Asynchronous Transfer Mode (ATM)

- A Connection-oriented network providing Quality of Service guarantees
  - – developed in the mid-1980's to combine packet switching with TDM
  - – had been hoped it would be the end-user to end-user network of choice
    - » to provide broadband services beyond ISDN
  - – TCP/IP Internet became the *de facto* standard instead
  - – ATM now used as part of the network infrastructure
    - » e.g. in enterprise networks, IP over ATM, Voice over ATM etc.
  - – claimed benefits :
    - » high performance *hardware switching*
    - » *dynamic bandwidth* allocation for bursty traffic
    - » *class-of-service* support for multimedia applications
    - » common LAN/WAN architecture via *LAN emulation*
    - » *scalability* in speed and network size

- An ATM network is a collection of endpoints, switches and interfaces :

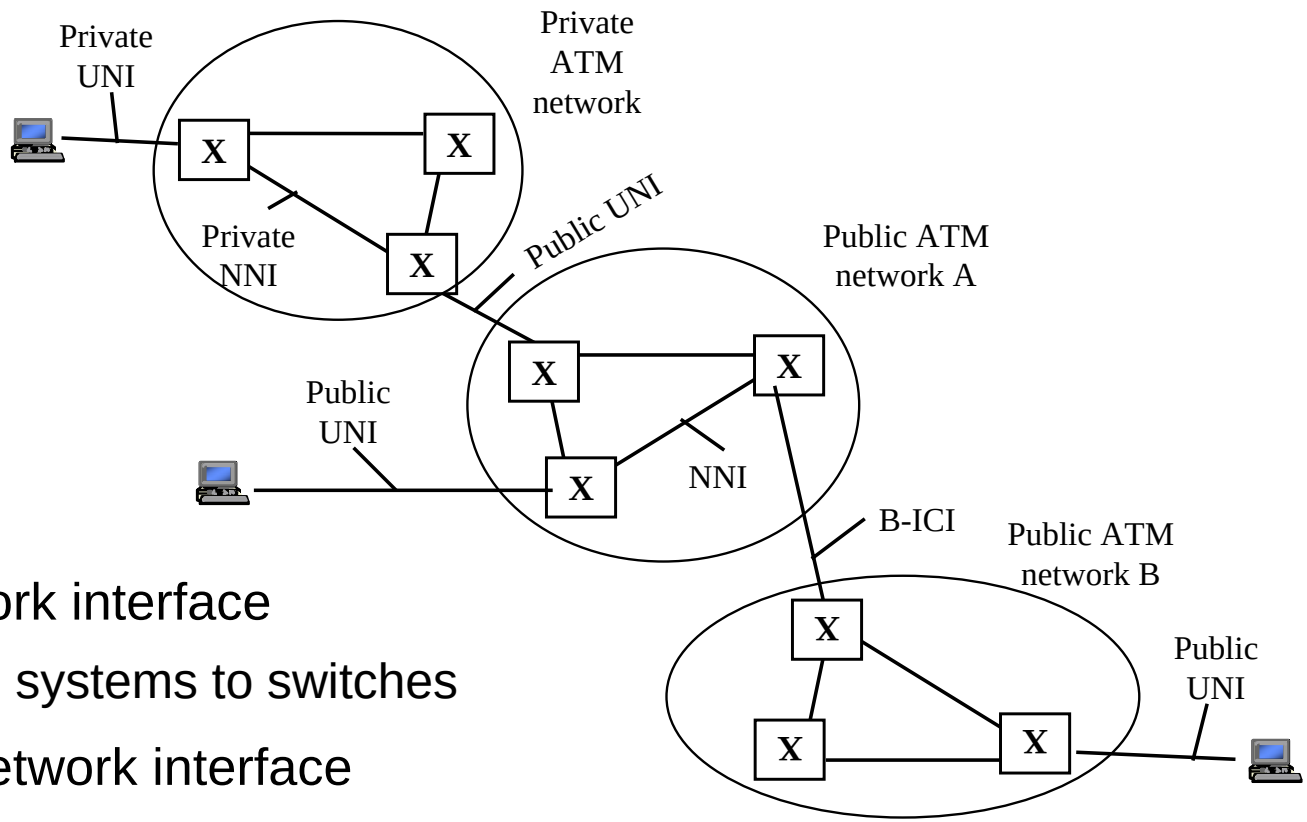

  - UNI : user-network interface
    - » connects end systems to switches
  - NNI : network-network interface
    - » connects two switches
  - B-ICI : broadband intercarrier interface
    - » connects two public switches from different service providers
  - provides Permanent Virtual Connections (PVCs) e.g. leased lines, and Switched Virtual Connections (SVCs) on demand

- ATM Technology
  - – all information formatted into fixed-length *cells* of 48 bytes + 5 bytes header
    - » ensures time-critical information not adversely affected by long packets
  - – header organised for efficient hardware switching
    - » carries payload type information, virtual-circuit identifiers, CRC etc.
  - – information from separate flows converted into cells and then multiplexed
    - » cells queued and transmitted according to some scheduling strategy
    - » no time-slot reservation as in TDM
      - cells flow asynchronously

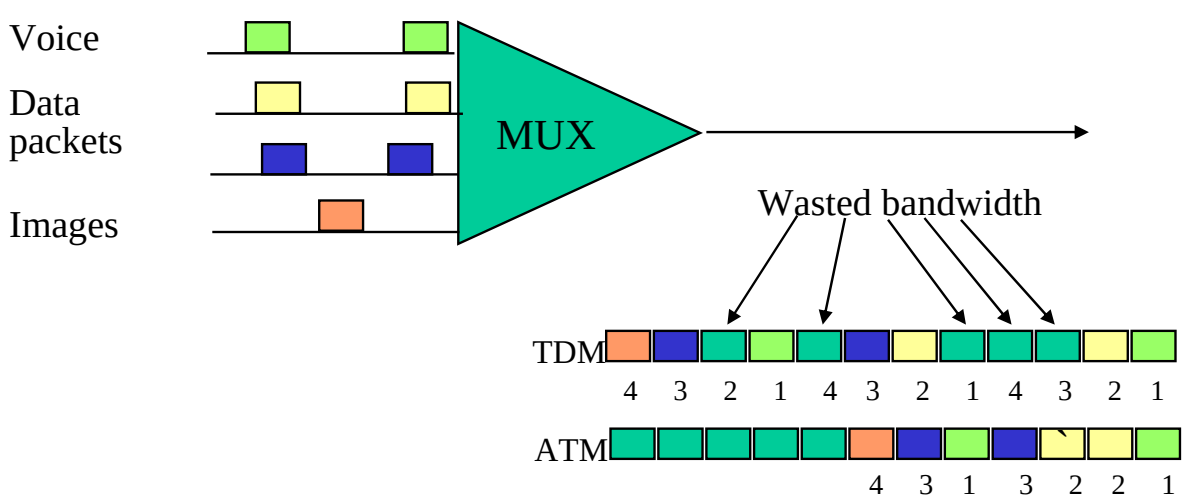

– ATM is connection-oriented

» a connection setup required prior to transfer of cells

» identifies a path through the network that can provide the required service

- a virtual channel connection (VCC)

» the VCC has a chain of local identifiers each used at a switch (VCIs)

- each input port to a switch uses its own private set of VCIs
- used to index into its routing table to find the next hop output port

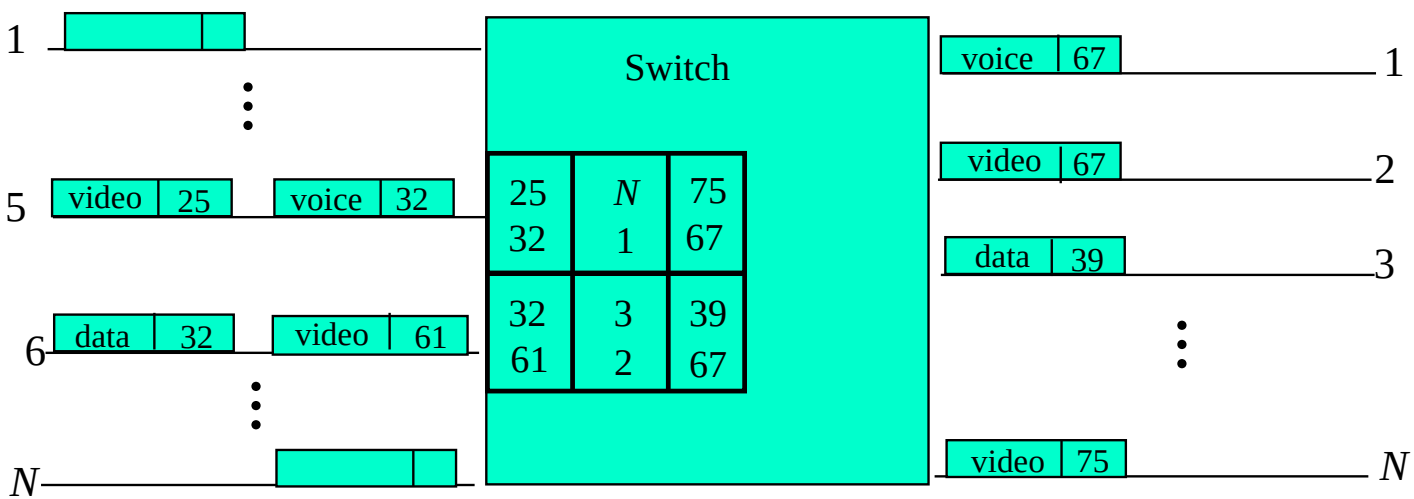

» high speed switching relatively straightforward, 155Mbps streams typical

- fixed sized cell with standard headers

– Aggregation into *virtual paths*:

» virtual channel connections (VCC) sharing a common path through the network

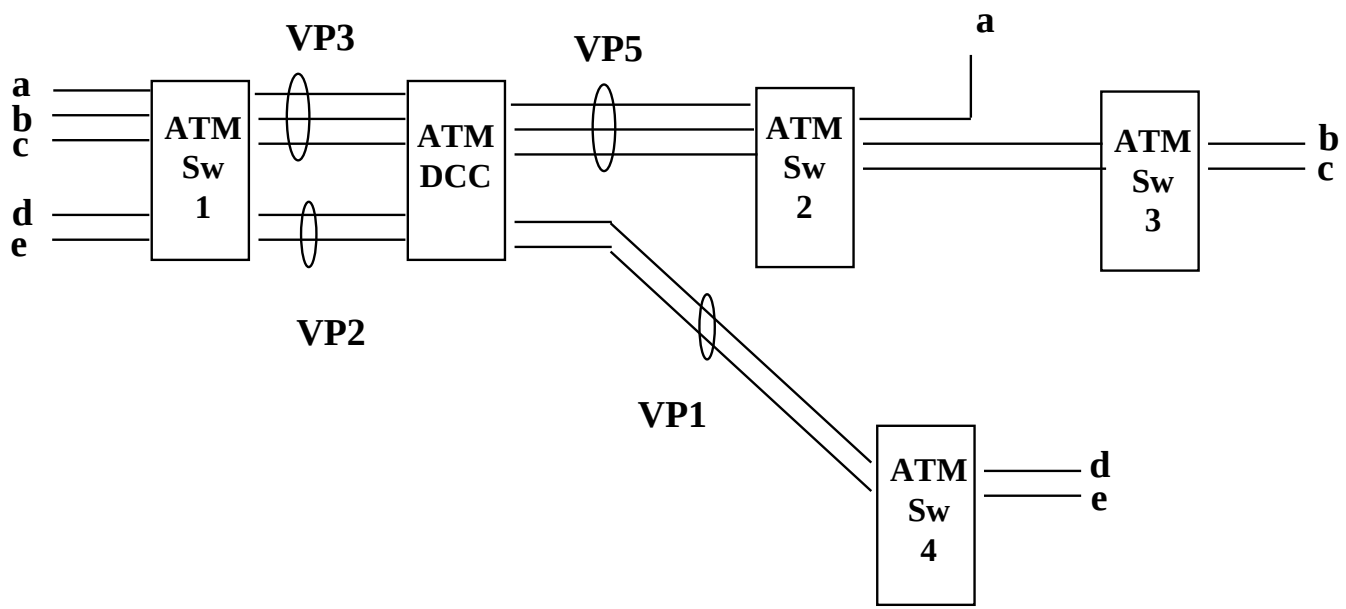

» VCCs a, b, c aggregated into VP3 after switch 1

- and to VP5 after the Digital Cross Connect to switch 2

» VCCs d, e aggregated into VP2 after switch 1

- and to VP1 after the DCC to switch 4

» VCCs b, c could have been aggregated after switch 2 to switch 3

» the DCC here switches aggregated virtual paths, not individual VCCs

– Two levels of identifiers

» Virtual Channel Identifier (VCI) and Virtual Path Identifier (VPI)

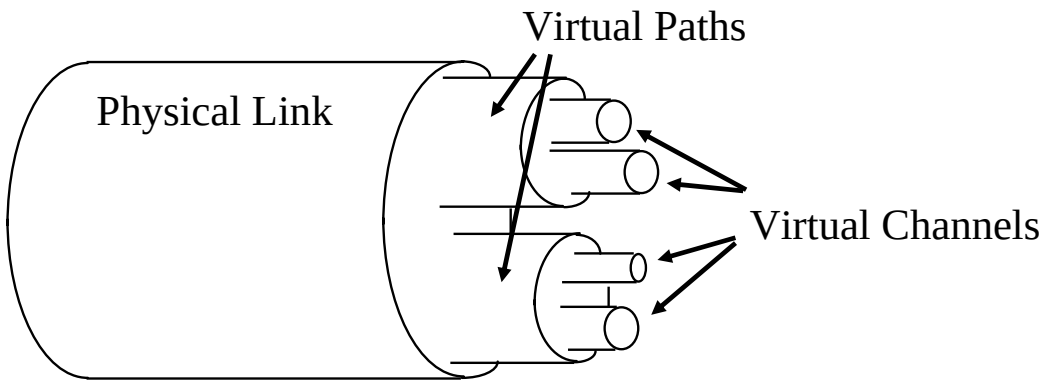

» all switches along a virtual path switch on the basis of the VPI alone

» VCIs only used at the end of a virtual path

– the VCI/VPI structure supports scalability to very large networks

– DCCs allow network path topology reconfiguration under software control

» similar to SONET ring reconfigurability

– network managers can allocate bandwidth to virtual paths

» and to any degree of granularity

– ATM Cells:

UNI cell

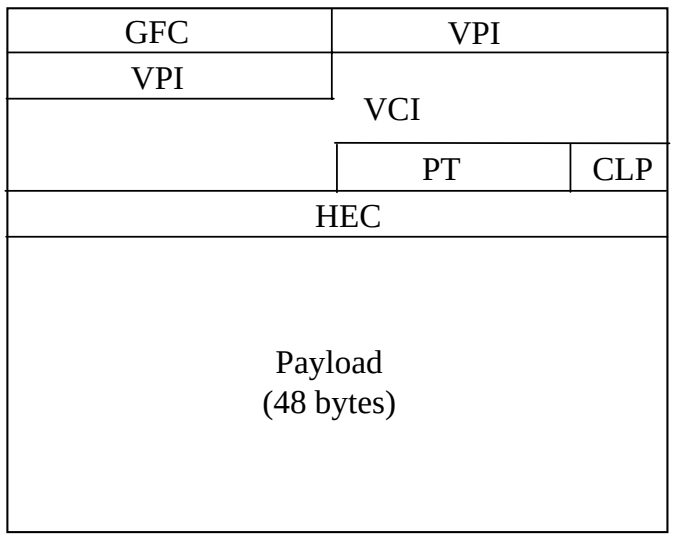

NNI cell

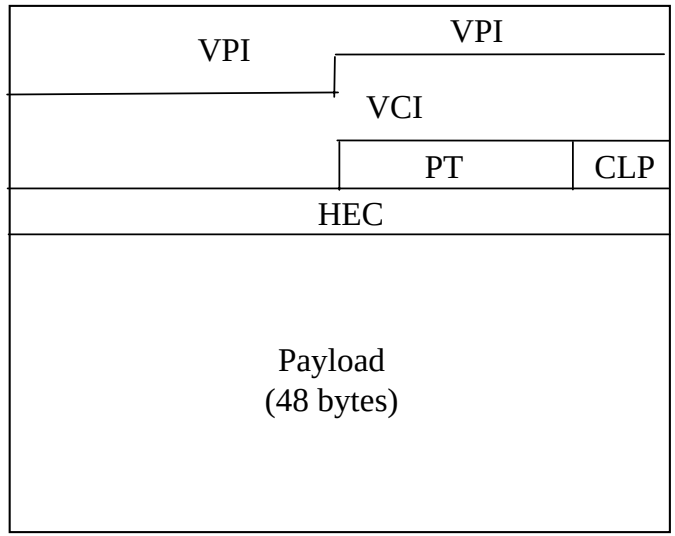

» GFC (Generic Flow Control) : local functions e.g. identifying multiple stations at end-points

- not used in practice

» no GFC in NNI header allows larger trunks between public ATM switches

» VPI : 8 or 12 bits; VCI : 16 bits

» PT (Payload Type) : payload contains user data/control data

» CLP (Cell Loss Priority) : lower priority cells get discarded first when congested

» HEC (Header Error Control) : CRC-8 (x^8+x^2+x+1) over first 4 bytes of header

- detection and/or single bit correction

- Reference Model:

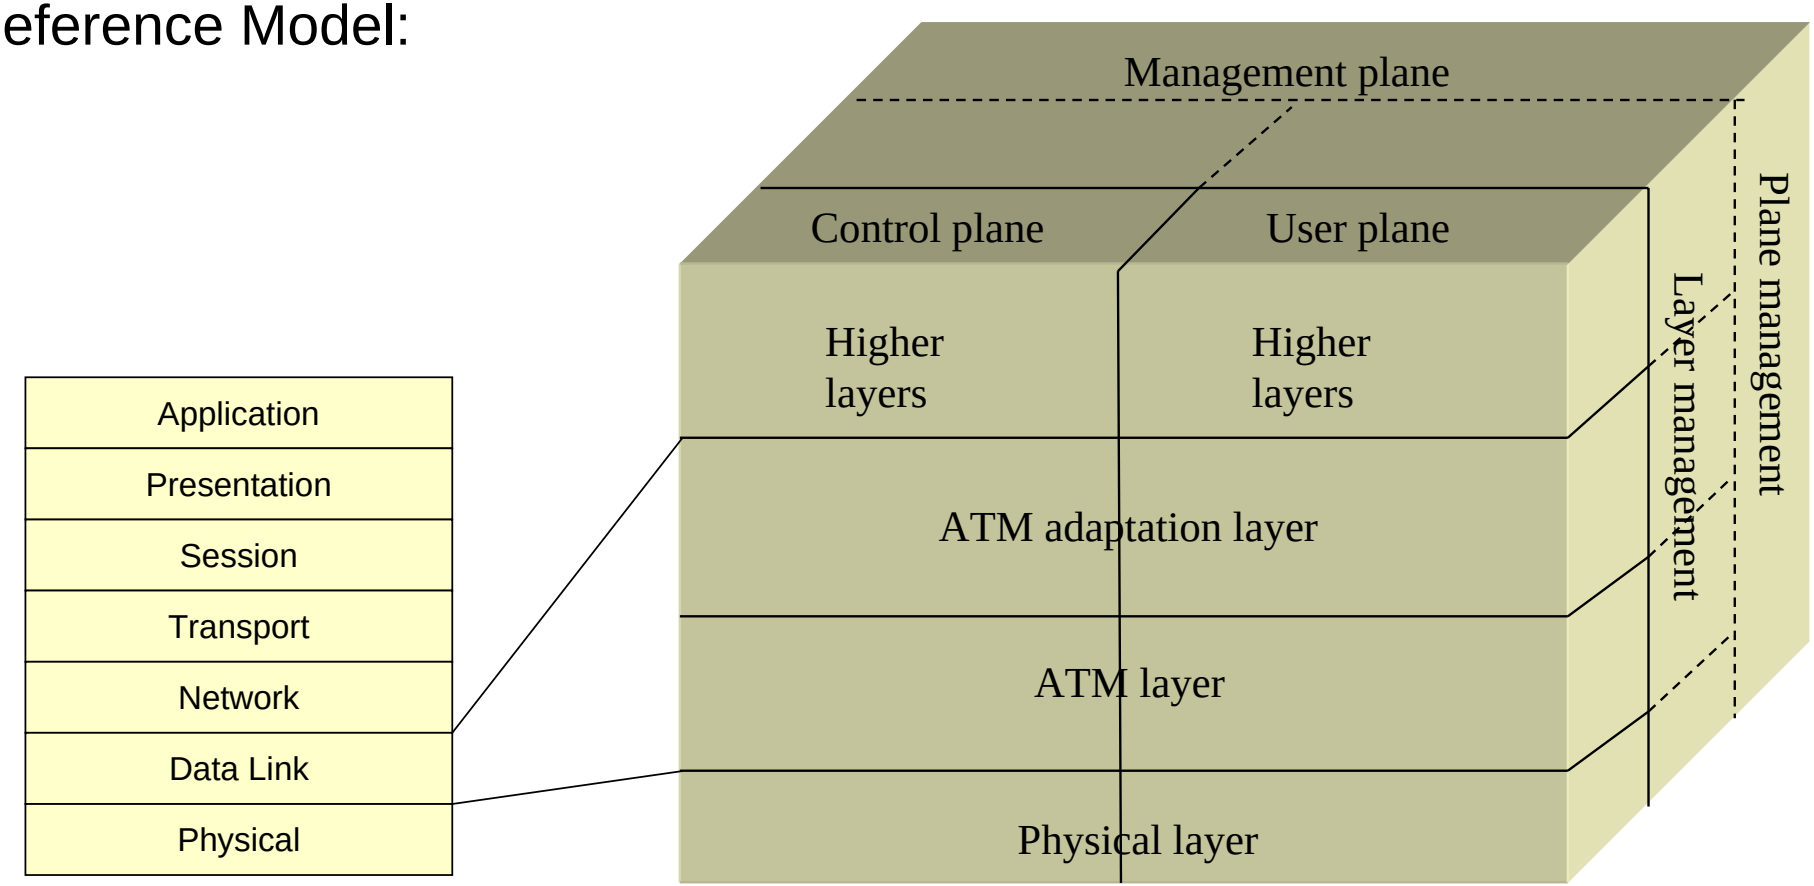

» user plane deals with data transfer, flow control, error recovery etc.

» control plane deals with signalling to set up, manage and release connections

» management plane deals with network resources and coordination of other planes

– Physical layer:

» transmission convergence sublayer

- ATM cell boundaries tracked
- checking of header checksums
- insertion and removal of idle cells

» physical medium dependent layer

- cells are converted into a bitstream
- cells packaged into frames for the medium e.g. SONET, FDDI, STP etc.

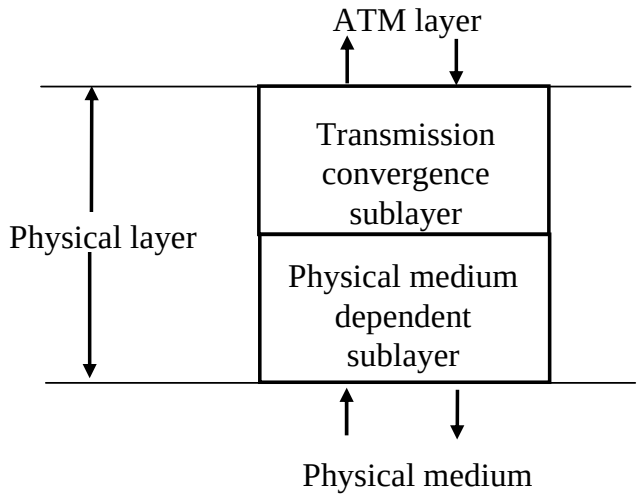

– ATM layer:

» concerned with sequenced transfer of cells in connections across the network

» accepts 48-byte blocks from the ATM Adaptation layer and adds 5-byte header

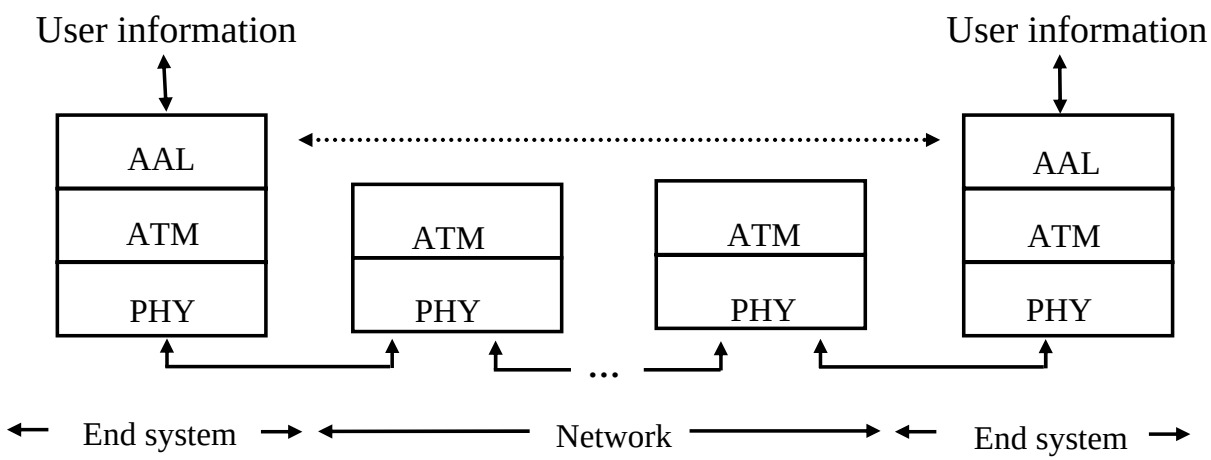

– ATM Adaptation layer :

» responsible for providing different applications with the appropriate support

- several AAL types are defined – AAL1 to AAL5

» converts higher level SDUs into 48-byte blocks for ATM layer

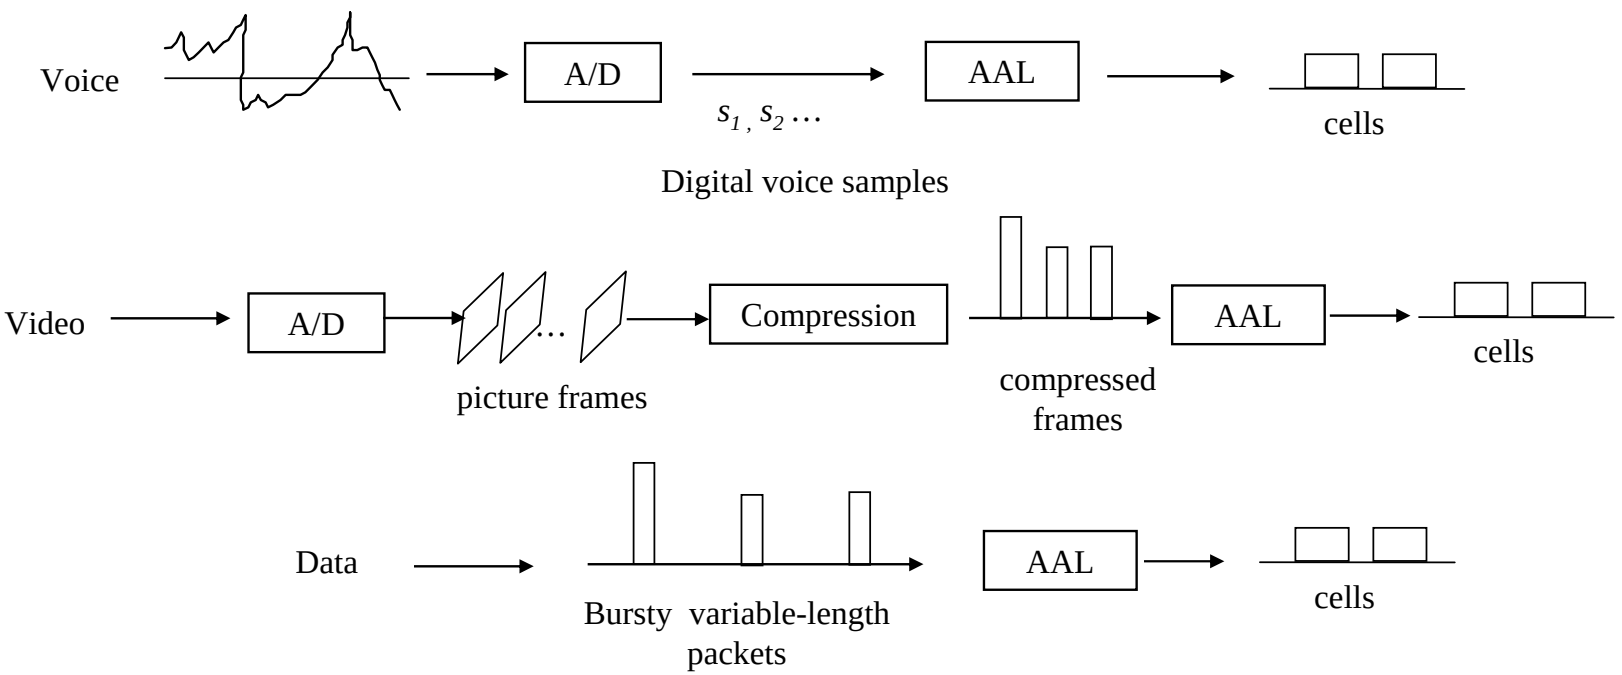

» provides support for the protocol layer directly above

- e.g. to provide a reliable stream service, if needed (not if TCP was the layer above)
- cell sequence numbering
- segmentation and reassembly support

- Service Classes
  - Constant Bit Rate (CBR)
    - » rate constant for the whole period of the connection
    - » for traffic sensitive to delay
      - e.g. voice, video, TV, circuit switching emulation
  - Variable Bit Rate, non-real-time (VBR-NRT)
    - » for bursty sources with no rigorous timing requirements
      - rate varies with time depending on rate at which application produces data
      - e.g. multimedia e-mail, transaction processing
  - Variable Bit Rate, real-time (VBR-RT)
    - » similar to VBR-NRT for applications sensitive to cell-delay variation
      - e.g. voice with speech activity detection
  - Available Bit Rate (ABR)
    - » allows sources to make use of any bandwidth that is still available
      - e.g. file transfer, e-mail
  - Unspecified Bit Rate (UBR)
    - » for everything else, including TCP/IP – lowest tariff

- Traffic Descriptors
  - parameters which the source must specify when negotiating a connection
  - traffic will be policed and enforced to conform with the parameters
  - Peak Cell Rate (PCR)
    - » rate in cells/second that the source is never allowed to exceed
    - » inverse of minimum inter-cell arrival time
  - Sustainable Cell Rate (SCR)
    - » average cell rate in cells/second produced by the source over a long period
  - Maximum Burst Size (MBS)
    - » number of consecutive cells that may be transmitted by the source at the PCR
    - » used as the bucket size parameter for the enforcement algorithm
  - Minimum Cell Rate (MCR)
    - » minimum cell rate in cells/second that the source is always allowed to send
  - Cell Delay Variation Tolerance (CDVT)
    - » tolerable level of cell delay variation in a given connection

- Negotiable Quality of Service parameters
  - Cell Loss Ratio (CLR)
    - » ratio of number of lost cells to total transmitted cells
      - lost in the network due to congestion and/or buffer overflow
    - » specified as an order of magnitude in the range 10^{-1} to 10^{-15}
    - » extent to which CLR can be negotiated depends on buffer allocation strategies available in the network
  - Cell Transfer Delay (CTD)
    - » delay experienced by a cell between network entry and exit
    - » includes propagation delays, queuing delays at intermediate switches and processing times at queuing points
    - » since cells each experience different delays, CTD specified by a probability density function
    - » maximum CTD can be negotiated by which some fraction $(1-\alpha)$ of the cells will be delivered
      - where α is some appropriately small value

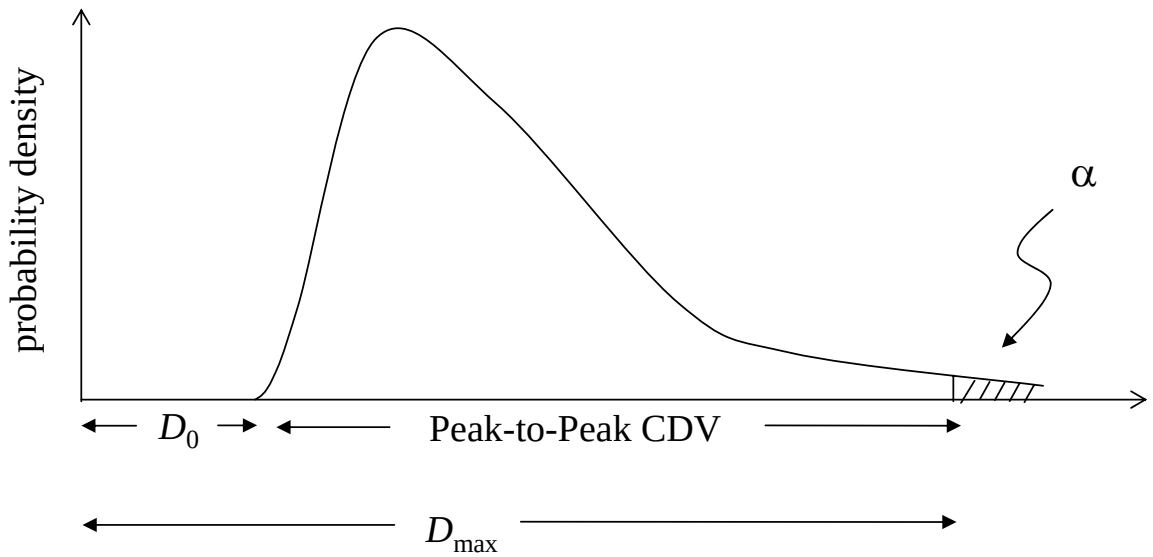

– Cell Delay Variation (CDV)

» the total delay encountered by cells in a connection

- excluding fixed delay D_0

» peak-to-peak delay can currently be negotiated

» network switches have only limited control of the variance of CTD values

- so range of negotiable CDV value is also limited

– relevant parameters for service classes:

Class of Service	CBR	VBR-NRT	VBR-RT	ABR	UBR
CLR	✓	✓	✓	✓	✕
CTD	✓	✕	✓	✕	✕
CDV	✓	✓	✓	✕	✕
PCR	✓	✓	✓	✕	✓
SCR	✕	✓	✓	✕	✕
MBS	✕	✓	✓	✕	✕
flow control	✕	✕	✕	✓	✕

– Connection Admission Control (CAC)

» ensures that new virtual connections are assigned to links with sufficient bandwidth to meet the committed levels of Quality of Service

- algorithms not standardised – left up to equipment suppliers

– Policing

» a Generic Cell Rate Algorithm (GCRA) defined

- equivalent to the leaky bucket scheme

- Connection establishment
  - setup, call proceeding, connect and connect acknowledgment messages:

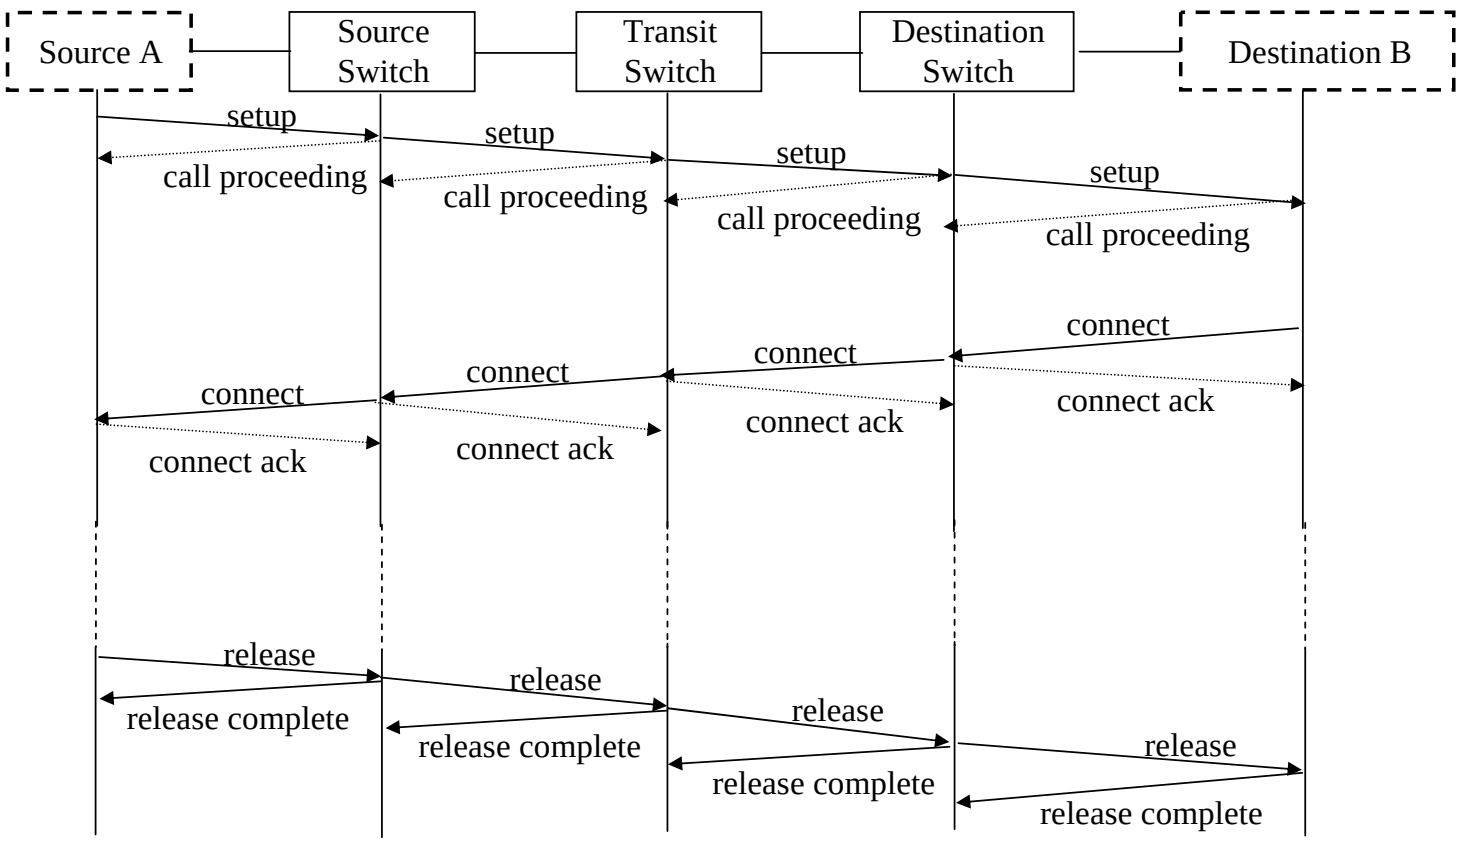

- ATM addressing
  - has its own system : telephony-oriented or private end-system addresses