

Why QoS will be needed in Metro Ethernets

Rainer Baumann, Ulrich Fiedler

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

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.

1 Introduction

Recent developments create new opportunities to converge data and telephony services. The emerging deployment of fiber fosters the proliferation of Metropolitan Gigabit Ethernets. Thus 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 [1]. Moreover, the burstiness of data traffic in these networks is known to be self-similar. This means that there is very little smooth out when aggregating the 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 popular operating systems [1]. 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 combined with an analytical evaluation to show that QoS requirements (in terms of delay and loss) as defined by the Metro Ethernet Forum [2] cannot be met even when data traffic utilization is as low as 1%.

2 Simulation Setup

We use OpNet modeler as our simulation environment. Modeler is a discrete event simulator that offers hierarchical network models and has a focus on layer two. We model the encapsulated TDM E1 telephony traffic that carries the signal of 32 telephony channels as a 5.1 Mbit/s stream. We generate the self-similar data cross traffic offline with a set of 100'000 superposed On/Off sources and inject the traffic into the simulation. On-times are heavy-tailed since the corresponding transfer size distribution is Pareto with 12 kByte average, 4.1 GByte maximum and tail index 1.2. Off-times follow a Pareto distribution with 1025 seconds average and tail index 1.2. We have verified that the data traffic generated has Hurst parameter 0.9. Since each switch has to meet the QoS requirements for delay and loss, we focus on representing the most congested switch (see figure 1 for the simulation topology). For QoS requirements, we refer to the Metro Ethernet Forum[2]. This essentially says that the fraction of E1 TDM frames that are delayed over 25ms or lost must be less than $8.75 \cdot 10^{-7}$. This is significantly more stringent than the requirements for VoIP (see [3]). Switches are configured to use FIFO queueing and tail drop. Buffer capacities are limited to 500 frames, 1000 frames, 1MB, 2MB, respectively. This choice is due to the Metro Forum's delay requirement of 25ms.

3 Overview of Results

First, we give an analytical evaluation to show that data frames in our simulation can overflow at frequencies that likely degrade the QoS of the E1 TDM traffic. Second, we present simulation results that account for effects that ameliorate the negative effect of frequent data frame 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.

The frequency with which data frames overflow in the switch can be estimated as follows: We take the transfer size distribution from the On/Off model to estimate which fraction of transfers cause buffer overflows¹. Then, we infer the expected frequency of data frame overflows given that the On/Off sources generate 106.6 transfers per second (see [5] for the math). For a buffer size of 1 MByte 0.056 percent of the transfers cause buffer overflows. From this we expect that data frames overflow every 16.7 seconds. 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.

Second, we present results of 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, degrades QoS. Limiting to 500 frames, 1000 frames, and 2MB only one out of 16 runs, respectively three out of 16 runs in the later two cases, meet the requirements for delay and loss. When limiting the buffer at 1MB, all 16 simulation runs meet the QoS requirement. However, this can be explained

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




Fig. 1. Topology




Fig. 2. Lock-out phenomenon

with the fact that 1500 byte large data frames are locked-out at the tail of a full buffer queue although there is remaining buffer capacity for 78 byte small E1 TDM frames (see figure 2). We denote that our simulation over-estimates the effect of this phenomenon since 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. For detailed results refer to [5].

4 Conclusion

In this paper, we have given an analytical evaluation and simulation results that indicate that Metro Forum’s QoS requirements for telephony traffic 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 the 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 prioritization 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.

References

1. Park, K., Willinger, W.: Self-Similar Network Traffic and Performance Evaluation. John Wiley & Sons, Inc. (2000)
2. Metro Ethernet Forum: Introduction to Circuit Emulation Services over Ethernet (2004)
3. Fiedler, U., Huang, P., Plattner, B.: Towards Provisioning Diffserv Intra-Nets. Lecture Notes in Computer Science **2092** (2001)
4. Fiedler, U., Plattner, B.: Using Latency Quantiles to Engineer QoS Guarantees for Web Services. Lecture Notes in Computer Science **2707** (2003)
5. Rainer Baumann and Ulrich Fiedler: Why QoS will be needed in Metro Ethernets. TIK Report 215, ETH Zurich (2005)