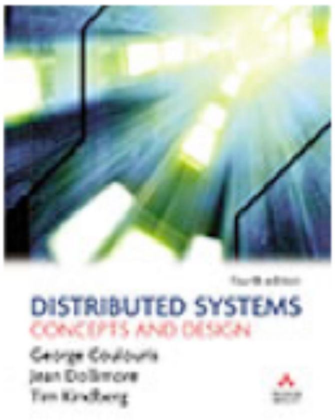

Archive material from *Edition 4* of Distributed Systems: Concepts and Design

Section CDK4–3.5.4

Originally published as pp. 124-127 of Coulouris, Dolllimore and Kindberg, Distributed Systems, Edition 4, 2005.

transmit promotional information to mobile phone users as they pass a store. A useful further reference on Bluetooth networking is the book by Bray and Sturman [2002].

Version 2.0 of the Bluetooth standard, with data throughputs up to 3 mbps – sufficient to carry CD-quality audio – was in the process of approval at the time of writing and improvements including a faster association mechanism and larger Piconet addresses were under development.

3.5.4 Asynchronous Transfer Mode networks

ATM has been designed to carry a wide variety of data including multimedia data such as voice and video. It is a fast packet-switching network based on a method of packet routing known as *cell relay*, which can operate much faster than conventional packet switching. It achieves its speed by avoiding flow control and error checking at the intermediate nodes in a transmission. The transmission links and nodes must therefore have a low likelihood of corrupting data. Another factor affecting the performance is the small, fixed-length units of data transmitted, which reduces buffer size and complexity and queuing delay at intermediate nodes. ATM operates in a connected mode, but a connection can only be set up if sufficient resources are available. Once a connection is

Figure 3.26 ATM protocol layers

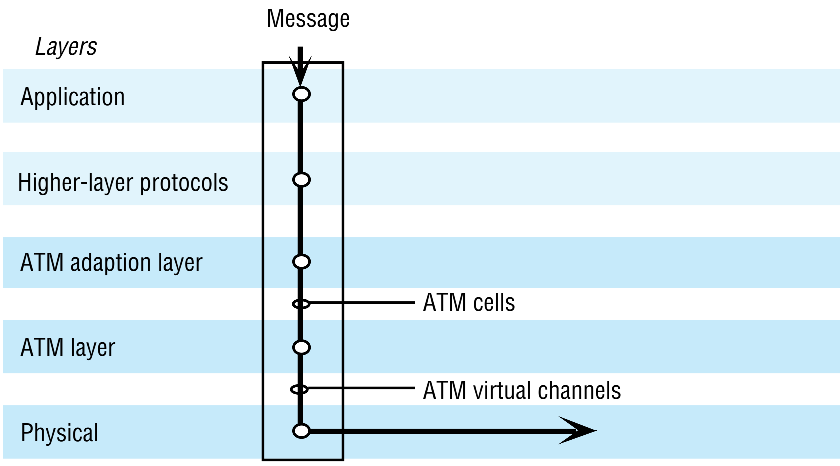

established, its quality (that is, its bandwidth and latency characteristics) can be guaranteed.

ATM is a data-switching technology that can be implemented over existing digital telephony networks, which were hitherto synchronous. When ATM is layered over a network of high-speed synchronous digital links such as those specified for the SONET Synchronous Optical Network [Omidyar and Aldridge 1993], it produces a much more flexible high-speed digital packet network with many virtual connections. Each ATM virtual connection provides bandwidth and latency guarantees. The resulting virtual circuits can be used to support a wide range of services with varying speeds. These include voice (32 kbps), fax, distributed systems services, video and high-definition television (100–150 Mbps). The ATM [CCITT 1990] standard recommends the provision of virtual circuits with data transfer rates of up to 155 Mbps or 622 Mbps.

ATM networks can also be implemented in *native mode* directly over optical fibre, copper and other transmission media, allowing bandwidths of up to several gigabits per second with current fibre technology. This is the mode in which it is employed on local and metropolitan area networks.

The ATM service is structured in three layers, represented by the darker panels in Figure 3.26. The *ATM adaptation layer* is an end-to-end layer implemented only at the sending and receiving hosts. It is intended to support existing higher-level protocols such as TCP/IP and X25 over the ATM layer. Different versions of the adaptation layer can provide a variety of different adaptation functions to suit the requirements of different higher-level protocols. It will include some common functions such as packet assembly and disassembly for use in building specific higher-level protocols.

The *ATM layer* provides a connection-oriented service that transmits fixed length packets called *cells*. A connection consists of a sequence of virtual channels within virtual paths. A *virtual channel* (VC) is a logical unidirectional association between two endpoints of a link in the physical path from source to destination. A *virtual path* (VP) is a bundle of virtual channels that are associated with a physical path between two switching nodes. Virtual paths are intended to be used to support semi-permanent

Figure 3.27 ATM cell layout

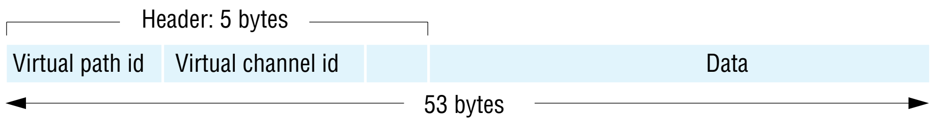

connections between pairs of end-points. Virtual channels are allocated dynamically when connections are set up.

The nodes in an ATM network can play three distinct roles:

- *hosts,* that send and receive messages;
- *VP switches,* that hold tables showing the correspondence between incoming and outgoing virtual paths;
- *VP/VC switches,* that hold similar tables for both virtual paths and virtual channels.

An ATM cell has a 5-byte header and a 48-byte data field (Figure 3.27). The full data field is always sent, even when it is only partially filled with data. The header contains an identifier for a virtual channel and an identifier for a virtual path, which together provide the information required to route the cell across the network. The virtual path identifier refers to a particular virtual path on the physical link on which the cell is transmitted. The virtual channel identifier refers to one specific virtual channel inside the virtual path. Other header fields are used to indicate the type of cell, its cell loss priority and the cell boundary.

When a cell arrives at a VP switch, the virtual path identifier in the header is looked up in its routing table to work out the corresponding virtual path identifier for the

Figure 3.28 Switching virtual paths in an ATM network

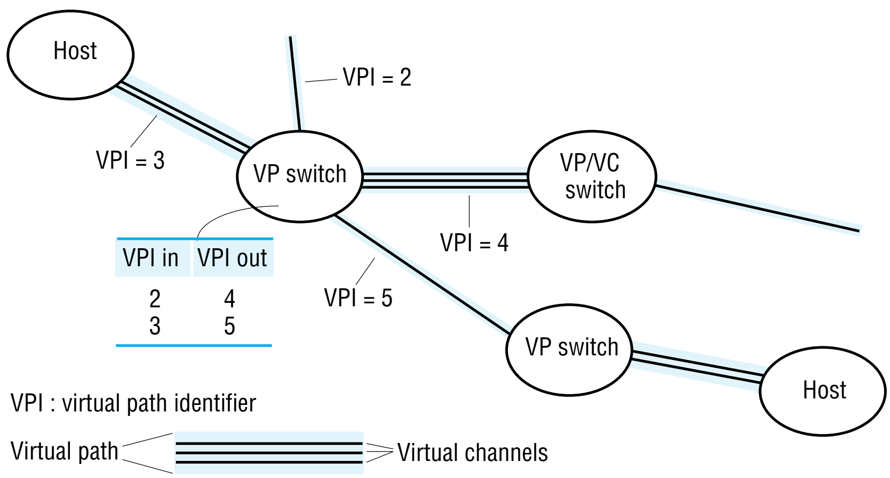

outgoing physical path; see Figure 3.28. It puts the new virtual path identifier in the header and then transmits the cell on the outgoing physical path. A VP/VC switch can perform similar routing based on both VP and VC identifiers.

Note that the VP and VC identifiers are defined locally. This scheme has the advantage that there is no need for global network-wide identifiers, which would need to be very large numbers. A global addressing scheme would introduce administrative overheads and would require cell headers and the tables in switches to hold more information.

ATM provides a service with low latency – the switching delay is about 25 microseconds per switch, giving, for example, a latency of 250 microseconds when a message passes through ten switches. This compares well with our estimated performance requirements for distributed systems (Section 3.2), suggesting that an ATM network will support interprocess communication and client-server interactions with a performance similar to, or better than, that now available from local area networks. Very high-bandwidth channels with guaranteed quality of service, suitable for transmitting streams of multimedia data at speeds up to 600 Mbps, will also be available. Gigabits per second are attainable in pure ATM networks.

3.6 Summary

We have focused on the networking concepts and techniques that are needed as a basis for distributed systems and have approached them from the point of view of a distributed system designer. Packet networks and layered protocols provide the basis for communication in distributed systems. Local area networks are based on packet broadcasting on a shared medium; Ethernet is the dominant technology. Wide area networks are based on packet switching to route packets to their destinations through a connected network. Routing is a key mechanism and a variety of routing algorithms are used, of which the distance-vector method is the most basic but effective. Congestion control is needed to prevent overflow of buffers at the receiver and at intermediate nodes.

Internetworks are constructed by layering a 'virtual' internetwork protocol over collections of networks linked together by routers. The Internet TCP/IP protocols enable computers in the Internet to communicate with one another in a uniform manner, irrespective of whether they are on the same local area network or in different countries. The Internet standards include many application-level protocols that are suitable for use in wide area distributed applications. IPv6 has the much larger address space needed for the future evolution of the Internet and provision for new application requirements such as quality of service and security.

Mobile users are supported by MobileIP for wide-area roaming and by wireless LANs based on IEEE 802 standards for local connectivity. ATM offers very high-bandwidth asynchronous communication based on virtual circuits with guaranteed quality of service.