

# On the Impact of IP Option Processing

Mattia Rossi University of Innsbruck (AT)  
Institute of Computer Science  
Csac3227@uibk.ac.at

Michael Welzl University of Innsbruck (AT)  
Institute of Computer Science  
michael.welzl@uibk.ac.at

## ABSTRACT

Oftentimes during special ping requests or protocol tests routers take much longer to process ip packets with set ip options than to process packets without options. This research was conducted in order to find out the average processing time discrepancy between those two types of ip packets and thus to come to a conclusion about the usability of ip packets with options. Even though the packets with options tend to be processed slower than the packets without options, this study shows that this delay is not as significant as sporadic measurements might indicate.

## 1. INTRODUCTION

This measurement shows the speed with which routers process packets with ip options. Even though ip options offer many possibilities to enhance packet routing or send special informations to hosts, they are hardly used because RTTs of those packets are too slow. This development might be due to results from sporadic measurements where the RTTs of packets with options are twice or even three times as long as RTTs of packets without options.

Using the *extended ping* program (further called *extping*) written by Gerhard Wahl, we pinged **5759** hosts with a wide ip address range in 2002 and pinged **27689** widely ranged hosts again in 2003. As a result we obtained 200 measured RTTs (100 per ping type). Since some hosts did not answer, the data vary. While only **4401** answered the second measurement (2003), the first measurement had a much better yield: **4427** out of **5759** hosts gave an answer. The smaller value in 2002 results from further measurements in the beta stage of *extping*, where all non-answering hosts were cancelled from the complete host list after a quick test run. This correction does not seriously affect the results, because the first measurement took place immediately after the completion of the program. *Extping* generates a result file which stores the the average, min, max, median RTTs

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.  
Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

and the variance of every path. The records are written once for packets without options (further called *NONE*) and once for packets with ip options (further called *NOP* because the “no operation” option is set). The NOP option was chosen because it is the option that needs the lowest additional processing resources. Based on the result file the program generates a summary with a straight average and median of all paths. However, this result is not exact because all paths have the same weight and some important values like the variance or the path length are ignored. We therefore wrote a statistical tool for better evaluation of the result file, giving each path a specific weight and subsequently applying the statistical formula for the weighted average. With this tool we obtained more exact values for the average RTTs per path length and host.

## 2. HOW THE PROGRAMS WORK

### 2.1 The *extping* program

In what follows, we briefly summarize the original documentation for *extping* which is enclosed with the software. The program consists of two parts: the *extended ping* and the *summary*. The aim of the program is to record the process time in routers for packets with ip options set. *extended ping* merges two known networking tools: *ping* and *traceroute*, while the *summary*-tool summarizes the gathered data.

*extended ping* works as follows: It sends two packets, continuously alternating between *NONEs* and *NOPs*, in the form of a ping to the host. The amount of pings may be set via the parameter *-p* when calling the program. If no number is given, it will automatically default to 10 pings; i.e. 10 *NONEs* and 10 *NOPs* are sent. A host is pinged every second. A sequence number identifies the answer packet for every ping. Before and after each ping, a traceroute to the host checks the path length and records the addresses of intermediate hosts (routers). At the end of every record a summary provides the main information about the pinged host. This summary is written to the result file, which will be explained later in this document.

The program was written in Linux and therefore needs a unix environment (it was also tested under Solaris) to compile and run as well as root rights because it creates a RAW-socket. It has to be run from a command line using some options:

|         |                 |   |
|---------|-----------------|---|
| -h      | help            | shows the available options                       |
| -n x    | number of hosts | sets the amount of hosts in the host list to ping |
| -p x    | pings per host  | sets the amount of pings per type and host        |
| -i file | in file         | the file containing the host list                 |
| -o file | out file        | the name of the outputfile                        |

An example of a command line call:

```
ext_ping -i hosts.dat -o results.dat -p 100
```

The results in this case are written to the file *results.dat* and 100 pings per type and host were done, i.e. 200 pings per host. The program creates 5 files at the end, but only 2 are really needed for further statistical processing. The *rtt\_log.dat* simply records every single ping and traceroute as described before. While it is the biggest output file, it is not really needed. The *optret.dat* records all answer packets which have come with an option set. The generation of this file was implemented only for debugging purposes. The *error.dat* records all encountered errors, such as non-answering hosts or ping or traceroute timeouts as well as the number of the host in the host file, the host name, the function where the error emerges, and the error string itself. From this file, it is possible to gather the amount of hosts which were not available for NOP pings. The timeouts are sorted between NOP and NONE timeouts. There are also other errors which may interrupt the ping sequence, but they are mostly irrelevant in the context of this report; NOP timeouts are relevant, as explained at the end of this document. The number of timeouts can easily be found with a simple shell command:

```
more error.dat | grep -c NOP
```

The two most important files are the result file and the *router.dat* file. The first one contains the summary for every host. The information is written in one line — some sort of a table is created. The following table shows the structure:

| column    | description  |
|-----------|--|
| NR        | Number of host in the host list                                    |
| HOST      | Name of the host in the host list                                  |
| CANONAME  | Canonic hostname   |
| IP        | Real ip adress of the host   |
| HOSTS     | Hops from source to destination (path length)                      |
| SUM(NONE) | Sum of the RTTs of all NONE pings                                  |
| SEQ       | Amount of NONE pigs  |
| MIN RTT   | Minimal RTT with NONE pings  |
| AVG RTT   | Average RTT with NONE pings  |
| MAX RTT   | Maximal RTT with NONE pings  |
| MEDIAN    | Median of the RTTs using NONE pings                                |
| VARIANCE  | Variance of the RTTs using NONE pings                              |
| R         | Contains an X if the reply to the NONE ping contained an ip option |
| SUM(NOP)  | Sum of the RTTs of all NOP pings                                   |
| SEQ       | Amount of NOP pings  |
| MIN RTT   | Minimal RTT with NOP pings   |
| AVG RTT   | Average RTT with NOP pings   |
| MAX RTT   | Maximal RTT with NOP pings   |
| MEDIAN    | Median of the RTTs using NOP pings                                 |
| VARIANCE  | Variance of the RTTs using NOP pings                               |
| R         | Contains an X if the reply to the NOP ping contained an ip option  |

The router file contains the paths to the hosts. The first column (as always) shows the number of the host in the host list as well as the ip adress and the path length. All hosts contained in the path are listed in the following columns.

The summary tool uses the result file to generate a summary. This tool is now obsolete because it does not take the different weights attributed to certain variables into account. Instead of the summary tool, a new statistic tool was used for the measurement.

The way the summary tool works explains its problems: the data collected in the result file are grouped according to the path length. These groups record the number of pinged hosts. This number is then split into two groups: NONE and NONE-A. The NONE column contains the summary for hosts that returned an answer with no option after a NONE ping. The other column shows the hosts that returned an answer with options. By adding the values of the two numbers, the whole amount results. In the same column the summary values are entered. There are 9 values: the top value and the value below show the absolute minimum and the absolute maximum RTT of the path length, followed by the average per-host minimum RTT. The next three values are the minimum, the average, and the maximum per-host average RTT. The third last value is the average per-host maximum RTT. The two last values are the average per-host median and the average per-host variance. The same is true for the NOP column. Finally, also the SUM group at the end of the summary file has the same structure, where the average per-host average RTT and the average per-host median RTT should be the significant values for the measurement. As explained above, however, no weight is given to path length or variance and the values are not correct.

## 2.2 The statistic program

The statistic program replaces the summary tool. It con-

tains three algorithms to weight the paths. The weights depend on the following values:

1. the variance
2. the path length
3. the frequency of occurrence of a certain router in the whole measurement

### 2.2.1 Why use different weights?

The results gathered with *extping* stem from the real Internet with all its problems, errors and congestions. Therefore we have to give more weight to the results that do not seem to be affected by congestions and thus do not falsify the measurement.

- The path length is the main risk factor for congestions: the longer a path, the higher the chance for a packet to experience congestion. Thus, the results with shorter routes get a higher weight.
- A high variance typically signifies problems encountered on a route: it shows a possible congestion on the route while pinging, a path change or other errors. We concluded that a result with a high variance should have less weight.
- If a path goes contains frequently used routers, the results may be falsified. If, for example, there are four paths containing four routers per path as shown below,

```
A → B → C → D
A → B → C → E
A → B → C → F
G → H → I → J
```

and routers A,B and C are extremely slow, the total per-router delay is increased substantially, even if all the other routers are extremely fast. Clearly, a router which occurs three times should not be three times as important as a single router. Thus, all routers get their own weight. The weight for each path results from the cumulative weight of all routers contained in the path.

### 2.2.2 Using the weight in the statistic tool

The statistic tool extracts data from the result file and the router file. The most relevant values of the result file are the variance and the path length. The first weight is calculated by inverting the square root of the variance. The square root has to be taken because of the high values the variances sometimes achieve (<350000). In this way we use the standard deviation instead of the variance, which gives us a more correct result. The weight for every host based on the standard deviation is calculated two times: once with the NOP variance and once with the NONES. This is important for the final result. In this way the weighted average can be found out correctly once for NOPS and once for NONES using different weights. All these weights are stored in a background array, which mirrors the host list in the result file. The number of replies with ip options set is a factor that does not directly relate to weighting but nevertheless plays an important role. Some hosts answered to NOP pings with packets with ip options set (not always NOP). Due to path asymmetry, the RTTs achieved by such pings are not

reliable because the packets could have been processed using the slow path in routers on their way back, and this path is not traceable. Therefore, all paths where hosts answered with ip options are ignored in the final result. For standard deviation weighting, the answerpackets with options are ignored as well. All values first are checked before being weighted — these checks reduce the program’s complexity. Then, the already known weights (once NONE, once NOP) of certain hosts are simply divided by their path length.

The next step is to compare the frequency of routers in each path. Due to the high number of comparisons, this procedure takes the most computing time and power. During this procedure, the router addresses are transferred from the router file into a matrix that mirrors the way they are written in the file. Then, the first router is compared it with all routers in the matrix. The number of similar routers found in the matrix is written to the corresponding field in another matrix, as well as in the fields where similar routers were found. In this way, taking the next router, the already checked routers are not going to be compared, and a lot of comparisons are saved. At the end a new matrix is created that contains values for every router and path. Based on the records of this matrix, the value for every path is calculated, all the values belonging to the path summed up, and the result inverted. The already entered weights (once NONE, once NOP) for every path are multiplied with the newly found value. Thus every path is weighted correctly, according to standard deviation, path length and the repetition of hosts. At the same time a counter is inserted to count the different hosts that were encountered by the packet.

The next step is to calculate the correct average per-host average RTTs and the correct average per-host median RTTs, with the statistical formula for the weighted arithmetic average [1]:

$$\bar{x} = \frac{\sum w_i * x_i}{\sum w_i}$$

This formula is applied once for every path length, and once for all hosts. It generates four outputs per path length and for the all-hosts result too; two (average and median) for the NOPS and two for the NONES. The really important values are the all-hosts results, four values describing the average RTT of an average host, processing pings with and without options.

## 3. THE MEASUREMENT

There were four measurements all in all: three immediately after the creation of *extping* in July / August 2002, and one in August / September 2003. The first two measurements were only test runs and are not really important, but still had a certain impact on the third measurement, and thus are to be considered. The first measurement took place during the beta stage of *extping*; it was a try with 10 pings per host and ping type.<sup>1</sup> The host list was taken from [3] and consists of **27689** hosts. The same list was used

<sup>1</sup>This test also included pings to discover how hosts would react in response to the historical “MTU” IP option described in RFC 1063. None of the “MTU”-pinged hosts replied.

in the 2nd and 4th measurement. The problem in the first measurements was that the NONE and NOP pings were not sent alternately but sequentially; i.e. first 10 NONE pings, then 10 NOP pings and so on. The problem in this case is that the NONE pings could be less or more affected by congestion than the NOP pings or vice versa. This behaviour was changed in the final version, as explained in the program section. Not all of the **27689** hosts were reachable: only **5475** hosts returned an answer. This number could be due to down times, firewalls which prohibit flooding, heavily congested routes and many other causes.

**Table 1: Measurement list**

|               |   |
|---------------|---|
| measurement 1 | test run 1, 10 pings sequentially (2002)                        |
| measurement 2 | test run 2, 10 pings alternating (2002)                         |
| measurement 3 | measurement run 1, 100 pings alternating, July/August 2002      |
| measurement 4 | measurement run 2, 100 pings alternating, August/September 2003 |

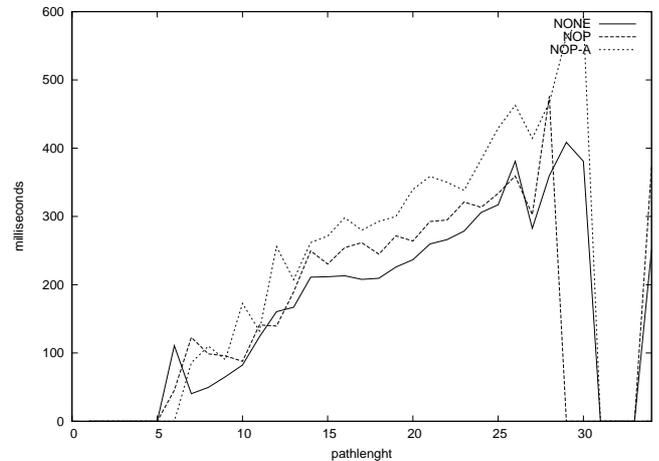
The second measurement used alternating pings. It was a test run too: as before, only 10 pings per host were sent. This time, **5399** hosts answered. For the third measurement using **100** pings per host, all hosts that did not answer during the first two measurements were removed from the host list. This was done because the first three measurements took place immediately after one another and required a substantial amount of time (the last one lasted much longer than the others); **4427** hosts answered. For the last measurement the initial host list was used again, with a yield of approximately **15%**, which translates into **4401** answering hosts.

This difference between the first / second and the fourth measurement may have different causes. Considering the immense growth of the Internet, there must have been great changes in a time span of a year. This means, that maybe many of the listed hosts might not exist any more, or they may be isolated, or at least unreachable from the testing point. Another cause could be the improved security on the net. Almost every user on the net now uses a firewall to close specific ports or to prevent flooding. Since our sequence of test pings could already be considered as a flood attack by some administrators, firewalls might forbid to complete the sequence. Unfinished sequences are the second most common errors after unreachable hosts.

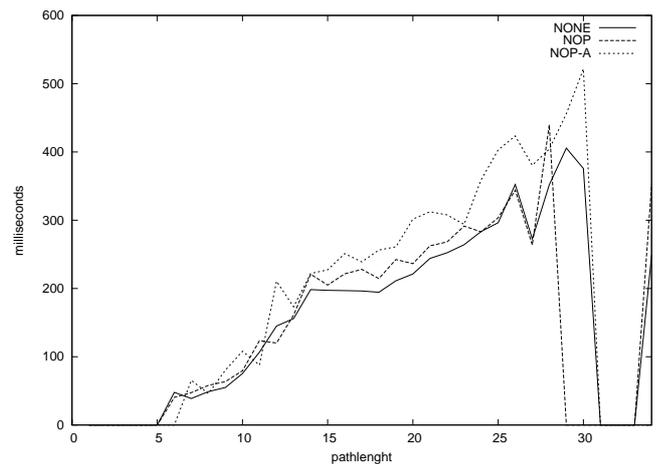
The traceroute could be a cause for breaking the sequence too. It works by sending UDP packets to the hosts to a random port, which should be unreachable in order to end tracerouting. If the port is open for use, on the other hand, the traceroute does not work, and all gathered data will be discarded, and the host logged in the *error.dat*. One other cause could be the still growing network traffic. The mass introduction of broadband connections everywhere in the world and many filesharing tools caused a higher use of the internet bandwidth but perhaps also increased congestion. Due to the introduction of easy to handle tools the amount of DDoS attacks reached its peak [2], and added to the congestion. In total, all of these factors can explain why we encountered lots of server downtimes and unreachable hosts.

### 3.1 Test runs

The main result of interest is not the number of responsive hosts, but the data we gathered from them. Figure 1 shows the average RTT per path length of the average NONE, NOP and NOP-A pings. The NOP-A are the pings that get an answer with ip options set.



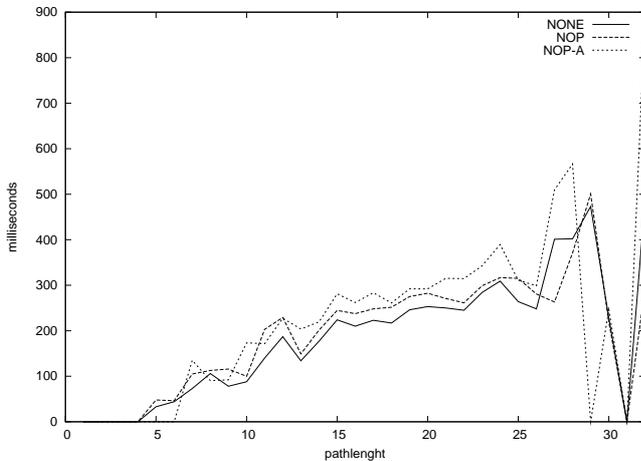
**Figure 1: The average per-host average results of the first measurement (10 pings sequential)**



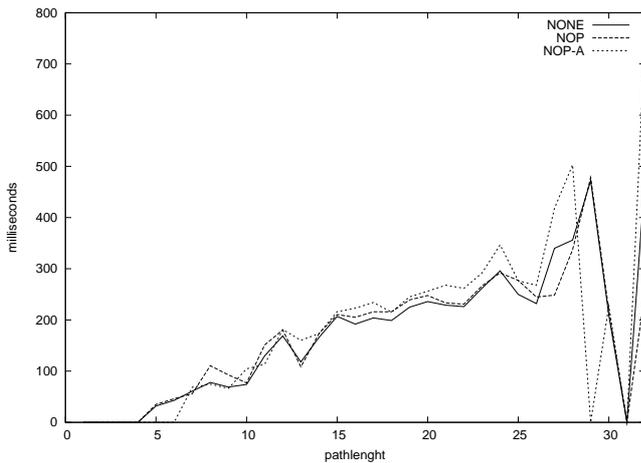
**Figure 2: The average per-host median results of the first measurement (10 pings sequential)**

Figure 2 already shows that the NOPs and the NOP-As are slower than the NONES, but only by about 50 ms on average. Compared to an average RTT of 200 ms of NONES, the RTT is only increased by a quarter. The NOP-As are twice as slow as the NOPs compared to the NONES, but there is great oscillation due to probable asymmetric paths, as explained earlier. The strange behaviour of the graph at path length 29 and higher is caused by only one host having such a long path. Therefore the splitting into NOP and NOP-A creates some visual confusion. The average per-host median shows more difference between NOP and NOP-A, and, in comparison, less difference between NOP and NONE. In this case, it is already imaginable that the final NOP results will not really differ from the NONES.

We have not statistically analyzed the result because 10 pings per host and the sequential processing do not fulfil the measurements' prerequisites. The second measurement that used alternating pings is not important for the final result either, but it may be interesting to have a look at the resulting graphs (figs. 3 and 4) for differences to the first measurement.



**Figure 3: The average per-host average results of the second measurement (10 pings alternating)**



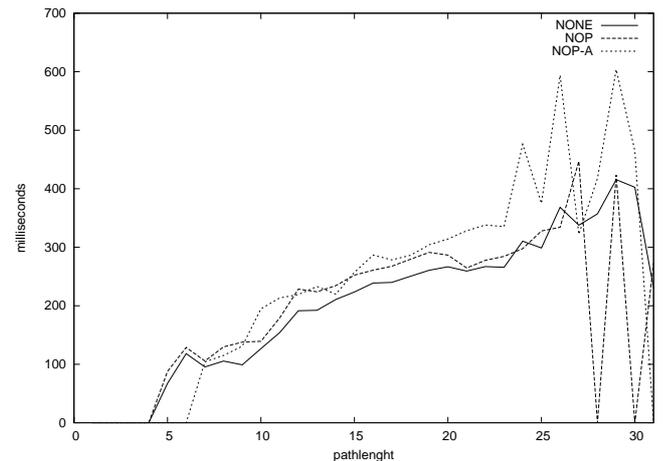
**Figure 4: The average per-host median results of the second measurement (10 pings alternating)**

As in the first measurement, the NOPs are slower than the NONES but faster than the NOP-As and the medians of NOPs and NONES are closer again. All in all, the alternating pings show less difference between the ping types, which appears to indicate that it was indeed a good idea to switch from sequential to alternating pings in our software.

### 3.2 Measurement runs

Let us now consider the really important results. The 100 ping run made in 2002 immediately after the first two runs, took place in Linz, Austria, and the 100 ping run made in summer 2003 took place in Innsbruck, Austria. In what follows, we will look at the average per-host average RTT,

average per-host median RTT and this time the average per-host standard deviation too.



**Figure 5: The average per-host average results of the third measurement (100 pings alternating)**

Figure 5 does not differ much from the ones shown earlier, but the values are more useful because of the  $2 \times 100$  pings per host. As in the other graphs, the last path lengths show some higher oscillations. This is caused again by a small number of hosts that have such long paths. Interestingly, in this graph, the NOPs are below the value of NONES slightly below and above path length 25. If we have a look at the standard deviations of NONE and NOP at these points (Fig. 7), we can see that the one of NONE is not much higher than the one of NOP. From this, we can conclude that it is not really probable that congestion was encountered — otherwise, the NOP pings would have a high standard deviation too. A path change is also improbable, because it would have affected the NOP as well. A possibility could be the switching between fast path and slow path. Most of the packets in the network are NONES that use the speed of fast path processing, and at the end they might “overload” the fast path, while the switched NOP packets might get every needed processing power without encountering such problems. So the fast path could perhaps become a “slow path” and vice versa — due to lack of knowledge in router architecture, we can only make wild guesses at this point.

The average per-host median graph again shows that the NOP lines are closer to the NONE than to the NOP-A. This happens because the median is less affected by outliers from congestion. As a result, we may say that the NOPs are more affected by anomalies than NONES. There are four occurrences where NOP values can be found below the NONE value (ignoring the two zero values at the high path lengths, which merely show that only NOP-A and NONE answers were received). Again, this means that NONES do not always have to be faster than NOPs; The high peaks of NONE this time show that most of the NONES were slower, but some few were really fast (this explains the low peaks in fig. 5), but were discarded by the median.

In fig. 7, there are four occurrences where the NOP values are below the NONES. Without any further knowledge of router architecture, we can only guess that there may sometimes be more fluctuations in the fast path than in the slow

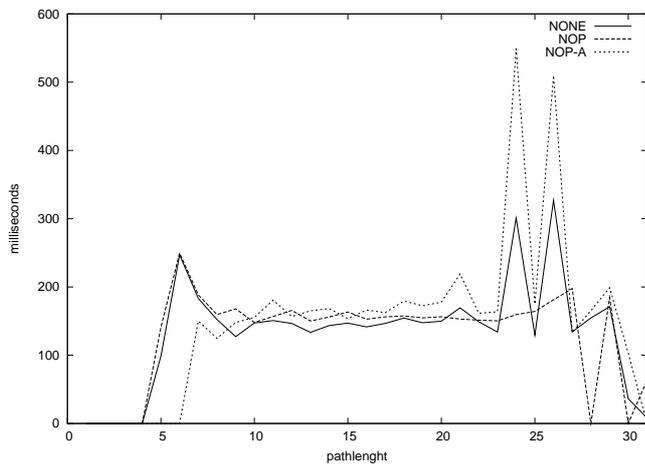


Figure 6: The average per-host median results of the third measurement (100 pings alternating)

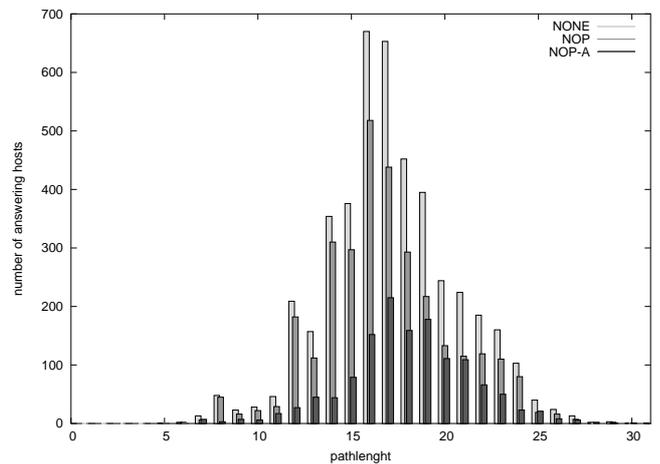


Figure 8: Answering hosts per path length (2002)

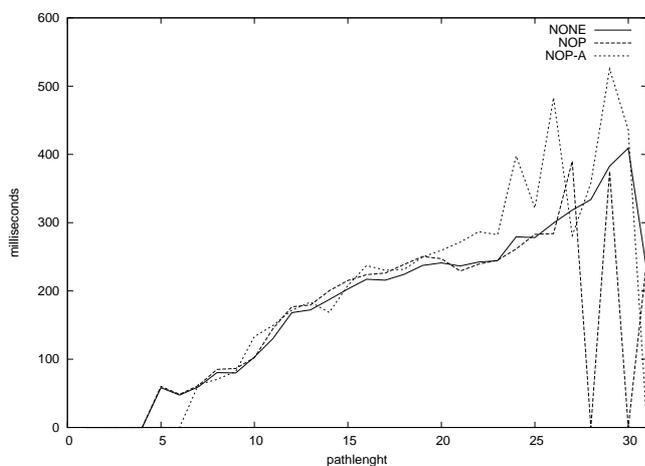


Figure 7: The average per-host standard deviation of the third measurement (100 pings alternating)

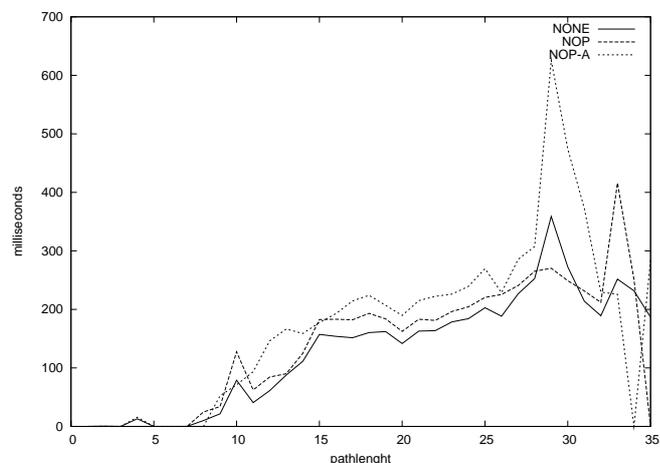


Figure 9: The average per-host average results of the latest measurement in 2003 (100 pings alternating)

path.

As we can see in Figure 8, there are less routers with a very high or a very low path length. The results obtained along very long paths often showed very high oscillations or strange values. These values are all filtered by the statistic tool mostly through the processing of the standard deviation. We considered it advisable to weight the number of answering hosts per path length as a negative factor even though long paths may seem to be more interesting because they involve more routers: the longer a path, the more disturbances of some kind can be expected.

The graphs of the latest measurement are very similar to all of the third measurement. The average per-host average RTT of every path length shows the NOPs again in the middle between NONEs and NOP-As.

While the average median of NOPs is nearer to the NONEs this time, the NOPs are sometimes lower than the NONEs again; this time however, there are less confusing standard deviation values than in the measurement before. In this measurement, the oscillation is really high — we probably did it at a time where our measurements encountered a lot

of cross traffic. The different behavior might also be the result of the different position (Innsbruck instead of Linz). However, the results show very similar behaviour.

If we compare the standard deviations from the 2002 measurement (fig. 7) with those of the 2003 measurement (fig. 11), we see that the NOPs are more or less at the same level, but the NONEs are much higher on average in 2002 than in 2003. This could mean that the first hop had different processing speeds. It might be possible that the first hop in the 2002 measurement did not differentiate between slow path and fast path processing, while the University of Innsbruck, where the measurement in 2003 took place, used different routers which switched between the fast path and slow path.

In figure 12 it can be seen again that there are less hosts with very high or very low path lengths. In this measurement there was a path length of 100, but some errors in the *extping* did not track this path length, and therefore it is not reflected in the graph. However, because of its high variance/standard deviation and high path length, it has no real importance for the measurement.

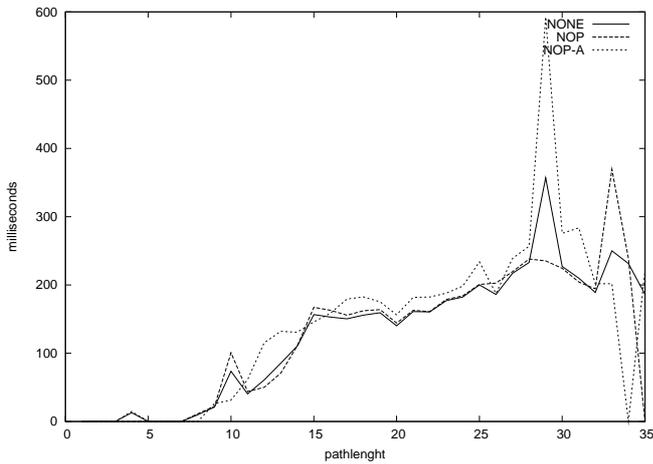


Figure 10: The average per-host median results of the latest measurement in 2003 (100 pings alternating)

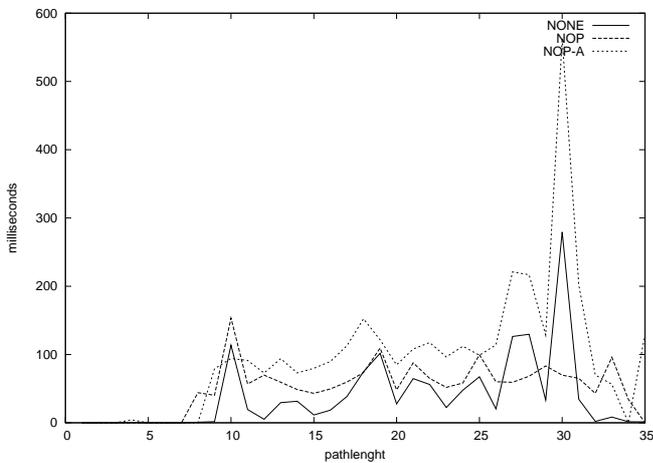


Figure 11: The average per-host standard deviation of the latest measurement in 2003 (100 pings alternating)

#### 4. STATISTICS

Now it is time to show the final results of the measurements. As described above, the statistic program generated two files: one for the measurement in 2002, and one for the latest one in 2003. The statistic result file is made up of a table where the weighted average per-host average and the weighted average per-host median RTTs of NONE and NOPs per path length can be found, and a complete weighted average per-host average and weighted average per-host median RTT for every ping type at the end (calculated by simply dividing through the number of hops). Those four values at the end of every file should be the final results of the measurement, and should serve as reference RTT for the needed packet type (tables 2 and 3); The next lines in the tables show the number of hosts and routers encountered during the ping process, and in this way give an idea of the extent of this measurement.

Figure 13 shows the final values for the 2002 measure-

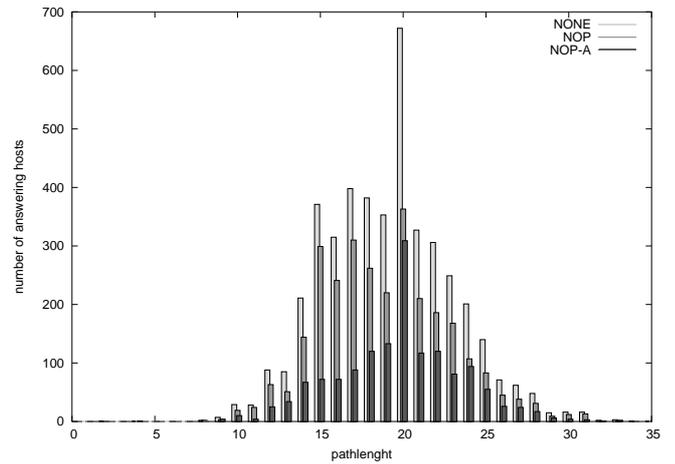


Figure 12: Answering hosts per path length (2003)

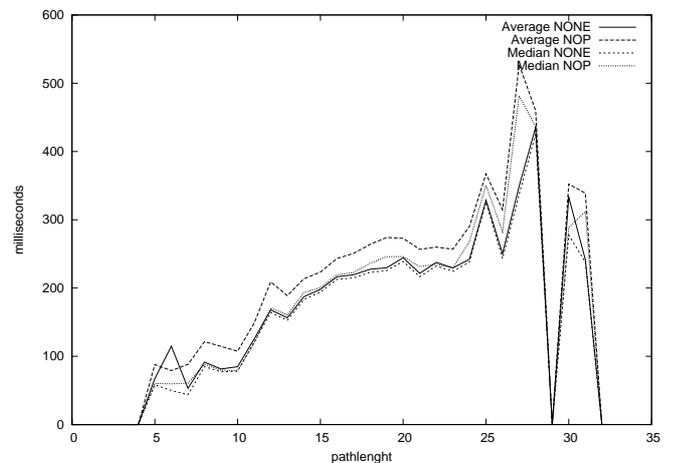


Figure 13: Averages and Medians per path length 2002

ment. This rather confusing graph shows that median and average of NONE pings are almost overlapping — but in general, the median line is slightly lower than the average. This slight difference is caused by the medians processed through standard deviation weighting two times (once for the gathering of the median itself and once by the statistical tool), in contrast to one time for the average (statistical only). The really important differences can be seen between the NONEs and the median NOP and average NOP. They are evidently higher than the NONEs and the average NOP is higher than the median NOP too; the fact that NOPs never fall below NONEs in this graph indicates that our statistics were a good choice and applied correctly. It does not come as a surprise that NOPs are slower than NONEs, but it is interesting to see that the difference between the average and the median of the NONEs is twice as great or even greater. This behaviour may be explained with the possibility, that packets with options were delayed more frequently than packets without options if there was congestion in the network. We measured that packets with options in the important medium path length range between 15 and 20 have a higher standard deviation than packets without (as

seen in Fig. 7). This could affect the whole measurement. From Fig. 8 we see, that most hosts have these path lengths, so it is possible that the median which ignores the anomalies eventually gets a lower value than the average which is affected by those anomalies.

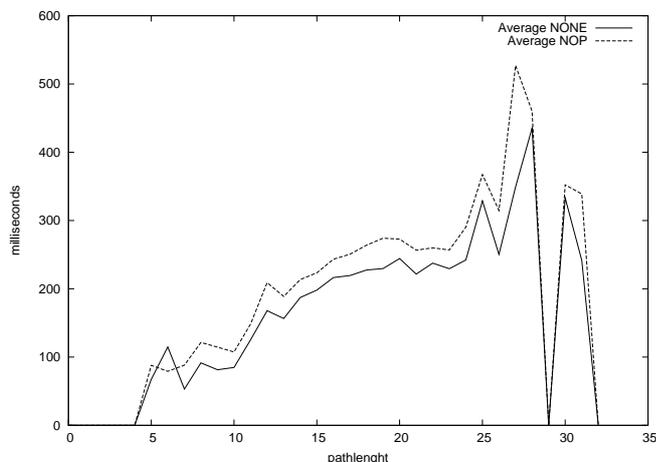


Figure 14: Averages per path length 2002

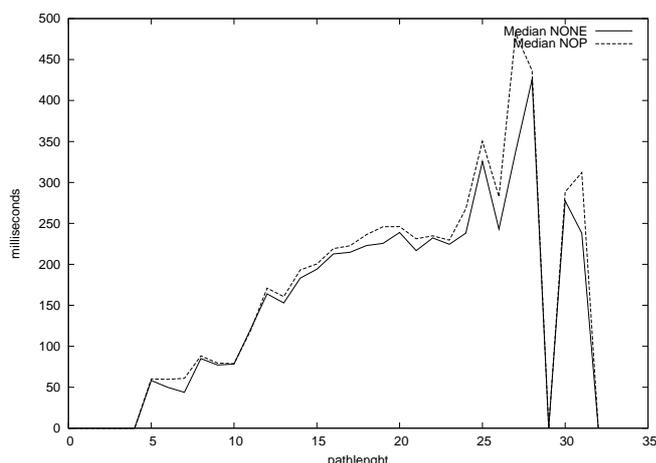


Figure 15: Medians per path length 2002

Figures 14 and 15 show a clearer difference between averages and medians of both pinging types. The NONE average exhibits some strange behaviours too. This is once again caused by the high oscillating values at extreme path lengths. The medians instead are quite stable, and differ by about 8 ms in average. The final results table (Table 2) shows the summarized final values.

We see that **10153** hosts and routers were hit: **4427** of them are the answering hosts of the host list, **5726** are routers encountered on the way to the hosts. But not all of the **4427** hosts were used; the interesting hosts are those which did not answer with options. Thus, if we subtract the NOP-A hosts (**1337**) from the complete host amount, **3090** hosts remain that were used for this statistic. If we do the same for the routers, **3226** really relevant routers remain. The difference between NOP and NONE is on average **19.55 ms** (median of 8.5, 30.6 on average). This means that the

Table 2: Measurement 2002

|   |                     |
|---|---------------------|
| weighted average per-host average NONE            | <b>198.521639</b>   |
| weighted average per-host average NOP             | <b>229.134040</b>   |
| weighted average per-host median NONE             | <b>194.188480</b>   |
| weighted average per-host median NOP              | <b>202.698017</b>   |
| processed hosts and routers                       | <b>10153</b>        |
| processed routers (with NOP-A)                    | <b>5726</b>         |
| processed routers (without NOP-A)                 | <b>3226</b>         |
| processed hosts only (NOP and NOP-A)              | <b>4427</b>         |
| NOP-A hosts                                       | <b>1337</b>         |
| hosts really interesting (NOP)                    | <b>3090</b>         |
| pinged hosts ( <i>with 10 pings measurement</i> ) | <b>5759 (27689)</b> |
| hosts which reacted to NONE only                  | <b>306 (4248)</b>   |

NOPs were processed only approximately **10%** slower than the NONEs.

**1332** hosts were lost during the ping process. This value is due to the unavailability of some hosts, for which various causes exist. Mostly, the ping process was interrupted because of a NONE timeout, that means, that the hosts did not answer NONE pings, a typical sign for a unreachable host. The more alarming value is that **306** of **5759** hosts were unreachable using NOP pings. This means that the host answered on NONE pings (otherwise a NONE timeout would have been reported), but probably did not react on NOP pings. This is a common problem with IP options. In fact, if we have a look at the original amount of pinged hosts (remember, the non-answering hosts of the 10 ping alternating test measurement not included in this measurement, therefore only 5760 hosts were pinged), we find **27689** pinged hosts and **4248** hosts that did not react to NOP pings. This value includes the **306** non-answering hosts to NOP, as explained above. It is possible that the hosts did not react to NOP and discarded the packets because of a firewall, but it is also possible be that routers discarded the packets on the way to the hosts (or in the backward direction, if the host included the option in its response). Unfortunately, we have no means to find out what exactly happened in these cases.

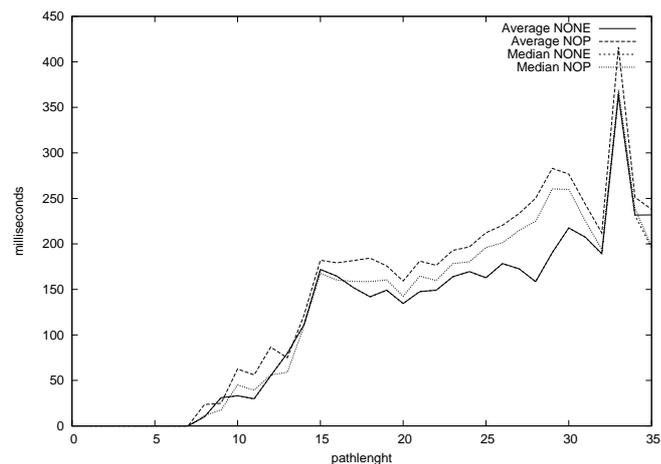


Figure 16: Averages and Medians per path length 2003

The measurement of 2003 has almost the same graph (Fig-

ure 16) as the one of 2002. Again average NONE and median NONE are almost overlapping, and median NOP and average NOP are higher than NONEs. This measurement however exhibits some strange behaviour: the median NOPs fall lower than the NONEs in three points. If we have a look at the standard deviation, we see that the NOPs is higher than the NONEs. This leads us to presume that congestion was not involved and once again, there may have been more fluctuations in the fast path processing delay of some routers than in the slow path.

Another anomaly is the peak at path length 33 for NONEs and NOPs, which probably was caused by a very slow router. The standard deviation at this path length is small, so a congestion may be excluded. We can see a high difference between NOP and NONE at the higher path lengths too — this only means that slow paths were encountered along the way. For the difference between median and average NOPs, the idea is the same as described in the statistics for the 2002 measurement. Figures 17 and 18 show the differences between NOP and NONE, once for the average, and once for the median.

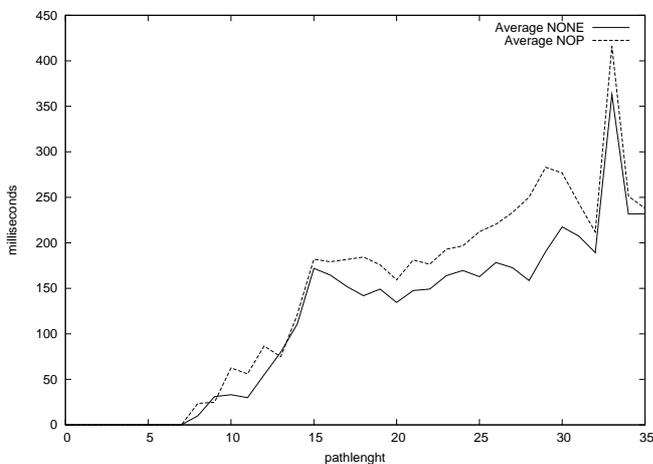


Figure 17: Averages per path length 2003

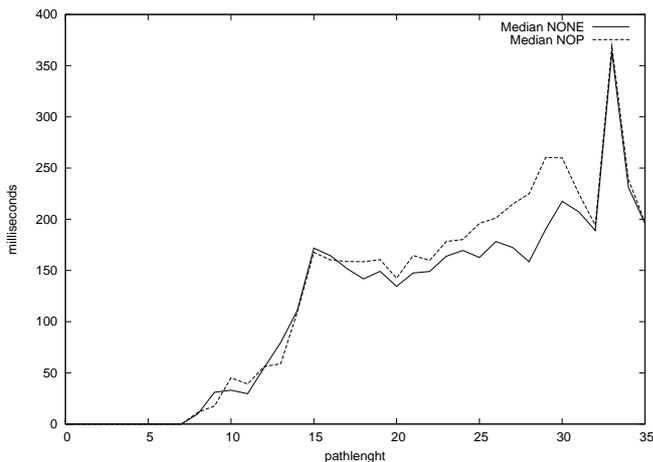


Figure 18: Medians per path length 2003

It can be seen again, that the differences between the av-

erages are higher than those between the medians. As explained before, the high difference between NONE and NOP at the higher path lengths, merely indicates the slower processing of packets with options. But even though the difference is high, it is an order of magnitude smaller than the RTT itself. This is an important conclusion. Table 3 shows the final summarized values for the measurement 2003.

Table 3: Measurement 2003

|                                      |                   |
|--------------------------------------|-------------------|
| weighted average average NONE        | <b>145.230768</b> |
| weighted average average NOP         | <b>165.180760</b> |
| weighted average median NONE         | <b>145.159222</b> |
| weighted average median NOP          | <b>147.551178</b> |
| processed hosts and routers          | <b>9595</b>       |
| processed routers (with NOP-A)       | <b>5194</b>       |
| processed routers (without NOP-A)    | <b>3041</b>       |
| processed hosts only (NOP and NOP-A) | <b>4401</b>       |
| NOP-A hosts                          | <b>1490</b>       |
| hosts really interesting (NOP)       | <b>2907</b>       |
| pinged hosts                         | <b>27689</b>      |
| hosts which reacted to NONE only     | <b>3507</b>       |

As we can see, the average and median NOP RTTs are almost equal as it is the case for the NONE. The difference between NONE and NOP is about **11.175 ms** on average (medina of 2.4, 19.95 on average). This means that the NOPs are approximately **7%** slower. The host statistic has a lower yield this time: **9595** hosts and routers were processed. This value results from the sum of **4401** hosts and **5194** routers. Here we can see, that the yield is lower than in the measurement of 2002. If we subtract the hosts that answered with ip options (**1490**) from the total number of hosts, we get a real yield of **2907** hosts and **3041** routers only. This is a lower value than in the measurement of 2002, but the complete processed host and router value is higher, quite contrary to the amount of NOP timeouts. In this measurement we pinged right from the start **27689** hosts, but only **3507** hosts did not answer NOP pings. Maybe we had better luck this time, and some more hosts or routers that usually discard NOP packets were just down for the pinging period — or maybe the packets did not encounter so many routers that throw away packets with options. However, we can still see that there are hosts and routers which discard packets with ip options. It can be that this happens accidentally, but it could even be that some routers work only on the fast path, without minding packets with options. Fortunately, there are still more hosts that answer to NOP than not.

## 5. DIFFERENCES

Some considerable changes happened between the measurement of 2002 and the one of 2003. As we have seen, the amount of hosts and routers were different. Figure 19 illustrates a comparison between the different ping types. We have not inserted any boxes showing the total amount of answering hosts because the hosts that answered NONE pings were the same as all answering hosts. That means that no answer had an IP option enabled (on NONE pings). Therefore, the NONE-A pings were ignored.

In Fig. 20 it can be seen, that the average and the median of the 2002 measurement are higher than the ones of 2003 for every ping type. This is probably due to the different

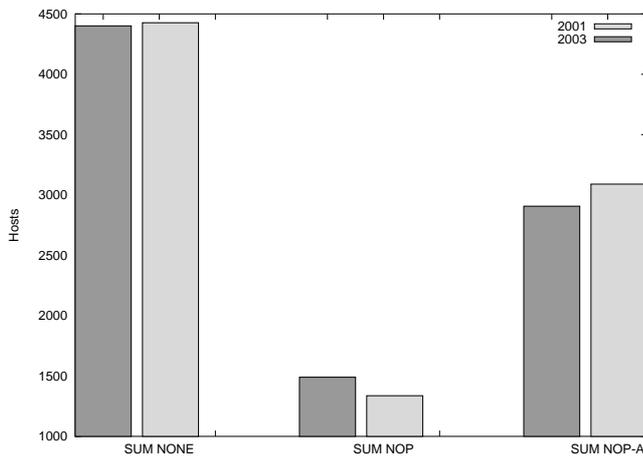


Figure 19: Difference of answering hosts per ping type between 2002 and 2003

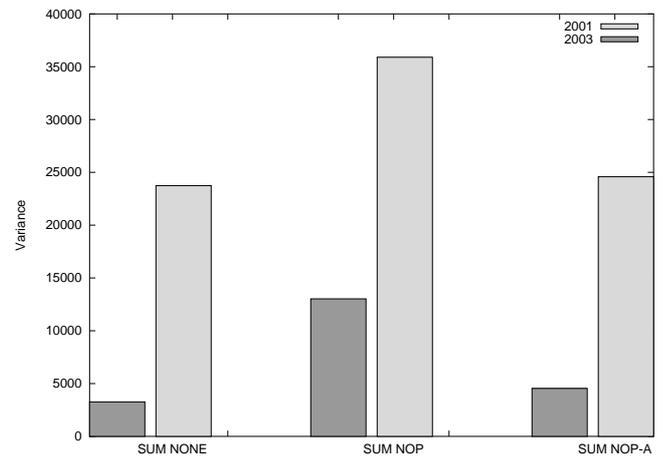


Figure 21: Variance differences between 2002 and 2003

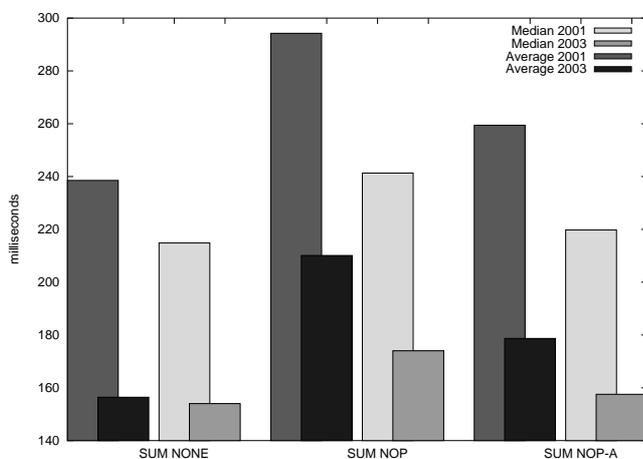


Figure 20: Average and median in 2002 and 2003

starting points — the measurement 2002 took place in Linz, but it did not pass through the highspeed network connection of the university. Instead, it passed through a possibly slower connection offered by a private provider. In 2003, a dedicated high speed Internet connection was used.

Figure 21 shows the differences between the variances. As expected, the values of 2002 are much higher. These differences do not affect the measurements, because the main interest lies in the difference between NONE and NOPs and not in the difference between the measurements themselves.

## 6. CONCLUSION

At the end we may say that even if we pinged only from two locations, we had a lot of results to work with; **27698** hosts were pinged, approximately **4400** answered, approximately **3000** gave an answer without options, and about **9500** took part in our measurements, either as destination host, or as router. About **3000** routers were processed while measuring the NOP differences, on routes that did not answer with options.

The final results show that the additional delay from slow

path processing was 10% of the RTT in the 2002 measurement and 7% in the 2003 measurement — which might indicate that the situation has in fact improved. In any case, the additional delay is clearly an order of magnitude smaller than the RTT. While these results cannot be seen as representative for the Internet, we can at least expect that some of the providers used standard “off the shelf” equipment, which would show similar behaviour in different places. Therefore, we believe that our results can provide some guidance when considering the use of IP options.

## 7. ACKNOWLEDGMENTS

The authors would like to thank Gerhard Wahl for creating the *extping* - program.

## 8. REFERENCES

- [1] W. Duerr and H. Mayer. *Wahrscheinlichkeitsrechnung und schliessende Statistik*. Hanser, Wien, 1987.
- [2] D. Moore, G. Voelker, and S. Savage. Inferring internet denial-of-service activity. 2001 USENIX Security, November 2001.
- [3] J. Padhye and S. Floyd. Identifying the tcp behavior of web servers. SIGCOMM2001, June 2001.