Evaluation of restoration mechanisms for future services using Carrier Ethernet

H. WESSING, M. BERGER, H. M. GESTSSON, H. YU, A. RASMUSSEN, L. BREWKA, S. RUEPP
DTU Fotonik
Technical University of Denmark
Bldg. 343, Oersteds Plads, 2800 Kgs. Lyngby
DENMARK
hewe@fotonik.dtu.dk

Abstract: In this paper, we evaluate and classify future service according to their requirements for delay, loss and bandwidth. The most demanding services include IPTV in different forms, hence IPTV is used as a representative for future services. Carrier Ethernet technologies are introduced with special focus on its OAM functionalities, and it is evaluated how IPTV performs in case of link failures on a Carrier Ethernet implementation. It is concluded that OAM update times of 10 ms is required to provide acceptable restoration performance in case of errors.

Key–Words: Carrier Ethernet, IPTV, MDI, PBT, MPLS-TE

1 Introduction

The next generation broadband services are expected to be the main driver for the future network with capacities well above 1 Gbps to the residential users. The emerging network possibilities merged with the requirements for multimedia contents give rise to a wide range of multimedia related applications like IPTV, Ultra HD IPTV, stereoscopic TV etc. In contrast to some of the driver services today, the multimedia services often require low delay to satify the requirements for interactivity. This places severe requirements to the access and the metropolitan area networks, where the user traffic is aggregated.

In the Metro Networks the consequence of network failures is severe due to the large number of affected users. Therefore, in relation to the Carrier Ethernet transport technologies, the restoration time for cable breaks etc. has to be analysed in order to determine the necessary buffer size at the receiver. On the other hand, large buffers increase the delay, which is a problem for delay sensitive applications like videoconferencing.

Hence, in this paper, we first in section 2 analyse the requirements to the network from future services and applications. In section 3 we discuss the Carrier Ethernet in Metro Networks. Then in section 4, we define the measures for quality quantification of IPTV and in section 5, we show the experimental setup and results in section 6. Finally in sections 7 we discuss the results and conclude the work

2 Services in the future

In the following several representative end user services are described. The bit-rate of the services varies from few kilobits per second (some simpler telemetric services) to 24 Gigabit per second (for immersive Ultra High Definition TV). Furthermore, the tolerable delay varies from tens of seconds (for e-mail) to a few tens of millisecond, the jitter from seconds to less than a millisecond, the packet loss from a few percent (on-line gaming) to less than a tenth of percent (thin client). The applications also have varying requirements on the traffic priority and security and may or may not allow mobility of the end-user. The services today includes video streaming and medium quality video conferencing, while some others, like remote home monitoring, location based services or Ultra High Definition Video are emerging services of the near or middle-term future. Finally, services like Web 3D and robotic assistant are more futuristic. In addition we see services like TV evolve to High Definition TV and further into Ultra High Definition TV. Drawbacks of this development include power consumption in the home and very complex control and management of traffic flows.

It is useful to classify the services according to their network requirements in terms of bandwidth requirements, packet loss, delay and jitter. This is done in Figure 1 for a selected number of the addressed services. The mapping of the services follows three main criteria:

- Delay: Here two classes are defined. Delay requirements below 150 ms and above 150 ms for one way delays. This corresponds to the definition from ITU on elastic and non-elastic services.

- Error-tolerance: This parameter is a combination of the acceptable packet loss and other parameters that transforms into packet loss and thus creates visible faults or glitches. Any application metric that produce these faults should be considered inside the ”error-tolerance” parameter. Initially, the threshold error rate is set to above or below 0,1% error induced packet loss.

- Bandwidth: The applications and services are ordered from low bandwidth to high bandwidth. These numbers from a few bps up to 24 Gbps for uncompressed Ultra HDTV distribution.

Some of the considered services are illustrated in Figure 1.

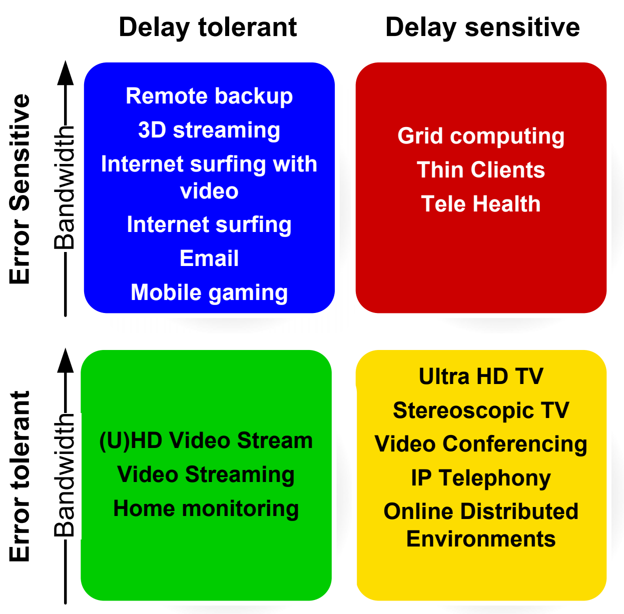

Figure 1: Classification of selected services today and tomorrow

It is clearly observed that the vast number of available and future applications have quite different requirements to the underlying network architecture.

Generally, the video services are among those with heavy requirements on bandwidth depending on display size and quality level. Normal internet video services are hard to categorise, thus it is assumed that they use medium bandwidth level. Unless interactivity is required these services have relaxed delay requirements. The future 3D services are expected to be highly bandwidth consuming.

The greatest network requirements stems from the future services in the ”entertainment” and the virtual environment area. Here, a number of high bandwidth demanding applications are matched with strict requirements on delay, making these very challenging to successfully implement. These service requirements introduce two main challenges. First, the raw capacity should be available in the network segment, and secondly, in order for a heart monitoring service to operate on the same network as an IPTV, there should be simple ways to prioritise traffic and thus make sure the large data amounts from the IPTV do not exhaust the monitoring flow.

Summarising, IPTV is by far the most demanding common service for end-users today, even if the delay requirements are far less strict than for telephony. As a rule of thumb it should be mentioned that a network that has been designed to and can carry IPTV in a satisfactory way will generally also have delays so small that VoIP is not a problem. Another rule of thump is that with fast Ethernet interfaces at the customer gateways, telephony does not normally need a special priority to separate it from ”normal” best effort Internet. This is because the telephony packages are small and the bandwidth for telephony is limited, combined with the fact that IPTV packets are very large. In other words, VoIP can live in a fine symbiosis with IPTV. Hence, focus has to be on delivering IPTV.

IPTV and in particular HDTV coded in MPEG-4 provides a trade-off between response time, bandwidth and perceived quality. More efficient video compression leads to longer response time (because of longer distance between I-frames) and lower bandwidth of the signal - and vice versa.

For commercial IPTV, the bandwidth of an MPEG-2 stream is in general between 5-10 Mbit/s; typically 9 Mbit/s. Using the same type of compression an MPEG-4 stream would require around 20 Mbit/s. Most likely a future MPEG-4 stream will be compressed to around 15 Mbit/s - at least when introduced. This is the background for the figures in Table 1

In conclusion, a triple play network should be designed so it supports the bandwidth required by IPTV and the delay required by telephony. For networks with symmetrical connections (or generally high upstream bandwidths) this enables future high-quality video conferencing and will probably enable most of the expected aggregated service scenarios envisioned in the future. As a rough rule of thumb, Carrier Ethernet based networks that are properly designed for IPTV will most probably also have delays short enough for supporting VoIP, gaming and most other expected services.

	Telephony	TV	Internet Access
Delay	150 ms in total - 50 ms for network	1 s	-
Bandwidth	80 kbit symmetric	9 Mbit/s	10 Mbit/s
Jitter	-	Buffer	Buffer
Packet loss	0,5%	Very little	0,5%
Response timer	5 s (dialup)	0.5 s (channel change)	50 ms ping time

Table 1: Network requirements for typical commercially triple play services

3 Carrier Ethernet in Metro Networks

The Metro Network is located between the access network, which might be based on GPON, EPON, xDSL or similar, and the core network. It is interconnecting the service providers access points and usually contains IPTV servers, if such services are offered. It is thus relevant to consider this network part and it resiliency mechanisms. This approach is similar to the part of the network of interest in a related paper [1], which investigate the network bottleneck of future services. A conceptual illustration of the network is shown in Figure 2, where the focus for this work is the metro part of the network.

One recent emerged technology in this network segment is Carrier Ethernet, which is basically covering technologies like MPLS-TP [2] and PBB-TE [3]. The motivation for this is to make packet network behave more like transport network and ease the migration from the existing networks. Basically, this means that the network follows the principle outlined in the ITU-T recommendation G.805, and the operation is focused on transport rather than services. The advantages are among others higher scalability, lower cost pr. bit, higher availability and better management.

3.1 OAM messages

One important aspect of this suite of technologies is their OAM information, i.e., the information used for constantly monitor the transmission, and in case of degradation or loss of signal initiate protection switching.

For the Provider Backbone Transport (PBT), the OAM information is specified through the ITU-T recommendation Y.1731. This includes ”heartbeat” continuity check signals, which operates with 7 possible rates, i.e., different OAM update intervals to indicate whether the link is operating. These rates are detailed in table 2

From table 2 it is obvious that the discussion in the remaining part of this paper should focus on the update intervals from 1 second and shorter.

Value	Interval
0	Invalid
1	3.33 ms
2	10 ms
3	100 ms
4	1 s
5	10 s
6	1 m
7	10 m

Table 2: Possible OAM update intervals

If a receiver does not received three consecutive OAM messages it determines that the link has failed and it will attempt protection switching depending on the setup.

The failure types comprise hard failures like cables cuts and node outages or soft failures like errors in the routing protocols or corrupted OAM messages. If multicast is implemented other error types are problems with the multicast tree.

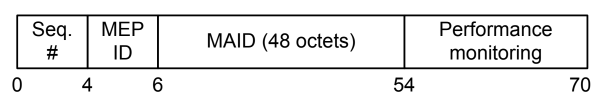

Figure 3: Continuity Check Message (CCM) format.

Figure 3 shows the structure of the Continuous Check Message (CCM). The message includes the following fields and functions:

Sequence number The bridge can keep a running tally of the number of CCM messages sent. If so then that number can be inserted into this field and transmitted to the peer which can then use it to track the number of CCMs lost dropped in the network.

Maintenance Endpoint Identifier Maintenance endpoint identifiers (MEPID) are used to identify individual MEPs and needs to be unique inside a maintenance association.

Maintenance Association Identifier Maintenance

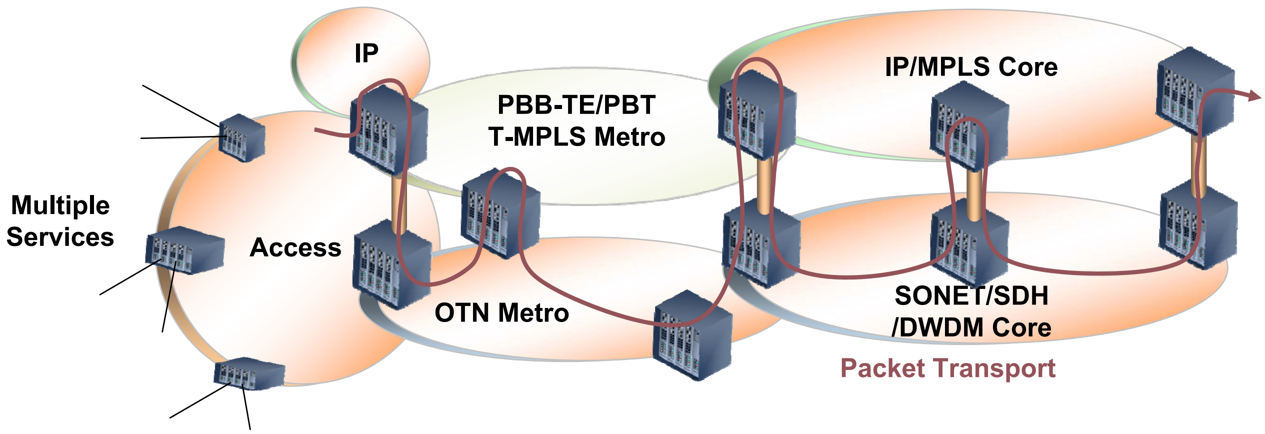

Figure 2: Carrier Ethernet Transport Network.

Association Identifiers (MAID) are used to create maintenance associations between a pair of MEPs.

Performance monitoring This field contains various counters that can be used to gauge the performance of the tunnel.

The OAM header also includes a number of CCM flags. They have three fields as shown in figure 4. First is the remote defect indicator (RDI) field, which is one bit. This bit is set on upstream CCMs to signal a defect on the downstream link. The second bit, marked Traff was added in 802.1Qay to indicate CCMs associated to Traffic Engineering Service Instances (TESIs). The last three bits in the flags field are used to indicate the interval at which the CCMs are sent as discussed above.

RDI	Traff	Reserved	Interval

Figure 4: Flags in the Continuous Check Message.

TESI are made up up two unidirectional co-routed Ethernet Switched Paths (ESPs) between two Backbone Edge Bridges (BEBs). The TESIs are grouped in pairs into traffic engineering protection groups, with one TESI acting as the working path and the other the protection path. Assigned to each TESI is a maintenance association, with the BEBs as maintenance endpoints (MEP). The MEPs exchange CCMs continuously at regular intervals. If a CCM is not received within three times the CCM interval from the reception of the last CCM then it is assumed that the path has failed and traffic is then switched to the protection path. The BEBs still continue to send CCMs on the failed path with the RDI bit set. Reception of a CCM with the RDI bit set on a path previously assumed to be in a working state indicates that an unidirectional fault has occurred and will also trigger a fail-over.

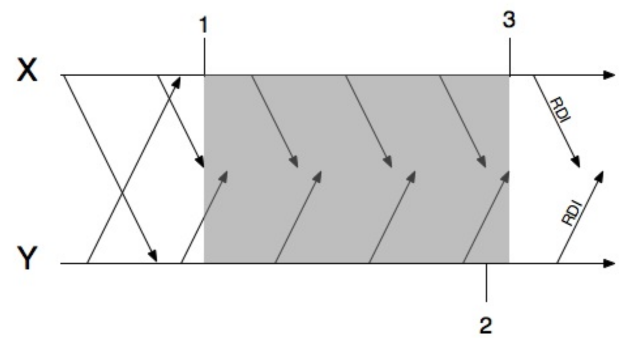

Figure 5: Remote Defect Indicator (RDI) flow in a failover situation

Figure 5 shows two nodes, X and Y, that form end-points of a tunnel and are exchanging CCMs as shown by the arrows. At the point marked ”1” in the figure a bi-directional failure occurs on the path leading to loss of messages going in both directions. Point marked ”2” is at 3.5 times the CCM interval after when Y last received a CCM from X. Y declares the link to have failed and switches traffic to the protection path. At point ”3”, 3.5 times the CCM interval after when X last received a CCM from Y, X declares the link to have failed. X then moves traffic to the protection path. All further CCMs transmitted on the working path will have the RDI bit set. Note that full service is not established on the protection path until after point ”3”. Between ”2” and ”3” node ”Y” is using the protection path while ”X” is still expecting traffic on the working path.

If the protection is set up as revertive, then traffic is switched to the working entity when both endpoints have received a CCM with RDI cleared. Fig-

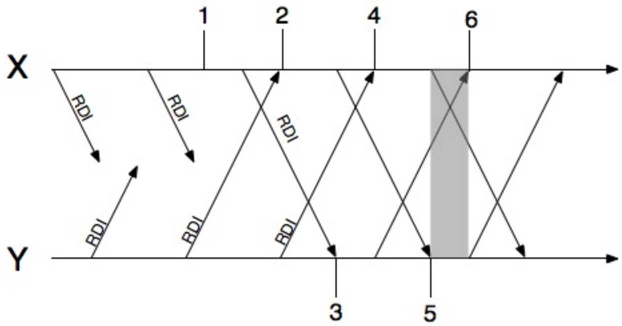

Figure 6: Remote Defect Indicator (RDI) flow when the connection is reverted. Note the short out of service in the shaded area

ure 6 is an example of this.

1. Continuity restored
2. Node X receives a CCM from Y with RDI set
3. Node Y receives a CCM from X with RDI set
4. Node X receives a CCM from Y with RDI set
5. Node Y receives a CCM from X without RDI. Begin to transmit and receive on working path.
6. Node X receives a CCM from Y without RDI. Switches to working path and service is restored.

The choice of the OAM update interval is a trade off between how fast the systems reacts to failures and the overhead imposed by the messages. The experimental studies reported in this paper evaluates the required OAM update time acceptable for an IPTV user. We have chosen IPTV as the main driver for Carrier Ethernet solutions, because most of the considered services in section 2 use IPTV in different variants. In the experimental evaluation we only focus on the one-way properties of the restoration process, however, the two way interactions can easily be derided from this.

4 Measuring video quality

In order to evaluate the quality of a video stream we need to define some definitions for acceptable video quality. This can be grouped into two major categories as shown in Figure 7.

The subjective evaluation is based on human perceived quality and mathematical or network evaluations based on measurable parameters.

Among the subjective evaluations, the Mean Opinion Score (MOS) is the most well known. Here a panel of participants are formed and they give a subjective quality score for the shown video and/or audio. MOS is inherited from the audio world and an expert panel scores the perceived stream from 1 to 5 which is from bad to excellent, respectively. This has mainly been used to compare the quality from different codec.

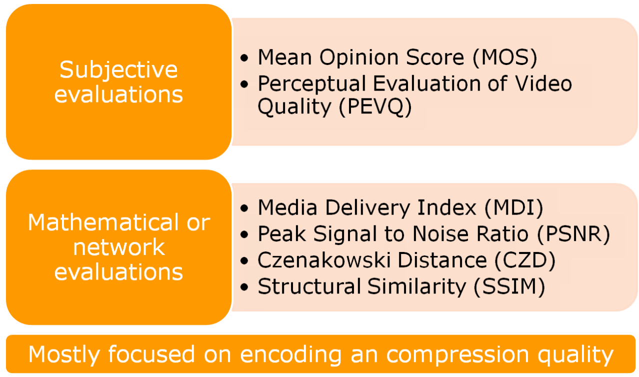

Figure 7: Evaluation of video and audio

More specifically the Perceptual Evaluation of Video Quality (PEVQ) uses a MOS evaluation for video perception. In addition to the overall MOS, other measurable parameters like PSNR are taken into account.

The second group includes the measurable parameters. First the MDI provides clear values of the delay and loss of a media stream.

In addition we have the Peak Signal to Noise Ratio (PSNR). Here the "signal" is the original uncompressed signal and the noise is the differences in the shown picture caused by the codec in the encoding phase. The Czenakowski Distance (CZD) is close to the PSNR and operates per pixel. Furthermore the Structural Similarity (SSIM) measures the differences between two pictures or two frames.

However, while these are very relevant tools for measuring compression quality they do not necessarily give a clear picture on the artifacts caused by the network layer. Only the MDI gives this information and even here the perceived quality can be very different depending on where or when the failure occurs and which type of video is shown.

Therefore we focus on a subjective evaluation backed by MDI measures. Thus, the Media Delivery Index (MDI) is used for monitoring the video stream and it is probably the most useful measure on the network layer and below. The results from the MDI can be used to indicate for a service provider the required buffer length in the video client. Off course this is also a trade off with the interactivity and the time it takes to provide control functions to the setup. Basically, MDI comprises measures of the delay, jitter and the loss, which are called the Delay Factor and the Media

Loss Rate. Hence an MDI value is described as:

$$MDI = (DF : MLR) \tag{1}$$

The Delay Factor (DF) can be used for the service provider to indicate how much buffer he or she should include in the solution, however he should usually stay within acceptable values of less than 50 ms. Otherwise the user will find zapping and other controls annoying. The Media Loss Rate (MLR) measures lost packets during a given time period, e.g., 1 sec. For HDTV the MLR should be no more than 0.05 percent.

Other studies in [4] and [5] have investigated the quality of IPTV.

5 Experimental setup

The objectives of the experimental studies are to evaluate the effect of the Carrier Ethernet signalling and monitoring. In our setup we use three Carrier Ethernet enable nodes and observe and measure the quality degradation when a failure is induced. We use the MDI value in combination with a subjective evaluation of a video stream. The experimental setup is shown in Figure 8.

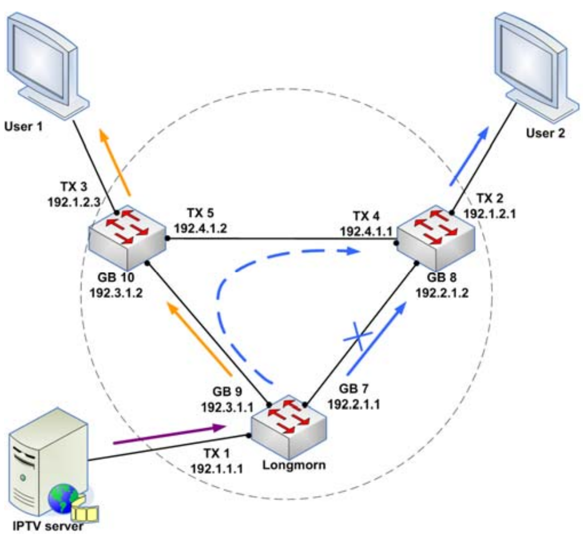

Figure 8: Experimental setup

We use three Carrier Ethernet nodes, a IPTV server and two monitors. A PBT channel with a working and a protecting path is defined and implemented in each of the CE switches. The blue solid path is our working path the dashed blue is the protecting path. The orange path is a reference path so that we can have two displays operating.

The Carrier Ethernet switches are three logical switches based on a TPACK Longmorn switch, which has the very powerful feature that it can instantiate several virtual switches as long as there are enough available ports. This makes is possible for us to use one physical switch to provide three logical switches. The network is established in a few basic steps. First, PBT tunnels are defined and for each node it is indicated whether this is a termination or a tunnelling node. Then working and protection paths are established for each flow. Finally, the virtual switches are instantiated and the device operates completely as three individual Carrier Ethernet PBT switches.

To evaluate the different OAM interval times this is adjusted as shown in Figure 9, as derived from the Carrier Ethernet Switch interface. We specifically evaluate the values 3 ms, 10 ms, 100 ms and 1 s.

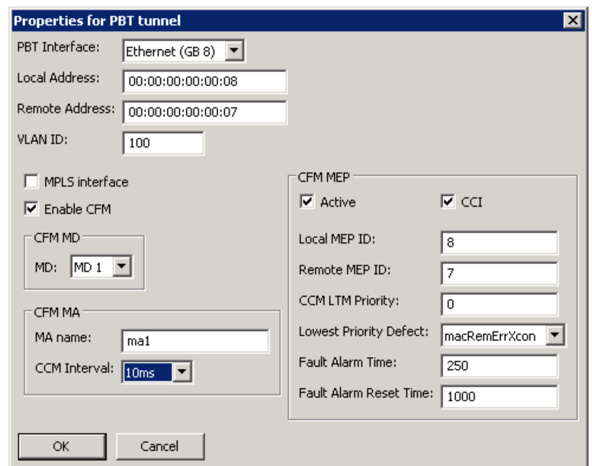

Figure 9: Adjusting OAM update interval

To measure the MDI values we use an Agilent N2X Solution. It is a very powerful measurement equipment that is extremely flexible due to its FPGA platform and it can be used to emulate the behaviour and flow characteristics of several thousand of triple play users.

6 Results

The experimental results are based on the subjective evaluation a video streaming over the network and the corresponding MDI values. Unfortunately, the video clips cannot be shown here on paper, however, the following results can be summarised as in Table 3.

It is seen from the subjective evaluation that the OAM update intervals of 10 msec or less is acceptable for a representative user. Examples of MDI measurements for the four different scenarios are shown in Figures 10, 11, 12 and 13.

In each graph the peaks are grouped two and two, the first peak is the situation when the failure occurs

Update Interval	Comments
1 sec	Long breaks and it seems that the stream looses synchronisation. Definitely not acceptable.
100 msec	Often this induces breaks in the stream, and severe artifacts are present. The result depends on the failure time with respect to the MPEG stream. The quality is not acceptable
10 msec	Minor artifacts observed. Rarely loss of sync or frames. Acceptable
3 msec	Very few artifacts. Definitely acceptable

Table 3: Subjective evaluation

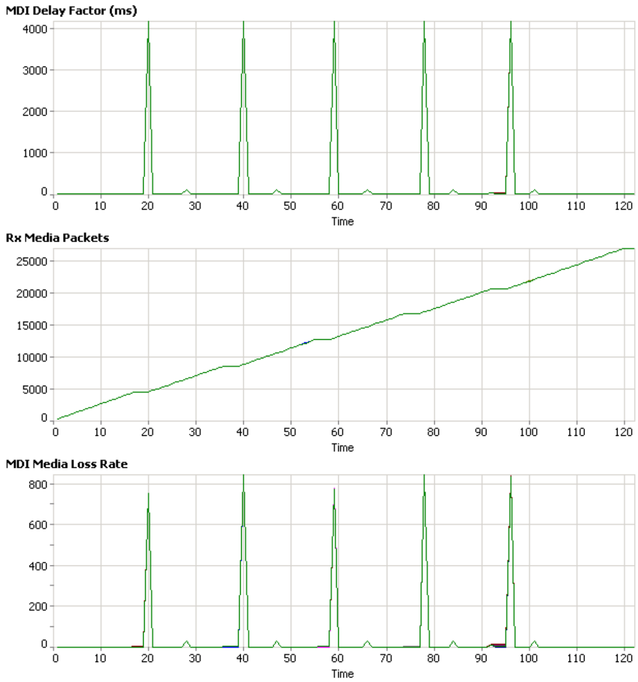

Figure 10: MDI for 1 sec update interval

and the restoration process is initiated. Then the second failure is because the original working path is again available and the system switches back to this.

Whereas the above described figures illustrates the behavior, a more in depth series of measurements were done. To measure the media loss rate the N2X was set up to record the MDI values at one second intervals whilst the network was forced into fail-over by disabling an interface on the working path. This was performed 100 times for each stream and repeated with CCM intervals set to 3:33 ms, 10 ms and 100 ms,

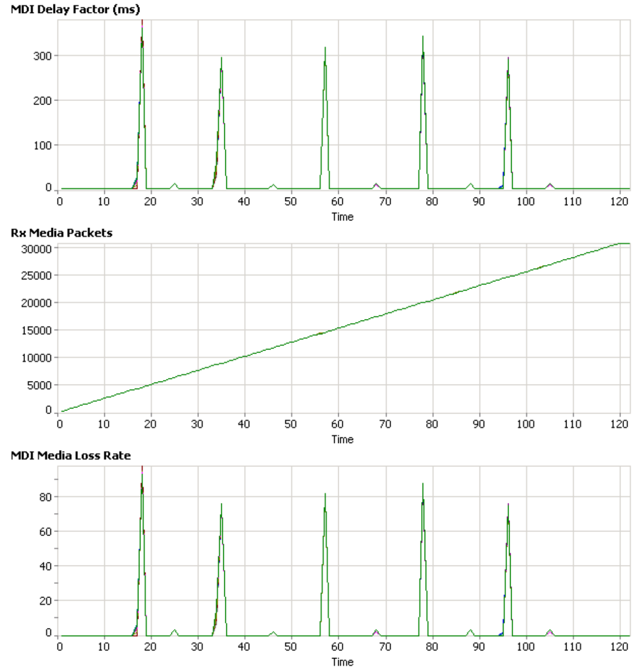

Figure 11: MDI for 100 msec update interval

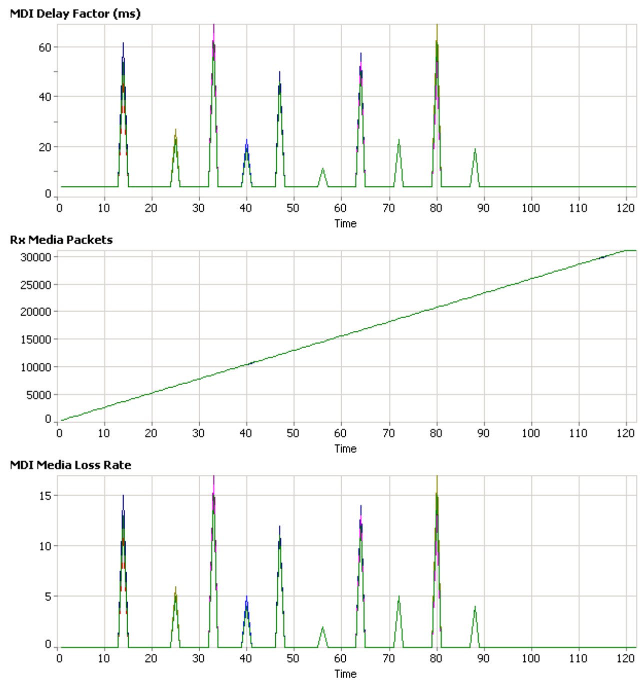

Figure 12: MDI for 10 msec update interval

the result of which can be seen in table 4 measures for a Standard Definition (SD), 720i High Definition (HD) and a 1080i HD stream.

Then 20 seconds of traffic from each stream was captured. The number of frames transmitted during each one second interval was found and averaged out. The result of this is as shown in table 5

	SD	720i HD	1080i HD
Frame Rate	204 frames/s	852 frames/s	847 frames/s
Frame Interval	4.90 ms	1.17 ms	1.18 ms

Table 5: Frame rates transmitted by the N2X

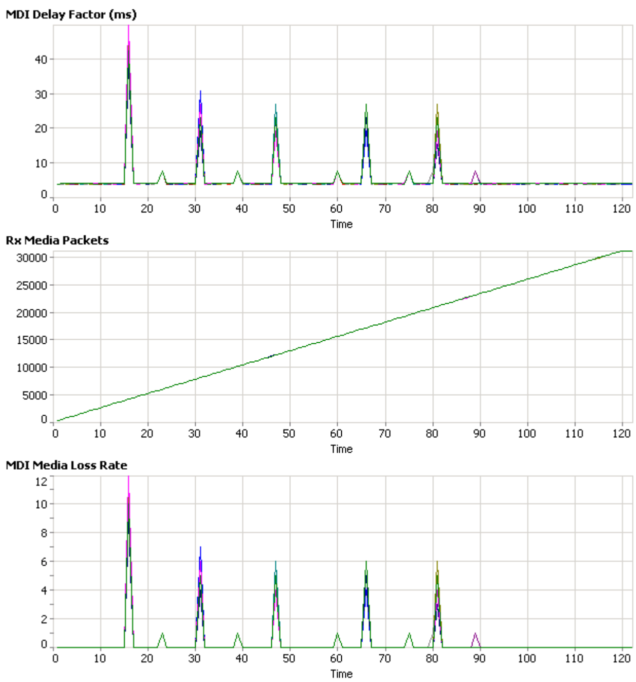

Figure 13: MDI for 3 msec update interval

	SD	720i HD	1080i HD
3.33 ms	2,6	11,0	10,7
10 ms	6,5	27,6	26,2
100 ms	59,8	248,5	245,4

Table 4: Media loss rate averaged over 100 trials using different CCM intervals

A figure for the service restoration time, that is the time elapsed from reception of the last frame before the failure until a frame is received after the fail-over can be found by multiplying the frame interval with the MLR value. The result from this and the average value for each interval is shown in table 6. While not the actual fail-over time, it should non the less give a good enough upper bound. The real fail-over time should always be lower than the number given by this measurement by an amount no more than one frame interval. It should be noted that those figures were obtained under ideal conditions, and nearly no latency due to cabling and queueing delays in intermediary nodes. In the field those figures are likely to be somewhat higher could fluctuate depending on where the occur and when they do so in relation to transmission of CCM frames.

Another aspect that is interesting to study is the effect of reversion on service availability. According to the 802.1Qay standard no service outage should be experienced due to traffic reverting to the working path. However, during testing it was found that this was generally not the case. Figure 15 shows the distribution of MLR values recorded for the two HD streams using a CCM interval of 3.33 ms. It shows that 80-88 percent of the reversions result in a MLR of 6 or 7 frames, which translates to a service outage of around 7 to 8 ms. When the CCM interval was set to 10 ms the distribution was as shown in figure 15. Here there can be seen two distinct groupings each involving two separate values. 33 to 37 percent resulted in a MLR of either 0 or 1 frame while 63 to 67 percent gave an MLR of either 9 or 10 frames. Results using other combinations also gave similar shown in those two figures.

It is believed that the unequal distribution of the reversion time is caused by some sort of synchronization of the CCM message delivery. Even though the switches act independently at the logical level, they are physically in one device and synchronization is thus possible.

The results and the subjective evaluation indicates that a service provider, who plans to use a Carrier Ethernet for providing IPTV services should use an OAM update interval of no more than 10 ms to ensure his customers an acceptable quality even in case of failures. Should he also provide interactive services, it might even be relevant to lower the interval to 3 ms, as this has a dramatic impact on the MDI quality measure.

Another study in [6] have investigated how MPLS can be used for restoration of streams; however, they are more focused on rerouting where in this paper we focus on protection.

7 Conclusion

In this paper we have discussed the future services in the internet with special focus on the Metro Network. We evaluated the expected future services, which we classified according to their network requirements. Furthermore, we examined different commercial IPTV triple play solution, and it became

CCM Interval	SD	720i HD	1080i HD	Average
3.33 ms	12.5 ms	12.9 ms	12.6 ms	12.7 ms
10 ms	31.8 ms	32.4 ms	31.0 ms	31.7 ms
100 ms	293.2 ms	291.7 ms	289.8 ms	291.6 ms

Table 6: Service restoration time

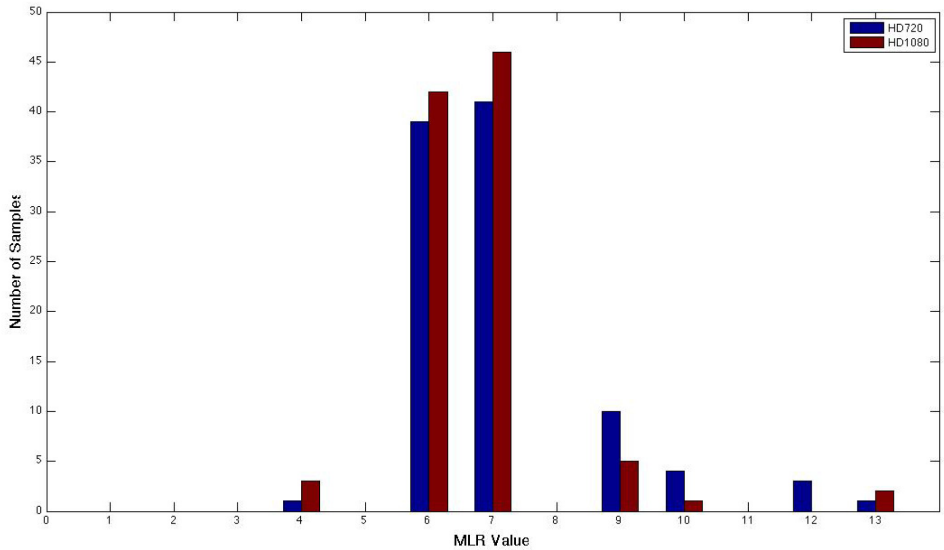

Figure 14: MLR caused by outage when reverting to working path. CCM interval set to 3.33 ms

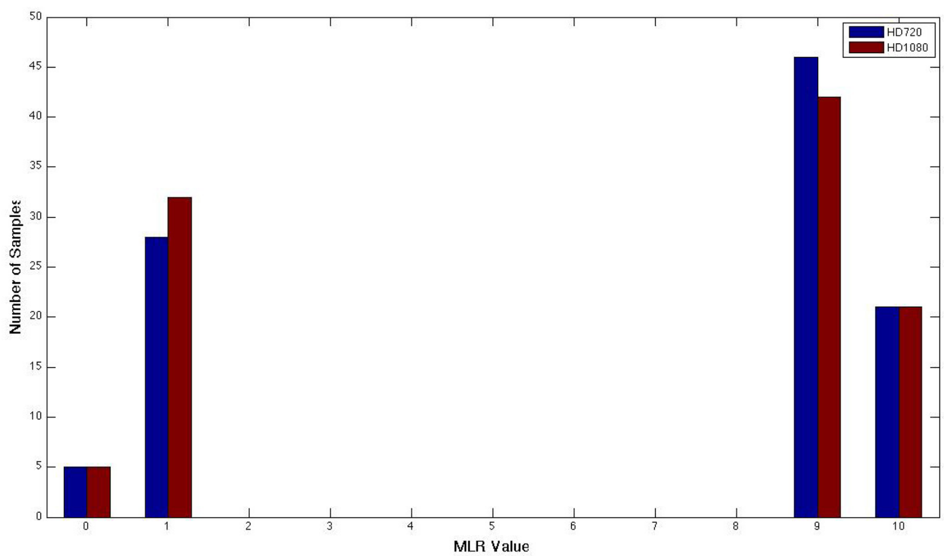

Figure 15: MLR caused by outage when reverting to working path. CCM interval set to 10 ms

obvious, that network capable to transport IPTV are also capable of transporting most other triple play and general future services. Hence, we chose IPTV as a representative candidate, and we showed how a Carrier Ethernet can be operated to support such application. We gave a detailed theoretical background for the remote defect indicators (RDI), and how they influence the restoration and reversion processes. Then, we gave examples of the measurements and provided a series of subjective evaluations performed with a panel in our lab facilities. In addition we calculated and measured the MDI values for different video types and related the results with theory. Specifically, we investigated the statistics of the reversion time in case automatic reversion is enabled. This showed some strange behavior, which we believed is caused by synchronization between the otherwise independent virtual switches. From the subjective evaluation and the measurements, we qualified a discussion for the most optimal update frequency for IPTV applications. In a specific PBT enabled Carrier Ethernet setup, it was thus indicated that an OAM update interval no longer than 10 ms provides acceptable user perceived quality.

Acknowledgements: This work was funded by the Danish Advanced Technology Foundation.

References:

[1] L. Kazovsky and N. Cheng, "Next generation broadband access networks: metro-access integration and optical-wireless convergence," *9th WSEAS International Conference on Evolutionary Computing (EC'08)*, p. 10, 2008.

[2] ITU-T, "G.8110.1, Architecture of Transport MPLS (T-MPLS) layer network," ITU-T Recommendation, 11 2006.

[3] IEEE, "IEEE 802.1Qay, Provider Backbone Bridge Traffic Engineering," IEEE, 2007.

[4] R. Bruzgiene, L. Narbutaite, and T. Adomkus, "Analysis of quality parameters influence to translation of IPTV service," *WSEAS Transactions on Systems and Control*, vol. 4, no. 11, pp. 551–560, 2009.

[5] I. Udroiu, C. Salisteanu, I. Caciula, and I. Tache, "Analyze of the MPEG-4 compressed streams," *WSEAS Transactions on Communications*, vol. 8, no. 9, pp. 1002–1011, 2009.

[6] V. Alarcon-Aquino, Y. Takahashi-Iturriaga, J. Martinez-Suarez, and L. Guerrero-Ojeda, "MPLS/IP analysis and simulation for the implementation of path restoration schemes," *WSEAS Transactions on Computers*, vol. 3, no. 6, pp. 1911–1916, 2004.