Asynchronous Transfer Mode (ATM)

Broadband ISDN (B-ISDN)

Petr Grygárek

ATM basic characteristics

- Integrates transfer of voice, video, data and other media using statistical multiplexing
  - multiplexes cells of fixed length
- Combines advantages of circuit-switched networks (constant delay, guaranteed capacity) and packet-switched networks (flexibility, efficiency for intermittent data transfers)
  - using (bidirectional) virtual circuits

ATM network

- implemented using polygonal infrastructure based on ATM switches
  - point-to-point lines allows to combine various media types and transfer rates
  - provides scalable bandwidth from a few Mbps to many Gbps
- ATM switches maintain switching tables for established virtual circuits

ATM network interfaces

- User-to-Network Interface (UNI)
- Network-to-Network Interface (NNI)
- Broadband intercarrier interface (B-ICI)

UNI, NNI and B-ICI

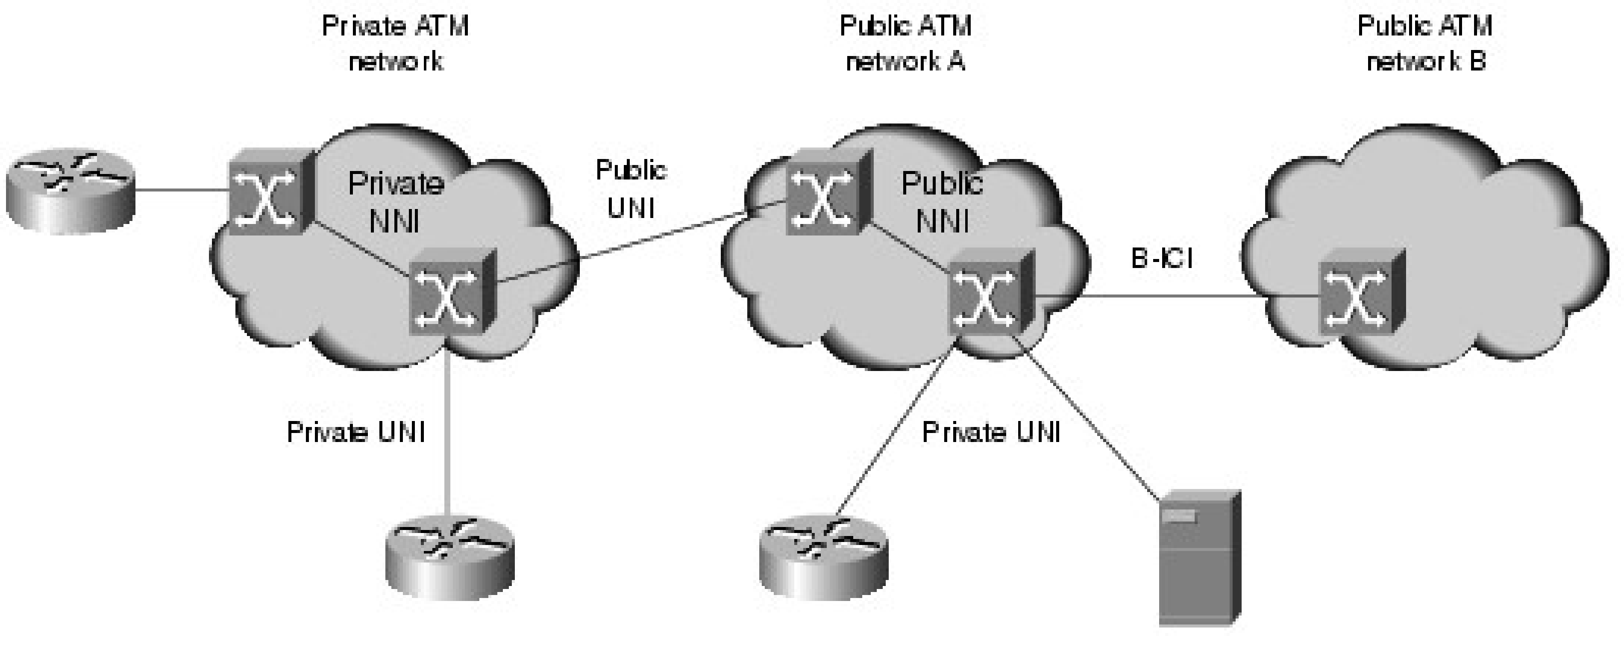

Picture from http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/atm.htm

User-to-Network Interface (UNI)

Connects ATM end systems to ATM switch

- Private UNI
  - connects ATM endpoint and a private ATM switch
- Public UNI
  - connects ATM endpoint or private switch to a public switch
- UNI 3.0, 3.1, 4
  - UNI 3.x doesn't support specification of QoS parameters of required switched virtual circuit
  - UNI 3.1 specification is based on the Q.2931

Network-to-Network Interface (NNI)

Connects two ATM switches

- Private NNI
  - connects two ATM switches within the same private organization
- Public NNI
  - connects two ATM switches within the same public organization
- Broadband intercarrier interface (B-ICI)
  - connects two public switches from different service providers

ATM virtual circuits

- Each cell carries virtual circuit identifier in it's header (VPI+VCI)
  - some VPI/VCIs reserved for signalling and management
- Permanent and switched virtual circuits (PVC, SVC)
  - PVC: manual setup, no network resiliency
  - SVC: signalling protocol needed
- Virtual circuits are QoS-enabled
  - separate QoS contract can be appointed in each direction

- Point-to-point (unidirectional/bidirectional) VCs
- Point-to-multipoint (unidirectional) VCs

ATM Endpoints

- (Workstations)
- Routers, LAN switches
- Digital service units (DSUs)
- Video coder-decoders (CODECs).
- …

Asynchronous transfer mode (statistical multiplexing) - principle and advantages

Advantages of statistical multiplexing

- Bursty character of data and video transfer
  - even 1:1000 ratio
  - inefficient to reserve constant bandwidth using TDM
- In ATM, data in timeslots marked with headers identifying data flow they belong to
  - continuous stream of cells flows on the medium
  - timeslots are available to data flows on demand

Statistical multiplexing principle

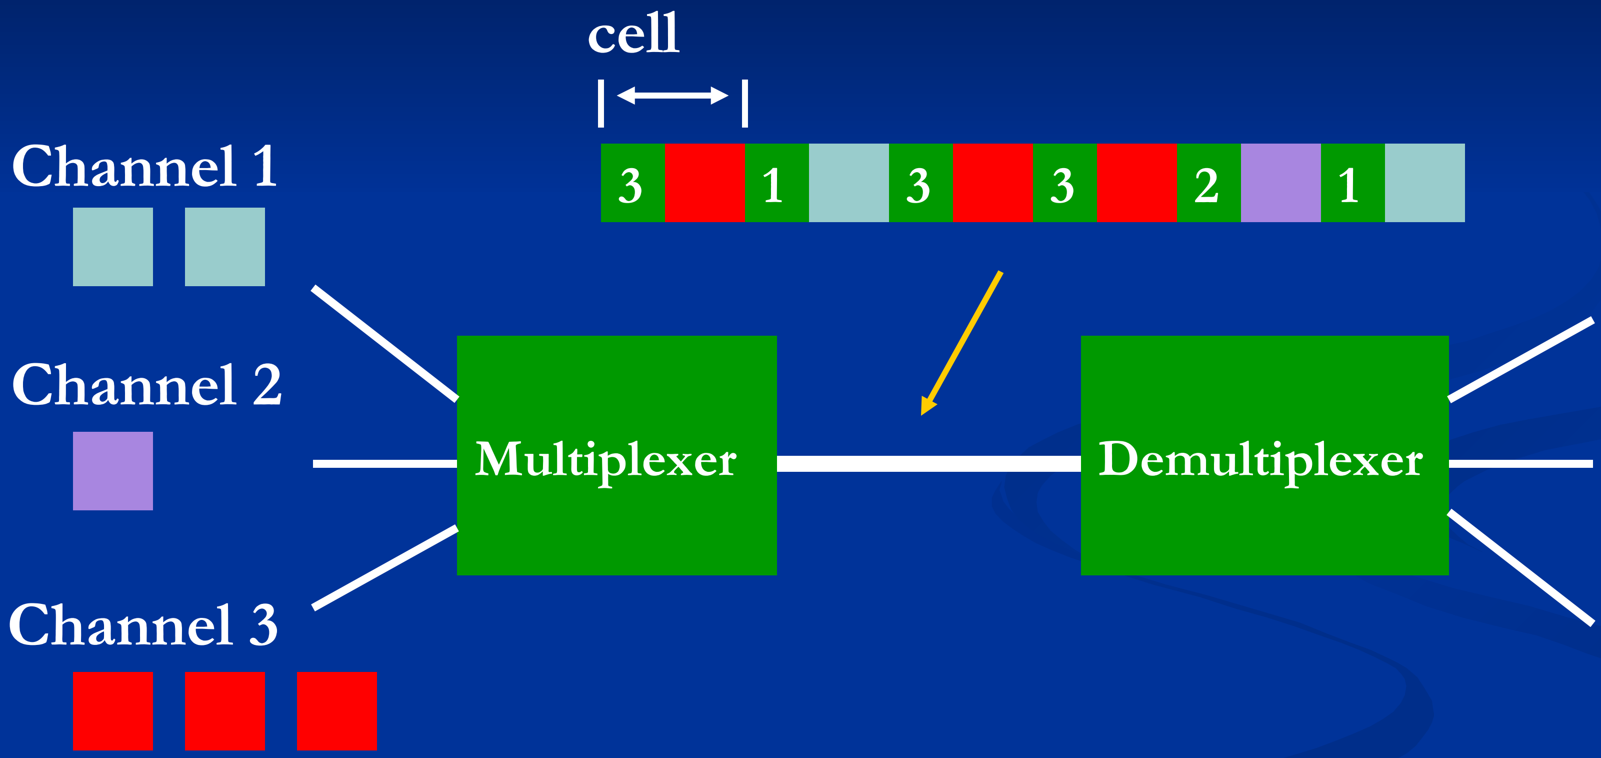

VPI and VCI

- VPI+VCI together identify virtual circuit
- virtual path
  - identified by virtual path identifier (VPI)
  - bundle of virtual channels, all of them can be switched transparently across the ATM network as a bundle (based on the common VPI)
- virtual channel
  - identified by the combination of a VPI and VCI

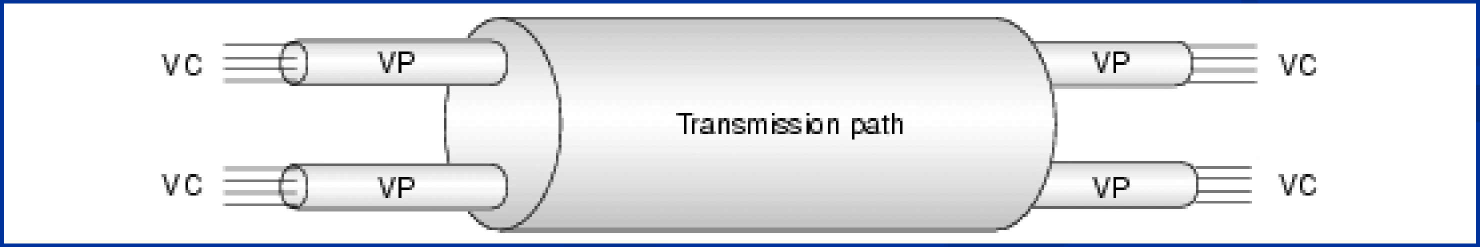

ATM Cell

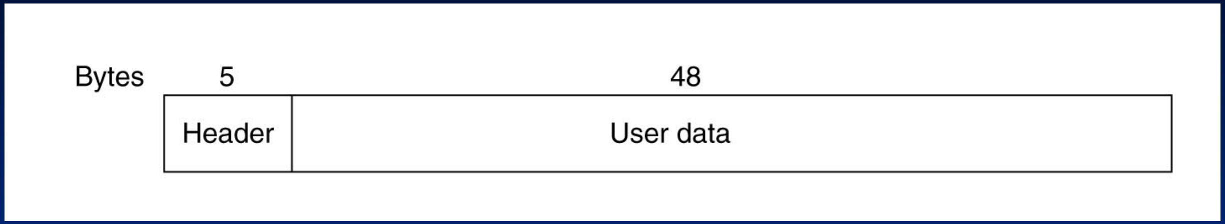

- Total cell length is 53B
- Fixed and short payload length (48B)
  - short and deterministic delay – priority cells can be sent without waiting for finish of transmission of previous long packets
  - well suited to transfer voice and video intolerant of delays and delay variation
  - 48 B is trade-off between 32 and 64B suggestions

ATM Cell header

- VPI, VCI - Virtual Path Identifier, Virtual Circuit Identifier
  - identify virtual circuit (local significance)
- PT - Payload Type
  - Bit 1: cell carries data / system control & management
  - Bit 2: congestion notification
  - Bit 3: indicates whether the cell is the last in a series of cells that represent a single frame (AAL5)
- CLP - Cell Loss Priority
- HEC - Header Error Control
  - can correct a single bit error in header
- Generic Flow Control (GFC)
  - Provides various local functions (typically not used)

Idle cells have reserved header bit pattern

Cell header on UNI and NNI (1)

Slight differences

- NNI header has more bits for VPI field - allows for larger trunks between public ATM switches
- Generic flow control field omitted on NNI

Cell header on UNI and NNI (2)

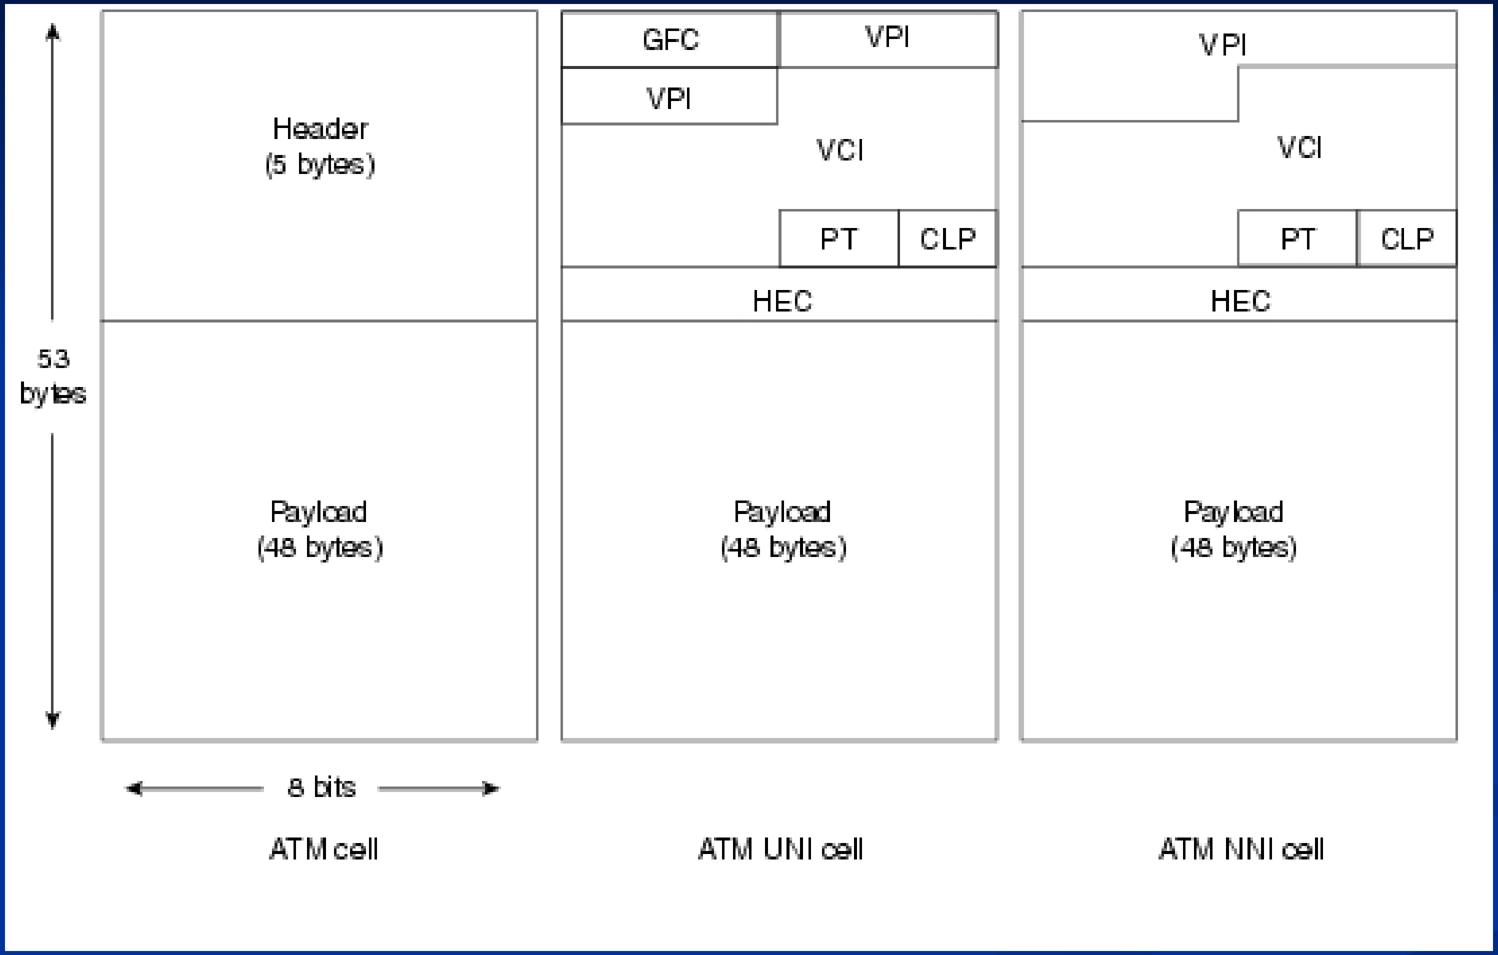

Picture from http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/atm.htm

ATM layers and reference model

ATM Reference model

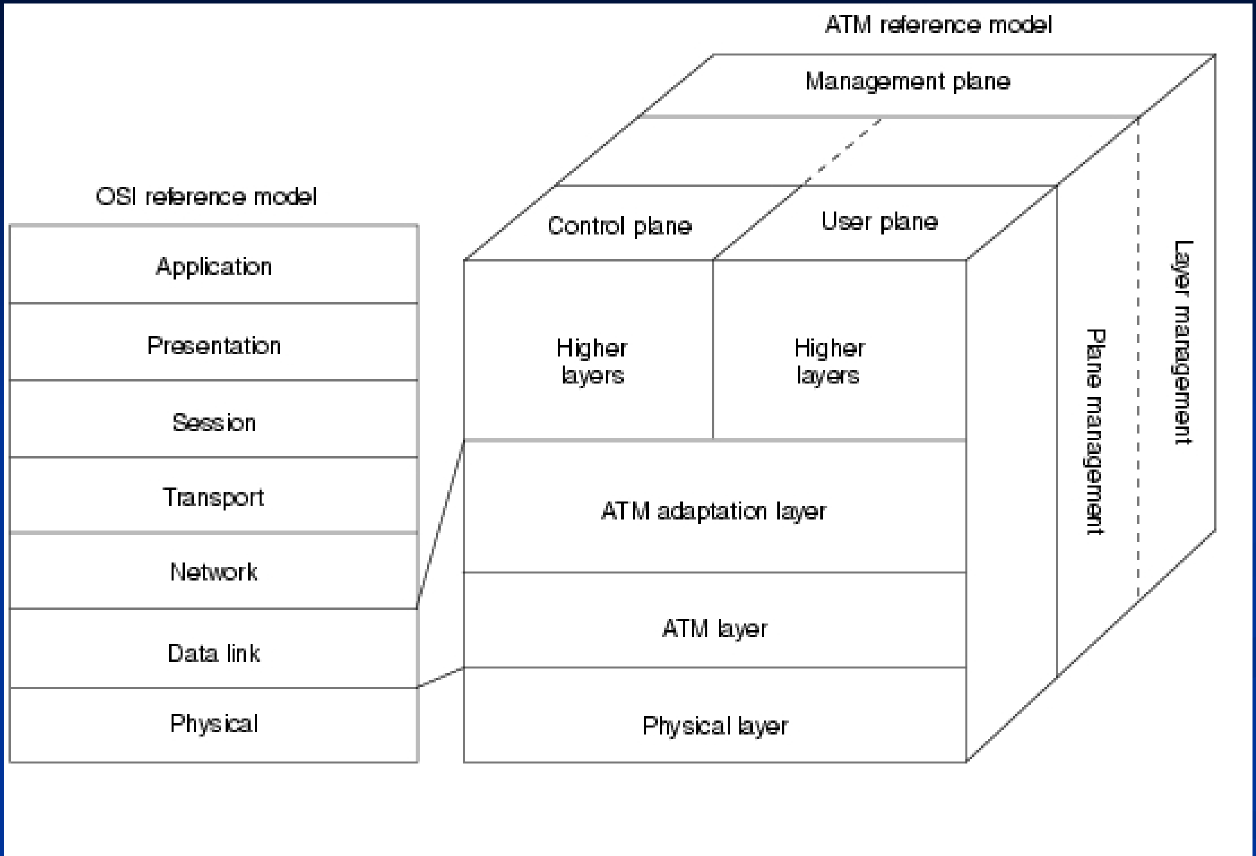

Picture from http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/atm.htm

Covers physical layer and part of data link layer

ATM Reference model planes

- **Control -** responsible for signalling
- **User -** responsible for data transfer
- **Management**
  - Layer management manages layer-specific functions, such as the detection of failures and protocol problems
  - Plane management manages and coordinates functions related to the complete system

ATM layers in detail

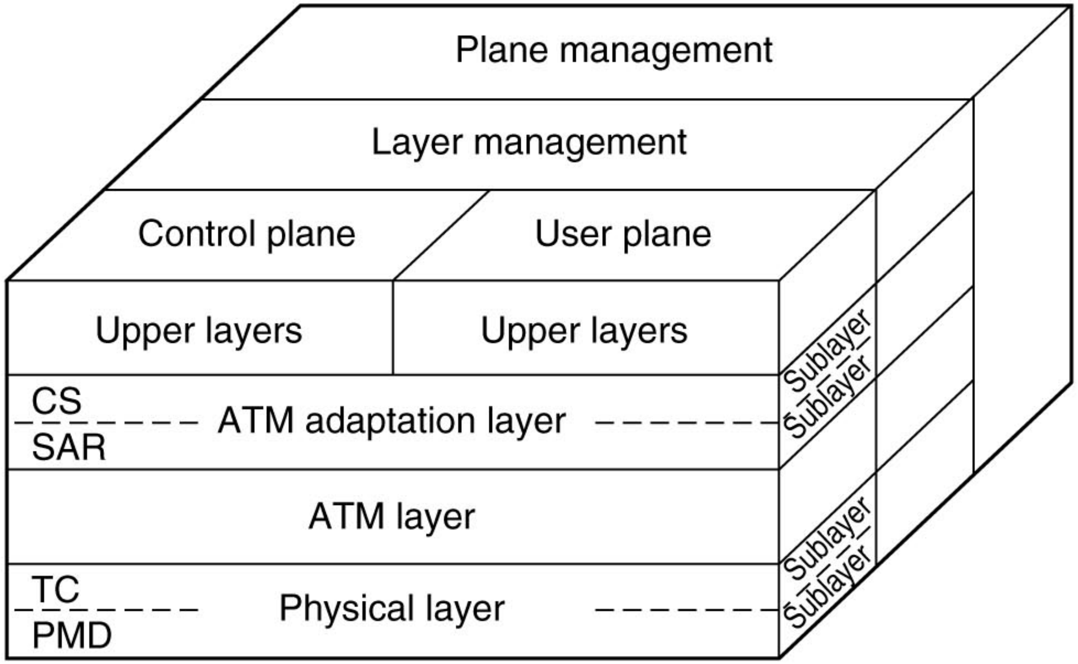

CS: Convergence sublayer
SAR: Segmentation and reassembly sublayer
TC: Transmission convergence sublayer
PMD: Physical medium dependent sublayer

Physical layer

- Manages medium-dependent issues
  - Converts cells into bitstream
    - and packages it into frames of underlying technology (SDH, E3, …)
  - Tracks for cell boundaries
- Divided into two parts
  - Physical medium-dependent (PMD) sublayer
    - synchronizes transmission and reception of bits according to timing information received from medium
    - specifies the physical medium and interface parameters
  - Transmission convergence (TC) sublayer
    - cell delineation
    - header error control (HEC) generation and verification
    - Insertion / removal of idle cells
    - Insertion of ATM cells into physical frames
- Uses various media, various transfer rates
  - plesiochronous hierarchy (E/T lines), SDH/SONET, …

ATM layer

- Switches cells according to switching table
  - forwarding + VPI/VCI rewrite
- Controls sharing of physical link by multiple VCs
  - negotiated VC QoS parameters have to be fulfilled
- Controls SVC connection establishment
  - connection setup (including setup request routing and switching table creation) decoupled from fast switching process
- VPI switch/ VCI+VCI switch

ATM Adaptation layer

- isolates higher-layer protocols from details of the ATM network
  - prepares user data for conversion into cells
  - segments data into 48-byte cell payloads
- Adapts existing upper-layer services to ATM
- Divided to Convergence Sublayer and Segmentation and Reassembly Sublayer
- Multiple adaptation layer types defined in for various user service classes (AAL1-AAL5)

AAL1

- connection-oriented service
- suitable for handling constant bit rate flows (CBR)
- requires timing synchronization between source and destination
  - depends on medium that supports clocking (SONET, SDH)
- Fills cell payload with samples and adds Sequence Number (SN) and Sequence Number Protection (SNP) to the first payload byte to ensure correct ordering

AAL2

- Used for variable bit rate traffic
  - for traffic with timing requirements like CBR, but bursty in nature
  - packetized voice or video
- Either real-time or non real-time
- uses 44 bytes of cell payload for user data and reserves 4 bytes of the payload to support the AAL2 processes

AAL3/4

- supports both connection-oriented and connectionless data.
  - closely aligned with Switched Multimegabit Data Service (SMDS)
    - service common in U.S.
- convergence sublayer creates PDU by prepending a beginning/end tag header to the frame and appending a length field as a trailer
- AAL 3/4 SAR appends PDU header:
  - Type – indicate whether the cell is the beginning, continuation, or end of a message
  - Sequence Number – for reassembly in correct order
  - Multiplexing Identifier – allows to interleave PDUs from different sources on the same virtual circuit
- AAL 3/4 SAR appends PDU trailer: CRC-10

SAR PDU becomes the Payload field of an ATM cell

AAL5

- supports both connection-oriented and connectionless data
  - used most commonly to transfer classical IP over ATM, LAN Emulation, …
- CS sublayer appends a variable-length pad and an 8-byte trailer to a frame
  - pad ensures the PDU falls on the 48-byte boundary of an ATM cell
  - trailer includes the length of the frame and a 32-bit CRC
    - CRC allows to detect bit errors, lost cells, or cells that are out of sequence
- SAR sublayer segments the CS-PDU into 48-byte blocks
  - no additional header added
    - simple, efficient
    - cells of single PDUs cannot be interleaved
  - last cell of a single frame marked with bit in PT header field

Summary of layer functions

OSI layer	ATM layer	ATM sublayer	Functionality
3/4	AAL	CS	Providing the standard interface (convergence)
		SAR	Segmentation and reassembly
2/3	ATM		Flow control Cell header generation/extraction Virtual circuit/path management Cell multiplexing/demultiplexing
2	Physical	TC	Cell rate decoupling Header checksum generation and verification Cell generation Packing/unpacking cells from the enclosing envelope Frame generation
1		PMD	Bit timing Physical network access

ATM QoS

ATM Traffic Classification

- Constant Bit Rate (CBR)
  - isochronous applications like digitalized voice (AAL1)
- Variable Bit Rate (VBR)
  - bursty data transfer (MPEG video)
- Available Bit Rate (ABR)
  - guarantees minimal flow bitrate, tries to reach better
  - implements flow-control
- Unspecified Bit Rate (UBR)
  - no obligatory bitrate

ATM QoS

- Traffic contract
  - specifies limits that describe the intended data flow
    - average sustained bandwidth, peak bandwidth, burst size , …
  - agreed during connection setup
- Traffic shaping
  - constrain data bursts, limit peak data rate, and smooth jitter so that traffic will fit within contract
  - Implemented at boundary user device using intelligent queueing
- Traffic policing
  - implemented by ATM ingress switch to enforce the contract
  - out-of-contract cells marked with CLP=1
    - (cell eligible for discard during congestion periods)

Examples of ATM VC QoS parameters

- Peak Cell Rate
- Sustained Cell Rate
- Minimum Cell Rate
- Cell Delay Variation
- Cell Loss Ratio
- Cell Transfer Delay
- Cell Error Ratio,Cell Misinsertion Ratio, Severely Errored Cell Block Ratio
  - not appointed, given by network nature

ATM Addressing

Usage of ATM addresses

- Identify ATM end systems
- Needed for SVC establishment

ATM Address Options

- ITU-T: dictates use of E.164 addresses
  - used for public ATM (B-ISDN) and some private networks
  - similar to telephone numbers
- ATM Forum: extended ATM addressing for private networks
  - Subnetwork Model of Addressing (overlay model)
    - decouples the ATM layer from any existing higher-layer protocols
      - requires an entirely new addressing scheme and routing protocol

Subnetwork Model of Addressing

- Each ATM system must be assigned an ATM address, in addition to any higher-layer protocol addresses
- requires an ATM address resolution protocol (ATM ARP) to map higher-layer addresses to their corresponding ATM addresses.
- Address format based on the structure of the OSI network service access point (NSAP) addresses – 20B

NSAP Address formats (private networks)

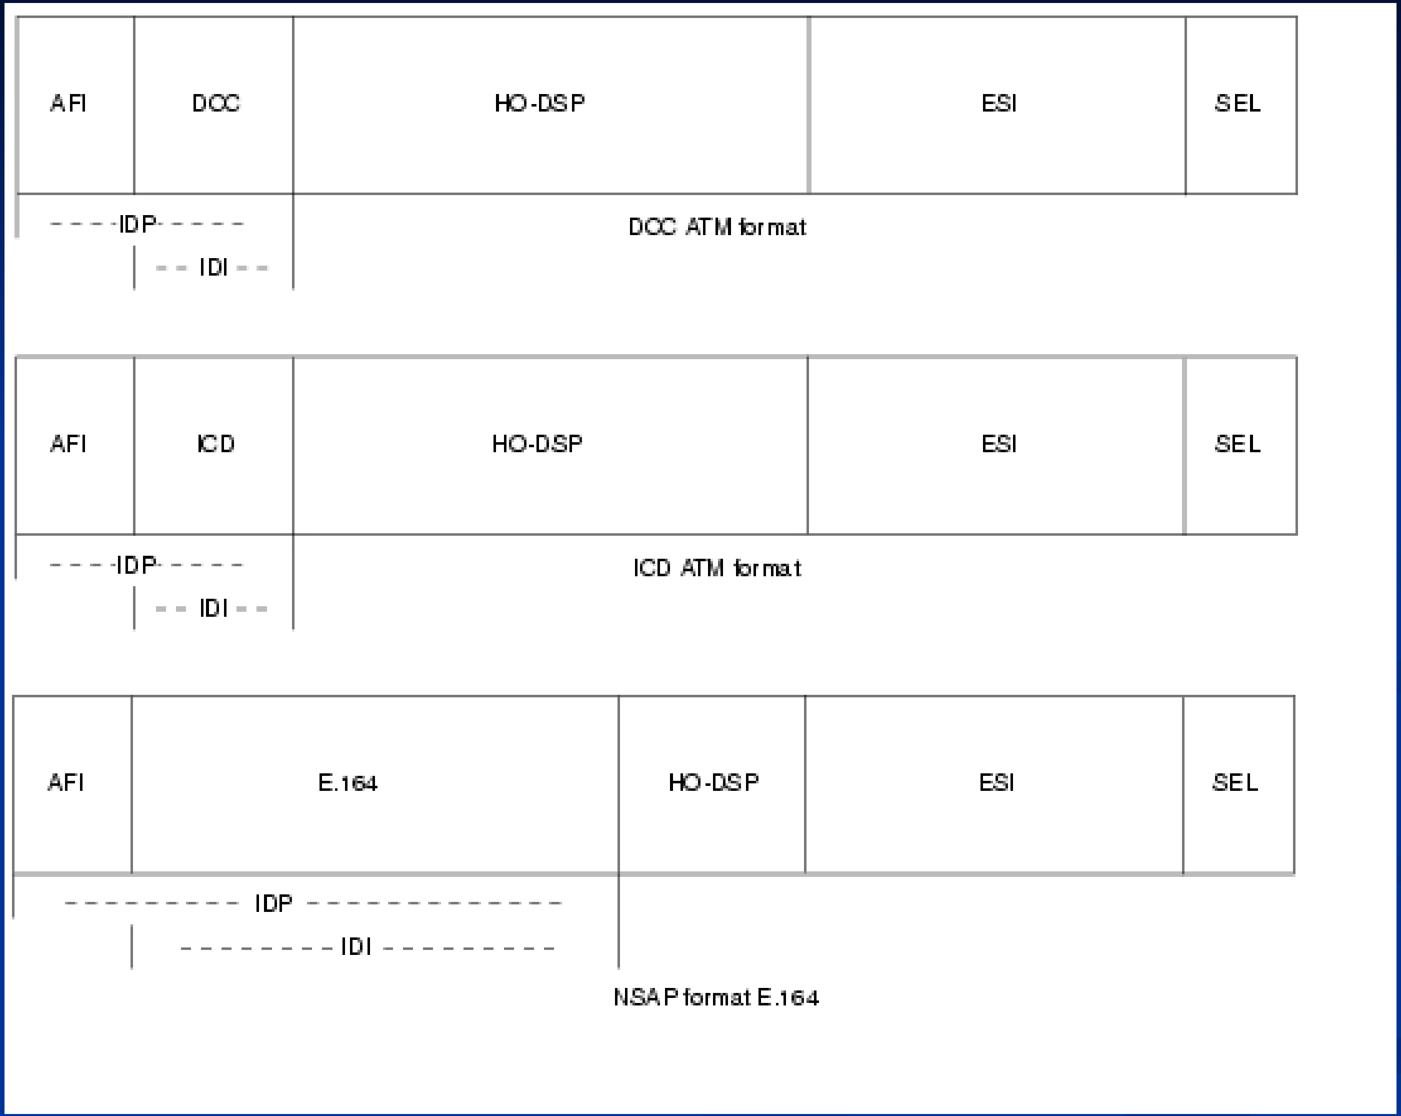

Picture from http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/atm.htm

Fields of NSAP address

- authority and format identifier (AFI)
  - identifies the type and format of the address (E.164, ICD, or DCC)
- initial domain identifier (IDI)
  - identifies the address allocation and administrative authority
- domain-specific part (DSP)
  - contains actual routing information

IDI structure

- NSAP-encoded E.164 format:
  - IDI contains E.164 number
- DCC format:
  - IDI is Data Country Code (DCC) - ISO 3166
  - addresses administered by the ISO National Member Body in each country.
- ICD format:
  - IDI is an international code designator (ICD)
  - allocated by the ISO 6523 registration authority (the British Standards Institute).
  - ICD codes identify particular international organizations.

NSAP address structure (practical view)

- First 13 bytes: NSAP prefix – addresses switch
  - AFI+High-Order Domain-Specific Part (HO-DSP)
  - HO-DSP combines the routing domain (RD) and the area identifier (AREA) of the ISO NSAP addresses
    - Supports more flexible, multilevel addressing hierarchy for prefix-based routing protocols.
- Next 6 bytes: End station identifier (ESI) - identify the ATM element attached to the switch
  - Each device attached to the switch must have a unique ESI value.
  - Corresponds to MAC address
- Last byte: Selector (SEL) - identifies the intended process within the device
  - used for local multiplexing within end station
  - has no network significance

Signalling and request routing

Building of SVC

- Messages similar to ISDN signalling
  - setup, call proceeding, connect, release, …
  - UNI and NNI signalling protocols are different
- UNI signalling requests are carried in a well-known default connection
  - VPI = 0, VCI = 5
- Request contains peer ATM endpoint address and QoS parameters required
- Switch-by-switch request forwarding and VC establishment
  - building of switching tables in ATM switches
  - routing needed to find way to the opposite endpoint
  - if no resources available on some path, request backtracks and tries to find another route (cranckback)
  - final destination either accepts or rejects the connection request
- Routes calculated based on topology information (PNNI routing protocol) and Generic Connection Admission Control (GCAC) algorithm
  - GCAC used to guess whether ATM switch accept another connection or not

Connection setup:
Routing and resource reservation (according to QoS contract required)

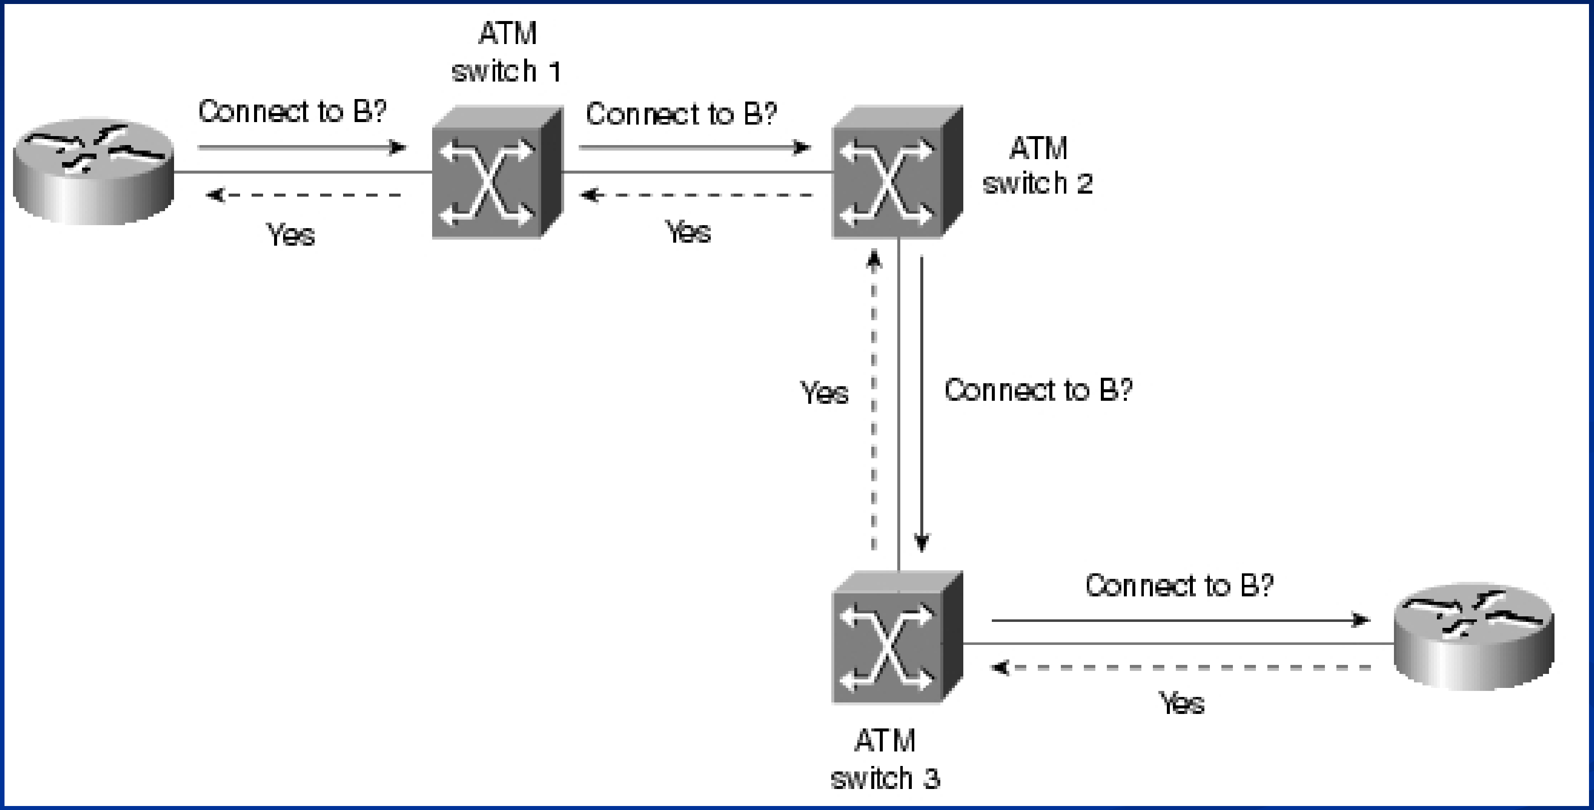

Picture from http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/atm.htm

Connection request routing protocols

- Interim Inter-switch Signalling Protocol
  - Static routing (manual route setup)
- Private-NNI
  - discovers ATM topology
  - link-state protocol (same principle as OSPF)
  - Multiple metrics, load balancing support
  - VCI=18

Routing protocols have to find alternate paths if shortest one cannot fulfill required QoS parameters because of insufficient resources

Integrated Local Management Interface (ILMI)

- Enables devices to determine status of components at the other end of a physical link and to negotiate a common set of operational parameters
  - UNI/NNI + version supported, ATM address prefix
  - Registering of endsystem MAC address
- Functionality roughly analogous to SNMP
- Uses VCI=16

Multicasting in ATM

Problems

- AAL5 does not provide a mechanism to interleave cells from different AAL5 packets on a single connection
  - all cells of AAL5 packets sent to a particular destination across a particular connection must be received in sequence to reassemble them correctly
- If point-to-multipoint connections were bidirectional, multiple leafs could send packets at the same time and receiving nodes would not be able to differentiate between cells of those individual packets

Multicasting - proposed solutions

- multicast server
- overlaid point-to-multipoint connections
- VP multicasting

Multicast server

- nodes wanting to transmit onto a multicast group set up a point-to-point connection with an external device - multicast server
- multicast server connected to all nodes wanting to receive the multicast packets through a point-to-multipoint connection
- The multicast server receives packets across the point-to-point connections and then retransmits them across the point-to-multipoint connection
  - Multicast packet sent from multicast server one after another, i.e. avoiding interleaving

Overlaid point-to-multipoint connection

- every node in the multicast group establishes a point-to-multipoint connection with each other node in the group
  - and becomes leaf in the equivalent connections of all other nodes
- requires a registration process for informing the nodes that join a group of the other nodes in the group so that the new nodes can form the point-to-multipoint connection
  - and be added into existing ones

Multicast server mechanism is more scalable in terms of connection resources but potential bottleneck and a single point of failure.

VP multicasting

- multipoint-to-multipoint VP links all nodes in the multicast group, and each node is given a unique VCI value within the VP
- Interleaved packets hence can be identified by the unique VCI value of the source.
- mechanism would require a protocol to uniquely allocate VCI values to nodes

Usage of ATM for data transfers

LAN Emulation - LANE

LAN Emulation

- Gives to stations attached via ATM the same capabilities that they normally obtain from legacy LANs
- Defines mechanisms for emulating either an IEEE 802.3 Ethernet or an 802.5 Token Ring LAN
  - does not attempt to emulate the actual MAC protocol of the specific LAN (CSMA/CD for Ethernet or token passing for IEEE 802.5).
  - does not support QoS
  - defines a service interface for network layer protocols that is identical to that of existing LANs
- LAN emulated on OSI layer 2
  - ELAN emulates one LAN segment
- LANE operation is transparent to ATM switches
  - uses only standard ATM signaling
  - no need of ATM switch modification
- LANE protocol maps MAC addresses to ATM addresses, so that LANE end systems can set up direct connections between themselves to forward data

LANE – the goal

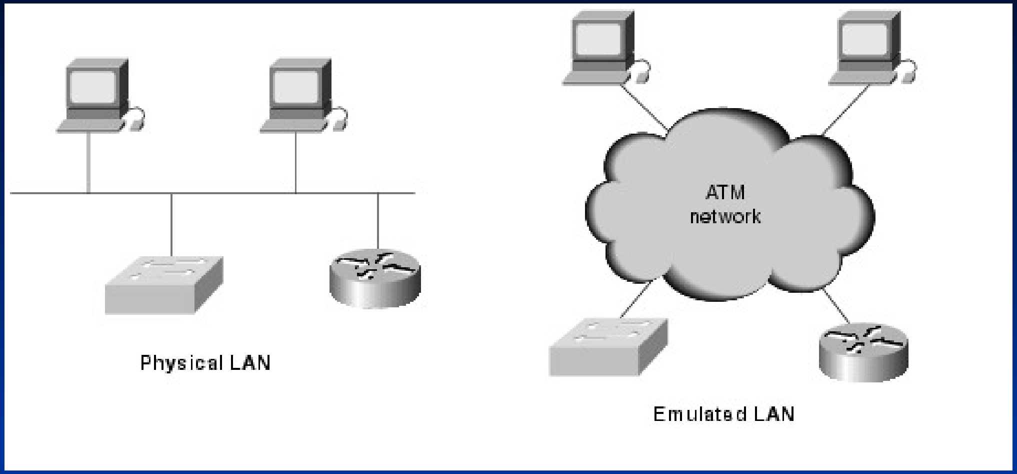

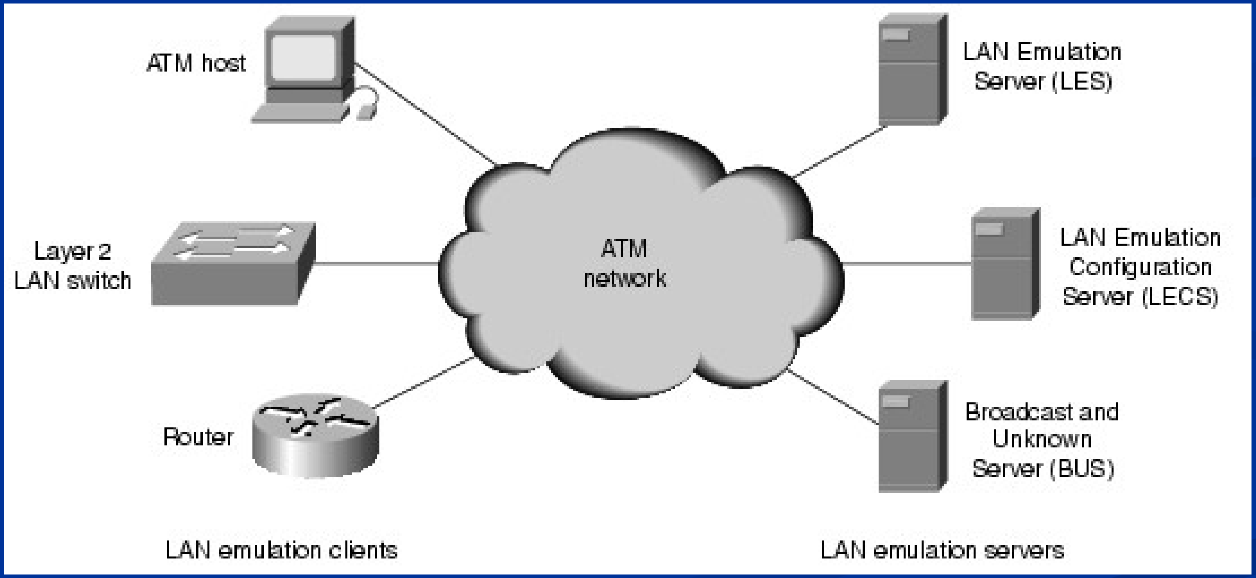

Picture from http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/atm.htm

LANE - Components

- LAN Emulation Client (LEC)
- LAN Emulation Server (LES)
- LAN Emulation Configuration Server (LECS)
- Broadcast and Unknown Server (BUS)

All of them implemented by SW in some ATM switch or LAN switch/router with ATM interface

LAN Emulation Client (LEC)

- entity in end system (station or router/switch) that performs
  - MAC-to-ATM address resolution
  - registration of MAC addresses behind it with the LAN Emulation Server
  - connection establishment to other LECs and data forwarding
- Provides a standard LAN interface to higher-level protocols on legacy LANs
- If end system connects to multiple ELANs, it has has one LEC per ELAN

LAN Emulation Server (LES)

- maintains a list of MAC addresses in the ELAN and the corresponding ATM addresses of LECs
- LECs registers MAC addresses behind them with LES
- LECs can ask LES for ATM address corresponding to given MAC address

One LES exists per ELAN

Broadcast and Unknown Server (BUS)

- used to flood unknown destination address and broadcast traffic to all end-systems behind all LECs within a particular ELAN
- Each LEC is associated with BUS for every ELAN

LAN Emulation Configuration Server (LECS)

- maintains a database of LECs and ELANs to which they belong
- accepts queries from LECs and responds with the appropriate ELAN parameters
  - ATM address of the LES

One LECS per administrative domain serves all ELANs within that domain.

LANE fault tolerance

- Basic structure contains single points of failure, lacks redundancy
- Support for multiple LECS and LES/BUS pairs per ELAN
  - proprietary solutions exist
  - ATM Forum also released a vendor-independent method of providing server redundancy

Example: LANE host to LANE LAN switch communication

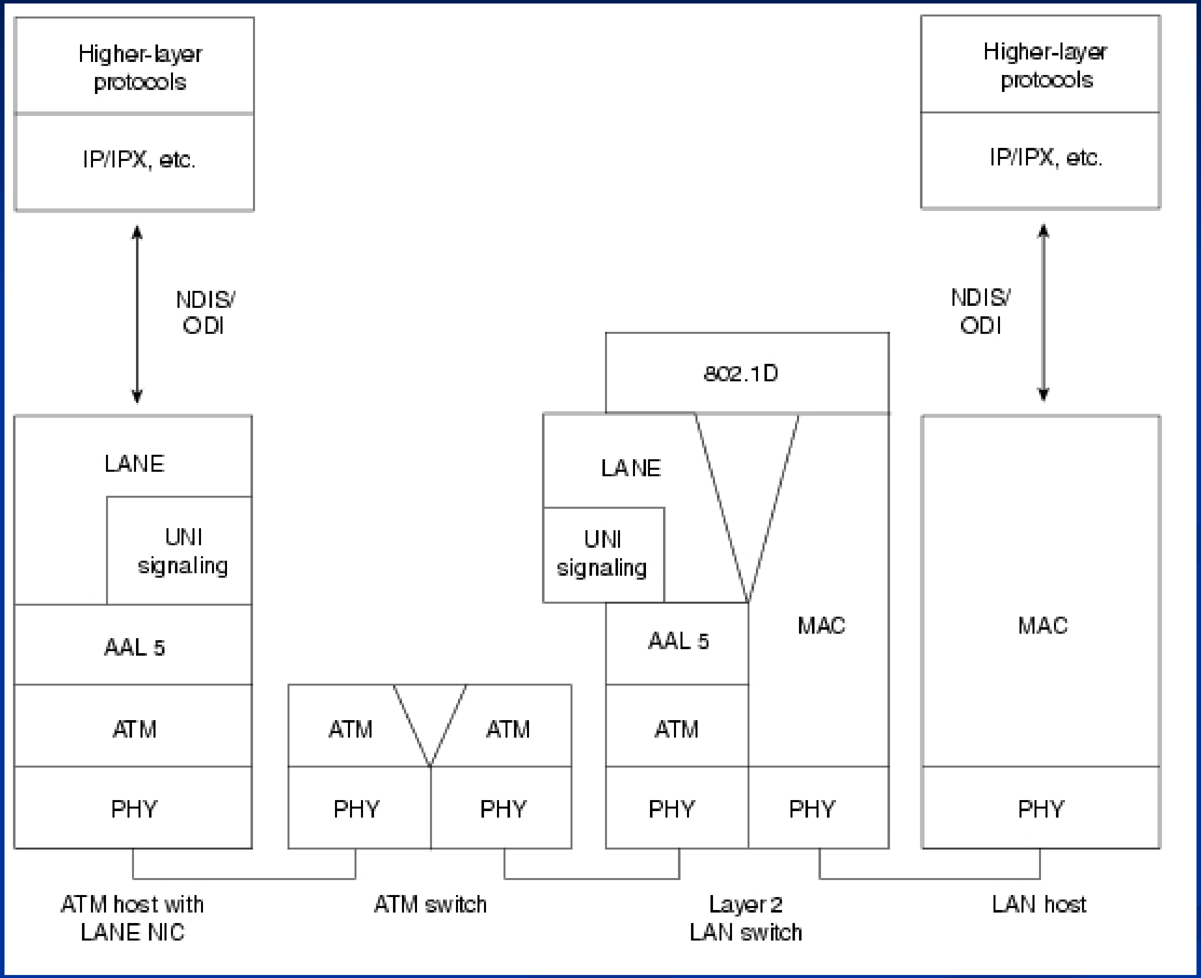

Picture from http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/atm.htm

LANE Connection Types

LANE: Data connections

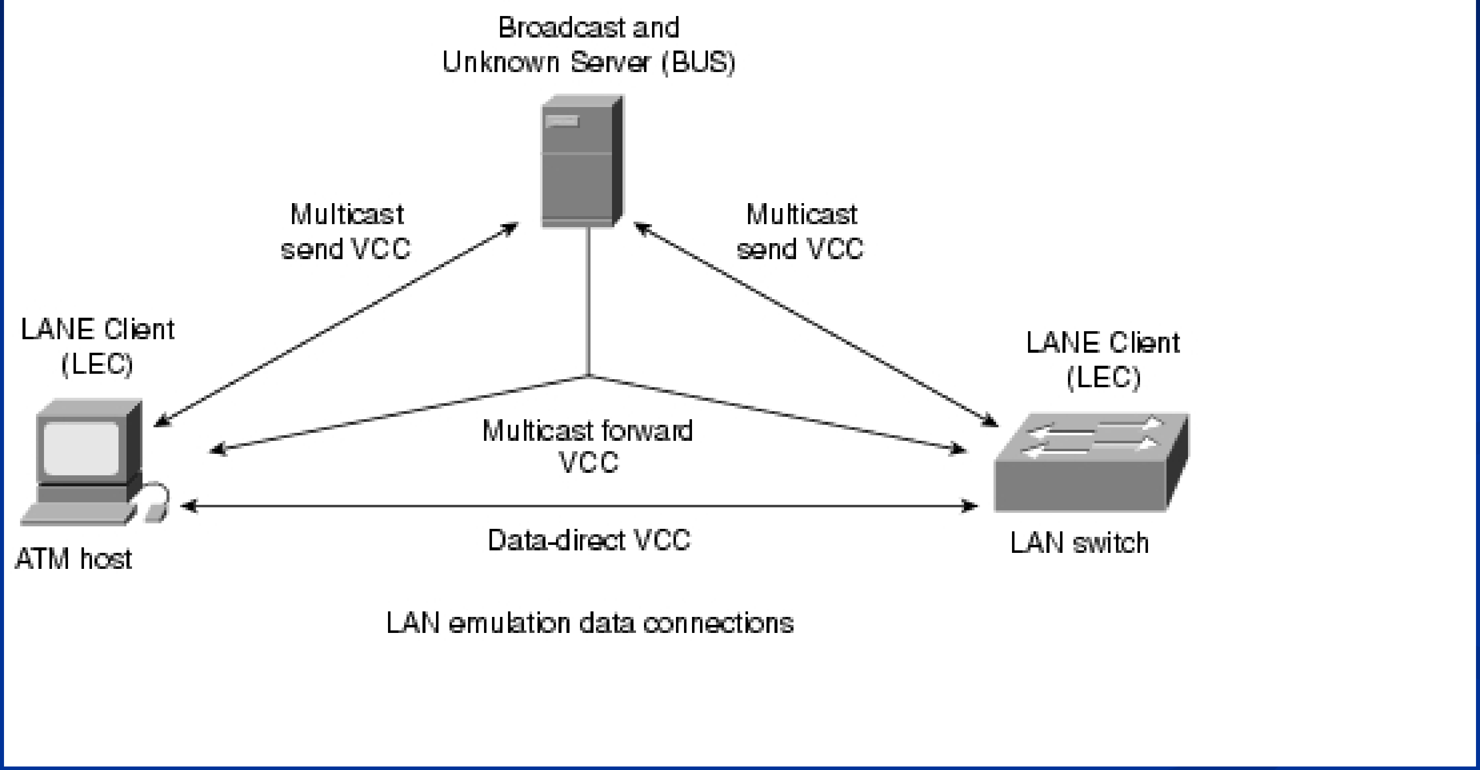

Picture from http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/atm.htm

Data connections

- Data-direct VCC
  - bidirectional point-to-point VCC between two LECs for data exchange between MAC addresses behind those LECs
- Multicast send VCC
  - bidirectional point-to-point VCC set up by LEC to BUS.
- Multicast forward VCC
  - unidirectional VCC set up to LEC from BUS
  - point-to-multipoint connection, with each LEC as a leaf

LANE: Control Connections

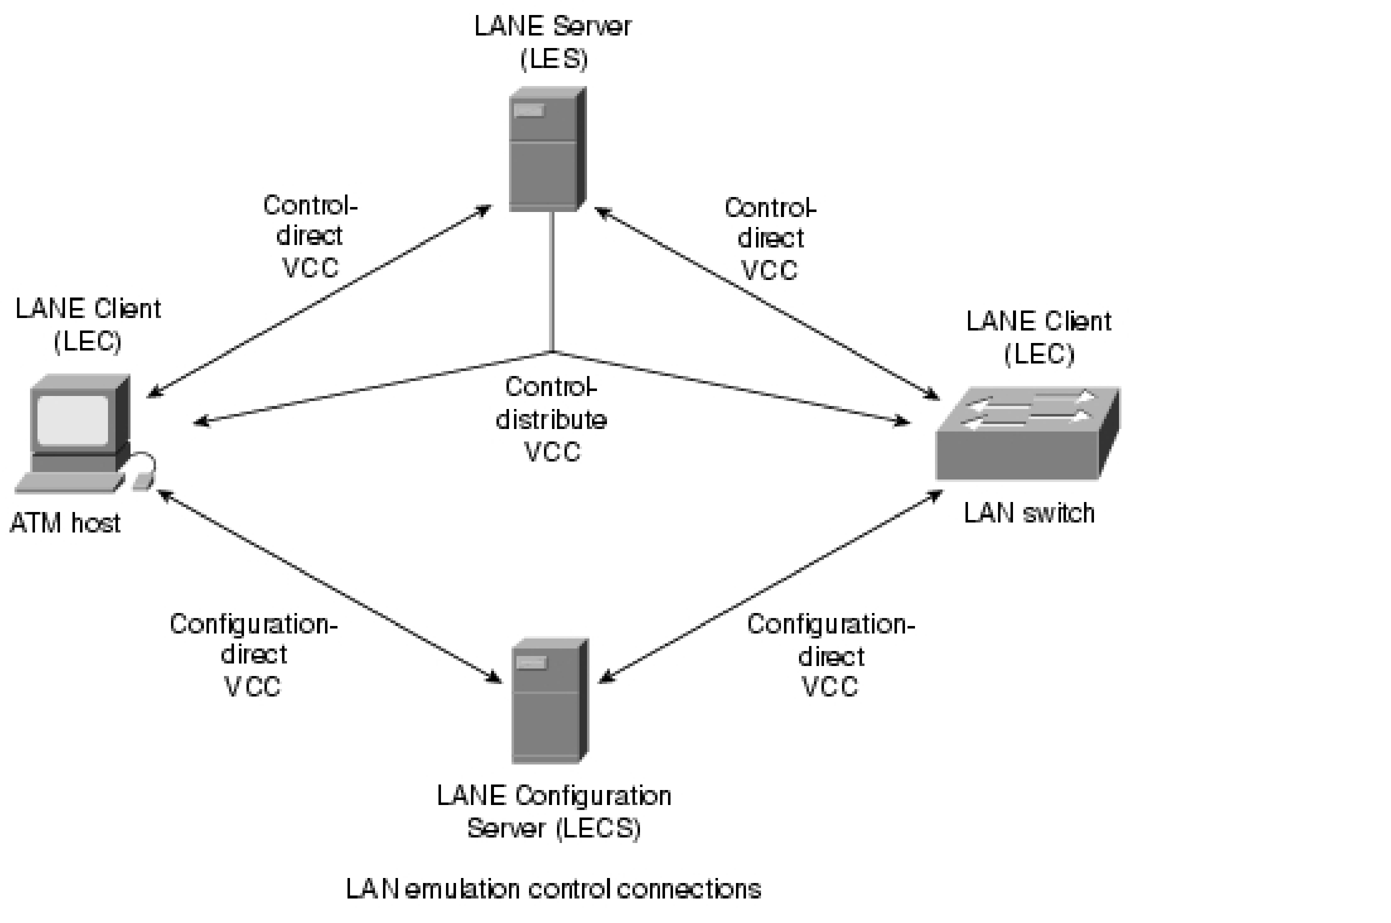

Picture from http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/atm.htm

Control Connections

- Configuration-direct VCC
  - bidirectional point-to-point VCC set up by LEC to LECS
- Control-direct VCC
  - bidirectional VCC set up by LEC to LES
- Control-distribute VCC
  - unidirectional VCC set up from LES back to LEC (typically point-to-multipoint connection).

LANE Operation

1. Initialization

- LEC finds the LECS to obtain required configuration information
  - ILMI procedure or well-known PVC (VPI = 0, VCI = 17).
- Based on LANE client ATM address, LECS returns ELAN parameters
  - LES ATM address, type of LAN being emulated (Ethernet/Token Ring), MTU, ELAN name

2. Registering LECs with LES

- LEC sets up the control-direct VCC to the LES
- LEC joins the LES and registers its own ATM MAC address and all MAC addresses behind it
- LES adds the LEC as a leaf of its point-to-multipoint control-distribute VCC

3. Finding BUS

- LEC asks LES for 0xFFFFFF MAC address
- LES responds with BUS ATM address
- LEC sets up a multicast-send VCC with the BUS
- BUS adds the LEC as a leaf on its point-to-multipoint multicast forward VCC

When sending data to BUS, each LEC appends it's own LEC ID to avoid sending broadcast back to end-systems behind itself

4. Data transfer

- Resolving the ATM address of the destination LEC based on MAC address of destination device
- Connection setup to remote LEC and actual data transfer
- Connection teardown after period of inactivity

MAC address to ATM address resolution

If LEC has a data frame to send to an unknown destination MAC address, it must discover the ATM address of the destination LEC through which the particular address can be reached

- LEC sends LE_ARP_REQUEST to LES
- LES responds with ATM address of destination LEC if it knows it
- If LES does not know the answer, it floods the LE_ARP_REQUEST to other registered LECs
- If some LEC responds, source LEC is informed with destination LEC's ATM address
- Data Direct VCC is established to destination LEC

To decrease delay, frame with unknown MAC address is sent to BUS in parallel with above described process and distributed to all LECs

Classical IP over ATM (CLIP)

Classical IP over ATM Architecture

- Defined in RFC 1577
- Uses ATM ARP to map IP addresses to ATM addresses
- Mappings maintained at ATM ARP server
  - ATM ARP server operates in scope of Logical IP Subnet (LIS)
  - Every node of LIS register it's ATM and IP address with ATM ARP server
- Communication between LIS has to pass router
- Typical MTU 9kB

Multiprotocol over ATM (MPOA)

Multiprotocol over ATM Usage

- Sometimes used together with LANE
- Provides a method of transmitting data between ELANs without need to pass through a router
  - Allows direct connections across subnets (ELANS)
- Only first packet has to be passed over router, direct (shortcut) connection is established after that

Communication without and with MPOA

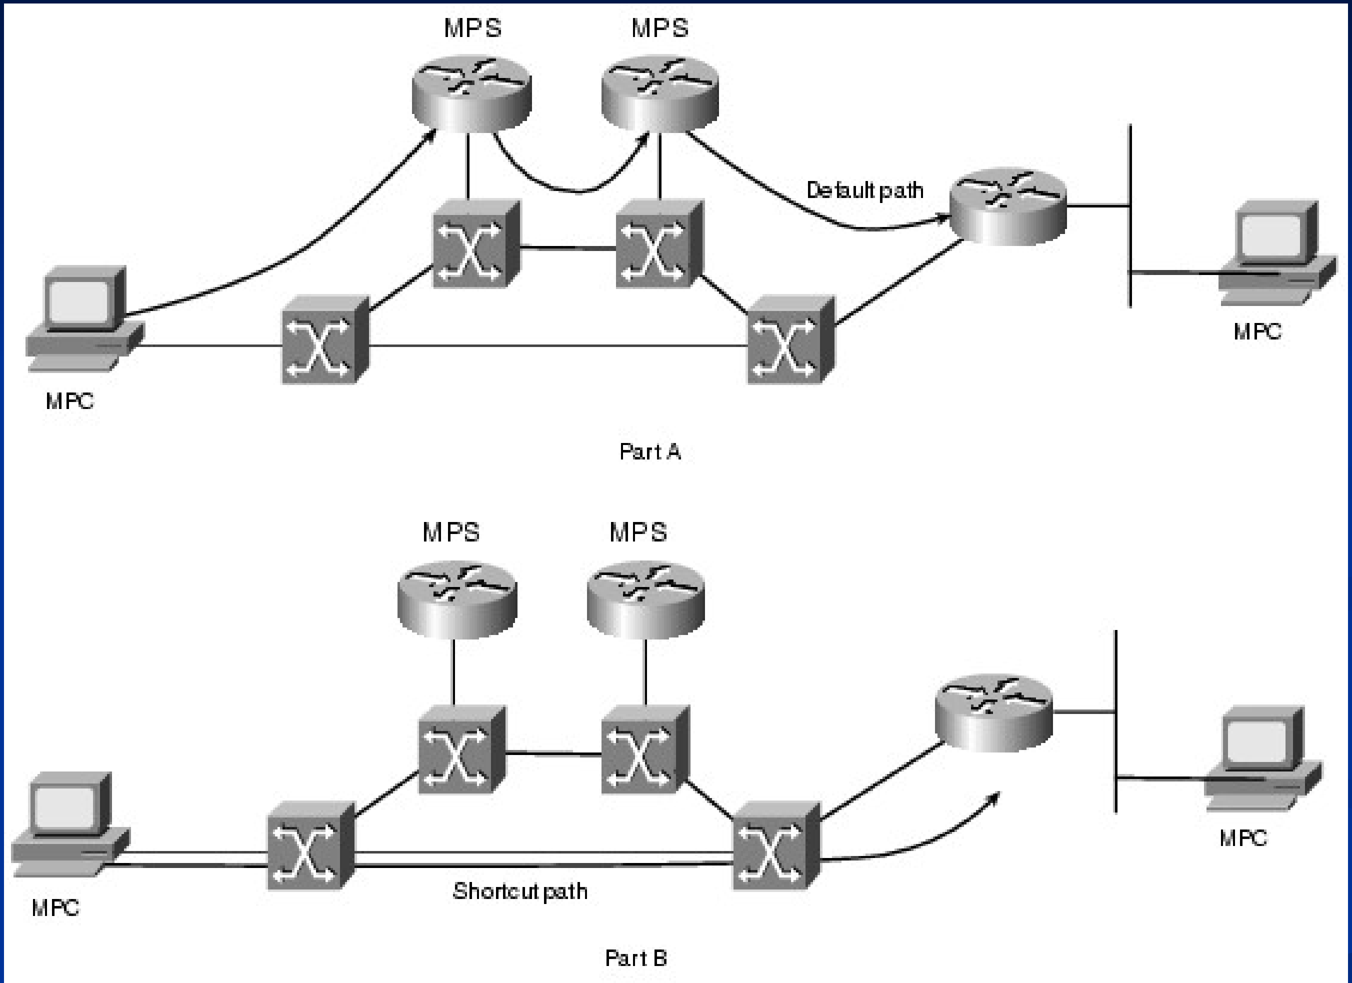

Picture from http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/atm.htm

MPOA Components

- MPOA clients
  - ATM edge devices (most commonly routers)
  - Capable to take information from MPOA servers how to build shortcut VCC and bypass router
- MPOA servers
  - inter-ELAN routers
  - inform MPOA clients how to build a shortcut VCC

MPOA Pros and Cons

- Reduces load of inter-ELAN routers
- Conserves resources of ATM switches
  - one direct connection instead of two separate connections through inter-ELAN router
- Breaks security
  - after packet causing direct connection setup passes the router, LANE client forwards all other traffic directly, bypassing router's ACLs

ATM position today

ATM position today

- Telco company backbone (some advantages over SDH)
  - very flexible
  - allows VBR VCs (videoconferencing)
  - used as infrastructure for Frame Relay service VCs
- Backbone of campus networks
  - Shifted out today and replaced by (10)Gigabit Ethernet
    - much more simpler administration, cheaper
  - Most of campus network designers decided not to invest to ATM any more
- Lot of technologies use ATM principles at physical layer
  - DSL
  - IEEE 802.16
    - (Wireless Local Loop, WiFi replacement for MANs)

ATM Standards

- ITU-T: Broadband-ISDN (B-ISDN)
  - high-speed transfer technology for voice, video, and data over public networks
  - technically, ATM has very little in common with ISDN
- ATM forum – www.atmforum.org
  - public and private ATM networks