Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Institut für
Technische Informatik und
Kommunikationsnetze

Why QoS will be needed in Metro Ethernets

This is the long version of [1]
TIK Report #215

Rainer Baumann, Ulrich Fiedler

Compute Engineering and Networks Laboratory
Swiss Federal Institute of Technology
ETH-Zentrum, Gloriastrasse 35
CH-8092 Zurich, Switzerland
{baumann,fiedler}@tik.ee.ethz.ch

7th April 2005

Abstract

Emerging Metro Ethernets create new opportunities to converge data and telephony services. However, to connect legacy telephony equipment, networks have to meet customary QoS requirements even at the presence of bursty data cross traffic. We show with analytical evaluation and simulation results that without appropriate mechanisms the QoS requirements cannot be met. Thus, mechanisms for priorization such are highly needed.

1 Introduction

Recent developments in transmission and switching technologies create new opportunities to converge data and telephony services. Given the emerging deployment of metropolitan Gigabit Ethernets [5], a major aspect of interest is to employ these networks to connect legacy private branch exchanges (PBX) and GSM base stations to the core telephony network. This protects investments in existing infrastructure and creates new revenues for network providers. However, to implement this convergence, these networks have to be configured in a way that customary QoS requirements for telephony, which are significantly more stringent than the ones for VoIP, can be met. This is a crucial problem given that the data cross traffic in these networks is bursty. This burstiness can cause excessive queueing delay and frame losses due to buffer overflow [8]. Moreover, the burstiness of data traffic in these networks is known to be self similar over a wide range of time scales. This means that there is very little smooth out when aggregating this traffic over time which in turn means that buffering has little effect in mitigating the burstiness. However, it is also known that the burstiness is caused by the heavy-tail in the distribution of transfer sizes which is known to be bounded due to the limitation of file sizes in the operating systems [8]. Since Gigabit Ethernets operate at very fast transmission rates, the problem whether metropolitan Gigabit Ethernet can accommodate the burstiness of expected traffic patterns becomes real.

Therefore in this paper, we review the problem whether it is reasonable to expect that TDM E1 telephony traffic can be run on Metropolitan Gigabit Ethernets. We conduct a simulation study to investigate whether the minimum QoS requirements (in terms of loss and delay) defined by the Metro Ethernet Forum [6] can be met. As a simulation environment we use OpNet modeler [7] which is widely used in both industry and academia. To generate the data cross traffic, we enhance OpNet with a model of

On/Off sources [8] that generates self-similar traffic. Switches are configured to use FIFO queueing and tail drop.

We give an analytical evaluation to show that data buffers in switches overflow at frequencies that may degrade QoS of telephony traffic. Then, we show simulation results indicating that the Metro Forums' QoS requirements cannot be met when switch buffers are configured to limit the number of frames as currently done in practice. Thus, we also configure the switch to limit the buffer size. Here, we find that with small buffer sizes of 1MB, QoS requirements may be met when data traffic utilization is low as 1%. We explain this finding with the fact that large data frames are locked-out at the tail of a full buffer queue where as there is still buffer capacity for small telephony frames. Moreover, we review our model of On/Off sources and explain why our simulation studies tend to over-estimate this effect.

We thus doubt that it is reasonable to rely on such effects when employing Metropolitan Gigabit Ethernets to transport telephony traffic. We conclude that Metropolitan Gigabit Ethernets without network QoS support are not suited for telephony traffic. It may thus be necessary to consider QoS mechanisms such as priorization as defined in IEEE 802.1p. This may lead to a break through in the deployment of QoS support mechanisms since Metro Ethernets are usually managed by a single provider, which significantly simplifies deployment.

The rest of this report is organized as follows: Section 2 briefly describes the simulation environment and reviews Metro Forum's QoS requirements. In section 3 we present an overview of our results and conclude in section 4.

2 Setup

For the simulation study, we use the OpNet modeler [7] as our simulation environment. Modeler is a discrete event simulator that offers hierarchical network models with a focus on layer 2. We inject the offline generated traffic of superposed On/Off sources into the simulation environment. Moveover, we have enhanced modeler for collecting detailed frame delay and loss statistics.

2.1 Workload Generation

The E1 TDM signal from a PBX or a GSM base stations carries 32 telephony cannel and consumes 2.048 Mbit/s. However, transporting this signal with the standard sample rate of 8kHz over Ethernet requires 5 Mbit/s due to header/trailer overhead of 46 Bytes (8 byte preamble, 14 byte Ethernet header, 20 byte IP header, 32 byte E1 signal, 4 byte CRC).

The data traffic is generated with a set of superposed On/Off sources. Sources either send frames (are in On state) or wait (are in Off state). Parameters for the On/Off sources as listed in table 1. This leads to an average utilization of 1%. We have employed variance time plots (not shown) to verify that the burstiness characteristic of the generated traffic is as expected, i.e. that the Hurst parameter is 0.9.

Table 1: *On/Off model to generate data traffic*

Parameter	Value
Number of On/Off sources	100'000
On-time: Transfer size Pareto-distributed	Avg. 12kByte, $\alpha = 1.2$ Max. 4.1GByte
Off-time Pareto-distributed	Avg. 1025 sec $\alpha = 1.2$

2.2 Topology

Since each switch has to be such that the QoS requirements for delay and loss can be met, we focus on representing the most congested switch on this path. This leads to the topology depicted in figure 1 where all links are 1000BaseZX, thus have a link bandwidth of 1GBit/s.

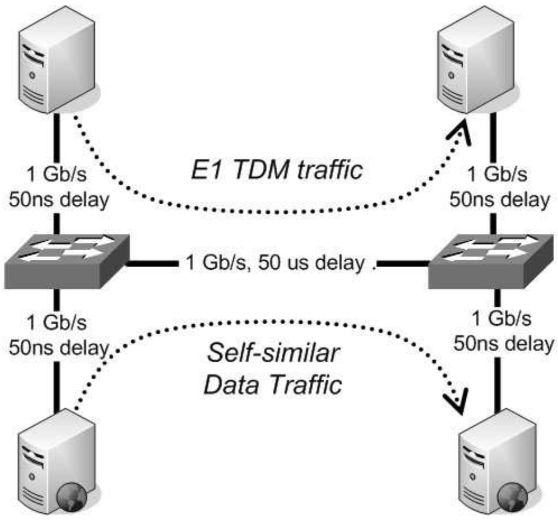

Figure 1: *Topology*

2.3 QoS Requirements

The Metro Ethernet Forum [6] essentially says that the fraction of E1 TDM frames that are lost or delayed over 25ms must be less than $8.75 \cdot 10^{-7}$. This is significantly more stringent than the requirements for VoIP (see [3]).

2.4 Switch Configuration

A widely used rule-of-thumb states that each link needs a buffer for around

$$RTT * C/MTU \tag{1}$$

frames where RTT is the average round-trip time of flows through the link, C is its data rate, and MTU is the size of the maximum transfer unit. In our case we estimate

$$RTT * C \qquad with \qquad 10ms * 1Gb/s = 1,25Mbytes \tag{2}$$

and the number of frames with 874 given that MTU is 1500Bytes. The Metro Forum's delay requirement of 25ms leads to maximum buffer size of 3.2MB. We thus consider the following configurations for limiting buffer capacity: 500 frames, 1000 frames, 1MB, 2MB.

3 Overview of Results

First, we employ theory to show that data frames in our simulation overflow the switch buffer at frequencies that likely lead to a violation of QoS for E1 TDM traffic. Secondly, we present simulation results that account for effects that ameliorate the negative effect of frequent buffer overflows on E1 TDM traffic such as the lock-out of large data frames in favor of small E1 TDM frames in the switch buffer (see figure 2).

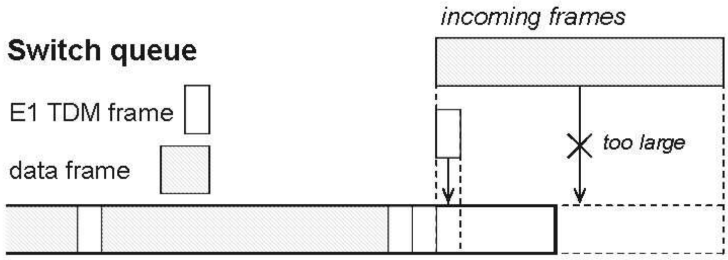

Figure 2: *The lock-out phenomenon*

3.1 Analytical Evaluation

The frequency with which data frames overflow in the switch can be estimated as follows: First, we take the transfer size distribution from the On/Off model to estimate which fraction of transfers causes buffer overflows. This fraction is essentially given with

$$p = F^{-1}(x) = k(1-x)^{-\frac{1}{\alpha}} \tag{3}$$

where F is the cumulative distribution function following a Pareto distribution[1] with

$$F(x) = 1 - (\frac{k}{x})^{-\alpha} \quad for \quad x\epsilon[k,\infty) \tag{4}$$

The location parameter k can be calculated using the shape parameter $\alpha = 1,2$ and the average transfer size $avg = 12kB$ [2]

$$k = \frac{\alpha - 1}{\alpha} avg \tag{5}$$

From equation 4, we infer the expected frequency of data frame overflows given that the On/Off sources generate 106.6 transfers per second. This number of transfers is calculated from the average utilization of the link and the average transfer size (see table 1) [4]. Results listed in table 2 indicate that QoS for E1 TDM traffic is violated. For a buffer size of 1 MByte 0.056 percent of the transfers cause buffer overflows since

$$1 - F^{-1}(1MB) = 0.056 \tag{6}$$

From this we expect that data frames overflow every 16.7 seconds. Bounding the tranfer size distribution at 4.1GB (see table 1) has no major impact on these results since

$$1 - F^{-1}(1MB) = 5.6 \cdot 10^{-2} \gg 2,5 \cdot 10^{-8} = 1 - F^{-1}(4.1GB) \tag{7}$$

For a buffer size of 2 MByte the corresponding numbers are 0.024 percent and 39.0 seconds. If E1 TDM frames overflow at comparable rates QoS is degraded.

3.2 Simulation Study

However, to account for effects that ameliorate the negative effect of frequent buffer overflows on E1 TDM traffic, we secondly present simulation results. We conduct 2:47 hours simulation runs for 16 seeds each. We find that limiting the number of frames in the buffer as currently done in most deployed switches leads to a violation of the QoS requirements. When limiting at 500 frames, only one out of 16 simulation runs meets the QoS requirements for delay and loss of E1 TDM frames. When limiting at 1000 frames only three out of 16 simulation runs meet the QoS requirements.

[1] see [4] why percentiles lead to meaningful expected values

Table 2: *Time to overflow switch buffer with data frames*

Queue Size	Corresponding Transfer Size Percentile	Expected overflow frequency
1MByte	99.944	every 16.7sec
2MByte	99.976	every 39.0sec

Therefore we alternatively limit the buffer by size. When limiting the buffer at 2 MB, only three out of 16 simulation runs meet the QoS requirements. All others violate the requirement for delay. When limiting the buffer at 1MB, we find that all 16 simulation runs meet the QoS requirement. Analyzing these, we find that this can be explained with the fact that 1500 byte large data frames are locked-out at the tail of a full buffer queue. At the same time there is still capacity for 78 byte small E1 TDM frames. However, our simulation tends to over-estimate the effect of this phenomenon. This is since (i) almost all data frames generated by the On/Off sources in our simulation are 1500 bytes large. This has to do with the way superposition is done. (ii) E1 TDM frames may in practice contain UDP and RTP headers which makes their size grow to 98 bytes instead of 78 bytes. (iii) E1 TDM frames may transport more than one sample increasing their size by 32 Bytes per additional sample. A summary of results can be found in figure 3.

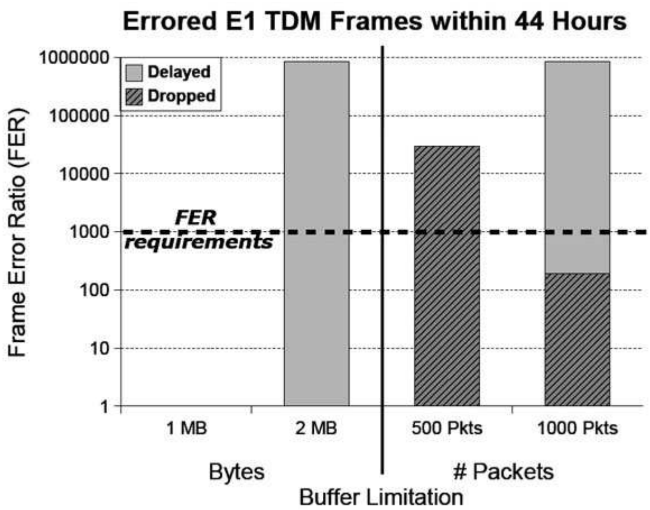

Figure 3: *Frame error ratio*

4 Conclusion

In this paper, we have given an analytical evaluation to show that data frames in switches can overflow at a frequencies that may degrade the QoS of telephony traffic. Moreover, we have shown simulation results that indicate that Metro Forum's QoS requirements cannot be met when switch buffers are configured to limit the number of frames as currently done in practice. Even when limiting the buffer size instead of number of frames, QoS may not be achieved despite the fact that large data frames are locked-out at the tail of a full buffer queue where as there is still buffer capacity for small telephony frames. We thus conclude that it is necessary to consider QoS support mechanisms such as priorization as defined in IEEE 802.1p. This may lead to a break through in deploying QoS support mechanisms, since Metropolitan Gigabit Ethernets are usually managed by a single provider, which significantly simplifies deployment.

5 Acknowledgment

We would like to thank Philipp Sager for his contributions in [9], our head of the research group Prof. Dr. Bernhard Plattner and our industry partner Siemens Schweiz. Additional acknowledge goes to KTI/CTI The Innovation Promotion Agency of the Swiss Confederation for fonding this project.

References

[1] Rainer Baumann and Ulrich Fiedler. Why QoS will be needed in Metro Ethernets. *Lecture Notes in Computer Science*, IWQoS'05, 2005.

[2] Anja Feldmann, Anna C. Gilbert, Polly Huang, and Walter Willinger. Dynamics of IP traffic: A study of the role of variability and the impact of control. In *SIGCOMM*, pages 301–313, 1999.

[3] Ulrich Fiedler, Polly Huang, and Bernhard Plattner. Towards Provisioning Diffserv Intra-Nets. *Lecture Notes in Computer Science*, 2092, 2001.

[4] Ulrich Fiedler and Bernhard Plattner. Using Latency Quantiles to Engineer QoS Guarantees for Web Services. *Lecture Notes in Computer Science*, 2707, 2003.

[5] Metro Ethernet Forum. Metro ethernet network - a technical overview, 2002.

[6] Metro Ethernet Forum. Introduction to Circuit Emulation Services over Ethernet, April 2004.

[7] OPNET Technologies Inc. OpNet Modeler 10.5A (PL3), 2004.

[8] K. Park and W. Willinger. Self-similar network traffic: An Overview. In *Self-Similar Network Traffic and Performance Evaluation*, chapter 1. Wiley-Interscience, NY, 2000.

[9] Philipp Sager. Does circuit emulation in metropolitan gigabit ethernets require service priority? Post Diploma Thesis NA-2005-02, February 2005.