

Connecting Base Stations over Metro Gigabit Ethernets

TIK Report 238

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

28th January 2006

Abstract

Emerging high-speed metropolitan Ethernets create new opportunities to save costs when converging data and telephony services. However, connecting GSM and UMTS base stations as well as private branch exchanges (PBX) over metropolitan Ethernets require these networks to meet QoS requirements in the presence of bursty data traffic. To investigate this problem, we have probed ETH's campus network for several weeks. This network is one of the best equipped campus networks in Europe and spans the metropolitan area of Zurich. Our results indicate that even this network does not offer sufficient QoS to connect base stations and PBXs. Specifically, the frame loss rate is at times high and thus violates the QoS requirement as stated by the Metro Ethernet Forum.

1 Introduction




Figure 1: Scenario of our study.

The emerging deployment of fiber enables the proliferation of metropolitan Gigabit Ethernets. This proliferation in turn creates new opportunities to save costs when converging data and telephony services. To ensure a smooth migration for telephony services, equipment manufacturers have a great interest to investigate whether such Ethernets can be employed to connect GSM base stations and PBX to the core telephony network (see Figure 1). This interest includes the investigation whether and with what configuration such Ethernets can ensure the quality of service required for connecting private branch exchanges (PBX) as well as GSM and UMTS base stations in the presence of bursty cross traffic.

Such investigations on quality of service can be done either with simulations or with active network probing. Active network probing is done by injecting traffic into the network to measure the QoS that

this test traffic would receive. Employing simulations, we have shown in [1] that the self-similarity property in data traffic is sufficient to prevent metropolitan Ethernets from meeting quality of the services requirements as specified by the Metro Ethernet Forum (MEF) [2]. Ref. [1] also considers that network traffic can only be approximated with self-similar processes since transfer sizes are bounded due to limitations in popular operating systems. To verify that the conclusion from the simulations apply to real networks, we have probed ETH’s network to test what quality of service encapsulated E1 traffic from base stations and PBX would receive. To perform our study, we have sent traffic into the network through access routers in different buildings throughout the city and we have received this traffic at access routers in other buildings. The routes for this traffic went through the network core and traversed either lightly loaded links, averagely loaded links, or highly loaded links. To conduct the probing measurements, we have implemented an infrastructure that includes generator, sink, and archiver applications. These applications run on top of RTAI [3], a real time Linux. Each route has been probed for more than 15 days. Our results give strong indications that traffic from base stations and PBX cannot meet MEF’s QoS requirements in terms of frame loss.

The rest of the paper is structured as follows: Section 2 reviews the Metro Ethernet Forums QoS requirements for telephony circuit emulation in Metro Ethernets. Section 3 discusses the methods of our study, i.e. reviews the topology of ETH’s network and our measurement infrastructure. In Section 4 we present probing measurement results before we conclude in Section 5.

2 QoS Requirements

Metric	QoS requirement
Jitter	max. 10 ms
Delay	max. 25 ms
Loss rate	max. $8.75 \cdot 10^{-7}$

Table 1: *Metro Forum’s QoS requirements.*

The Metro Ethernet Forum (MEF) [2] is a consortium of ISPs and network equipment manufacturers that specify quality of service requirements for telephony circuit emulation over metropolitan Gigabit Ethernets. These requirements pertain (i) to a requirement for frame jitter, i.e. an upper bound to the time elapsed between receiving successive frames, (ii) to a requirement for the maximum delay and (iii) to a requirement for the frame loss rate within an observation interval of “30 hours or longer”. Numbers are listed in Table 1. Moreover, MEF specifies an additional composite requirement that the number of frames that violate any of the three above stated requirements must be bounded. This requirement is called the frame error rate, and is defined as

$$FER = JVR + DVR + LR. \tag{1}$$

where

FER: frame error rate

JVR: jitter violation rate

DVR: delay violation rate

LR: loss rate

MEF specifies that this frame error rate is less than $8.75 \cdot 10^{-7}$.

3 Methods

The investigation of whether metropolitan Ethernets can be employed to connect base stations and PBXs can be done either passively or actively. Passive measurements mean that the traffic in the network

is observed without injecting probe traffic. However, drawing conclusions from passive measurements would require excessive simulation studies of what might happen if traffic of base stations or PBXs is added to this network. For such a simulation study refer to [1]. In this paper, we use active measurements to inject traffic into ETH’s network, which spans the metropolitan area of Zurich connecting institutes and research facilities; we study the QoS that this traffic receives.

3.1 Network Topology

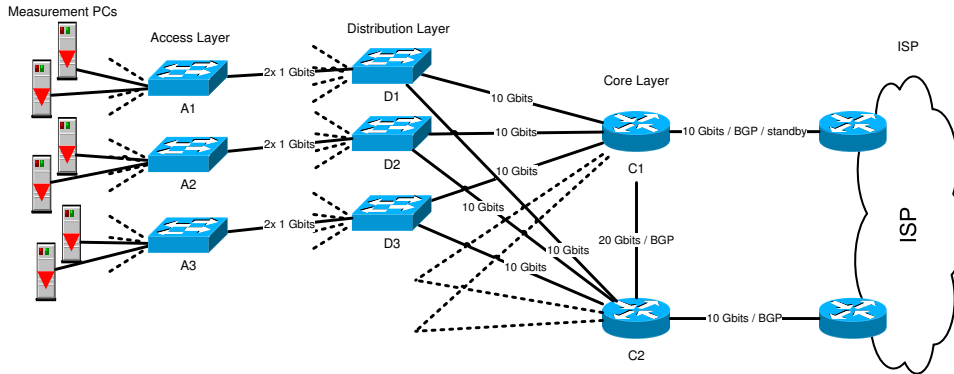


Figure 2: *Topology of the ETH Zurich network (abstraction).*

The topology of the probed network is depicted in Figure 2. This topology follows the CISCO recommendations on using a three level hierarchy and a dual core [4]. The core is connected to an ISP which connects ETH to the outer world. To perform probing measurements we have sent probing traffic into the network through access routers in different buildings throughout the city and we have received this traffic at access routers in other buildings. The routes for this traffic went through the network core and traversed either lightly loaded links, averagely loaded links, or highly loaded links as listed in Table 2 (forward direction). The backward direction has also been probed.

Probing	Route
A1 → A3	A1 → D1 → C2 → D3 → A3
A1 → A2	A1 → D1 → C2 → D2 → A2
A2 → A3	A2 → D2 → C2 → D3 → A3

Table 2: *Routes of probing traffic (forward direction).*

Link	Utilization Up	Utilization Down
A1 ↔ D1	0.3%	0.5%
A2 ↔ D2	0.2%	0.2%
A3 ↔ D3	0.4%	1.0%

Table 3: *Avg. util. for access-distribution layer links.*

The average utilization of the links on the probed routes are listed in Table 3 and Table 4. The probing traffic had a rate of 8000 frames per second and included Ethernet, VLAN, IP, UDP headers as well as 32 bytes dummy payload. The resulting bandwidth of the probing traffic is 5.248Mbit/s . The observation intervals were between 15 and 29 days, i.e. significantly longer than the 30 hours minimum specified by the MEF. Different routes were probed during disjunct observation intervals.

Link	Utilization Up	Utilization Down
D1↔C2	0.2%	0.5%
D2↔C2	0.5%	1.5%
D3↔C2	3.1%	5.1%

Table 4: Avg. util. for distribution-core layer links.

3.2 Probing Infrastructure

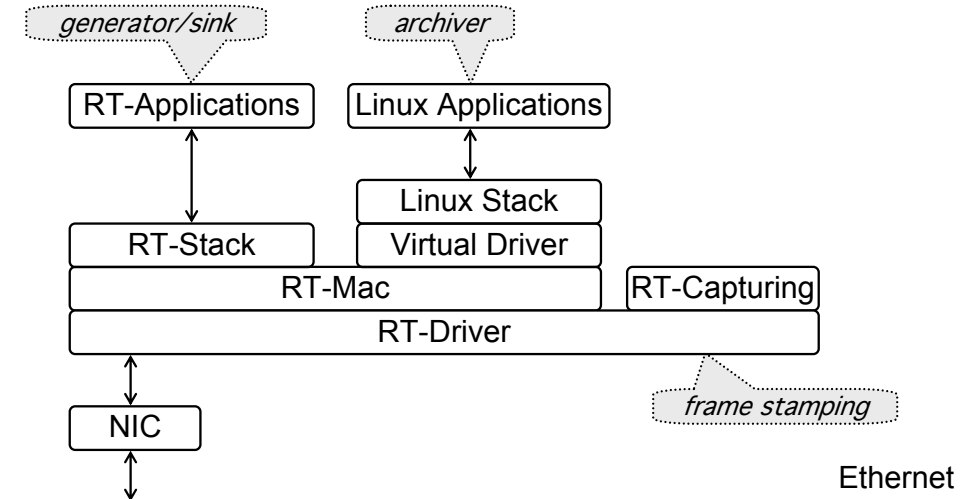


Figure 3: Measurement tasks and protocol stack.

Pasztor and Veitch show that a combination of off-the-shelf PCs and specialized real time software is sufficient to infer one-way network latencies with accuracy $\pm 10\mu s$ from time stamps written into Ethernet frames, when sending and receiving probing traffic [5]. We thus conceptually followed their approach and have employed off-the-shelf workstations. On these workstations we have run real time Linux and a real time protocol stack to implement the probing generator, sink and archiver (see Figure 3). To perform the probing measurement, the generator injects Ethernet frames into the network at the sender side, right after the RT-Driver of the Ethernet card has stamped the frame with the sending time. In addition to the send-time stamp, the generator also provides frame sequence numbers to detect frame duplicates, drops and reordering. At the receiver side, the RT-Driver of the Ethernet card stamps incoming frames with a receiver time stamp right after the frame arrives, and it forwards them to the real-time sink. Then the non real-time archiver writes the sinked frame into a log file which resides on the workstation's hard disk. Once a day, this log file is transferred to a tape archive. Finally, an offline task analyzes this log file for the quality of service associated with the archived probing traffic.

As hardware we have used Dell precision 340 workstations which are equipped an Intel Pentium III processor with $1GHz$, $512MB$ RAM and a 3COM 3C595 Fast EtherLink network interface card. For the real time Linux, we chose RTAI [3], which is known for good real time performance, and available under a GNU public license. The interrupt latency that limits the real time performance depends on the frequency of the PCI 8254, which runs at $1193180Hz$, and is as short as $0.8\mu s$. For the real time protocol stack, we chose RTnet [6], a hard real-time networking stack for Linux/RTAI. RTnet is a slightly modified Linux protocol stack which offers atomic send and receive time stamps based on the TSC processor register (see [5] for a discussion on accuracy).

To calibrate clock skew between sender and receiver side, we have directly connected the workstations with a crossover cable for 12 hours. The accuracy of the calibration is listed in Table 5.

Avg. delay	$1.30\mu s$
Std. dev. of delay	$0.51\mu s$
Max. delay	$13.55\mu s$
Sample size	$3.3 \cdot 10^8$

Table 5: Accuracy of measurement infrastructure.

4 Results

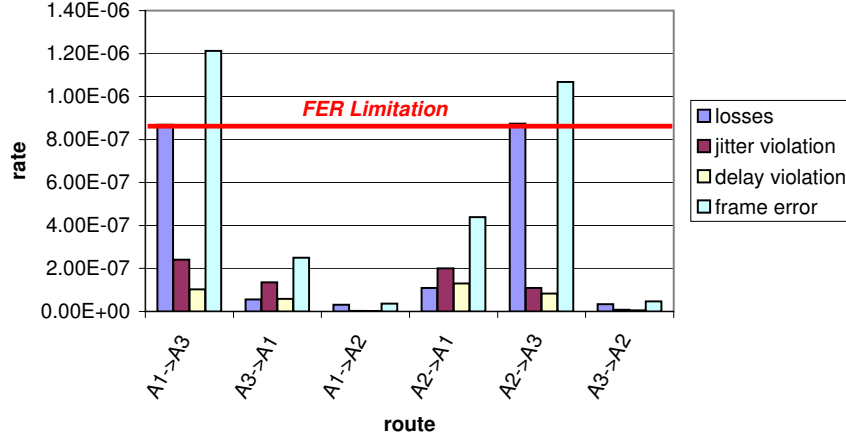


Figure 4: Quality of service received by probes (full duration)

Table 6 lists the path of the probing measurements, the sample size in term of number of frames and the duration. Figure 4 depicts the quality of service which the probing traffic received over the full

Path	Sample size	Duration [days]
A1→A3	$1.04 \cdot 10^{10}$	15.05
A3→A1	$1.04 \cdot 10^{10}$	15.05
A1→A2	$2.03 \cdot 10^{10}$	29.37
A2→A1	$2.03 \cdot 10^{10}$	29.37
A2→A3	$1.19 \cdot 10^{10}$	17.22
A3→A2	$1.19 \cdot 10^{10}$	17.22

Table 6: Summary of probes conducted.

measurement duration. Specifically, this figure lists the loss rate, jitter violation rate, delay violation rate and frame error rate for each of the path. The frame error rate limitation (FER limitation) bar depicts the Metro Ethernet Forum’s overall quality of service requirement. Comparing the shown results with the utilizations on the traversed links (see Table 2 for the routes of the probing traffic and Table 3 and Table 4 for the link utilizations), we infer that the high average utilization on the link from C2 to D3 presumably leads to frame loss rates that violate MEF’s requirement. To further investigate this violation, we have looked at the temporal distribution of frame losses. Figure 7 shows the temporal distribution of frame losses on the path that experienced a violation of MEF’s requirement, i.e. $A2 \rightarrow A1$ and $A1 \rightarrow A3$, respectively. >From this figure, we see that the loss rate may be acceptable for several days before heavy data bursts cause a violation of MEF’s frame loss rate requirement.




Figure 5: *a) frame loss A1 to A3 (6 min bins) / b) frame loss A2 to A3 (6 min bins)*




Figure 6: *(c) frame losses A1 to A3 (1 hour bins) / (d) frame losses A2 to A3 (1 hour bins)*




Figure 7: *(e) frame losses A1 to A3 (1 day bins) / (f) frame losses A2 to A3 (1 day bins)*
Temporal distribution of frame losses on path with high loss

5 Conclusion

We have probed ETH's campus network with traffic that emulates encapsulated E1 traffic from GSM/UMTS base stations and PBXs. To our knowledge, we are the first to publish probing data of large scale measurements in network that spans a metropolitan area. We have found evidence that MEF's QoS requirements for connecting base stations and PBX over Metropolitan Ethernets are occasionally violated. This violation occurred although the most highly loaded network link included in our measurements had an average utilization of 5.1%. As a consequence, we conjecture that networks like the ETH's campus network, which is among the best equipped campus networks in Europe, cannot easily take over telephony services without changes in the network. This finding confirms simulation results published in [1], which show that the self-similarity property in data traffic is sufficient to prevent metropolitan Ethernets from meeting MEF's QoS requirements even though the average network utilization is low. Thus, metropolitan area networks will presumably need some prioritization scheme to be ready to take over traffic from base stations or PBX. The software of our measurement infrastructure is freely available on request to the authors. In the future, we plan to probe other networks and to check whether we can infer conditions when MEF's QoS requirements for connecting base stations and PBX over Metropolitan Ethernets are violated.

6 Acknowledgment

We would like to specially thank Derk Valenkamp and his team from the Informatikdienste ETH Zurich for providing us with access to the network and for their great support. Additional acknowledgment goes to the head of our research group, Prof. Dr. Bernhard Plattner, our industry partner Siemens Schweiz and KTI/CTI, the Innovation Promotion Agency of the Swiss Confederation for funding this project.

References

- [1] Rainer Baumann and Ulrich Fiedler, "Why QoS Will Be Needed in Metro Ethernets," in *IWQoS 2005: Proc. of the 13th Workshop on Quality of Service*. January 2005, vol. 3552, pp. 379–381, Springer LNCS.
- [2] Metro Ethernet Forum, "Technical Specification MEF3, Circuit Emulation Service Definitions, Framework and Requirements in Metro Ethernet Networks," April 2004.
- [3] "RTAI, Real-Time Application Interface," www.rtai.org.
- [4] CISCO Systems Inc., "Hierarchical campus design-at-a-glance," 2004, available via www.cisco.com.
- [5] A. Pasztor and D. Veitch, "A precision infrastructure for active probing," in *Proc. of PAM*, 2001.
- [6] "RTnet, Hard Real-Time Networking for Linux/RTAI," www.rts.uni-hannover.de/rtnet.