11. Traffic management in ATM

Literature

1 W. Stallings (1998)
   - “High-Speed Networks: TCP/IP and ATM design principles”
   - Prentice Hall, New Jersey

2 ATM-Forum, Technical Committee
   - “Traffic Management Specification, Version 4.0”
   - April 1996

3 ITU-T, Study Group 13
   - “Recommendation I.371: Traffic Control and Congestion Control in B-ISDN”
   - July 1995

Contents

Traffic management

- Problems:
  - Traffic is **random** in nature (varying unpredictably)
    - every now and then, **congestion** occurs (unavoidably)
  - Traffic sources may behave “badly”
- **Traffic management** is needed to
  - achieve the required QoS and performance under these circumstances
  - protect the network and other users against badly behaving traffic sources
- Two approaches to “manage” congestion
  - **predictive** methods to **avoid** congestion (before it occurs)
  - **reactive** methods to **alleviate** and **remove** congestion (after it has occured)

Examples on traffic management

- Telephone network (circuit switching):
  - only predictive methods:
    - call admission control
    - resource reservation (bandwidth)
- X.25 (connection oriented packet switching):
  - only predictive methods:
    - call admission control
    - resource reservation (buffers)
    - window-based open-loop flow control
- IP (connectionless packet switching):
  - only reactive methods:
    - window-based closed-loop flow control

RTT * C

- The product of round-trip time (RTT) and transmission rate (C) determines
  - the amount of information transmitted before the first feed-back from the network and the destination (to adjust the transmission rate)
- In broadband wide-area-networks (WAN) this product can be very large $\Rightarrow$ reactive methods alone are insufficient
- Example:
  - Assume that
    - distance between two users is 1500 km
    - transmission rate is C = 100 Mbps
  - The two-way propagation delay will be
    - 2*1500/300,000 s = 0.01 s
  - Thus, the product of RTT and C is at least
    - 0.01*100,000,000 bits = 1,000,000 bits = 1 Mbit

Contents

Acronyms

- ATM — Asynchronous Transfer Mode
- B-ISDN — Broadband Integrated Services Digital Network
- CCITT — International Telephone and Telegram Consultative Committee (1992: CCITT → ITU-T)
- ITU-T — International Telecommunication Union - Telecommunication Standardization Sector

History

- Traditionally: dedicated networks for different services
  - For example: telephone, telex, data, broadcast networks
  - Optimized for the corresponding service
- Need for integration of all services into a single ubiquitous network
  - “One policy, one system, universal service” (T. Vail, AT&T’s first president)
  - Early 80’s: Research on Fast Packet Switching started
- Answer from the “Telecom World” : **B-ISDN**
  - 1985: B-ISDN specification started by Study Group SGXVIII of CCITT
  - 1988: Approval of the first B-ISDN recommendation (I.121) by CCITT
- Chosen implementation method: **ATM**
  - 1990: ATM chosen for the final transfer mode for B-ISDN by CCITT
  - 1991: ATM Forum founded
    - to accelerate the development of ATM standards
    - to take into account the needs of the “Computer World”

Design objectives

- Flexibility
  - integration of different types of traffic (like voice, data and video) into a single network
  - supporting any transmission rate (even variable)
  - combination of best properties of circuit and packet switching
- Guaranteed Quality of Service (QoS)
  - end-to-end delay, delay variation, cell loss ratio
- Scalability
  - WAN/LAN
- Fastness
  - both in switching and transmission $\Rightarrow$ broadband

ATM

- Connection oriented:
  - virtual connections set up end-to-end before information transfer
  - resource reservation possible but not mandatory
- All information carried in short, fixed-length packets (**cells**)
  - along the route chosen for that virtual connection
  - statistical multiplexing at nodes
  - identifier label (local address) at cell header
  - no error detection/recovery for the information field

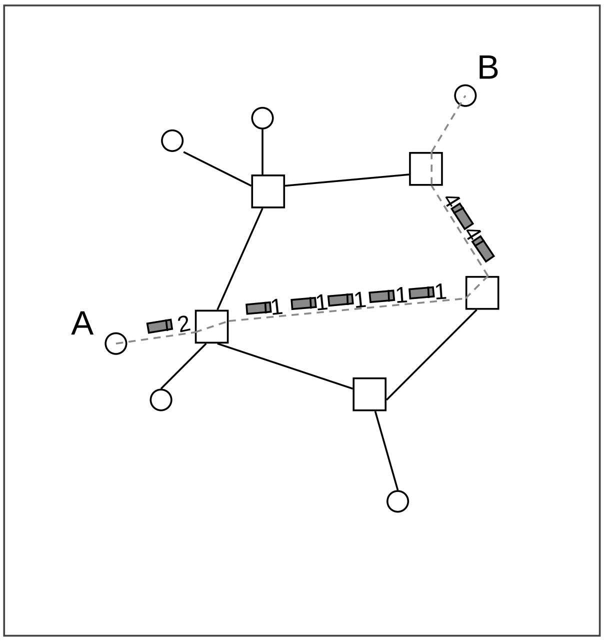

Technical choices

- Connection oriented
  $\Rightarrow$ resource reservation possible
  $\Rightarrow$ guaranteed QoS possible

- Packet switching
  $\Rightarrow$ statistical multiplexing
  $\Rightarrow$ flexibility (any bit rate possible) and efficiency (high utilization)

- Packets small and fixed-length (cells)
  $\Rightarrow$ cell switching
  $\Rightarrow$ fast switching

Cell

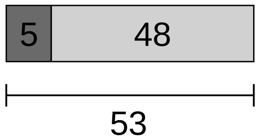

- **Cell = short, fixed-length packet**
  - Total length = **53 bytes** (octets)
  - **Header**: 5 bytes
    - **GFC**, generic flow control (4 [0] bits at UNI [NNI])
    - **VPI**, virtual path identifier (8 [12] bits ⇒ 256 [4096] values)
    - **VCI**, virtual channel identifier (16 bits ⇒ 65,536 values)
    - **PT**, payload type (3 bits)
    - **CLP**, cell loss priority (1 bit)
    - **HEC**, header error control (8 bits)
  - **Information field**: 48 bytes
    - compromise (Europe: 32 bytes; USA: 64 bytes; (64 + 32)/2 = 48)
    - carried transparently, without error detection/recovery

Virtual connections

- Basic connection: **Virtual Channel Connection** (**VCC**)
  - identified by the **VPI/VCI** pair (24 [28] bits) in **each** cell header
    - max. 16,777,216 [268,435,456] VCCs per physical link
  - VPI and VCI fields are **local** labels
    ⇒ reuse possible in different physical links ⇒ scalability
- Aggregated connection: **Virtual Path Connection** (**VPC**)
  - identified by the **VPI** field (8 [12] bits) in **each** cell header
    - max. 256 [4096] VPCs per physical link
  - consists of the VCCs with the same VPI

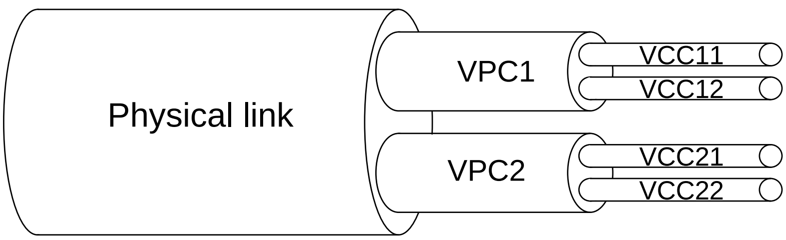

Virtual paths

- Pros
  - faster connection establishment
  - easier network management
  - differentiated QoS possible
  - virtual networks possible
- Cons
  - reduced statistical multiplexing gain

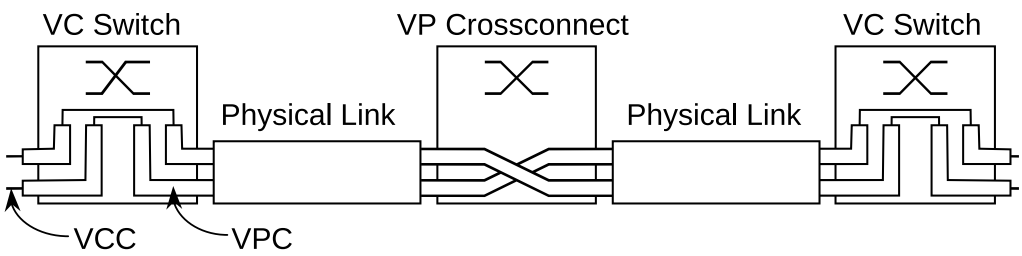

Alternative: Internet

- Based on **connectionless packet switching**
- Originally designed for **non-real-time computer communications**
  - file transfer
  - remote use of central computers
  - e-mail
- No guaranteed Quality of Service
  - only ”Best Effort”
- But currently: ubiquitous
  - TCP/IP protocol suite supported by
    - (virtually) any computer (including your own PC) and
    - very many applications

Contents

- Introduction
- ATM technique
- Service categories
- Traffic contract
- Traffic and congestion control in ATM
- Connection Admission Control (CAC)
- Usage Parameter Control (UPC)
- ABR flow control

Traffic sources

- **Traffic source** =
  end-system application generating the cells of an ATM connection
- **Real-time** traffic sources
  - e.g. interactive voice and video
  - require tight delay and delay variation constraints
  - stream (rate-oriented) traffic
  - constant-bit-rate (CBR) vs. variable-bit-rate (VBR) traffic sources
    - voice and video can be either CBR-coded or VBR-coded
- **Non-real-time** traffic sources
  - e.g. data messaging (such as E-mail) and data or image retrieval (such as WWW) applications
  - does not require tight delay and delay variation constraints
    but typically benefits from shorter delays
  - elastic (unit-oriented) traffic
    - without any natural “rate”

Quality of service

- **Connection level** quality (also called **Grade of Service**, **GoS**) :
  - low connection blocking probability required by all connection types
- **Cell level** quality (also called **Quality of Service**, **QoS**):
  - For real-time traffic sources:
    - End-to-end delay and delay variation are important
      - information (arriving at the destination) is usable only if the end-to-end delay is less than some (strict) constraint
      - in addition, end-to-end delay variation should be negligible
    - Certain transmission rate should be guaranteed
    - Instead, a certain degree of information corruption or loss is tolerable
  - For non-real-time traffic sources:
    - most important: uncorrupted and lossfree information transfer
    - benefits from (but no strict constraints on) shorter delay
    - benefits from (but no strict constraints on) higher transmission rate

Service categories (1)

- Due to differences in the nature of traffic sources and also in the required QoS they impose, connections are divided into different **service categories**
  - ATM-Forum terminology [2]: ‘Service category’
  - ITU-T terminology [3]: ‘ATM transfer capability’
- Functions, such as routing, connection admission control, and resource allocation are, in general, structured differently for different service categories
- Service category is chosen during the connection establishment phase
  - more detailed description of the traffic characteristics and the QoS requirements is given by a set of traffic and QoS parameters defined in the **traffic contract**

Service categories (2)

- **Service categories** (according to ATM Forum [2]):
  - **CBR** = constant bit rate
    - real-time, quaranteed QoS
  - **VBR-rt** = variable bit rate, real-time
    - real-time, quaranteed QoS
  - **VBR-nrt** = variable bit rate, non-real-time
    - non-real-time, quaranteed QoS
  - **ABR** = available bit rate
    - non-real-time, no absolute QoS quarantees
  - **UBR** = unspecified bit rate
    - non-real-time, no QoS quarantees at all

Service categories (3)

- **CBR** (Constant Bit Rate)
  - mainly for real-time constant-bit-rate traffic sources requiring tightly constrained cell transfer delay and delay variation
    - e.g. cbr-coded interactive voice and video, circuit emulation
  - traffic characterized and policed in terms of Peak Cell Rate (PCR)
  - guaranteed QoS in terms of cell transfer delay (CTD), cell delay variation (CDV) and cell loss ratio (CLR)

Service categories (4)

- **VBR-rt** (Variable Bit Rate, real-time)
  - for real-time variable-bit-rate traffic sources requiring tightly constrained cell transfer delay and delay variation
    - e.g. vbr-coded interactive voice and video
  - traffic characterized and policed in terms of Peak Cell Rate (PCR), Sustainable Cell Rate (SCR) and Maximum Burst Size (MBS)
  - guaranteed QoS in terms of CTD, CDV and CLR
- **VBR-nrt** (Variable Bit Rate, non-real-time)
  - for non-real-time variable-bit-rate traffic sources expecting a low cell loss ratio
    - e.g. vbr-coded video retrieval
  - traffic characterized and policed in terms of PCR, SCR and MBS
  - guaranteed QoS in terms of CLR

Service categories (5)

- **ABR** (Available Bit Rate)
  - for non-real-time elastic traffic sources expecting a low cell loss ratio
    - e.g. transfer of large bulks of data (large files, images, video clips, etc.)
  - traffic initially specified by the maximum required bandwidth (PCR) and the minimum usable bandwidth (MCR)
  - however, traffic sources should be able to adjust their rate according to a specific feed-back control mechanism (ABR flow control)
  - guaranteed Minimum Cell Rate (MCR)
  - no absolute QoS guarantees (“low” cell loss ratio promised, fairness as a target)
- **UBR** (Unspecified Bit Rate)
  - for non-real-time elastic traffic sources without any QoS requirements
    - e.g. TCP/IP data traffic
  - no rate or QoS guarantees at all (“laissez-faire” service)
  - relies totally on upper level traffic control (such as TCP)

Contents

Traffic Contract

- During the connection establishment phase, the user and the network negotiate a **Traffic Contract**
  - concerns traffic (that is: transmitted cells) at the interface between the user and the network (**UNI**)
- Traffic Contract specifies
  - the service category of the connection
  - the traffic characteristics of the connection
    - given by a set of **traffic parameters** and **delay tolerances** (PCR, SCR, MBS, MCR, CDVT, BT)
  - Quality of Service (QoS) guaranteed for the connection by the network
    - given by a set of **QoS parameters** (maxCTD, peak-to-peak CDV, and CLR)

Traffic parameters

- **PCR** = $1/T$ = peak cell rate (cells per second)
  - max. instantaneous transmission rate
  - defined for all connections
- **SCR** = $1/T_s$ = sustainable cell rate
  - max. rate at which the traffic source can transmit cells continuously
  - defined for VBR connections
- **MBS** = maximum burst size
  - max. number of cells that can be sent consecutively at the PCR
  - defined for VBR connections
- **MCR** = minimum cell rate
  - max. rate at which cells are guaranteed to be carried by the network
  - defined for ABR connections

Delay tolerances

- **CDVT** = τ = cell delay variation tolerance
  - an upper bound for cell “clumping” due to ATM layer functions (independent of the traffic source), e.g.
    - cell multiplexing (slotted system!)
    - insertion of PHY layer overhead or OAM cells
  - defined for all connections
- **BT** = τ_s = burst tolerance
  - $\mathrm{BT} = (\mathrm{MBS} - 1)(T_s - T)$ where
    - $T_s = 1/\mathrm{SCR}$
    - $T = 1/\mathrm{PCR}$
  - defined for VBR connections

QoS parameters (negotiated)

- **maxCTD** = maximum cell transfer delay
  - maximum end-to-end cell transfer delay (more precisely: 1-α quantile)
  - end-to-end cell transfer delay (CTD) includes
    - fixed part: propagation, transmission and pure processing delays
    - randomly varying part: synchronization and queueing delays
  - defined for CBR and VBR-rt connections (that is: for real-time connections)
- **peak-to-peak CDV** = cell delay variation
  - difference between maxCTD and the fixed part of CTD
  - defined for CBR and VBR-rt connections (that is: for real-time connections)
- **CLR** = cell loss ratio
  - CLR = nr of lost cells / total nr of transmitted cells
  - defined for CBR and VBR connections

maxCTD and peak-to-peak CDV

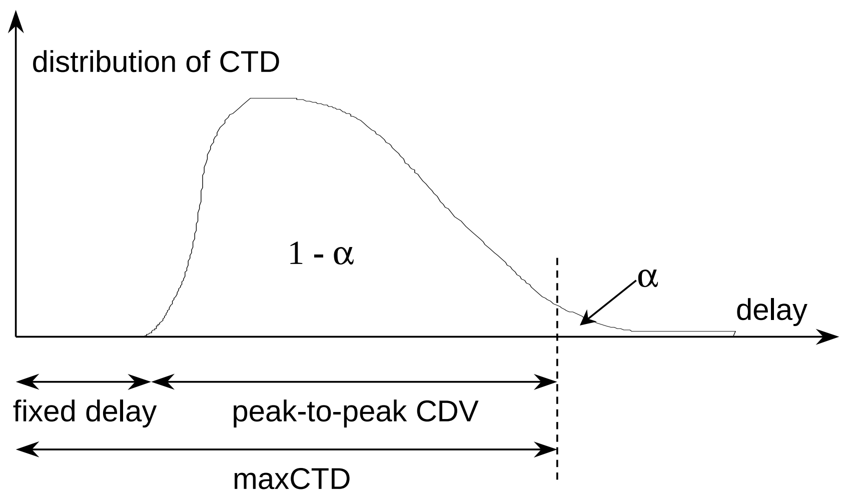

Source: [2]

ATM Service Category Parameters

<table>
<tr><td rowspan="2"></td><td colspan="5">ATM Layer Service Category</td></tr>
<tr><td>CBR</td><td>VBR-rt</td><td>VBR-nrt</td><td>UBR</td><td>ABR</td></tr>
<tr><td>Traffic parameters:</td><td colspan="5"></td></tr>
<tr><td>PCR, CDVT</td><td colspan="5">specified</td></tr>
<tr><td>SCR, MBS</td><td>n/a</td><td colspan="2">specified</td><td colspan="2">n/a</td></tr>
<tr><td>MCR</td><td colspan="4">n/a</td><td>specified</td></tr>
<tr><td>QoS parameters:</td><td colspan="5"></td></tr>
<tr><td>Peak-to-peak CDV</td><td colspan="2">specified</td><td colspan="3">unspecified</td></tr>
<tr><td>MaxCTD</td><td colspan="2">specified</td><td colspan="3">unspecified</td></tr>
<tr><td>CLR</td><td colspan="3">specified</td><td colspan="2">unspecified</td></tr>
</table>

Contents

- Introduction
- ATM technique
- Service categories
- Traffic contract
- Traffic and congestion control in ATM
- Connection Admission Control (CAC)
- Usage Parameter Control (UPC)
- ABR flow control

Time scales for ATM control functions

Response time:	Traffic control functions:	Congestion control functions:
Long term (hours - days)	• Resource Management using Virtual Paths	
Connection duration (secs - mins)	• Connection Admission Control (CAC)	
Round trip propagation time (ms)	• Fast Resource Management	• Explicit Forward Congestion Indication (EFCI) • ABR Flow Control
Cell insertion time (µs)	• Usage Parameter Control (UPC) • Priority Control • Traffic Shaping	• Selective Cell Discard • Frame Discard

Source: [1]

Traffic control functions (1)

- Resource management using virtual paths
  - grouping of similar VCCs
  - differentiated QoS for different VCC groups
  - VPC performance objectives determined by the most demanding VCC QoS requirement
- Connection admission control (CAC)
  - to protect the network from excessive loads
  - during the connection establishment phase, the network decides whether the requested connection can be established or should be rejected
  - when established, the traffic characteristics and the QoS requirements of the connection are specified in a traffic contract
  - if the connection behaves as specified in the traffic contract (compliant connection), the network should guarantee the required QoS

Source: [2]

Traffic control functions (2)

- Usage parameter control (UPC)
  - to protect the other connections against badly behaving connections
  - during the information transfer phase, the network monitors the connection whether the traffic conforms to the traffic contract (traffic policing)
  - the non-conforming cells are either discarded or tagged as low-priority cells (CLP = 1)
- Traffic shaping
  - to smooth out a traffic flow and reduce cell clumping
  - users can ensure that the traffic conforms to the traffic contract
  - network may also shape the traffic (within the limits of the traffic contract)

Source: [2]

Congestion control functions (1)

- Explicit forward congestion indication (EFCI)
  - the EFCI bit in the cell header is used to warn the end systems (first the destination, and then the source) about (impending) congestion
  - a network element in a congested state (or if it is impending) may set the EFCI bit (EFCI = 1) --- but cannot never reset it (EFCI = 0)
  - reasonable if end-systems cooperate reducing their cell rates
- ABR flow control
  - in the ABR service, the source adapts its rate to changing network conditions
  - ABR flow control is based on specific resource management (RM) cells
    - each RM cell include the **explicit cell rate** (ER) field
  - instead of one bit EFCI information (included in every cell), a congested network element can express, by the ER field, an upper bound for the rate at which the source is permitted to transmit cells

Source: [2]

Congestion control functions (2)

- Selective cell discard
  - traffic sources may tag less important cells as low-priority cells
  - UPC can tag non-conforming cells as low-priority cells
  - a congested network element (beyond UPC) may selectively discard low-priority cells or cells belonging to a non-compliant connection
- Frame discard
  - If a network element needs to discard cells, it is in many cases more efficient to discard at the frame (packet) level rather than at the cell level
  - Frame discard may be used whenever it is possible to delineate frame boundaries

Source: [2]

Contents

Connection admission control (CAC)

- During the connection establishment phase, CAC determines
  - whether the requested connection can be accepted or not
  - the traffic parameters needed by UPC
  - allocation of network resources (including routing decision)
- Resource allocation is due to the operator itself
  - no method is standardized (but is open for competition)
  - should be based on mathematical traffic models
- A new connection can be accepted only if
  - the network has resources to support the new connection while at the same maintaining the agreed QoS level of existing connections
- Rejections mean a higher connection blocking probability
  - the problem (of lacking resources) is pushed up to a higher time scale
  - the only (permanent) solution is to increase the network resources

Resource allocation

- CBR
  - bandwidth allocated according to PCR
- VBR
  - bandwidth allocated according to so called **Effective Cell Rate** (ECR)
  - ECR is determined by the operator
    - $SCR \leq ECR \leq PCR$
    - ECR depends not only on the traffic characteristics and the QoS requirements of the traffic source but also on the network resources (e.g. link capacities)
- ABR
  - bandwidth allocated according to MCR
- UBR
  - no bandwidth allocation at all

Example

- Consider a VPC dedicated to a single service category
  - C = VPC's PCR (capacity of logical link)
  - Existing connections: $i = 1,\ldots n$
  - New connection request of the same type: $n + 1$
- The new connection is accepted if …
  - CBR: … $\mathrm{PCR}_1 + \ldots + \mathrm{PCR}_{n+1} \leq C$
  - VBR: … $\mathrm{ECR}_1 + \ldots + \mathrm{ECR}_{n+1} \leq C$
  - ABR: … $\mathrm{MCR}_1 + \ldots + \mathrm{MCR}_{n+1} \leq C$
  - UBR: … always!

Statistical multiplexing of independent VBR sources

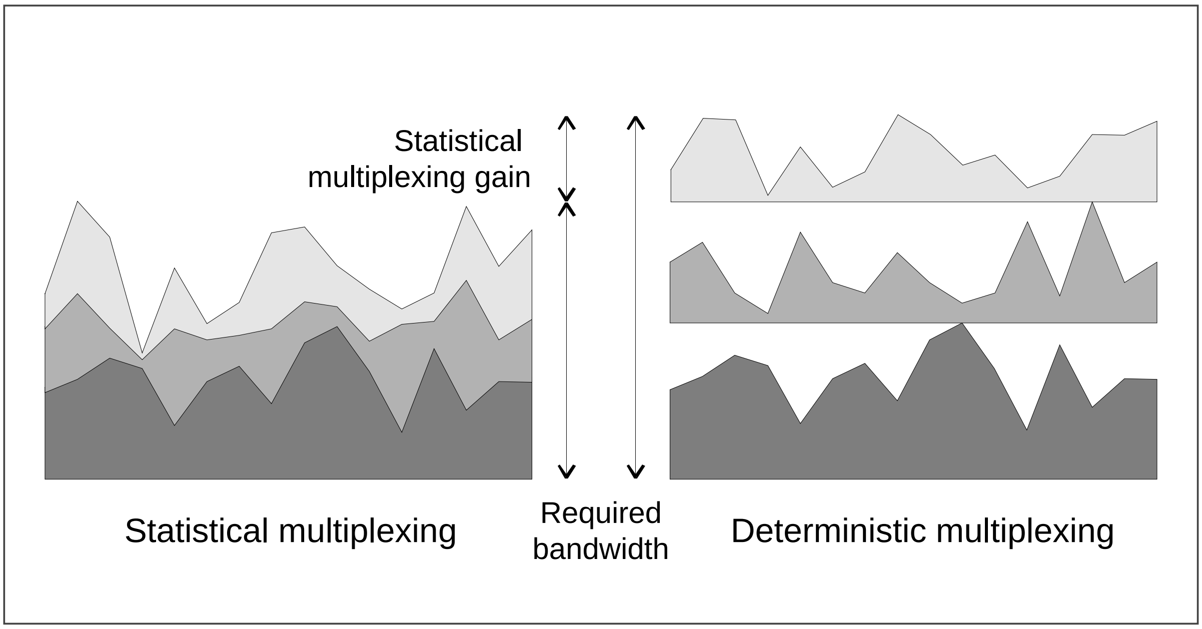

Effective cell rate (1)

- Consider n **independent** and identical VBR traffic sources
  - Let R_i denote the instantaneous cell rate of source i
  - Assume that the corresponding VBR connections are carried by a VPC of capacity C
- Using so called **bufferless** model, we approximate the cell loss ratio by the probability $P\{\Sigma_i R_i > C\}$
- Now, given ε, we determine $e = \mathrm{ECR}_i$ so that $P\{\Sigma_i R_i > ne\} = \varepsilon$

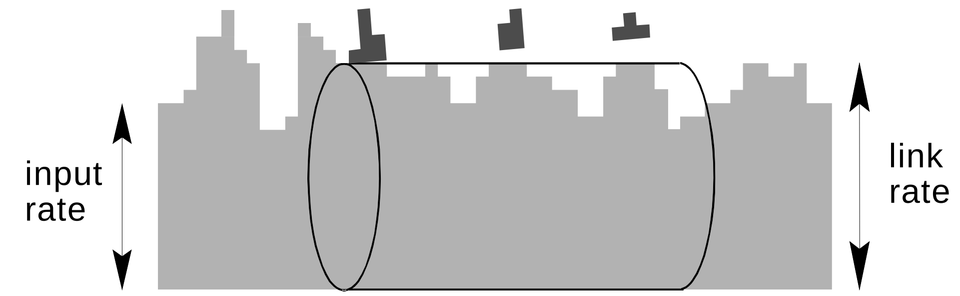

Effective cell rate (2)

- Denote
  - $m = E[R_i]$, $v = \mathrm{Var}[R_i]$
- Due to independency,
  - $M := E[\Sigma_i R_i] = nm$, $V := \mathrm{Var}[\Sigma_i R_i] = nv$
- By the Central Limit Theorem,
  - $\Sigma_i R_i \approx \mathrm{N}(M,V) = \mathrm{N}(nm,nv)$
- Thus,

$$P\{\sum_i R_i > ne\} = \varepsilon$$
$$\Leftrightarrow \quad ne = M + z_{1-\varepsilon}\sqrt{V} = nm + z_{1-\varepsilon}\sqrt{nv}$$
$$\Leftrightarrow \quad e = m + z_{1-\varepsilon}\sqrt{\frac{v}{n}}$$

  - Here z_p refers to the p-quantile of the N(0,1) distribution

Contents

Usage Parameter Control (UPC)

- During the information transfer phase, UPC controls the cell flow of the connection (at the UNI interface) by
  - passing the conforming cells
  - tagging (CLP = 0 $\rightarrow$ CLP = 1) or discarding the non-conforming cells
- The conformance of cells is determined by the **Generic Cell Rate Algorithm** (GCRA)
  - GCRA(I,L) is specified by the following two time parameters:
    - I = Increment = $1/R$, where R is the target rate (under policing)
    - L = Limit = the maximum negative deviation from the “scheduled” time
- Peak cell rate policing:
  - GCRA(1/PCR,CDVT) = GCRA(T,τ)
- Sustainable cell rate policing:
  - GCRA(1/SCR,BT) = GCRA(T_s,τ_s)

GCRA(I,L)

- Notation:
  - t_k = theoretical arrival time of the kth cell
  - a_k = arrival time of the kth cell
- Initialization at time a_1 :
  - $t_1 = a_1$
- Algorithm run at times a_k
  - $k = 1,2,...$

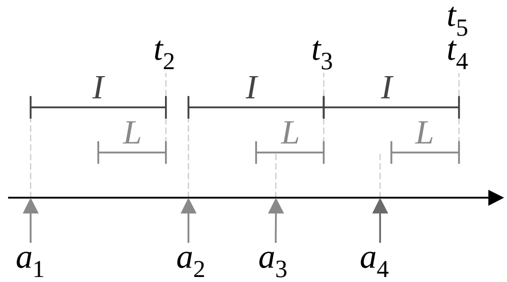

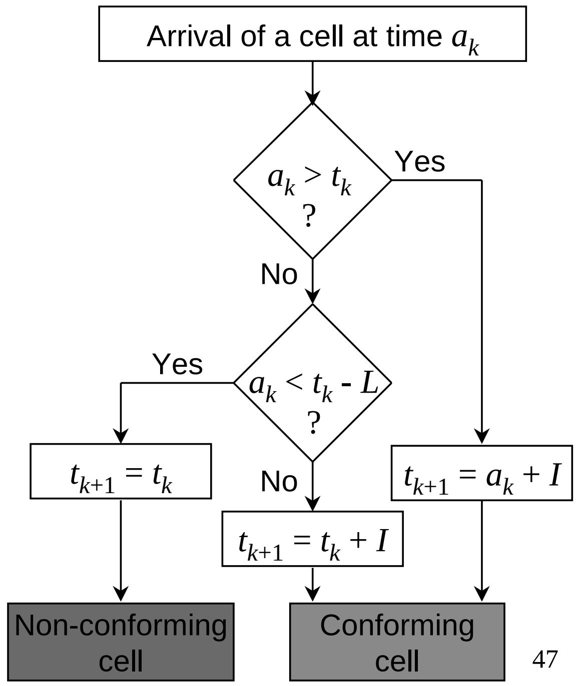

Contents

“Best Effort” services

- Guaranteed QoS for CBR and VBR connections ⇒ low utilization
- Available free capacity can be offered to the (lower priority) ABR and UBR connections (so called “best effort” services)
- Because of no explicit resource reservation for ABR and UBR connections,
  - essentially larger buffers are needed
  - fairness is an important aspect

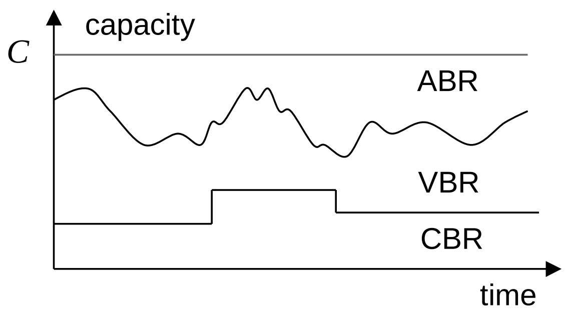

ABR flow control

- Target:
  - to achieve a high link utilization and a low cell loss ratio at the same time
- Traffic of an ABR connection initially specified by
  - the maximum required bandwidth (PCR) and
  - the minimum usable bandwidth (MCR)
- To adjust the total rate of ABR connections to fit with the available rate (left over by CBR and VBR connections), the transmission rate of each ABR connection is controlled continuously by the network (closed-loop control)
  - the allowed (instantaneous) transmission rate of a single connection can have any value between MCR and PCR during the connection holding time
  - implementation by means of specific Resource Management (RM) cells
- This control method is clearly reactive

Rate based ABR flow control

- Source transmits RM cells after each *n* ordinary cells
  - telling its current rate (CCR) and the rate at which it wishes to transmit (ER)
- Switches along the (fwd) route of the connection read this information
  - they compute the fair share of the link capacity
  - they can also reduce the value of the ER field
- Destination returns the RM cell to the source
  - it can also reduce the value of the ER field
- Switches along (bwd) the route of the connection read this information
  - they can still reduce the value of the ER field
- Source adjust its current rate according to the feedback information

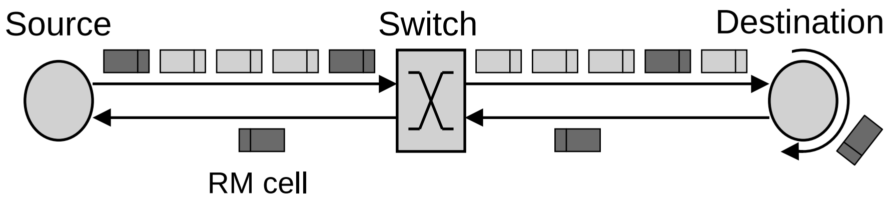

THE END