Universiteit Utrecht

rTPL: a Package to Measure Network Performance from an End User Viewpoint

WFI-2000-01

Hans Blom Cees de Laat
Workgroup Computational Physics
Utrecht University

February 22, 2000

1 Introduction

The rTPL package (Remote Throughput, Ping and Load) is a package intended to do periodically net performance measurements tests between a set workstations to obtain the network performance from a user viewpoint. The performance measurements are consisting of roundtrip and throughput measurements between a number of hosts. Default all pairs in the hosts set are used, but it is also possible to select the host pairs for the tests. Also the machine load of the participating workstations is measured. In this way performance lost can be related to heavily machine load.

Because we are interested in the network performance from a user viewpoint and not in the maximum possible capacity of a network, the measurement parameters are configured to default values and the test durations are limited: various long term statistics will blur the short time fluctuations.

The tests are performed by scripts which are written in the generally used scripting language Perl, which is available for many operating systems. The scripts are used to parse, sample and organize the results. The real performance measurements are executed by commonly used Internet tools invoked by the Perl scripts.

The presentation of the results is Web based and dynamic: the net performance data are stored in ZIP compressed plain text files which are accessible from a Web server. There are various files such that a user can be offered several views at the data, including several time based averages. The file data are read into the Web browser of a user by a Java Applet. The HTML scripting language JavaScript is used to display the data in various tables. The Applet can also be used to present the data in plot form.

The rTPL package is only available for Unix platforms. The main reason is that the so called remote shell mechanism, used to obtain the results from the participating hosts, although implemented on Windows NT/9X, is not standard and not generally available. There is also a standard form of the package which only executes roundtrip tests from the local host to a specified set of remote hosts. This basic package is available for both Unix and Windows NT/9X. The basic version is not further discussed in this report.

The roundtrip tests via the ping service is also used by the HEP Internet Monitoring Project (*PingER*) [1]. This is an initiative to measure the performance of the Internet used by the High Energy Physics Research community. However, we believe that, especially for networks with a larger capacity also throughput measurements are required to become a better overview of the network performance.

This report is organized as follows. Section 2 gives a more detailed description of this package. Also the software to install and the user requirements are discussed here. Section 3 describes some examples of net performance measurements which are currently running. Section 4 contains

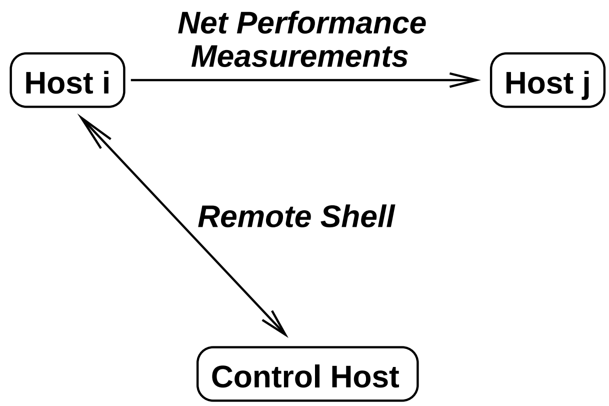

Figure 1: *The control host starts the measurements at all hosts i by means of remote shells. The host i performs the measurements to the hosts j The results are send back by host i to the control host.*

a proposal for a start setup of a performance test at the TEN-155 network.

2 Description and Requirements

2.1 Measurement procedure

In this section a description of the rTPL package (Remote Throughput, Ping and Load) is given in more detail. As already explained in the introduction, this package does periodically net performance measurements between a set of hosts. The measurements are performed by a, so called, control host. This workstation starts the net performance measurements at each of the participating workstations with a, so called, remote shell command[1]. Also the results of the measurements are send back via the remote shell command. This process is schemed in figure 1.

At each host i of the set hosts, described above, the following net performance measurements are executed:

Throughput. Formal definition from RFC 1224 [2]: "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

[1] A remote shell makes it possible to control a remote process (nearly) in the same way as it was running at the local host. To be able to do this it connects the Standard I/O streams from the remote host to the local host using sockets. Also the relevant signals are transported to the remote 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. A modern implementation of the remote shell is the *Secure Shell* (`ssh` command) which offers an encrypted connection with the remote host.

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 or Ping. This Internet application is described in RFC 2151 [3] 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) [4] 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 subsection 2.2 for a description of the data files.

The tests are periodically started with the so called `crontab`[2] command. A sample `crontab` input file, which is specific for a hosts set, will be generated at the installation of that set. When the previous test is not ended when a new one is started by `crontab`, the new test terminates. However, to prevent deadlocks, after some terminations the new version kills the old one and continues. Timeouts are used in the test scripts to prevent deadlocks.

[2]With the `crontab` command, a user can start programs periodically at specific dates and or times using the `cron` daemon (server). It is available at most Unix systems. There also exists a public domain, source version.

2.2 Presentation results and data files

The presentation of these results is Web based: a Java Applet[3] is used to load the data from the files into the memory of the Web browser from a user analyzing the results. The functionality to read (ZIP compressed) data files from a Web browser is implemented in Java. Therefore, the Applet is required for access to 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. From the JavaScript code direct calls to Applet methodss are applied to obtain the required data for the HTML tables to display. The direct call of an Applet method is a standard JavaScript feature which is (and should be) supported by most Web browsers. The Applet is also used to make various plots of the data. The plots are displayed in a new Applet window.

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

- The data of the last 7 days.
- The data during a week. There are several week data files available. The # weeks where data are stored, can be selected by the user, however, the maximum period is a year.
- 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, and determined for the days of the week. They are averaged during a quarter. The data for the quarters are stored during a year.
- The mean values, calculated at the periodic measurement times, and determined for the workdays of the week (Monday – Friday). They are averaged during a month. The data for the months are stored during a year.

The user can download the presented data file to the local host. He/she can also download a special version of the RTPL package which can be used to display the downloaded data files off line. This version only contains the Web files from the package.

2.3 Test modes

The net performance tests can be executed in the modes explained below. These modes may be used in all possible combinations. Each selected mode

[3] A Java is an architecture independent application, written in the Java programming language which runs in the so called JAVA virtual machine of the Web browser.

leads to a different presentation of the tables and plots to the user. The modes to use can be selected at the configuration of the set of hosts by the user.

- Do the tests across a Quality-of-Service route between the hosts and via the “usual” Internet route. Alternatively only one route between each host pair can be selected.
- Do tests between each host pair only one-directional or do the tests two-directional.

2.4 Platforms

The RTPL package is in principle available for all Unix platforms. It has been and is used with: BSD/OS (Free BSD), OSF1 (Digital Unix), IRIX64, HP-UX, Linux (RedHat), SunOS. For other platforms modification may be needed, especially for the script `host_net_test` which executes the measurements at host *i*. However, these modifications are in general simple: they are caused by small differences in the options and output format of the programs, used by the scripts, at the different Unix platforms. This is the price one has to pay for using standard system programs. Fortunately the programs used by the scripts are a fundamental part of the system. Therefore, there are hardly any modifications in user interfaces of these programs during system upgrades.

2.5 Software

Below the software is listed which should be installed at the control host and the test hosts (see figure 1). Also download URL's are specified.

Perl 5 The scripting language Perl can be downloaded from `http://www.perl.com/`.

Netperf (*test hosts only*) The Netperf throughput measurement package can be downloaded from `http://www.netperf.org/netperf/NetperfPage.html`.

Web server (*control host only*) A HTTPD server is used to present the results via the Web. Note that no special rights at the Web server are needed to display the results. A popular Web server is the *Apache* server. It can be downloaded from `http://www.apache.org/httpd.html`. It is no limitation to run the Web server as a regular user.

2.6 Requirements

The following requirements should be fulfilled to be able to run the package.

Roundtrip Parameter	*Value*
The packet size	64 bytes
The # packages which are send	40
The initial # packages not used in the measurements	2

Table 1: *Roundtrip parameters used in all examples.*

Throughput Parameter	*Value*
The test duration	10 s
The send size	32768 bytes
The local socket send / receive buffer sizes	32768 bytes
The remote socket send / receive buffer sizes	32768 bytes

Table 2: *Throughput parameters used in all examples.*

- A regular user account at all participating hosts without special privileges.
- The test hosts should allow the user at the control host to run remote shells without the need to specify a password.
- The user at the control host should be allowed to use the `crontab` command.

3 Examples

In this section some examples are given of net performance measurements which are currently running. First the used measurement parameters are described.

3.1 Measurement parameters

As has already been mentioned in section 1: it is our intention to get an impression of the network performance from a users point of view. Therefore, there are quite ordinary and default values used for the parameters of the roundtrip and throughput measurements presented in the sequel of this section. The parameters are the same for all data shown. The used roundtrip (throughput) parameters are specified in table 1 (2).

Title	*Full Name*	*Location*
	English Home Page	
VU	Vrije Universiteit Amsterdam	Amsterdam
	`http://www.vu.nl/English/`	
UvA	Universiteit van Amsterdam	Amsterdam
	`http://www.uva.nl/english/index.html`	
RUL	Universiteit Leiden	Leiden
	`http://www.leidenuniv.nl/index_e.html`	
TUD	Delft University of Technology	Delft
	`http://www.tudelft.nl/matrix/home.en.cfm`	

Table 3: *Dutch universities where the DAS clusters are located. Please note that the <u>Title</u> columns refer to the titles used in the net performance Web pages and the <u>Full Name</u> columns refer to the university names used in the English home pages.*

3.2 DAS Cluster

3.2.1 Setup

The DAS cluster[4] is a super computer formed by a large number of Intel processors running Free BSD or Linux and locally connected by Myrinet, a popular Gigabit LAN. There are DAS Clusters operational at the Dutch universities which are specified in table 3.

The DAS clusters are inter-connected by an ATM network. Often there also exists a route via the "usual" Internet. The connection to the network is served by gateway hosts which are running Free BSD or RedHat Linux. There was a need in the DAS community to monitor the network performance a.o. to compare the ATM connection to the "usual" Internet route.

This is realized by installing the test host scripts of the rTPL package at the gateways of the DAS clusters at each of the four universities. A Digital Unix host from the "Department of Computational Physics" at the "Faculty of Physics & Astronomy", Utrecht University, is used as the control host. The HTTPD server of the faculty is used for the Web[5] presentation. This could be done without extra actions, because the directory with results is also accessible from the host where the HTTPD server resides.

The performance measurements were executed via the ATM network and also via the "usual" Internet route. However, not all gateways did

[4]See `http://www.cs.vu.nl/das/` for a general description of the DAS cluster.

[5]See `http://www.phys.uu.nl/~blom/net_perf/das/Table` for the DAS net performance monitor.

Site	*# Ping's*	*# Tput's*	*Time*	*Active*
VU	1	5	90 s	5.0 %
UvA	1	4	80 s	4.4 %
RUL	1	2	60 s	3.3 %
TUD	0	0	0 s	0.0 %
Total	3	11	230 s	12.8 %

Table 4: *Specify the time [s] each site is busy with the network performance tests. Also the percentage active has been given. Note that the measurements are repeated each half an hour.*

have an Ethernet interface which was accessible from the WAN. Therefore, some tests across the regular Internet had to be skipped. Skipped tests are marked in the result tables with "---". The tests were performed only in one direction. The reason is that the network performance in the ATM network is expected to be rather symmetric.

3.2.2 Measurement times

At the DAS Cluster tests in one direction are performed, however, both an ATM and a "usual" Internet route is used. This implies that for n sites there are $n(n-1)$ roundtrip measurements, the same amount of throughput measurements and off course n machine load calculations. Due to the one-directional tests the load per site is asymmetric: site k ($k = 1, \ldots, n$) has to do $(n-k)$ roundtrip / throughput measurements.

The DAS Cluster consists of four sites, however, the tests across the VU – UvA connection via the "usual" Internet is skipped, because there is no route available. This implies that there are in total 11 roundtrip and 11 throughput measurements to be done. At each site the roundtrip measurements to the other sites are performed simultaneously, while the throughput measurements are performed serially. Using the parameters from subsection 3.1, this leads for all sites to a total roundtrip measurement time of 120 seconds and a total throughput measurement time of 110 seconds, so this gives a total measurement time of 230 seconds. The tests are performed each half an hour. There is never more than one site at a time executing the tests, so in 12.8 % of the time a site from the cluster is active with the tests.

Table 4 also shows the measurement time differentiated for each site.

3.2.3 Results

The results of December 15, 1999 are shown here by including some of the Web available HTML tables. Please note that the selected day is typical: the results do not vary much during time, but off course there is a difference between the weekends and the working days. We do not show here the long term statistics which are also available (page 4; subsection 2.2), because there are not much short time fluctuations in the network here. Figure 2 (figure 3) shows the HTML table of the the minimum (average) roundtrip values [ms] between the VU site and the other sites. Figure 4 shows the throughput values [Mbit/s] for those connections. The tables are fragments of the corresponding Web pages. In both tables the background of the table cell reflects the value of the roundtrip or throughput value: black for the minimum and white for the maximum. These pseudo gray values may help to find tendencies more easily. Please note that normally pseudo colors are used, but for the sake of reproduction we limited ourself to gray values.

From the figures 2, 3 and 4 the following observations can be made:

- The ATM minimum roundtrip values give a reasonable reflection of the geographical distance. The reason therefore, is that all ATM routes require one hop.

- There is not much difference in minimum roundtrip values between the ATM connections (columns entitled "*ATM*") and the "usual" Internet route (columns entitled "*Inet*"), however, the differences and variations are larger for the average roundtrip values, especially during daytime. This suggests some congestion at the "Inet" route.

- The ATM throughput values are constant and close to the expected maximum of about 5 Mbit/s which follows a.o. from the # reserved cells. We do not know why the throughput of the VU – RUL connection is somewhat lower. The DAS sites are described in table 3.

- The variations in the "usual" Internet throughput values, reflect the load of the Internet during the day. But the ATM throughput values remain rather constant during the day. This probably is an indication that the capacity of the ATM connections is not yet fully used. These observations are typical. Note that for distributed computing short roundtrip times and constant bandwidth are important for message passing.

3.3 Bandwidth IPP Jülich – FOM Nieuwegein

3.3.1 Setup

It is the intension to control plasma physics experiments, situated at the IPP institute in Jülich, near Achen, Germany, by the FOM institute in Nieuwegein, near Utrecht, the Netherlands (see below for a short description of the sites). Therefore, it has to be investigated if their is sufficient network performance available between these two sites. The route between the sites is provided by the SURFnet, TEN-155 and DFN networks.

Date	Time	UVA		RUL		TUD	
		ATM	Inet	ATM	Inet	ATM	Inet
15/12/1999	23:00:05	1.200	---	4.400	4.400	5.000	5.300
15/12/1999	22:00:03	1.200	---	4.400	4.400	5.000	5.200
15/12/1999	21:00:04	1.200	---	4.400	4.400	5.000	5.700
15/12/1999	20:00:05	1.200	---	4.400	4.400	5.000	5.600
15/12/1999	19:00:03	1.200	---	4.400	4.400	5.000	5.800
15/12/1999	18:00:06	1.200	---	4.400	4.400	5.000	5.300
15/12/1999	17:00:07	1.200	---	4.400	4.400	5.000	5.600
15/12/1999	16:00:04	1.200	---	4.400	4.400	5.000	6.400
15/12/1999	15:00:05	1.200	---	4.400	4.400	5.000	5.700
15/12/1999	14:00:05	1.200	---	4.400	4.400	5.000	6.200
15/12/1999	13:00:04	1.200	---	4.400	4.400	5.000	5.600
15/12/1999	12:00:04	1.200	---	4.400	4.400	5.000	5.900
15/12/1999	11:00:05	1.200	---	4.400	4.400	5.000	5.400
15/12/1999	10:00:05	1.200	---	4.400	4.400	5.000	5.400
15/12/1999	09:00:04	1.200	---	4.400	4.400	5.000	5.200
15/12/1999	08:00:03	1.200	---	4.400	4.300	5.000	5.400
15/12/1999	07:00:03	1.200	---	4.400	4.400	5.000	5.400
15/12/1999	06:00:04	1.200	---	4.400	4.400	5.000	5.300
15/12/1999	05:00:04	1.200	---	4.400	4.400	5.000	5.200
15/12/1999	04:00:05	1.200	---	4.400	4.400	5.000	5.500
15/12/1999	03:00:05	1.200	---	4.400	4.400	5.000	5.400
15/12/1999	02:00:05	1.200	---	4.400	4.400	5.000	5.400
15/12/1999	01:00:05	1.200	---	4.400	4.400	5.000	5.300
15/12/1999	00:00:05	1.200	---	4.400	4.400	5.000	5.500

Figure 2: *The HTML table with the minimum roundtrip values [ms] between the VU site (table 3) and the other sites. It is a fragment of the corresponding Web page. The background of the roundtrip table cell is a reflection of its intensity.*

Date	Time	UVA		RUL		TUD	
		ATM	Inet	ATM	Inet	ATM	Inet
15/12/1999	23:00:05	1.234	---	4.468	4.924	5.092	7.679
15/12/1999	22:00:03	1.200	---	4.476	6.534	5.129	9.043
15/12/1999	21:00:04	1.224	---	4.508	15.866	5.258	18.868
15/12/1999	20:00:05	1.234	---	4.450	8.432	5.103	10.155
15/12/1999	19:00:03	1.221	---	4.421	7.447	5.095	11.419
15/12/1999	18:00:06	1.255	---	4.408	7.776	5.216	11.942
15/12/1999	17:00:07	1.203	---	4.458	7.955	5.097	13.208
15/12/1999	16:00:04	1.205	---	4.453	10.189	5.208	13.319
15/12/1999	15:00:05	1.226	---	4.411	9.158	5.239	14.981
15/12/1999	14:00:05	1.200	---	4.489	6.758	5.105	12.695
15/12/1999	13:00:04	1.226	---	4.421	6.095	5.129	11.068
15/12/1999	12:00:04	1.205	---	4.403	5.139	5.105	9.586
15/12/1999	11:00:05	1.203	---	4.408	7.021	5.095	10.529
15/12/1999	10:00:05	1.203	---	4.418	4.987	5.095	7.900
15/12/1999	09:00:04	1.239	---	4.442	4.863	5.095	7.174
15/12/1999	08:00:03	1.213	---	4.453	8.303	5.092	9.045
15/12/1999	07:00:03	1.200	---	4.426	4.866	5.095	6.208
15/12/1999	06:00:04	1.216	---	4.437	4.621	5.174	7.447
15/12/1999	05:00:04	1.200	---	4.432	4.505	5.155	7.263
15/12/1999	04:00:05	1.200	---	4.468	5.776	5.579	8.974
15/12/1999	03:00:05	1.237	---	4.416	4.939	5.397	7.368
15/12/1999	02:00:05	1.213	---	4.408	5.792	5.537	8.613
15/12/1999	01:00:05	1.203	---	4.400	6.232	5.405	8.195
15/12/1999	00:00:05	1.213	---	4.413	5.376	5.111	7.695

Figure 3: *The HTML table with the average roundtrip values [ms] between the VU site (table 3) and the other sites. It is a fragment of the corresponding Web page. The background of the roundtrip table cell is a reflection of its intensity. Here the background colors are logarithmic scaled; linear is the default.*

Date	*Time*	UVA		RUL		TUD	
		ATM	Inet	ATM	Inet	ATM	Inet
15/12/1999	23:00:05	4.88	---	4.01	5.12	4.91	4.11
15/12/1999	22:00:03	4.88	---	4.01	4.67	4.91	4.87
15/12/1999	21:00:04	4.88	---	4.02	5.03	4.91	3.32
15/12/1999	20:00:05	4.88	---	4.02	2.18	4.92	3.4
15/12/1999	19:00:03	4.88	---	4.02	4.12	4.91	3.08
15/12/1999	18:00:06	4.88	---	4.02	3.99	4.91	3.17
15/12/1999	17:00:07	4.88	---	4.02	2.65	4.91	1.3
15/12/1999	16:00:04	4.88	---	4.02	2.8	4.9	2.01
15/12/1999	15:00:05	4.87	---	4.02	3.85	4.86	.14
15/12/1999	14:00:05	4.88	---	4.02	4.74	4.86	1.62
15/12/1999	13:00:04	4.87	---	4.02	3.44	4.9	1.28
15/12/1999	12:00:04	4.88	---	4.02	4.79	4.91	1.57
15/12/1999	11:00:05	4.88	---	4.02	4.08	4.9	2.59
15/12/1999	10:00:05	4.88	---	4.02	2.65	4.9	6.04
15/12/1999	09:00:04	4.88	---	4.02	4.44	4.9	6.61
15/12/1999	08:00:03	4.88	---	4.02	5.6	4.91	6.75
15/12/1999	07:00:03	4.88	---	4.02	4.9	4.91	7.53
15/12/1999	06:00:04	4.88	---	4.02	5.36	4.91	4.67
15/12/1999	05:00:04	4.87	---	4.02	5.43	4.91	4.98
15/12/1999	04:00:05	4.88	---	3.98	5.78	4.91	5.79
15/12/1999	03:00:05	4.88	---	3.98	4.77	4.9	4.46
15/12/1999	02:00:05	4.87	---	3.98	4.19	4.91	4.6
15/12/1999	01:00:05	4.88	---	3.98	4.37	4.91	4.21
15/12/1999	00:00:05	4.88	---	3.98	4.05	4.91	6.15

Figure 4: *The HTML table with the throughput values [Mbit/s] between the VU site (table 3) and the other sites. It is a fragment of the corresponding Web page. The background of the throughput table cell is a reflection of its intensity.*

To be able to differentiate between the congestion sources in the involved networks, the following test setup with the sites, specified in table 5, has been configured. The host at the FOM site (table 5) is used as control host. However, this host does not possess a HTTPD server. Therefore, the results are copied to the IPP host for the Web[6] presentation. The possibility to copy the results to another host is a feature of the package. In contradiction to the DAS monitor (subsection 3.2) here are no extra Quality-of-Service test connections used, but on the other hand the measurements are performed in two directions. This leads to a somewhat different presentation of the data in the tables.

3.3.2 Measurement times

In this configuration there are tests in two directions between each site pair. This implies that in contradiction to subsection 3.2.2 the situation per site is symmetric here. When there are n sites, for each site there are $(n-1)$ roundtrip / throughput measurements and off course one machine load measurement, so for all sites there are $n(n-1)$ roundtrip / throughput measurements and n machine load calculations to be done.

In this setup there are nine sites participating. No tests are performed at the connections RUS – ZAM and RUS – ZELAS, so for all sites there are maximal 28 roundtrip / throughput measurements and 7 machine load calculations to be executed. Also per site there are (maximal) 6 roundtrip / throughput measurements to be done. All roundtrip measurements are performed simultaneously. This implies that each site is (maximal) 100 seconds busy with the measurements. The tests are repeated each hour, therefore, the average load for each site is 2.8 %.

3.3.3 Results

As in the DAS results (subsection 3.2.3) December 15, 1999 is selected as a typical day, despite the fact that the data are more varying during longer time frames, caused by recent tuning of the TEN-155 network. Therefore, the long term statistics (page 4; subsection 2.2) are not yet valid here. However, the data of this day are still typical for a working day.

In the following paragraphs some results will be presented 1) in the form of HTML tables, 2) as plots derived from the Applet plot window. Both are fragments of the net performance Web pages. In these paragraphs the sites are labeled with the titles specified in table 5. For reproduction reasons all results are here also presented as gray values.

Figure 5 presents the HTML table of the minimum roundtrip values [ms] at December 15, 1999 between the UU-36 site and correspondingly figure 6

[6]See `http://ipp277.ipp.kfa-juelich.de/~blom/net_perf/ipp/Table` for the IPP – FOM network monitor.

Title	*Full Name*	*Location*
	(English) Home Page	*Rate [Mbit/s]*
	Site Description	
IPP	Institut für Plasmaphysik	Jülich
	`http://www.fz-juelich.de/ipp/`	10
	This site, situated at the FZJ (Forschungszentrum Jülich) contains the plasma physics experiments to control.	
ZELAS	Central Electronics Laboratory	Jülich
	`http://www.fz-juelich.de/zel/zel_e.htm`	10
	This site, also situated at the FZJ, has been add to make a comparison possible with the host at the IPP and to be able to detect net congestion of the FZJ's LAN.	
ZAM	Central Institute for Applied Mathematics	Jülich
	`http://www.fz-juelich.de/zam/`	100
	The host at this site is near the router at the FZJ. It has been add 1) to measure the influence of congestion at the LAN; 2) it has a 100 Mbit/s interface.	
RUS	Rechenzentrum der Universität Stuttgart	Stuttgart
	`http://www.uni-stuttgart.de/Rus/`	10
	This site has been add to detect possible congestion at the DFN network.	
SARA	Academic Computing Services Amsterdam	Amsterdam
	`http://www.sara.nl/`	100
	The host at this site is located near the SURFnet backbone to detect congestion at this network.	
FOM	FOM-Institute for Plasma Physics Rijnhuizen	Nieuwegein
	`http://www.rijnh.nl/`	10
	This site is the location where the plasma physics control room should be situated.	
UU-36	Institute of Computational Physics	Utrecht
	`http://www.phys.uu.nl/~wwwfi/`	100
	This site has been add to detect congestion at the SURFnet backbone and for the reason and for the reason of its 100 Mbit/s interface.	

Table 5: *Sites participating in the network performance tests for network route IPP – FOM. The* Title *columns are again referring to the titles used in the Web pages.*

shows the percentage roundtrip packets lost. Figure 7 shows the table of the throughput values [Mbit/s] at that day for those connections. Note that in the headers of these tables the title “>> *Site*” (“<< *Site*”) stands for the connection: UU-36 to (from) the site entitled *Site*.

To make the tendency in the throughput data displayed in figure 7 more clear, figure 8 displays the plot of the throughput data between the UU-36 site and the other sites with a 100 Mbit/s interface (SARA and ZAM) at December 15, 1999 and the day before.

In the following two figures the results of the 100 Mbit/s interfaces are compared with the 10 Mbit/s interfaces. Figure 9 compares the performance differences between 100 – 100 and 100 – 10 connections by displaying the throughput data between the ZAM site (100 Mbit/s) and the UU-36 (100 Mbit/s) and FOM (10 Mbit/s) sites. Reversely, figure 10 compares 10 – 10 with 10 – 100 connections by presenting the throughput data between the IPP site (10 Mbit/s) and the FOM (10 Mbit/s) and UU-36 (100 Mbit/s) sites.

From figure 5, displaying the minimum ping values from / to site UU-36 (see table 5) the following can be concluded:

- There is no significant difference in roundtrip times during the 24 hours the values are shown. This appears to be typical.

- Most roundtrip times are symmetric in both directions. However, the difference in roundtrip times for the connections UU-36 >> FOM and UU-36 << FOM is remarkable. A possible explanation may be a different route for both connections.

- The roundtrip time to / from the RUS site in Stuttgart is shorter than the roundtrip time to the ZAM and IPP sites in Jülich. Off course this is explained by the fact that the TEN-155 ATM link is between Amsterdam and Frankfurt. The DFN route between Frankfurt and Jülich goes along Köln, so this route probably is longer.

From figure 7, showing the throughput values for the connections to / from UU-36 for all sites, and from figure 8, displaying the plots from the throughput results for the connections to / from UU-36 for the sites with a 100 Mbit/s interface, there follows:

- During day time there is a considerable performance degradation for the connections from UU-36 to the German sites when the throughput values during the night are compared with the values obtained during the day.

- There is no performance lost during daytime for the Dutch connections between UU-36 and FOM.

Date	*Time*	>> IPP	<< IPP	>> ZAM	<< ZAM	>> RUS	<< RUS	>> SARA	<< SARA	>> FOM	<< FOM
15/12/1999	23:00:12	18.100	17.550	16.800	17.000	14.800	15.000	2.190	2.100	2.310	4.165
15/12/1999	22:00:12	17.900	17.550	17.400	16.000	14.100	15.000	2.170	2.200	2.420	4.245
15/12/1999	21:00:11	18.600	17.586	16.700	17.000	14.500	15.000	2.200	2.300	2.420	4.196
15/12/1999	20:00:08	18.800	17.586	17.400	17.000	14.900	15.000	2.280	2.300	2.420	4.966
15/12/1999	19:00:16	18.800	17.550	17.100	17.000	14.800	15.000	2.180	2.300	2.470	4.328
15/12/1999	18:00:08	18.800	18.525	17.400	18.000	15.000	15.000	2.220	2.300	2.530	4.475
15/12/1999	17:00:08	19.200	17.586	19.100	17.000	15.400	16.000	2.370	2.300	2.410	4.489
15/12/1999	16:00:07	19.100	18.525	17.600	18.000	15.700	16.000	2.270	2.200	2.480	4.699
15/12/1999	15:00:08	19.700	18.525	17.300	17.000	15.400	17.000	2.350	2.200	2.490	4.786
15/12/1999	14:00:07	18.800	17.550	17.800	17.000	16.200	16.000	2.280	2.300	2.620	4.194
15/12/1999	13:00:09	19.100	18.525	17.500	18.000	16.300	17.000	2.350	2.400	2.540	4.107
15/12/1999	12:00:07	19.000	18.525	18.100	18.000	15.700	17.000	2.630	2.100	2.760	4.709
15/12/1999	11:00:08	18.700	18.525	17.600	17.000	16.300	16.000	2.380	2.200	2.500	4.314
15/12/1999	10:00:07	18.600	17.550	16.800	16.000	14.700	16.000	2.310	2.300	2.460	4.547
15/12/1999	09:00:07	18.600	17.550	16.500	16.000	14.600	15.000	2.170	2.200	2.390	4.425
15/12/1999	08:00:07	17.600	16.609	16.300	16.000	14.400	14.000	2.260	2.300	2.470	3.163
15/12/1999	07:00:08	17.800	17.550	16.000	16.000	14.100	14.000	2.230	2.200	2.660	4.547
15/12/1999	06:00:08	17.700	16.575	16.100	16.000	13.700	14.000	2.250	2.300	2.680	4.597
15/12/1999	05:00:07	17.600	17.550	15.900	16.000	13.700	14.000	2.180	2.300	2.540	4.210
15/12/1999	04:00:08	17.600	17.550	16.500	16.000	13.900	14.000	2.240	2.200	2.440	4.780
15/12/1999	03:00:09	17.700	17.550	16.800	16.000	13.800	14.000	2.200	2.100	2.330	4.246
15/12/1999	02:00:08	17.900	17.550	16.500	16.000	14.100	14.000	2.270	2.200	2.540	4.671
15/12/1999	01:00:07	18.100	16.575	16.500	16.000	14.400	14.000	2.290	2.200	2.520	4.357
15/12/1999	00:00:07	18.400	17.550	17.200	17.000	14.100	15.000	2.220	2.200	2.540	4.060

Figure 5: *The HTML table with the minimum roundtrip values [ms] between the UU-36 site and the other sites. The ZELAS site is not listed here, because at that moment there were no data available.*

Date	Time	>> IPP	<< IPP	>> ZAM	<< ZAM	>> RUS	<< RUS	>> SARA	<< SARA	>> FOM	<< FOM
15/12/1999	23:00:12	5.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
15/12/1999	22:00:12	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
15/12/1999	21:00:11	0.000	0.000	0.000	0.000	0.000	0.000	0.000	2.500	0.000	0.000
15/12/1999	20:00:08	2.500	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
15/12/1999	19:00:16	5.000	0.000	2.500	0.000	0.000	0.000	0.000	0.000	0.000	0.000
15/12/1999	18:00:08	2.500	0.000	0.000	0.000	0.000	2.500	0.000	0.000	0.000	0.000
15/12/1999	17:00:08	2.500	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
15/12/1999	16:00:07	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
15/12/1999	15:00:08	0.000	0.000	0.000	2.500	0.000	0.000	0.000	0.000	0.000	0.000
15/12/1999	14:00:07	0.000	2.500	5.000	0.000	2.500	2.500	0.000	0.000	0.000	0.000
15/12/1999	13:00:09	0.000	0.000	0.000	2.500	0.000	2.500	0.000	0.000	0.000	0.000
15/12/1999	12:00:07	0.000	7.500	5.000	7.500	2.500	2.500	0.000	2.500	5.000	0.000
15/12/1999	11:00:08	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
15/12/1999	10:00:07	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
15/12/1999	09:00:07	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
15/12/1999	08:00:07	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
15/12/1999	07:00:08	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
15/12/1999	06:00:08	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
15/12/1999	05:00:07	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
15/12/1999	04:00:08	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
15/12/1999	03:00:09	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
15/12/1999	02:00:08	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
15/12/1999	01:00:07	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
15/12/1999	00:00:07	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Figure 6: *The HTML table with the percentage roundtrip packets lost between the UU-36 site and the other sites. The ZELAS site is not listed here, because at that moment there were no data available.*

Date	*Time*	>> IPP	<< IPP	>> ZAM	<< ZAM	>> RUS	<< RUS	>> SARA	<< SARA	>> FOM	<< FOM
15/12/1999	23:00:12	3.18	3.83	3.24	8.33	3.9	7.34	15.51	20.17	5.96	6.14
15/12/1999	22:00:12	2.74	2.61	2.9	6.31	4.31	7.62	15.6	21.56	5.98	4.74
15/12/1999	21:00:11	3.38	3.35	3.23	5.37	3.94	7.89	12.69	14.7	6.03	6.36
15/12/1999	20:00:08	3.2	2.04	3.39	4.48	4.65	8.1	15.18	18.47	5.61	5.53
15/12/1999	19:00:16	2.01	2.37	.18	5.12	3.96	8.07	14.87	21.78	5.98	6.42
15/12/1999	18:00:08	1.36	1.56	2.3	4.48	3.95	7.84	14.72	17.54	5.94	5.53
15/12/1999	17:00:08	2.36	2.66	1.97	4.69	2.98	7.62	14.35	16.08	5.48	6.37
15/12/1999	16:00:07	2.88	2.25	3.86	5.34	2.52	6.72	13.37	5.26	5.51	5.9
15/12/1999	15:00:08	2	1.05	2.39	3.38	4.44	4.89	12.67	4.44	6.08	5.73
15/12/1999	14:00:07	1.15	2.08	1.28	3.59	4.33	3.78	13.19	6.3	5.68	5.92
15/12/1999	13:00:09	.1	1.99	1.73	2.95	2.94	2.27	12.37	6.62	6.17	4.58
15/12/1999	12:00:07	2.54	1.32	2.48	1.55	3.03	4.22	2.53	2.82	5.14	3.23
15/12/1999	11:00:08	2.03	3.93	3.87	7.94	3.86	7.95	13.23	9.27	5.66	6.27
15/12/1999	10:00:07	3.53	4.65	4.37	8.09	3.45	7.83	14.46	14.35	5.53	5.76
15/12/1999	09:00:07	4.23	6.66	8.69	11.35	3.81	8.07	15.32	20.1	6.29	6.29
15/12/1999	08:00:07	4.24	5.07	10.83	12.15	5.04	8.07	15.73	23.74	6.33	5.19
15/12/1999	07:00:08	4.44	7.95	13.56	12.06	5.18	8.08	16.25	23.74	6.21	6.6
15/12/1999	06:00:08	4.46	6.35	13.79	12.51	5.18	8.12	16.62	22.94	6.02	6.7
15/12/1999	05:00:07	4.18	6.44	12.34	11.76	4.87	7.89	14.36	23.89	5.7	5.04
15/12/1999	04:00:08	4.46	5.79	11.88	12.19	5.16	8.12	17.02	23.39	5.81	6.42
15/12/1999	03:00:09	4.36	2.79	12.24	12.06	5.15	8.12	15.18	21.91	6.39	6.32
15/12/1999	02:00:08	4.45	8.07	9.58	12.13	5.76	8.11	15.23	22.16	6.15	6.23
15/12/1999	01:00:07	4.16	2.15	8.38	10.6	5.07	8.11	15.21	22.65	6.08	6.41
15/12/1999	00:00:07	2.86	2.26	5.2	7.59	4.34	8.05	15.61	21.44	6.26	6.4

Figure 7: *The HTML table with the throughput values [Mbit/s] between the UU-36 site and the other sites. The ZELAS site is not listed here, because at that moment there were no data available.*

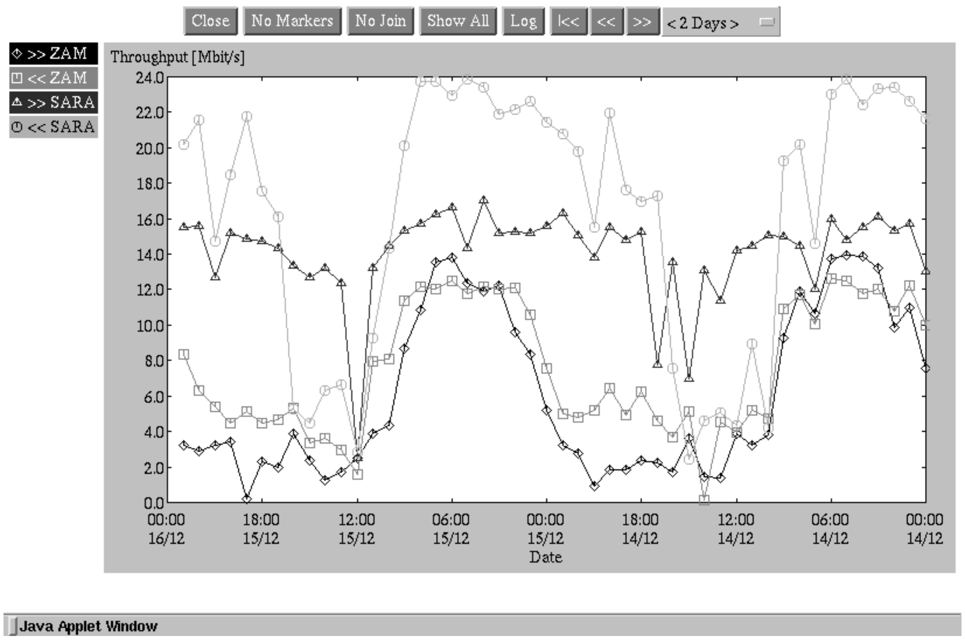

Figure 8: *Plot of the throughput data between the UU-36 site and the other sites with a 100 Mbit/s interface (SARA and ZAM) at December 15, 1999 and the day before. See table 5 for a description of the sites.*

- There is a bit performance lost for the connection UU-36 >> SARA, however, there is also considerable lost for SARA >> UU-36. This asymmetry is typical. The reason herefore is not clear, but the relatively limited processing power of the SARA site may have some influence.

- Also during daytime there is an asymmetry in the connections between UU-36 and SARA. This may be caused by the fact that there are more active processes at the UU-36 site, while there are few processes active at the SARA site. The processes at UU-36 are a.o. using the network. This implies that the maximum available throughput at UU-36 for one process is somewhat limited. Note that the performance Web pages showed a load of one for UU-36 and zero load for SARA.

- The performance degradation for the UU-36 – SARA connections occurs in the time frame 10:00 – 19:00, so this may be caused mainly by local office hours. However, for the UU-36 – ZAM connections this time frame is 10:00 – 02:00. This may indicate that also traffic to / from the US does play a role here.

- Without sites near the backbone of the TEN-155 network it is difficult

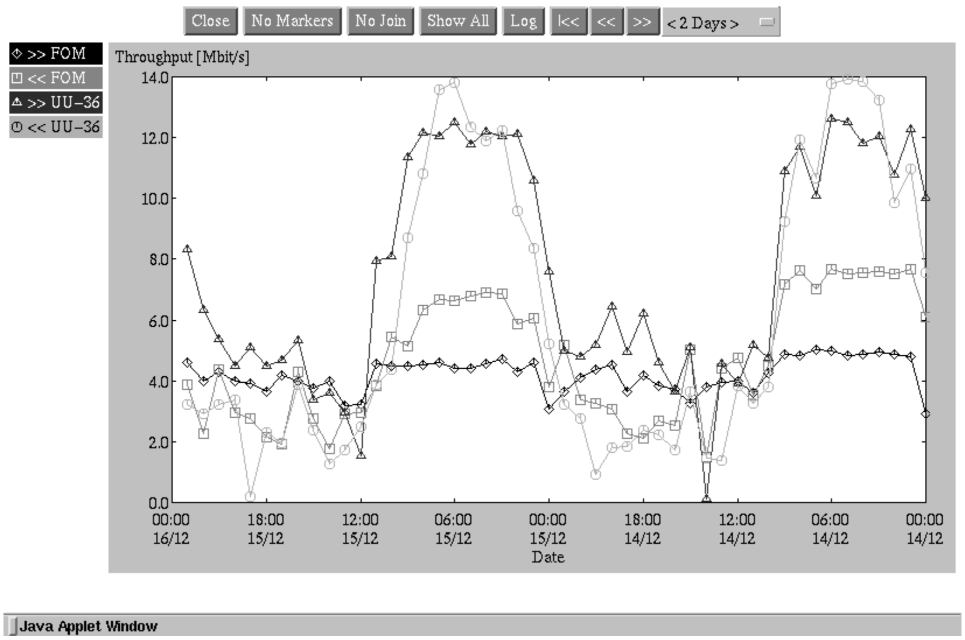

Figure 9: *Compare 100 – 100 with 100 – 10 connections by displaying the throughput data between the ZAM site (100 Mbit/s) and the UU-36 (100 Mbit/s) and FOM (10 Mbit/s) sites.*

to analyze at which part of the route the congestion did take place.

From figure 6, showing the percentages roundtrip packets lost for the connection from / to UU-36, the following can be concluded:

- The most roundtrip packets are lost at 12:00. This is in good agreement with the largest performance dip in the throughput values (figure 7).
- There # roundtrip packets lost from UU-36 to the German sites are larger than to the Dutch sites. This is not surprisingly because more routers are involved. However, a # packets lost larger than about 5 % resolution are quite considerable.
- There are no packets lost during the night. This suggests that the package lost is induced by congestion and not, for instance, by a bad configuration of a router.

From figure 9, which compares the performance difference between 100 – 100 and 100 – 10 connections by showing the throughput values between the ZAM site and the UU-36 and FOM sites, the following can be concluded:

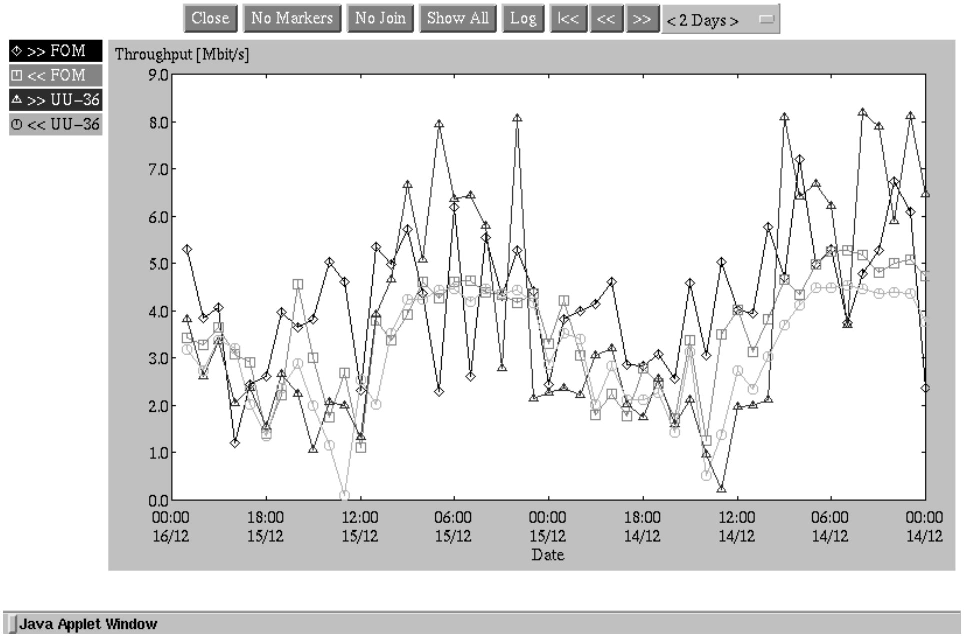

Figure 10: *Compare 10 – 10 with 10 - 100 connections by displaying the throughput data between the IPP site (10 Mbit/s) and the FOM (10 Mbit/s) and UU-36 (100 Mbit/s) sites.*

- During daytime there is not much performance difference between the connections: the connection ZAM >> UU-36 performs slightly better than ZAM >> FOM. In the other direction the difference in performance is less profound, but the ZAM << FOM connection performs a fraction better.

- During periods with congestion at daytime the performance of the connections from ZAM are better than in the reverse direction. During nighttime without congestion there is no significant difference in the connections with UU-36. However, in contradiction to the previous item the connection from FOM performs better than in the reverse direction.

From figure 10, where the performance difference between 10 – 10 and 10 – 100 connections are compared using the throughput values between the IPP site and the FOM and UU-36 sites, there can be concluded:

- There is no significant difference between the 10 – 10 and 10 – 100 connections.

- There are more fluctuations in the throughput values for the connections from IPP to both sites than in the reverse direction.

3.4 Conclusions

The examples discussed in the subsections 3.2 and 3.3 lead to the following general conclusions:

- The RTPL package gives a reasonable insight into the performance of a network from a user point of view without heavily processor and network loads. Long term statistics may help to average out fluctuations, at least when the network is stable.
- The Netperf package gives a reasonable direct insight into the available throughput of a network for the end user. Average roundtrip and roundtrip package lost information also give some insight into it, however, to make it valuable a more advanced evaluation of the results as a.o used in [1] is required. However, we have the impression that these methods are always less direct than `netperf` measurements. Therefore, at least for short term measurements, we prefer the Netperf package.
- At the IPP – FOM setup (subsection 3.3) a mixture is used of sites with 100 Mbit/s and 10 Mbit/s interfaces. In the sequel we prefer tests with only 100 Mbit/s interfaces, because they give more clear information of performance dips by the results of their throughput measurements.
- In the IPP – FOM setup (subsection 3.3) we have the impression that more direct information at the TEN-155 backbone would be valuable to make the interpretation of the results still more clear. We understand very well that workstations at the backbone have a limited capacity for extra load. Therefore, in subsection 4.2 we propose a setup which is intended to limit the load as much as possible.

4 Proposal Start Setup TEN-155

In this section we make a proposal for a start setup for a performance test with as subject the TEN-155 network. Because this test will interfere with the measurements at the IPP – FOM setup (subsection 3.3), these tests will be stopped during the TEN-155 test period.

In first instance only workstations at the NRN backbones (possible also off-backbone workstations) will be used. In the sequel also TEN-155 on-backbone workstations will be added.

4.1 Initial Setup

After some inquiries at the TEN-155 community, it appears that a good starting setup would be between DFN, SURFnet and JANET. One of the

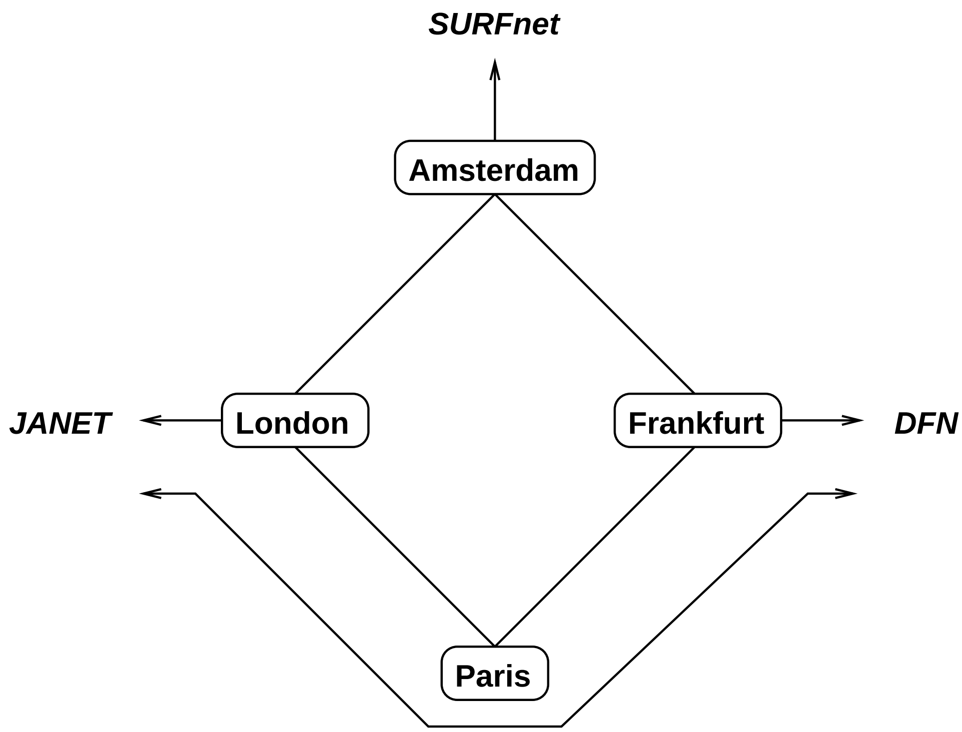

Figure 11: *Proposed start setup for net performance tests at the TEN-155 network. Note that the route between JANET and DFN goes via France.*

reasons is that these are the three points around the current ring UK – NL – DE – FR where most of the traffic flows. Traffic between JANET and DFN goes via France. This setup is drawn in figure 11.

So we would suggest to start the TEN-155 performance tests with three workstations at or near the backbones of the NRN's. Because the # sites is limited we would suggest to test the full matrix at an hourly period. This implies two roundtrip and two roundtrip measurements for each site. As explained in section 3, the roundtrip tests are executed simultaneously, so using the parameters of subsection 3.1 this means an engagement of each site with the package of 60 seconds. With the suggested test period of an hour, this is a relative load of 1.7 % which is rather moderate. Considering that only default socket parameters are used at the throughput measurements we do not expect congestion problems. The tests at the DAS Cluster (subsection 3.2) are running longer than a year and the tests at the IPP – FOM setup for about half a year without any problems.

		UK		NL		DE	
		NRN	TEN	NRN	TEN	NRN	TEN
UK	NRN	−	∗	∗	−	∗	−
	TEN	∗	−	−	∗	−	∗
NL	NRN	∗	−	−	∗	∗	−
	TEN	−	∗	∗	−	−	∗
DE	NRN	∗	−	∗	−	−	∗
	TEN	−	∗	−	∗	∗	−

Table 6: *Proposed test scheme were each NRN (TEN) site executes performance tests with the other NRN (TEN) sites and with the local TEN (NRN) site. Note that '∗' implies do a test and '−' implies skip the test.*

4.2 Limited use of hosts at the TEN-155 backbone

As stated before in subsection 3.3 the information of workstations at the backbone of the TEN-155 network would be very valuable. On the other hand we understand very well that the extra capacity of these workstations for such tests is limited. Therefore, we would suggest a topology where each:

- Each NRN workstation is tested together with the TEN-155 workstation in the same country and with all other NRN workstation.
- Each TEN-155 workstation is tested together with the NRN workstation in the same country and with the TEN-155 workstations to which there are routes configured. Please note that in the setup, described in subsection 4.1, this implies that each UK, NL and DE workstation is tested with all others.

This setup is visualized in table 6.

With the scheme of table 6 each NRN or TEN workstation is tested together with three others. This means for each workstation a measurement time of 70 seconds. With periodic tests at each hour this is 1.9 % of each host. This is in our opinion still quite reasonable.

References

[1] R. L. A. Cottrel, C. A. Logg, and D. E. Martin, "What is the Internet doing? Performance and reliability monitoring for the HEP Community," *Computer Physics Communications*, vol. 110, pp. 142–148, May 1998.

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

[3] 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.

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