Net Performance Measurements IPP – FOM – UU: Focused at IPP

Hans Blom Mark Jansen Cees de Laat
Workgroup Computational Physics
Utrecht University

February 2, 2000

1 Introduction

To check if the network supplies sufficient bandwidth between IPP Jülich and FOM Nieuwegein for the remote participation services, some performance tests were executed between these sites and some other, control sites situated at the network providers along the route and at Utrecht University (UU). The sites are described more precisely further on in section 3.

The measurements are performed with a package called RTPL (Remote Throughput, Ping and Load) developed by the Institute of Computational Physics at Utrecht University. Several, basic aspects of the network performance are measured by this package. In fact the real measurements are performed by standard software. The package uses scripts to invoke the net performance software. These scripts are also responsible for the sampling and storage of the results. The presentation of the results is a separate part. It is Web based and dynamic: the user can select various views at the data. In section 2 the package will be described in greater detail.

At the moment the RTPL is only available for the Unix operating system. The reason is that the mechanisms to do the measurements at remote sites (remote shells) is a standard for Unix but not for Windows NT/9X: there are implementations of remote shell servers for Windows, but they are in general proprietary or not very secure. The used net performance software and scripting language are however available for Windows NT/9X. Please note that the Web presentation is platform independent.

2 Description

The network performance between the sites at IPP, FOM and UU is measured with a package called RTPL (Remote Throughput, Ping and Load). This package does periodically net performance measurements between a

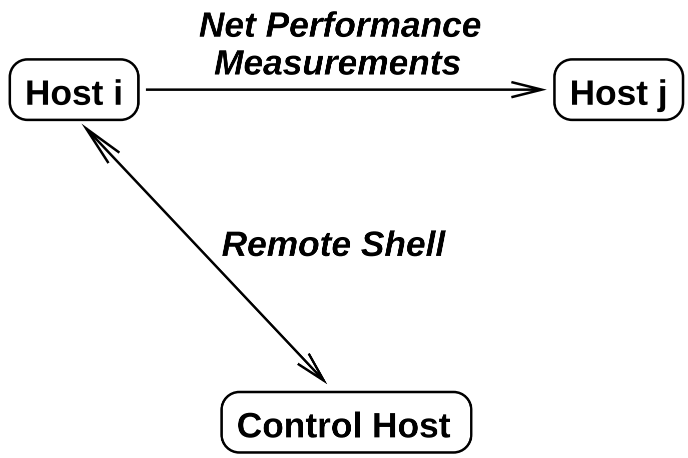

Figure 1: *The control host starts the measurements at all hosts i by means of remote shells. The host i does the measurements to the hosts j The results are send back by host i to the control host.*

set of hosts which can be specified by the user. A, so called, control host is used to start the net performance measurements at each of the participating hosts with a, so called, remote shell command. Also the results of the tests are send back via the remote shell command. See figure 1.

Please note that a remote shell offers the possibility to start a command at a remote host. It also connects the I/O streams of the command at the remote host to the process, which started the remote shell, in a transparent way: that is for the process it seems that the command is running at the local host Security aspects are solved in various ways, depending from the remote shell command used. However in all cases there is no need to specify a user / password to get access to the remote host: this would deny a non-interactive run of the command at the remote host.

At each host i of the set the following performance tests are executed:

Throughput. Formal definition from RFC 1224 [1]: "The maximum rate at which none of the offered frames are dropped by the device". It is a way to quantify the traffic flow which can be handled by a network connection. Default it is measured for the connections from the current host i to all other hosts, but it is also possible to skip connections. The throughput is measured with the public domain command `netperf`.

Roundtrip. This Internet application is described in RFC 2151 [2] by the paragraph quoted below:

> Ping, reportedly an acronym for the Packet Internetwork Groper, is one of the most widely available tools bundled

with TCP/IP software packages. Ping uses a series of Internet Control Message Protocol (ICMP) [3] Echo messages to determine if a remote host is active or inactive, and to determine the round-trip delay in communicating with it.

The roundtrip time quantifies the response offered by a network connection. It will be measured, before the throughput, across the same connections as the throughput. The roundtrip time is measured with the system command `ping`.

Load. This is expressed here as # fully active processes at a host. It is no network quantity, but it may help to explain unexpected performance decreases. The load is measured at the current host i, using the system command `uptime`.

The sampling of the results at the control host and the measurements at all hosts, participating in the tests, are performed by scripts in the scripting language *Perl*. Perl is a powerful scripting language which is available for many platforms, including Unix, Windows NT/9X and MacOS.

The Perl script at the control host collects the results of the measurements for each host i and stores the results in ZIP compressed data files. The ZIP compression is used to reduce disk space and download time (see below).

The presentation of these results is Web based: a JAVA Applet is used to load the data from the files into the memory of the Web browser from a user looking at the results. Please note that an Applet is an architecture independent application, written in the JAVA programming language which runs in the so called JAVA virtual machine of the Web browser. The functionality to read (ZIP compressed) data files from a Web browser is a JAVA feature. This implies that access to the data is only possible via the Applet which had read the data.

The HTML scripting language JavaScript is used to dynamically present the user various HTML tables of the data. That is: the user selects a view at the data and the HTML code is generated on demand by JavaScript. JavaScript applies direct calls to Applet methods to obtain the required data for the HTML table to display. A direct call of Applet methods by JavaScript is a common functionality of most Web browsers. Further on the Applet can also be used to present various plots of the data to the user. The plots are displayed in a new window.

The following data files are available to be viewed via the Web:

- The data of the last 7 days.
- For each week of the last half year a data file is available.
- The week mean values from the last year.

- The day mean values from the last year.
- The mean values, calculated at the periodic measurement times, for the days of the week, averaged during a quarter. The data are stored during a year.
- The mean values, calculated at the measurement times, for the workdays of the week, averaged during a month. The data are stored during a year.

3 Sites IPP – FOM – UU

Below a description of the participating sites for the particular setup IPP – FOM – UU will be given. The same titles are used as in the Web pages displaying the results. The following hosts are used here:

- **IPP**; bandwidth: 10 Mbit/s; location: Jülich.
  This host is situated at the "Institut für Plasmaphysik", "Forschungszentrum Jülich". The IPP will participate in the remote participation services, discussed in this report.
- **ZELAS**; bandwidth: 10 Mbit/s; location: Jülich.
  This host is situated at the "Central Electronics Laboratory", FZJ. This host has been add to make a comparison possible with the host at the IPP and to be able to measure the inluences local net congestion.
- **ZAM**; bandwidth: 100 Mbit/s; location: Jülich.
  This host is situated at the "Zentralinstitut für Angewandte Mathematik", FZJ. It is situated near the router and has a larger bandwidth available than the IPP host. Hence its participation in the tests.
- **RUS**; bandwidth: 10 Mbit/s; location: Stuttgart.
  This host is situated at the "Rechenzentrum der Universität Stuttgart". It has been add to check for congestion at the "Deutsche Forschungsnetz" (DFN).
- **SARA**; bandwidth: 100 Mbit/s; location: Amsterdam.
  This host is situated at the "Academic Computing Services Amsterdam". It is situated near the backbone of the Dutch SURFnet network.
- **FOM**; bandwidth: 10 Mbit/s; location: Nieuwegein.
  This host is situated at the "FOM-Institute for Plasma Physics". The FOM is also a participant in the remote participation services.
- **UU-36**; bandwidth: 100Mbit/s; location Utrecht.
  This host is situated at the "Institute for Computational Physics", Utrecht University.

The topology of the network between these sites is shown in figure 2.

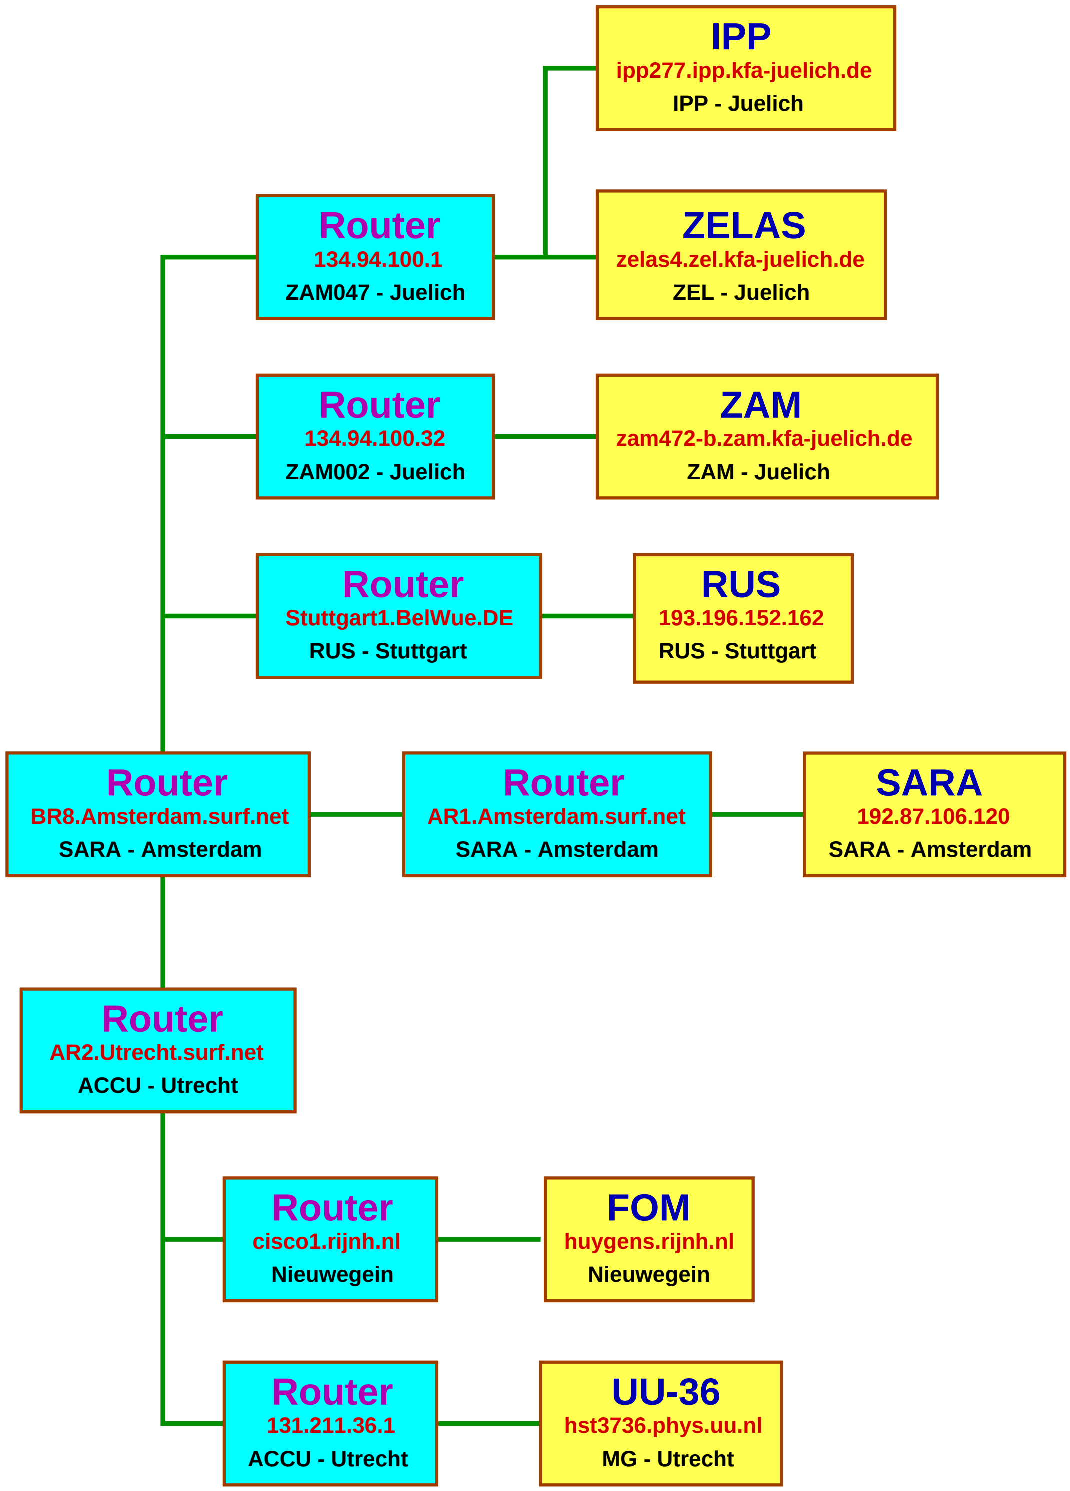

Figure 2: *Topology sites participating in the net performance tests.*

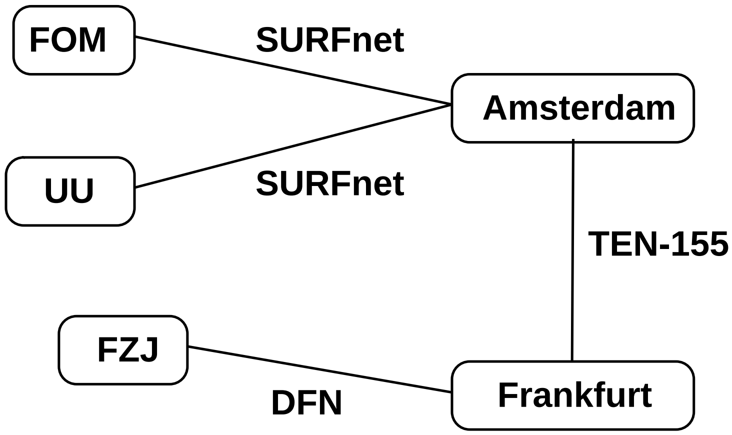

Figure 3: *Topology involved providers.*

4 Measurements IPP – FOM – UU

4.1 Introduction

In this version of the package the network performance was monitored between the sites described in section 3. In the comparison of the results we will focus to the connections from IPP and ZAM to FOM, SARA and UU. The topology of the involved providers is shown in figure 3. Its is expected that a bandwidth of 10 Mbit/s is required for the connection with the FZJ. In our analysis we will use this values as a target.

During the month January 2000, the network seems to be rather stable, although the throughput values during daytime in the first week of January were higher than in the other weeks. The reason probably is that in the first week of January many people still did have holidays. Nevertheless, besides the typical daily data we also now have the opportunity to present mean data.

As typical daily data we selected the two days from Wednesday 26/01/2000 23:00 back in history until and included Tuesday 25/01/2000 00:00. These days show a rather typical behavior, however, incidentally during daytime the throughput values between Germany and the Netherlands were still somewhat worse than shown in these pages.

For the mean data the workday mean values averaged during January 2000 were selected. Please note that the mean values are calculated at the periodic measurement times, such that still a distribution across the day is obtained. See section 2, page 3.

Please note that there are also mean values available for the days of the week, averaged at an quarterly basis. But we feel that after one month the statistic is still insufficient: the average at each measurement time and at

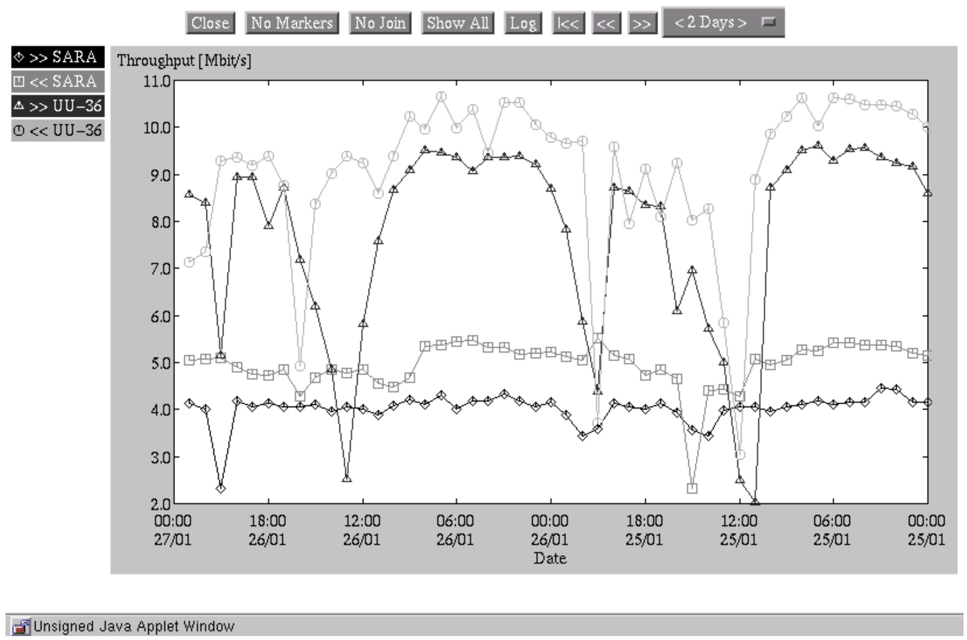

Figure 4: *This plot shows the throughput values between the ZAM host and the SARA and the UU-36 hosts. All these hosts have 100 Mbit/s interfaces. Typical data from Wednesday 26/01/2000 and Tuesday 25/01/2000 are presented here.*

each day of the week is only based at four or five individual throughput values.

We do not present roundtrip results here. the reason is that we found no, or hardly any, unusual observations: no unusual long roundtrip times and no unusual package lost.

In the following sub sections we will make several comparisons between the performances of the sites.

4.2 Typical daily throughput

In this subsection the typical data of Wednesday 26/01/2000 and Tuesday 25/01/2000 are presented.

In figure 4 the throughput measurements between the ZAM host and the SARA and UU-36 hosts are given. These hosts are selected here, because all have an interface with a bandwidth of 100 Mbit/s. It is the intension of this plot to check for the speed of the 100 Mbit/s interface at Jülich. Figure 5 displays the throughput measurements between the IPP host and the FOM and UU-36 hosts. Note that the IPP and FOM hosts have 10 Mbit/s interfaces, in contradiction to the 100 Mbit/s interface of UU-36. This plot

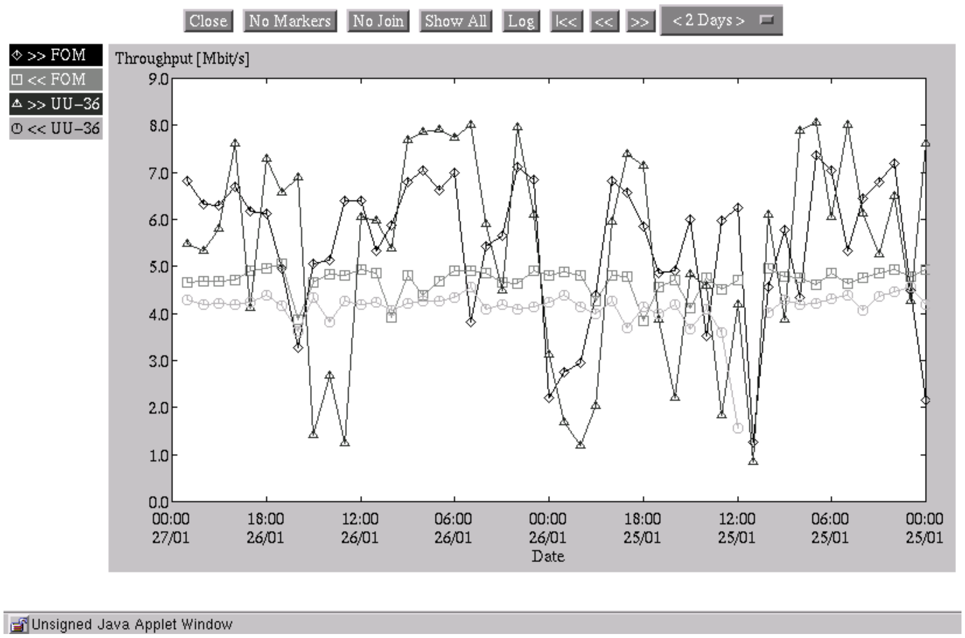

Figure 5: *This plot shows the throughput values between the IPP host and the FOM and UU-36 hosts. The first two hosts have 10 Mbit/s interfaces, the latter has a 100 Mbit/s interface. Typical data from Wednesday 26/01/2000 and Tuesday 25/01/2000 are presented here.*

is intended to compare the 10 Mbit/s interface of IPP with the 100 Mbit/s interface of ZAM. To make a comparison between the two DFN routes to Jülich and Stuttgart, figure 6 shows the throughput data between the FOM host and the IPP (Jülich) and RUS (Stuttgart) hosts. All hosts have 10 Mbit/s interfaces.

From these figures the following conclusions can be drawn:

1. In figure 4 we see rather deep daily performance decreases (12:00 – 18:00) especially in the connection ZAM – UU-36. The decrease in the direction ZAM >> UU-36 is deeper and a bit earlier than in the other direction.

2. The performance decreases (figure 4) are less clear found at in the connection ZAM – UU-36. However, the overall performance of this connection is much lower. The reason therefore is unclear.

3. In figure 4 we also see performance lost during the evening (19:00 – 23:00). Probably due to Internet browsing from people at home.

4. In figure 5 the same performance dips appear as in figure 4, however, due to the smaller bandwidth of the interfaces they are less profound.

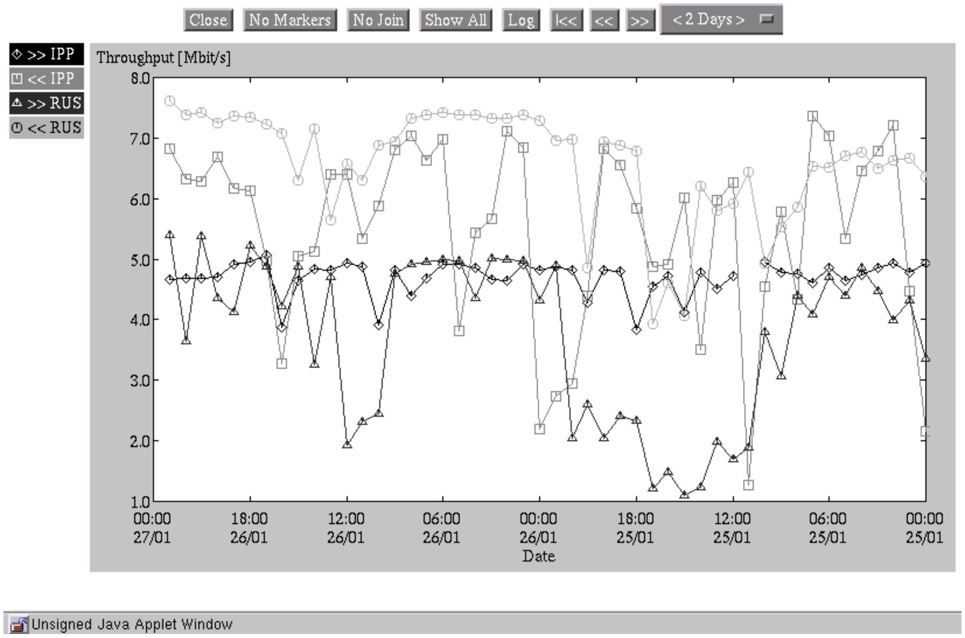

Figure 6: *This plot displays the throughput values between the FOM host and the IPP and RUS hosts. All hosts have 10 Mbit/s interfaces. Typical data from Wednesday 26/01/2000 and Tuesday 25/01/2000 are presented here.*

5. In figure 5 it also that a higher throughput can be reached at the connection IPP – UU-36, than at IPP – FOM, although the performance dips are deeper and even below the IPP – FOM connection.

6. Comparing in figure 6 the connection FOM << IPP with FOM << RUS, it appears that the performance in the connections from Jülich, also found in the figures 4 and 5 is less profound in the connection to Stuttgart. However, there is a considerable performance dip in the connection FOM >> RUS.

7. From all figures it follows that the throughput from the FOM to Germany is lower than in the other direction, while for host UU-36 the reverse is valid.

4.3 Average workday throughput in January

In this subsection the workday (Mon – Fri) average values during January are displayed. The mean values are calculated at the periodical measurement times. More precisely, this means that a throughput measurement executed at *hh*:*mm* will be added to the mean value for *hh*:00. The mean values are

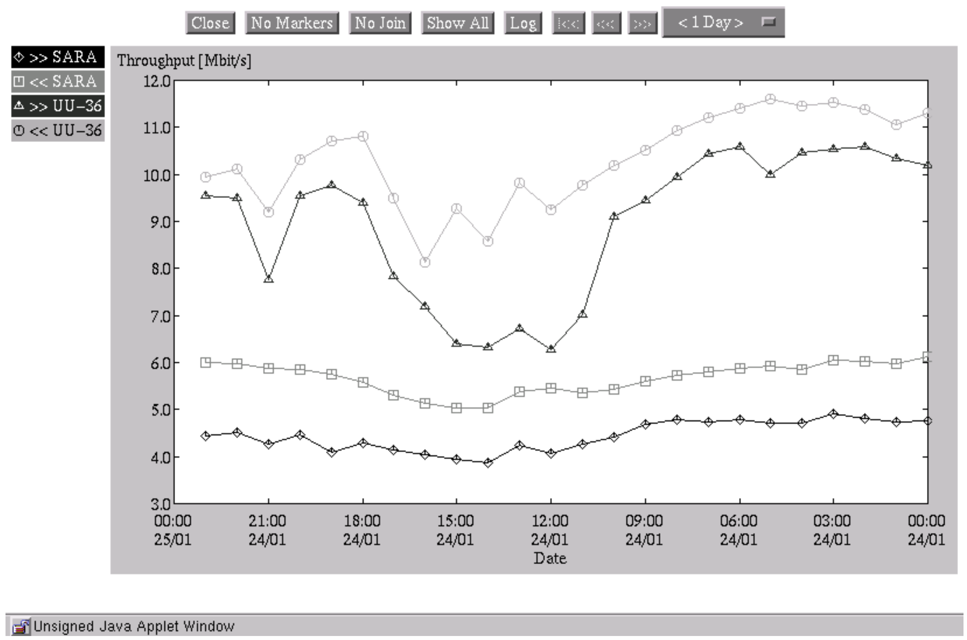

Figure 7: *This plot shows the throughput workday mean values between the ZAM host and the SARA and the UU-36 hosts. All these hosts have 100 Mbit/s interfaces. The data are obtained in January 2000. This plot is the averaged equivalent of figure 4.*

averaged about twenty measurements. The following workday mean data are shown below.

In figure 7 the throughput workday mean values for the measurements between the ZAM host and the SARA and UU-36 hosts are presented. All hosts have an interface with a bandwidth of 100 Mbit/s. This plot is the averaged equivalent of figure 4, but only one day is displayed here. Figure 8 displays the throughput workday mean values between the IPP host and the FOM and UU-36 hosts. The IPP and FOM hosts have 10 Mbit/s, in contradiction to the 100 Mbit/s interface of UU-36. This plot is the averaged equivalent of figure 5. To make a comparison between the two DFN routes two Jülich and Stuttgart, figure 9 shows the throughput workday mean values between the FOM host and the IPP (Jülich) and RUS (Stuttgart) hosts. All hosts have 100 Mbit/s interfaces. This plot is the averaged equivalent of figure 6.

From these figures the following conclusions can be drawn:

1. The performance decreases for the connection ZAM – UU-36 are in the plot with average values (figure 7) less deep than in the daily plot (figure 4). A part of the explanation is that in the beginning of January

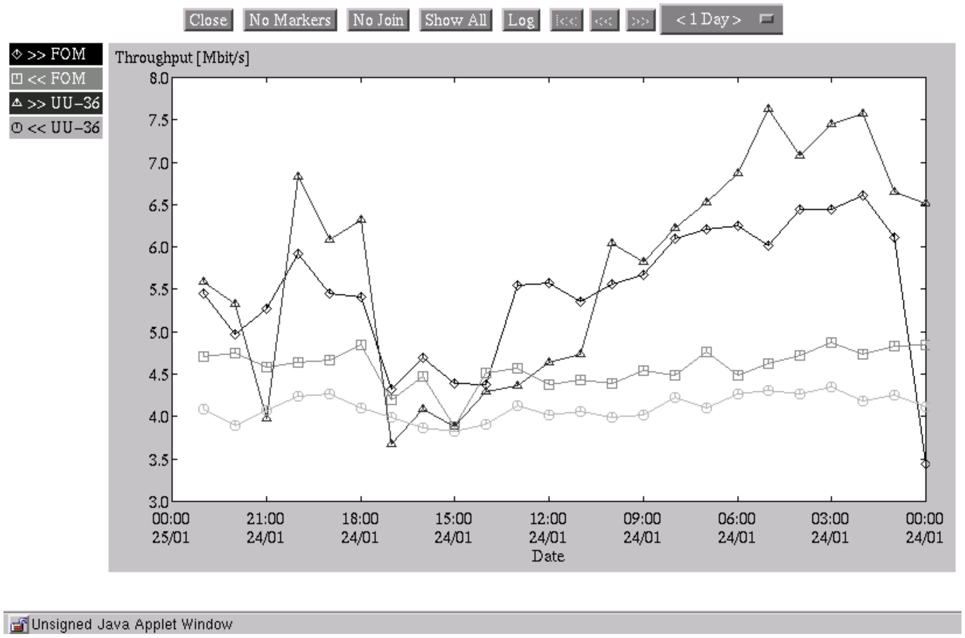

Figure 8: *This plot shows the throughput workday mean values between the IPP host and the FOM and UU-36 hosts. The first two hosts have 10 Mbit/s interfaces, the latter has a 100 Mbit/s interface. The data are obtained in January 2000. This plot is the averaged equivalent of figure 5.*

the performances during daytime were better. Another explanation is that at Mondays and Friday network performances seems to be better during daytime. The quarterly mean data for the days of the week will make this more clear. A third possible explanation is that the daily performance decreases are not always at the same time of day. This implies that the performance dip is broader, but not so deep.

2. The daily data plot of the connections IPP – FOM and IPP – UU-36 (figure 5) is in good agreement with the plot of the workday mean data (figure 8). However, in the latter plot the fluctuations are clearly averaged out. About the same explanations as in item 1 can be given here.

3. The same conclusions for the comparison between the daily data plot of the connections FOM – IPP and FOM – RUS (figure 6) and the workday plot (figure 9) as in item 2 can be drawn.

4. The performance dip during daytime in the connection FOM >> RUS (figures 6 and 9) are also found in the other connections to the RUS host.

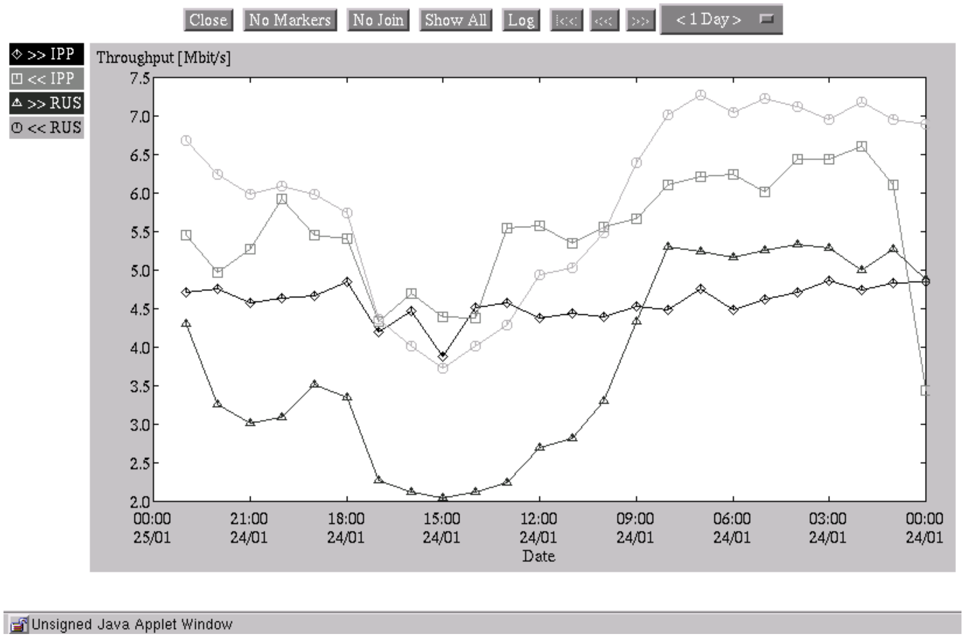

Figure 9: *This plot displays the throughput workday mean values between the FOM host and the IPP and RUS hosts. All hosts have 10 Mbit/s interfaces. The data are obtained in January 2000. This plot is the averaged equivalent of figure 6.*

5 Conclusions

From the measurements in section 4 the following, general conclusions can be given:

1. The for remote control required bandwidth is not obtained during the busiest hors during daytime. However, the workday average show that the mean performances decreases are not so large as may be expected from the daily values.

2. The hosts with 100 Mbit/s interfaces clearly give better better results than the hosts with 10 Mbit/s interfaces. Hence the importance to upgrade everywhere the LAN's to 100 Mbit/s.

3. Participation of hosts at the TEN-155 would help very much in the detection of congestion takes places in the NRN's or in the TEN network. Thereby it is not so important that the TEN hosts have a "fast" connection with the TEN network, but that they make it possible to "isolate" the TEN section from the NRN's, hence making a better analysis of the results possible.

References

[1] S. Bradner, “Benchmarking terminology for network interconnection devices,” RFC 1242, Internet Engineering Task Force, July 1991. URL: `http://www.ietf.org/rfc/rfc1224.txt`.

[2] G. Kessler and S. Shepard, “A primer on Internet and TCP/IP tools and utilities,” RFC 2151, Internet Engineering Task Force, June 1997. URL: `http://www.ietf.org/rfc/rfc2151.txt`.

[3] J. Postel, “Internet Control Message Protocol,” RFC 792, Internet Engineering Task Force, September 1981. URL: `http://www.ietf.org/-rfc/rfc0792.txt`.