Connecting Base Stations over Metro Gigabit Ethernets

Rainer Baumann and Ulrich Fiedler
Computer Engineering and Networks Laboratory
ETH Zurich, Switzerland

***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 over metropolitan Ethernets require these networks to meet stringent QoS requirements in the presence of bursty data traffic. To investigate this problem, we have probed ETH's campus network, which spans the metropolitan area of Zurich, for several weeks. From our results, we infer that lightly-loaded metropolitan Gigabit Ethernets with average utilizations below 1% presumably have a potential to carry traffic from GSM/UMTS base stations.**

I. Introduction

Recent developments create new opportunities to converge data and telephony services. The emerging deployment of fiber fosters the proliferation of Gigabit Ethernets in metropolitan areas. Thus, a major aspect of interest is to employ these networks to connect UMTS and GSM base stations to the core telephony network. However, to pave the way for this convergence to happen, it needs to be researched whether Gigabit Ethernets can indeed meet the necessary quality of service (QoS) requirements. These requirements are significantly more stringent than the ones for VoIP as given in [1]. Meeting the QoS requirement may be difficult since data cross traffic is inherently bursty. This burstiness, which has been subject to many investigations [2], can cause excessive queueing delay and frame losses due to buffer overflow. Since Gigabit Ethernets operate at very fast transmission rates, the perspective that these networks can accommodate much of the burstiness of the cross traffic becomes real.

In [3] we report on network simulations that show that the burstiness property generally encountered in data cross traffic prevents Gigabit Ethernets from meeting quality of the services requirements in terms of frame delay and loss as specified by the Metro Ethernet Forum (MEF) [4] (see section II for details). Thus, to make a link between simulation and reality, we in this paper address the problem, whether a real metropolitan network that is utilized with less than 1% in average can meet MEF's QoS requirement when carrying traffic from GSM/UMTS base stations.

Our method of investigation for this problem is probing the operational campus network of ETH Zurich. We have injected probing traffic into the network that is comparable to the encapsulated E1 traffic, that UMTS and GSM bases stations generate, and have measured the QoS that this traffic perceived. The rate of probing has been $5.248 Mbit/s$. To conduct the probing measurements, we have implemented an infrastructure that includes generator, sink, and archiver applications. These applications run on top of RTAI [5], a real time Linux. We have systematically connected the generator and sink/archiver to access routers in various buildings throughout the city. We have made sure that the associated routes of the probing traffic went from an access router to a distribution router, further on to a core router and back. To establish the relation between the average network utilization and QoS perceived by probes, we focused our investigation on six routes. Most links on these routes were lightly loaded with an average utilization of less than 1%. One link had been highly loaded with an average utilization above 5%[1]. Each route has been probed for at least 15 consecutive days. Our results give strong indications that probing traffic can meet MEF's QoS requirements. The only vio-

[1]We denote that this average utilization of 5% can lead to temporary frame losses due to the high traffic burstiness.

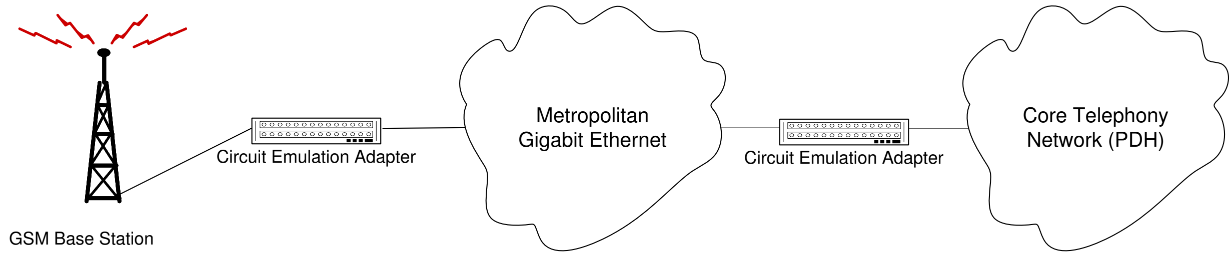

Fig. 1. Scenario of our study.

lations we have found were on routes that included the most highly loaded network link. Thus, lightly-loaded metropolitan Gigabit Ethernets with average utilizations below 1% can presumably be employed to connect UMTS and GSM base stations.

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

II. QoS Requirements

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

TABLE I
METRO FORUM'S QoS REQUIREMENTS.

The Metro Ethernet Forum (MEF) [4] 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 I. 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 $8.75 \cdot 10^{-7}$.

III. Related Work

To our knowledge, we are the first to publish any QoS measurement in relation with employing packet-switched networks to forward GSM/UMTS base station's traffic. We conjecture that this comes from the fact that reporting QoS of operational networks is a highly sensitive issue. We thus focus on reporting related work on a more general level.

Simulation results from our group on employing packet-switched networks to forward GSM/UMTS base station's traffic are reported in [3]. This paper shows that connecting GSM/UMTS base stations via packet-switched networks may not be possible for average network utilization higher than 1% due to cross traffic burstiness characteristics. Burstiness characteristics in packet-switched networks, which is the basis for this type of simulations, is excessively discussed in [2], [6] and [7].

Simulations and analytic evaluations on the QoS for VoIP in IP networks, which is somehow related are reported in [1], [8], [9].

IV. METHODS

Network probing to investigate whether metropolitan networks 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 is difficult since carrying GSM/UMTS traffic significantly consumes bandwidth which in turn has an impact on the networks performance. Therefore, we have decided to perform active probing. We actively inject traffic into the ETH Zurich network. We then study the QoS that this traffic perceives.

A. Network Topology

The topology of the ETH Zurich network, which spans the metropolitan area, is depicted in Figure 2. This topology follows the CISCO recommendations on using a three level hierarchy and a dual core [10]. The core layer routers are 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 have been as listed in Table II. The average utilization of the links on the routes has been as listed in Table III and Table IV. We emphasize that both forward and backward direction have 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 II

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 III

AVG. UTIL. FOR ACCESS-DISTRIBUTION LAYER LINKS.

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

TABLE IV

AVG. UTIL. FOR DISTRIBUTION-CORE LAYER LINKS.

The probing traffic has had a rate of 8000 frames per second and included Ethernet, VLAN, IP, UDP headers for encapsulation as well as 32 bytes dummy payload. The resulting rate of probing traffic has been $5.248 Mbit/s$. The observation intervals were between 15 and 29 days, i.e. significantly longer than the 30 hours minimum specified by the MEF. Different routes had been probed during disjunct observation intervals.

B. Probing Infrastructure

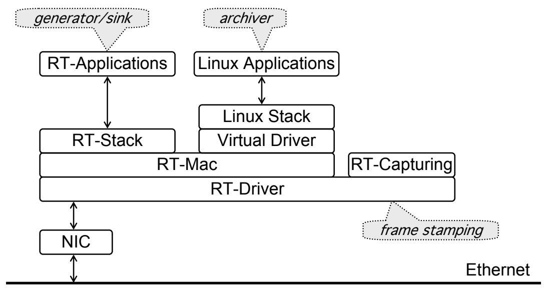

Fig. 3. Measurement tasks and protocol stack.

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 [11]. 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 Ethernets 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

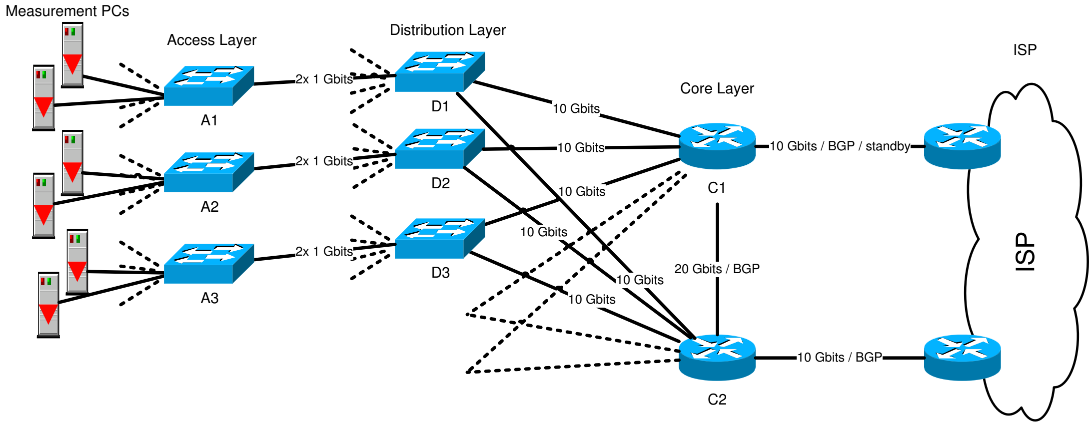

Fig. 2. Topology of the ETH Zurich network (abstraction).

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 off-line 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 with 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 [5], 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 [12], 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 [11] 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 V.

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 V

ACCURACY OF MEASUREMENT INFRASTRUCTURE.

V. RESULTS

Table VI 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 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 average utilizations on the traversed links (see Table II for the routes of the probing traffic and Table III and Table IV 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 investi-

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 VI

SUMMARY OF PROBES CONDUCTED.

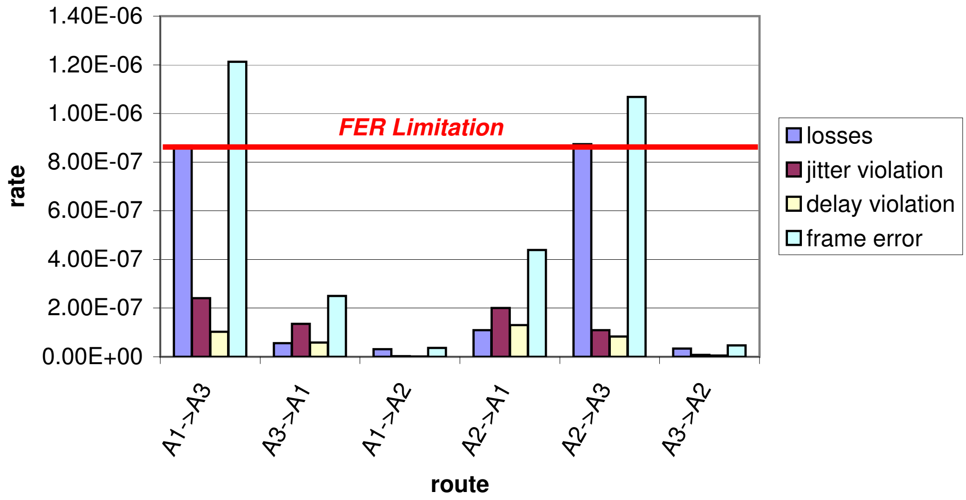

Fig. 4. Quality of service received by probes (full duration)

gate this violation, we have looked at the temporal distribution of frame losses. Figure 5 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 infer that the loss rate may be acceptable for several days before heavy data bursts cause a violation of MEF's frame loss rate requirement. Upgrading this most utilized link will likely results in fulfillment of the loss rate requirement. However, such upgrades in an operational network are beyond the scope of this work.

VI. Conclusion

We have probed the operational campus network of ETH Zurich for several weeks to investigate whether this network offer sufficient QoS to carry the traffic of GSM/UMTS base stations. To our knowledge, we are the first to publish probing data of large scale measurements in network that spans a metropolitan area. We found strong indications that probing traffic can meet the QoS requirements as stated by the MEF. The only violations we have found were on routes that included the most highly loaded network link. Upgrading this most utilized link will presumably results in complete fulfillment of MEF's QoS requirement. Thus, we conclude that lightly-loaded metropolitan Gigabit Ethernets with average utilizations below 1% presumably have a potential to carry traffic from GSM/UMTS base stations.

In the future, we plan (i) to probe other metropolitan networks and (ii) to conduct a combined measurement and simulation study to derive guidelines for provisioning metropolitan networks that carrying traffic from GSM/UMTS base stations.

Moreover, we denote that the probing software employed to perform the measurements is available on request to the authors.

VII. 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] U. Fiedler, P. Huang, and B.Plattner, "Towards Provisioning Diffserv Intra-Nets," in *Proceedings of IWQoS'01*, Karlsruhe, Germany, June 2001, pp. 27–43, Springer.

[2] K. Park and W. Willinger, *Self-Similar Network Traffic and Performance Evaluation*, Wiley-Interscience, NY, 2000.

[3] 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.

[4] Metro Ethernet Forum, "Technical Specification MEF3, Circuit Emulation Service Definitions, Framework and Requirements in Metro Ethernet Networks," April 2004.

[5] "RTAI, Real-Time Application Interface," www.rtai.org.

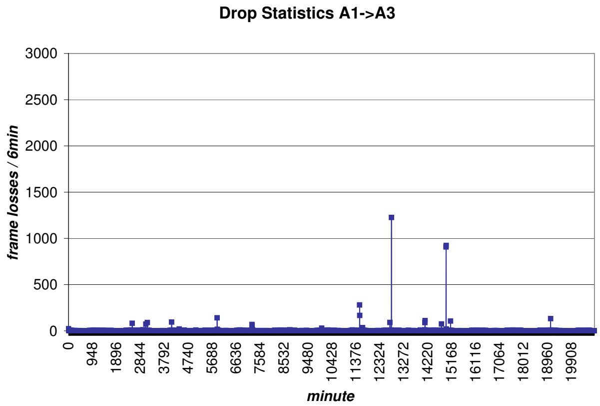

(a) frame loss A1 to A3 (6 min bins)

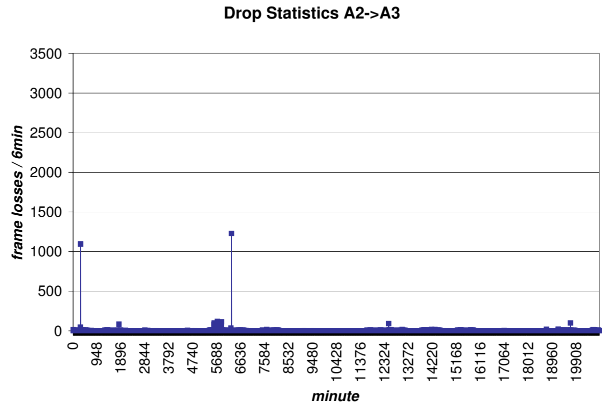

(b) frame loss A2 to A3 (6 min bins)

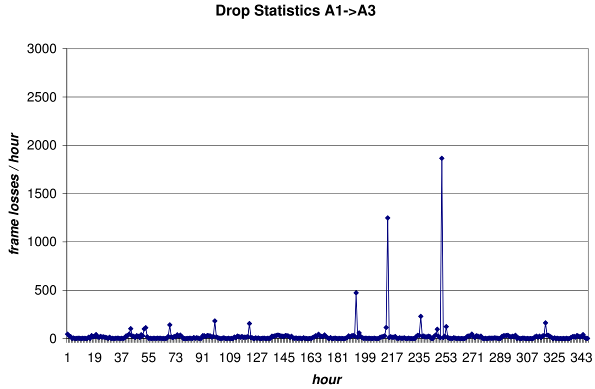

(c) frame losses A1 to A3 (1 hour bins)

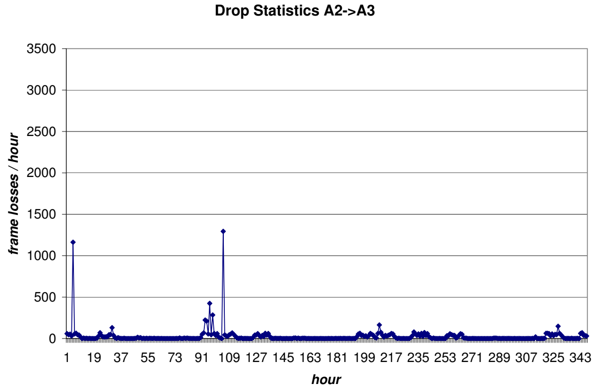

(d) frame losses A2 to A3 (1 hour bins)

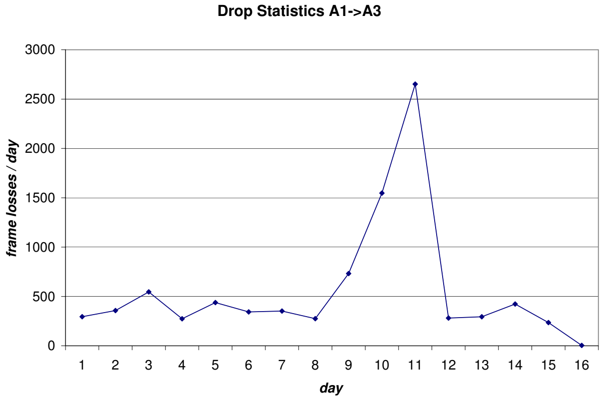

(e) frame losses A1 to A3 (1 day bins)

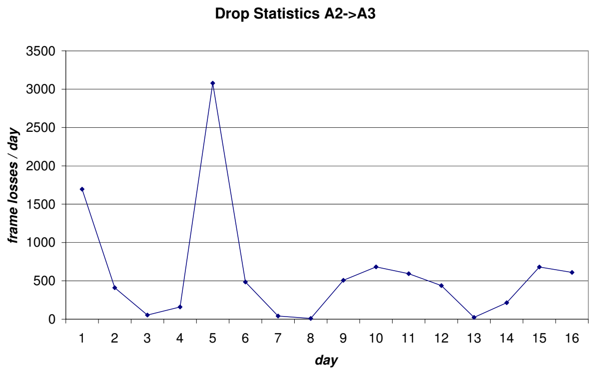

(f) frame losses A2 to A3 (1 day bins)

Fig. 5. Temporal distribution of frame losses on path with high loss

[6] A. Downey, "Evidence for Long-Tailed Distributions in the Internet," in *Proc. of ACM SIGCOMM Internet Measurement Workshop*, Nov. 2001.

[7] Mart Molle Thomas Karagiannis and Michalis Faloutsos, "Long-range dependence, ten years of internet traffic modeling," in *IEEE Internet Computing*. September 2004, pp. 1089–7801, IEEE Computer Society.

[8] M. Karam and F. Tobagi, "Analysis of the Delay and Jitter of Voice Traffic Over the Internet," in *Proceedings of INFOCOM'01*, 2001, pp. 824–833.

[9] G. Kim et. al., "Qos provisioning for voip in bandwidth broker architecture: A simulation approach," in *Proceedings of the communication networks and distributed systems modeling and simulation conference (CNDS)'01*, Phoenix, Arizona, USA, Jan. 2001.

[10] CISCO Systems Inc., "Hierarchical campus design-at-a-glance," 2004, available via www.cisco.com.

[11] A. Pasztor and D. Veitch, "A precision infrastructure for active probing," in *Proc. of PAM*, 2001.

[12] "RTnet, Hard Real-Time Networking for Linux/RTAI," www.rts.uni-hannover.de/rtnet.