Slide supporting material

Lesson 5: ATM Networks, 1st part

Giovanni Giambene

Queuing Theory and Telecommunications: Networks and Applications
2nd edition, Springer

Introduction

ATM Introduction

- **Broadband ISDN** (B-ISDN) was defined in 1990 by the ITU-T Recommendation I.150.
- Asynchronous Transfer Mode (ATM) is the name of a **layer 2 protocol** that characterizes the B-ISDN network.
- The basic transmission unit is a packet of fixed length, named **cell**.
  - The cell is composed of a **payload** of 48 bytes and a **header** of 5 bytes.
  - The cell header of 5 bytes contains all information to support the ATM protocol.
  - The payload of an ATM packet (cell) is transparently-managed by the network: there is no error control and flow control at intermediate nodes, but only end-to-end.
- With ATM, the transmission on the links uses a form of **asynchronous time division multiplexing (i.e., no rigid slot assignment from frame to frame)**.

ATM Introduction (cont'd)

- **Multimedia traffic classes** can be managed by the ATM network. Each traffic class is described in terms of the bit-rate behavior and has guaranteed some **Quality of Service (QoS) parameters** (e.g., mean packet delay, packet loss rate, delay jitter, etc.).
- The ATM network is **connection-oriented**: use of virtual circuits at layer 2 (switching is performed at layer 2).
- Due to the connection-oriented nature of an ATM network, before a sender and a receiver can exchange data, an end-to-end path must be established by means of a **set-up procedure**.
  - During the set-up phase, an end-to-end path is established and it is verified [**Connection Admission Control (CAC)** protocol] that the resources on the links of the path are sufficient to support the new traffic guaranteeing for it (as well as for the already-active connections) the contractual QoS levels.

ATM Introduction (cont'd)

- The ATM technology is quite expensive and not widely employed.
  - ATM is still a valid option for access networks, but not for backbone ones.
  - The ATM protocol (layer 2) is used for the broadband Internet access through the twisted-pair medium of the telephone network (**Asynchronous Digital Subscriber Line, ADSL**).
  - The ATM technology is also used for some geostationary **satellite networks**.

ATM: Multiplexer and Switch

- An ATM network is formed of two different network element types:
  - **Switches**
  - **Multiplexers/demultiplexers (Digital Subscriber Line Access Multiplexer, DSLAM).**

ATM Multiplexer (DSLAM)

- A multiplexer typically allows passing from low utilization input lines to high utilization output lines, i.e., a traffic concentrator that exploits the **statistical multiplexing gain** in the presence of bursty traffic sources.

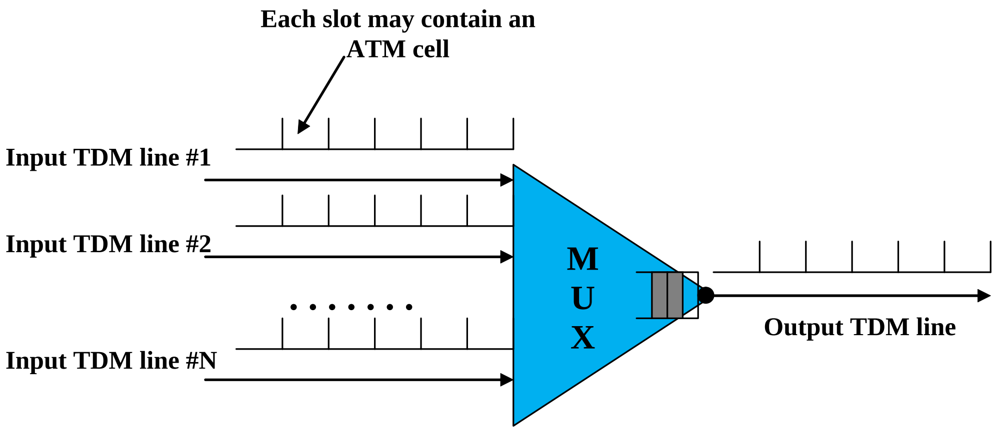

ATM Switch

- A switch connects TDM input lines to TDM output lines.
  - Each packet of a given input line must be analyzed by the processor of the switch.
  - The **virtual path description** in the cell header permits to switch the packet on the appropriate output link on the basis of suitable **switching tables**.
  - Different switch technologies are available. In general, we can consider that, internally to the switch, there are **buffers at input lines or at output lines**.

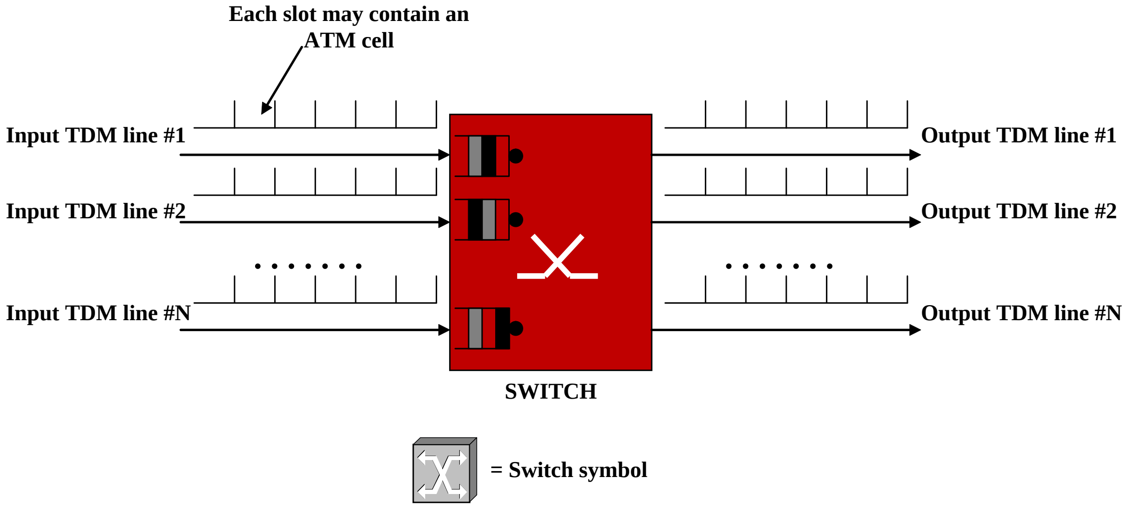

The data in the **switching tables** are either permanently defined or updated during the virtual circuit set-up phase.

ATM Network

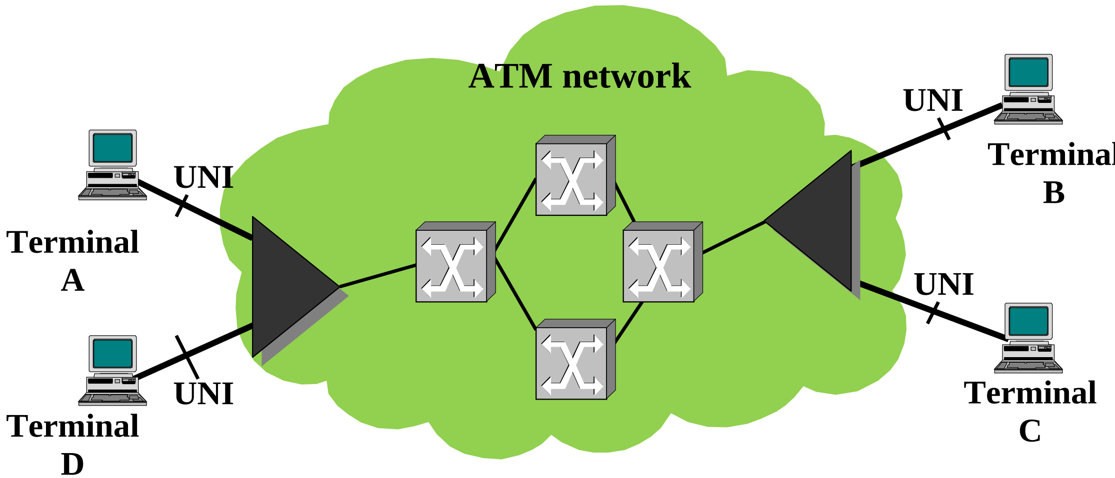

- Multiplexers and demultiplexers are used at the edge of the network. Switches are core elements.
- There are two types of interfaces:
  - The **User-to-Network Interface (UNI):** the interface for the access of the user to the network.
  - The **Network-to-Network Interface (NNI):** the interface between to internal network elements.

Virtual Path and Virtual Channel

- The cell header contains the description of the **virtual path** characterized by means of two fields: Virtual Path Identifier (VPI) and Virtual Channel Identifier (VCI).

An **ATM connection** is characterized by a (VPI, VCI) couple on each link.

These values may be changed at each switch.

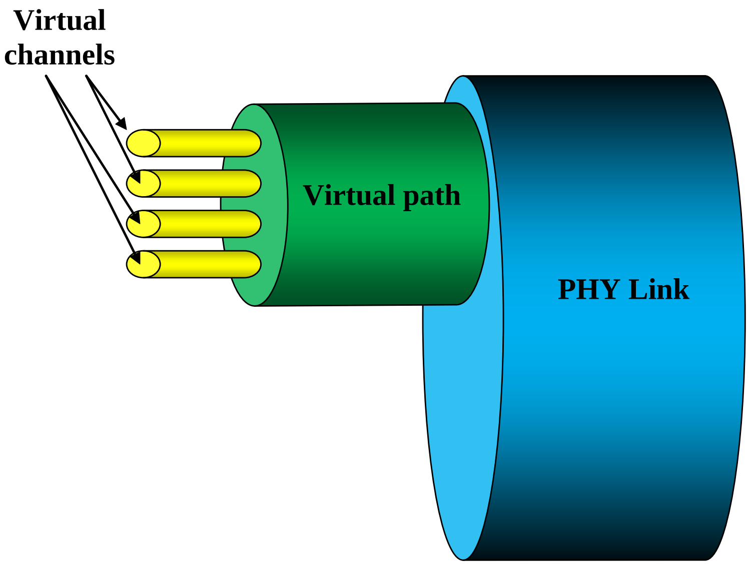

- The physical transmission links employed by ATM are typically based on **optical fibers (SDH)**.

Protocol Stack

- ITU-T Recommendation I.321 characterizes the ATM protocol stack.
- The ATM protocol stack is three-dimensional, with three planes:
  - **User plane,** for the end-to-end transfer of information traffic;
  - **Control plane**, for signaling traffic needed to admit a new connection, for the maintenance of the connection and, finally, for the release of the connection;
  - **Management plane**, for operation and maintenance functions and for the coordination of the different planes.

Both user and control planes are characterized by two (stacked) protocols at layer 2: **ATM Adaptation Layer (AAL) and the proper ATM layer**.

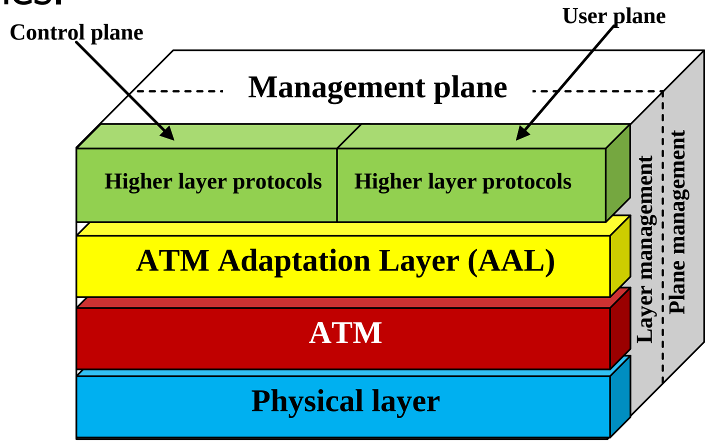

ATM Cell

- In previous data networks (i.e., X.25 and frame relay) the switched unit was a packet of variable length (e.g., in frame relay the maximum packet -frame- length is 4096 bytes).
- In the ATM case, a fixed-length packet, called 'cell' (5-byte header and 48-byte payload), has been defined as a result of a **complex standardization process** that took different aspects into account, such as:
  - Efficient utilization of transmission resources;
  - End-to-end delay to transfer a cell;
  - Routing / switching complexity;
  - Delay to cross a node.
- A PDU received from higher layer protocols is **fragmented** into many ATM cells (the last cell is only partially used thus causing a loss of efficiency).

ATM Cell (cont'd)

- The **cell header reduces the transmission efficiency**, since header bits do not carry information (H = number of bytes of the cell header; P = number of bytes of the cell payload). The efficiency of the cell η can be expressed as:

$$\eta = \frac{P}{P+H} = \frac{48}{53}$$

  - On a link with a physical capacity of 155 Mbit/s, about 14.6 Mbit/s (10 %) are lost due to the cell headers; this is a considerable capacity that has to be used to support the ATM protocol.

- **The use of a fixed-length packet permits to reduce the delays encountered at the queue for the transmission on a link.**
  - This is based on the comparison of the queue delay for M/M/1 and M/D/1 queuing systems with same traffic intensity ρ.

Message Switching versus Cell Switching

- Let E[X] denote the mean packet/cell transmission duration. In the ATM case X is constant.

- The mean delay T experienced at a transmission buffer for the packet/cell transmissions (M/M/1 vs. M/D/1) can be expressed by means of the **Pollaczek-Khinchin formula**:

$$T = E[X] + \frac{\lambda E[X^2]}{2[1-\lambda E[X]]} = \begin{cases} E[X] + \dfrac{\lambda \{E[X]\}^2}{2[1-\lambda E[X]]}, & \text{for M/D/1} \\ E[X] + \dfrac{2\lambda \{E[X]\}^2}{2[1-\lambda E[X]]}, & \text{for M/M/1} \end{cases}$$

More details on this formula will be provided in Lesson No. 7.

- **The use of cells of fixed-length (M/D/1 case) permits to reduce the queuing delays** with respect to packets of variable length (M/M/1 case) with the same traffic intensity.

ATM Cell Structure

- The cell structure is defined in ITU-T Recommendation I.361.

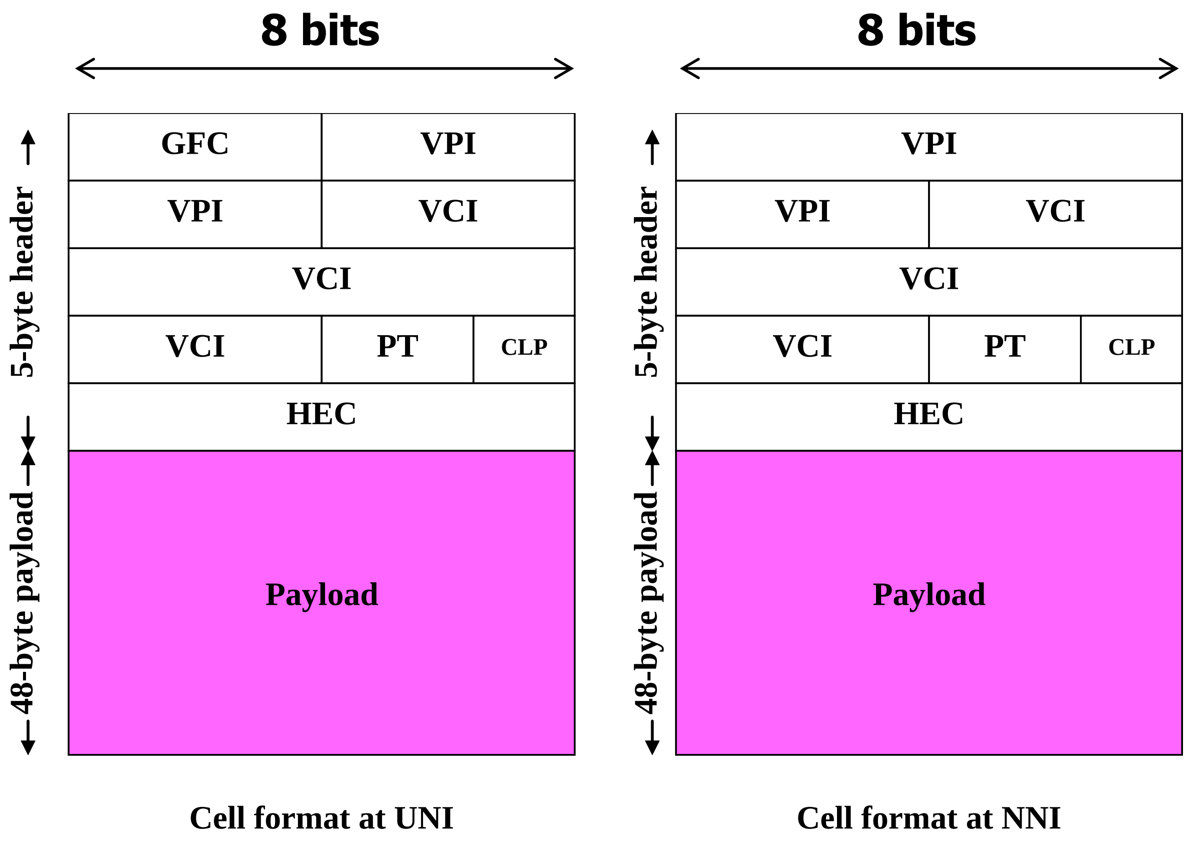

ATM Cell Structure (cont'd)

- GFC (Generic Flow Control) is present in the UNI case, but not present in the NNI one. GFC is used to support a flow control scheme for the input traffic of the user towards the network.
- VPI identifies a virtual path between two nodes; VPI is a field of 8 bits for the UNI cell or of 12 bits for the NNI cell: **2^8 or 2^{12} VPI paths**.
- VCI identifies the virtual channels within a virtual path; VCI is a field of 16 bits (both UNI and NNI cell format): **2^{16} VCI/path**.
- PTI (Payload Type Identifier) field is 3-bit long and is used to describe the type of cell and to transport some control information.
  - PTI describes the content of the cell payload, among the following cases: information data, Operation, Administration, and Maintenance (OAM), Resource Management (RM) signaling.

ATM Cell Structure (cont'd)

- CLP (Cell Loss Priority) flag denotes whether the cell has low (CLP = 1) or high (CLP = 0) priority.
  - **Low priority cells can be dropped if switch queues are congested**. CLP bit can be set either by the sender to differentiate the priority among different cells or by the access node in case that the connection violates its **traffic contract** with the network. The **CLP bit has a similar role to the DE bit in the packet header of frame relay.**
- HEC (Header Error Control) field of one byte makes the **parity check of just the cell header at each hop** (the PHY layer on the basis of 32 bits of the cell header generates the last 8 parity bits of the same cell header).
  - Due to the high reliability of the transmission medium (optical fiber) it is not convenient to check the integrity of the entire cell (this task will be performed only end-to-end). Only the header is verified by means of HEC: if the cell header is **correct (or with a single error that is corrected)** the cell is further forwarded, otherwise the cell is discarded.
  - The HEC code is also used to find the **appropriate cell synchronism in a received stream of ATM cells (SDH)**. The correlation in the header bits introduced by HEC is almost unique in the cell (it is unlikely that the same **correlation on 40 bits** is verified in another position of the cell). Such characteristic is used when the ATM traffic stream has to be extracted from complex physical layer multiplexed streams.
  - **HEC is typically recomputed at each ATM node**.

ATM Protocol Stack

ATM Protocol Stack

- **PHY layer** is divided into two sub-layers: Physical Medium (PM) in charge of physical layer-related functions (electro-optic conversion of bits and bit timing) and Transmission Convergence (TC) that generates the cell HEC field. HEC is recomputed at each node since VPI/VCI may change.
- **ATM layer** operates the following functions:
  - Flow control at the UNI by means of the GFC;
  - Generation of the first 4 bytes of the ATM cell header;
  - Translation of the VPI&VCI fields from input to output of a switch;
  - Multiplexing (and demultiplexing) of the cells of different VPIs and VCIs (connections) on the same stream.
- **AAL layer** is operated only end-to-end and not at intermediate ones. The AAL layer is sub-divided into two different sub-layers: Segmentation And Reassembly (SAR) and Convergence Sublayer (CS). AAL layer has the following tasks:
  - End-to-end transfer of messages of various lengths with cells of fixed length (segmentation/reassembly);
  - Management of erroneous cells and lost cells;
  - Flow control and congestion control;
  - Timing of the transported flow;
  - Multiplexing of different traffic flows of a device on the same ATM connection.

ATM Protocol Stack (cont'd)

- **SAR functions** are as follows:
  - In transmission, SAR divides the PDUs received from the CS sub-layer into smaller units (SAR-SDUs) that, with some added control, form the SAR-PDUs that fit with the cell payload length (segmentation); in reception, SAR re-obtains the PDU for the CS sub-layer.
  - SAR introduces bits in the payload of each cell that, depending on the AAL type, have a different function. For instance, cell numbering, PDU length in cells, etc.
  - In some AAL cases, SAR performs the Cyclic Redundancy Check (CRC) on the information bits.

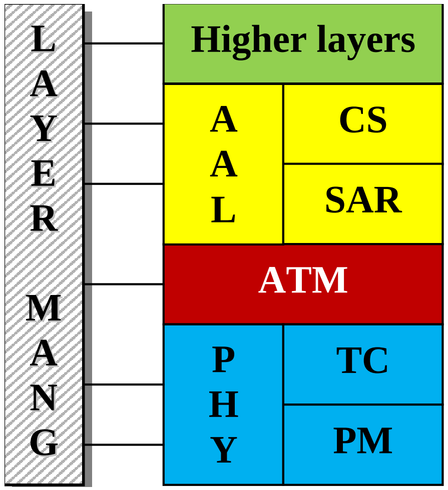

- **CS function** is to manage the higher layer PDUs for the different services supported, thus providing to SAR a CS PDU, including header and trailer control bits.

ATM and ITU-T Traffic Classes

- Traffic classes are differentiated on the basis of time criticality of the information transfer, bit-rate behavior, and type of connection. The **ITU-T Recommendations of the I.363.x series** describe the AAL characteristics for ITU-T **traffic classes A, B, C, and D**.

<table>
<tr><th>Class A</th><th>Class B</th><th>Class C</th><th>Class D</th></tr>
<tr><td colspan="2">Real-time traffic</td><td colspan="2">Non-real-time traffic</td></tr>
<tr><td>Constant bit-rate traffic</td><td colspan="3">Variable bit-rate traffic</td></tr>
<tr><td colspan="3">Connection-oriented services</td><td>Connection-less services</td></tr>
</table>

- We will see later how these generic definitions of traffic classes have been implemented in ATM by means of ATM Adaptation Layers (AAL).

ATM Traffic Classes and AAL

- **AAL1** is used for Class A for the support of services with circuit emulation (dedicated end-to-end circuit). AAL1 is used for Constant Bit-Rate (CBR) real-time traffic for audio, video, and, in general, isochronous applications.
- **AAL2** is used for class B for real-time Variable Bit-Rate (rt-VBR) and for connection-oriented traffic. AAL2 can be used for voice and video packet services. AAL2 allows multiplexing of different AAL2 flows on the same ATM connection with given VPI and VCI fields by means of suitable flow identifiers.
- **AAL3 and AAL4** have practically the same characteristics. They can be used for Class C and Class D, that is non-real-time Variable Bit-Rate (nrt-VBR) traffic for connection-oriented (e.g., frame relay) or connection-less services. The AAL3/AAL4 protocol allows multiplexing of different flows on the same connection by means of suitable flow identifiers.
- **AAL5** has been conceived to simplify AAL3/AAL4; it is the most simple and efficient adaptation protocol for supporting services for all the different traffic classes except CBR. AAL5 does not support the multiplexing of different AAL flows.
- All AAL levels introduce some degree of overhead explained as follows.

ATM Traffic Classes and AAL (cont'd)

<table>
<tr><th></th><th>Class A
(circuit emulation)</th><th>Class B
(packetized voice/video)</th><th>Class C
(connection-oriented data)</th><th>Class D
(connection-less, Datagram)</th></tr>
<tr><td rowspan="2">Type of AAL</td><td rowspan="2">AAL1</td><td rowspan="2">AAL2</td><td>AAL3</td><td>AAL4</td></tr>
<tr><td colspan="2">AAL5</td></tr>
<tr><td>Timing between source and destination</td><td colspan="2">Real-time</td><td colspan="2">Non-Real-Time</td></tr>
<tr><td>Bit-rate</td><td>Constant</td><td colspan="3">Variable</td></tr>
<tr><td>Connection mode</td><td colspan="3">Connection-oriented</td><td>Connection-less</td></tr>
</table>

AAL1 Example

- The 48-byte cell payload format for AAL1:
  - The Sequence Number (SN) field is used for denoting lost cells;
  - The Sequence Number Protection (SNP) is a code used to protect the SN field.
  - The remaining 47 bytes form the effective cell payload with AAL1 and the fragmentation unit operated by the SAR sub-layer in this case.

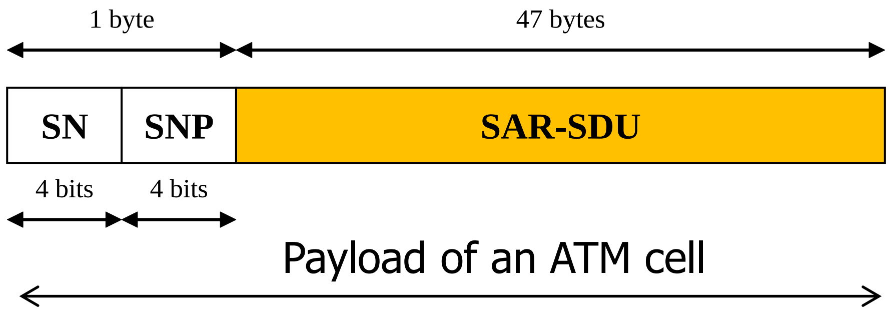

AAL1 overhead of 1 byte

AAL2 Example

- The AAL2 internal protocol architecture entails:
  - Service Specific Conversion Sublayer (SSCS),
  - Common Part Sublayer (CPS).
- SSCS receives the higher layer PDU and formats a **CPS packet** (see next slide picture) to be included in the CPS PDU. Such PDU becomes the payload of the underlying ATM layer cell.
  - The Channel IDentifier (CID) field is a logical identifier of the virtual connection to which this information unit belongs.
  - The Length Indicator (LI) field denotes the length of the CPS packet; the default value considered here is 45 bytes so as to fit the CPS packet in just one CPS PDU that represents the cell payload with AAL2.
  - The User-to-User (UUI) field is used to convey end-to-end user data or to support OAM operations.
  - The Header Error Control (HEC) is a code to protect the first 19 bits of the CS PDU.

AAL2 Example (cont'd)

- CPS PDU format:

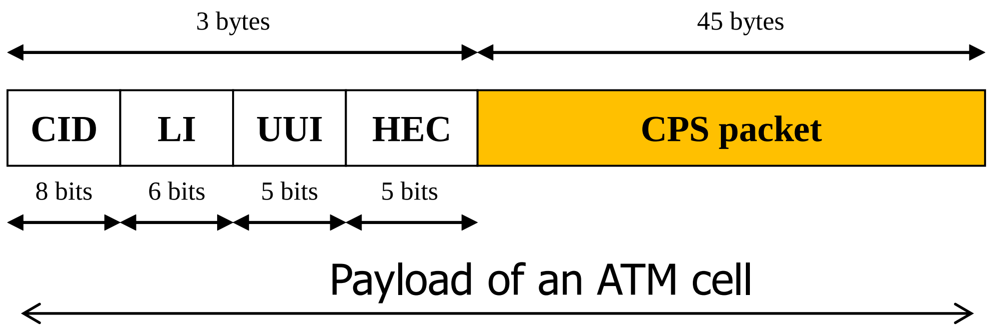

AAL2 overhead of 3 bytes

AAL5 Example

- AAL5 uses a cumulative overhead at the CS PDU level: **a 8-byte trailer is added**.
  - For simplicity, the CS PDU is made of a length multiple of 48 bytes (a PAD field is used to adjust properly the length).
  - The CS PDU payload has a variable length up to 65535 bytes to support IP packets (RFC 1577); such length is coded by the Length (LEN) field in the trailer.
  - The CRC field is used for revealing errors on the entire CS PDU.

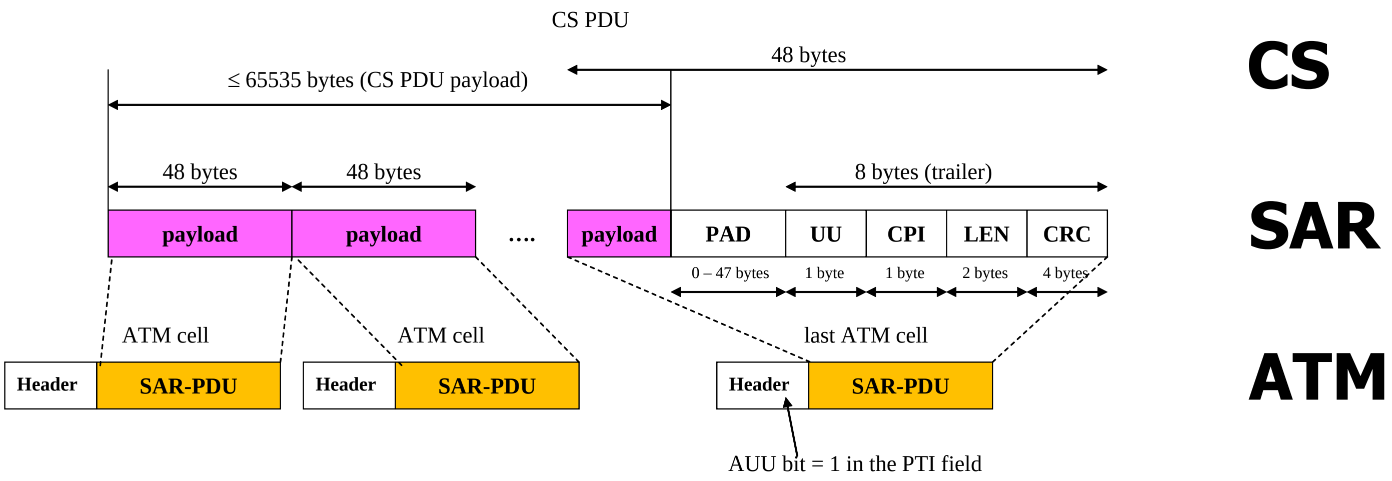

ATM Switches

ATM Switches Architectures

- A switch operates at layer 2 (ATM layer) and realizes the virtual circuit switching by receiving a cell on an input port with a given VPI+VCI and by switching it (according to routing instructions defined in the path set-up phase) to an output port with, in general, a new VPI+VCI couple (so that cell HEC changes).
- Two possible switch architectures can be considered:
  - **ATM cross-connect**: a cell changes only its VPI from input to output,
  - **Typical ATM switch**: a cell changes both VPI and VCI from input to output.

Cross-Connect Switch

- The ATM cross-connect switch can be considered as a first and simplified implementation of an ATM switch since it can switch at most 4096 (= 2^{12}, VPI field contains 12 bits) input virtual circuits.
- In the example below, VPI = 10 is switched to VPI = 40 and VPI = 20 is switched to VPI = 30.

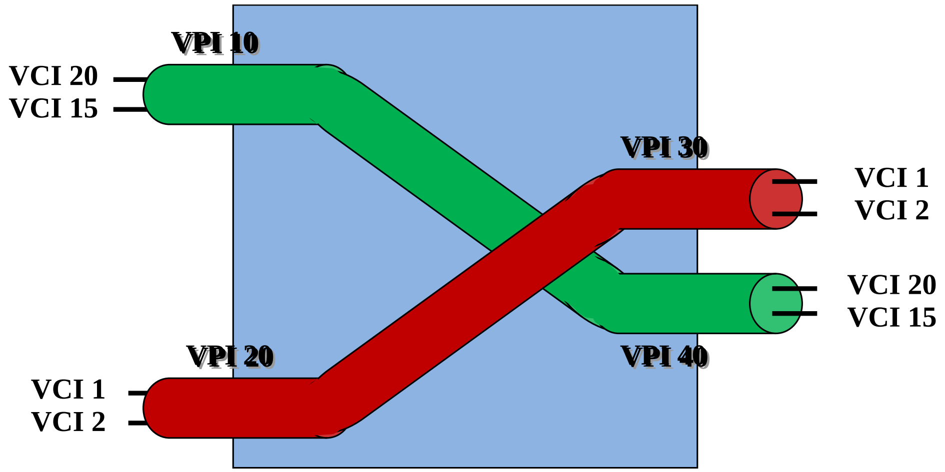

Switching a VPI means to switch all its channels (VCIs)

Classical Switch

- In the example below, the input VCI = 20, VPI = 10 is switched to the output VCI = 1, VPI = 30.

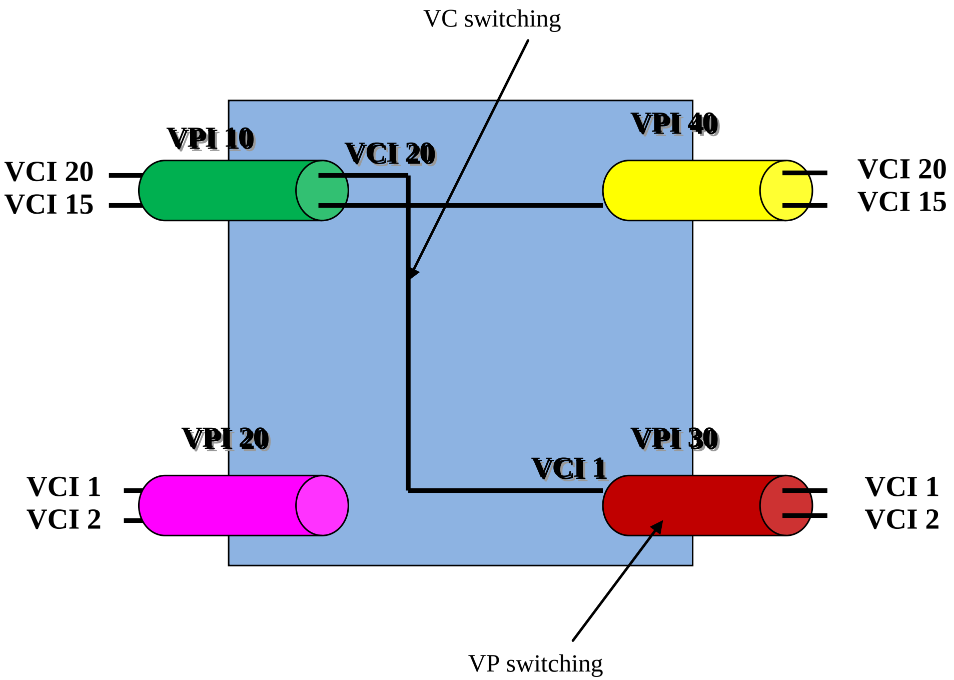

Details on the Switch Architecture

- Three major factors have a large impact on the implementation of the ATM switch architecture:
  - The high speed at which the switch has to operate (from 150 Mbit/s, 2 Gbit/s, etc.).
  - The statistical behavior of the ATM flows crossing the switch.
  - Switching elements use **routing tables** (these tables are almost pre-complied to minimize the switching complexity at call set-up). Input and output ports are associated with (VPI, VCI) couples.

Input buffering uses a dedicated buffer on each input port. Buffers manage cells in a First-Input, First-Output (**FIFO**) basis. A cell at the top of an input buffer may be blocked due to repeated conflicts on the destined output port with cells from other buffers. Such cell blocks all the other cells in the same buffer even if they could be delivered without conflicts to their output lines (**Head-Of-Line blocking, HOL**).

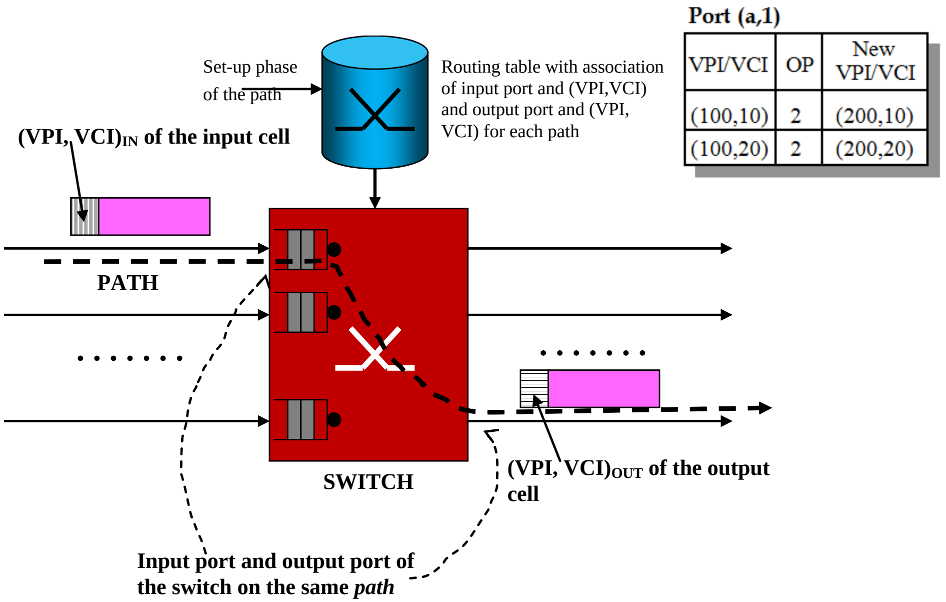

Thank you!

giovanni.giambene@gmail.com

Lesson No. 8 contains the second part of this Lesson.