TECHNICAL REPORT

End-to-End Measurements on Performance Penalties of IPv4 Options

Pierre Fransson and Andreas Jonsson

Technical Report

Department of Computer Science and Electrical Engineering
Division of Computer Communication

End-to-End Measurements on Performance Penalties of IPv4 Options

Pierre Fransson
`pierre@sm.luth.se`

Andreas Jonsson
`aj@sm.luth.se`

Division of Computer Science and Networking
Department of Computer Science and Electrical Engineering
Luleå University of Technology
SE - 971 87 Luleå

February 18, 2004

Abstract

IP version 4 specifies options that extends the basic IP header and also allow new functions to be added to IP without breaking existing implementations. Since options must always be inspected routers, it is generally believed that routers prioritize ordinary packets over packets carrying options in a way that significantly impacts the performance of options. This article presents an experiment based on end-to-end probing using UDP packets, capturing the performance penalties associated with IPv4 options. Analysis of experiment results quantifies the impact of IPv4 options on forwarding performance in terms of delay, jitter and loss rate. From the analysis it can be concluded that there is a slight increase in delay and jitter and severely increased loss rate.

1 Introduction

IPv4 options are a generic and simple way of transmitting per packet information pertinent to network layer components, e.g., routers and hosts [15]. Besides the options in the specification, such as loose source routing, record route and Internet timestamp, IPv4 options have recently and in the past been proposed for several other custom extensions. A small

and non-exhaustive list of these extensions include: conveying information between different layers in the networking stack [9], providing quality of service in ad hoc networks[10], propagating the priority of a real-time media stream to a receiver [11], as method of enabling Active Networks [16, 5] and for multicast schemes [12, 14].

The existence of these proposals show that even though ip-options were conceived over two decades ago, there still exists the need for their functionality.

Unfortunately IPv4 packets containing options are reputed to be treated as second class citizens in the networking world, supposedly causing packets containing options to suffer from increased delay, increased jitter and dropped packets, compared to packets without options. Despite these alleged shortcomings of IP options there does not appear to exist any recent measurements of their impact.

The generally referred to reason, as to why packets containing options are exposed to greater delay, jitter and drops, is that routers are unable to handle options using fast-path forwarding. A reason for this is that when router hardware was developed in the end of the 1980's, IP options were not commonly used. Therefore router designers optimized for the common case, i.e., packets without options, thereby reducing complexity. Furthermore since options were

rarely used, there did not seem to be any reason to handle options at line speed [4].

One example is Cisco routers, which do not process options in hardware. As options may be used for denial of service on those routers, Cisco has introduced the possibility to drop or ignore packets with options [2].

Firewalls are also problematic, since these tend to be configured to drop everything except the most common types of packets. It can be noted, however, that this practice will always be an obstacle when deploying new functionality in the Internet.

This is a study of the performance difference packets with options experience compared to packets with no options. Data from network measurements based on UDP packets with and without options is analyzed with various statistical methods. If ICMP ping packets are used instead, it would be possible to study a wider area of the Internet. ICMP packets are, however, often forwarded differently than other packets, which would lessen the validity of the study.

The rest of this article is structured as follows. Section 2 describes the measurement method and the measured network. Section 3 presents the results and the statistical analysis of the data. Section 4 discusses an alternative to options. Finally, section 5 concludes the article.

2 Testing the Impact of IP-Options

This section presents a test consisting of experiments designed to discover characteristics of the difference in forwarding performance caused by the presence of an option. First, the metric of interest is described and a method for measuring this. This is followed by a description of the measured network and details on how statistical data is gathered. The data set and the software used for gathering the data can be downloaded from the web [1].

2.1 A Method for Measuring Difference in Forwarding Performance

The measurements captures *time in transit* of packets with and without options. Time in transit is the time between the moment a packet is sent until the moment it is received. The procedure to measure time in transit for a probe packet sent by host S to host R is as follows: Let t_s be the local time of S immediately before passing the packet to the *send* system call and t_r the local time of R immediately after the *receive* system call returns the packet. The resulting time in transit, $t = t_r - t_s$, includes a potentially large error caused by the fact that S and R do not share a clock.

There are two types of probe packets, one which includes IP-options (packets of this type are called *O-packets*) and one without (called *N-packets*). A *sample* is obtained by sending one N-packet and one O-packet and subtracting the times in transit. (This value is also called *delay-difference*.) This yields an accurate measurement of the difference, since the errors caused by S and R having different clocks cancel out each other.

Both O-packets and N-packets are UDP packets with payloads consisting of a unique sequence number (32 bits) and the send time, t_s , (64 bits). O-packets have the IP header length set to 6 words and bit 160 to 191 (i.e., the options) in the IP-header set to 0. The first octet among the options should be interpreted as “end of option list” and should not require any processing by routers [15]. N-packets includes 32 bits padding to be equally long as the O-packets (totally 44 octets, including headers).

2.2 The Network

Fifteen hosts from the resilient overlay network [3] and the network that connects them constitutes the subset of the Internet that has been measured. The nodes and the paths between the nodes is part of the public Internet. There is one host in Sweden, one in Canada, one in the Netherlands, and the rest of the hosts are at different locations in the USA. The network paths covered by the measurements traverse networks of several different service providers, such as: ATT Canada Telecom Services, ATT Worldnet,

Sprint (Sprintlink DK), AOL Transit Data Network, Level 3 Communications, and e.spire Communications. Several of the networks traversed during the test were large university networks, such as: Nordunet, Sunet, Surfnets, Geant, MIT, Harvard and University of Utah. The hosts and their locations are listed in table 1. Although the resilient overlay network provides features for routing at the application layer, this is not used in the test.

Table 1: The hosts used in the test

host	location
ron0	MIT
ron2	Utah
ron3	Denver?
ron4	Cambridge, MA
ron5	Luleå
ron11	Cornell
ron15	MIT
ron18	US
ron22	MIT
ron23	US (ATT)
ron24	NL
ron31	US (Teleglob)
ron33	Canada
ron35	New Jersey
ron36	San Diego

2.3 How Statistical Data is Gathered

An *experiment* involves a pair of hosts and consists of a set of *samples*. Within each sample two packets are sent as close in time as possible in randomized order. The randomization of send order is vital for the statistical analysis. Without it, packet transmissions cannot be said to be independent. Randomization ensures that systematic errors and network peculiarities do not bias the results.

During an experiment, samples are sent according to a Poisson process with intensity 0.0278 (about 2 samples per minute) to produce temporally unbiased measurements [13]. The receiver stores the send time, t_s , the receive time t_r and the sequence number of each received packet in a log file for later analysis.

Data was gathered under a time period of 174 hours starting at Apr 18 2003. All pairs among the 15 hosts ran two experiments (one in each direction) resulting in 210 data sets. The next section presents the analysis of these.

3 Results

In this section the data sets collected by the probing application, are subject to statistical analysis with respect to delay, jitter and loss rates.

Based on the widely accepted notion that packets carrying options receive inferior treatment on the Internet, being forced to traverse the slow path through routers[4, 2], we initially made the following assumptions regarding delay, jitter and loss rate:

- O-packets are delayed to a greater extent than N-packets. Most likely to such an extent that they are rendered useless for real-time applications.
- O-packets are subjected to larger jitter than N-packets. Probably to such an extent that they become useless for real-time applications.
- O-packets are subjected to higher loss rates than N-packets. Presumably to such an extent that they are hazardous to use for any application.

Regarding the two first assumptions, delay and jitter, the results show that for this *test* there is indeed a penalty of extra delay and an increase in jitter for O-packets. The increase in delay can be confirmed for all experiments except one at a very high level of confidence. The increase in jitter has not been verified using a statistical test, but comparisons of the *Inter Quantile Ranges* (see 3.3) show that for a majority of the experiments, the dispersion of O-packets is much larger than the dispersion of N-packets. While these differences exist for this test, it is not, however, certain that the absolute differences are so great as to render O-packets useless for real-time applications. The largest median of the delay-differences, see section 3.2, is 28 milliseconds and the majority of IQRs for O-packets are below 40 milliseconds. While extra delay never is a good thing, it is not certain that these values are large enough to be harmful.

Regarding the loss rates, the results do confirm our initial assumption. While not all experiments show signs of greater loss-rate for O-packets compared to N-packets, a substantial amount do show a large increase in loss rate. The most important finding is that many of the experiments exhibit exceptionally high loss-rates (up to 100%) for O-packets. While it is not possible to draw the conclusion that this behavior would be repeated in other tests, it is not possible to say that it could not occur. If the number of experiments with exceptionally high loss rates from this test are anything to judge by, then it would most likely be very unsound to advocate the use of options in any networking proposal that could see widespread use.

The observed differences in delay, jitter and packet losses corroborate the notion that a significant number of routers handle IP options via the “slow path”. The results show that the penalty for taking the slow path is not negligible and can in the case of packet losses be severely detrimental to any application using IP options.

3.1 Analysis of Result Data

The data sets resulting from each experiment are statistically analyzed by various scripts (which can be obtained from the web [1]). The results of the analysis are presented in the tables 2–9. All tables are presented in the same format. The rows represent receiving hosts and the columns represent sending hosts. The elements in the table are the resulting statistical values for each experiment.

To ensure that the statistics from each experiment is based on a large enough sample size, the tables only include values where the sample size is at least 500. The only exceptions to this are the tables concerning loss-rate statistics, since these statistics are only based on the quantities of lost packets and not on values derived from the samples themselves.

On average, each experiment has resulted in a data set of roughly 17,000 samples. For the host *ron24*, fewer samples were received (approximately 8,000) due to the host being down for a prolonged period of time. Unfortunately, all experiments gathered by *ron23* were lost due to a reconfiguration of the host

after the *test* had been concluded. All tables therefore exhibit an absence of data at the row corresponding to *ron23*. However, since data sent from *ron23* is present in the tables, the corresponding row remains in the table.

The tables relating to loss-rates, only present statistics from 196 data sets, since the data sets received by *ron23* were lost. For all other tables the number of presented statistics is 166, due to the removal of data sets with less than 500 samples.

3.2 Analysis of Delay

To determine if O-packets experience greater delay than N-packets, the difference in delay for each sample was examined. Since the packets in a sample is sent back-to-back, the networking conditions experienced should more or less identical. As mentioned earlier the order between O-packets and N-packets within a sample were randomized to avoid dependencies due to run order. Also, the time intervals between samples were randomly selected from an exponential distribution to ensure independence between all sample-pairs. For convenience, the difference in delay between the O-packet and the N-packet for each sample is hereafter referred to as *delay-difference*. Throughout this section, subscript *N* is used for variables corresponding to N-packets and subscript *O* for variables corresponding to O-packets.

Initially the delay model in (1) was assumed. The delay experienced by a packet (y) is determined by an average delay (μ) and a random error (ϵ), which is a normally distributed and independent random variable. The resulting delay-difference for each sample ($y_O - y_N$) would therefore also be a normally distributed and independent random variable. It was, however, quickly determined that the assumption of normality did not hold (i.e., the random variables ϵ_O and ϵ_N were not normally distributed). The delay histograms of nearly all experiments showed clear signs of being skewed and long-tailed. The delay histograms bear close resemblance to the delay histograms found in [6], where parts of the RIPE NCC measurements of end-to-end delay were analyzed. The histogram in Fig. 1 and the normal probability plot in Fig. 2 are typical of most of the ex-

Figure 1: Histogram of delay-difference

Figure 2: Normal probability plot of delay-differences

periments and clearly show that the assumptions of normality are violated. It should be noted that although Fig. 1 can be considered as a typical example of the general shape of the distributions, there exists large variations in the location of central tendency, dispersion and the extent to which the distributions are skewed and long-tailed.

$$\begin{aligned} y_O &= \mu_O + \epsilon_O \\ y_N &= \mu_N + \epsilon_N \end{aligned} \quad (1)$$

Because the delay distributions of the experiments are rather heavily skewed and long-tailed, the mean is not an appropriate measure of the central tendency of the delay-difference [8]. A better measure of central tendency is instead the median, M . Evaluation of the delay-difference was therefore performed using the median for each experiment. It is easy to realize that if there on average are no differences in delay between O-packets and N-packets, the median would be very close to zero and about 50% of the delay-differences would have negative signs and the other half should have positive signs.

Table 2 shows the medians of the delay-difference for all experiments. While most experiments have a median value of a few milliseconds, there are a num-

ber of experiments with substantially higher medians. Experiments with medians greater than 20 milliseconds are highlighted.

In order to determine whether the observed medians indicate a statistically significant delay-difference, we have used the non-parametric sign-test [8] to test the hypothesis in (2) for each experiment. If we are able to reject the null hypothesis at a high confidence level, we can conclude (at the given level of confidence) that O-packets are indeed delayed to a greater extent for an individual experiment.

$$\begin{aligned} H_0 : M &= 0 \\ H_1 : M &> 0 \end{aligned} \quad (2)$$

Table 3 shows the probabilities for each experiment, that the the null hypothesis, H_0 , is correct. A value of 0 indicates a probability value (P-value) less than $2.2 \cdot 10^{-16}$, the smallest value attainable given the precision of the statistical calculations. As can be seen from the table only three experiments have a P-value greater than $2.2 \cdot 10^{-16}$. For all other experiments the null hypothesis can be rejected at an incredibly high level of confidence. Of the other three experiments, there is only one experiment where H_0

Table 2: Medians of delay-differences for all experiments [milliseconds]

	ron0	ron2	ron3	ron4	ron5	ron11	ron15	ron18	ron22	ron23	ron24	ron31	ron33	ron35	ron36
ron0	-	23.2	7.85	3.40	4.98	4.16	4.99	9.26	-	6.09	5.15	-	4.50	13.8	28.0
ron2	20.9	-	2.68	2.39	19.9	16.9	3.94	2.26	2.61	14.9	22.7	-	4.90	4.68	8.62
ron3	5.20	2.49	-	2.95	13.0	3.60	4.69	-	1.76	19.1	13.0	3.24	4.73	2.12	-
ron4	2.29	2.98	3.25	-	3.30	3.35	2.83	3.03	-	13.5	2.80	-	9.42	8.41	5.68
ron5	3.99	21.7	3.81	4.82	-	3.03	4.28	28.0	-	20.7	3.41	4.43	16.7	25.7	27.0
ron11	8.99	17.5	4.08	3.30	7.35	-	1.77	3.66	-	12.8	8.75	4.72	2.78	18.5	22.0
ron15	3.37	1.97	-	-2.43	1.72	2.46	-	1.24	1.73	10.7	2.21	0.73	2.20	5.42	3.94
ron18	4.12	1.02	-	1.41	7.59	-	0.28	-	-	16.6	-	6.62	14.3	14.0	8.11
ron22	3.52	2.21	1.56	-	5.24	2.61	3.35	-	-	-	-	-	2.37	-	-
ron23	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
ron24	3.39	22.4	11.4	2.83	3.64	1.77	10.5	27.2	-	18.4	-	2.80	15.7	7.72	27.1
ron31	3.70	-	3.27	2.51	7.74	2.82	3.36	2.66	-	15.7	7.19	-	2.10	9.57	10.6
ron33	24.0	18.6	16.4	15.3	22.9	2.90	4.71	16.0	-	12.9	-	-	-	5.96	26.5
ron35	24.0	13.4	23.4	22.5	24.4	19.6	9.56	8.21	-	22.5	-	8.84	7.95	-	20.0
ron36	25.6	8.11	7.69	6.51	24.9	21.3	6.93	6.05	-	19.8	27.3	3.22	8.56	15.5	-

cannot be rejected at any reasonable level of confidence.

In other words, all but one experiment show at a tremendous level of confidence that packets carrying options experience higher delay than ordinary packets. On the other hand the extra incurred delay is small.

3.3 Analysis of Jitter

It is possible to determine if O-packets experience higher jitter than N-packets, by examining the difference in dispersion between the two packet types. The obvious choice to measure dispersion would be the sample variance. Variance, however, is difficult to interpret for distributions that are both skewed and long-tailed. A better measure is therefore the use of the Inter Quantile Range (IQR), defined as the difference between the 3rd and the 1st quartiles, Q_3 and Q_1 respectively. A benefit of using the IQR is that it is an intuitive metric; half of the packets are “closer to each other” (measured in time) than the value of the IQR and the other half of the packets are “further apart” (also measured in time) than the value of the IQR.

Table 4 shows the IQRs for O-packets. The high-

lighted values, 15.7% of all experiments (26/166), have IQR values greater than 50 milliseconds, which can be regarded as a rather large delay if real-time applications are considered. For the N-packets 4.8% (8/166) have an IQR that is above 50 milliseconds. These are the experiments with IQRs above 500 milliseconds, all received by the host *ron11*. This shows that while O-packets often seem susceptible to higher jitter, there are still networking conditions that can cause substantially larger variations for both packet types.

A more comprehensive comparison of the IQR of both packets types is shown in table 5, which shows the quotients of the IQR for O-packets and the IQR for N-packets (IQR_O/IQR_N). The IQR quotients span the range from 192 to 0.97. The highlighted values, quotients greater than 10, account for 28.9% of all the values. Only one quotient is less than 1.0, where the IQR for O-packets is actually less than the IQR for N-packets. Lastly, quotients in the range 1.0 – 2.0 and 2.0 – 10.0 account for 15.7% and 54.8% respectively. All in all, quotients greater than 2.0 account for 83.4% of the experiments.

Finally, an important part of analyzing jitter is to examine how bad the worst case delays are and how

Table 3: Probability that H_0 is correct. Zero-values indicate a probability less than $2.2 \cdot 10^{-16}$

	ron0	ron2	ron3	ron4	ron5	ron11	ron15	ron18	ron22	ron23	ron24	ron31	ron33	ron35	ron36
ron0	-	0	0	0	0	0	0	0	-	0	0	-	0	0	0
ron2	0	-	0	0	0	0	$2.34 \cdot 10^{-12}$	0	0	0	0	-	0	0	0
ron3	0	0	-	0	0	0	0	-	0	0	0	0	0	0	-
ron4	0	0	0	-	0	0	0	0	-	0	0	-	0	0	0
ron5	0	0	0	0	-	0	0	0	-	0	0	0	0	0	0
ron11	0	0	0	0	0	-	0	0	-	0	0	0	0	0	0
ron15	0	0	-	1.00	0	0	-	0	0	0	0	0	0	0	0
ron18	0	0	-	0	0	-	0	-	-	0	-	0	0	0	0
ron22	0	0	0	-	0	0	0	-	-	-	-	-	0	-	-
ron23	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
ron24	0	0	0	0	0	0	0	0	-	0	-	0	0	0	0
ron31	0	-	0	0	0	0	0.02092	0	-	0	0	-	0	0	0
ron33	0	0	0	0	0	0	0	0	-	0	-	-	-	0	0
ron35	0	0	0	0	0	0	0	0	-	0	-	0	0	-	0
ron36	0	0	0	0	0	0	0	0	-	0	0	0	0	0	-

Table 4: Inter Quantile Ranges of O-packets for all experiments [milliseconds]

	ron0	ron2	ron3	ron4	ron5	ron11	ron15	ron18	ron22	ron23	ron24	ron31	ron33	ron35	ron36
ron0	-	25.9	52.1	5.85	15.1	9.07	27.9	20.4	-	21.8	14.3	-	12.0	24.7	29.1
ron2	25.1	-	31.3	1.57	28.0	21.1	11.1	6.62	1.52	24.4	24.7	-	11.5	17.6	12.9
ron3	42.8	3.29	-	18.6	75.6	54.0	125	-	2.78	49.0	65.7	50.1	89.0	46.3	-
ron4	6.08	15.4	40.1	-	8.39	15.7	15.6	11.4	-	26.5	10.5	-	22.5	21.0	10.4
ron5	15.2	28.7	88.5	18.8	-	24.7	47.8	85.8	-	43.1	29.1	28.1	43.9	60.1	31.7
ron11	16.7	20.8	655	605	18.2	-	698	607	-	630	16.5	625	624	618	23.9
ron15	3.64	1.40	-	2.42	3.27	13.2	-	8.79	1.18	25.1	4.81	4.76	2.60	18.0	9.30
ron18	22.9	17.6	-	17.6	39.2	-	54.8	-	-	40.7	-	46.0	32.8	35.8	19.7
ron22	8.91	3.95	9.97	-	34.8	22.8	44.7	-	-	-	-	-	9.99	-	-
ron23	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
ron24	11.2	24.4	62.1	11.9	29.8	7.17	48.1	78.8	-	40.1	-	24.5	39.4	24.3	26.7
ron31	13.1	-	47.7	6.14	27.0	20.9	4.86	39.0	-	31.4	24.8	-	2.40	24.2	23.7
ron33	28.8	1.69	90.7	19.9	32.1	20.4	52.2	30.8	-	32.2	-	-	-	27.2	27.7
ron35	32.3	23.8	71.8	26.1	38.0	32.8	62.6	26.6	-	38.0	-	21.9	24.1	-	27.5
ron36	27.2	12.6	33.3	13.1	30.5	22.9	62.2	20.7	-	32.1	27.5	7.34	25.1	29.9	-

Table 5: Quotient of IQR (IQR_O/IQR_N) for all experiments

	ron0	ron2	ron3	ron4	ron5	ron11	ron15	ron18	ron22	ron23	ron24	ron31	ron33	ron35	ron36
ron0	-	21.6	40.0	2.32	4.08	6.01	4.49	9.64	-	1.83	3.38	-	8.15	2.56	18.2
ron2	57.2	-	54.2	1.39	22.6	97.3	2.82	8.66	4.04	2.10	6.14	-	31.1	1.80	16.6
ron3	36.1	4.37	-	12.0	77.2	6.02	31.8	-	1.39	4.03	16.4	20.1	192	4.90	-
ron4	1.48	23.4	53.1	-	5.16	1.42	1.56	8.80	-	1.66	2.36	-	28.4	2.07	6.36
ron5	3.34	6.42	82.2	11.0	-	1.36	11.1	60.4	-	3.72	6.90	10.5	40.4	6.23	7.32
ron11	33.4	89.7	1.06	1.00	5.47	-	1.01	1.02	-	1.00	4.41	1.02	1.02	1.02	26.7
ron15	5.98	2.92	-	1.24	2.80	1.20	-	8.52	2.08	1.74	1.30	1.01	5.13	2.15	9.12
ron18	2.83	2.19	-	2.07	4.85	-	3.58	-	-	2.32	-	5.54	4.09	2.79	2.45
ron22	6.73	6.84	1.13	-	22.3	2.09	4.96	-	-	-	-	-	7.99	-	-
ron23	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
ron24	3.05	6.20	17.3	3.46	7.31	1.89	9.81	18.8	-	2.90	-	3.88	10.2	3.81	7.44
ron31	11.0	-	19.5	3.79	10.5	2.43	1.04	15.4	-	2.62	3.89	-	0.97	2.48	6.88
ron33	8.28	2.28	172	11.9	31.4	2.24	12.8	34.5	-	2.77	-	-	-	2.88	28.4
ron35	3.12	2.20	6.65	2.33	3.45	2.56	2.82	2.34	-	2.09	-	1.95	2.49	-	2.68
ron36	31.1	16.6	35.1	6.25	9.14	19.5	3.10	12.8	-	2.16	7.79	2.07	28.2	3.00	-

these are distributed over time. This is especially important when considering the impact on real-time applications. By examining the proportion of O-packets that are heavily delayed compared to N-packets (e.g., > 100 milliseconds), it is possible to get an impression of how a real-time application would cope with the extra delay. It is also important to understand how these heavily delayed packets are distributed over time since this also affects how an application copes with large delays. For instance, while a real-time application can be assumed to cope with large delays if they occur in isolation and far apart. This is not the case if long continuous sequences of packets are heavily delayed.

Each entry in table 6 shows the lower bound of the 5% largest delay-differences for each experiment. The lower bound of the 5% largest delay-differences exceeds 100 milliseconds in 62% of the experiments (highlighted in the table); the lower bound of the 10% largest delay-differences exceeds 100 milliseconds in 25.3% of the experiments; the lower bound of the 20% largest delay-differences exceeds 100 milliseconds in 4.2% of the experiments. It is obvious from the table 6 and the numbers presented above that a substantial amount of O-packets are delayed

in excess of 100 milliseconds, when compared to N-packets. The question of how these packets are distributed therefore needs to be answered. Fig. 3 shows a typical example of how the 5%–0.1% largest delay-differences¹ are distributed in time. As can be seen in the figure, there does not seem to exist any apparent pattern that would suggest an uneven distribution of the largest delay-differences. The rest of the experiments show similar distributions, indicating that large delay-differences are spread out well, not occurring in long continuous sequences.

3.4 Analysis of Loss rate

The most significant impact of using IPv4 options is if the likelihood of packets being lost is increased. If this is the case, options not only become unsafe to use for real-time applications but also for all other applications.

Table 7 and table 8 shows the proportion of lost packets, for O-packets and N-packets respectively. The highlighted values in table 7 indicate experiments with loss rates of 15% or higher, accounting

¹The 0.1% largest delay-differences (18 values) have been omitted from the figure, in order to provide better resolution. In total Fig. 3 shows approximately 840 values

Table 6: 5% worst case delay-differences for all experiments [milliseconds]

	ron0	ron2	ron3	ron4	ron5	ron11	ron15	ron18	ron22	ron23	ron24	ron31	ron33	ron35	ron36
ron0	-	122	218	65.5	164	87.3	71.4	153	-	82.1	98.8	-	173	171	114
ron2	111	-	215	11.7	158	72.6	11.2	94.7	111	43.6	98.0	-	134	128	46.6
ron3	221	67.4	-	205	223	202	213	-	8.47	208	217	206	235	204	-
ron4	54.0	47.2	201	-	118	21.2	11.9	94.6	-	40.9	65.4	-	145	119	36.3
ron5	160	151	235	127	-	129	114	212	-	138	140	139	189	195	149
ron11	93.6	71.7	204	15.5	137	-	24.6	110	-	46.7	84.2	87.3	133	141	74.7
ron15	57.7	14.1	-	3.95	114	11.4	-	89.4	61.1	38.0	55.2	1.22	128	120	31.0
ron18	139	113	-	98.7	163	-	88.2	-	-	113	-	173	168	164	97.7
ron22	120	91.9	6.46	-	162	112	77.9	-	-	-	-	-	163	-	-
ron23	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
ron24	98.0	98.1	212	74.9	139	73.0	117	199	-	97.2	-	100	155	129	92.9
ron31	92.4	-	202	46.7	140	62.5	4.03	173	-	63.5	107	-	128	125	90.1
ron33	184	136	246	142	192	139	128	177	-	136	-	-	-	177	149
ron35	172	141	229	130	184	140	123	163	-	130	-	123	177	-	133
ron36	104	46.2	202	44.0	136	70.3	75.1	102	-	62.4	94.9	24.8	165	132	-

Figure 3: Plot of the 5%–0.1% largest delay-differences

for 22.9% of all experiments (many of them close or equal to 100%). This can be compared to the highest loss rate for N-packets which is only 3.5%, highlighted in table 8. In table 7, 14.3% of the experiments exhibit loss-rates of 99.9% or higher for O-packets, and while most of these experiments (17 out of 28) either emanate from or are received by the host *ron22*, there is still a large proportion that are not directly coupled to *ron22*. Continuing the categorization, 5.6% of all the experiments exhibit loss-rates between 49% to 99.9%, and 3.1% of the experiments exhibit loss-rates between 15% and 49%. Although, as mentioned earlier, many of the experiments containing large loss-rates are related to one host, *ron22*, the rest of the large loss-rates are evenly spread out. This is easily realized by noting that even if *ron22* is removed from table 7, there is only one sender (*ron35*) and one receiver (*ron5*) that would not be associated with any experiments experiencing loss-rates greater than 15%.

A comparison of the loss-rates of O-packets and N-packets is shown in table 9, which contains the quotient between the number of lost O-packets and the number of lost N-packets ($\#lost_O/\#lost_N$). The highlighted values are the quotients that are larger

Table 7: Proportion of lost O-packets for all experiments [%]

	ron0	ron2	ron3	ron4	ron5	ron11	ron15	ron18	ron22	ron23	ron24	ron31	ron33	ron35	ron36
ron0	-	0.22	1.16	0.34	1.00	0.27	0.48	0.68	100	0.19	0.38	99.4	0.19	0.46	0.16
ron2	0.53	-	0.08	0.11	0.78	0.22	0.16	0.54	1.30	0.18	0.03	100.0	0.04	0.11	0.03
ron3	0.99	2.24	-	0.08	0.75	0.50	0.15	100	0.28	0.14	0.57	0.07	0.01	0.09	100
ron4	1.11	0.67	0.25	-	0.84	0.58	0.16	0.55	100	0.17	0.37	100	0.23	0.83	0.03
ron5	1.28	0.91	0.65	0.79	-	0.78	0.73	1.23	100	0.87	0.33	0.70	0.79	0.83	0.95
ron11	0.53	0.15	50.0	2.13	0.76	-	2.58	2.60	100	2.46	0.43	51.7	2.16	2.14	0.07
ron15	2.74	2.54	100	1.42	2.31	2.11	-	2.02	1.82	1.72	3.56	1.63	1.66	1.76	1.42
ron18	0.54	0.42	100	0.09	0.78	98.0	0.14	-	100	0.21	100	0.03	0.04	0.16	0.03
ron22	3.71	2.40	0.19	100.0	1.49	0.56	0.22	100.0	-	100.0	100	100	0.14	100.0	100
ron23	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
ron24	0.23	0.01	0.39	49.4	0.40	0.42	0.94	1.20	100	49.2	-	0.61	50.0	0.74	0.19
ron31	1.33	100	0.20	0.32	0.85	0.61	0.46	0.78	100	0.34	0.70	-	0.25	0.30	49.4
ron33	1.70	49.7	71.7	0.06	0.95	0.49	1.46	0.46	100	0.12	100	100.0	-	0.07	0.02
ron35	1.36	0.56	50.0	0.12	0.75	0.64	0.29	0.65	100	0.24	100.0	0.07	0.12	-	0.13
ron36	18.5	18.2	0.15	0.22	18.8	18.6	31.5	0.59	100	0.25	20.1	0.13	0.18	0.23	-

Table 8: Proportion of lost N-packets for all experiments [%]

	ron0	ron2	ron3	ron4	ron5	ron11	ron15	ron18	ron22	ron23	ron24	ron31	ron33	ron35	ron36
ron0	-	0.29	0.38	0.38	0.38	0.44	0.60	0.81	0.58	0.26	0.46	0.41	0.33	0.49	0.35
ron2	0.63	-	0.08	0.06	0.06	0.20	0.16	0.55	0.44	0.17	0.03	0.10	0.07	0.12	0.02
ron3	1.22	3.09	-	0.02	0.03	0.43	0.15	0.59	0.20	0.15	0.15	0.05	0.02	0.11	0.01
ron4	1.33	0.06	0.05	-	0.04	0.46	0.15	0.53	0.16	0.18	0.26	0.03	0.03	0.14	0.02
ron5	0.69	0.08	0.02	0.02	-	0.12	0.16	0.49	0.20	0.20	0.13	0.06	0.03	0.15	0.04
ron11	0.66	0.11	2.21	2.24	0.16	-	2.58	2.79	2.38	2.48	0.36	2.03	2.19	2.36	0.07
ron15	2.31	1.67	1.74	1.36	1.56	2.00	-	2.01	1.82	1.74	3.58	1.58	1.59	1.76	1.43
ron18	0.61	0.08	0.02	0.02	0.07	0.66	0.14	-	0.17	0.21	0.23	0.03	0.02	0.16	0.02
ron22	0.87	3.06	0.13	0.09	0.09	0.54	0.21	0.59	-	0.22	0.28	0.24	0.11	0.25	0.09
ron23	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
ron24	0.24	0.00	0.13	0.22	0.15	0.42	0.51	0.68	0.18	0.41	-	0.20	0.36	1.04	0.18
ron31	1.57	0.20	0.19	0.19	0.17	0.68	0.49	0.73	0.34	0.26	0.31	-	0.14	0.19	0.21
ron33	1.56	0.13	0.04	0.03	0.07	0.45	0.22	0.54	0.18	0.15	0.18	0.07	-	0.09	0.03
ron35	0.95	0.15	0.06	0.11	0.09	0.66	0.30	0.66	0.25	0.24	0.19	0.06	0.13	-	0.14
ron36	0.58	0.05	0.07	0.02	0.10	0.13	0.18	0.50	0.18	0.11	0.17	0.00	0.01	0.14	-

than 2.0, accounting for 44.9% of all the experiments. The values that are both highlighted and italicized are loss rates that are higher than 2.0, but which were not highlighted in table 7. These values account for 21.9% of all the experiments, indicating that O-packets do not only have higher loss rates when drastic amounts of losses occur (here meaning greater than 15%) but also in cases of less severe loss rates.

Finally, an important question that arises, is if the drastic loss rates ($> 15\%$) seen for O-packets, manifest themselves as long continuous sequences of losses or if losses are evenly distributed. If losses occur as long continuous sequences, a possible explanation for the high loss rates might be that traffic was rerouted during a period, traversing a router that discards all O-packets. Another explanation could be that excessive load triggers a state in some routers resulting in loss of all packets for a prolonged period of time. In either case it is clear that long continuous sequences of losses is very hard for any application to deal with. Should however losses be evenly spread out over the entire transmission, this indicates that O-packets might simply incur a higher loss rate, which is evenly distributed in time. It turns out that both types of loss patterns can be observed. Fig. 4 shows an example of long continuous loss sequences, resulting in 31.5% loss of O-packets. The topmost graph shows the occurrences of received and lost packets plotted against the packet sequence numbers. The bottommost graph shows the cumulative sum of all the occurring losses. As can be seen there is a sudden change in the cumulative sum close to packet sequence number 24000, showing clearly that losses are not spread out evenly over the entire transmission. Fig. 5 in turn, shows an example of losses occurring evenly resulting in 49.7% loss of O-packets. Of the experiments with packets losses in the range of 15% to 80%, two experiments can be categorized as long continuous sequences of losses and the other 13 experiments, as evenly distributed losses.

3.5 Discussion of Results

It is important to note, that while it is possible to perform statistically significant tests on each individ-

Figure 4: Plot of long continuous loss sequence

Figure 5: Plot of evenly distributed losses

Table 9: Quotient between # of lost O-packets and # of lost N-packets

	ron0	ron2	ron3	ron4	ron5	ron11	ron15	ron18	ron22	ron23	ron24	ron31	ron33	ron35	ron36
ron0	-	0.76	<i>3.03</i>	0.88	<i>2.61</i>	0.60	0.80	0.83	174	0.72	0.84	244	0.58	0.94	0.44
ron2	0.84	-	1.00	1.80	<i>13.2</i>	1.09	1.00	0.98	<i>2.99</i>	1.07	1.00	1048	0.58	0.90	1.25
ron3	0.81	0.73	-	<i>3.50</i>	<i>21.8</i>	1.14	1.00	171	1.39	0.92	<i>3.83</i>	1.33	0.67	0.84	8687
ron4	0.84	<i>10.4</i>	<i>5.50</i>	-	<i>20.7</i>	1.27	1.04	1.03	615	0.94	1.41	2894	<i>8.00</i>	<i>5.76</i>	1.67
ron5	1.85	<i>11.9</i>	<i>28.2</i>	<i>34.5</i>	-	<i>6.43</i>	<i>4.70</i>	<i>2.54</i>	492	<i>4.44</i>	<i>2.55</i>	<i>12.2</i>	<i>22.8</i>	<i>5.54</i>	<i>23.7</i>
ron11	0.80	1.39	22.6	0.95	<i>4.75</i>	-	1.00	0.93	42.0	0.99	1.20	25.4	0.99	0.91	1.00
ron15	1.19	1.52	57.5	1.04	1.48	1.06	-	1.01	1.00	0.99	0.99	1.03	1.04	1.00	0.99
ron18	0.88	<i>5.46</i>	4333	<i>4.00</i>	<i>11.3</i>	149	0.96	-	585	0.97	436	1.20	1.75	1.00	1.67
ron22	<i>4.24</i>	0.78	1.55	1150	<i>16.2</i>	1.04	1.05	169	-	459	360	424	1.32	396	1173
ron23	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
ron24	0.95	1/0	<i>2.91</i>	226	<i>2.75</i>	1.00	1.86	1.75	562	120	-	<i>3.12</i>	138	0.71	1.07
ron31	0.85	508	1.06	1.70	<i>4.97</i>	0.90	0.93	1.07	298	1.30	<i>2.27</i>	-	1.72	1.56	230
ron33	1.08	394	1795	2.00	<i>12.7</i>	1.09	<i>6.68</i>	0.84	546	0.81	541	1467	-	0.87	0.50
ron35	1.43	<i>3.84</i>	796	1.11	<i>8.25</i>	0.97	0.98	0.99	405	1.00	518	1.18	0.87	-	0.96
ron36	31.7	338	<i>2.17</i>	<i>13.0</i>	193	146	171	1.18	560	<i>2.20</i>	120	23/0	<i>31.0</i>	1.60	-

ual experiment, this does not hold for tests aimed at comparing the outcome of different experiments to each other. The reason for this is that the different experiments cannot be guaranteed to be independent from each other, a critical requirement for most statistical tests. It is in fact quite likely that several of the flow generated by the experiments share routers that may be of consequence with regards to delay, jitter and loss rate. Furthermore, the hosts in the test cannot be said to have been completely randomly selected from all hosts on the Internet. The participating hosts have been selected from the set of available ron-nodes (which we were kindly allowed to use). Neither has selection amongst ron-nodes been completely random. We have for instance strived to avoid hosts that were close to each other with regards to the number of intermediate routers, since it could potentially have been problematic to detect differences in delay if the delays were very small to begin with. Finding out the parts of the end-to-end paths that are shared and which of those parts that contribute to the differences in delay, jitter and packet loss is a difficult task, as it requires information about internal network details which network owners and administrators are often reluctant to share such infor-

mation.

However, for the sake of this investigation the above issues are not a problem. The intention of this investigation is not to model the behavior of IPv4 options on the Internet as whole, a task which is of course much greater and perhaps even infeasible. Instead we focus on the observations that can be made for this test, arguing that while the observed problems cannot be proven to manifest themselves throughout the Internet, it cannot be ruled out that these problems would not be encountered elsewhere on the Internet.

4 An Alternative to IP-Options

It is clear that the use of built in IP options for customizing the IP header results in degraded forwarding performance. An alternative is to use custom extension headers instead. To achieve this, a chain of extension headers follows the IP header. The protocol number field in the IP header is redefined to indicate the type of the next header. The format could be identical to IP version 6 extension headers [7].

For backwards compatibility, the use of this feature

must be negotiated beforehand. There is an unused bit in the IP-header which can be used for this negotiation. RFC791 specifies that this bit must be set to zero and that its value must be ignored [15]. Hence, assuming that implementations have followed the specification, this bit can be set to one to indicate support for extension headers. This is a good use of the unused bit, as it provides opportunity for arbitrary extensions of the IPv4 header.

Firewalls will still be a problem. There is no reasonable way to allow compliancy with current firewall practice². Instead, either firewall configurations must be updated or worked around with transport layer IP-tunnels.

5 Conclusions

The belief that packets carrying options (O-packets) are subject to inferior treatment compared to ordinary packets (N-packets) has been investigated. This is of interest as functionality provided by options are important for the development of new functions in IP.

The analysis of delay, jitter and losses for O-packets shows that the use of options is indeed problematic. For the presented experiments, we can with high level of confidence state that the delay is higher for O-packets than for P-packets, although the extra delay is small. O-packets also experience higher jitter. While the extra delay and jitter may be tolerable for most application, the extra losses are not. 22.9% of the experiments exhibit loss rates of 15% or higher for O-packets, many of them close or equal to 100%. This can be compared to the highest loss rate for N-packets which is only 3.5%.

This study discourages the use of options for introducing new functions in IP. Extension headers is an alternative format for introducing new functions, which should not cause the same problems with legacy routers. The unused bit in IP can be used for negotiating the use of extension headers.

²Firewall compliancy requires that extension headers are placed after the transport header, which also means that some additional heuristics is needed to indicate their presence.

Acknowledgments

The authors would like to thank David G. Andersen for providing access to the resilient overlay network.

References

- [1] Reference to web page removed for purpose of blind review.
- [2] Ip options selective drop. Cisco whitepaper.
- [3] Resilient overlay networks. <http://nms.lcs.mit.edu/ron/>.
- [4] Steven M. Bellovin. Email correspondence, March 2002.
- [5] Samrat Bhattacharjee, Kenneth L. Calvert, and Ellen W. Zegura. An architecture for active networking. In *Proceedings of High Performance Networking (HPN'97)*, White Plains, New York, USA, may 1997.
- [6] C. J. Bovy, H. T. Mertodimedjo, H. Uijterwaal, and P. Van Mieghem. Analysis of end-to-end delay measurements in internet. In *Proceedings of the Passive and Active Measurements Workshop (PAM2002)*, Ft.Collins, mar 2002.
- [7] S. Deering and R. Hinden. Internet protocol, version 6 (ipv6) specification. Request for Comments 2460, IETF, dec 1998.
- [8] Jean Dickinson Gibbons. *Nonparametric Statistical Inference*, volume 65 of *Statistics, textbooks and monographs*. Marcel Dekker, Inc., New York, 2nd edition, 1985.
- [9] Lars Åke Larzon, Ulf Bodin, and Olov Schelén. Hints and notifications. In *Proceedings of Wireless Communications and Networking Conference, WCNC'02*, Orlando, Florida, March 2002.
- [10] Chiang Lee, Chih-Horng Ke, and Chao-Chun Chen. Improving location management for mobile users with frequently visited locations. *Performance Evaluation*, 43(1):15–38, June 2001.

- [11] T. Nakajima. Nps:user-level real-time network engine on real-time mach. Proceedings in Workshop, 1994.
- [12] Christos Papadopoulos, Guru M. Parulkar, and George Varghese. An error control scheme for large-scale multicast applications. In *Proceedings of IEEE Infocom '98*, volume 3, pages 1188–1196, March 1998.
- [13] V. Paxson, G. Almes, J. Mahdavi, and M. Mathis. Framework for ip performance metrics. Request for Comments 2460, IETF, may 1998.
- [14] Radia J. Perlman and Suchitra Raman. Techniques for making ip multicast simple and scalable. In *Networked Group Communication, NGC'99*, volume 1736 of *Lecture Notes in Computer Science*, pages 270–285, Pisa, Italy, nov 1999. Springer.
- [15] Jon Postel. Internet Protocol. Standard RFC791, STD0005, Internet Engineering Task Force, September 1981.
- [16] David J. Wetherall and David L. Tennenhouse. The ACTIVE IP option. In *Proceedings of the Seventh ACM SIGOPS European Workshop*, Connemara, Ireland, September 1996.