FORE
SYSTEMS

The 53-Byte Cell

V1.0: Geoff Bennett

ATM = Cells

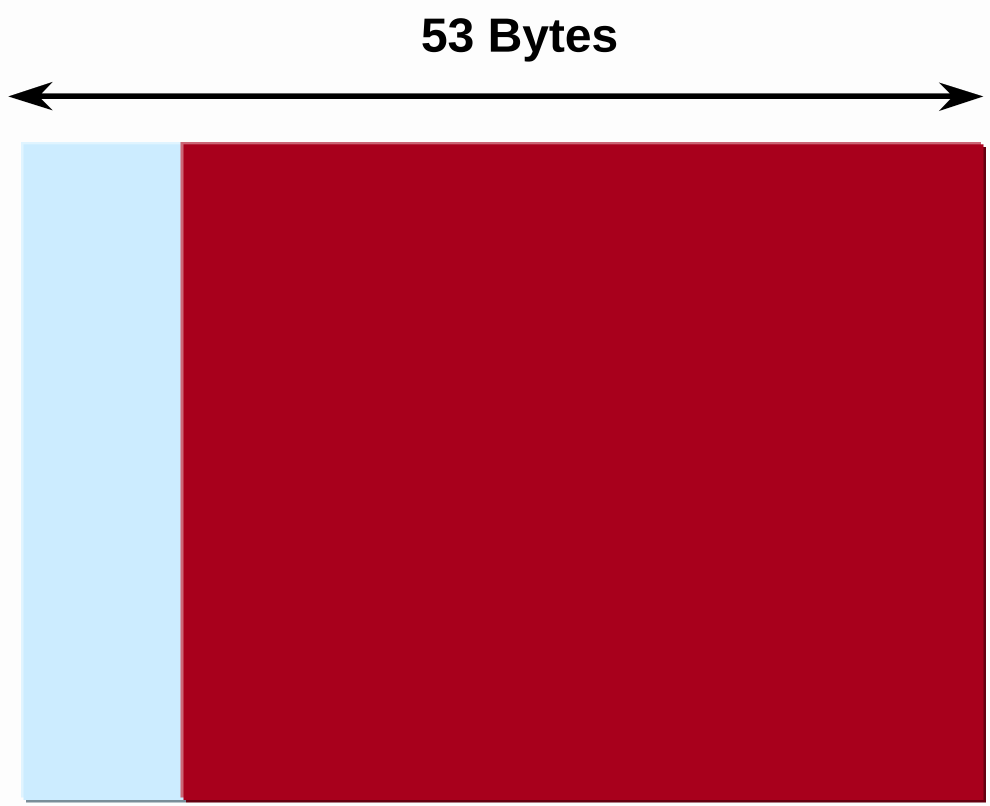

One of the most significant differences between conventional technologies and ATM is that ATM uses simple, fixed length information units called cells. These cells are 53 bytes long, and this module describes the reason they are this size, and the internal structure of the cell.

What is a Cell?

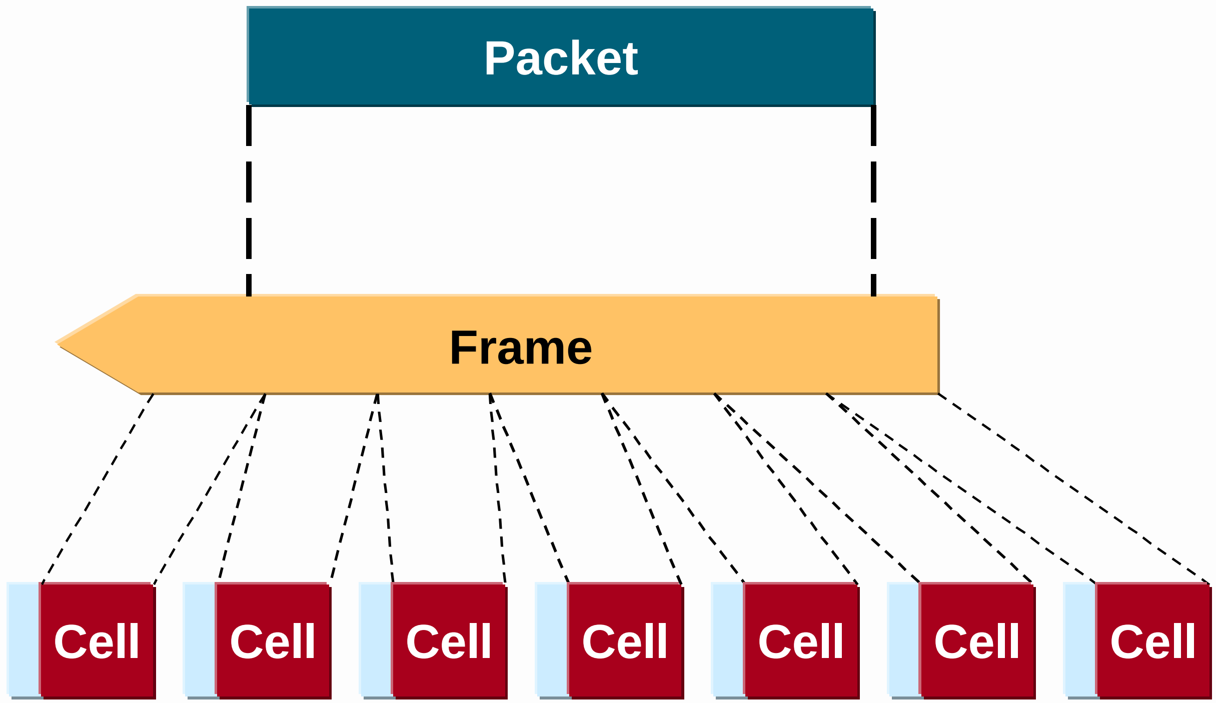

Most of the information units that pass around our networks share two things in common. First, they're variable in size (to adapt efficiently to the amount of data being sent). Second, their maximum size is pretty big. These units in the Data link level are called ***frames***, and examples include Ethernet frames, Token Ring frames and Frame Relay frames. At the Network layer, these units are called ***packets***. Examples include IP or IPX packets. Typically there is a one-for-one mapping of a packet into a frame (another way of saying that an IP packet would normally be carried in one Ethernet frame).

Cells are regularly sized, small information units. If the data requires more than one cell, then it is broken down into cell-sized payloads and packed into multiple cells.

Why Use Cells?

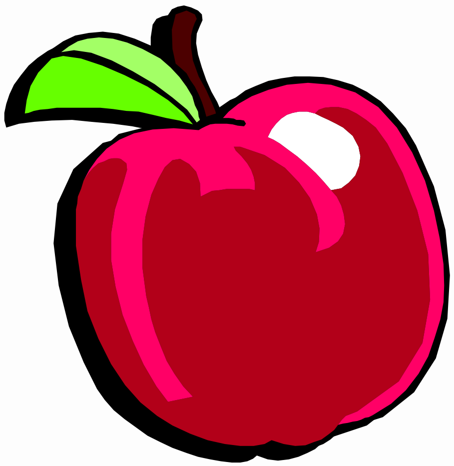

or

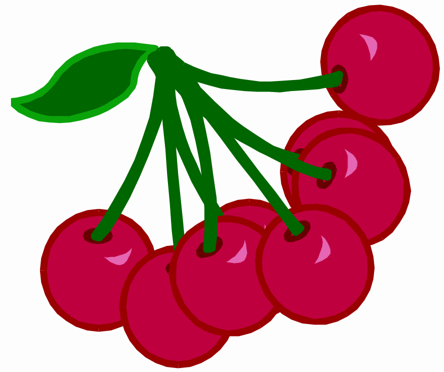

?

Which is easier to share with your friends, and apple or a bunch of cherries?

What’s this got to do with networking?

Why Use Cells?

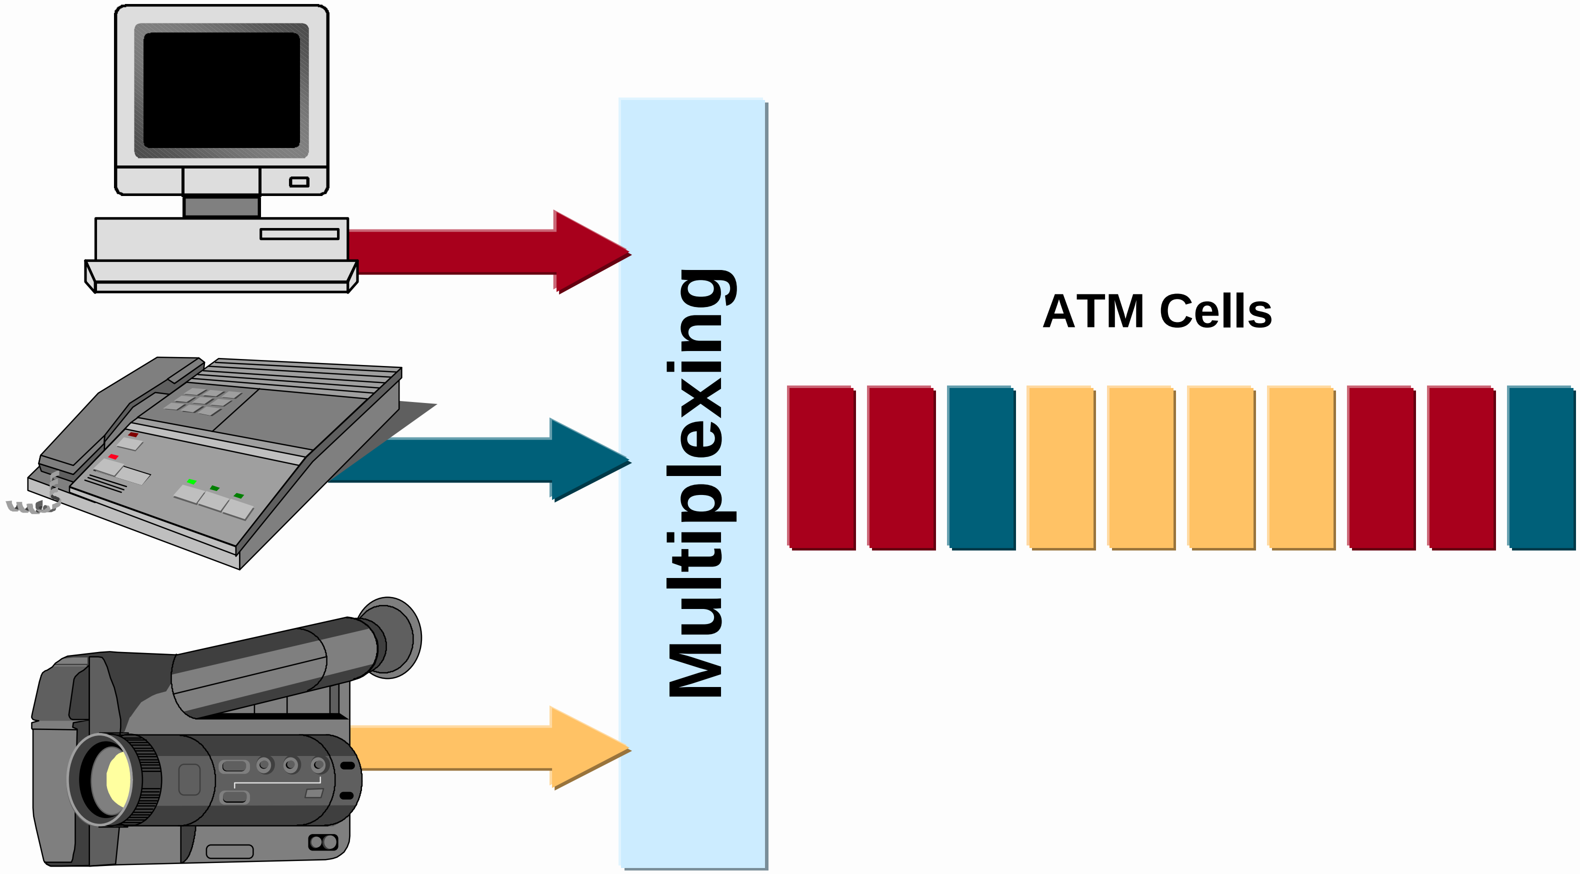

The devices we use to merge, or share information streams perform a function called multiplexing. It's easier to multiplex together separate information streams if the units that we use to carry information are regular in size, and small compared to the transmission speed we're using (like cherries). Smaller units mean that it's possible for higher priority information to "jump the queue" past lower priority information without complex slicing up of the information units.

For bigger information units we need devices to slice up the information, just like trying to share an apple.

Why 53 Bytes?

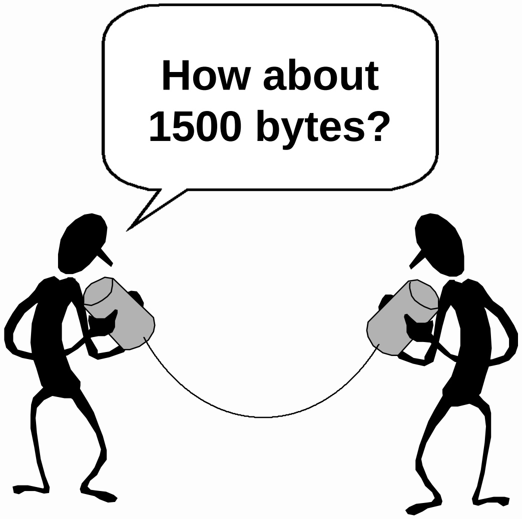

Given that it's a good idea to use cells, what size should they be?

Back in the mid 80's this was a topic of hot debate in the ATM development community, particularly the CCITT.

Everyone knew that a cell needed a certain amount of information to identify where it was supposed to go (and some other information that I'll cover later). This ***cell header*** was non-payload information, or ***overhead***, as far as the user of the network is concerned, and should be kept to a minimum. So there was an argument for large cells.

In fact one group suggested a 1500 byte payload size so ATM could carry Ethernet frames intact.

Why 53 Bytes?

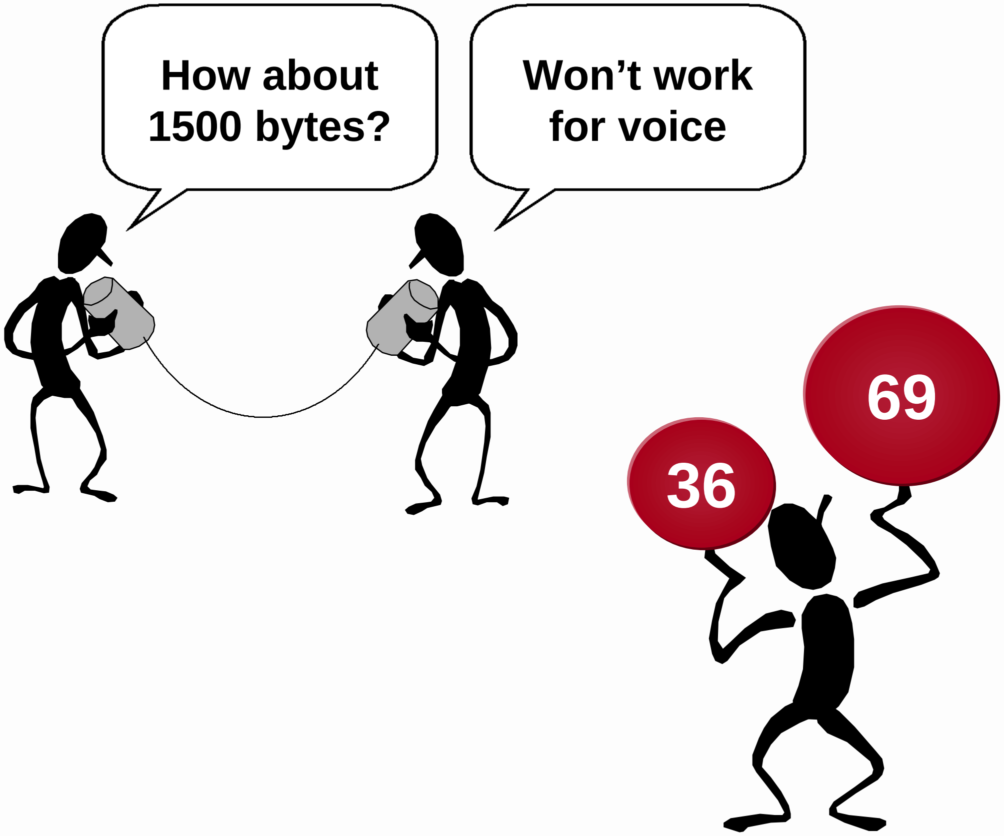

The problem with cells this large is that they take some time (about 12 milliseconds) to pass through a link at 1Mbps. This would cause problems when sharing data with voice connections, and so the idea was dropped.

Two smaller sizes were proposed. A 32-byte payload (with a 4-byte header) by the non-US members, and a 64-byte payload (with a 5-byte header) by the US members.

Members outside the USA wanted the smaller payload to make sure that cell transit delays would be kept below the level that would require echo cancellation equipment for ATM telephony systems.

Why 53 Bytes?

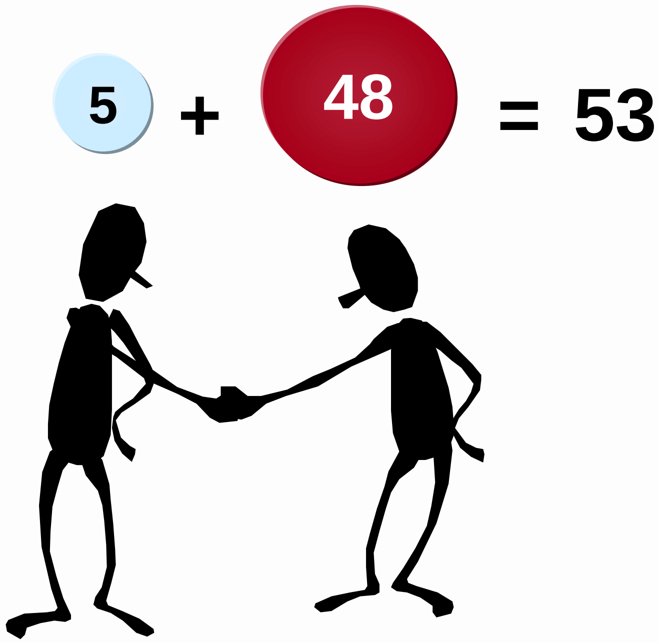

The compromise was to choose a payload halfway between 32 and 64...48 bytes. The 5-byte cell header is added to give the famous 53-byte cell.

Many people criticize the number (it is, after all, a prime number and bears no relationship to any known CPU register structure) but frankly the cell size is irrelevant as long as it's standard.

This decision was made in 1988 by the CCITT (now called the ITU-T).

The Cell Structure

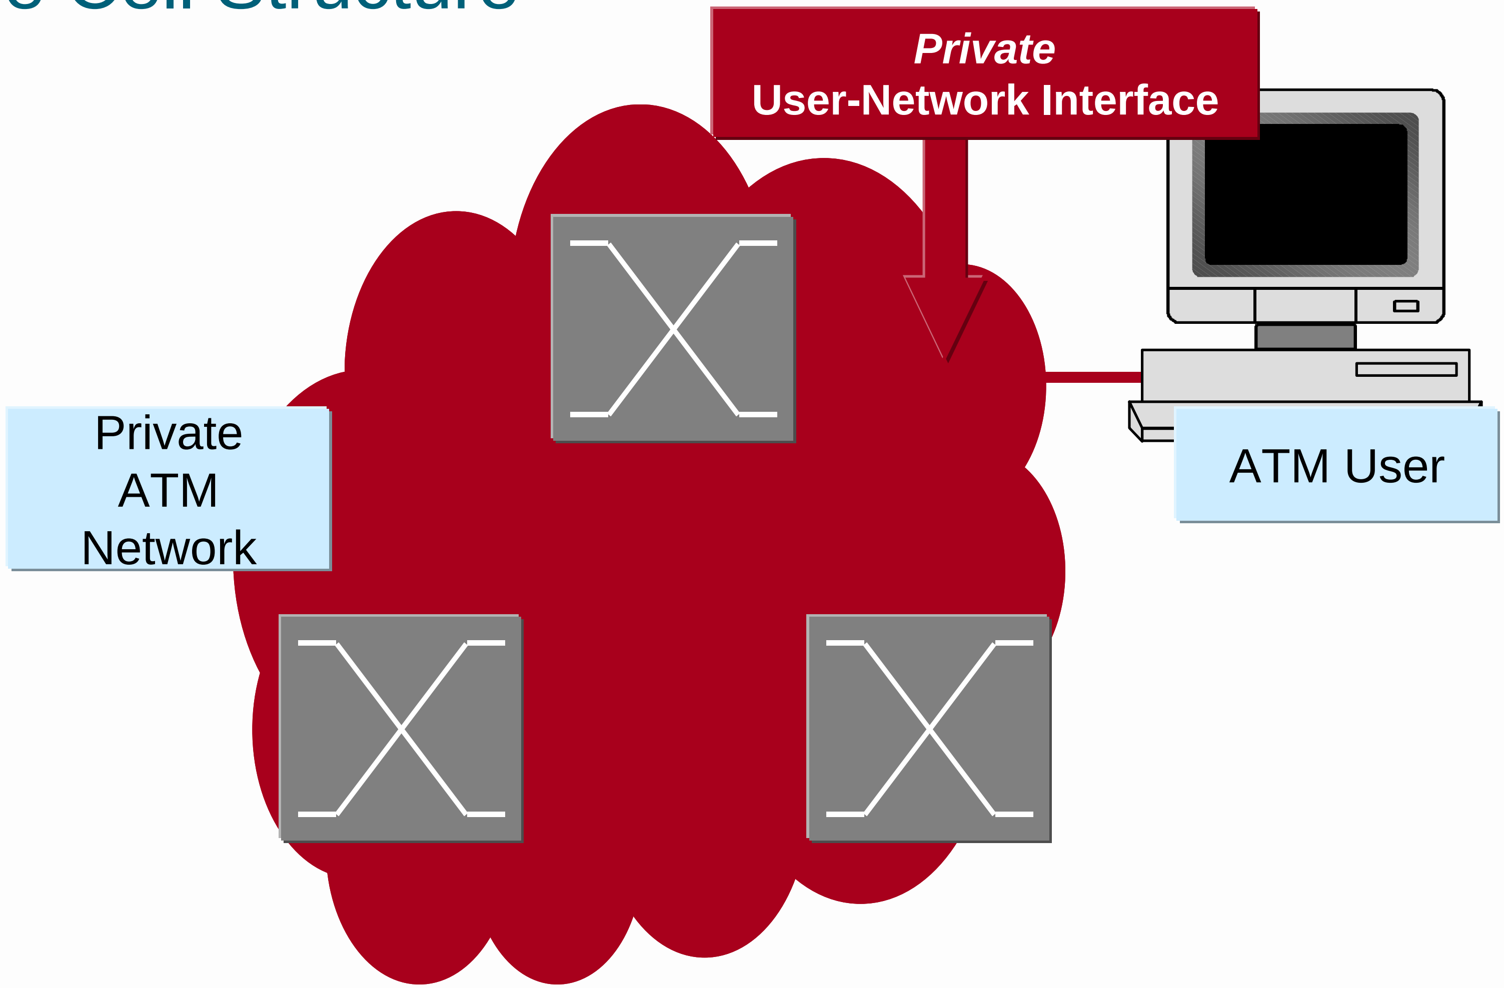

So what does this famous cell look like?

I'll describe the cell structure for the User-Network Interface (UNI). This is a bit of jargon that's explained in another module, but in simple terms it represents the link from an ATM end system (like an ATM-attached workstation or edge switch) and the first switch in the network.

(NOTE: Recently this UNI structure has also been adopted within a PNNI private network).

The Cell Structure

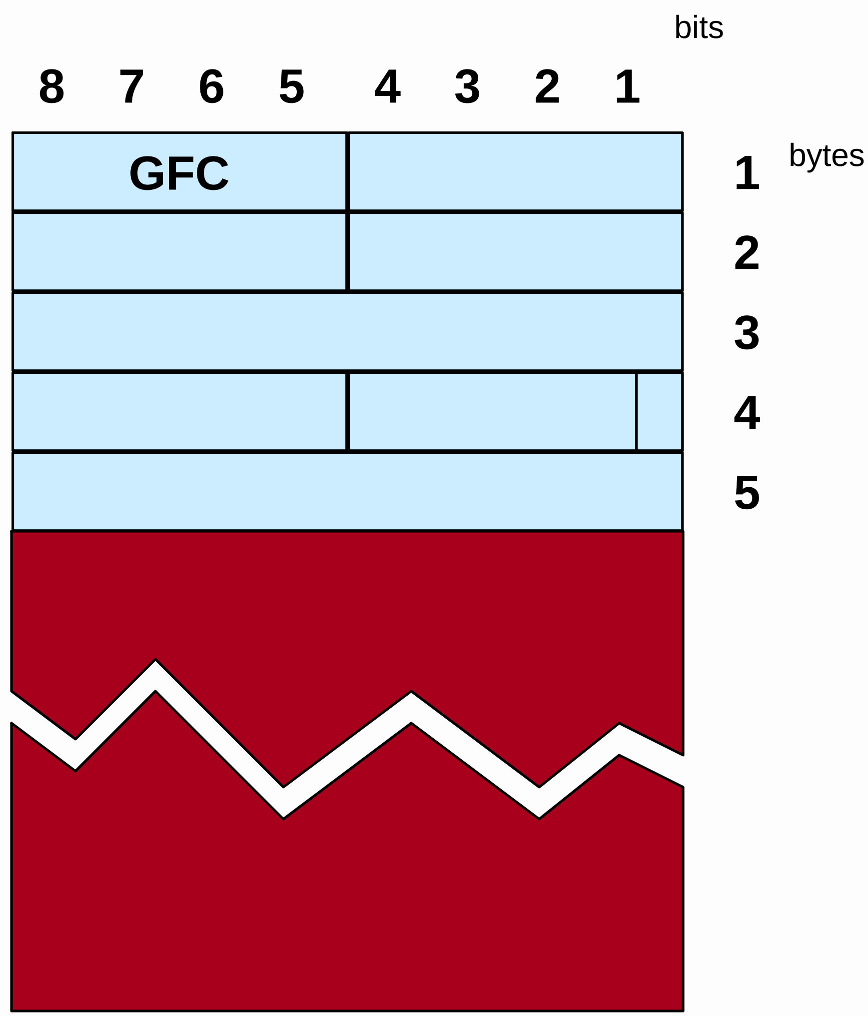

The first part of the header, for an UNI cell only, is the generic flow control (GFC) field. GFC was put into the header to cope with the idea of several subscribers sharing an ATM connection. So far this hasn't happened in ATM networks (which are switched to every connection), and so the GFC field remains un-used. It's possible with the work going on for wireless ATM that GFC may be defined further in the future.

The Cell Structure

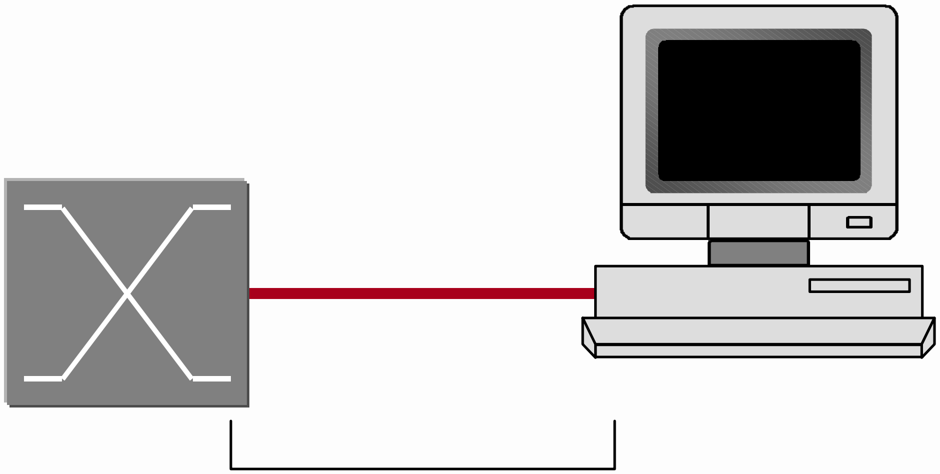

Scope of the local cell address

VPI : VCI

8-bits　　16-bits

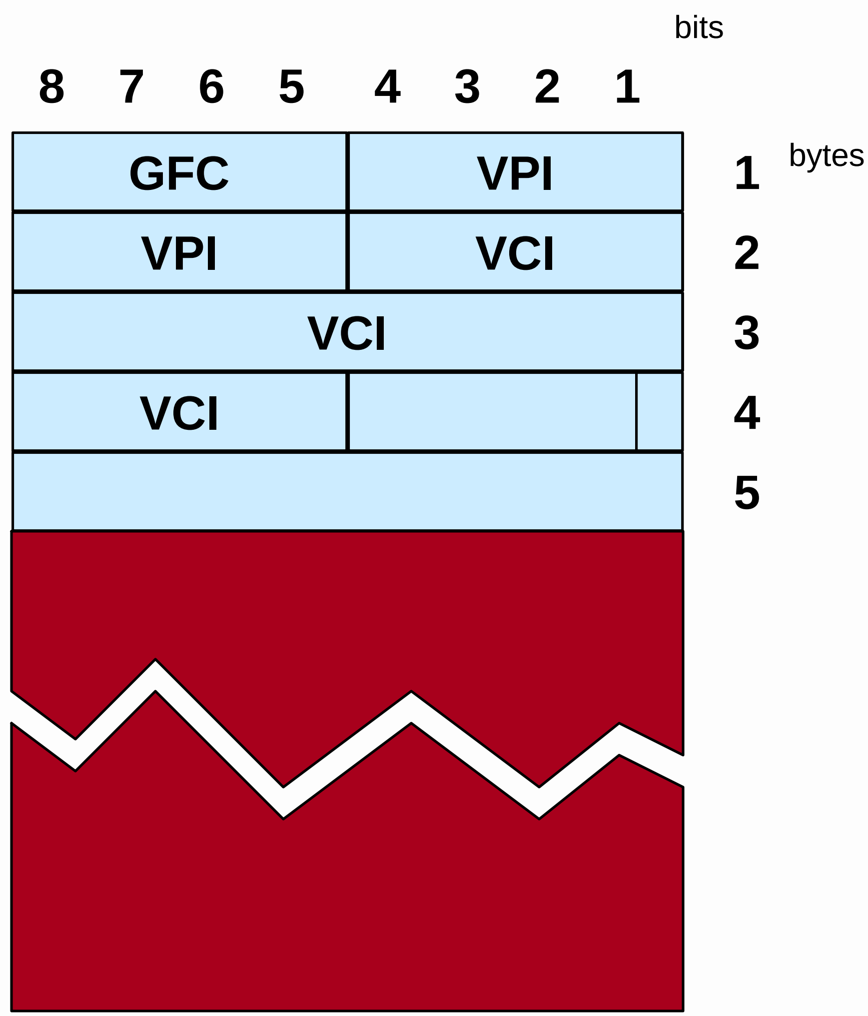

The next two complete fields represent a low level, local address on the ATM cell. This address is “local” to a given link. In other words, the local address may change on the cell as it moves through the network.

The complete address is made up of two components, the Virtual Path Identifier (VPI) and the Virtual Channel Identifier (VCI).

The VPI is 8 bits long, and is intended as a simple aggregation mechanism to allow “bundles” of Virtual Channels to be directed through the network with a minimum of processing.

The VCI refers to an individual, unidirectional connection into a given end system.

The Cell Structure

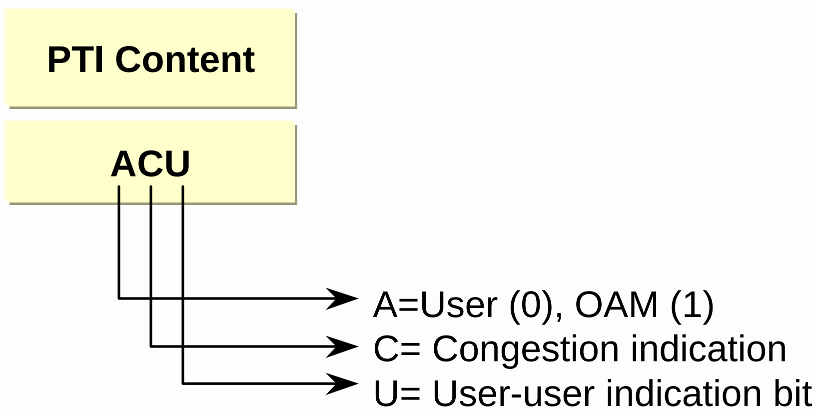

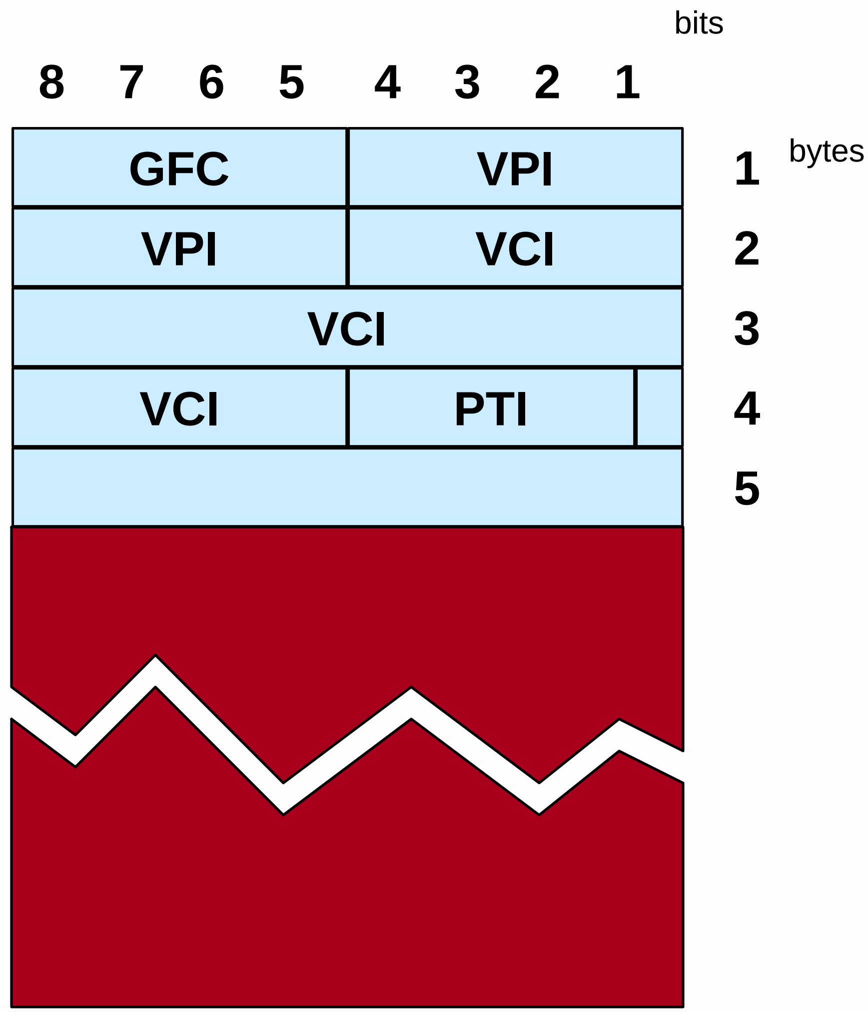

The Payload Type Indicator (PTI) field is three bits long and is used to indicate if this cell is a ***user cell*** (carrying a payload that needs to be switched onwards through the network), or a management (OAM) cell (which may terminate in the switch or edge device, depending on the kind of management cell).

The Congestion Indication bit is used in an ABR flow control mechanism called EFCI.

The user-user indication bit is used to identify high level PDUs (like IP packets) within the cell stream. It's used in congestion management schemes such as Packet Level Discard, which are discussed in the Quality of Service modules.

The Cell Structure

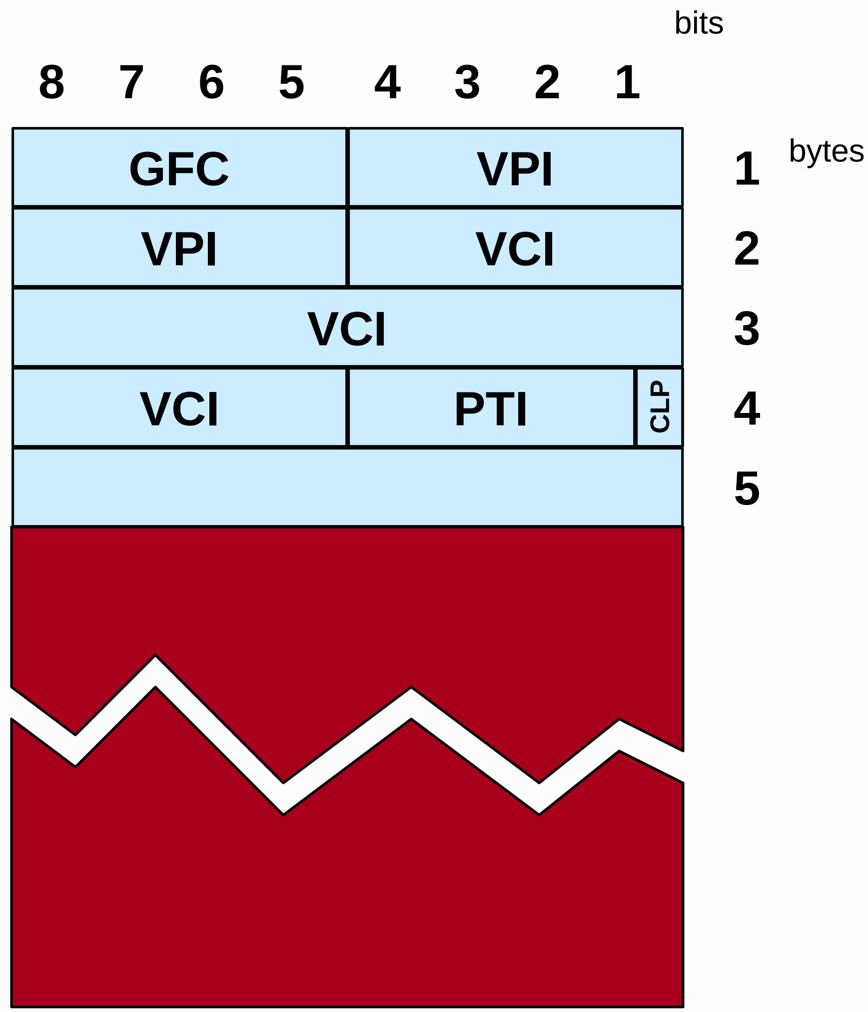

For user cells, the final field of PTI has a special meaning. It's known as the Cell Loss Priority (CLP) bit, and indicates that the cell has broken its traffic contract with the network.

ATM Traffic Policing mechanisms (also discussed in the Quality of Service modules) make use of this bit.

The Cell Structure

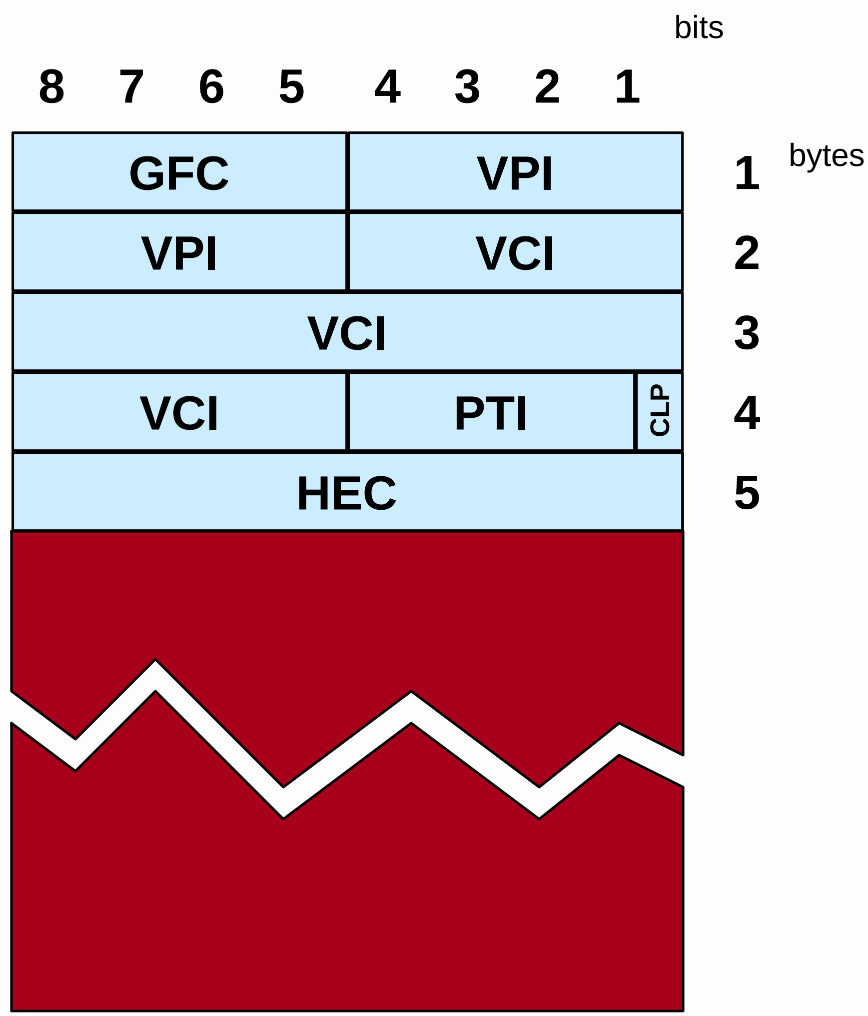

The final field in the header is the Header Error Check (HEC) field. This is a standard error check mechanism that allows the ATM device to correct for a single bit error in the header, or to detect a multiple bit error.

The Cell Structure

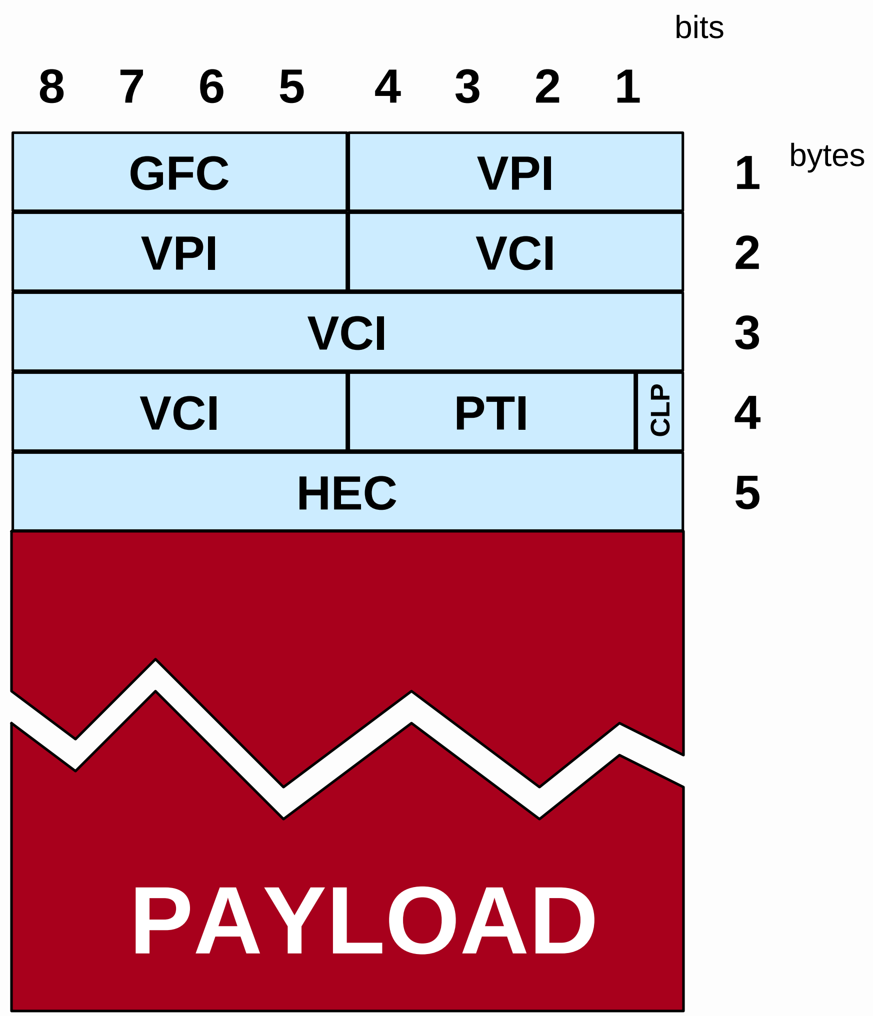

The 48-byte payload of the cell is completely free-format. It either consists of tightly packed bytes of information from higher level PDUs, or of specifically formatted smaller information chunks, such as voice samples.