CS4/MSc
Computer Networking

Lecture 10:

Asynchronous Transfer Mode Networks

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
  - Conceived to be the end-to-end, fully-digital network of choice providing 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

ATM Technology

- 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*

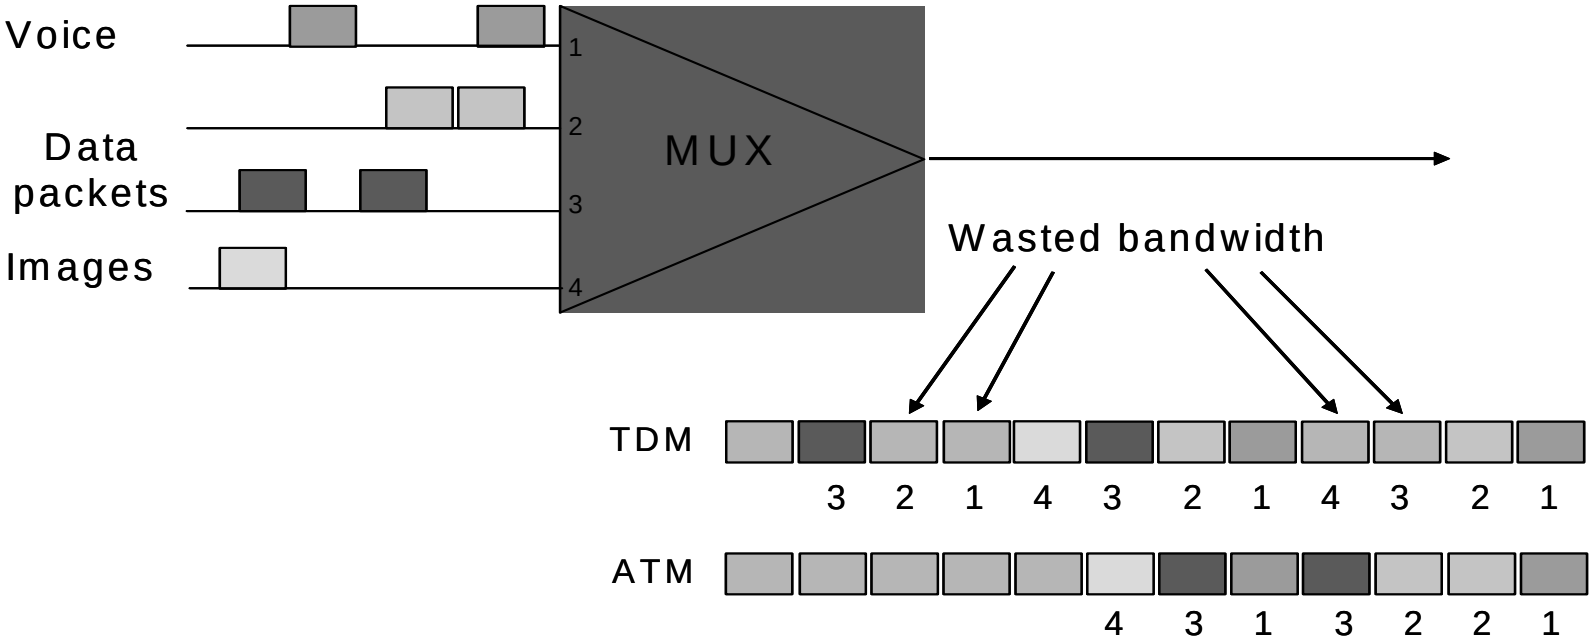

ATM Virtual Circuits

- Connection setup procedure sets up a path: 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
  - Facilitates very fast switching: no search for destination address needed

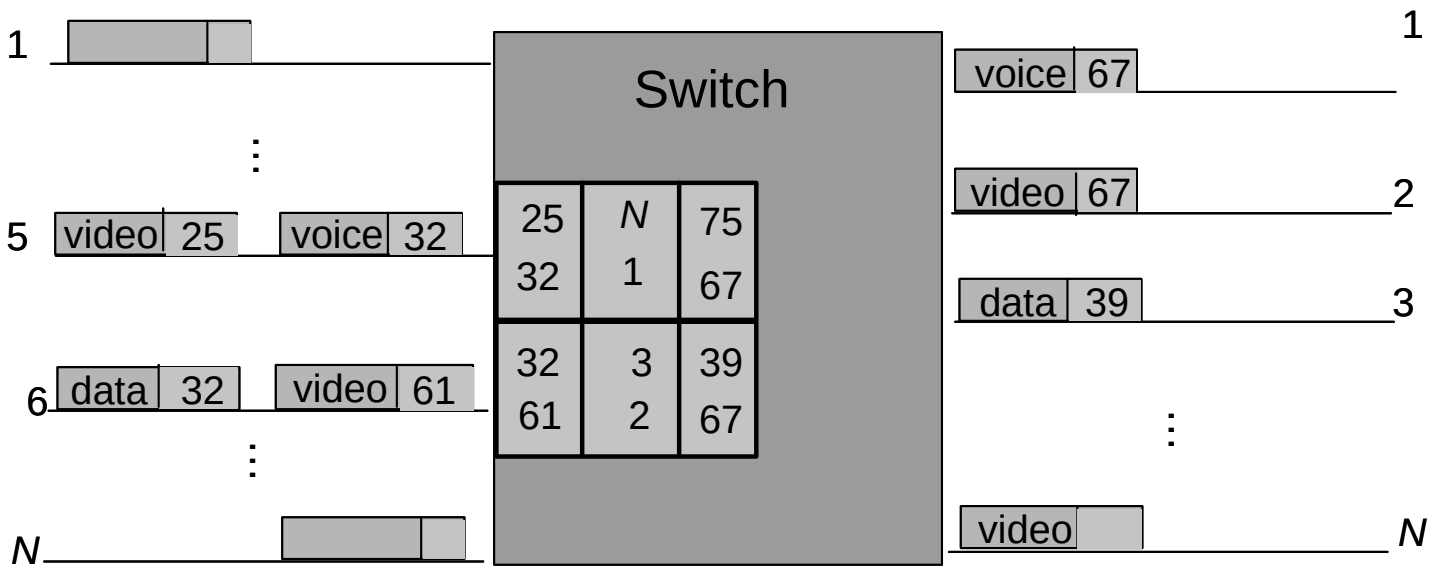

ATM Virtual Circuits

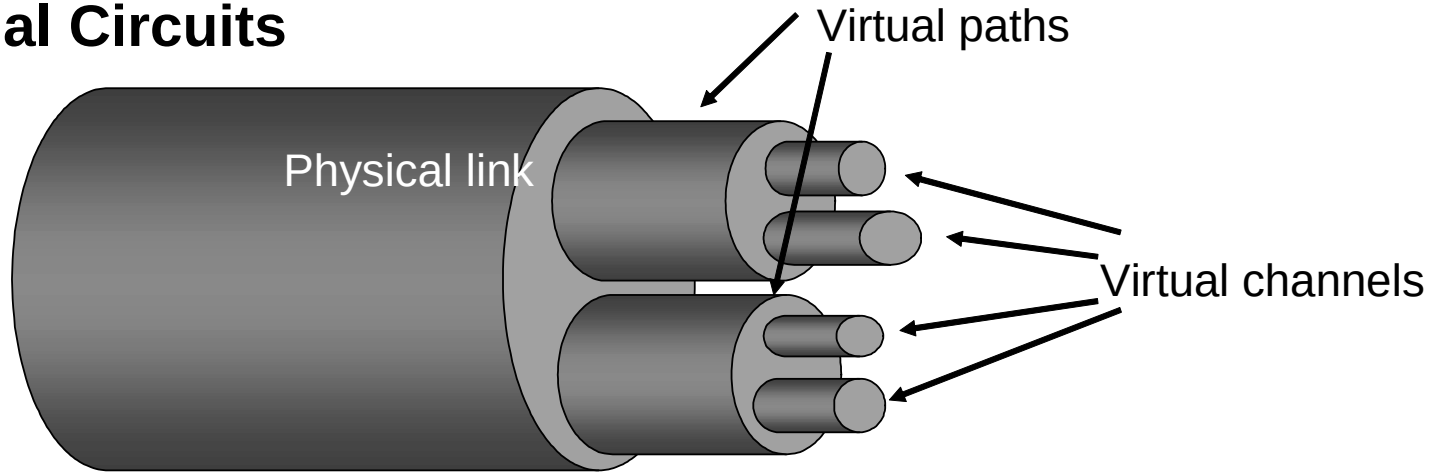

- Virtual channel connections (VCC) sharing a common path through the network are aggregated into *virtual paths* (VP)
  - Similar to SONET tributaries but with flexible/variable bit rates
- ATM cells have two levels of identifiers
  - Virtual Channel Identifier (VCI) and Virtual Path Identifier (VPI)
- Switching on the basis of the VPI first, if no VP exists, switch on VCI
  - VCIs only used at the end of a virtual path
- The VCI/VPI structure supports scalability to very large networks
- Network managers allocate bandwidth arbitrarily by controlling the scheduling priorities of individual VPs, VCs

ATM Network Interfaces

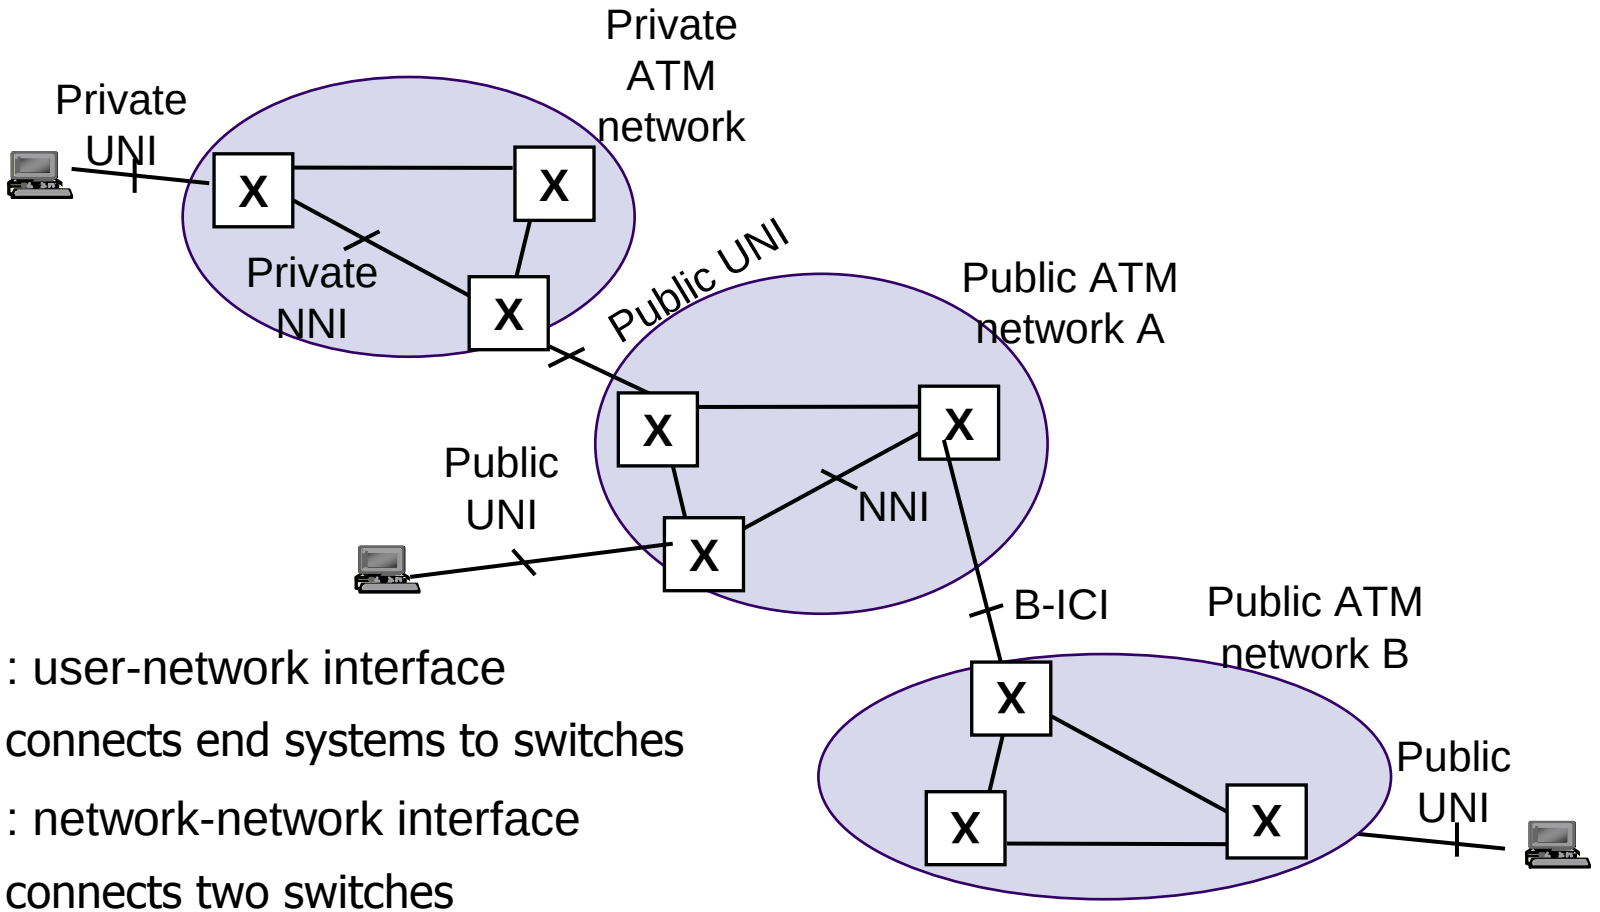

- UNI : user-network interface
  - » connects end systems to switches
- NNI : network-network interface
  - » connects two switches
- B-ICI : broadband intercarrier interface
  - » connects two public switches from different service providers
- ATM provides Permanent Virtual Connections (PVCs) e.g. leased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

ATM cell

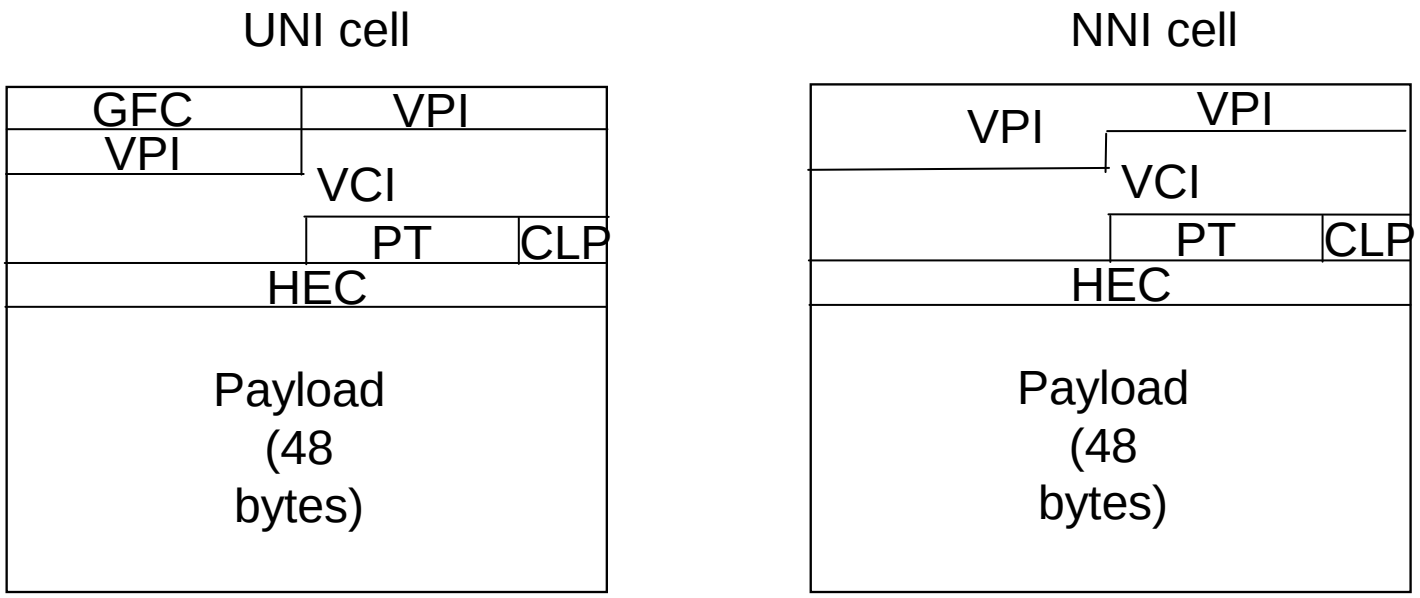

- GFC (Generic Flow Control) – not used in practice
- No GFC in NNI header allows larger trunks between public ATM switches
- VPI : 8 or 12 bits; VCI : 16 bits
- PT (Payload Type) : payload contains user data/control data
- CLP (Cell Loss Priority) : lower priority cells get discarded first when congested
- HEC (Header Error Control) : CRC-8 (x^8+x^2+x+1) over first 4 bytes of header
  - » detection and/or single bit correction

BISDN Reference Model

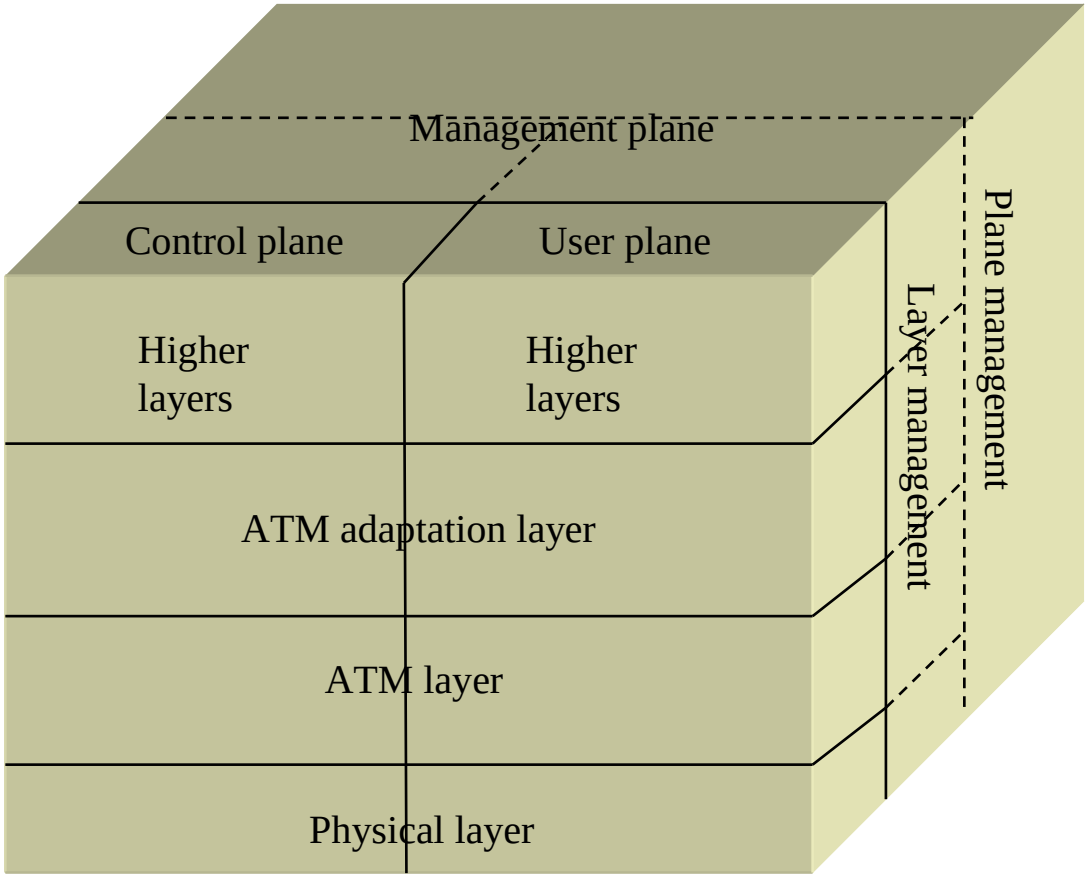

- User plane deals with data transfer, flow control, error recovery etc.
- Control plane deals with signalling to set up, manage and release connections
- Management plane deals with network resources and coordination of other planes

User plane ATM layers

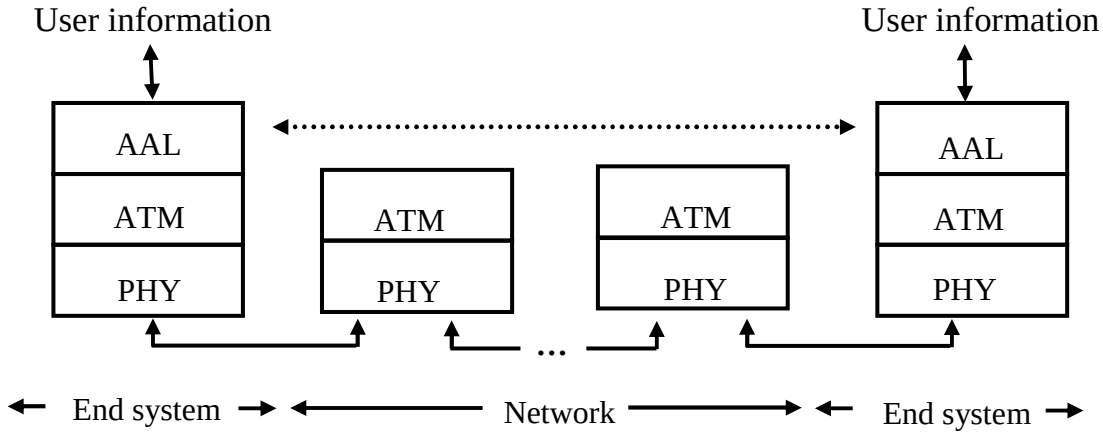

- Physical layer:
  - Transmission convergence sublayer
    - » ATM cell boundaries tracked
    - » checking of header checksums
    - » insertion and removal of idle cells
  - Physical medium dependent layer
    - » cells are converted into a bitstream
    - » cells packaged into frames for the medium e.g. SONET, FDDI, STP etc.

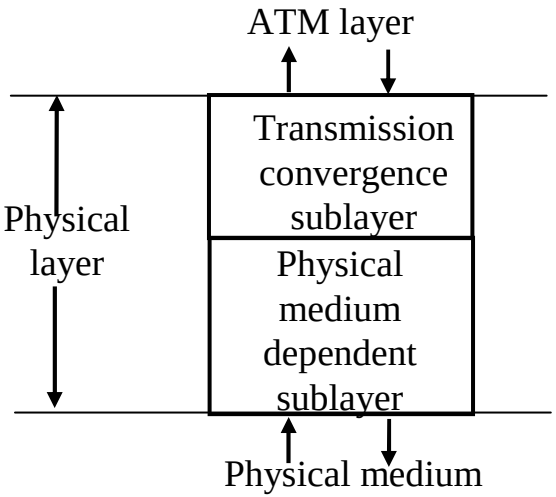

- ATM layer:
  - Concerned with sequenced transfer of cells in connections across the network
  - Accepts 48-byte blocks from the ATM Adaptation layer and adds 5-byte header

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

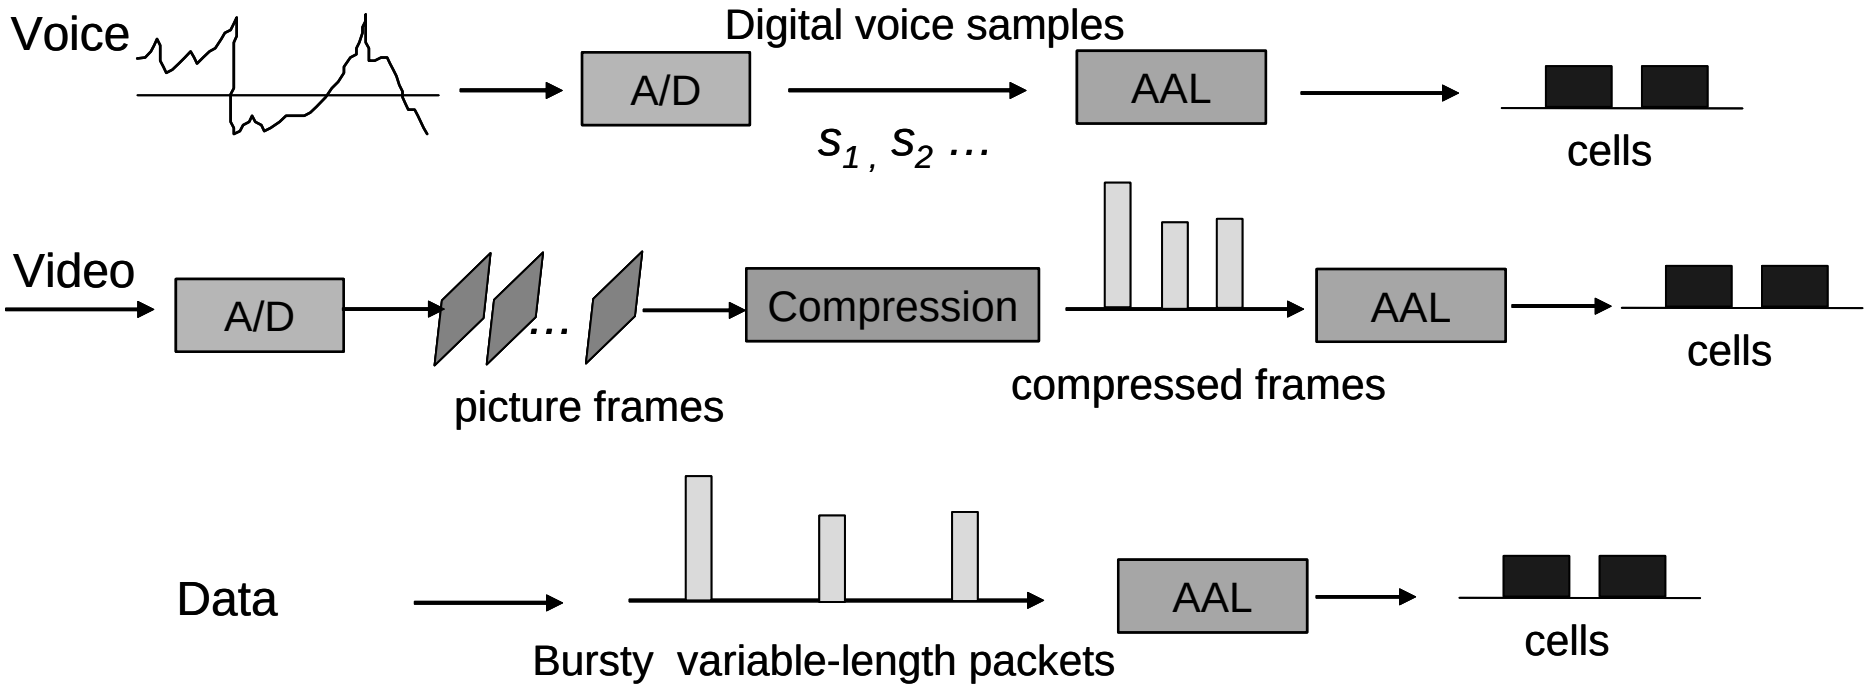

- Provides support for the protocol layer directly above
  - » e.g. to provide a reliable stream service, if needed (not if TCP was the layer above)
  - » cell sequence numbering
  - » segmentation and reassembly support

ATM Service Categories

- 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, compressed video
- 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

ATM Traffic Descriptors

- Parameters specified by the source 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

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

Quality of Service Parameters

Six QoS parameters are defined

Three are intrinsic to network performance and are *not* negotiated during connection setup:

- *Cell error ratio:* fraction of delivered cells that contain bit errors
- *Cell mis-insertion ratio:* average number of cells/second that are misdelivered
- *Severely errored cell block ratio*: ratio of severely errored cell blocks to total transmitted cell blocks
  - *Severely errored cell block* – M or more out of N cells are lost, in error, or misdelivered. M, N defined by network provider

Negotiable QoS 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 (bounded by CLR)

Negotiable QoS Parameters

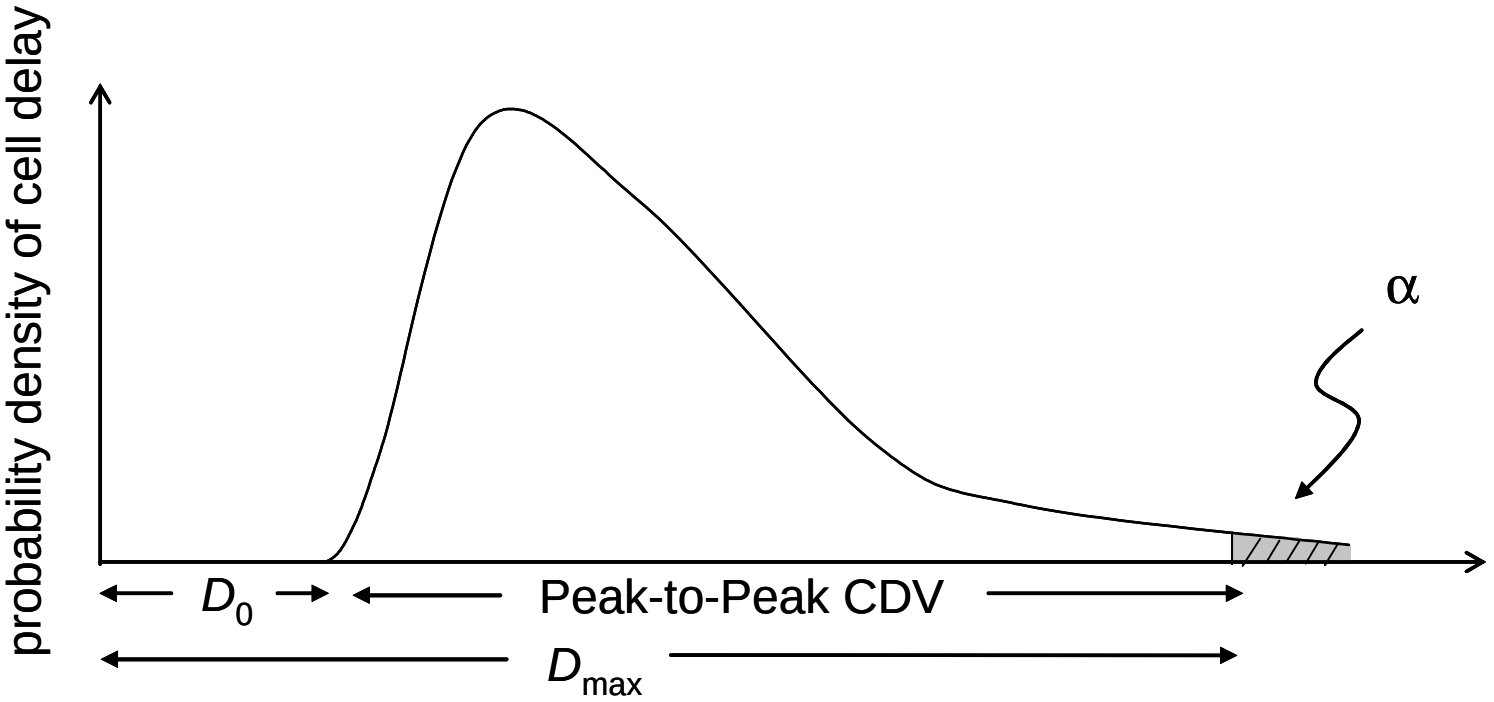

- Cell Delay Variation (CDV)
  - The total delay encountered by cells in a connection
    - » excluding fixed delay D_0
  - Peak-to-peak delay can currently be negotiated
  - Network switches have only limited control of the variance of CTD values
    - » so range of negotiable CDV value is also limited

Relevant Parameters for Service Classes

	Class of Service	CBR	VBR-NRT	VBR-RT	ABR	UBR
QoS	CLR	✓	✓	✓	✓	✗
	CTD	✓	✗	✓	✗	✗
	CDV	✓	✗	✓	✗	✗
Traffic descriptor	PCR	✓	✓	✓	✓	✓
	SCR	✗	✓	✓	✗	✗
	MBS	✗	✓	✓	✗	✗
	flow control	✗	✗	✗	✓	✗

Traffic Contract, Call Admission, Policing

- *Traffic contract*: includes the ATM service category, the traffic descriptors, and the QoS requirements
- *Connection Admission Control* (CAC): accept or reject connection request?
  - 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: *Usage parameter control (UPC)* is the process of enforcing the traffic agreement at the UNI
  - QoS guarantees are valid only if the user traffic conforms to the connection contract
  - a Generic Cell Rate Algorithm (GCRA) defined
    - » equivalent to the leaky bucket scheme
  - Non-conforming cells can be tagged (CLP=1) or dropped

ATM Signalling

- Means for dynamically setting up and releasing connections in ATM
- Involves message exchange across:
  - User-Network-Interface
  - Network-Network Interface
  - Broadband Inter-Carrier Interface
- Signaling messages use the services of the SAAL layer in the control plane

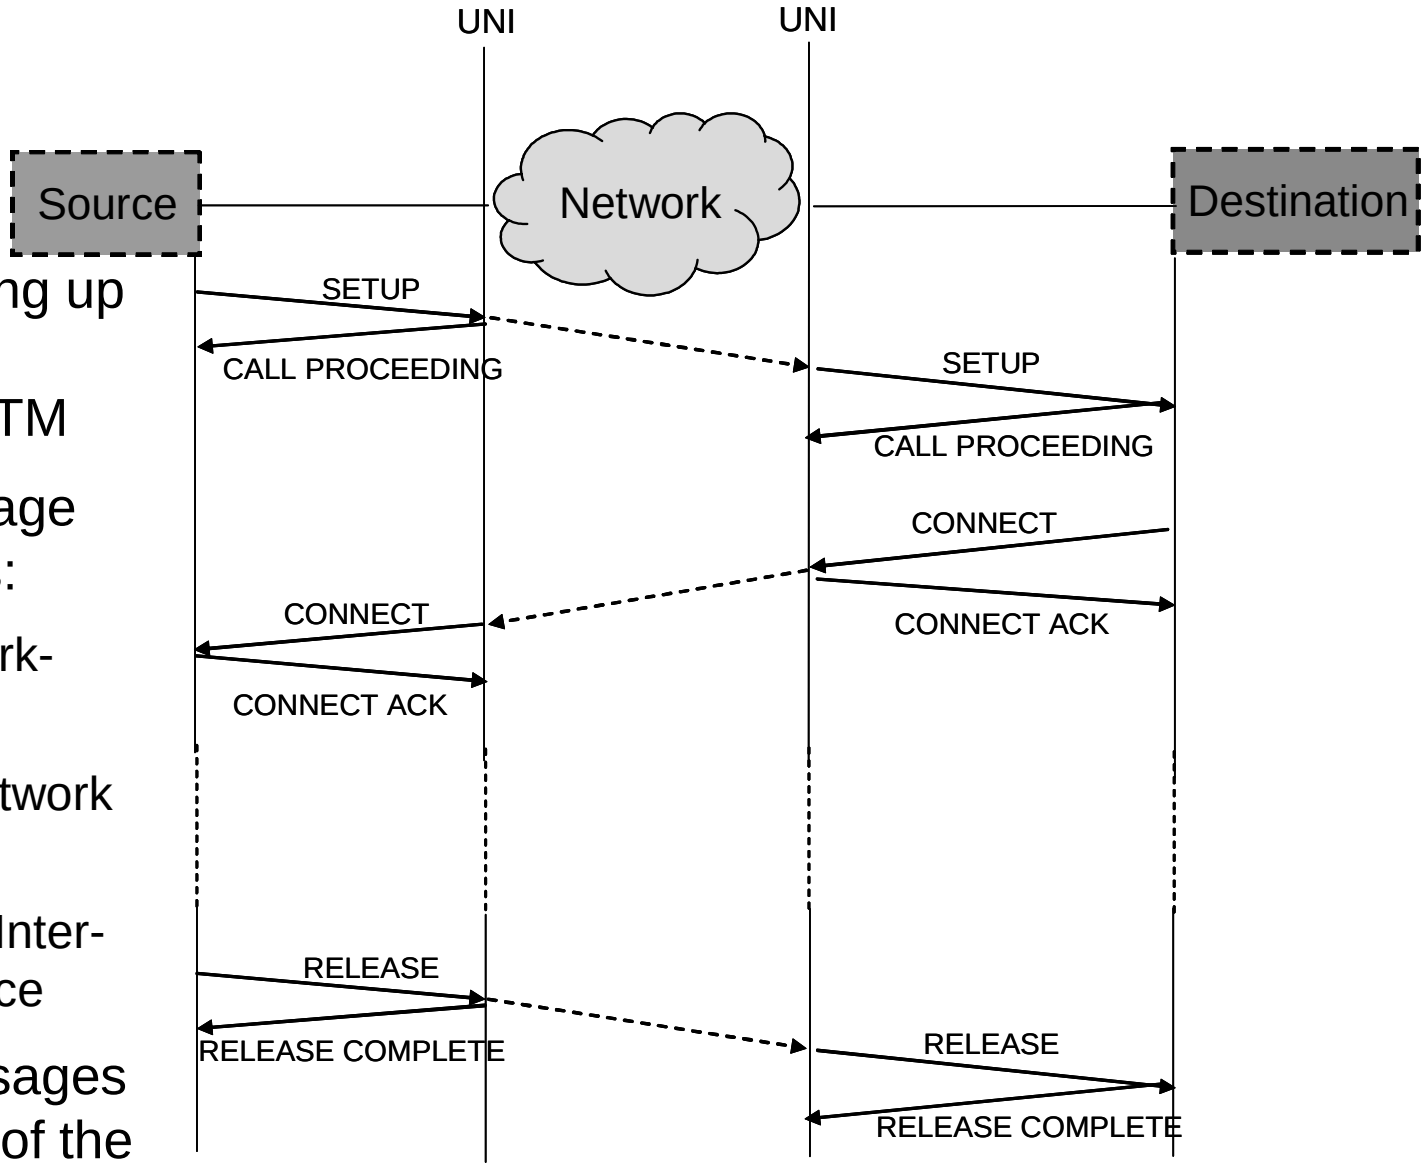

PNNI Routing

- ATM Forum developed PNNI for use
  - between private ATM switches (Private Network Node Interface)
  - between group of private ATM switches (Private Network-to-Network Interface)
- Routing: exchange of network topology information between nodes
  - Required for setting up connections
  - Not for forwarding cells
- PNNI routing
  - For intra-domain and inter-domain routing
  - Link-state approach: each node has network topology
  - Introduces hierarchy in the ATM network that provides a switch:
    - » Detailed routing information in its immediate vicinity
    - » Summary information about distant destinations

PNNI Routing Hierarchy

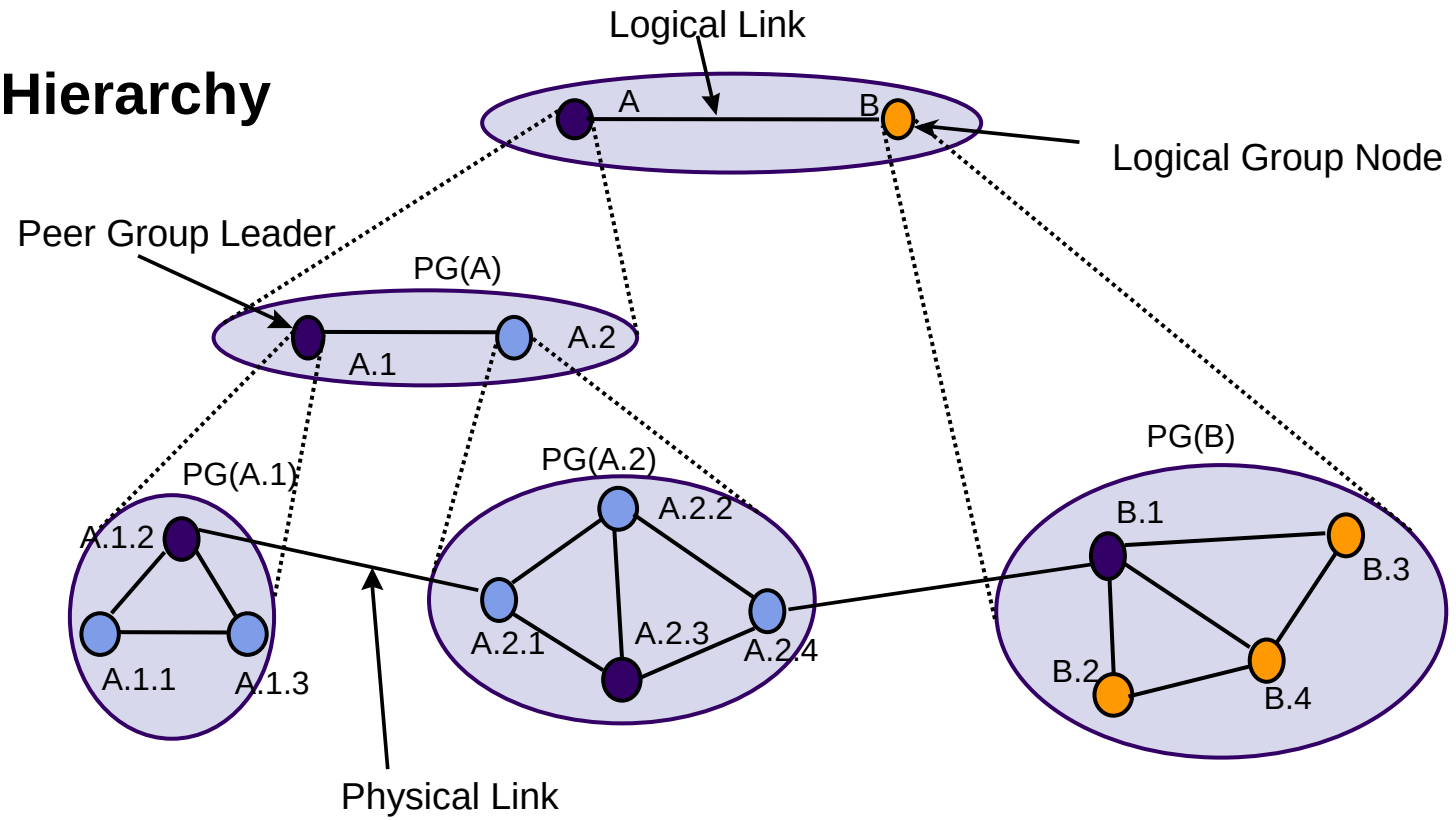

- *Peer Group*: collection of nodes that maintain an identical view of the group
- *Logical Group Node:* abstract representation of a peer group at a higher level in the routing hierarchy
- *Peer Group Leader*: node in peer group that executes functions of LGN for the PG
  - Summarizes topology info within the PG
  - Injects summary info into higher order groups and into the PG

PNNI Routing

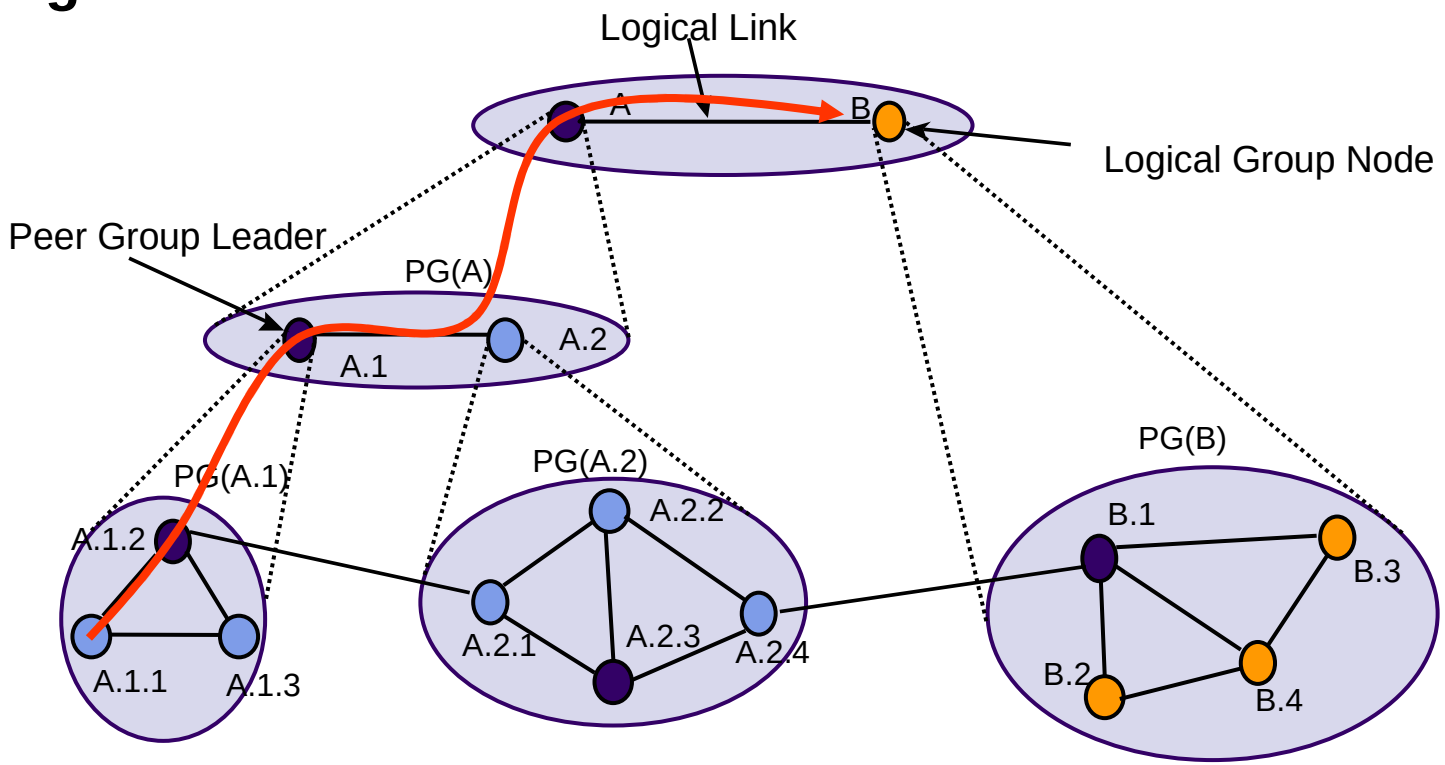

- PNNI source node specifies entire path across its PG using *designated transit list (DTL)*
- Rest of path specified using higher levels in the hierarchy
- Example: station in A.1.1 requests path to B.3
- Path: (A.1.1, A.1.2, A.2, B)