

Poznan Supercomputing and Networking Center

Network measurement tools tests. Part II

Document version: draft

Prepared by:

Michal Przybylski
Szymon Trocha

July 2001

1. Contents

1.	Contents	2
2.	Objective	4
3.	Tools tested.....	5
4.	Devices used in the testbed.....	6
5.	Test topology	6
6.	The tests outcome	7
6.1	<i>CLink tests</i>	7
6.1.1	Command line	7
6.1.2	Measured bandwidth vs. configuration parameters and network load.	7
a)	empty network	7
b)	25% network load	7
c)	50% network load	8
d)	80% network load	8
e)	Induced maximum peak network load	8
6.1.3	Outcome	8
6.2	<i>Netperf tests</i>	9
TCP	9
6.2.1	Command line	9
6.2.2	Measured bandwidth vs. configuration parameters and network load.	9
a)	Empty network	9
b)	25% network load	9
c)	50% network load	10
d)	80% network load	10
e)	Induced maximum peak network load	10
6.2.3	Outcome	10
UDP	11
6.2.4	Command line	11
6.2.5	Measured bandwidth vs. configuration parameters and network load.	11
a)	Empty network	11
b)	25% network load	11
c)	50% network load	12
d)	80% network load	12
e)	Induced maximum peak network load	12
6.2.6	Outcome	12

6.3	<i>Pathchar tests</i>	13
6.3.1	Command line	13
6.3.2	Measured bandwidth vs. configuration parameters and network load.	13
	a) Empty network	13
	b) 25% network load	13
	c) 50% network load:.....	14
	d) 80% network load:.....	14
	e) Induced maximum peak network load	14
6.3.3	Outcome	14
6.4	<i>Pathrate tests</i>	15
6.4.1	Command line	15
6.4.2	Measured bandwidth vs. configuration parameters and network load.	15
	a) Empty network (25 measurement for each command line).....	15
	b) 25% network load (25 measurement for each command line):.....	15
	c) 50% network load (25 measurement for each command line):.....	15
	d) 80% network load (25 measurement for each command line):.....	16
6.4.3	Outcome	16
7.	Summary	17

2. Objective

The aim of the second part of the measurement tests was to work out the guidelines for the use of the previously tested tools. The first part of the tests [1] showed that some of the tools performed well in various network configuration, while others most probably required a special configuration to get acceptable measurements.

Another point was to check the influence of the existing network traffic on the tool accuracy. In this test the optimal configuration has been chosen for each tool and the tests have been performed in the presence of the different intensity network load (Ethernet frames generated by the network analyzer). All measurements have been observed and the tools parameter tuned, in order to achieve the most accurate measurement.

All the tests have been performed, with an assumption, that the user is not aware the network speed, configuration and topology (it means that the network itself is considered as a “black box” with measurement devices connected at the sides).

As some of the tools had problems with measurements on the switched Ethernet this phase involved use of a very simple network topology, consisting of just a switch and a hub.

3. Tools tested

Clink (tool: www.caida.org/Pathchar, doc. www.cs.colby.edu/~downey/pathchar)

(Single; measures the latency and bandwidth of internet links by sending UDP packets and measuring round-trip time – similar to ping and traceroute – just using more packets; determines the capacity of links)

Netperf () (tool: <http://www.netperf.org/netperf/NetperfPage.html>)

(Netperf is a benchmark that can be used to measure the performance of many different types of networking. It provides tests for both UDP and TCP throughput)

Pathchar (tool: www.caida.org/Pathchar, doc. www.cs.colby.edu/~downey/pathchar)

(Single; pathchar is a new tool, written by Van Jacobson at Lawrence Berkeley Laboratory (LBL), that tries to infer the characteristics of individual links along an Internet path by measuring the round trip time of packets sent from a single host. An alpha version of pathchar is available for FreeBSD, Linux, OSF and Solaris)

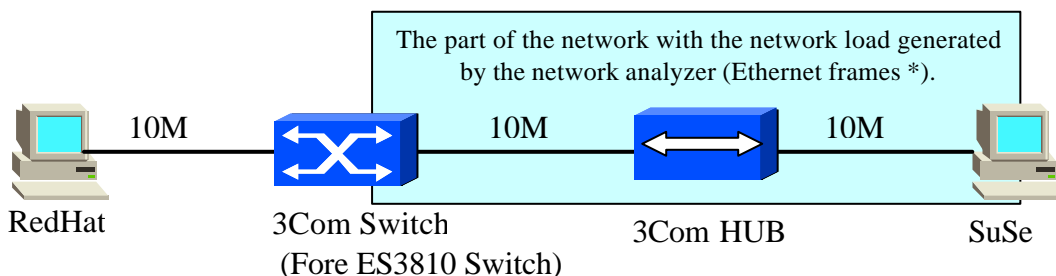
Pathrate (bw_meter) (tool: www.cis.udel.edu/~dovrolis/bw_meter.html)

(Client-Server; pathrate is a tool for measuring the capacity of a path from sender S to receiver R. The capacity of the path is the maximum IP-layer throughput that a flow can get from S to R.)

4. Devices used in the testbed

1. Linux Suse - 6x86 PR166+ CPU running at 133 MHz, 3Com 3C905B Fast Ethernet card (100 Mbit/s full duplex)
2. Linux RedHat Guinness – Pentium Celeron 700, 3Com 3C905B Fast Ethernet card (100 Mbit/s full duplex)
3. 3Com SuperStack II 2700 Ethernet switch
4. Fore ES3810 Fast Ethernet/Ethernet Switch
5. 3Com OfficeConnect 8/TPC Ethernet hub
6. Siemens K1102 CSMA-CD network analyser/traffic generator.

5. Test topology



* as the traffic generator had only an option to generate one size of frames at a time, we were unable to generate the network load with the given frame size distribution (i.e. simulating the real network load). The Ethernet frame size, used to load the network was equal to 1500 bytes.

All the test were performed on the RedHat machine, as [1] showed some influence of the slower processor on the measurement quality.

6. The tests outcome

6.1 CLink tests

6.1.1 Command line

1. `./clink -i10 -r1 -q4 -s96 10.10.0.2`
2. `./clink -i10 -r1 -q8 -s96 10.10.0.2`
3. `./clink -i10 -r1 -q4 -s48 10.10.0.2`
4. `./clink -i10 -r1 -q8 -s48 10.10.0.2`

6.1.2 Measured bandwidth vs. configuration parameters and network load.

- For the reasons described in the CLink documentation, best bandwidth estimate does not necessarily fall between low and high estimate.
- Measurement time depends on the run-time parameters chosen.
- The numbers in the first column of all tables indicate the command line used to obtain given measurements.
- The repetitiveness tests have been performed for configuration considered “the best”. This test consisted of ten runs of the tool with the same set of parameters. Then, the average measured capacity and its range (highest estimate – lowest estimate) has been calculated.

a) empty network

Tests repetitiveness

Repetitive tests (10 times) showed measurement range variation at the level of tens of kbit/s (less than 1% of the link bandwidth). The measured bandwidth capacity was about 72% of the real link speed.

Measured bandwidth vs. configuration parameters

No.	#probes	#sizes	low bw estimate [Mbit/s]	best bw estimate [Mbit/s]	high bw estimate [Mbit/s]	measurement time [min:ss]
1	4	16	7.262	7.240	7.315	1:19
2	8	16	7.267	7.251	7.275	2:23
3	4	31	7.203	7.246	7.245	2:19
4	8	31	7.258	7.269	7.291	4:23

* gray color indicates the configuration used for the repetitiveness tests.

b) 25% network load

Tests repetitiveness

Repetitive tests (10 times) showed measurement range variation at the level of max. 300 kbit/s (less than 3% of the link bandwidth). The measurement bandwidth capacity was close to that on the empty network - at about 72% of the real link speed.

Measured bandwidth vs. configuration parameters

No.	#probes	#sizes	low bw estimate [Mbit/s]	best bw estimate [Mbit/s]	high bw estimate [Mbit/s]	measurement time [min:ss]
1	4	16	6.256	7.089	7.948	0:58
2	8	16	7.173	7.254	7.251	2:10
3	4	31	6.138	6.940	6.578	2:03
4	8	31	7.176	7.230	7.202	4:10

* gray color indicates the configuration used for the repetitiveness tests

c) 50% network load

Tests repetitiveness

Repetitive tests (10 times) showed measurement range variation at the level of 700 kbit/s (7% of the link bandwidth). The bandwidth estimate was much lower than in the empty network – reporting about 50% of the real link speed.

Measured bandwidth vs. configuration parameters

No.	#probes	#sizes	low bw estimate [Mbit/s]	best bw estimate [Mbit/s]	high bw estimate [Mbit/s]	measurement time [min:ss]
1	4	16	4.967	5.978	5.474	0:59
2	8	16	4.022	4.578	4.619	2:09
3	4	31	4.790	5.790	6.691	2:03
4	8	31	4.687	5.216	4.793	4:08

* gray color indicates the configuration used for the repetitiveness tests

d) 80% network load

Tests repetitiveness

Repetitive tests (10 times) showed measurement range variation at the level of more than 4000 kbit/s (more than 40% of the link bandwidth). The bandwidth estimate was very inaccurate – reporting from 55% to 97% of the real link speed.

Measured bandwidth vs. configuration parameters

No.	#probes	#sizes	low bw estimate [Mbit/s]	best bw estimate [Mbit/s]	high bw estimate [Mbit/s]	measurement time [min:ss]
1	4	16	4.115	5.448	5.270	1:19
2	8	16	4.078	4.364	5.435	2:14
3	4	31	4.627	4.104	4.680	2:03
4	8	31	4.421	4.385	5.094	4:07

* gray color indicates the configuration used for the repetitiveness tests

e) Induced maximum peak network load

The CLink does not produce a significant network load. The maximum induced network load was less than 60kbit/s (0.6 %).

6.1.3 Outcome

The bandwidth estimation provided by the tool is not very accurate, but the tool stability for over-provisioned network (with the network load < 10%-20%) is very high. With the increasing network load the tool stability decreases, as well as the accuracy.

CLink was not very accurate measuring fast links (it was reporting less than 50% of the link bandwidth – [1]).

The parameters used for the measurement seems not to have any special influence on the output, as long as they cover sufficient number of probes and sizes.

6.2 Netperf tests

TCP

6.2.1 Command line

1. `./netperf -H 10.10.0.2 -l 60 -n 1 -t TCP_STREAM -f k -- -m 250 -s 32768 -S 32768`
2. `./netperf -H 10.10.0.2 -l 60 -n 1 -t TCP_STREAM -f k -- -m 500 -s 32768 -S 32768`
3. `./netperf -H 10.10.0.2 -l 60 -n 1 -t TCP_STREAM -f k -- -m 1000 -s 32768 -S 32768`
4. `./netperf -H 10.10.0.2 -l 60 -n 1 -t TCP_STREAM -f k -- -m 1500 -s 32768 -S 32768`

6.2.2 Measured bandwidth vs. configuration parameters and network load.

- The numbers in the first column of all tables indicate the command line used to obtain given measurements.
- The repetitiveness tests have been performed for configuration considered “the best”. This test consisted of ten runs of the tool with the same set of parameters. Then, the average measured capacity and its range (highest estimate – lowest estimate) has been calculated.

a) Empty network

Tests repetitiveness

Repetitive tests (10 times) showed measurement range variation at the level of approximately 60 kbit/s (less than 1% of the link bandwidth). The average bandwidth estimate showed about 57% of the real link speed.

Measured bandwidth vs. configuration parameters

No.	msg. size	bandwidth estimate [Mbit/s]
1	250	5.022
2	500	5.586
3	1000	5.715
4	1500	6.153

* gray color indicates the configuration used for the repetitiveness tests

b) 25% network load

Tests repetitiveness

Repetitive tests (10 times) showed measurement range variation at the level of approximately 60 kbit/s (less than 1% of the link bandwidth). The average bandwidth estimate showed about 54% of the real link speed.

Measured bandwidth vs. configuration parameters

No.	msg. size	bandwidth estimate [Mbit/s]
1	250	4.801
2	500	5.394
3	1000	5.357
4	1500	5.848

* gray color indicates the configuration used for the repetitiveness tests

c) 50% network load

Tests repetitiveness

Repetitive tests (10 times) showed measurement range variation at the level of approximately 100 kbit/s (1% of the link bandwidth). The average bandwidth estimate showed about 53% of the real link speed.

Measured bandwidth vs. configuration parameters

No.	msg. size	bandwidth estimate [Mbit/s]
1	250	4.544
2	500	5.270
3	1000	5.251
4	1500	5.806

* gray color indicates the configuration used for the repetitiveness tests

d) 80% network load

Tests repetitiveness

Repetitive tests (10 times) showed measurement range variation at the level of approximately 100 kbit/s (1% of the link bandwidth). The average bandwidth estimate showed about 51% of the real link speed.

Measured bandwidth vs. configuration parameters

No.	msg. size	bandwidth estimate [Mbit/s]
1	250	4.548
2	500	5.143
3	1000	5.125
4	1500	5.615

* gray color indicates the configuration used for the repetitiveness tests

e) Induced maximum peak network load

The Netperf TCP tests produce a heavy network load. The maximum induced network load was up to 84 % of the link bandwidth

6.2.3 Outcome

The Netperf TCP stream test is a very stable tool – the measurement range variation is less than 1% of the link capacity (for 10Mbit/s link). Unfortunately, the bandwidth estimate falls much below the real link speed.

This problem does not exist in the point-to-point or a switched network connection, with no background traffic, TCP tests performed very good, giving very accurate measurement [1].

The potential use of the Netperf implies generation of a heavy network load, which may influence other flows.

What is interesting, the results seem to be independent of the network load.

UDP

6.2.4 Command line

1. `./netperf -H 10.10.0.2 -l 60 -n 1 -t UDP_STREAM -f k -- -m 250 -s 32768 -S 32768`
2. `./netperf -H 10.10.0.2 -l 60 -n 1 -t UDP_STREAM -f k -- -m 500 -s 32768 -S 32768`
3. `./netperf -H 10.10.0.2 -l 60 -n 1 -t UDP_STREAM -f k -- -m 1000 -s 32768 -S 32768`
4. `./netperf -H 10.10.0.2 -l 60 -n 1 -t UDP_STREAM -f k -- -m 1500 -s 32768 -S 32768`

6.2.5 Measured bandwidth vs. configuration parameters and network load.

- The numbers in the first column of all tables indicate the command line used to obtain given measurements.
- The repetitiveness tests have been performed for configuration considered “the best”. This test consisted of ten runs of the tool with the same set of parameters. Then, the average measured capacity and its range (highest estimate – lowest estimate) has been calculated.

a) Empty network

Tests repetitiveness

Repetitive tests (10 times) showed measurement range variation at the level of single kbit/s (less than 1%). The average bandwidth estimate showed about 89% of the real link speed.

Measured bandwidth vs. configuration parameters

No.	msg size	bandwidth estimate [Mbit/s]
1	250	7.389
2	500	8.190
3	1000	8.939
4	1500	8.786

* gray color indicates the configuration used for the repetitiveness tests

b) 25% network load

Tests repetitiveness

Repetitive tests (10 times) showed measurement range variation at the level of single kbit/s (less than 1%). The average bandwidth estimate showed about 88% of the real link speed.

Measured bandwidth vs. configuration parameters

No.	msg size	bandwidth estimate [Mbit/s]
1	250	7.258
2	500	8.139
3	1000	8.788
4	1500	8.672

* gray color indicates the configuration used for the repetitiveness tests

c) 50% network load

Tests repetitiveness

Repetitive tests (10 times) showed measurement range variation at the level of single kbit/s (less than 1%). The average bandwidth estimate showed about 87% of the real link speed.

Measured bandwidth vs. configuration parameters

No.	msg size	bandwidth estimate [Mbit/s]
1	250	7.241
2	500	8.128
3	1000	8.722
4	1500	8.657

* gray color indicates the configuration used for the repetitiveness tests

d) 80% network load

Tests repetitiveness

Repetitive tests (10 times) showed measurement range variation at the level of single kbit/s (less than 1%). The average bandwidth estimate showed about 87% of the real link speed.

Measured bandwidth vs. configuration parameters

No.	msg size	bandwidth estimate [Mbit/s]
1	250	7.246
2	500	8.124
3	1000	8.747
4	1500	8.652

* gray color indicates the configuration used for the repetitiveness tests

e) Induced maximum peak network load

The Netperf UDP tests produce very heavy network load. The maximum induced network load was up to 100 % of the link bandwidth.

6.2.6 Outcome

The Netperf UDP stream test is a very stable tool – the measurement range variation was less than 1% of the link capacity (for 10Mbit/s link). The bandwidth estimate usually showed about 85%-90% of the original link speed, regardless of the network load.

6.3 Pathchar tests

6.3.1 Command line

1. `./pathchar 10.10.0.2 -n -q 4 -Q +16 -m 1500`
2. `./pathchar 10.10.0.2 -n -q 4 -Q +32 -m 1500`
3. `./pathchar 10.10.0.2 -n -q 8 -Q +16 -m 1500`
4. `./pathchar 10.10.0.2 -n -q 8 -Q +32 -m 1500`

6.3.2 Measured bandwidth vs. configuration parameters and network load.

- The numbers in the table represent just a single sample measurement with given run-time parameters. In some tests (especially with the heavy network load) the average bandwidth estimate measured in the repetitiveness tests may be different.
- The measurement time depends on the run-time parameters set – the more probes & sizes, the longer measurement time.
- The numbers in the first column of all tables indicate the command line used to obtain given measurements.
- The repetitiveness tests have been performed for configuration considered “the best”. This test consisted of ten runs of the tool with the same set of parameters. Then, the average measured capacity and its range (highest estimate – lowest estimate) has been calculated.

a) Empty network

Tests repetitiveness

Repetitive tests (10 times) showed measurement range variation at the level of 300 kilobits per second (3% - in the interval between 9.3 and 9.5 Mbit/s). The average bandwidth estimate showed about 94% of the real link speed.

Measured bandwidth vs. configuration parameters

No.	#probes	#sizes	bandwidth estimate	measurement time [mm:ss]
1	4	16	9.5	00:51
2	4	32	9.4	01:54
3	8	16	9.5	01:51
4	8	32	9.4	03:43

* gray color indicates the configuration used for the repetitiveness tests

b) 25% network load

Tests repetitiveness

Repetitive tests (10 times) showed measurement range variation at the level of 600 kilobits per second (6% - in the interval between 9.1 and 9.6 Mbit/s). The average bandwidth estimate showed about 95% of the real link speed.

Measured bandwidth vs. configuration parameters

No.	#probes	#sizes	bandwidth estimate	measurement time [mm:ss]
1	4	16	9.6	00:51
2	4	32	9.6	01:53
3	8	16	9.5	01:51
4	8	32	9.4	03:45

* gray color indicates the configuration used for the repetitiveness tests

c) 50% network load:

Tests repetitiveness

Repetitive tests (10 times) showed measurement range variation at the level of 2000 kilobits per second (20% - in the interval between 5.3 and 7.3 Mbit/s). The average bandwidth estimate showed about 65% of the real link speed.

Measured bandwidth vs. configuration parameters

No.	#probes	#sizes	bandwidth estimate	measurement time [mm:ss]
1	4	16	5.5	00:50
2	4	33	5.0	01:57
3	8	16	8.9	01:52
4	8	33	7.1	03:47

* gray color indicates the configuration used for the repetitiveness tests

d) 80% network load:

Tests repetitiveness

Repetitive tests (10 times) showed measurement range variation at the level of 3200 kilobits per second (32% - in the interval between 9.8 and 13 Mbit/s). The average bandwidth estimate showed about 110% of the real link speed.

Measured bandwidth vs. configuration parameters

No.	#probes	#sizes	bandwidth estimate	measurement time [mm:ss]
1	4	16	11.0	00:52
2	4	32	8.5	01:52
3	8	16	5.7	01:52
4	8	32	6.1	03:52

* gray color indicates the configuration used for the repetitiveness tests

e) Induced maximum peak network load

The Pathchar tests produce network load less than 1% of the link bandwidth.

6.3.3 Outcome

The Pathchar tool measures *available* bandwidth. The measurement is quite accurate on an empty network. But the tests performed in the presence of the background traffic show that (for the low network load < 25%) pathchar measures rather the maximum link bandwidth. With the higher network load the tool accuracy falls to the unacceptable level.

6.4 Pathrate tests

6.4.1 Command line

1. `./pathrate_rcv 10.10.0.2 100000 100`
2. `./pathrate_rcv 10.10.0.2 100000 250`
3. `./pathrate_rcv 10.10.0.2 100000 500`
4. `./pathrate_rcv 10.10.0.2 100000 1000`
5. `./pathrate_rcv 10.10.0.2 100000 2000`

6.4.2 Measured bandwidth vs. configuration parameters and network load.

- Bandwidth shown in the tables below is a concatenation of the measured intervals. Example – if a number of measurements with resolution set to 1Mbit/s reported bandwidth interval 8.9 – 9.9 Mbit/s and another measurements reported 9.2 – 10.2 Mbit/s, the interval shown in the table will be 8.9 – 9.2 Mbit/s.
- % of incorrect measurements shows the relative number of measurements reporting totally incorrect bandwidth for the link (100Mbit/s). For this case, some more detailed description has been provided in the summary.
- Test repetitiveness has been checked through multiple (25) tests with the same number of parameters.

a) Empty network (25 measurement for each command line)

No.	range [kbit/s]	resolution [kbit/s]	lower bw. interval boundary [Mbit/s]	upper bw. interval boundary [Mbit/s]	% of incorrect measurements
1	100000	100	10.4	10.6	8%
4	100000	250	10.4	10.7	56%
7	100000	500	9.4	10.0	48%
10	100000	1000	9.1	10.2	0%
13	100000	2000	8.6	11.0	4%

b) 25% network load (25 measurement for each command line):

No.	range [kbit/s]	resolution [kbit/s]	lower bw. interval boundary [Mbit/s]	upper bw. interval boundary [Mbit/s]	% of incorrect measurements
1	100000	100	9.7	10.4	36%
4	100000	250	9.5	10.6	28%
7	100000	500	9.3	10.5	36%
10	100000	1000	8.9	10.1	44%
13	100000	2000	8.1	10.6	12%

c) 50% network load (25 measurement for each command line):

No.	range [kbit/s]	resolution [kbit/s]	lower bw. interval boundary [Mbit/s]	upper bw. interval boundary [Mbit/s]	% of incorrect measurements
1	100000	100	9.7	9.8	12%
4	100000	250	9.5	9.8	32%
7	100000	500	9.3	9.8	16%
10	100000	1000	8.8	10.0	26%
13	100000	2000	7.9	10.6	8%

d) **80% network load (25 measurement for each command line):**

No.	range [kbit/s]	resolution [kbit/s]	lower bw. interval boundary [Mbit/s]	upper bw. interval boundary [Mbit/s]	% of incorrect measurements
1	100000	100	9.7	10.3	24%
4	100000	250	9.4	9.8	28%
7	100000	500	9.3	9.9	16%
10	100000	1000	8.6	10.2	32%
13	100000	2000	7.8	10.5	12%

Tests repetitiveness

The repetitiveness of all the correct measurements were quite high, usually indicating two or three bandwidth intervals close to each other (i.e. 8.9-9.9, 9.0-10.0 and 9.1-10.1 for 1000kbit/s resolution). Each time the measurement was done, the output had to be analyzed to eliminate obviously incorrect measurements (more information in summary)

Induced maximum peak network load

Pathrate generates the peak load of maximum 29% of the link bandwidth (average peak load – 15%-20%, depending on the run-time parameters).

6.4.3 Outcome

The Pathrate tool measures the path capacity. The tool's bandwidth estimate is very accurate and almost does not depend on the network load. The only problem with the tool is that the measurement sometimes indicates ***obviously*** unacceptable bandwidth. This may be eliminated by the experienced user (maybe even it is a software bug – this needs to be investigated).

6.5 Pathrate 2.2.1 tests

6.5.1 Command line

```
./pathrate_rcv
```

Pathrate 2.2.1 is an improved version of Pathrate, which does not need any run-time parameters.

6.5.2 Measured bandwidth vs. configuration parameters and network load.

Performed tests did not show any influence of the network load on the reported bandwidth range.

Network load	Real link bandwidth	Low bandwidth estimate [Mbit/s]	High bandwidth estimate [Mbit/s]	Measurement time [mm:ss]
0%	10	9.6	10.0	01:25
25%	10	9.4	9.8	27.00
80%	10	9.4	9.7	27.00
0%*	100	95	105	00:30

(*) measurement done on a single fast Ethernet switch in full duplex mode.

Tests repetitiveness

The tool output is very stable. The reported bandwidth values were concentrated in the interval $\pm 5\%$ of the link bandwidth.

6.5.3 Outcome

Pathrate 2.2.1 is a very stable and accurate tool for the bandwidth measurement. The tests showed little or no influence of the network load on measured bandwidth. Measurement resolution does not exceed 5% of the path bandwidth.

7. Summary

Testing and evaluating of the measurement tools is not an easy task. Many different elements may influence the measurement taken, as most of the tested tools use a statistical approach – generally by measuring the “round trip time”. Experiments showed, that even a small change in the network configuration (replacement of the 10Mbit/s 3Com 2700 switch with the 10/100 Mbit/s Fore ES-3810 switch) caused measurement difference reaching 10% - while the total link speed remained the same. Different types of network devices may differently influence the measurement. The same effect exist while using different machines – ie: different operating systems or network interfaces, processors, number of services etc. The potential user should be aware of these problems and should be able to perform the analysis of the measurement, to evaluate its quality.

To evaluate the tool, we need to know how the measurement is taken. Let’s try to shortly analyze all of the used tools.

1. CLink – this program measures the RTT of the generated UDP packets (of different sizes). Based on that measurement, the bandwidth estimate is given in the form of interval (minbw, maxbw) and the proposed “best_estimate”. In order to establish the given link bandwidth, packets are send with a specified TTL. This technique implies that the measurement of the link usually contains an error, introduced by the measurement of the previous link. This error distributes and aggregates on the subsequent links. It may be confusing with the large number of hops. CLink does not work well with the 100Mbit/s and faster networks – as previous tests proved [1] the measurement was very inaccurate.
2. Netperf – this program generates a UDP or TCP stream at full possible network speed. Both of these tests show the maximum path bandwidth that a given type flow can get. For the total link bandwidth measurement, the UDP flow seems to give better results under all conditions. The obtained measurement is stable and almost independent from the network load. The disadvantage of this method is the induced network load – UDP stream uses all of the link bandwidth, starving other flows.
3. Pathchar – works just like the traceroute, measuring round trip time for packets of different size. Regarding to the author, it “does a fairly good job... on links 10Mb/s or slower”. This has been confirmed during the test, where the tool performed quite well. The only problem arose with the network load of 50% and more – the tool accuracy was not good anymore. The author’s hint is to use more probes – but this causes longer measurement time (more than 180 sec.).
4. Pathrate – the measurement mechanism of that tool is quite complicated [7]. The tool requires two input parameters – bandwidth of the slower network interface card on sender or receiver and required measurement resolution (precision). If we know the interface card bandwidth – the measurement is more often very accurate. But if we do not know that value (say we have 10Mbit/s card) and assume the default parameter value (100Mbit/s), surprisingly we may get the bandwidth estimate ranging 100Mbit/s.
5. Pathrate 2.2.1 – the measurement mechanism is similar to that of the Pathrate, but greatly improved. The tool reported proper bandwidth values under different conditions.

Our experiments let us to investigate the importance of the run-time parameters as well as repetitiveness of the measurements.

Run-time parameters

CLink and Pathchar: the user should select parameters in such a way to get sufficient number of probes with many different packet sizes (we would suggest to have more than a 100 measurements, ie. 10 probes at 15 sizes = 150 measurements).

Generally – the more probes the better estimate, though after a certain level we cannot really observe any improvement.

For Netperf the best output was for the UDP test with message size close to the MTU size (1500 bytes). The measurement time (if more than 60s) does not really influence the output. To ensure proper measurement we would suggest setting this parameter to at least 60s.

The Pathrate requires two parameters – interface bandwidth and test resolution. The meaning of the bandwidth parameter has been shown above, and the test resolution is responsible for the test length and final precision. Usually increasing resolution causes the test to last longer and, in some conditions, increases number of incorrect measurements.

Pathrate 2.1.1 does not require any run-time parameters.

Accuracy

- CLink: the tool does not produce very accurate measurements – the reported bandwidth is approx. 72% of the real bandwidth for empty network and approx. 43% of the real bandwidth for the network with 80% load.
- Netperf: regardless from the network load, the tool reports bandwidth at the level of 87-89% of the real bandwidth (with the UDP test).
- Pathchar: (available bandwidth) for the low network load, the tool shows quite accurate measurement at the level of 95% of the real link speed. With the increasing network load the output is not acceptable.
- Pathrate: with the proper configuration the tool accuracy is very high. Careful output analysis is required.
- Pathrate 2.1.1: very high tool accuracy.

Repetitiveness

- CLink: acceptable with the network load lower than 50% of the link (measurement variation range < 7% of the link bandwidth)
- Netperf: measurement variation range < 1% of the link bandwidth, regardless the network load.
- Pathchar: good performance on the low network load (less 25%) – (measurement variation range < 3% of the link bandwidth)
- Pathrate: if the bandwidth parameter is set properly - very good performance. Careful output analysis for interfaces with unknown speed required.

- Pathrate 2.1.1: very high measurement repetitiveness

Recommendations

In the authors opinion, taking into consideration recommended measurement granularity [2] the Netperf and Pathrate 2.2.1 are the best tools to measure a network path bandwidth (bottleneck bandwidth). Both of them have some disadvantages (like a heavy network load caused by Netperf and the measurement time of the Pathrate) but their measurements are quite accurate and repeatable under most conditions.

If the use of over provisioned network (OvNet) is taken into consideration and for slow links, Pathchar may be used as well.

References

1. "Network measurement tools test", M. Przybylski, Sz. Trocha, Poznan Supercomputing and Networking Center, May 2001
http://qos.man.poznan.pl/files/measurement_full.pdf
2. QoS monitoring and SLS auditing, V.Reijs
3. Tools for measuring the SLS metric, V. Reijs, July 12, 2002.
http://www.heanet.ie/Heanet/projects/nat_infrastruct/nettools.html
4. Clink documentation: www.cs.colby.edu/~downey/pathchar
5. Netperf documentation: <http://www.netperf.org/netperf/NetperfPage.html>
6. Pathchar documentation: www.cs.colby.edu/~downey/pathchar
7. Pathrate (*bw_meter*) documentation: www.cis.udel.edu/~dovrolis/bwmeter.html