Payload evaluation: efficient Ethernet qualification on bit-transparent copper and fibre lines

How can transmission errors in the last mile be detected? Experts in the field of telecommunications measuring technology continually confront this question. Over the last 20 years, instrument manufacturers have rolled out a wide range of mechanisms, counters and graphics for visualising errors better and more rapidly on DSL. Up to now, the emphasis has been on ensuring that no packets are lost on the DSL line segment. With initially 100 Mbit/s, the downstream Ethernet, which essentially also constitutes the last mile, never used to be the bottleneck of the connection – but that is now changing.

Ever increasing data rates

Today, end-user Ethernet devices of private customers are equipped with a 1000Base-T interface, and modern DSL, G.fast and FTTx accesses now deliver 500 Mbit/s and more, depending on the line length. High-performance VDSL2 accesses that offer several hundred Mbit/s using vectoring and super vectoring (VDSL2 profile 35b) are already being rolled out throughout Europe. The data rate will virtually double by bundling these accesses, a technology known as VDSL bonding. Such plans even exist for G.fast, and the first chip fabricators are already offering corresponding solutions.

In the business-customer segment, by contrast, there are still a number of applications that do not require such high data rates. For this purpose, network operators are offering these customers, in addition to DSL, Ethernet accesses with a guaranteed data rate of 5, 10, 50 or 100 Mbit/s – often over great distances and thus using other technologies that enable the bit-transparent transmission of Ethernet frames over long distances. Examples of this include SHDSL and SDH lines, in which small data streams are bundled into larger ones (multiplexing) and transmitted synchronously over great distances, which today is generally realised optically (fibre).

This means that in future, simply testing the data rate at the DSL interface alone will not be enough. To ensure that the required and contractually warranted data rate is consistently and reliably available to the applications, the entire last mile including the Ethernet line segment must be tested; in the business sector, this includes the bit-transparent transmission segments.

Three-step testing process

The detection of possible bit errors that frequently occur during transmission on the line in as timely a manner as possible is of particular importance when commissioning, maintenance and in the case of outages. These bit errors generally take the form of “flipped” bits, bits whose state is unintentionally changed from 0 to 1 (or vice versa) during transmission. This is often caused by noise or crosstalk.

Such bit errors lead to lost packets and quality deterioration in time-critical UDP or RTP transmissions such as television (IPTV) and telephony (VoIP), and throughput-choking overhead (extra data) due to resending of packets in simple data transmissions (TCP). To detect these bit errors early on, testing should be carried out on the lower levels of the ISO/OSI model – ideally layer 2.

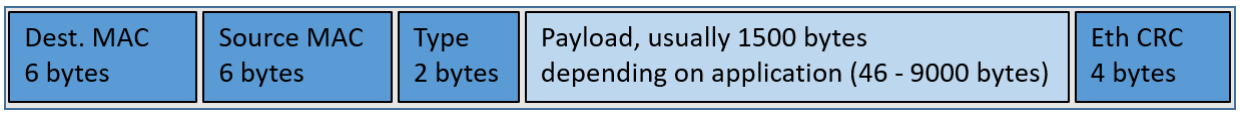

Figure 1: Typical Ethernet frame (simplified)

First testing step: Ethernet FCS

The easiest way to test a specific Ethernet segment is to generate Ethernet frames (traffic generator). However, in order to obtain a qualitative assessment, these frames must be returned to the transmitter at the end of the tested segment in a loop, so that the transmitter is the receiver.

In this step, it is sufficient when the loop has just enough intelligence to be able to swap the destination and source MAC addresses. Often, a switch with loop function or a simple loop box is sufficient for this purpose. A second tester can of course also be added (see Fig. 2).

When the frame is sent from the transmitter to the loop, a checksum (see Ethernet CRC in Fig. 1) – also called an Ethernet FCS (in this example FCS A) – is calculated and transmitted with the frame, while a separate Ethernet FCS is generated in the layer-2 loop via the incoming frame. Ideally, this is identical to the received Ethernet FCS (FCS A), the incoming frame is rated “OK” (see Fig. 3) and processed. In a true layer-2 loop, only the destination and source MAC addresses are swapped, so that in principle the same Ethernet FCS (in this example FCS A') is generated for the returned frame.

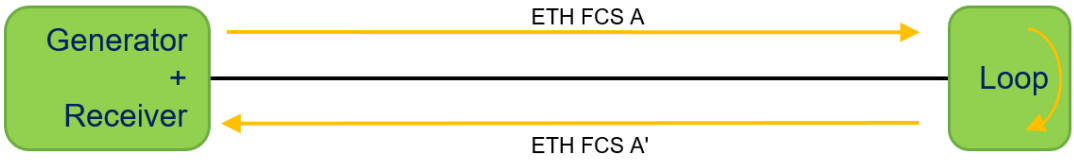

Figure 2: Traffic generator and loop, with the same ETH FCS sent in both directions.

When the traffic generator is started, the test device continually generates frames with a frame size configured beforehand. In principle, any frame size is between 64 and 10,232 bytes is possible – ideally, however, the size is selected so that it roughly corresponds with the subsequent application. If this is not clear, it is a good idea to generate different sizes. If "normal" internet traffic is expected, as is typical in the private-customer segment, a size of 1,518 bytes is ideal. Additionally, the speed with which the frame is cyclically generated and output to the line must be defined using the parameter "Bandwidth". This must be set exactly to the speed that is to be warranted to the customer or remote station in production operation. Generally, speeds of 10, 100 or 1000 Mbit/s are offered, but often "odd" speeds such as 5, 25 or 50 Mbit/s as well.

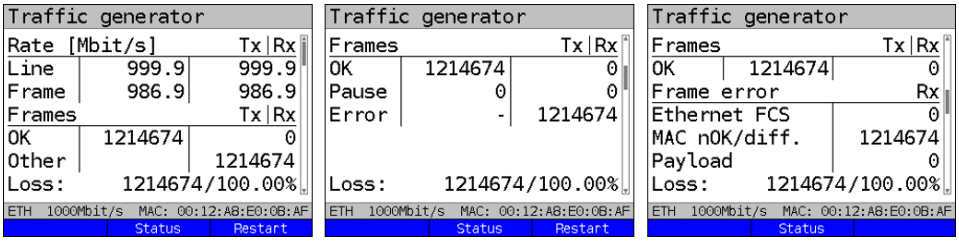

Figure 3: Example test results with errors identified in the display of the testing instrument.

The transmission mode can also be defined if necessary. This determines whether the traffic is generated for a defined period or indefinitely, or whether a previously defined number of frames is to be sent. The follow-up time defines how long the system waits for delayed frames. A time of 3 seconds has proven useful here. With these settings and the loop, the quick test with the traffic generator on layer 2 rapidly reveals any problems. The data rate is shown in the display of the testing instrument both during and after the test (see "Line rate" in Fig. 3). This immediately reveals whether the desired speed was achieved. Additionally, a counter shows the total number of OK frames – ideally all frames transmitted and received.

In the example shown here (see Fig. 3), 100% of the transmitted (Tx) frames were OK, but no OK frames were received (Rx). This means that 100% of all frames were rejected as bad or corrupted.

What happened here? In the counters of the instrument display, the number of transmitted (Tx) frames is shown as identical to the number of "other frames" (see Fig. 3, 1st image). A second look (see Fig. 3, 2nd image) reveals that these are not pause frames which keep transmission from overrunning the recipient, but 100% corrupted frames, called frame errors.

Second testing step: MAC address

These frame errors are broken down precisely in the further display (see Fig. 3, 3rd image). The number of frames with a bad Ethernet FCS and the frames with a MAC address that is either bad or unknown is easy to read in the display. In this example, all transmitted (Tx) frames were received (Rx) with a bad MAC address.

The simplest explanation for this error is that the loop did not swap the MAC addresses, so that all sent (Tx) frames were returned with an unknown address. This is likely a layer-1 loop which passes through the bits 1:1, or a layer-2 loop in which swapping of the MAC addresses was suppressed.

If a bit error had randomly falsified the MAC address used during the transmission of several thousand frames, for instance through noise, crosstalk or attenuation, the counter "MAC nOK/unknown" would only show the value 1, and not, as in the example, the total of all transmitted (Tx) frames.

To ensure that errors in the MAC address are always detected and displayed, the testing instrument must enable testing of the MAC address. For reasons of simplicity, the instrument only looks for its own MAC address in the source, destination or both fields, in the event that layer-2 traffic is generated using a loop plug (100BT) or the user does not know exactly what type of loop (layer 1 or 2) is being used.

In principle, the MAC address can change in tests on higher layers due to the line routing, so that one cannot assume that the destination MAC address when transmitting (Tx) actually corresponds to the source MAC address when receiving (Rx). Consequently, the address of the remote station is never verified. Any errors occurring here are detected by the Ethernet FCS.

Special issue: bit transparency

However, the bit-transparent transmissions described above, for example via SHDSL lines, have a unique feature that makes further testing necessary. On a bit-transparent transmission line, the bits are transmitted unaltered as a kind of bit stream with no defined beginning or end. The content of the transmission is irrelevant to the line – it is ultimately just a series of zeros and ones.

These lines are often used with technologies that enable transmission over long distances. Although these technologies have error-correcting mechanisms, they are generally unable to detect and correct flipped bits.

Thus, if a bit error occurs in the payload on these lines due to long line lengths (attenuation) or external interference (noise) from adjacent cores or devices, neither of the test methods described above will be able to detect it.

The reason for this is relatively simple: to minimise redundancy, the Ethernet FCS (in this example FCS A) is not transmitted together with the useful data on the bit-transparent line; rather, the 4 bytes are simply stripped before transmission and then recalculated and supplemented at the end of the line (see Fig. 4). Additionally, the technology does not support concurrent transmission of the preamble, the SFD byte or the inter-packet gap (IPG).

If a bit is flipped during transmission here, this is not detected at the end of the bit-transparent line segment. The new checksum, which includes this error, is treated as the correct checksum and is used in the further transmission. A remote station, e.g. a layer-2 loop, now receives the frames with the correct checksum (FCS B) and loops this back. The frames are retransmitted via the bit-transparent line, where they can again be corrupted by bit errors, and assigned another newly generated FCS (now FCS C) at the end of the line. The evaluating receiver, here identical to the transmitting traffic generator, receives these looped frames, checks the Ethernet FCS and looks for its own MAC address. In the example shown here, it assesses the packet as OK and forwards it for further processing. The bit error remains undetected.

As in this test the actual content, i.e. the useful data, is not processed by downstream protocols, the occurrence of errors in transmission on this bit-transparent line are not detected.

Instead of answering the question as to how many bit errors occurred, it is much more important to determine how many frames are lost over what period.

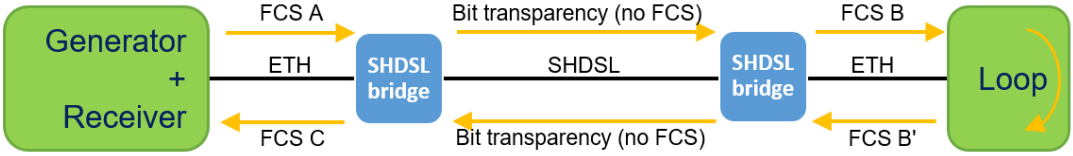

Figure 4: Ethernet transmission via a bit-transparent line without Ethernet FCS.

The line is determined to be OK. However, once this line is used in real operation, errors and transmission problems that entail data loss and reduced bandwidth can occur repeatedly.

Third test step: payload evaluation

To enable the correctness of the payload to be tested following transmission via a bit-transparent line, a third test, independent of the Ethernet FCS, is thus required to rule out possible errors. For this purpose, the payload is written with an own test pattern whose structure is comparable to that of the Ethernet frame. The first step is the insertion of an own header that contains a time stamp and a packet number. This enables the receiving device (which was also the transmitter, see Fig. 4) to immediately detect how long this frame was in transit and whether all frames are received in the correct order. The last 4 bytes of the payload are reserved for a checksum that contains the last 4 bytes of the sum of all bytes of the header and the "payload with pattern".
The part that is not used for the header and checksum is filled with the bit pattern "00 01 02 03 04 05 … FA FB FC FD FE FF". This pattern is repeated until the entire remaining payload ("orange payload" in Fig. 5) is filled in. Now, it immediately becomes apparent when a bit is flipped on a bit-transparent line. The frame is immediately identified as nOK and rejected (see Fig. 3, 3rd image).

However, the combination of the three tests:

- 1. Ethernet FCS,
- 2. MAC address evaluation,
- 3. Payload evaluation

ensures that every error – even a single bit error – is detected and displayed. However, as Ethernet is a packet-oriented transmission method, it is not the number of bit errors that is decisive, but the packet loss, i.e. the number of bad or lost frames, as every bad frame is rejected.
If bit errors occur in bursts, one or more adjacent frames can be bad. However, if they occur with regularity, then individual frames are repeatedly being lost over a longer period.
A counter that counts the seconds in which packet losses occur can be correlated with the total loss to determine whether disruptions occur in bursts or on a recurring basis. This enables deductions as to the cause. If the number of errored seconds is high and at the same time the packet loss is significant, the disruptions are recurring or constant. If the number of errored seconds is high and at the same time the packet loss is significant, the disruptions are recurring or constant. If the number of errored seconds is lower but the packet loss is still significant, the disruption is burst-type and thus transient. The ration of the

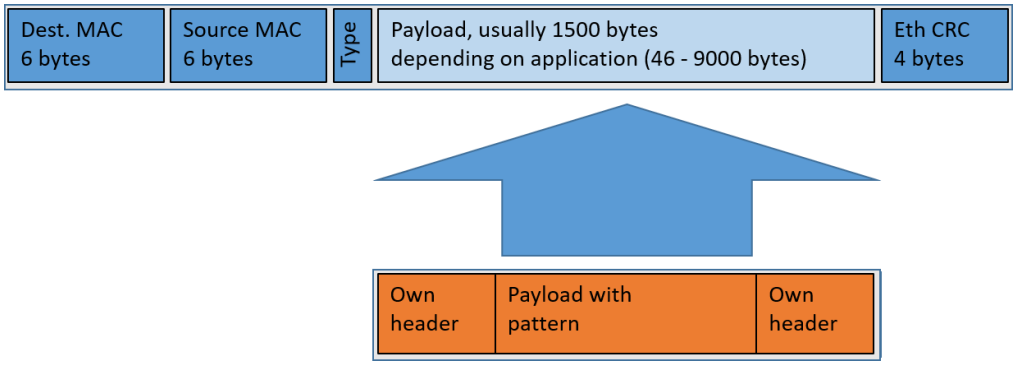

Figure 5: The payload of the transferred frames is filled with an own header, a pattern and an own CRC.

Payload evaluation vs. BERT

A Bit error rate test (BERT) familiar from ISDN does not represent a more effective alternative, as each BERT assumes a continuous bit stream, which is never the case for Ethernet. Although it is possible to generate a continuous bit stream using a typical bit pattern, this would not simulate a real Ethernet transmission. A BERT on this basis would only identify the absence or number of bit errors, but would not indicate whether an Ethernet transmission with a set speed of e.g. 1 Gbit/s and a specific packet size could be used error-free in real operation.
The number of frames arriving too soon or too late is thus not relevant for assessing quality.

errored seconds to the total run time is known as the errored seconds rate, and is expressed as a percentage. This makes it easy to recognise whether a transmission is good or lossy. Whether a rate is good or bad, on the other hand, depends on the service level agreement, but ideally maximum rates should be obtained in commissioning and maintenance – which is why a traffic generator test with payload evaluation prior to release of an Ethernet line is not just useful, but recommended.

This article was written by Dennis Zoppke in cooperation with the editorial office of the professional journal „net".