

Asynchronous Transfer Mode

By Scott Midkiff
ECpE/CS 5516, VPI
Spring 1997

(modified by Marc Abrams for Spring 1998)

Welcome to ATM!

- Loooooooooooooooooooooots of acronyms!
- A lot of what's in ATM came from the phone and switching worlds, not the LAN world, so the terminology may be unfamiliar.
- OSI model doesn't really apply!

Asynchronous Mode Transfer

- ATM (asynchronous transfer mode) is ...
 - Something of a revolution in both local and wide area networks
 - ◆ Innovations for LANs: scalable, switching-only, quality-of-service guarantees
 - ◆ Innovations for WANs: scalable, cell (packet)-switched, bandwidth-on-demand
 - Both a type of switching and a set of standards
 - Becoming widely used, but is not without real or proposed competition
 - ◆ 100-Mbps and 1-Gbps Ethernet
 - ◆ 100-Mbps FDDI
 - ◆ DS3, SMDS, Frame relay services

WAN Transmission: Digital Hierarchy

North America and Japan* (rate in Mbps)

<u>Name</u>	<u>Carrier</u>	<u>Rate</u>	<u>Channels</u>
DS0		0.064	1
DS1	T1	1.544	24
DS1C	T1C	3.152	48
DS2	T2	6.312	96
DS3	T3	44.736	672
DS4	T4	274.176	4032

* a similar hierarchy exists in Europe (CEPT n , E n)

WAN Transmission: Fiber Optics and SONET

- Fiber optic transmission
 - AT&T completed first transcontinental fiber system in 1986
 - Nearly all high-density trunks are now fiber
- SONET (Synchronous Optical NETwork) hierarchy

<u>Name</u>	<u>Rate (Mbps)</u>	<u>Equivalent</u>
OC-1	51.84	28 DS1 or 1 DS3
OC-3	155.52	3 OC-1
OC-12	622.08	12 OC-1 or 4 OC-3
OC-48	2.488	48 OC-1
OC-192	9.953	

WAN Transmission: ISDN (1)

- Integrated Services Digital Network (ISDN) provides two standard types of service
 - Basic rate (basic service)
 - Primary rate (primary service)
- *Basic Rate Interface* (BRI): 2B+D
 - Delivers ISDN services over standard a twisted-pair telephone line to subscribers
 - Two 64 Kbps B channels -- Voice, circuit-switched connection, access to a packet switched network
 - One 16 Kbps D channel -- User packet data at 16 Kb/s and call control signals

WAN Transmission: ISDN (2)

- *Primary Rate Interface (PRI)*
 - Delivers ISDN services to digital PBXs, LANs, hosts, ...
 - Divides 1.544 Mbps T1 into:
 - ◆ 23 B channels, each at 64 Kbps
 - ◆ One D channel for messaging, also at 64 Kbps

B-ISDN Concept (1)

- B-ISDN: Broadband Integrated Services Digital Network
 - Designed to provide similar services as ISDN
 - B-ISDN is substantially more capable than ISDN

- B-ISDN versus ISDN (or N-ISDN for narrowband ISDN)
 - N-ISDN uses existing telephone network (copper pairs); B-ISDN uses optical fiber
 - N-ISDN is primarily circuit switched; B-ISDN uses only packet switching (specifically ATM switching)
 - N-ISDN uses predefined channel rates; B-ISDN uses virtual channels without prespecified rates (bandwidth on demand)

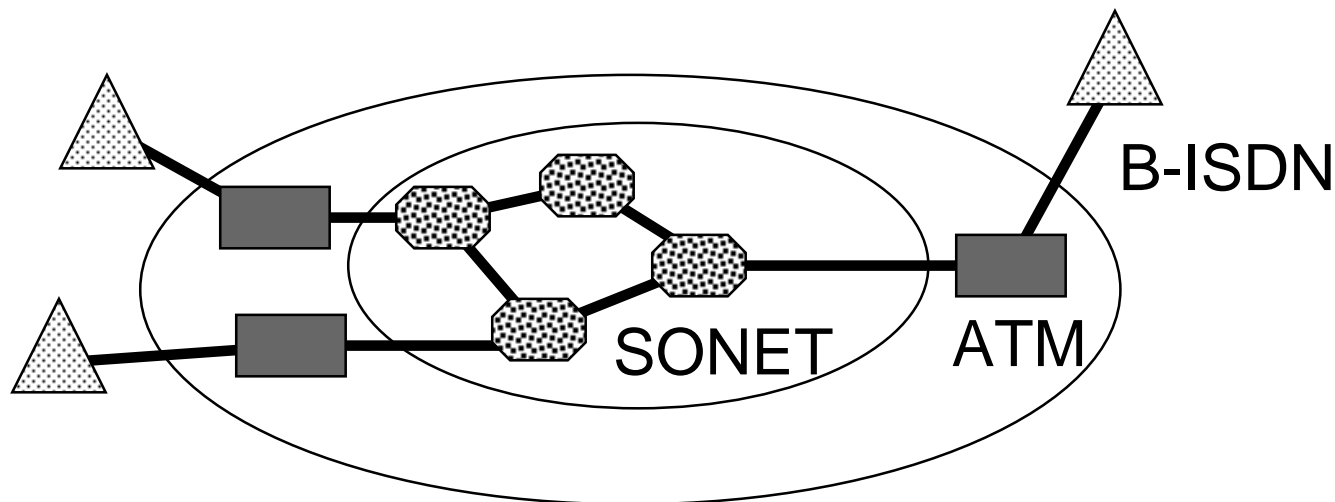
B-ISDN Overview (2)

- Integrated network for data, voice, video
 - Single system needed for the local loop
 - Single system for WANs and LANs
 - Economies of scale for collection of services
 - However, must support varied traffic types
 - Many political, economic, social issues

- Changes from traditional data and voice networks
 - Low processing complexity critical for high data rates
 - Voice/video require fixed bit rate/delay, but data networks have varying delays/data rates
 - Data rates may vary (1 bps to 100's of Mbps); but voice networks support constant data rates

B-ISDN Components

- *B-ISDN* services to users
- *Asynchronous Transfer Mode* (ATM) for switching
- *Synchronous Digital Hierarchy* (SDH) or *Synchronous Optical Network* (SONET) for transmission



SONET

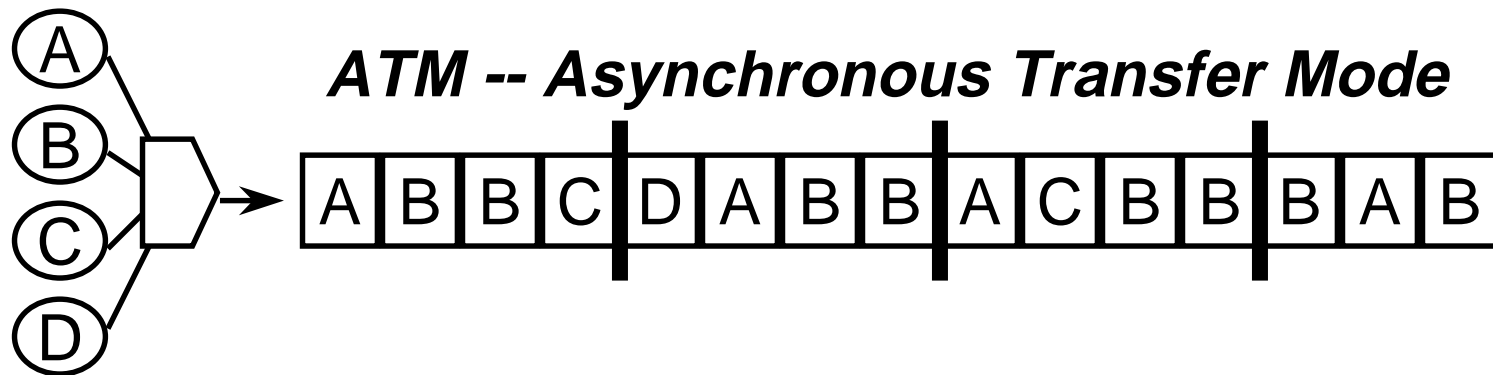
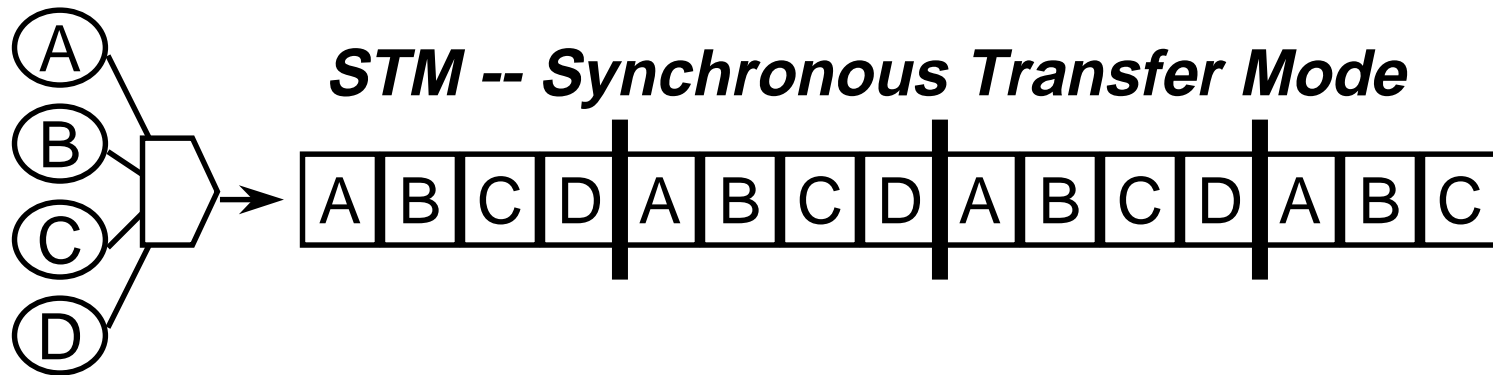
- A synchronous protocol (unlike ATM or packet switching)
- A physical protocol (so it's can be used under ATM)
- Used for optical fiber
- Frame of 810 bytes
- Doesn't use "headers" as we're used to.
- Does use "pointers" to demultiplex traffic sent over one fiber path.
- Circuit switched: Allows a stream of ATM packets to be sent long distances purely at physical layer without ATM switches in the path

ATM Overview

- ATM now exceeds BISDN in relevance
 - ATM first deployed in the LAN environment
 - ATM WANs dominated by data and video applications

- ATM versus Synchronous Transfer Mode (STM)
 - ◆ STM is circuit-switched; ATM is packet-switched
 - ◆ STM would require all channels to use some standard multiple of 64 Kbps; ATM is more flexible
 - High complexity if many standard rates
 - Inefficiencies if few standard rates
 - ◆ ATM benefits from high-speed packet switching -- very large scale integration (VLSI) circuits
 - Cell processing
 - Switching

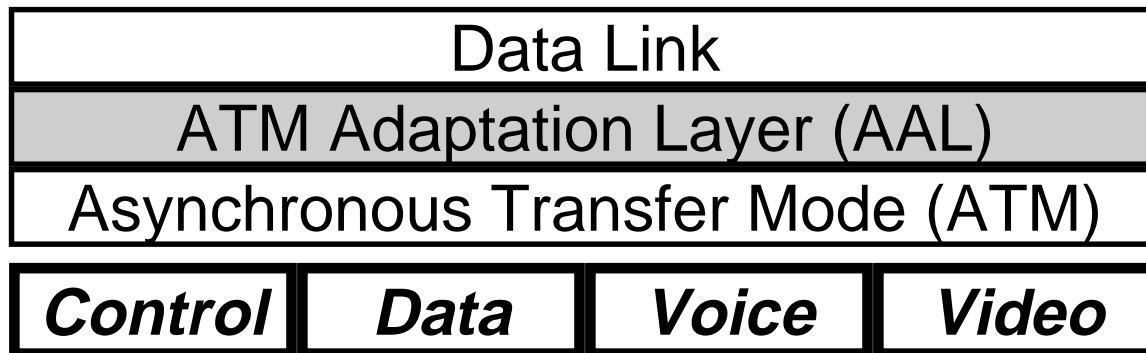
ATM versus STM



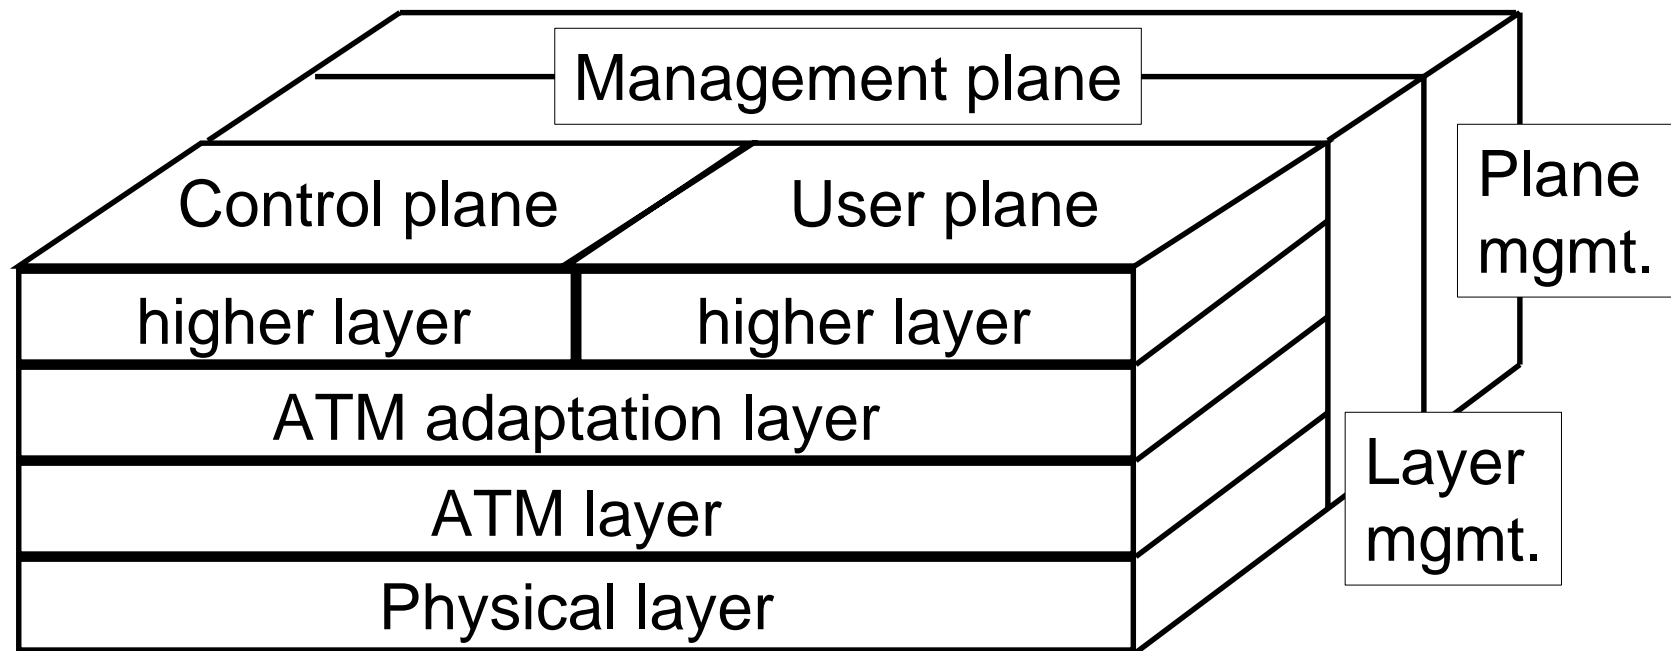
ATM Functionality

- ATM is difficult to map onto the OSI Reference Model
 - Performs layer 3 functions, such as switching
 - Used as a data link layer in some applications

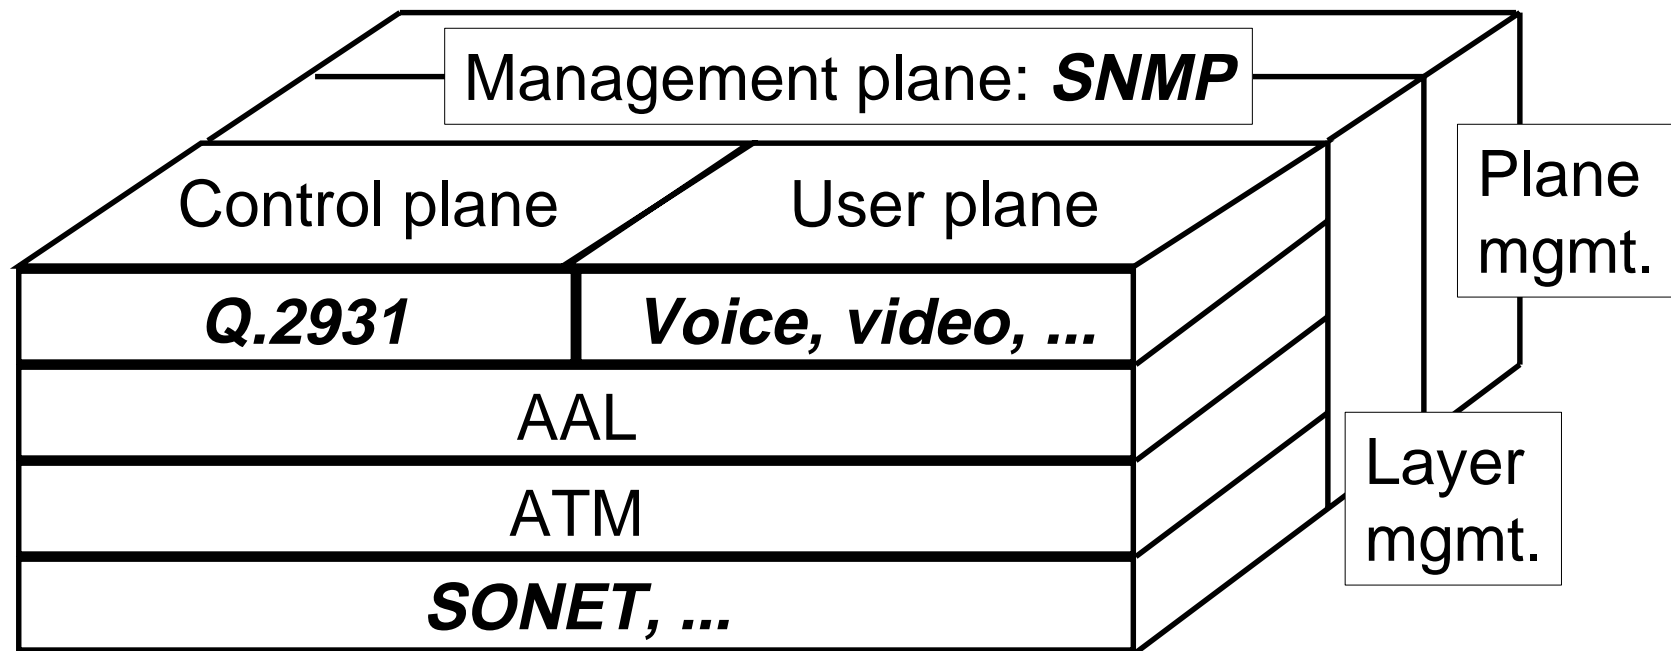
- An *Adaptation Layer* above ATM
 - at sender, divides messages or bit streams from above into ATM *cells* (packets)
 - at receiver, reconstructs messages or streams
 - thus performs some layer 4 functions



B-ISDN Protocol Reference Model



Reference Model: Another View



Q.2931: for call setup/release

Management plane: uses management protocol like SNMP

ATM Basics (1)

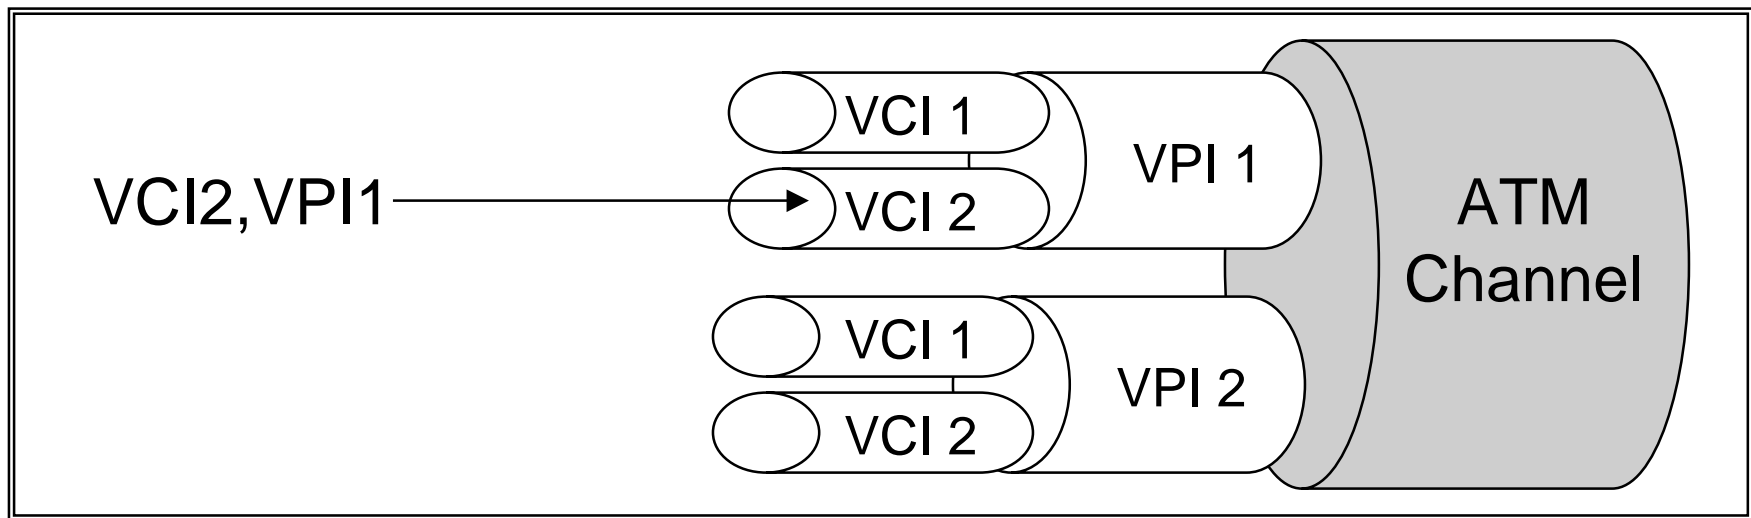
- ATM is a connection-oriented (virtual circuit) network
 - Permanent virtual circuit (PVC) -- connections and paths through the network are established when network is established
 - Switched virtual circuit (SVC) -- connections and paths through the network are established on an as-needed basis

- Connection-oriented eliminates need for
 - Source and destination addresses in header
 - Sequence number for resequencing

- Error & flow control are only done on an end-to-end basis if needed either by
 - application or
 - special “signalling”

ATM's Virtual Circuits

- Two multiplexing hierarchies:
 - Virtual Path Identifier (VPI):
 - ◆ Correspond to something that a switch switches
 - ◆ A bundle of Virtual Channels
 - Virtual Channel Identifier (VCI):
 - ◆ Correspond to sessions in an application
- Virtual Circuit is VPI + VCI

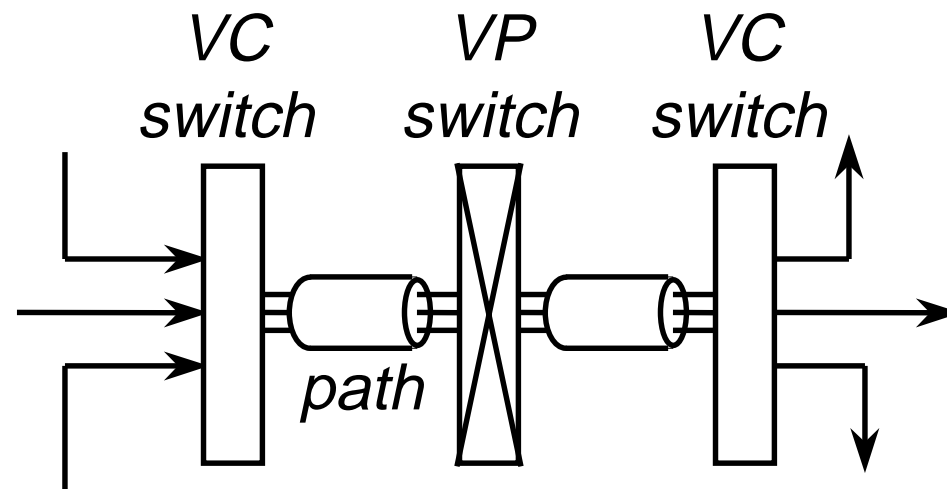


Virtual Channels and Virtual Paths (1)

- Virtual channel identifier (VCI) -- 16-bit field in header
 - Identify a virtual channel on a link between two ATM switches
 - Up to $2^{16} = 64K$ different channels can be carried over one link

- Virtual path identifier (VPI) -- 8- or 12-bit field in header
 - Identify a path over which the VCI does not change
 - Used for semi-permanent connections
 - Up to $2^8 = 256$ or $2^{12} = 4K$ paths can be carried over one link
 - There can be 64K VCIs over one VPI

Virtual Channels and Virtual Paths (2)



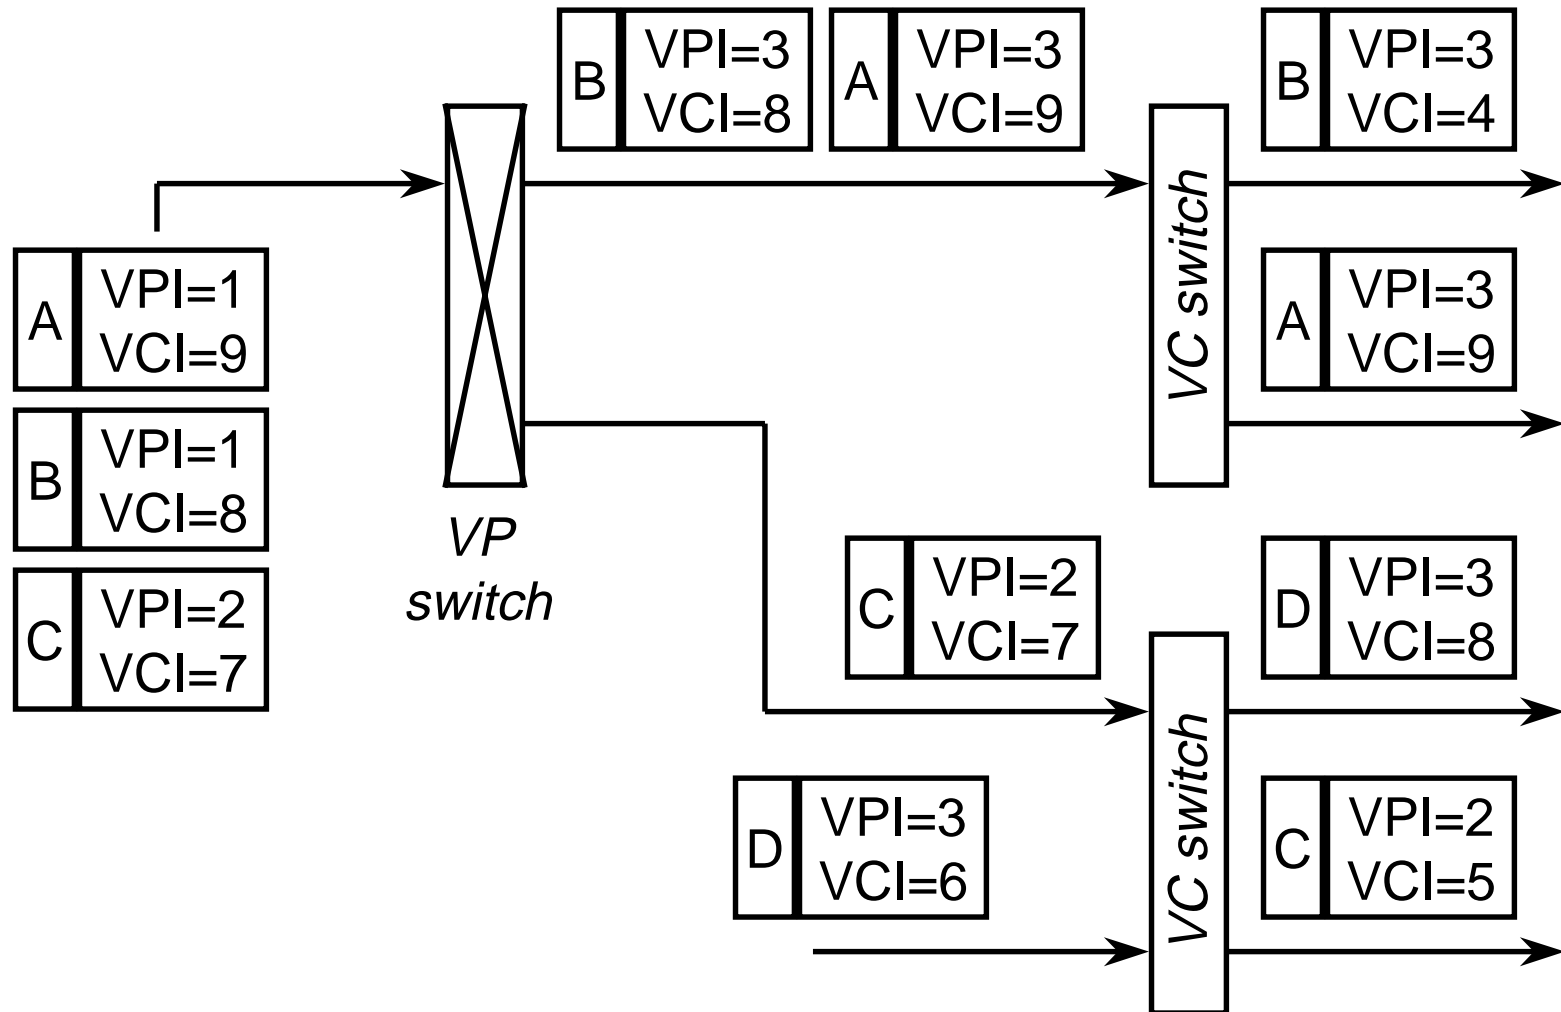
Simple case:

- Same VCI is used at source, destination
- Multiple paths (hence VPIs are used)

Complex case:

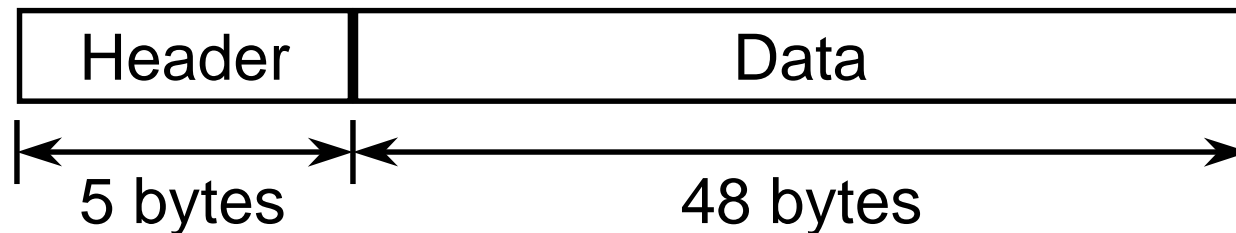
- Multiple VCIs are used (e.g., cross MCI/Sprint boundary)

Virtual Channels and Virtual Paths (3)



ATM Basics (2)

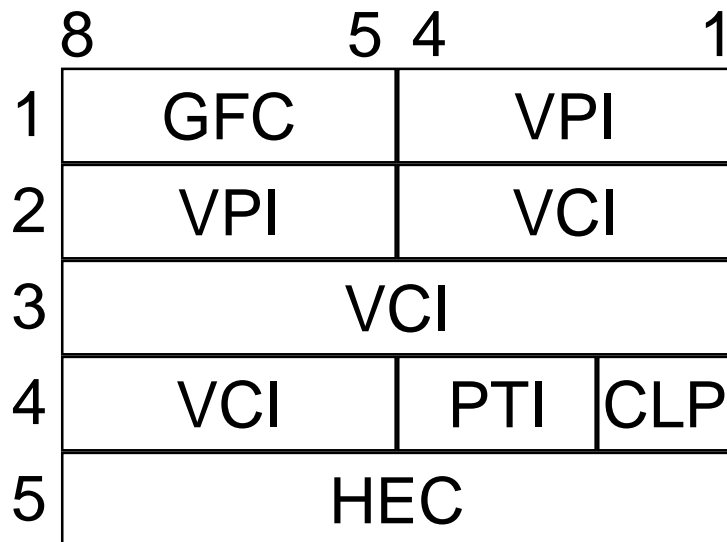
- ATM's basic unit of transfer is a fixed-length cell



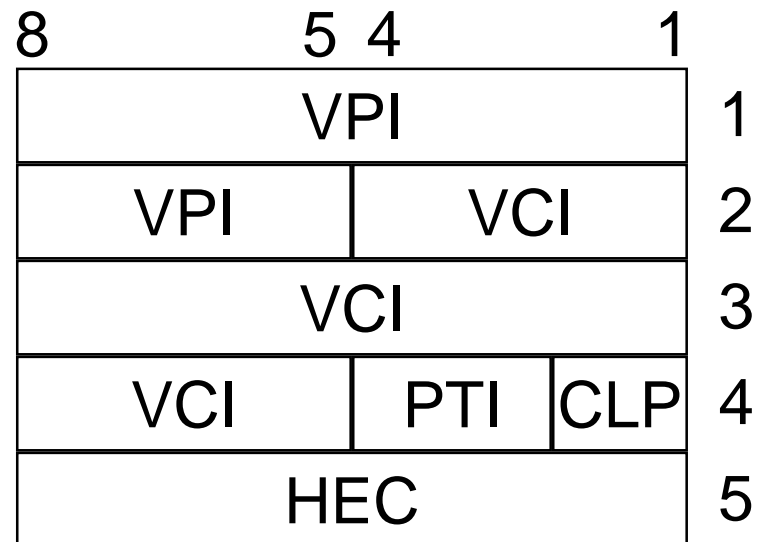
- Fixed cell size simplifies switches (processing in hardware rather than software)
- Small size minimizes packetization delay for voice transmission
- Larger cell size would be more efficient for data -- due to per packet processing and header overhead

ATM Header Format (1)

- There are two header formats:
 - User network interface (UNI): user to subnet
 - Network node interface (NNI): internal subnet



UNI Header Format



NNI Header Format

ATM Header Format (2)

- The address field (channel identifier) consists of two subfields:
 - Virtual channel identifier (VCI): 16 bits
 - Virtual path identifier (VPI): 8 bits in UNI format, 12 bits in NNI format

- VCI and VPI together identify the virtual connection
 - 24 bits in UNI format
 - 28 bits in NNI format

- A special reserved address indicates “unassigned” or “idle” cells which carry no data but are needed to fill up a cell slot in transmission

ATM Header Format (3)

- 4-bit Generic Flow Control (GFC):
 - Can be used by the user to multiplex data from multiple applications or devices onto the access link to the network

- 3-bit PTI: Payload type indicator:
 - user data versus network control information

- 1-bit CLP: Cell loss priority bit

- 8-bit Header error control (HEC):
 - CRC check over the header (does not include data)
 - Not only detects, but corrects single bit errors
 - One bit error in VCI is deadly!

Reference Model -- Layers and Sublayers

Convergence	CS	AAL
Segmentation and reassembly	SAR	
Generic flow control Cell VPI/VCI translation Cell multiplex and demultiplex		ATM
Cell rate decoupling HEC header sequence generate/verify Cell delineation Transmission frame adaptation Transmission frame generate/recovery	TC	PHY
Bit timing	PM	
Physical medium		

CS:
*convergence
sublayer*

SAR:
*segmentation
and reassembly*

TC:
*transmission
convergence*

PM:
*physical
medium*

HEC: header
error checking

Quality of Service (QoS) and Cell Loss

■ Quality of service

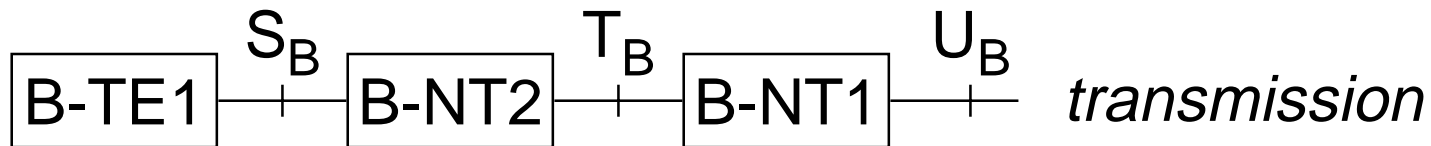
- Cell loss
- Cell delay
- Cell delay variation

■ ATM standard provides for a cell loss priority (CLP) bit in the header; allows a user to identify two levels of priorities for each cell

- CLP = 0 for high priority cells (e.g., voice)
- A switch can discard CLP=1 cell (e.g., data)
- Network can monitor high priority and low priority traffic and compare to negotiated bandwidth

Network Termination and Generic Flow Control

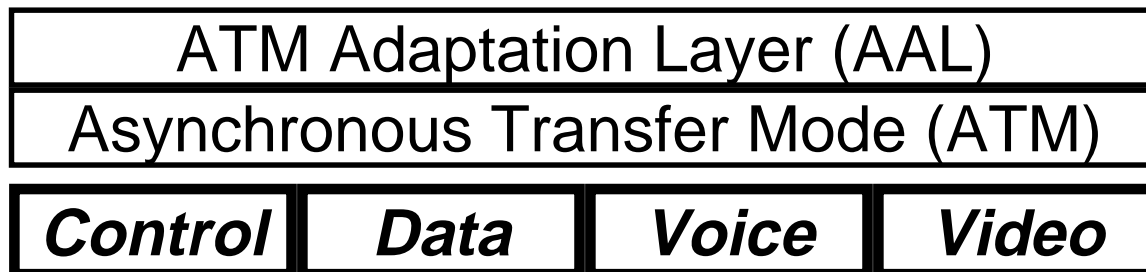
- Customer premises equipment, or Broadband Terminal Equipment-1 (B-TE1), connects to ATM network at Broadband Network Termination-1 (B-NT1) or Broadband Network Termination-2 (B-NT2)



- Multiple devices may share access to the S_B interface
- Each device should be ensured fair access to the interface
- Generic flow control is intended to support contention for the local interface

ATM Adaptation Layer -- AAL (1)

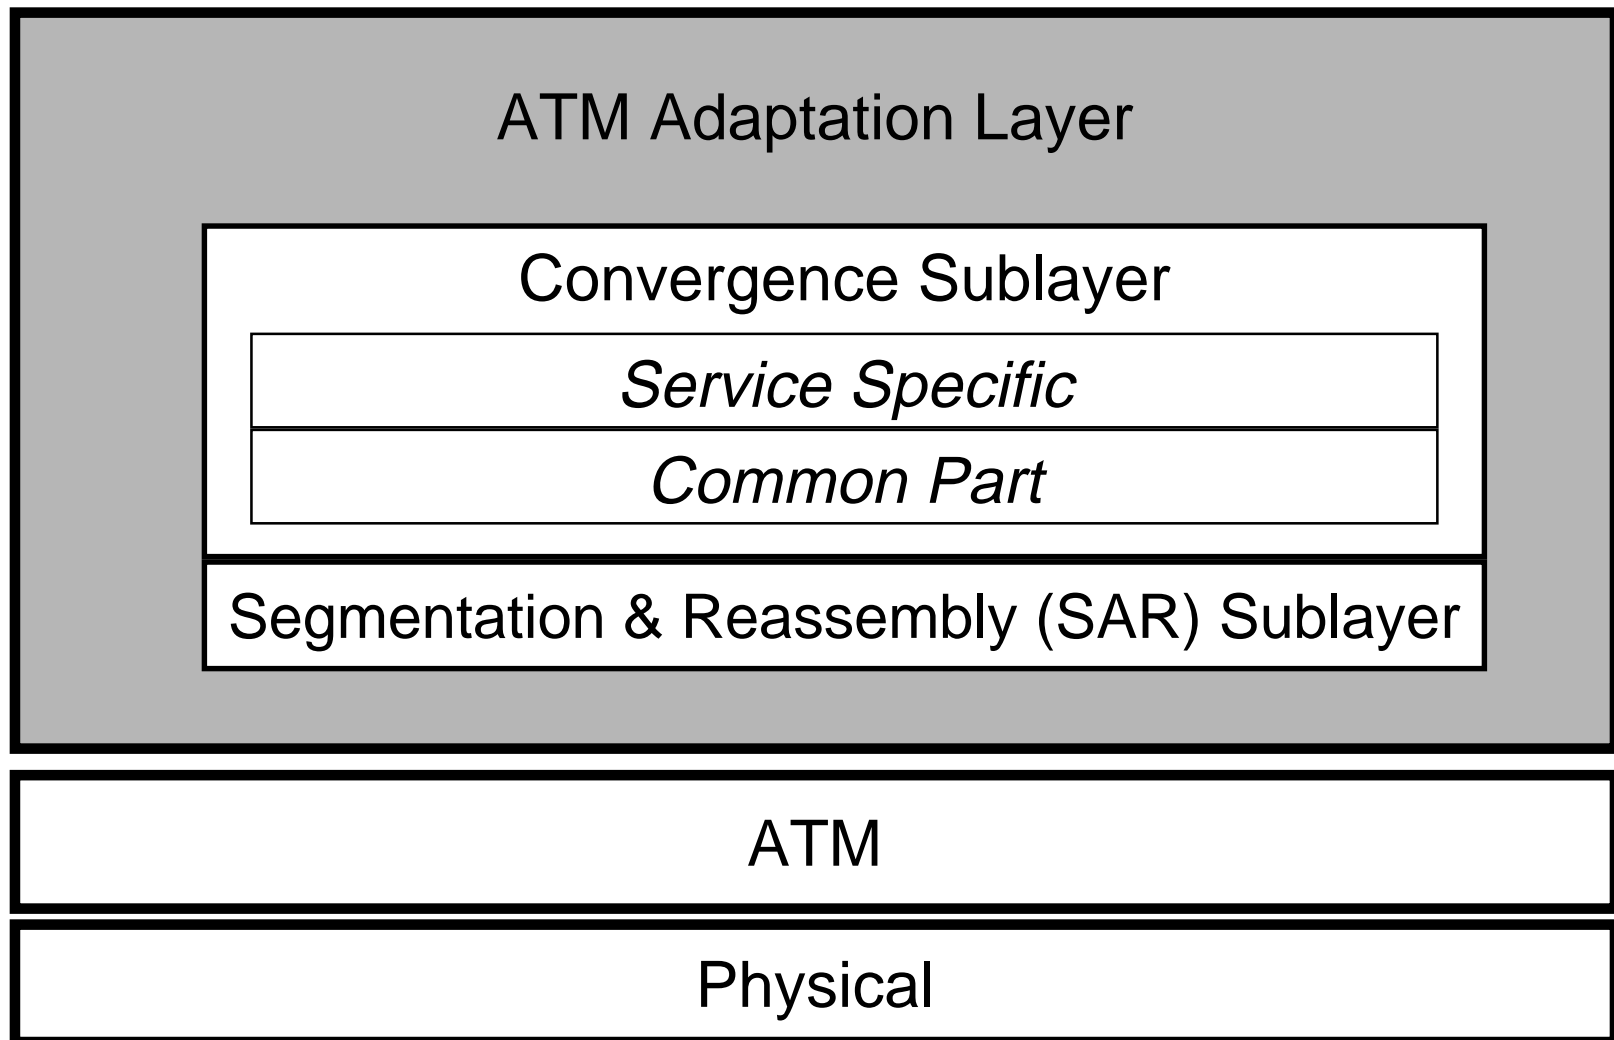
- The ATM Adaptation Layer (AAL) is responsible for breaking incoming source data into 48-byte pieces and reconstructing the data at the receiver
 - Used at entry and exit from network
 - If ATM is viewed as a network layer, then AAL is a transport layer



ATM Adaptation Layer -- AAL (2)

- Operation depends on the type of source traffic:
 - *Class 1* Constant bit rate (CBR) traffic,
e.g. 64 Kbps voice and fixed-rate video
 - *Class 2* Variable bit rate (VBR) traffic to be delivered
with fixed delay, e.g. compressed and
packetized voice, video
 - *Class 3/4* Non-real-time data in messaging or
streaming modes (may be connection-oriented)
 - *Class 5* Similar to class 3/4, but without multiplexing
or error detection (always connection-oriented)
- Other classes being defined, e.g. for MPEG-2 video

Adaptation Layer Sublayers (1)



Adaptation Layer Sublayers (2)

- Convergence Sublayer (CS)
 - Upper-layer frames are basic data units
 - Concerned with flow control and error recovery for Class 3 (connection-oriented traffic)

- Segmentation and Reassembly (SAR) Sublayer
 - Segments of upper-layer frames are basic data units
 - Concerned with segmenting frames at source and reconstructing frames at destination

Class 1 AAL

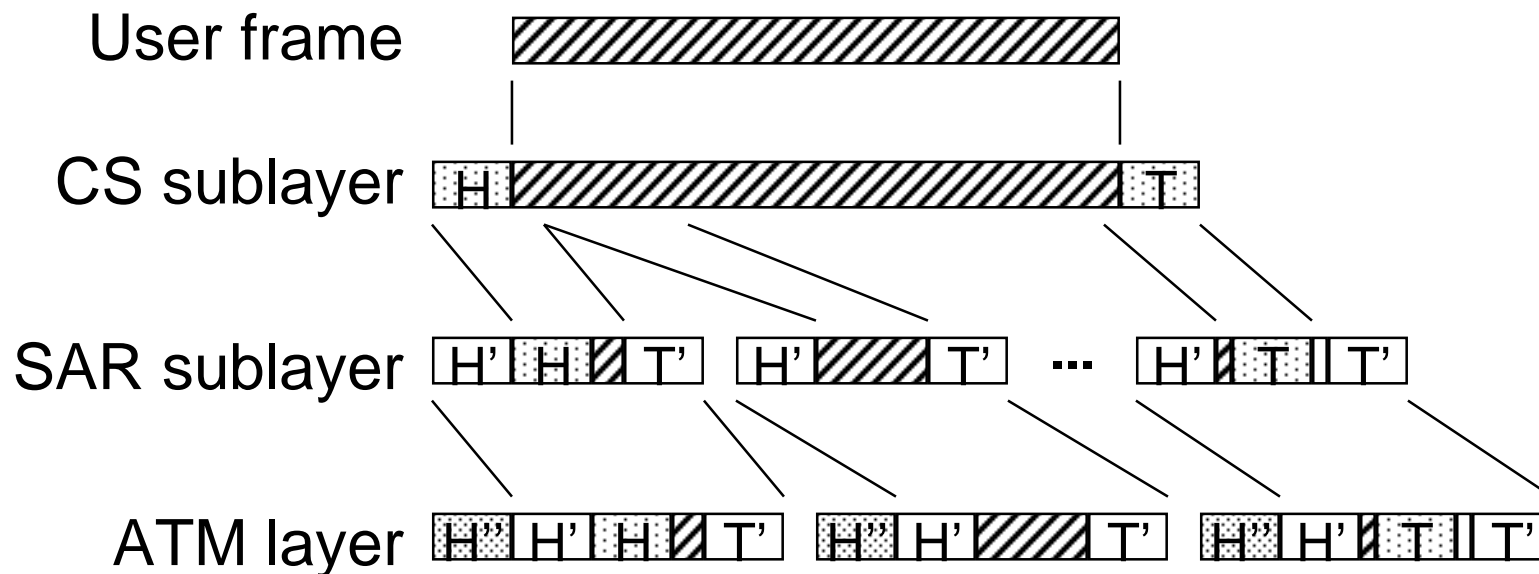
- Class 1 traffic is constant bit rate traffic and has no framing structure on input
- SAR sublayer must segment incoming data stream
- Occasional errors can be tolerated, but these need to be detected to maintain framing
- One-byte header (no trailer):
 - 1-bit convergence sublayer indicator
 - 3-bit segment number
 - 4-bit CRC

Class 2 AAL

- Class 2 traffic is similar to Class 1 traffic except that there is a frame structure on input
- Occasional errors can be tolerated, like Class 1
- Some Class 3/4 type functions needed to maintain input frame structure

Class 3/4 AAL (1)

- Non-real-time data
- Connection is coordinated by peer communication between the CS and SAR sublayers using header and trailers



Class 3/4 AAL (2)

- Segmentation and reassembly SAR sublayer
 - Two-byte header
 - ◆ *Segment Type* (2 bits): Indicates if this is first segment of a CS frame, middle segment, last segment, or an entire CS frame
 - ◆ *Sequence Number* (4 bits): Used to check for dropped or misdirected cells (cells delivered in order)
 - ◆ *Reserved/Multiplexing Identifier* (10 bits): May be used for multiplexing multiple user sessions on a single connection

Class 3/4 AAL (3)

- Segmentation and reassembly SAR sublayer (continued)
 - Two-byte trailer
 - ◆ *Length Indicator* (6 bits): For last segment and single segment SAR PDUs with less than 44 bytes of data
 - ◆ *CRC* (10 bits)

Class 3/4 AAL (4)

- Common part convergence sublayer -- CPCS
 - Header
 - ◆ *Common part indicator (CPI)* (1 byte): indicates unit for length and size fields (possible future uses for multiplexing identifier allocation and operations/maintenance)
 - ◆ *Beginning tag (Btag)* (1 byte): Sequence number placed in header and trailer
 - ◆ *Buffer allocation size indication (BAsize)* (2 bytes): buffer requirements for this CPCS-PDU

Class 3/4 AAL (5)

- Common part convergence sublayer -- CPCS (continued)
 - In PDU
 - ◆ *Padding (PAD)* (0 to 3 bytes): padding in header
 - Trailer
 - ◆ *Alignment (AL)* (1 byte): puts trailer on 32-bit boundary
 - ◆ *End tag (Etag)* (1 byte): matches Btag value
 - ◆ *Length (Length)* (2 bytes): length of CPCS-PDU

Class 5 AAL (1)

- Class 5 AAL is similar to class 3/4 AAL
 - Lower overhead
 - No error detection in the SAR (just at CPCS)
 - No multiplexing
 - No buffer allocation information (BA field in class 3/4)
 - Supports connection-oriented only -- class 3/4 AAL supports both connection-oriented and connectionless

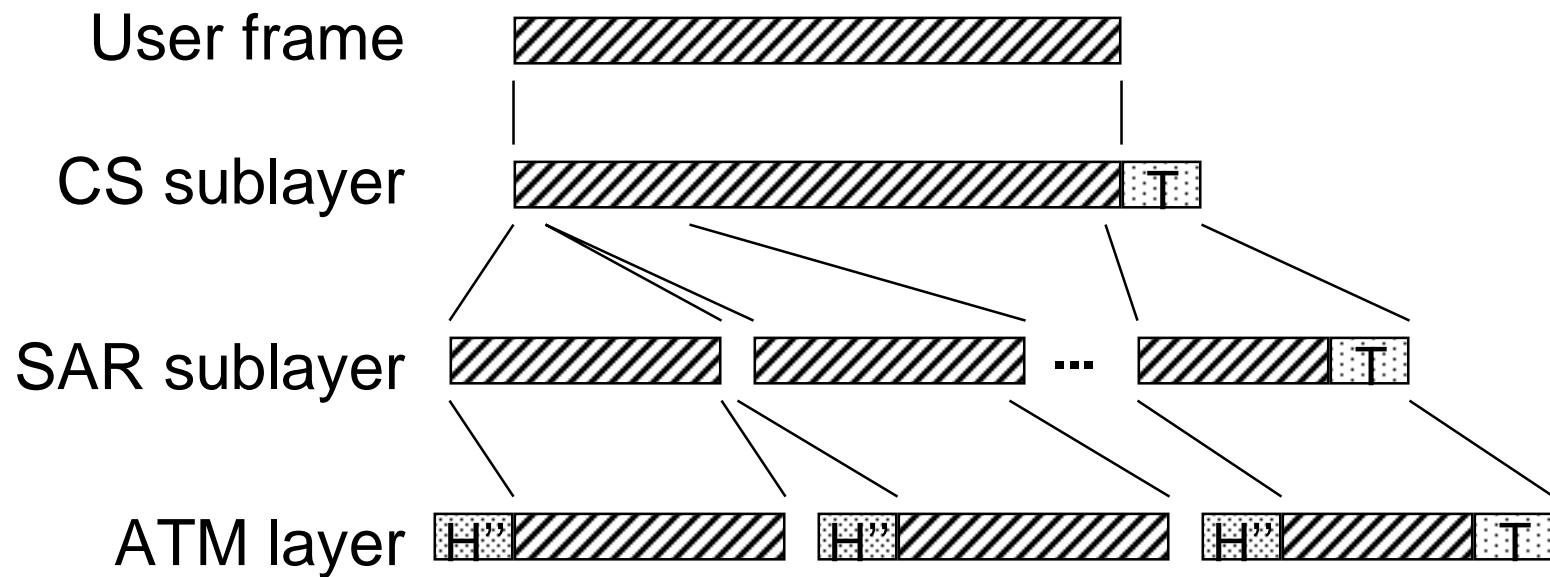
- Promoted by the ATM Forum to better match the needs of
 - LAN equipment manufacturers
 - High data rate, connection-oriented data service users

Class 5 AAL (2)

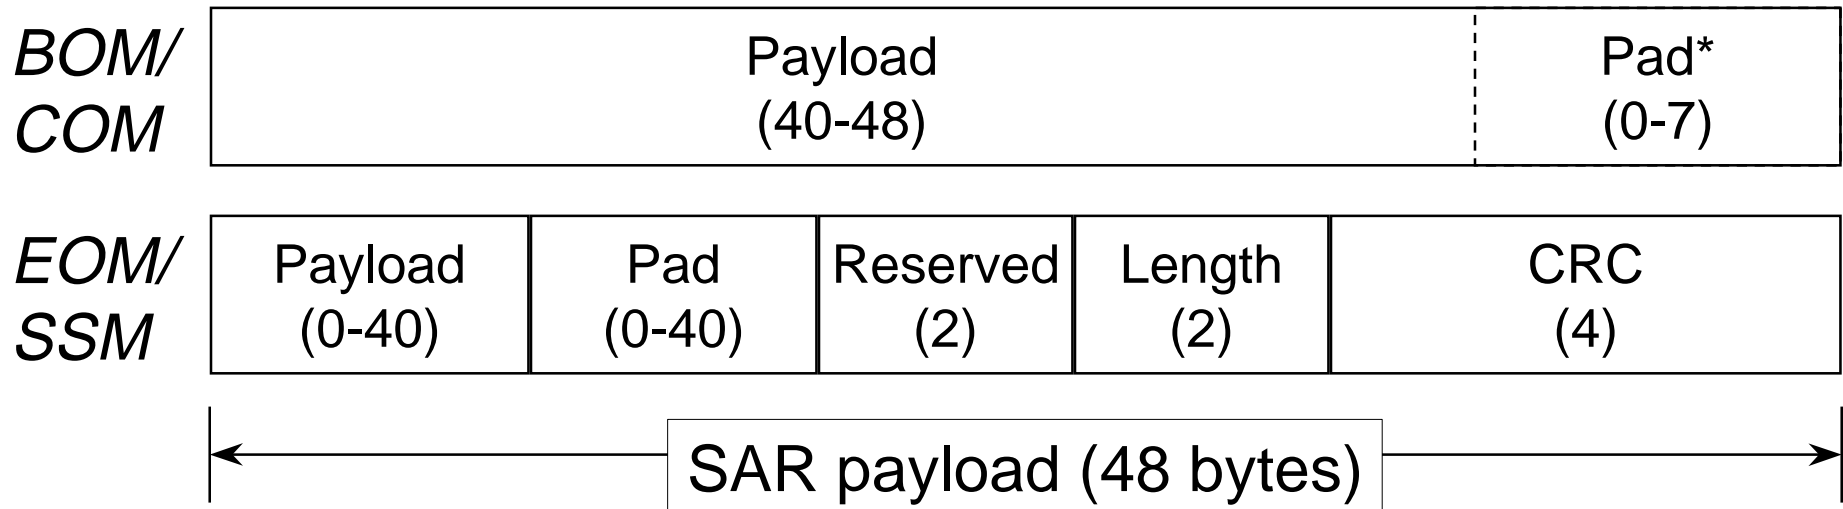
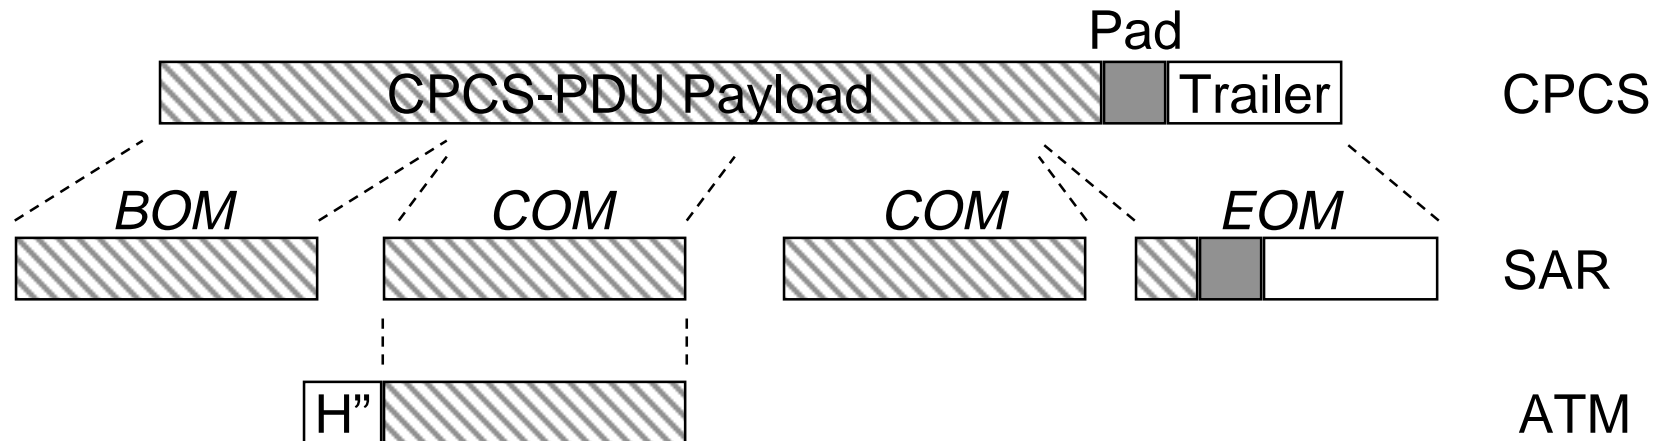
- The only SAR information for each SAR-PDU is
 - BOM/COM (beginning/continuation of message) -- AUU=0
 - EOM/SSM (end of message/single segment message) -- AUU=1

- AAU (ATM-user-to-ATM-user) indication is in one bit of the ATM header, specifically payload type (PT) field

Class 5 AAL (3)



Class 5 AAL (4)



Class 5 AAL (5)

- AAU indicator is in ATM header PT field
 - AAU = 0: BOM/COM
 - AAU = 1: EOM/SSM
- Length is number of bytes in CPCS PDU payload
- Reserved (2 bytes) may be:
 - CPCS User-to-User indication (1 byte)
 - Comon Part Indicator (1 byte)
- BOM/COM padded only on next to last SAR-PDU (i.e., next is EOM) and following SAR-PDU (the EOM) has no payload

ATM Congestion Control

- Even with high data rates supported by a B-ISDN network, it is inevitable that congestion will occur
- Three elements of congestion control:
 - 1) Rate, burstiness, and quality of service are negotiated by user and network per connection
 - 2) Network can monitor (at CS of AAL) rate and burstiness of user traffic
 - 3) Priority bit in ATM header (CLP) indicates cells that can be dropped

Negotiated Rate, Burstiness, QoS

- The user and network can agree upon a data rate, burstiness of data, and quality of service (QoS) to be provided by the network
- Effects of constant bit rate traffic is fairly easy to determine, however, the effects of variable bit rate traffic is a complex and open issue
- Required quality of service varies for broadband services
 - Delivery latency
 - Variance of interframe delay
 - Probability of frame loss
 - ...

Monitoring and Regulating Incoming Traffic

- Network must monitor incoming traffic to ensure that it meets agreed upon limits
- Could accept excess traffic, if bandwidth is available, but mark it so that it can be discarded in the event of congestion
 - E.g., using the cell loss priority (CLP) bit in the ATM header
- Other schemes possible

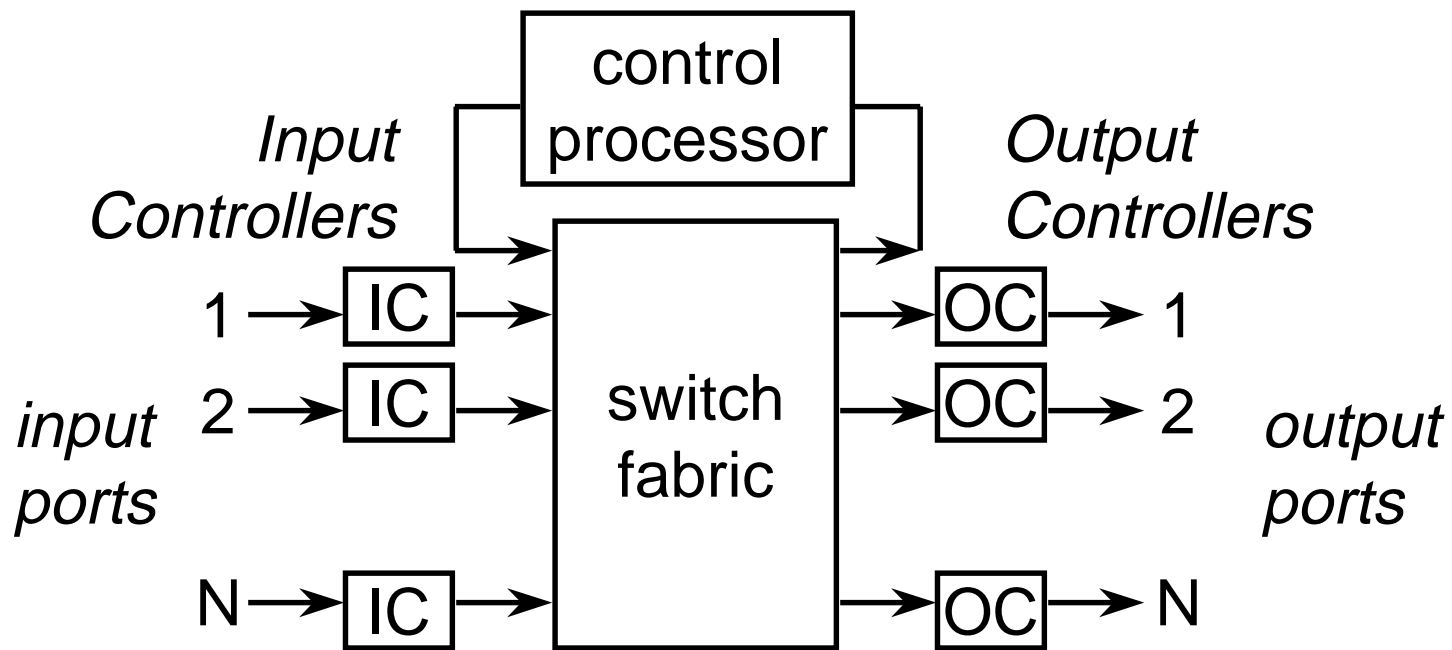
Cell Loss Priority Bit

- Priority bit in ATM header indicates cells that may be discarded
- Can be used by the network to mark traffic beyond a negotiated limit
- Can be used by a user to mark data that can be lost, e.g. video frames

ATM Switches

- A single 155 Mbps connection can produce
 $155 \text{ Mb/s} \div [(8 \text{ b/B}) \times (53 \text{ B/cell})] \approx 366,000 \text{ cells/sec}$
or $2.7 \mu\text{s/cell}$
- Switches need to range from a few input ports to thousands of input ports
- VLSI and parallel implementation critical to achieve throughput

General Structure of an ATM Switch (1)

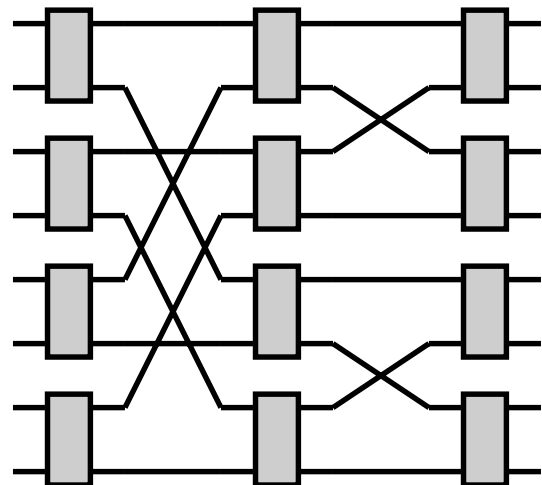


General Structure of an ATM Switch (2)

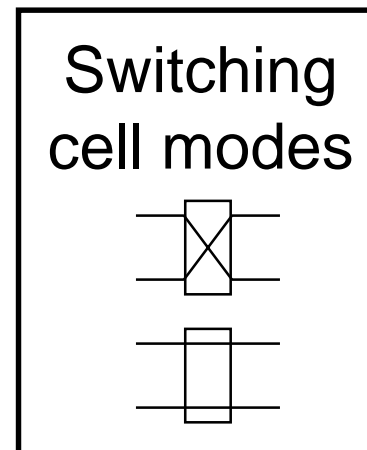
- The control processor is used only for high-level functions such as connection establishment and release, bandwidth allocation, maintenance, and management
- All switching done by input controllers (ICs), switch fabric, and output controllers (OCs)
- Cell headers are aligned and IC, OC, and switch fabric operations are synchronous
- Input controllers translate a cell's VCI, VPI, or VPI/VCI to an output port -- contention may occur
- Design issues: switch fabric topology, cell buffer locations, contention mechanisms

Example Switching Fabric Topology

- Banyan network provides a single path from each input to each output
- An $n \times n$ switch is constructed from $d \times d$ crossbar switches
- Requires $\lceil \log_d n \rceil$ columns of $\lceil n/d \rceil$ switches, or $\lceil \log_d n / \log_d p \rceil$ columns of $\lceil n/p \rceil$ chips for chips with p inputs

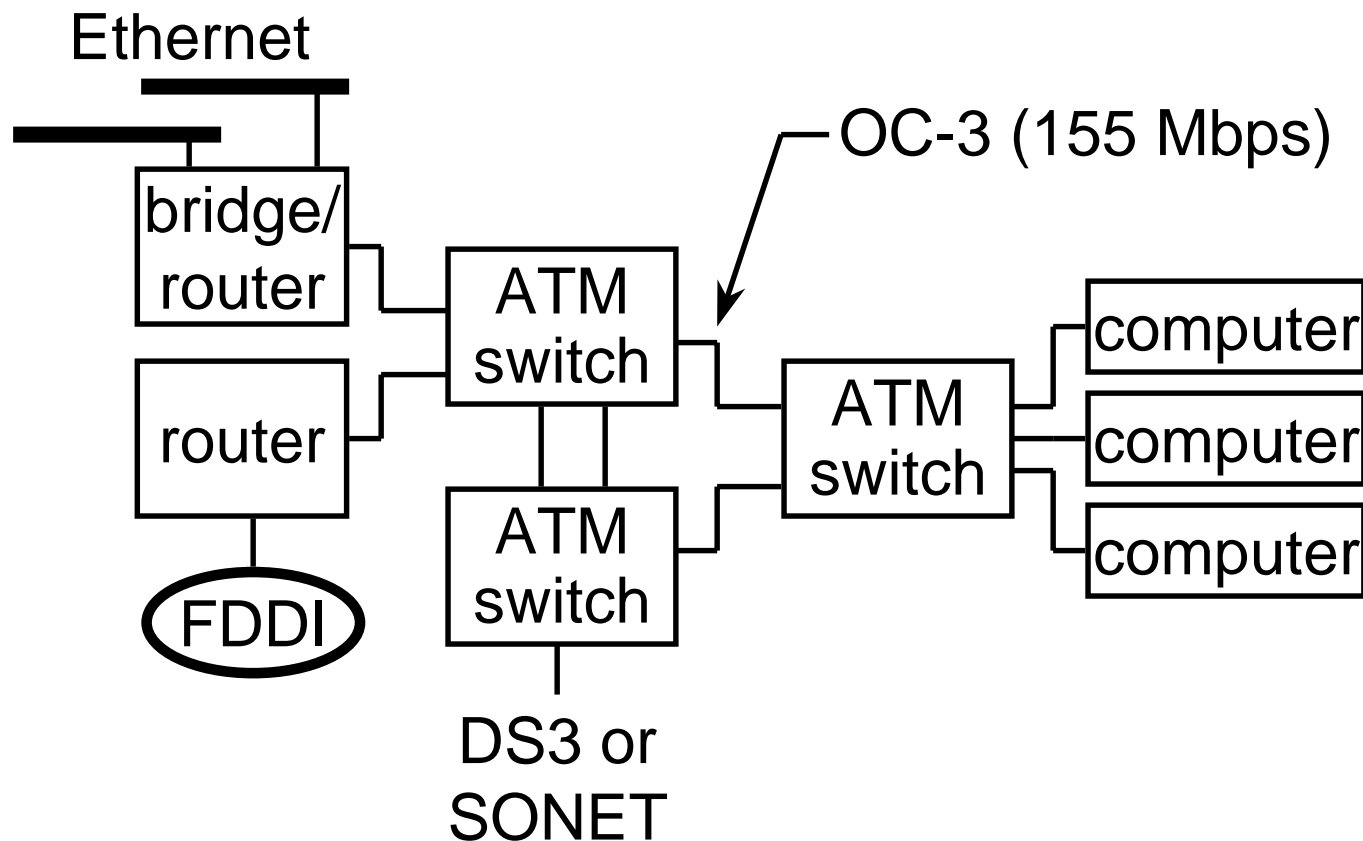


$n=8, d=2$



ATM in Local Area Networks

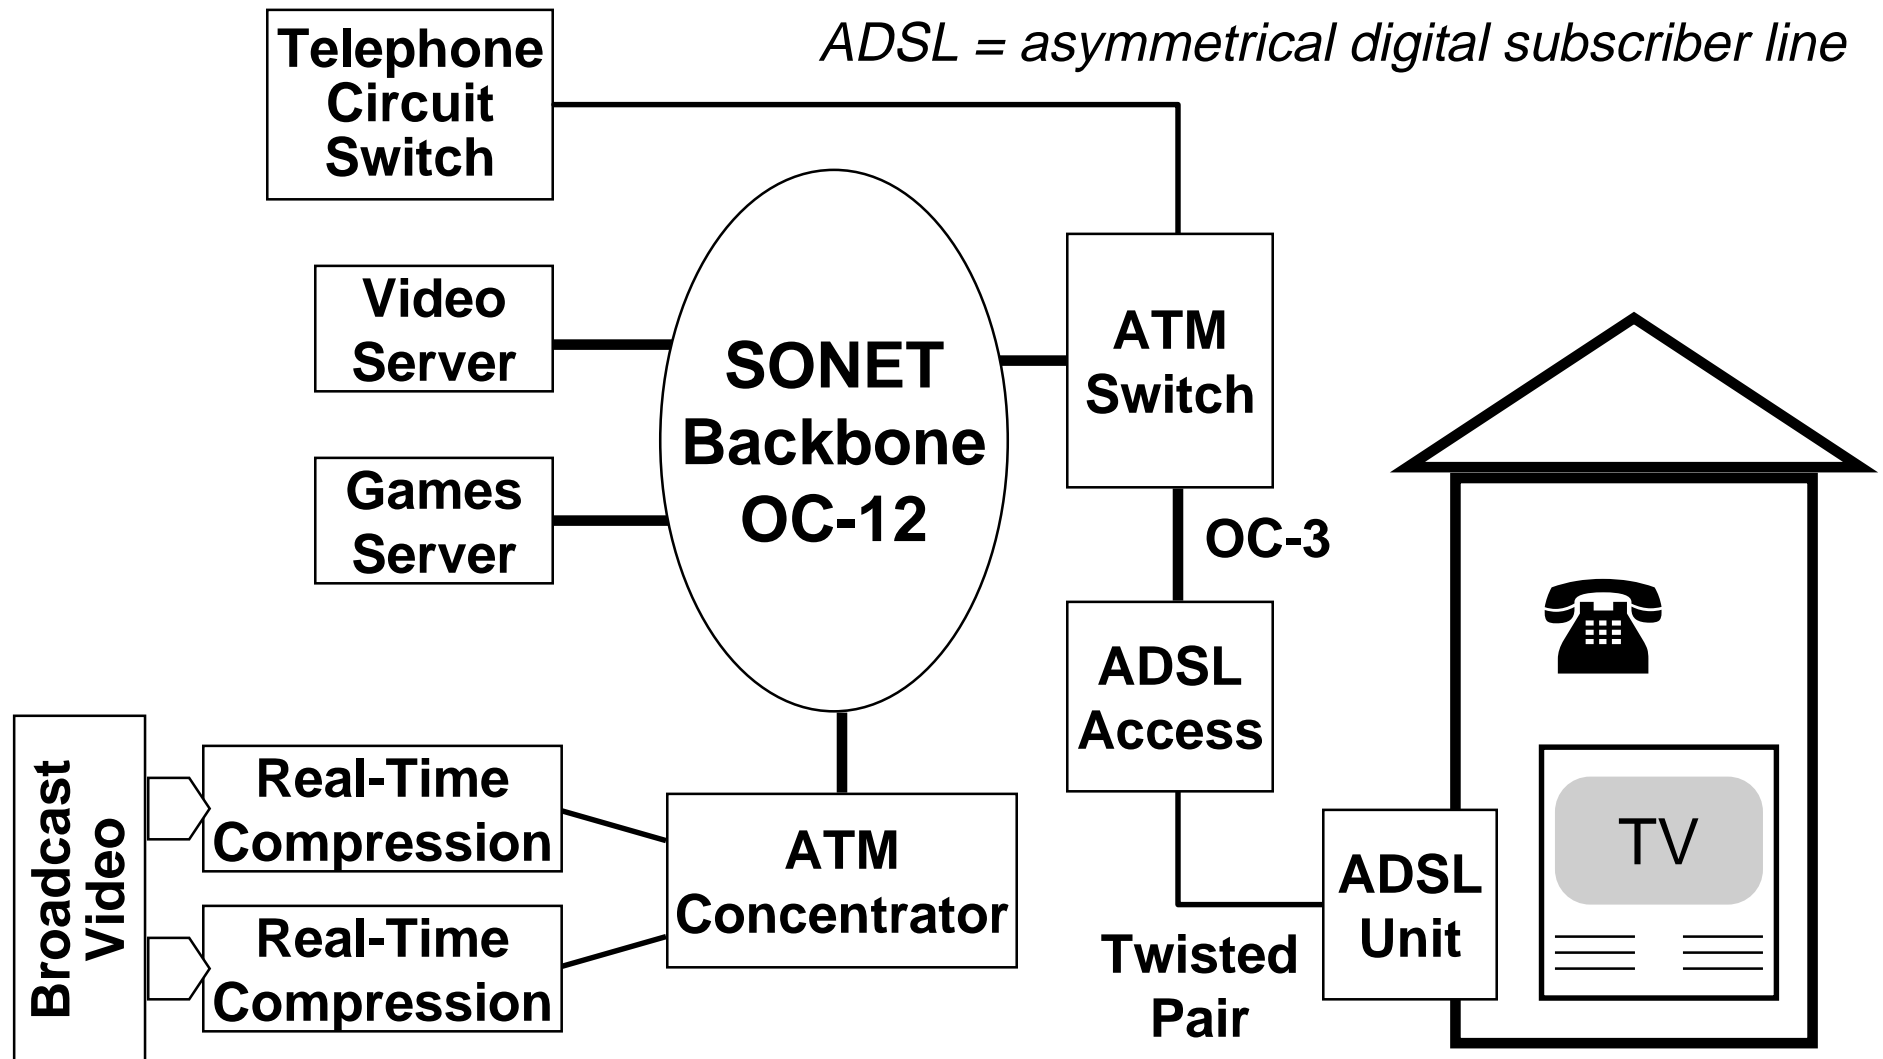
- ATM switches are available now for LAN applications



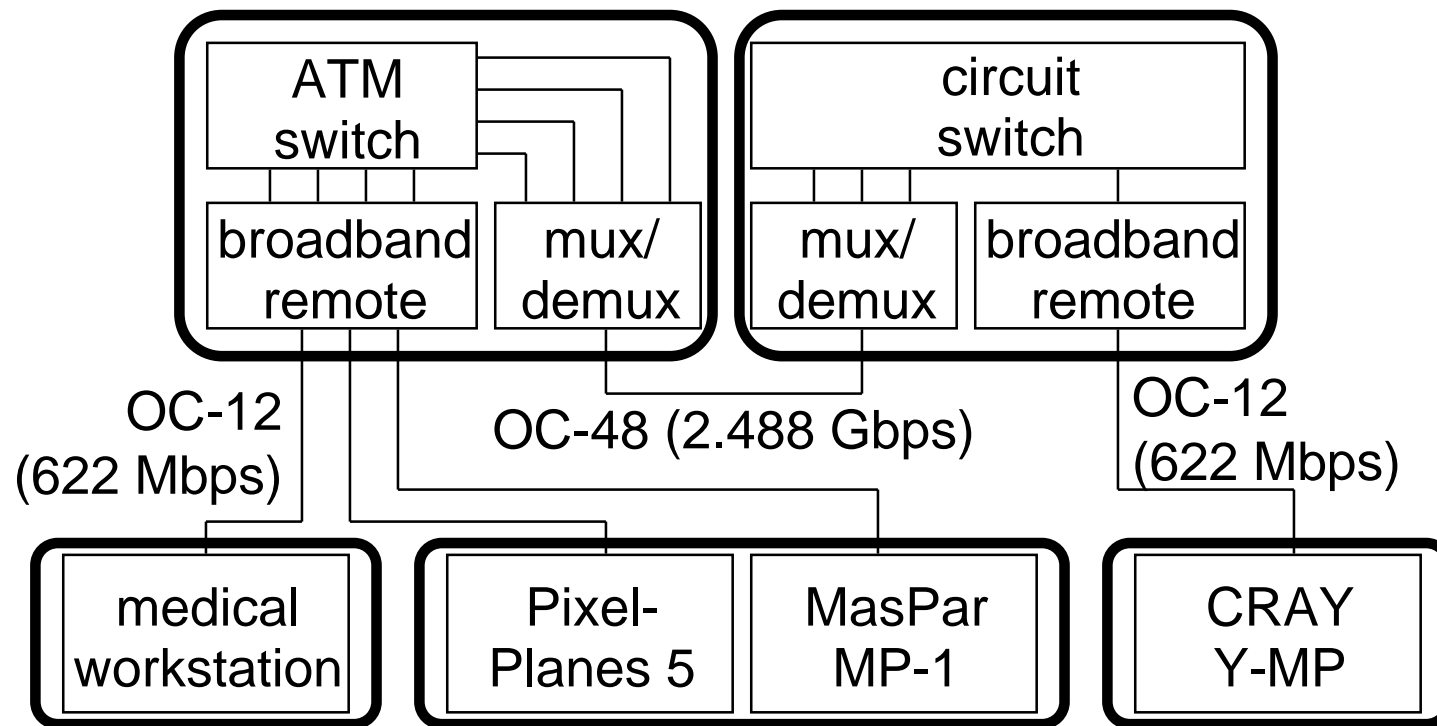
ATM WAN Applications

- LAN-to-LAN internetworking
- Video services
 - Switched access television
 - Video-on-demand
- High resolution still image transfer
 - Medical imaging
- Multimedia services
 - Multimedia conferencing/class
 - Multimedia databases
 - Multimedia mail

Video-on-Demand Network



VISTAnet Network Configuration



D. S. Stevenson and J. G. Rosenman, "VISTAnet gigabit testbed," *IEEE J. Selected Areas in Communications*, vol. 10, no. 9, pp. 1413-1420, Dec. 1992.