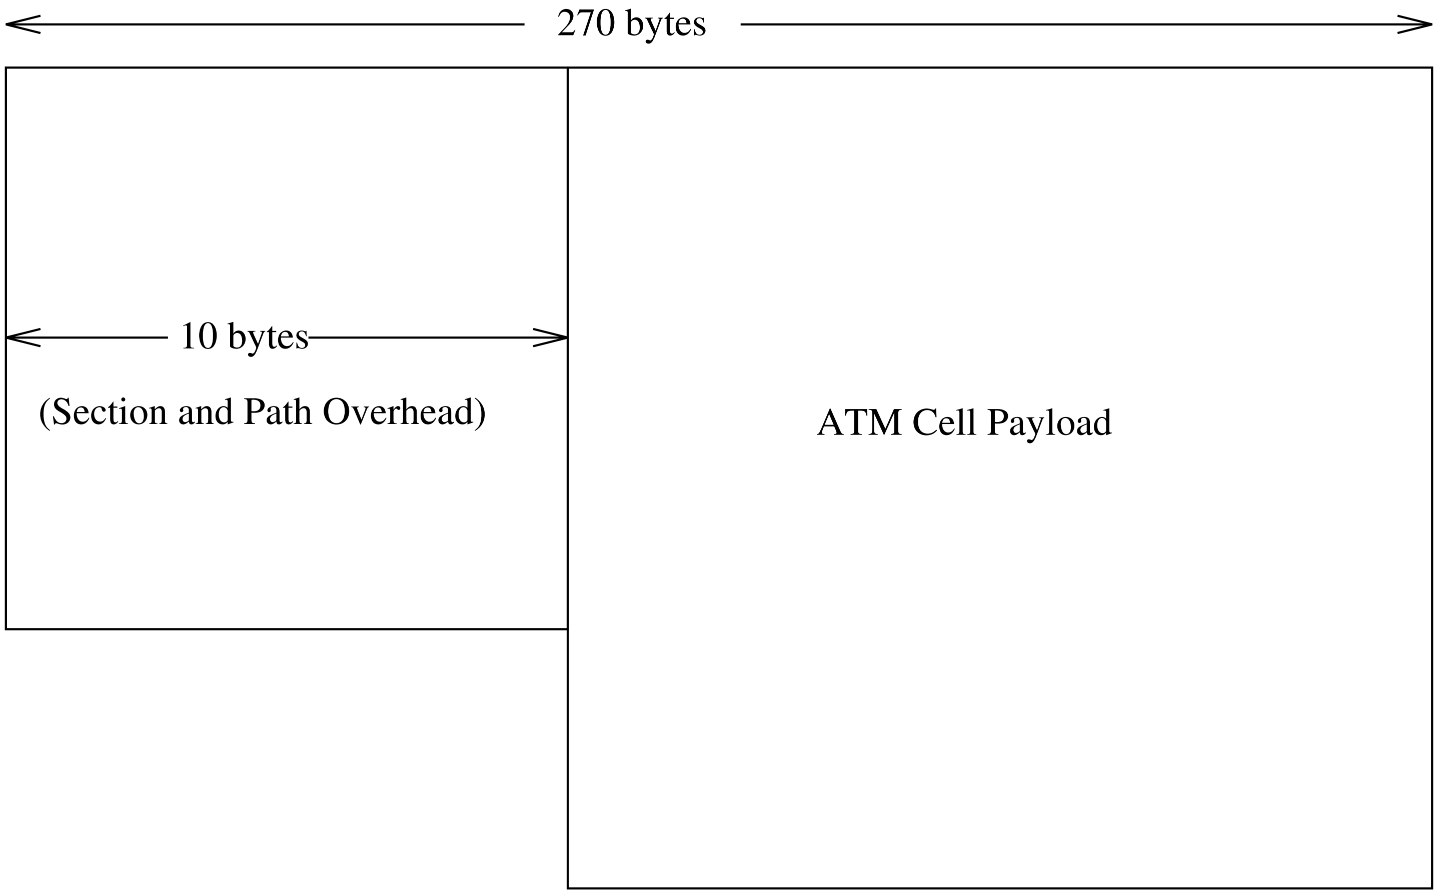

*** A Much Simplified SONET OC3 FRAME.

Jitter Caused by AAL5

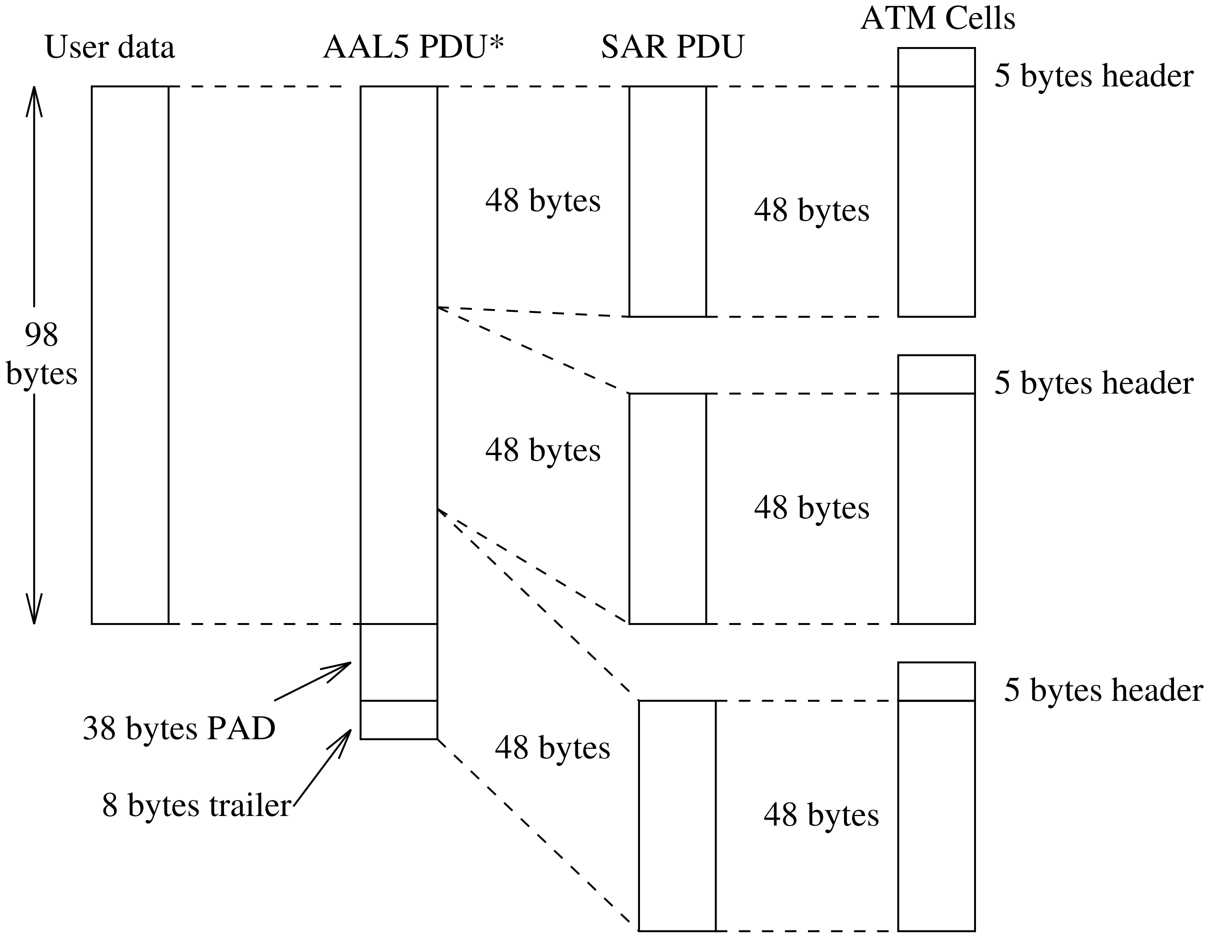

*PDU = Protocol Data Unit

Jitter Is A 'Fact-of-Life' in A Network

It's all a matter of granularity of time:

– If it's within a cell time — who cares.

– If it's at most X cell times — an upper bounded jitter.

– Otherwise — we are in trouble.

IV. Mapping RT Communication onto ATM Service Categories

Issues to be considered:

– Use of CBR: Possible schedulability loss.

– Use of VBR: To guarantee the same amount of real time traffic, we need to use extra network bandwidth.

– Use of Virtual Path, instead of Virtual Channels:
Do our own scheduling within a VPC.

— User Level Multiplexing.

Related and dependent on the CAC and cell scheduling algorithms used in the network.

Real-Time Periodic Communication is not necessarily CBR

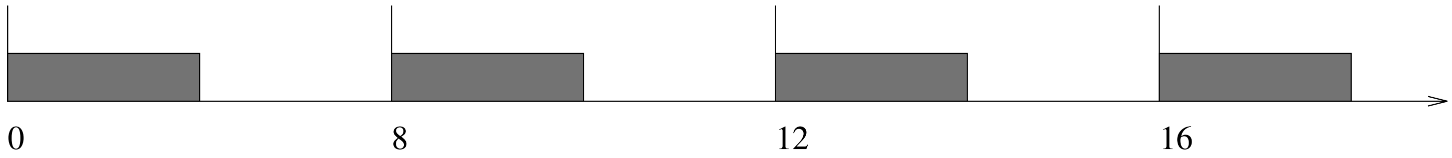

A real-time periodic channel

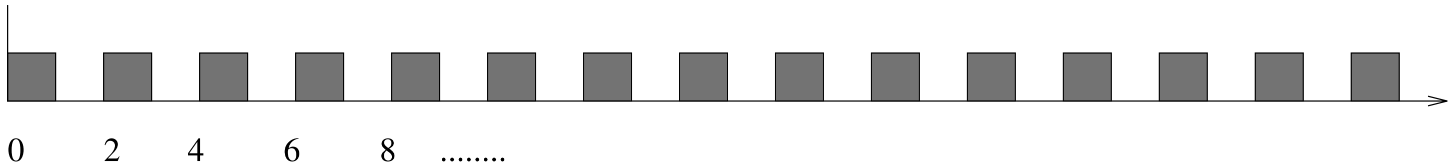

A CBR channel with equal bandwidth reservation

Issues to be considered:

* CTD
* Deadline < period

Can We Specify a Real-Time Channel That Conforms to rt-VBR?

A real-time bursty traffic source that wishes to emit bursts at peak rate (PCR) periodically.

We must calculate traffic descriptor values such that the NIC will transmit conforming traffic, i.e.,

— Our traffic must pass the dual Leaky Bucket GCRA(1/PCR, 0) and GCRA(1/SCR, τ) test.

Can We Specify a Real-Time Channel That Conforms to rt-VBR? (cont.)

If T_I is the minimum message (cell burst) interarrival time, and MBS is the message size in cells, and we must choose SCR and τ at least large enough such that

MBS $= 1 + \lfloor min(T_I - 1/SCR, \tau)/(1/SCR - 1/PCR) \rfloor$

then the traffic will be conforming to GCRA(1/PCR, 0) and GCRA(1/SCR, τ).

NOTE: The traffic patterns conforming to GCRA(1/PCR, 0) and GCRA(1/SCR, τ) is not unique.

Example: Mapping Real-Time Periodic Communication onto rt-VBR

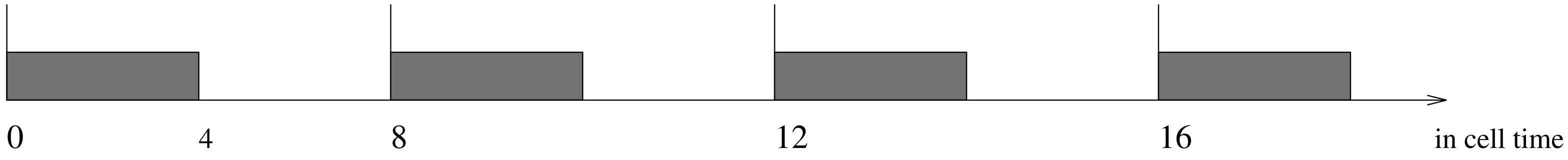

A real-time periodic channel

To map to a rt-VBR channel:

PCR = link speed

SCR = 0.5 link speed

MBS = 4

CTD = end-to-end propagation delay

tau = (MBS-1)(1/SCR - 1/PCR)

Burst (message) interarrival time = MBS*(1/SCR - 1/PCR) + 1/PCR

Question:

* What is the impact on network utilization? (Depending on switch scheduling algorithm)

Cell Scheduling and CAC

Given

– traffic descriptors,
– QoS parameters and
– conformance definitions

of a connection,

the network must allocate resources

– buffer space
– cell transmission rate/priority
– delay bounds

CAC depends on cell scheduling algorithms.

Cell Scheduling Disciplines

	Delay bounds	Firewall Property	Implementation
FCFS	1	No	No special support
SP	Fixed	Per priority	Per priority queueing
DP	Arbitrary	Per connection	Per cell priority Per cell sorting

FCFS = First Come First Serve

SP = Static Priority

DP = Dynamic Priority

Not All DP Algorithms Are The Same

Rate-based vs EDF-based

Rate-based:

– E.g., Virtual Clock, PGPS ...
– Individual rate per connection.
– Coupling of reserved rates and guaranteed delay bounds.
– The priority value only depends on the arrival time of the previous cell.

EDF-based:

– Individual rate and deadline per connection.
– Decoupling of reserved rates and guaranteed delay bounds.
– The priority value depends on the position of a cell within a message.

Observation1:

Delay of a Connection Does Not Imply Total Buffer Requirement

E.g., EDF can cause large delay of some channels, but does not necessarily increase the requirement of the total buffer space in a switch.

Observation 2:
Rate Control Minimizes The Total Amount of Buffers Inside the Network

If B is the worst case traffic upper bound within time I, with rate control, the total buffer needed is B, while without rate control, we need n*B, where n is the number of switches.

Observation 3:

SP + Per-VC Queueing + Many Bits Representing Priorities $\neq$ DP

It is more than just the number of bits. It needs a per cell priority calculation.

V. Limitations of ATM

Rate-based cell scheduling inherently couples delay bounds and bandwidth reservation.

— Leads to inefficient usage of network bandwidth for real-time connections that requires little bandwidth but also a very tight delay bound.

— Use User Level Multiplexing, or use Virtual Path to bundle resource allocation might be the solution.

V. Limitations of ATM (cont.)

ATM is connection oriented

— Thus EVERY connection must be admitted!

— VCs are independent, switches cannot relate them if they are set up separately.

— There are applications that can efficiently **share** bandwidth among a set of known VCs.

— Solutions may need to change how current signaling is defined.

ATM Is Still In Its Infancy
Technical Challenge 1

ATM network architecture that **efficiently** provides the end-to-end QoS guarantee fo rt-VBR:

(1) Distributed CAC algorithms.

(2) Switch cell scheduling and resource/buffer allocation.

Technical Challenge 2

Middleware that closes the gap between applications and ATM services.

– End system is still the major bottleneck in both performance and functionality.

– Programming models and abstractions that effectively utilize ATM.

– Efficient real-time transport protocols on top of ATM Native Services.

– QoS mapping algorithms and User Level Multiplexing schemes.

The middleware must be **configurable** and **selectable**.

R&D Experiemental System Examples
Mercuri ATM Testbed – Honeywell

Multimedia and real-time control applications.

Experiments with first generation ATM products, e.g., FORE ASX100.

— Throughput, delay, jitter ...

— Reported major bottleneck in the end system support.

R&D Experiemental System Examples
MidART — Mitsubishi Electric Corp

An on-going R&D experimental system — **Mi**ddleware and network **A**rchitecture for **R**eal-**T**ime industrial applications.

Second generation ATM products are used
– E.g., FORE ASX200BX switch. 3 service levels per output port, each with 127 priorities, giving a total of 381 service priorities per output port.

Incorporating multimedia information into distributed industrial plant control applications.

Focusing on a client-server programming model and QoS mapping strategies.

Real-time LAN vs Internet WAN

In LAN, propagation delay is negligible, end-station latency in comparison is large.

In WAN, coast-to-coast propagation delay is large, in order of 30 – 50 msecs, fibre is a buffer.

Related Research in IETF

Two IETF (Internet Engineering Task Force) subworking groups

(1) Integrated Services (intserv)

(2) Resource Reservation Setup Protocol (RSVP)

Working on providing real-time services, QoS guarantees, to internet applications.

IETF RSVP and InterServ

http://www.ietf.cnri.reston.va.us/home.html

RSVP protocol is the internet signaling for IP flow resource reservation.
InterServ specifies services models as the Traffic Management in ATM Forum does.

Integrated service models:

- Guaranteed Quality of Service
- Controlled-Load Network Element Service
- Predictive Quality of Service
- Controlled Delay Quality of Service

RSVP and InterServ for IP

– Receiver-based, connectionless.

– Possibly many receivers.

– Receiver requests QoS and provides traffic parameters.

– May modify QoS at will.

– Use IP multicast distribution. Full multicast, bidirectional.

– Soft state, dynamic QoS change.

– The source advertises what it "can" send, e.g., 10 Mbps. A receiver can reserve only 2 Mbps out of the 10 Mbps according to its capability.

Internet must deal with scalability problem in data distribution to a globe-wide population.

ATM

– Sender-based, connection oriented.

– Point-to-multipoint, unidirectional multicast.

– Same QoS on all branches of a multicast tree.

– To support Internet style multicast has been a very difficult issue for ATM.

Remember: IP datagram is variable length. ATM cells are fixed length.

Conclusion

ATM is here.

We must understand the design space to effectively build real-time systems on ATM networks.

There are still R&D challenges.

References

1. V. Bansal, R. Siracusa, J. Hearn, G. Ramamurthy, and D. Raychaudhurt. Adaptive QoS-Based API for ATM Networking. In Proceedings of the 5th International Workshop on Network and Operating System Support for Digital Audio and Video, April 1995.

2. ATM Forum Technical Committee. Traffic Management Specification 4.0, April 1996.

3. R. Cruz. A Calculus for Network Delay, Part I: Network Elements in Isolation and Part II: Network Analysis. IEEE Transactions on Information Theory, 37(1):114–141, January 1991.

4. D. Ferrari and D. Verma. A Scheme for Real-Time Channel Establishment in Wide-Area Networks. IEEE Journal on Selected Areas in Communications, 8(3):368–379, April 1990.

5. P. Goyal, S. Lam, and H. Vin. Determining End-to-End Delay Bounds in Heterogeneous Networks. In Proceedings of the 5th International Workshop on Network and Operating System Support for Digital Audio and Video, April 1995.

6. P. Goyal, H. Vin, C. Shen, and P. Shenoy. A Reliable, Adaptive Network Protocol for Video Transport . to appear in Proceedings of the 15th IEEE INFOCOM, March 1996.

7. M. Grossglauser, S. Keshav, and D. Tse. The Case Against Variable Bite Rate Service. In Proceedings of the 5th International Workshop on Network and Operating System Support for Digital Audio and Video, April 1995.

References

8. E. W. Knightly, D. E. Wrege, J. Liebeherr, and H. Zhang. Fundamental Limits and Tradeoffs of Providing Deterministic Guarantees to VBR Video Traffic . In Proceedings of ACM SIGMETRICS, Ottawa, Canada 1995.

9. S.-K. Kweon and K. G. Shin. Traffic-Controlled Rate-Monotonic Priority Scheduling of ATM Cells. to appear in Proceedings of the 15th IEEE INFOCOM, March 1996.

10. S. Lam and G. Xie. Burst Scheduling Networks: Flow Specification and Performance Guarantees . In Proceedings of the 5th International Workshop on Network and Operating System Support for Digital Audio and Video, April 1995.

11. H. Lauer, A. Ghosh, and C. Shen. A General Purpose Queue Architecture for an ATM Switch. In proceedings of First Annual Conference on Telecommunications R&D in Massachusetts, volume 6, pages 17–22, October 1994.

12. H. Lauer, C. Shen, R. Osborne, J. Howard, Q. Zheng, M. Takegaki, H. Shimakawa, and I. Mizunuma. Digital Audio and Video in Industrial Systems . In Proceedings of the 5th International Workshop on Network and Operating System Support for Digital Audio and Video, April 1995.

13. C. L. Liu and J. W. Layland. Scheduling Algorithms for multiprogramming in a Hard-Real-Time Environment. Journal of ACM, 20(1), 1973.

14. D. E. McDysan and D. L. Spohn. ATM Theory and Application. McGraw-Hill Series on Computer Communications. 1995.

References

15. A. Parekh. A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks. PhD thesis, MIT, 1992.

16. A. Raha, S. Kamat, and W. Zhao. Admission Control for Hard Real-Time Connections in ATM LANs . to appear in Proceedings of the 15th IEEE INFOCOM, March 1996.

17. J. Salehi, Z. Zhang, J. Kurose, and T. Towsley. Supporting Stored Video: Reducing Rate Variability and End-to-End Resource Requirements through Optimal Smoothing . To appear in Proceedings of ACM SIGMETRICS, 1996.

18. J. Turner. New Directions in Communications (or Which Way to the Information Age?). IEEE Communications Magazine, 25(8):8–15, October 1986.

19. G. Xie and S. Lam. Delay Guarantee of Virtual Clock Server. IEEE/ACM Transactions on Networking, 3(6), December 1995.

20. H. Zhang and D. Ferrari. Rate-Controlled Static-Priority Queueing. In Proceedings of IEEE INFOCOM, 1993.

References

21. H. Zhang and D. Ferrari. Improving Utilization for Deterministic Service In Multimedia Communication . In Proceedings of International Conference on Multimedia Computing and Systems, May 1994.

22. H. Zhang and S. Keshav. Comparison of Rate-Based Service disciplines . In Proceedings of ACM SIGCOMM, Sept. 1991.

23. H. Zhang and E. Kightly. A New Approach to Support Delay-Sensitive VBR Video in Packet-Switched Networks . In Proceedings of the 5th International Workshop on Network and Operating System Support for Digital Audio and Video, April 1995.

24. L. Zhang. VirtualClock: A New Traffic Control Algorithm for Packet Switching Networks. In Proceedings of ACM SIGCOMM, pages 19–29, August 1990.

25. Z. Zhang, Z. Liu, J. Kurose, and D. Towsley. Call Admission Control Schemes under the Generalized Processor Sharing Scheduling . To appear in Proceedings of Telecommunication Systems, May 1996.

26 Q. Zheng, K. Shin, and C. Shen. Real-Time Communication in ATM Network. In Proceedings of 19th Annual Local Computer Network Conference, pages 156–165, Minneapolis, Minnesota, October 1994.

Mercuri ATM Testbed References

1. Huang, J. and Pavan, A., "The Mercuri Multimedia Laboratory at Honeywell", to appear in IEEE MultiMedia, Summer, 1996.

2. Tsang, R., Du, D.H.C. and Pavan, A., "High Speed Network Support for Video-Based Distributed Applications", to appear in ACM/Springer Verlog Multimedia Systems Journal, Second/Third Quarter, 1996.

3. Pavan, A., Guha, A., et. al., "Experimental Evaluation of Real-Time Support on the Mercuri Wide Area ATM Testbed", Proceedings of the 20th IEEE Conference on Local Computer Networks, Minneapolis, October, 1995.

4. Guha, A., Pavan, A., et. al., "Controlling the Process with Distributed Multimedia", IEEE MultiMedia, Summer 1995.

5. Guha, A., Pavan, A., et. al., "Supporting Real-Time and Multimedia Applications on the Mercuri Testbed", IEEE Journal on Selected Areas in Communications, Vol. 13, No. 4, May, 1995.

6. Guha, A. and Pavan, A., "Real-Time Support of Continuous and Variable Bit Rate Traffic on an ATM Network", Proceedings of the IEEE International Conference on Distributed Computing Systems (ICDCS), Poznan, Poland, June, 1994.

7. Guha, A., Agrawal, M.B. and Pavan, A., "A Real-Time Network Architecture for Time-Critical Applications", Proceedings of the IEEE Workshop on the Role of Real-Time Multimedia/Interactive Computing, Raleigh-Durham, North Carolina, November, 1993.

8. Guha, A., Agrawal, M.B. and Rundquist, R., "Multimedia Integrated Real-Time Control Systems", Proceedings of the First IEEE Workshop on Real-Time Applications, May, 1993.