Asynchronous Transfer Mode

Asynchronous Transfer Mode (**ATM**) is a telecommunications standard defined by ANSI and ITU (formerly CCITT) for digital transmission of multiple types of traffic, including telephony (voice), data, and video signals in one network without the use of separate overlay networks.[1][2] ATM was developed to meet the needs of the Broadband Integrated Services Digital Network, as defined in the late 1980s,[3] and designed to integrate telecommunication networks. It can handle both traditional high-throughput data traffic and real-time, low-latency content such as voice and video. ATM provides functionality that uses features of circuit switching and packet switching networks. It uses asynchronous time-division multiplexing,[4][5] and encodes data into small, fixed-sized network packets.

In the OSI reference model data link layer (layer 2), the basic transfer units are generically called frames. In ATM these frames are of a fixed (53 octets or bytes) length and specifically called *cells*. This differs from approaches such as IP or Ethernet that use variable sized packets or frames. ATM uses a connection-oriented model in which a virtual circuit must be established between two endpoints before the data exchange begins.[5] These virtual circuits may be either permanent, i.e. dedicated connections that are usually preconfigured by the service provider, or switched, i.e. set up on a per-call basis using signaling and disconnected when the call is terminated.

The ATM network reference model approximately maps to the three lowest layers of the OSI model: physical layer, data link layer, and network layer.[6] ATM is a core protocol used in the SONET/SDH backbone of the public switched telephone network (PSTN) and in the Integrated Services Digital Network (ISDN), but has largely been superseded in favor of next-generation networks based on Internet Protocol (IP) technology, while wireless and mobile ATM never established a significant foothold.

Contents

IBM Turboways ATM 155 PCI network interface card

Protocol architecture

If a speech signal is reduced to packets, and it is forced to share a link with bursty data traffic (traffic with an abnormally large number of packets over a brief period of time, such as could occur during a large scale emergency and the cellular network has become oversubscribed) then no matter how small the speech packets could be made, they would always encounter full-size data packets. Under normal queuing conditions the cells might experience maximum queuing delays. To avoid this issue, all ATM packets, or "cells," are

Marconi ForeRunnerLE 25 ATM PCI network interface card

the same small size. In addition, the fixed cell structure means that ATM can be readily switched by hardware without the inherent delays introduced by software switched and routed frames.

Thus, the designers of ATM utilized small data cells to reduce jitter (large changes in the packet round trip time or one-way time, in this case) in the multiplexing of data streams. Reduction of jitter (and also end-to-end round-trip delays) is particularly important when carrying voice traffic, because the conversion of digitized voice into an analogue audio signal is an inherently real-time process, and to do a good job, the decoder that does this needs an evenly spaced (in time) stream of data items. If the next data item is not available when it is needed, the codec has no choice but to produce silence or guess – and if the data is late, it is useless, because the time period when it should have been converted to a signal has already passed.

At the time of the design of ATM, 155 Mbit/s synchronous digital hierarchy (SDH) with 135 Mbit/s payload was considered a fast optical network link, and many plesiochronous digital hierarchy (PDH) links in the digital network were considerably slower, ranging from 1.544 to 45 Mbit/s in the US, and 2 to 34 Mbit/s in Europe.

At 155 Mbit/s, a typical full-length 1,500 byte (12,000-bit) data packet, sufficient to contain a maximum-sized IP packet for Ethernet, would take 77.42 µs to transmit. In a lower-speed link, such as a 1.544 Mbit/s T1 line, the same packet would take up to 7.8 milliseconds.

A queuing delay induced by several such data packets might exceed the figure of 7.8 ms several times over, in addition to any packet generation delay in the shorter speech packet. This was considered unacceptable for speech traffic, which needs to have low jitter in the data stream being fed into the codec if it is to produce good-quality sound. A packet voice system can produce this low jitter in a number of ways:

- Using a playback buffer between the network and the codec, one large enough to tide the codec over almost all the jitter in the data. This allows smoothing out the jitter, but the delay introduced by passage through the buffer require echo cancellers even in local networks; this was considered too expensive at the time. Also, it increased the delay across the channel, and made conversation difficult over high-delay channels.
- Using a system that inherently provides low jitter (and minimal overall delay) to traffic that needs it.
- Operate on a 1:1 user basis (i.e., a dedicated pipe).

The design of ATM aimed for a low-jitter network interface. However, "cells" were introduced into the design to provide short queuing delays while continuing to support datagram traffic. ATM broke up all packets, data, and voice streams into 48-byte chunks, adding a 5-byte routing header to each one so that they could be

reassembled later. The choice of 48 bytes was political rather than technical.[7] When the CCITT (now ITU-T) was standardizing ATM, parties from the United States wanted a 64-byte payload because this was felt to be a good compromise in larger payloads optimized for data transmission and shorter payloads optimized for real-time applications like voice; parties from Europe wanted 32-byte payloads because the small size (and therefore short transmission times) simplify voice applications with respect to echo cancellation. Most of the European parties eventually came around to the arguments made by the Americans, but France and a few others held out for a shorter cell length. With 32 bytes, France would have been able to implement an ATM-based voice network with calls from one end of France to the other requiring no echo cancellation. 48 bytes (plus 5 header bytes = 53) was chosen as a compromise between the two sides. 5-byte headers were chosen because it was thought that 10% of the payload was the maximum price to pay for routing information.[3] ATM multiplexed these 53-byte cells instead of packets which reduced worst-case cell contention jitter by a factor of almost 30, reducing the need for echo cancellers.

Cell structure

An ATM cell consists of a 5-byte header and a 48-byte payload. The payload size of 48 bytes was chosen as described above.

ATM defines two different cell formats: user–network interface (UNI) and network–network interface (NNI). Most ATM links use UNI cell format.
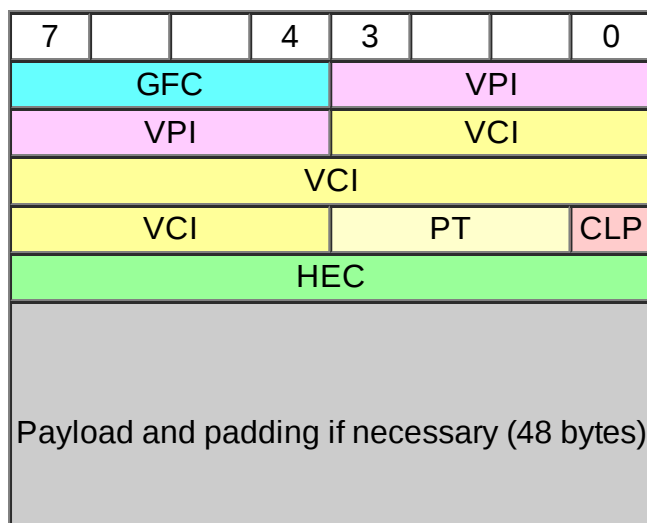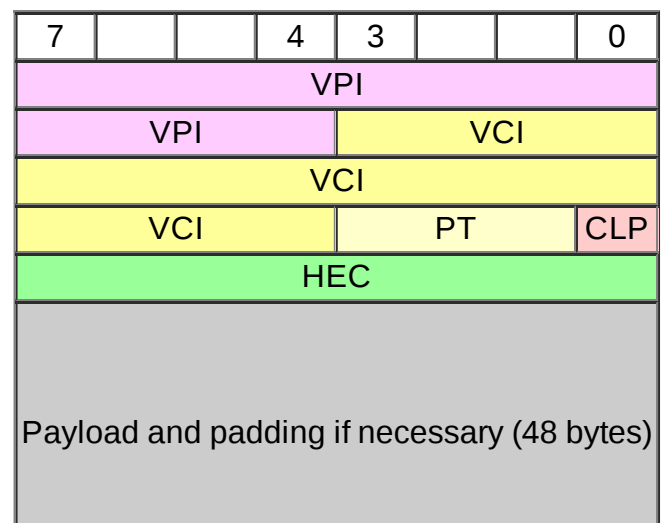
Diagram of a UNI ATM cell

7			4	3			0
GFC				VPI			
VPI				VCI			
VCI							
VCI				PT			CLP
HEC							
Payload and padding if necessary (48 bytes)							

Diagram of an NNI ATM cell

7			4	3			0
VPI							
VPI				VCI			
VCI							
VCI				PT			CLP
HEC							
Payload and padding if necessary (48 bytes)							

GFC = The generic flow control (GFC) field is a 4-bit field that was originally added to support the connection of ATM networks to shared access networks such as a distributed queue dual bus (DQDB) ring. The GFC field was designed to give the User-Network Interface (UNI) 4 bits in which to negotiate multiplexing and flow control among the cells of various ATM connections. However, the use and exact values of the GFC field have not been standardized, and the field is always set to 0000.[8]
VPI = Virtual path identifier (8 bits UNI, or 12 bits NNI)
VCI = Virtual channel identifier (16 bits)
PT = Payload type (3 bits)

PT bit 3 (msbit): Network management cell. If 0, user data cell and the following apply:
PT bit 2: Explicit forward congestion indication (EFCI); 1 = network congestion experienced

PT bit 1 (lsbit): ATM user-to-user (AAU) bit. Used by AAL5 to indicate packet boundaries.

CLP = Cell loss priority (1-bit)
HEC = Header error control (8-bit CRC, polynomial = $X^8 + X^2 + X + 1$)

ATM uses the PT field to designate various special kinds of cells for operations, administration and management (OAM) purposes, and to delineate packet boundaries in some ATM adaptation layers (AAL). If the most significant bit (MSB) of the PT field is 0, this is a user data cell, and the other two bits are used to indicate network congestion and as a general purpose header bit available for ATM adaptation layers. If the MSB is 1, this is a management cell, and the other two bits indicate the type. (Network management segment, network management end-to-end, resource management, and reserved for future use.)

Several ATM link protocols use the HEC field to drive a CRC-based framing algorithm, which allows locating the ATM cells with no overhead beyond what is otherwise needed for header protection. The 8-bit CRC is used to correct single-bit header errors and detect multi-bit header errors. When multi-bit header errors are detected, the current and subsequent cells are dropped until a cell with no header errors is found.

A UNI cell reserves the GFC field for a local flow control/submultiplexing system between users. This was intended to allow several terminals to share a single network connection, in the same way that two Integrated Services Digital Network (ISDN) phones can share a single basic rate ISDN connection. All four GFC bits must be zero by default.

The NNI cell format replicates the UNI format almost exactly, except that the 4-bit GFC field is re-allocated to the VPI field, extending the VPI to 12 bits. Thus, a single NNI ATM interconnection is capable of addressing almost 2^{12} VPs of up to almost 2^{16} VCs each (in practice some of the VP and VC numbers are reserved).

Service types

ATM supports different types of services via AALs. Standardized AALs include AAL1, AAL2, and AAL5, and the rarely used AAL3 and AAL4. AAL1 is used for constant bit rate (CBR) services and circuit emulation. Synchronization is also maintained at AAL1. AAL2 through AAL4 are used for variable bitrate (VBR) services, and AAL5 for data. Which AAL is in use for a given cell is not encoded in the cell. Instead, it is negotiated by or configured at the endpoints on a per-virtual-connection basis.

Following the initial design of ATM, networks have become much faster. A 1500 byte (12000-bit) full-size Ethernet frame takes only 1.2 µs to transmit on a 10 Gbit/s network, reducing the need for small cells to reduce jitter due to contention. Some consider that this makes a case for replacing ATM with Ethernet in the network backbone. The increased link speeds by themselves do not alleviate jitter due to queuing. Additionally, the hardware for implementing the service adaptation for IP packets is expensive at very high speeds. Specifically, at speeds of OC-3 and above, the cost of segmentation and reassembly (SAR) hardware makes ATM less competitive for IP than Packet Over SONET (POS);[9] because of its fixed 48-byte cell payload, ATM is not suitable as a data link layer *directly* underlying IP (without the need for SAR at the data link level) since the OSI layer on which IP operates must provide a maximum transmission unit (MTU) of at least 576 bytes. SAR performance limits mean that the fastest IP router ATM interfaces are STM16 - STM64 which actually compares, while as of 2004 POS can operate at OC-192 (STM64) with higher speeds expected in the future, limits based on segmentation and reassembly (SAR).

On slower or congested links (622 Mbit/s and below), ATM does make sense, and for this reason most asymmetric digital subscriber line (ADSL) systems use ATM as an intermediate layer between the physical link layer and a Layer 2 protocol like PPP or Ethernet.[10]

At these lower speeds, ATM provides a useful ability to carry multiple logical circuits on a single physical or virtual medium, although other techniques exist, such as Multi-link PPP and Ethernet VLANs, which are optional in VDSL implementations. DSL can be used as an access method for an ATM network, allowing a DSL termination point in a telephone central office to connect to many internet service providers across a wide-area ATM network. In the United States, at least, this has allowed DSL providers to provide DSL access to the customers of many internet service providers. Since one DSL termination point can support multiple ISPs, the economic feasibility of DSL is substantially improved.

Virtual circuits

A network must establish a connection before two parties can send cells to each other. In ATM this is called a virtual circuit (VC). It can be a permanent virtual circuit (PVC), which is created administratively on the end points, or a switched virtual circuit (SVC), which is created as needed by the communicating parties. SVC creation is managed by signaling, in which the requesting party indicates the address of the receiving party, the type of service requested, and whatever traffic parameters may be applicable to the selected service. "Call admission" is then performed by the network to confirm that the requested resources are available and that a route exists for the connection.

Motivation

ATM operates as a channel-based transport layer, using VCs. This is encompassed in the concept of the virtual paths (VP) and virtual channels. Every ATM cell has an 8- or 12-bit virtual path identifier (VPI) and 16-bit virtual channel identifier (VCI) pair defined in its header.[11] The VCI, together with the VPI, is used to identify the next destination of a cell as it passes through a series of ATM switches on its way to its destination. The length of the VPI varies according to whether the cell is sent on the user-network interface (on the edge of the network), or if it is sent on the network-network interface (inside the network).

As these cells traverse an ATM network, switching takes place by changing the VPI/VCI values (label swapping). Although the VPI/VCI values are not necessarily consistent from one end of the connection to the other, the concept of a circuit *is* consistent (unlike IP, where any given packet could get to its destination by a different route than the others).[12] ATM switches use the VPI/VCI fields to identify the virtual channel link (VCL) of the next network that a cell needs to transit on its way to its final destination. The function of the VCI is similar to that of the data link connection identifier (DLCI) in frame relay and the logical channel number and logical channel group number in X.25.

Another advantage of the use of virtual circuits comes with the ability to use them as a multiplexing layer, allowing different services (such as voice, frame relay, n* 64 channels, IP). The VPI is useful for reducing the switching table of some virtual circuits which have common paths.[13]

Types

ATM can build virtual circuits and virtual paths either statically or dynamically. Static circuits (permanent virtual circuits or PVCs) or paths (permanent virtual paths or PVPs) require that the circuit is composed of a series of segments, one for each pair of interfaces through which it passes.

PVPs and PVCs, though conceptually simple, require significant effort in large networks. They also do not support the re-routing of service in the event of a failure. Dynamically built PVPs (soft PVPs or SPVPs) and PVCs (soft PVCs or SPVCs), in contrast, are built by specifying the characteristics of the circuit (the service "contract") and the two end points.

ATM networks create and remove switched virtual circuits (SVCs) on demand when requested by an end piece of equipment. One application for SVCs is to carry individual telephone calls when a network of telephone switches are inter-connected using ATM. SVCs were also used in attempts to replace local area networks with ATM.

Routing

Most ATM networks supporting SPVPs, SPVCs, and SVCs use the Private Network Node Interface or the Private Network-to-Network Interface (PNNI) protocol to share topology information between switches and select a route through a network. PNNI is a link-state routing protocol like OSPF and IS-IS. PNNI also includes a very powerful route summarization mechanism to allow construction of very large networks, as well as a call admission control (CAC) algorithm which determines the availability of sufficient bandwidth on a proposed route through a network in order to satisfy the service requirements of a VC or VP.

Traffic engineering

Another key ATM concept involves the traffic contract. When an ATM circuit is set up each switch on the circuit is informed of the traffic class of the connection.

ATM traffic contracts form part of the mechanism by which "quality of service" (QoS) is ensured. There are four basic types (and several variants) which each have a set of parameters describing the connection.

1. CBR - Constant bit rate: a Peak Cell Rate (PCR) is specified, which is constant.
2. VBR - Variable bit rate: an average or Sustainable Cell Rate (SCR) is specified, which can peak at a certain level, a PCR, for a maximum interval before being problematic.
3. ABR - Available bit rate: a minimum guaranteed rate is specified.
4. UBR - Unspecified bit rate: traffic is allocated to all remaining transmission capacity.

VBR has real-time and non-real-time variants, and serves for "bursty" traffic. Non-real-time is sometimes abbreviated to vbr-nrt.

Most traffic classes also introduce the concept of Cell-delay variation tolerance (CDVT), which defines the "clumping" of cells in time.

Traffic policing

To maintain network performance, networks may apply traffic policing to virtual circuits to limit them to their traffic contracts at the entry points to the network, i.e. the user–network interfaces (UNIs) and network-to-network interfaces (NNIs): usage/network parameter control (UPC and NPC).[14] The reference model given by the ITU-T and ATM Forum for UPC and NPC is the generic cell rate algorithm (GCRA),[15][16] which is a version of the leaky bucket algorithm. CBR traffic will normally be policed to a PCR and CDVt alone, whereas VBR traffic will normally be policed using a dual leaky bucket controller to a PCR and CDVt and an SCR and Maximum Burst Size (MBS). The MBS will normally be the packet (SAR-SDU) size for the VBR VC in cells.

If the traffic on a virtual circuit is exceeding its traffic contract, as determined by the GCRA, the network can either drop the cells or mark the Cell Loss Priority (CLP) bit (to identify a cell as potentially redundant). Basic policing works on a cell by cell basis, but this is sub-optimal for encapsulated packet traffic (as discarding a single cell will invalidate the whole packet). As a result, schemes such as partial packet discard (PPD) and early packet discard (EPD) have been created that will discard a whole series of cells until the next packet

starts. This reduces the number of useless cells in the network, saving bandwidth for full packets. EPD and PPD work with AAL5 connections as they use the end of packet marker: the ATM user-to-ATM user (AUU) indication bit in the payload-type field of the header, which is set in the last cell of a SAR-SDU.

Traffic shaping

Traffic shaping usually takes place in the network interface card (NIC) in user equipment, and attempts to ensure that the cell flow on a VC will meet its traffic contract, i.e. cells will not be dropped or reduced in priority at the UNI. Since the reference model given for traffic policing in the network is the GCRA, this algorithm is normally used for shaping as well, and single and dual leaky bucket implementations may be used as appropriate.

Reference model

The ATM network reference model approximately maps to the three lowest layers of the OSI reference model. It specifies the following layers:[17]

- At the physical network level, ATM specifies a layer that is equivalent to the OSI physical layer.
- The ATM layer 2 roughly corresponds to the OSI data link layer.
- The OSI network layer is implemented as the ATM adaptation layer (AAL).

Deployment

ATM became popular with telephone companies and many computer makers in the 1990s. However, even by the end of the decade, the better price/performance of Internet Protocol-based products was competing with ATM technology for integrating real-time and bursty network traffic.[18] Companies such as FORE Systems focused on ATM products, while other large vendors such as Cisco Systems provided ATM as an option.[19] After the burst of the dot-com bubble, some still predicted that "ATM is going to dominate".[20] However, in 2005 the ATM Forum, which had been the trade organization promoting the technology, merged with groups promoting other technologies, and eventually became the Broadband Forum.[21]

Wireless or mobile ATM

Wireless ATM,[22] or mobile ATM, consists of an ATM core network with a wireless access network. ATM cells are transmitted from base stations to mobile terminals. Mobility functions are performed at an ATM switch in the core network, known as "crossover switch",[23] which is similar to the MSC (mobile switching center) of GSM networks. The advantage of wireless ATM is its high bandwidth and high speed handoffs done at layer 2. In the early 1990s, Bell Labs and NEC[24] research labs worked actively in this field. Andy Hopper from Cambridge University Computer Laboratory also worked in this area.[25] There was a wireless ATM forum formed to standardize the technology behind wireless ATM networks. The forum was supported by several telecommunication companies, including NEC, Fujitsu and AT&T. Mobile ATM aimed to provide high speed multimedia communications technology, capable of delivering broadband mobile communications beyond that of GSM and WLANs.

Versions

One version of ATM is ATM25, where data is transferred at 25 Mbit/s.[26]

See also

- VoATM

References

1. Telcordia Technologies, *Telcordia Notes on the Network*, Publication SR-2275 (October 2000)
2. ATM Forum, The User Network Interface (UNI), v. 3.1, ISBN 0-13-393828-X, Prentice Hall PTR, 1995, page 2.
3. Ayanoglu, Ender; Akar, Nail. "B-ISDN (Broadband Integrated Services Digital Network)" (http://repositories.cdlib.org/cpcc/2/). Center for Pervasive Communications and Computing, UC Irvine. Retrieved 3 June 2011.
4. "Recommendation I.150, B-ISDN Asynchronous Transfer Mode functional characteristics" (http://www.itu.int/rec/dologin_pub.asp?lang=e&id=T-REC-I.150-199902-I!!PDF-E&type=items). ITU.
5. McDysan (1999), p. 287.
6. McDysan, David E. and Spohn, Darrel L., *ATM : Theory and Application*, ISBN 0-07-060362-6, McGraw-Hill series on computer communications, 1995, page 563.
7. D. Stevenson, "Electropolitical Correctness and High-Speed Networking, or, Why ATM is like a Nose", Proceedings of TriCom '93 (https://www.amazon.co.uk/dp/0306444860), April 1993.
8. "ATM Cell Structure" (https://technet.microsoft.com/en-us/library/cc976978.aspx). Retrieved 13 June 2017.
9. "Asynchronous Transfer Mode (ATM)" (http://www.123seminarsonly.com/Seminar-Reports/2013-01/49988992-Asynchronous-Transfer-Mode.pdf) (PDF).
10. Carpenter, Brian E.; Crowcroft, Jon (1997). "Prospects for Internet Technology" (https://doi.org/10.1088%2F0967-1846%2F4%2F2%2F002). *Distributed Systems Engineering*. **4** (2): 78–86. Bibcode:1997DSE.....4...78C (https://ui.adsabs.harvard.edu/abs/1997DSE.....4...78C). doi:10.1088/0967-1846/4/2/002 (https://doi.org/10.1088%2F0967-1846%2F4%2F2%2F002). S2CID 2837224 (https://api.semanticscholar.org/CorpusID:2837224).
11. Cisco Systems *Guide to ATM Technology* (https://indigothemes.com/wikipedia-contribution/techgd.pdf) (2000). Section "Operation of an ATM Switch". Retrieved 2 June 2011.
12. Cisco Systems *Guide to ATM Technology* (https://indigothemes.com/wikipedia-contribution/techgd.pdf) (2000). Section "ATM Cell Header Formats". Retrieved 2 June 2011.
13. "What is VPI and VCI settings of broadband connections?" (http://www.techlineinfo.com/what-is-vpi-and-vci-settings-of-broadband-connections/). *Tech Line Info*. Sujith. Retrieved 1 July 2010.
14. ITU-T, *Traffic control and congestion control in B ISDN*, Recommendation I.371, International Telecommunication Union, 2004, page 17
15. ITU-T, *Traffic control and congestion control in B ISDN*, Recommendation I.371, International Telecommunication Union, 2004, Annex A, page 87.
16. ATM Forum, The User Network Interface (UNI), v. 3.1, ISBN 0-13-393828-X, Prentice Hall PTR, 1995.
17. "Guide to ATM Technology for the Catalyst 8540 MSR, Catalyst 8510 MSR, and LightStream 1010 ATM Switch Routers" (https://www.cisco.com/application/pdf/en/us/guest/products/ps1891/c2001/ccmigration_09186a00800f4b5d.pdf) (PDF). *Customer Order Number: DOC-786275*. Cisco Systems. 2000. Retrieved 19 July 2011.
18. Steve Steinberg (October 1996). "Netheads vs Bellheads" (https://www.wired.com/wired/archive/4.10/atm.html). *Wired*. **4** (10). Retrieved 24 September 2011.
19. "What's in store for FORE?" (https://books.google.com/books?id=IxgEAAAAMBAJ&lpg=PA12). *Network World*. 16 September 1996. p. 12. Retrieved 24 September 2011.

20. "Optical Ethernet firms brave stormy industry seas" (https://books.google.com/books?id=GBwE AAAAMBAJ&pg=PA14). *Network World*. 7 May 2001. p. 14. Retrieved 24 September 2011.
21. "About the Broadband Forum: Forum History" (https://web.archive.org/web/20111009003434/ht tp://www.broadband-forum.org/about/forumhistory.php). Archived from the original (http://www.b roadband-forum.org/about/forumhistory.php) on 9 October 2011. Retrieved 24 September 2011.
22. Wireless ATM (https://archive.is/20130615202256/http://connectedplanetonline.com/wireless/m ag/wireless_wireless_atm_debate/)
23. Book on Wireless ATM Networks - (http://www.barnesandnoble.com/w/wireless-atm-and-ad-ho c-networks-c-k-toh/1112737100?ean=9780792398226)Chai Keong Toh, Kluwer Academic Press 1997
24. WATMnet: a prototype wireless ATM system for multimedia personal communication, D. Raychaudhuri,at.al (https://www.researchgate.net/publication/3233809_WATMnet_a_prototype _wireless_ATM_system_for_multimedia_personalcommunication)
25. "Cambridge Mobile ATM work" (https://web.archive.org/web/20150625230533/http://www.cl.ca m.ac.uk/research/dtg/attarchive/radio/). Archived from the original (http://www.cl.cam.ac.uk/rese arch/dtg/attarchive/radio/) on 25 June 2015. Retrieved 10 June 2013.
26. Inc, IDG Network World (26 March 2001). *Network World* (https://books.google.com/books?id= EBwEAAAAMBAJ&q=%2225M+bit/sec+ATM+to+the+desktop%22&pg=PA98). IDG Network World Inc.

- Black, Uyless D. (1998). *ATM—Volume III: Internetworking with ATM* (https://archive.org/details/ atm00blac). Toronto: Prentice Hall. ISBN 0-13-784182-5.
- De Prycker, Martin (1993). *Asynchronous Transfer Mode. Solutions for Broadband ISDN*. Prentice Hall.
- Joel, Amos E., Jr. (1993). *Asynchronous Transfer Mode*. IEEE Press.
- Golway, Tom (1997). *Planning and Managing ATM Network* (https://archive.org/details/planning managing0000mino). New York: Manning. ISBN 978-0-13-262189-2.
- McDysan, David E.; Darren L. Spohn (1999). *ATM Theory and Applications*. Montreal: McGraw-Hill. ISBN 0-07-045346-2.
- Neelakanta, P. S. (2000). *A Textbook on ATM Telecommunications, Principles and implementation*. CRC Press. ISBN 0-8493-1805-X.
- ATM Cell formats- Cisco Systems (https://archive.is/20121206000842/http://www.cisco.com/uni vercd/cc/td/doc/product/atm/c8540/12_0/13_19/trouble/cells.htm)
- "Asynchronous Transfer Mode (ATM)" (https://web.archive.org/web/20071029234552/https://w ww.cisco.com/web/psa/technologies/index.html?c=268435599). Cisco Systems. Archived from the original (http://www.cisco.com/web/psa/technologies/index.html?c=268435599) on 29 October 2007.

External links

- "ATM forum" (https://web.archive.org/web/20050701081559/http://www.atmforum.com/). Archived from the original (http://www.atmforum.com/) on 1 July 2005.
- ATM Info and resources (http://www.telecomspace.com/vop-atm.html)
- ATM ChipWeb - Chip and NIC database (http://www.chipweb.de/atm/)
- *A tutorial from Juniper web site* (http://www.juniper.net/techpubs/software/erx/junose61/swconfi g-link/html/atm-config2.html)
- ATM Tutorial (https://web.archive.org/web/20070114172002/http://www2.rad.com/networks/200 4/atm/main.htm)
- "Asynchronous Transfer Mode Switching" (https://web.archive.org/web/20180131030149/http:// docwiki.cisco.com/wiki/Asynchronous_Transfer_Mode_Switching). *DocuWiki*. Cisco Systems.

Archived from the original (http://docwiki.cisco.com/wiki/Asynchronous_Transfer_Mode_Switch ing) on 31 January 2018.