

Analyzing campus traffic using the meter-MIB

Remco Poortinga*, Remco van de Meent*, Aiko Pras*

Abstract— The University of Twente, which is the only campus university in the Netherlands, connects all students living on the campus via 100 Mbit/s switched Ethernet links to the campus network. This student network not only interconnects all students among themselves, but also provides, via a 300 Mbit/s link, access to other parts of the university network as well as the external Internet. Since the 300 Mbit/s link is a potential bottleneck, the university has defined a policy to limit the total amount of traffic a single student is allowed to exchange to 50 Gigabyte per week. To check if students do not exceed this limit, a number of interface MIB variables within the student's access switches are monitored on a periodic base. Although monitoring these MIB variables is good enough to detect which students exchange more than 50 Gigabyte weekly, the approach fails to reliably determine which students are the top users of the potential bottleneck, which is the 300 Mbit/s link. The question has therefore been raised whether it would be feasible to directly measure all traffic on that link to precisely determine which student is using which portion of the link's capacity and for what purpose (downloading information from the teaching departments, browsing the Internet etc.). Although NeTraMet, which is an implementation of the IETF's meter MIB, would in theory be useable for such measurements, it was not clear whether it would be fast enough to analyse each packet on a 300 Mbit/s full-duplex link. This paper therefore discusses the experiments that have been performed to find the limits of Netramet, and investigate whether the tool would be useful to solve real management problems. In our experiments Netramet was running on a standard PC.

Index Terms—Flow metering, Network measurements, NeTraMet, IF-MIB.

I. INTRODUCTION

RECENT years have seen advancements in network technology. Switched (fast) Ethernet to the desktop is now ubiquitous. Gigabit for the backbone is introduced at high speed. With these advances in network technology, getting an insight of traffic flows in the network is becoming increasingly difficult. Also the University of Twente, the Netherlands, had a growing need to measure the traffic on one of their high-speed links.

* Remco Poortinga is with the Telematica Instituut, Enschede, The Netherlands. e-mail: Poortinga@telin.nl

* Remco van de Meent is with the University of Twente, Enschede, The Netherlands. e-mail: meent@cs.utwente.nl

* Aiko Pras is with the University of Twente, Enschede, The Netherlands. e-mail: pras@ctit.utwente.nl

A. The CAMPUSnet

All 2000 rooms of the student dorms on the campus of the University of Twente have a 100 Mbit/s full duplex connection to the university's network ('UniNet'). A number of students (living outside the campus) and employees of the university are connected to the university network via ADSL and Cable. The CIV is the organization that manages these and almost all other network facilities available at the university.

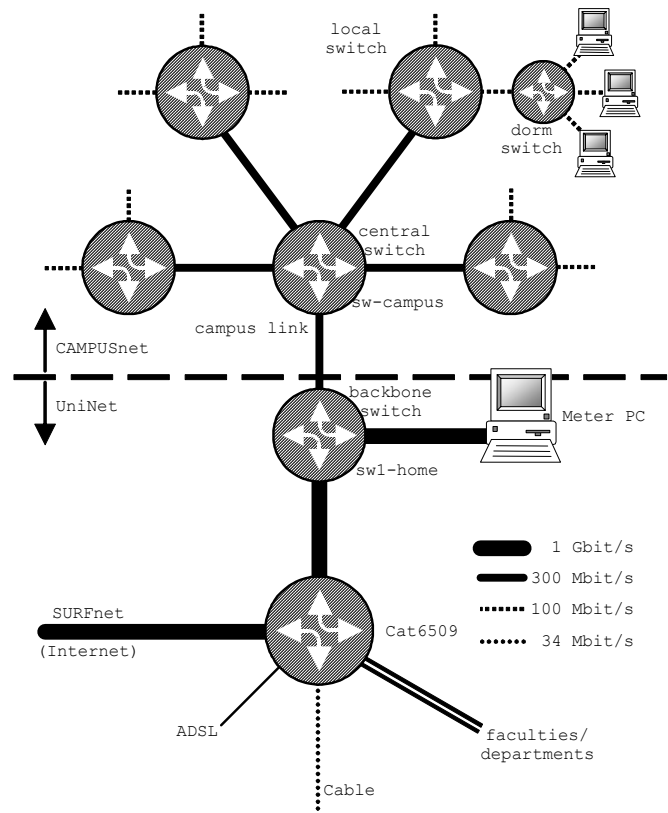


Fig. 1. Part of the university network setup.

Figure 1 shows a schematic representation of a part of the university network. The area above the dashed line is the part of the network that provides connectivity to the student dorms. This part of the network is known as the CAMPUSnet. The CAMPUSnet itself is a fully switched 100 Mbit/s Ethernet with a hierarchical structure.

Every host on the CAMPUSnet connects to a dedicated port on the nearest switch. In general, there is one switch in every student dormitory.

These dorm switches connect to a local switch with a 100

Mbit/s link. Every dorm thus shares a 100 Mbit/s link. The four local switches connect to a central switch (*sw-campus* in Figure 1) via 300 Mbit/s links. The CAMPUSnet connects to the rest of the university network through 3 trunked 100 Mbit/s full duplex channels (*campus link*), providing a 300 Mbit/s full duplex connection, to a backbone switch (*sw1-home*). This trunked link connects the CAMPUSnet to the rest of the university network (and via the university network to the Internet). A number of different servers, such as a student mail server, (not shown in the picture) connect to the backbone switch as well. Some 2000 hosts connect to the CAMPUSnet.

B. Problem Statement

All students connected to the CAMPUSnet share the 300 Mbit/s campus link to the rest of the university network; this link is therefore a relatively scarce resource. Although there are more shared links in the network, this one has to be shared by all 2000 hosts on the CAMPUSnet. For a fair distribution of this resource, the CIV has set a traffic limit of 50 Gigabyte per week per host. This means that the campus link is 'overbooked' by a factor of approximately 5 to 1.

When needed, bottlenecks can be removed by upgrading the relevant equipment or link capacity. The necessity of an upgrade depends on the relevance of network traffic. There is no fixed policy in determining what relevant network traffic is; but in general traffic between the CAMPUSnet and the different university departments is deemed more important than between the CAMPUSnet and the global Internet. The campus link is more likely to be upgraded in case of saturation if a large part of the traffic traversing the link is for communication between the university and the CAMPUSnet.

In order to be able to enforce the limit, the CIV monitors network usage statistics of all hosts connected to the CAMPUSnet. Since every host on the CAMPUSnet has a dedicated port on a switch, the monitoring process consists of scripts reading 64bit MIB counters associated with these ports (see chapter II). In general, the CIV enforces the limit only in times of overloading (e.g. due to excessive usage by a limited number of people).

Students can check their network usage and relative ranking by accessing a web page that displays their daily amount of traffic, as well as an anonymized list of top-users of the network.

Although these measurements provide information on the total amount of traffic sent and received by a host, it does not provide information about the destination of the traffic. As an extreme example: a host could generate or receive more than 50 Gigabyte of traffic on the CAMPUSnet alone, not placing

any load on the campus link at all. In this case enforcing the 50 Gigabyte limit could be considered unfair since the limit intends to ensure fair use of the campus link, not to restrict local traffic on the CAMPUSnet itself. This does not mean that saturating the network with local traffic is allowed, but there is less need to restrict it.

One of the questions therefore is whether setting a limit on the volume internally is a viable method of ensuring fair use of the campus link. To determine this, the CIV wanted to know if the top talkers (in volume of traffic) on the CAMPUSnet are also top talkers on the 300 Mbit/s campus link. In addition the CIV wanted more information on the different destinations of the traffic.

C. Approach

In order to determine the destination of the traffic on the campus link, information needs to be gathered. We used NeTraMet (Network Traffic Meter) [1] to measure and classify traffic across the external link.

A meter PC connects to a gigabit port on the backbone switch (*sw1-home* in Figure 1). By replicating the 300 Mbit/s traffic from the campus link on the Gigabit link, the meter PC can monitor the traffic without disturbing it. The meter PC runs NeTraMet [1] and a MySQL database for storing measurement results.

II. IF-MIB

The standard method used by the CIV to measure traffic volume for every host is by reading IF-MIB counters [2].

The monitoring process consists of scripts reading the 64bit counters `IF-MIB::ifHCInOctets` and `IF-MIB::ifHCOutOctets` once a day for every port to which a host is connected. The counter values read from the IF-MIB are then used to calculate the total volume of traffic for every host over the past 7 days. Figure 2 shows the top talkers in the period from September 28 to October 5 2001 as determined by this method, split into two different directions: out means traffic sent by a host on the CAMPUSnet. IP-addresses used in Figure 2 were changed to protect the privacy of the students involved.

As can be seen in Figure 2, the host that exchanged most of the traffic exceeded the maximum volume set by the CIV by a factor of 4.




Fig. 2. Top-talkers over a one-week period, as determined by reading IF-MIB variables within the dorm switches.

III. NETRAMET

NeTraMet [1] is the first implementation of the RTFM architecture [3], [4]. It is a programmable network traffic meter that classifies observed traffic into different flows based on the 'instructions' (called rule sets) it gets from a manager. Rule sets can be written in the SRL language [5]. NeTraMet can capture the packets from the network directly, but optionally it can get its input from routers using NetFlow (NetFlowMet) or LFAP (LfapMet). We used the standard version of NeTraMet (4.4b10 release), which captures packets directly, running on a dedicated meter PC.

A. Configuration

The meter PC connects to a Gigabit Ethernet port on the backbone switch sw1-home. The port is configured as a monitor port for the 300 Mbit/s campus link.

For administration and configuration purposes, the meter PC has a second (Fast Ethernet) connection to the network, so as not to influence the measurements. Traffic on this second link does not cross the campus link and is therefore invisible to the meter PC itself. Table 1 shows the configuration of the meter PC.

TABLE I
METER PC CONFIGURATION

Component	Specifications
CPU	Pentium-III 1 GHz
Mainboard	Asus CUR-DLS (64 bit, 66 MHz PCI)
Hard disk	60 Gigabyte, UDMA/66
Operating system	Debian Linux, 2.4.5-ac13 kernel, memory mapped I/O
Network interfaces	1Gbit Intel Pro/1000T (Intel driver) 10/100Mbit Intel 82259 LAN (onboard)
Main memory	512 MB reg. SDRAM
Video card	4 MB ATI RAGE-XL (onboard)

B. (Additional) tools

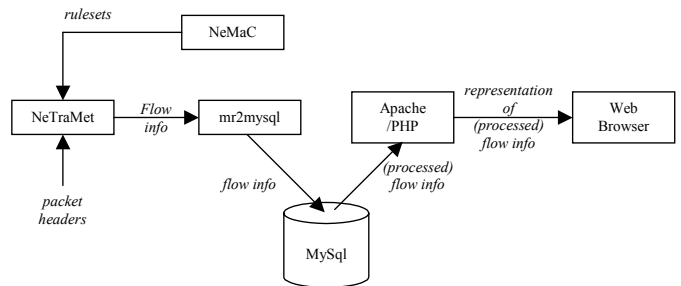


Fig. 3. Measurement set-up.

Several other tools, besides NeTraMet, were used for these experiments. Figure 3 shows the relationships between the different tools that have been used.

To facilitate further processing; information produced by NeTraMet was stored in a MySQL database running on the same computer. Reading the information and storing it in the MySQL database was done by mr2mysql [6]. The information stored in the database by mr2mysql includes:

- Source and Destination IP addresses.
- Transport type.
- Port numbers (if any)
- The rule set associated with the flow.
- The 'FlowKind' (as determined by the rule set)
- The meter that observed the flow
- The number of packets and bytes transmitted and received within that flow.
- The first time the flow was observed (relative to the meter start and converted to an absolute time).
- The last time the flow was active.

The presence of each of these attributes is determined by the rule sets used by the meter.

Mr2mysql is a custom meter reader that reads flow information from a meter at configurable intervals. For every new flow detected by the meter an entry is created in the database. If information about a particular flow is changed since the last time the meter was checked, the relevant flow information in the database is updated. Mr2mysql automatically detects new rule sets downloaded to the meter and incorporates them in the information retrieval process.

Besides updating flow information in the flow database, mr2mysql can also keep a log of changes for every flow by adding an entry containing the current statistics of a flow. Since mr2mysql will create such a 'snapshot' entry every time it checks the meter, it is possible to track changes over time (see for example Figures 4 through 7). Whether or not mr2mysql should take snapshots of a flow is determined by the value of 'FlowKind' for that flow.

We used a combination of Apache, PHP scripts, and the JpGraph library [7] to process the data stored in the database and present it in a graphical form via the web.

C. Rule sets

To obtain the information about the amount of traffic and the different destinations of the traffic across the campus link, we made a number of rule sets. Rule set 1 counts all traffic passing the link:

```
if SourcePeerType == IPv4
{
    if SourcePeerAddress == (campus)
    {
        store FlowKind := 128;
        count;
    }
}
SET all;
```

This rule results in NeTraMet classifying all IPv4 traffic from the CAMPUSnet into one flow. By setting the most significant bit of FlowKind, mr2mysql is instructed to take 'snapshots' of the flows.

```
if SourcePeerType == IPv4
{
    if SourcePeerAddress == (campus)
    {
        if DestPeerAddress == (UniNet)
        {
            if DestPeerAddress == (adsl)
            {
                store FlowKind := 193; # 'A'+128
            }
            else if DestPeerAddress == (cable)
            {
                store FlowKind := 195; # 'C'+128
            }
            else
            {
                store FlowKind := 207; # 'O'+128
            }
        }
    }
}
```

```
    }
}
else
{
    store FlowKind := 216; # 'X'+128
}
count;
}
```

SET destinations;

Rule set 2 was used to determine the various possible source and destinations of interest for these measurements:

- Internal: Traffic with a destination address within the university's address range (i.e. staying within the university network). Internal traffic is subdivided into:
 - ADSL ('A'): Traffic between CAMPUSnet and people connected to the university network using ADSL.
 - Cable ('C'): Traffic between CAMPUSnet and people connected to the university network using a cable modem.
 - Other/local ('O'): Traffic between CAMPUSnet and university departments.
- External ('X'): Traffic with a destination address outside of the university's address range (i.e. 'the global Internet').

Rule set 3 determines the top talkers on the campus link:

```
if SourcePeerType == IPv4
{
    if SourcePeerAddress == (campus)
    {
        save SourcePeerAddress;
        if DestPeerAddress == (UniNet)
        {
            store FlowKind := 201; # 'I'+128
        }
        else
        {
            store FlowKind := 216; # 'X'+128
        }
        count;
    }
}
SET topten;
```

This rule set uses the source IP address (which has to be on the CAMPUSnet) as an attribute to a flow, so that the traffic can be attributed to the different hosts on the CAMPUSnet.

Running this rule set results in a list of all active hosts on the CAMPUSnet. A list of top talkers can later be obtained by the appropriate SQL commands to the database. The rule set makes a distinction between external and internal traffic; where internal means traffic staying within the address range of the university network.




Fig. 4. Bit rate on the campus link over a one-week period.




Fig. 5. Packet rate on the campus link over a one-week period.




Fig. 6. Ratio between traffic out and in on the campus link over a one-week period.




Fig. 7. Average packet-size on the campus link over a one-week period.

D. Measurement results

The measurement results presented are the results of measurements over a period from September 24 until October 5 2001. Figure 4 shows the bit rate on the campus link over this one-week period. The bit rate of the traffic coming from the CAMPUSnet towards the university network ('out') is consistently higher than the bit rate of the traffic in the reverse direction. We did expect an asymmetric bit rate, but we were surprised by the direction of this asymmetry.

If we take a look at the packet rates over the same period (Figure 5) we can see that the rates in both directions are

roughly the same, i.e. the number of packets per second going out is roughly equal to the number of packets coming in. Figure 6 shows the ratio of outgoing traffic over incoming traffic, both for packet and bit rate, again displaying the roughly equal packet rate in both directions and the asymmetric bit rate. The roughly equal packet rates and differing bit rates are reflected in the average packet sizes in either direction shown in Figure 7. The average size of a packet traveling from the CAMPUSnet towards the university network is approximately 1000 bytes. Packets traveling in the opposite direction are approximately 300 bytes in size.




Fig. 8. Top-talkers on the campus link over a one-week period, as determined by NeTraMet.




Fig. 9. Top-talkers on the campus link over a one-week period, differentiated for different destinations.

Figure 8 shows a bar graph of the 10 most active hosts on the campus link over a one-week period. IP addresses in this figure are anonymized the same way as for Figure 2 so that the information from the two figures can be related.

The traffic is split in two directions, where ‘out’ means traffic sent by the host on the CAMPUSnet. Figure 9 shows the same hosts over the same period, but differentiated for two different destinations:

- I. UniNet: traffic from the CAMPUSnet host to another host on the university network.
- II. Internet: traffic from the CAMPUSnet host to a host

somewhere outside of the university network (i.e. the Internet).

For both destinations the volume for each direction is also shown. It is interesting to observe that for the UniNet category data is primarily transferred from the student dorms to the other networks within the university, instead of the other direction.

As can be seen in figures 8 and 9 the most active host caused some 150 Gigabyte of traffic in one week; almost all of which was with hosts outside of the university network.




Fig. 10. Figure 2 and Figure 8 combined.




Fig. 11. Traffic measured by reading IF-MIB variables within the dorm switches, compared to traffic measured on the campus link via NeTraMet.

IV. NETRAMET AND IF-MIB MEASUREMENTS COMPARED

One of the questions the measurements should answer is whether measuring traffic on the switches produces a good estimate of the traffic on the campus link. By plotting measurements, over the same period, from both NeTraMet and the measurements as done by the CIV (i.e. reading the IF-MIB counters present in the switches) in one figure we can determine whether this is the case.

Figure 10 combines both measurements (as shown in figures 2 and 8) in one figure. It shows the volume of traffic per host as reported by the IF-MIB and by NeTraMet. The difference between the two is caused by traffic that remains on the CAMPUSnet itself and is therefore not visible to NeTraMet but is reported by the IF-MIB (i.e. the difference between the two is the local traffic on the CAMPUSnet). The hosts are ordered descending by the IF-MIB reported traffic.

Interestingly enough the top-talker as determined by the IF-MIB produced nearly 200 Gigabytes of traffic in total, but a relatively modest 25 Gigabytes of that traffic went outside of

the CAMPUSnet and was measured by NeTraMet. The top-talker as determined by NeTraMet (see also Figure 8) takes second place in figure 10.

Figure 11 expands on Figure 10 by showing the results for the first 500 hosts. The graph suggests that there is no, or only a weak, relationship between the amount of traffic that remains on the CAMPUSnet itself and the amount of traffic that crosses the campus link.




Fig. 12. Scatter plot of internal and external traffic (dependency removed).

One of the questions was whether measuring traffic with the IF-MIB provided a useful metric for ensuring fair use of the shared campus link. Note that there is a relation between the traffic measured by NeTraMet and the IF-MIB since all traffic passing the shared campus link is also measured by the IF-MIB. The reverse need not be the case, since traffic that is measured by the IF-MIB includes traffic that remains on the CAMPUSnet itself and is therefore not measured by NeTraMet. Measurements by the IF-MIB thus indicate the upper bound for particular hosts, e.g. if the IF-MIB reports a usage of 25 Gigabyte for a particular host; that host can not have used more than those 25 Gigabytes on the shared link.

What we want to determine is whether there is a relationship between the amount of internal traffic (local traffic on the CAMPUSnet) for a host and the amount of external traffic (passing the shared campus link) for that host. If there is a relation between the two the measurements from the IF-MIB might be used as an indication for the traffic over the shared link, instead of just indicating the upper bound.

The amount of traffic staying on the CAMPUSnet itself is determined by subtracting the traffic measured on the campus link from the traffic measured by the IF-MIB. Effectively it splits up the total amount of traffic in an 'internal' part (i.e. local traffic on the CAMPUSnet) and an 'external' part (communication with hosts outside the CAMPUSnet). By combining the internal and external traffic for all hosts in a

scatter plot we can see if there is an obvious relation between the two. Figure 12 shows the resulting scatter plot of all active hosts in the measurement period. There is a moderate correlation of $r=0.53$ between the 'internal' volume (traffic local to the CAMPUSnet) and 'external' volume (traffic crossing the shared campus link). The 'heavy' users seem to generate either a lot of internal or a lot of external traffic. The correlation for the top one hundred users is -0.46 .

Figure 12 is divided into three regions:

- I. Hosts that produced less than 50 Gigabytes of traffic in total.
- II. Hosts that produced more than 50 Gigabytes of traffic in total, but less than 50 Gigabytes on the shared link.
- III. Hosts that produced more than 50 Gigabytes of traffic on the shared link.

Hosts in the second category (some 5%) are the ones that would be reported by the IF-MIB for producing too much traffic, but not by NeTraMet.

Figures 13 and 14 show the cumulative curves for the percentage of hosts responsible for a percentage of traffic, both internal and external. This is done separately for traffic sent (or 'out') and traffic received (or 'in') to prevent local traffic being counted twice (local traffic sent by one host is local traffic received for another). If the campus link were perfectly shared, these graphs would show a straight diagonal line at a 45° angle for the external traffic. The solid lines show that for external traffic between 10% and 15% of the hosts are responsible for 80% of the traffic in either direction. For internal traffic the distribution is very different between traffic sent and received. The graphs show that a lot of traffic is sent by a relatively small percentage of hosts but that reception of traffic is distributed more evenly over a larger percentage of the hosts, suggesting that a limited number of hosts serve a lot of others locally.




Fig. 13. Percentage of hosts responsible for percentage of traffic sent.




Fig. 14. Percentage of hosts responsible for percentage of traffic received.




Fig. 15. Different destinations of traffic on the campus link.

Figure 15 shows the different destinations of the traffic over a one-week period as measured on the campus link. The figure shows that most of the traffic goes to destinations outside of the university network. Some 10% is traffic for users connected via the university's cable modem network. Although ADSL traffic is almost non-existent, this will change soon as ADSL is currently being introduced to all students and employees living outside the campus. During the measurement period there was only a small group of test users using ADSL.

V. NETRAMET LIMITS

Using NeTraMet to perform these measurements gave an opportunity to determine the limits of NeTraMet in terms of the maximum network traffic it can handle. In order to measure this, the CPU load on the machine over time was recorded using `vmstat(8)` and stored in the database with a Perl script.

By plotting the traffic load vs. the load on the CPU in a scatter plot the maximum traffic load the system can handle can be estimated.




Fig. 16. Estimating maximum packet rate the system can handle.

Figure 16 shows the system load plotted versus the number of packets per second passing the meter for two cases. With the top 'line' all three rule sets were running on the meter, as well as the MySQL database and the mr2mysql tool. The maximum number of packets per second this system can then handle is approximately 60 thousand. Due to the relatively large packet sizes, 1000 bytes for outward traffic and 300 for inward traffic, this equals roughly 350 megabits per second.

For the bottom 'line' only the first rule set (counting all traffic) was running, as well as the database and the mr2mysql tool. In this case the upper limit is about 90 thousand packets per second. By running the database and the mr2mysql tool on a separate computer, the upper limit was increased to nearly 100 thousand packets per second.

The relation between packets per second and the system load is linear except in overload situations. Figure 16 shows such an overload situation with a number of points at and near a 100% load; the result of a flood attack resulting in an overload situation that lasted for several minutes.

VI. CONCLUSION

The university monitors traffic for every host on the CAMPUSnet by reading certain IF-MIB counters within the access (dorm) switches. The question discussed in this paper is whether information gathered via this method can be used to ensure fair use of the backbone link that connects the CAMPUSnet to the rest of the university's networks as well as the Internet. Wouldn't it be possible to get a better insight in the traffic over this scarce resource by monitoring all traffic passing this backbone link via tools like NeTraMet? Are such tools capable to measure traffic even at speeds of several hundreds of Megabits per second?

From our measurements we learned that the current method of reading IF-MIB counters within the access switches is not really adequate to predict the load that an individual student places on the shared backbone link that connects the CAMPUSnet to the external world. To determine how much load individual students put on the backbone link, that link should be measured directly via tools like NeTraMet.

Our measurements also showed (Figures 13 and 14) that the 'vital few and trivial many' (also known as the 'Pareto Principle' or the '80-20' rule) applies to the shared backbone link and the amount of traffic sent on the CAMPUSnet. However, this inequality rule does not apply to the amount of traffic received locally on the CAMPUSnet.

We found that NeTraMet, running on an ordinary PC, was capable to reliably capture and analyze all packets flowing over the campus link. Depending on the number of rule sets in use; it could easily handle 60 thousand packets per second, which is equal to roughly 350 Mbit/s.

ACKNOWLEDGMENT

We would like to thank the CIV for making these

measurements possible by allowing us to do these measurements and for providing assistance whenever necessary.

This paper is based on the work performed within Work Unit 5 of the Internet Next Generation project [8]. This project is part of the Dutch Gigaport program, and sponsored by the Telematica Instituut (TI).

The following organizations are member of the Internet Next Generation project: the University of Twente (UT), Ericsson ETM, Ericsson EMN, KPN Research and TI. We would like to thank all members of the project for their contributions.

REFERENCES

- [1] N. Brownlee, "NeTraMet & NeMaC reference manual, Version 4.3," *Information Technology Systems & Services, The University of Auckland*, New Zealand, June 1999. <http://www2.auckland.ac.nz/net/Accounting/ntmref.pdf>
- [2] K. McCloghrie, F. Kastenholz, "The Interfaces Group MIB," RFC2863, June 2000.
- [3] N. Brownlee, "Traffic Flow Measurement: Meter MIB," RFC2720, October 1999.
- [4] N. Brownlee, C.Mills, G.Ruth, "Traffic Flow Measurement: Architecture," RFC2722, October 1999.
- [5] N. Brownlee, "SRL: A Language for Describing Traffic Flows and Specifying Actions for Flow Groups," RFC2723, October 1999.
- [6] "mr2mysql: meter reader to MySQL," <http://mr2mysql.sourceforge.net>.
- [7] "JpGraph 1.5 - An OO Graph library for PHP4," <http://www.aditus.nu/jpgraph/>.
- [8] "Internet Next Generation project," <http://ing.ctit.utwente.nl/WU5/>.