

Ethernet in the First Mile: an opportunity for xDSL?

*Michaël Beck, Erwin Six, Eric Borghs, Piet Vandaele,
Jeroen Jacobs, Adrian Mihanța, Thierry Pollet
Alcatel Research & Innovation
Francis Wellesplein, 1
B-2018 Antwerpen, Belgium
tel. +32 3 2404131
fax +32 3 2404886
michael.beck@alcatel.be*

Abstract: Three different transport modes have been specified and standardized for use in various Digital Subscriber Line (xDSL) systems: Asynchronous Transfer Mode (ATM), Synchronous Transfer Mode (STM) and more recently Packet Transfer Mode (PTM). Although ATM and STM have successfully been used for the transport of IEEE 802.3 MAC Frames (more commonly called Ethernet frames), PTM is generally believed to be more appropriate for this purpose because it is inherently packet-oriented. The industry interest in native packet transport over xDSL links has greatly increased since the start of IEEE's "Ethernet in the First Mile" project, which covers this topic among others. In this paper we present a critical study of several aspects of native packet transport over xDSL links, which shows that native Ethernet may not be ready to conquer the last mile.

Keywords: Ethernet, ATM, xDSL

1. Introduction

"Ethernet in the First Mile" (EFM) is the informal name of the IEEE 802.3ah Task Force, whose aim is to define a physical layer specification to support the use of the 802.3 MAC [1] in subscriber access networks. One particular network topology the Task Force has chosen to address, is "point-to-point copper". The EFM standard for copper would run over a single pair of voice-grade copper, as opposed to the existing Ethernet standards which run over multiple pairs of qualified wire (e.g. two pairs of CAT-3 of length ≤ 100 m for 10/100BASE-T). The application space for EFM-copper would include both the local loop of the telephone network and in-building wiring in MxU's and hospitality buildings.

Over the past few years, Digital Subscriber Line technologies have successfully met the demand for high-speed internet access over existing voice-grade copper, both in residential and in business environments. Unlike cable modems, DSL modems most often rely on ATM as

their data link protocol, a choice that was inspired by the vision of DSL as a multi-service multi-media platform. ATM is indeed very well equipped to handle multiple streams of data, each with different bandwidth and quality of service (QoS) requirements.

ITU-T, the Standardization branch of the International Telecommunication Union, has recently added a specification for a "Packet Transfer Mode" (PTM) to its VDSL Recommendation [2]. This mode is designed to provide a low-overhead, transparent way of transporting packets over DSL links, as an alternative to using ATM. The corresponding Transport Convergence layer (PTM-TC) is specified in a sufficiently generic way to allow any kind of packets (IP, PPP, Ethernet, MPLS, ...) to run over it, provided that the higher layers adapt to a PTM-TC specific γ -interface.

From a DSL point of view, the EFM initiative would be an excellent opportunity to put the PTM-TC to use, and to define an Ethernet-over-DSL Adaptation Layer that maps the PHY-interface of the Ethernet MAC to the PTM-TC γ -interface, as shown in Figure 1.

This paper consecutively discusses the current status of packet-related DSL standards, the main differences between ATM-TC and PTM-TC, and finally some performance issues.

2. Standards Overview

2.1 Existing DSL Standards

The xDSL recommendations of the ITU-T are built up along a layered model defined in [3], as shown in Table 1. The interface between the "Transport Protocol Specific Transport Convergence" (TPS-TC) layer and the higher layer protocols is called γ -interface. This interface is defined for the three data plane TPS-TCs (ATM-TC, STM-TC and PTM-TC), as well as for the Embedded Operations Channel (EOC-TC).

ATM-TC	STM-TC	PTM-TC	EOC-TC
Physical Medium Specific – Transmission Convergence (PMS-TC)			
Physical Medium Dependent (PMD)			

Table 1. Layered Model of ITU-T xDSL Standards

In the transmit path, the ATM-TC layer accepts 53-byte cells from the ATM Entity above, and performs HEC calculation, payload scrambling, and idle cell insertion before handing the data to the PMS-TC layer. In the receive path, the ATM-TC layer checks the HEC (this check is also used to determine cell delineation), removes idle cells, and descrambles the payload. The specification of the γ -interface (and of its preferred implementation, Utopia) is such that the ATM-TC layer controls the timing of cell transmissions. The advantage of this approach is the possibility to initialize the physical layers to any bit-rate (determined by the physics of the copper loop and the business model of the access provider), where the higher layers are limited by the pace of the lower ones.

The PTM-TC Layer, defined by the ITU-T in Annex H of [2], is intended to provide transparent transport of any kind of packetized data over a DSL link, and as such can be used for EFM. The corresponding γ -interface allows reception and delivery of packets, with appropriate signals for flow control on an octet basis and indication of anomalies. As in the ATM-TC, the flow control is triggered by the lower layers, which eliminates the need for large buffers in the PTM-TC layer and makes the layer independent of the packet size. Within the PTM-TC layer in the transmit path, packets are encapsulated in HDLC (according to [4]), with start/stop transmission. In the receive path, the HDLC encapsulation is removed and the 16-bit CRC is checked. The PTM-TC layer relays all received data to the higher layers, signaling the result of the CRC check at the end of the packet, or alternatively the detection of an abort sequence.

The STM-TC Layer will not be discussed in this article. For completeness, we mention that the physical layers (PMD and PMS-TC) exist in different flavors: ADSL (Asymmetric DSL) is characterized by discrete-multitone modulation and a frequency plan that allows a large downstream band and a small upstream; SHDSL (Single-Pair High-Speed DSL) delivers symmetric bit-rates; VDSL (Very High Bitrate DSL) is meant to deliver both symmetric and asymmetric bitrates over short loops.

2.2 Ethernet in the First Mile

The 802.3ah Task Force has identified several objectives. These include support of three subscriber access network topologies and physical layers: point-to-point copper at speeds of at least 10 Mbps full-duplex up to at least 750 m; point-to-point optical fiber over a single fiber at a speed of 1000 Mbps up to at least 10 km; and point-to-multipoint fiber at a speed of 1000 Mbps up to at least 10 km. The project will also define operations, adminis-

tration, and maintenance (OAM) for EFM. In this paper only the point-to-point copper access network topology is considered.

Apart from the rate/reach objective already mentioned, it is an objective that the point-to-point copper physical layer (PHY) shall recognize spectrum management restrictions imposed by operation in public access networks, including recommendations from NRIC-V, ANSI T1.417-2001, and frequency plans approved by ITU-T SG15/Q4, T1E1.4 and ETSI/TM6.

The VDSL physical layer specified in [2] and [5] fulfills these copper access objectives.

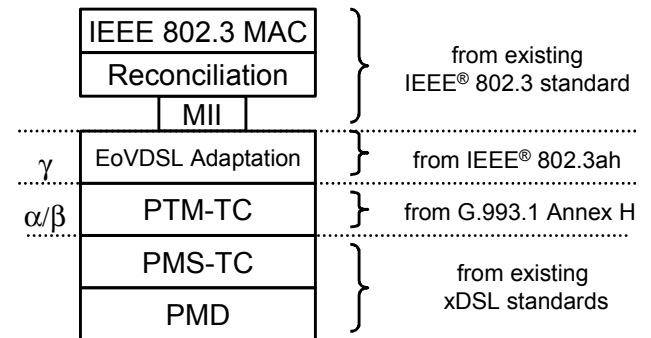


Figure 1. Ethernet-over-xDSL Functional Stack

3. ATM/PTM Comparison

3.1 Encapsulation and Error Detection

Error detection and correction in DSL systems are essentially provided by the PMS-TC layer; they consist of Reed-Solomon (RS) coding, in some cases complemented with interleaving and Trellis coding. The channel as seen by the TPS-TC layer is therefore highly non-Gaussian; errors will get through the RS protection only if they arrive in bursts. Such a burst will typically affect a limited number of packets concentrated in time, so the overall damage is limited. This explains why there is no need for sophisticated error correction mechanisms in the TPS-TC layer.

ATM encapsulation provides a single byte of redundancy, calculated over the bytes in the header (Header Error Check, HEC). If the HEC fails at the receiver side, the entire cell is discarded. In principle, a single bit error in the header can be corrected, but this capability is of little use on DSL links given the fact that single bit errors are highly unlikely to occur. ATM provides no additional protection for the payload by itself, but the different adaptation layers used to transport data over ATM do. AAL5 provides a 4-byte CRC calculated over the user payload and the AAL5 fields.

HDLC encapsulation, as used in the PTM-TC, provides two bytes of redundancy, calculated with the CRC-CCITT algorithm over the entire content of the HDLC frame (header + payload). If the CRC fails at the receiver side, this failure is signaled to the higher layers.

Traditional Ethernet PHYs use high-overhead codes that protect relatively small chunks of data (e.g. 4B/5B, 8B/10B, 64B/66B). A key concern in their design, was to protect control codewords from being misinterpreted in the event of a single or double bit error. In our opinion, these codes are not particularly useful on a channel with error correction, precisely because of the burstiness of the errors.

3.2 Flow Control

The γ -interface of the ATM-TC is designed to keep the flow control in the hands of the lower layers. This approach matches perfectly with the implementation of the γ -interface as a Utopia-interface. The same philosophy inspired the specification of the γ -interface of the PTM-TC, where the only difference is the fact that the flow is controlled on an octet basis instead of a cell basis.

The Ethernet equivalent of a Utopia-interface would be the Media Independent Interface (MII), as specified in [1]. This interface works at a fixed clock frequency (2.5 MHz for 10 Mbps operation, 25 MHz for 100 Mbps), and is unable to transmit anything less than an entire packet. This is clearly in contradiction with the octet-per-octet approach of the γ -interface. It is easy to see how this leads to problems in practical situations, where the DSL system will not be able to transmit and receive at exactly 10 Mbps or 100 Mbps. In the transmit path, if the DSL system runs faster than the MII, the HDLC encapsulation process will run dry and insert idle bytes which cause the frame to break up. If the DSL system is slower than the MII, a large transmit buffer is needed to store the data coming from the MII. In the receive path, a similar argument shows that more buffers will be needed for reliable operation [6].

A more elegant solution involves a change to the specification of the MII, which would add a set of signals (preferably by using a combination of existing pins) that allows the PHY to perform octet-based flow control over the transmit and receive paths. A MAC equipped with such a “Variable Bit-Rate MII” (VMII) could be made backward compatible with existing PHYs, by leaving the flow control signals in their “on” positions. This interface, which mimics the behavior of the underlying γ -interface, has been proposed to the EFM Task Force by the authors, [7]. It has the advantage of eliminating the need for flow control buffers in the PHY layers, allowing higher densities in PHY implementations.

3.3 Loop Aggregation

The term “Loop Aggregation” is used here to denote the use of several physical DSL links as a single logical data path carrying a stream of packets. The same functional result can be achieved by demultiplexing a stream of MAC frames at the MAC Client level, a mechanism which is known as “Link Aggregation”. An important distinction between the two, is the fact that loop aggregation fragments individual frames in smaller parts, which are then dropped into the different PTM-TCs, whereas link aggregations divides the flow of frames into

conversations (streams of frames with the same source and destination address), which are then sent to different links. Thanks to the generic nature of the PTM-TC layer, a frame fragment can be transmitted in exactly the same way as a MAC frame. The fragmentation and re-combination functions must be implemented in the layers above the PTM-TC.

In ATM, data packets are already fragmented by AAL5 in 48-byte cells, which normally travel over the same physical channel. When Inverse Multiplexing for ATM (IMA) is used above the ATM-TC layer, the exact same functionality can be offered as described above.

3.4 Scalability

One of the well-known problems with the MII is the fact that it consumes a lot of pins. On multi-port devices, every single port requires a complete instance of the MII. To overcome this problem, several vendors have specified and implemented alternative interfaces such as the Reduced MII (RMII) and the Serial MII (SMII). But even then, every additional port on a device requires at least two additional pins, with the corresponding flow control buffers. This can be a disadvantage for Ethernet-based solutions, especially when a large number of ports have to be provided per chip.

The situation is quite different in existing ATM-based DSL implementations. The Utopia interface is used as a bus, such that multiple internal ports on the device can communicate through the same set of pins.

When an Ethernet bridge receives a frame for a host it doesn't know, it will broadcast the frame on all ports. In a large network, this mechanism will lead to huge forwarding tables, and high network usage due to the broadcasts.

VLAN-tagging can be used to solve this problem, by isolating users from each other; every user/provider pair would be considered to be one virtual LAN. For a large number of users, this approach will quickly use up all available VLAN identifiers.

Label-switched technologies such as MPLS and ATM do not require broadcasts, but as for VLANs, there is an additional need to set up paths before user data can flow.

4. Performance Issues

4.1 Overhead

Of all the differences between ATM and PTM, the overhead is probably the easiest to quantify. HDLC encapsulation as used in the PTM-TC introduces one flag octet (on average), one address octet, one control octet and two CRC octets, which adds up to five octets. In addition to this, HDLC encapsulation changes every occurrence of the octets 0x7D and 0x7E in a two-octet sequence, to avoid interpreting data octets as closing flags. This causes an average increase of 0.78% for random data. Applying HDLC encapsulation to packets within the Ethernet length range (64–1522 octets) results in an expansion of 1.11%–8.59%.

The recommended protocol used to transport Ethernet frames over ATM is AAL5, [8]. It consists of adding a protocol trailer (1 octet of User–User information, 1 octet CPI, a 2-octet length field, and a 4-octet CRC), and a padding to make sure the total packet length is an integer multiple of 48 octets (the payload size of an ATM cell). The actual ATM overhead consists of 5 octets per cell. Applying this encapsulation to packets within the Ethernet length range results in an expansion of 12.67%–63.90% (using an average padding of 23 octets).

4.2 Quality of Service

The differences in QoS between cell-based and packet-based systems can be divided in two groups: bitrate related aspects and latency related aspects.

As for bitrate, ATM has the advantage of being connection-oriented, and thus being able to negotiate a service contract at set-up time. Such a contract may specify a certain minimum bitrate which must be supported. Packet-based systems are connection-less, which means that services sharing the same medium must compete with each other for bandwidth. This implies that the bandwidth requirements of different services can only be met if the user sets up a responsible mix of these services. Otherwise, a bandwidth control mechanism must be implemented above layer 2.

The latency problem results from the fact that an entire packet must be received for its content to become available to an application. In Ethernet terms, “entire packet” means up to 1500 bytes, while in ATM the unit of transferred data is only 53 Bytes. On a typical ADSL link, running at 1 Mbps, the time to transfer 1500 bytes is 12 ms. It is clear that queuing alone could consume a large portion of the delay budget of a time-critical application such as a voice connection.

5. Conclusions

The ATM-TC and PTM-TC layers defined for Digital Subscriber Line systems both have their advantages and disadvantages for the transport of Ethernet frames on the local loop. At low bitrates, latency becomes an important issue, which is best dealt with by ATM, because of its small PDU size. When different types of service are converged on the same channel, ATM is best equipped to guarantee the QoS requirements of each individual service. At high bitrates latency becomes somewhat less important, and the reduced overhead of PTM is a distinct advantage, provided that the necessary OAM&P tools are in place. The functionality of the ATM-TC and PTM-TC layers is comparable in terms of mechanisms for frame delineation, error detection and rate decoupling. The flow control mechanisms are also similar, but due to the larger granularity of Ethernet frames, larger buffers will be required than in the case of ATM.

Acknowledgments

Alcatel’s “Ethernet in the First Mile” project is conducted with support of the Flemish Region – IWT.

References

- [1] IEEE 802.3: “*IEEE Standard for Information Technology – Telecommunications and Information Exchange between Systems – Local and metropolitan area networks -- Specific requirements – Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications*”, 2000.
- [2] G.993.1: “*Draft VDSL Recommendation Foundation Document*”, ITU-T, 2002.
- [3] G.995.1: “*Overview of Digital Subscriber Line (DSL) Recommendations*”, ITU-T
- [4] ISO/IEC 3309: “*Information Technology – Telecommunications and Information Exchange Between Systems – High-level data link control (HDLC) procedures – Frame Structure*”
- [5] T1E1.4/2000-009: “*Very-high-speed Digital Subscriber Line (VDSL) Metallic Interface, Part 1: Functional Requirements and Common Specification*”, T1E1.4, 2000.
- [6] Arthur Marris: “*MAC-PHY Rate Adaptation Proposal*”, EFM/jan02/marris_1_0102.pdf, January 2002.
- [7] Michael Beck: “*The Rate Adaptation Problem*”, EFM/nov01/beck_1_1101.pdf, November 2001.
- [8] Martin De Prycker: “*Asynchronous Transfer Mode: Solution for Broadband ISDN*”, Prentice Hall, 1995.

References shown as EFM/<dir>/<filename> can be found on the website of the IEEE 802.3ah taskforce: <http://www.ieee802.org/3/efm/public>.

Glossary

<i>ATM:</i>	Asynchronous Transfer Mode
<i>EOC:</i>	Embedded Operations Channel
<i>OAM&P:</i>	Operations, Administration, Maintenance and Provisioning
<i>PTM:</i>	Packet Transfer Mode
<i>PMD:</i>	Physical Medium Dependent
<i>PMS-TC:</i>	Physical Medium Specific – Transmission Convergence
<i>TPS-TC:</i>	Transport Protocol Specific – Transmission Convergence