Measuring Networks Using IP Options

Pietro Marchetta, Valerio Persico, Giuseppe Aceto, Alessio Botta, and Antonio Pescapé

Abstract

Injecting purposely-created probing traffic makes it possible to detect the presence and location of performance issues or faults, to reveal the topology of the network, and to investigate related properties. While researchers and network operators still rely on traditional tools (e.g., traceroute or ping) to shed light on the Internet, here we present six novel measurement techniques based on IP options. We show how IP options can still offer unforeseen ways to augment the knowledge about networks, potentially presenting both new threats and new opportunities for different stakeholders.

Introduction

The radically distributed ownership of the Internet dramatically weakens our ability to measure and monitor this global infrastructure. Although many users still rely on traditional tools such as *traceroute* and *ping*, in recent years there have been new measurement techniques exploiting *active probing* [1–4]. This paradigm has profitably been used for studying many hot networking topics: fault diagnosis, troubleshooting, broadband performance evaluation, topology discovery, available bandwidth estimation, and recently, monitoring large-scale events such as Internet outages caused by censorship and natural disasters [5]. Indeed, new approaches using innovative probe packets (hereafter *probes*) have been presented in literature to improve network measurement techniques and tools and overcome their limitations.

In this article, we focus on active measurement techniques leveraging IP options that were introduced in 1981 with the RFC791 to increase the functionalities offered by the network layer of the TCP/IP stack. In more detail, we focus on how the use of IP options in unforeseen ways can reveal supposedly unavailable information or help investigate, mitigate, or solve long-unresolved issues in networking. To this aim, we analyze milestone contributions in literature, and present six novel measurement techniques based on the timestamp and record route options, showing their effectiveness in augmenting the knowledge gathered through active monitoring approaches.

Pietro Marchetta and Valerio Persico are with University of Napoli "Federico II."

Giuseppe Aceto, Alessio Botta, and Antonio Pescapé are with the University of Napoli "Federico II" (Italy) and with NM2 srl

Digital Object Identifier: 10.1109/MNET.2017.1600070NM

IP Options: Background and Use in Internet Measurements

By equipping packets with IP options, the sender asks the devices along the path to take specific actions based on information local to them or carried in the IP option (type and status). IP options do not benefit from full support from network operators and device manufacturers due to security concerns (RFC7126), and they can impact the performance of Internet traffic [6, 7]. More than two dozens IP option types have been defined (RFC891 and related), of which nine have been officially deprecated due to low rates of adoption (RFC6814). Here we focus on the two IP options most adopted and investigated in measurement literature: *record route* and *timestamp*, whose support has been experimentally assessed (Table 1a).

The IP record route option (RR) provides a way to record the route traversed by a packet toward its destination, and represents the first (and only) path-tracing approach included in the Internet standards. When receiving a packet equipped with this option, a network device inserts one of its IP addresses in the option body if room is available. The RR option can contain no more than nine address slots, due to the maximum size of the IP header. Hereafter, X_{RR}, with X in {ICMP, UDP, TCP} stands for a probe equipped with the RR-option.

The IP timestamp option (TS) has different variants, identified by the value of the *flag* field. With flag 0, each device on the path is requested to insert a 32-bit timestamp. With flag 3, four different addresses are prespecified in the option and a device enters a 32-bit timestamp only if it owns the first unmarked address. In all the variants, the overflow field counts the number of devices along the path that could not insert a timestamp due to the lack of space. Hereafter, $X\cdot_{TS0}$ symbolizes a probe with flag 0, while $X\cdot_{TS3|ABCD}$ stands for a probe equipped with flag-3 TS option prespecifying the addresses A, B, C, D.

Several milestone papers in the literature use IP-option-equipped probes for Internet measurements. The most important ones are summarized in Table 1b–g. Indeed, IP options helped researchers to investigate, mitigate, or solve long-unresolved issues related to, e.g., topology discovery (tracing paths, performing alias resolution) and troubleshooting (quantifying deviations from standard forwarding, remotely inferring the status of network devices). Note that in several of these works, IP options are exploited beyond the specific purpose they have been introduced for. In this way they can reveal information that the network administrator did not mean to expose or that a researcher had no other way to access.

IP-Options-Based Internet Measurement Techniques

In this section, inspired by the works presented in Table 1 and the related promising results, we briefly present six novel Internet measurement techniques based on IP options. An overall picture of these contributions and their main findings is reported in Table 2. More details on the techniques are reported in [8–10], where additional information about

Aim	Options	Source reference and description
A. Evaluating the support to IP options.	RR, TS0, NOP	R. Fonseca *et al.*, 2005, "IP Options are Not an Option," Technical Report, Electrical Engineering and Computer Sciences, University of California at Berkeley. The authors quantify the support to IP-options measuring wide-area paths using probes with and without IP options. They find that the core of the network drops very few packets with options, while the vast majority of drops occur at the edge of the network, concentrated in a minority of the ASs.
	TS3|AAAA	W. de Donato *et al.*, 2012, "A Hands-On Look at Active Probing Using the IP Prespecified Timestamp Option," *Proc. Passive and Active Measurement Conf. 2012*. The authors investigate the responsiveness of hosts to several kinds of probes, with and without the IP prespecified timestamp option. They find that IP options may significantly impact the responsiveness to the probes, also observing non RFC-compliant behaviors in a non negligible amount of replies.
B. Enhancing path tracing.	RR	R. Sherwood and N. Spring, 2006, "Touring the Internet in a TCP Sidecar," *Proc. 6th ACM SIGCOMM Conf. Internet Measurement 2006*. Pioneering work demonstrating the utility of IP options in Internet measurements. By using Traceroute and the RR option, the authors gather additional information about the traversed paths, potentially identifying load balancers, anonymous routers, multiple interfaces of the traversed routers, etc.
	RR	R. Sherwood *et al.*, 2008, Discarte: A Disjunctive Internet Cartographer," *Proc. ACM SIGCOMM 2008 Conf. Data Communication*. The authors face the non-trivial task of aligning Traceroute traces with those obtained with the RR option, by adopting disjunctive logic programming. The objective is to cross-validate results obtained by the two techniques to improve the accuracy of the inferred topology. Although of great interest, the proposed approach is computationally complex and hard to replicate.
C. Reverse path tracing.	RR, TS3|A	E. Katz-Bassett *et al*., 2010, "Reverse Traceroute," *Proc. 7th USENIX Conf. Networked Systems Design and Implementation.* Network path tracing techniques such as Traceroute provide no information on how the traffic is routed along the reverse path, i.e. the path connecting the destination back to the source. The authors propose an approach based on spoofed probes, multiple vantage points, and both the RR and TS options to trace the reverse path. Successfully tracing the reverse path is extremely helpful to troubleshoot, infer AS-level connections, and measure the properties of the network links.
D. Assessing link symmetry.	RR, TS3|ABA	H. V. Madhyastha, 2008, "An Information Plane for Internet Applications." Ph.D. thesis presenting the use of options to identify which links are traversed symmetrically. More specifically, the author adopts the prespecified TS option to understand if a link connecting two subsequent nodes A and B on a network path is traversed symmetrically by the traffic exchanged between a source S and B. Evidence is collected about the traversal of A–B link crafting probes from S and purposely choosing the ordering of prespecified addresses in the TS option flag 3.
E. Inference of router statistics.	TS3|ABCD	A. D. Ferguson and R. Fonseca, 2010, "Inferring Router Statistics with IP Timestamps," *Proc. ACM CoNEXT Student Wksp*. The authors uncover bounds on the rate of UDP traffic carried by Cisco 3600-series routers and the start and end of multicast traffic carried by 6500-series Catalysts, not requiring any control on the tested devices. This work demonstrates that using IP options may provide additional information on the status of routers, e.g. uncovering CPU-intensive operations like forwarding multicast traffic.
F. Detection of non destinationbased forwarding.	RR	T. Flach *et al.*, 2012, Quantifying Violations of Destination-Based Forwarding on the Internet," *Proc. 2012 ACM Internet Measurement Conf.* Each router is supposed to select the next hop on the path toward the destination exclusively based on the destination of the packet. However, increasingly common mechanisms such as load balancing, MPLS, and default routing represent a deviation from this paradigm. The authors quantify such deviation exploiting the RR option, discovering that ~29% of observed routers violate the destination-based forwarding.
G. Alias resolution.	TS3|AAAA, TS3|ABCD	J. Sherry *et al.* 2010, "Resolving IP Aliases with Prespecified Timestamps.," *Proc. 10th ACM SIGCOMM Conf. Internet Measurement*. The authors address the alias resolution problem, i.e. the problem of gathering under a unique identifier the address part of the same network device. In fact, failing the association between a device and the related IP addresses significantly impacts the accuracy of the inferred router-level topology. A technique based on the prespecified variant of the TS option is proposed, which is able to identify a significant amount of addresses in alias not recognized by state-of-the-art techniques.

TABLE 1. Using IP options in Internet Measurements: literature review.

the experiments done can also be found. We purposely present techniques with completely different goals to demonstrate the versatility of IP options as a tool to empower active measurements in general. Since IP options are not supported universally, all the techniques presented are designed to be complementary to the state-of-the-art approaches (those not relying on IP options).

To foster further experimentation with probes equipped with IP options, we have released the techniques as either stand-alone tools or programming libraries for crafting packets (see http://traffic.comics.unina.it).

Alternative Path Tracing

Traceroute is the *de facto* standard tool to trace the network path — i.e. discover the hops traversed — toward a destination. Operators and researchers heavily rely on it for several aims, e.g. to locate failures, measure latencies, etc. This solution is known to provide incomplete or inaccurate information due to several reasons including hidden routers [9, 11], ICMP filtering and rate-limiting, etc. Several optimizations and variants have been proposed over the years that are more robust, accurate, and efficient than the original version (Table 1b). However, the very basic mechanism, i.e. limiting the TTL of the injected packets to elicit ICMP time exceeded messages from the traversed routers, remained unchanged since its introduction in 1989. Hence, there is no chance for traceroute to discover devices that do not decrement the TTL of the forwarded packets. Moreover, when a traversed device resets the TTL to a high value, the probe packet can safely reach the destination, causing the last portion of the path to be totally invisible.

IP option	Application	Description	Our main findings
Malformed RR, Malformed TS0	Alternative path tracing	Discovering the IP addresses of the traversed routers by not relying on the TTL field.	Reporting interfaces and routers not listed by state-of-the-art solutions.
RR	Reducing redundancy among vantage points	Grouping the vantage points whose traffic reaches a network destination through the same ingress point of the targeted AS.	Discovering 99% and 98.5% of nodes and links of the topology of interest with a reduction of 35% of measurement traffic.
TS0	Detecting middleboxes	Detecting devices managing IP options but not decrementing the TTL.	Detecting middleboxes not revealed by state-of-the-art solutions.
TS0	Path length estimation	Counting the number of routers traversed by the traffic sent to a destination.	Providing consistent results even when the standard TTL-based solution provides misleading results.
TS3|AAAA, TS3|ABCD	Alias resolution	Grouping under the same identifier the IP addresses owned by the same network device.	Solving the alias resolution even for unresponsive addresses for which state-of-the-art solutions fail.
TS3|AAAA	Network device fingerprinting	Disclosing information about the kind of network device.	Distinguishing different router brands, based on specific peculiarities when managing the prespecified TS option.

TABLE 2. Six applications of measurement techniques presented in this paper: description and main findings.

Router family	Interfaces providing *x* timestamps					Total interfaces (100%)
	x = 0	*x* = 1	*x* = 2	*x* = 3	*x* = 4	
Cisco	47.6%	42.1%	4.2%	–	6.1%	16,149
Juniper	1.2%	0.2%	0.1%	–	98.6%	3,532
Others	88.5%	–	0.6%	–	11.0%	720

TABLE 3. Fingerprinting network devices. Routers of different brands provide different amounts of timestamps.

Inspired by pioneering works (Table 1b and Table 1c), we developed a novel technique for tracing Internet paths, exploiting the ICMP parameter problem message instead of the ICMP time exceeded message, thus radically departing from previous variations on traceroute. According to RFC1122 and RFC791, routers and hosts send back an ICMP parameter problem message in different cases (e.g., incoming packet discarded and no other ICMP message covers the problem, option check failed, overflow field of the TS option overflows, etc.). Our technique causes the packet to be discarded and an ICMP parameter problem message to be sent to the source from a specific hop (i^{th} in the following). To achieve this goal, we purposely craft malformed RR and TS (flag 0) options according to three different methods:

- Cut record route (CRR): Setting the *length* field of the RR-option to have in the option body enough space for $i - 1$ IP addresses, and only 3 B available for the i^{th} one.
- Cut timestamp (CTS): Setting the *length* field of the TS-option to have in the option body enough space for $i - 1$ timestamps, and only 3 B available for the i^{th} one.
- Overflow in overflow (OV2): Setting the *overflow* and *pointer* fields of the TS-option to make the option body appear as full, and the overflow to cause the *overflow in overflow* condition after i increments.

We point the reader to [8] for more details on the technique, and discuss hereafter some notable experimental results.

We investigated the effectiveness of these methods performing a large-scale measurement campaign [8]:

•The upper bound of devices that can be traced by our technique is shown in Fig. 1a, as the ratio of devices generating ICMP parameter problem over the ones managing the non-malformed option, for both TS and RR options. Malforming the TS (RR) option triggered ICMP parameter problem replies from 61 percent (62.5 percent) of the devices managing the option in each path, on average.

•The pairwise comparison of the effectiveness of the three methods is reported in Fig. 1b. OV2 proved to be the most effective, reporting most of the interfaces and routers collected by CRR and CTS. In more detail, OV2 alone is able to discover 96.5 percent (97.0 percent) of the interfaces (routers) while OV2 and CTS make up the best pair, discovering 99.6 percent (99.0 percent) of the interfaces (routers) together, with respect to those discovered by the three methods together.

•Compared to the Multipath Detection Algorithm (MDA) of Paris-traceroute [3], OV2 was able to report additional interfaces and routers, the gain being higher for paths that traverse devices that reset the TTL field. In particular, OV2 listed at least one interface (router) not reported by MDA in 31.3 percent (26.7 percent) of the scanned paths.

These results confirm that our alternative path-tracing technique provides additional and complementary information with respect to traditional approaches, for instance, in the presence of devices purposely configured to limit path tracing in corporate networks [12].

Reducing Redundancy among Vantage Points

In the previous section we used a malformed RR option to obtain more information regarding a path. Here we show how the regular RR option can be leveraged to reduce measurement overhead by detecting and avoiding redundancy among vantage points (inter-monitor). Discovering the topology of a remote AS network without having access to or control over it is challenging [1, 4]. A straightforward bruteforce approach is launching traceroute toward destinations inside the AS of interest from several vantage points. However, this approach proved to be very inefficient [4], e.g. using several vantage points instead of a single one is unlikely to uncover more internal nodes or links in networks with only one logical ingress point.

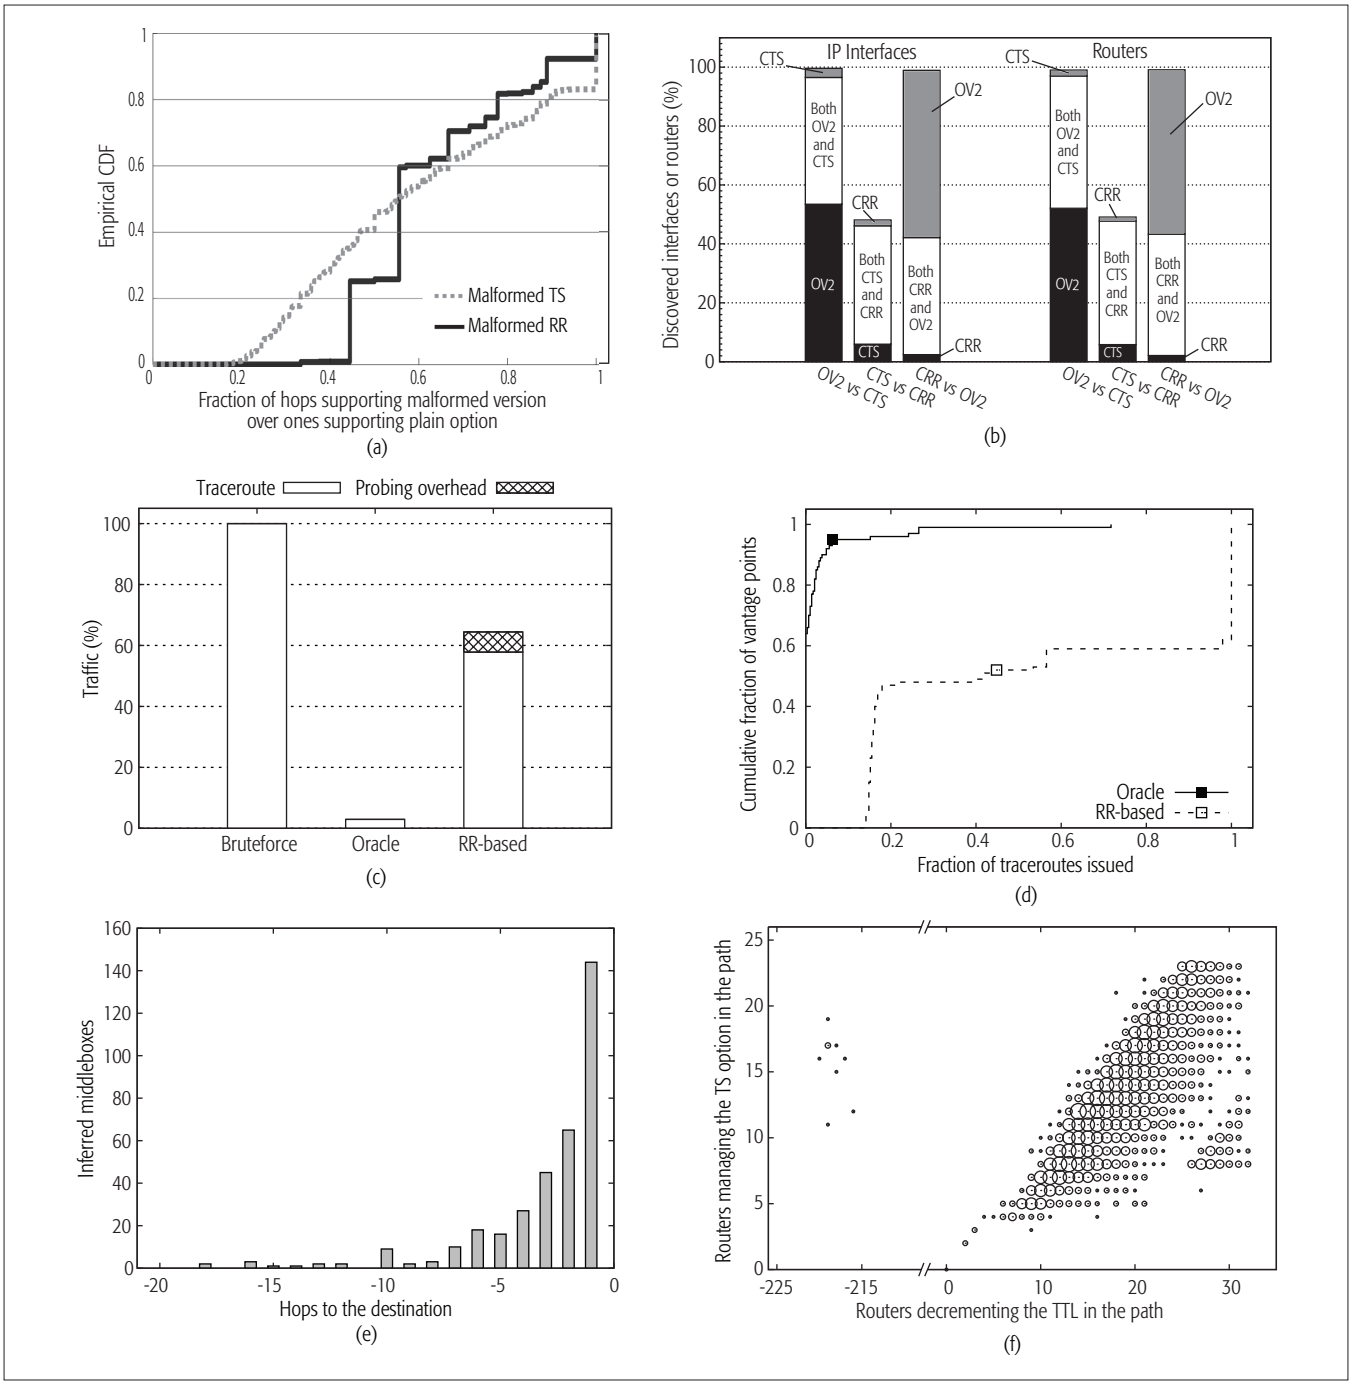

FIGURE 1. Results achieved with IP-options based measurements: a) devices generating ICMP Parameter Problem messages when managing malformed RR and TS options; b) pairwise comparison of the proposed methods at interface and router level (percentages are calculated over all the interfaces and routers discovered by the three methods together); c) total induced load; d) traceroutes issued per VP; e) middlebox detection and location; f) path length estimation — TTL-based vs TS-based approach (marker radius logarithmically grows with occurrences).

Approaches in the literature tried to mitigate this problem [4], e.g. coordinating the different vantage points to stop tracing when a common node is encountered along the path toward the same destination [4], or using data from previous topology mappings to assign destinations to vantage points avoiding tracing the same common path [13].

Our technique based on IP options aims to cluster vantage points that reach a destination inside the AS of interest through the same ingress point. Specifically, we use the RR option to identify vantage points whose paths toward a given destination D converge within the first nine hops and before entering the targeted network. To do this, we:

1) Launch a $ICMP_{RR}$ probe toward D from each vantage point.
2) Identify addresses belonging to the target AS using IP-to-AS mapping and remove them from the collected RR traces.
3) Build a graph $G(V, E)$ where V is the set of addresses contained in the RR traces and E represents adjacent addresses in each RR sequence.
4) Augment G with additional edges between

According to the RFC3234, a middlebox is any intermediary box performing functions apart from standard functions of an IP router on the data path between a source host and destination host. Detecting middleboxes is at the same time of utmost importance when troubleshooting network paths, and non-trivial, as they may be configured to not appear in traceroute traces.

each originating vantage point and the first address in the related RR sequences, and with the vantage points appearing as isolated nodes in *G* due to lack of received replies from *D*.

5) Analyze the graph *G*, where each connected component is a tree-like structure, in which the leaves are vantage points and the root represents a convergence point located outside the targeted network. Redundant vantage points are identified as part of the same connected component.

To verify the effectiveness of this technique, we performed an experimental campaign on the GARR network (AS137). This network has 12 different ingress points. We randomly selected 291 destinations in 20 different subnets, steadily responsive to ping. We used 75 Planetlab nodes to issue $ICMP_{RR}$ followed by traceroute toward each destination *D*. We compared our solution with two alternatives: *bruteforce*, using all the vantage points to target each destination, and *oracle*, querying a hypothetical omniscient oracle able to inspect the result of a traceroute *before* its collection (since discovering all the links of the targeted topology is a set-covering problem known to be NP-complete, the oracle adopts a greedy approach).

Comparing with bruteforce approach, we experimentally found that the oracle only requires 2.8 percent of the probes to reconstruct the same topology, while the proposed technique is able to discover 99 percent and 98.5 percent of the nodes and links, respectively, only injecting about 64 percent of the traffic (Fig. 1e). In terms of traceroutes, the oracle shows that theoretically each vantage point should issue no more than 71 percent of the bruteforce approach (Fig. 1d). With our technique, each vantage point issued only 54 percent of the traceroutes, on average.

In conclusion, the technique based on IP options uncovers the targeted network with accuracy comparable to the bruteforce approach, but significantly reducing the overall load imposed on the network. Path-related measurements (like the ones in the previous and next sections) can be performed with much less probing overhead leveraging the knowledge gained through such a technique.

Detecting Middleboxes

According to RFC3234, a middlebox is any intermediary box performing functions apart from standard functions of an IP router on the data path between a source host and destination host. Detecting middleboxes is at the same time of utmost importance when troubleshooting network paths, and non-trivial, as they may be configured to not appear in traceroute traces. Active probing methods have proven able to reveal some classes of middleboxes. Detal *et al.* [11] proposed an extension of traceroute that sends IP packets containing TCP segments with limited TTL values, and analyzes the packet encapsulated in the elicited ICMP time exceeded message in search of any modifications, potentially revealing the presence of middleboxes. As reported in Table 1b, Sherwood *et al.* proposed an option-based approach, using the RR option to possibly identify middleboxes. This approach proved effective but suffers from the limited exploring range allowed by the RR option.

We propose a novel technique using the TS option to mitigate this limitation. Such a technique counts both the devices managing the TS option and those decrementing the TTL in a given subpath. For each subpath with more routers managing the option than those decrementing the TTL, it infers the presence of middleboxes. In the measurement process, first the path to the destination is traced with a modified traceroute injecting UDP_{TS0} packets with increasing TTL values. Then, a procedure is applied to detect and locate middleboxes by inspecting each possible portion of the paths containing more routers managing the option than those decrementing the TTL.

This technique can give either a hint or strong evidence of the presence and location of a middlebox, due to implementation differences in the processing of options. Moreover, this technique is complementary to state-of-art solutions considered in this article, such as tracebox [11], that can be easily extended with it. We refer to [9] for further details on the technique and add notable results hereafter.

To evaluate this solution, we traced the paths toward 25K stably responsive random destinations with our TS-option-equipped traceroute (flag 0), observing more than 45K addresses. We found that in about 0.2 percent of the two-hop subpaths there are more than two devices managing the TS option between the two consecutive routers decrementing the TTL, which is clear evidence of the presence of middleboxes. These cases have been confirmed with measurements from several vantage points and are referred to as *unique two-hop subpaths* (i.e. distinct IP address couples). Surprisingly, we detected subpaths containing up to four consecutive invisible devices, which means that entire portions of the network may be invisible to state-of-the-art approaches. Figure 1c shows that these devices appear few hops before the targeted destinations: a likely position for gateways, thus supporting the validity and real-world usefulness of the proposed technique.

Path Length Estimation

In the previous section we used TS0 to detect middleboxes. Here we show experimental results from leveraging the same option to improve estimates of a basic metric associated with a network path, i.e. its *length*, in terms of the number of hops or devices that are traversed. An estimation of this metric provides rough information when predicting latency, its variation in time is evidence of possible routing changes, and it helps reduce the probing overhead when tracing the path [12]. A common approach for estimating the path length is inferring the number of devices that have decremented the TTL of the probe along the forward path [12]. Such an approach can be inaccurate due to a number of issues mainly related to the different TTL processing actually implemented by routers on the Internet [9, 11].

Our novel technique is based on UDP_{TS0} probes requesting each traversed router to either

insert a timestamp into the option body or increment the overflow field by one. From the ICMP port unreachable reply, we count how many routers managed the option along the forward path, summing the number of timestamps inserted and the number of overflow increments. Such a technique can be used simultaneously with the TTL-based approach.

We estimated the path length with both techniques toward 20K destinations in 3,732 ASs from our university laboratory. Figure 1f compares the path length estimations obtained. Although the option-based technique tends to underestimate the length, there are observed paths containing more devices managing the option than those decrementing the TTL. A strong correlation between the two results is evident: the Pearson coefficient is about 0.81. Relying on the classic approach may also lead to impossible or misleading results with negative path-length estimation. This may happen when a router along the path overrides the TTL field. In these cases, the presented technique provided consistent results, thus enhancing and augmenting state-of-the-art techniques.

Alias Resolution

Alias resolution is the process of identifying interfaces, i.e. IP addresses, belonging to the same router. It is required to convert IP-level topology data discovered by traceroute into a more useful router-level topology, and constitutes a critical step in producing Internet topology maps. Focusing on active probing, alias resolution techniques can be grouped into distinct families based on their rationale: source address (e.g., *Mercator*, *Iffinder*, *PalmTree*); shared IP ID counter (e.g., *Ally*, *Radargun*, *Midar*); and IP-option based (e.g., *Motu*, implementing the approach presented in Table 1g). Active techniques are usually accurate, but *incomplete* because many routers do not respond to the probes.

Pythia, the novel alias resolution technique we have proposed, is designed to inject TS-option-equipped UDP probes into the network and to inspect the ICMP destination unreachable returned. It is designed to work with routers providing four timestamps when probed with $UDP_{TS3|AAAA}$, where A is an address of the device, i.e., with any-interface stamping routers. Pythia performs two phases: preliminary test and alias resolution. The first phase is aimed at splitting the set of candidate addresses into three subsets:

- *Unresponsive* addresses
- *Compliant* addresses, providing four timestamps
- *Non-compliant* addresses, providing less than four timestamps

Non-compliant addresses are discarded at the end of the first phase. The second phase consists in probing each compliant address A with probes like $UDP_{TS3|ABCD}$, where B, C, and D belong to either compliant or unresponsive addresses. The second phase leverages the fact that the router owning the compliant address A, will also provide timestamps for B, C, and D if these addresses belong to it. Note that only the router owning A is allowed to insert its timestamps, because of the destination address A also being pre-specified in the first position. Moreover, when a timestamp out of the four requested is not provided (e.g. the third one), the alias inference process takes into account that timestamps cannot be inserted out of order (i.e. the fourth prespecified address may still be an alias) and schedules the following probes accordingly. We point the reader to [10] for more details on the technique, and report hereafter some notable experimental results.

Alias resolution is the process of identifying interfaces — i.e., IP addresses — belonging to the same router. It is required to convert IP-level topology data discovered by traceroute into a more useful router-level topology, and constitutes a critical step in producing Internet topology maps.

We found that Pythia performs better than other approaches according to several metrics:

- Pythia has a higher applicability, making it 4.5 and 11 times more applicable than Motu and Palmtree, respectively.
- Pythia has a notable gain in the percentage of IP pairs properly aliased/dealiased, which grows from 8.5 percent (3.4 percent) to 37.8 percent comparing Motu (Palmtree) with Pythia.
- At the same time, IP pairs wrongly aliased/dealiased proved to increase by less than 1 percent.

Pythia was able to classify 46.7 percent of all the pairs not classified by Motu, taking the correct decision in 97.9 percent of these cases. Compared to the state of the art, Pythia shows interesting advantages:

- Different from all the other techniques, Pythia is able to tell if a given address B is in alias or not with another address A even if B does not reply to active probing at all.
- Besides the probing of the preliminary phase (showing linear complexity), Pythia requires a single probe to infer if two addresses are in alias or not, whereas other techniques require at least two or three probes.
- To the best of our knowledge, Pythia is the only active probing technique able to identify up to four addresses in alias with a single probe. To reach the same result, traditional techniques would require testing of six different pairs of addresses.

Network Device Fingerprinting

In the previous section we used a TS3|AAAA probe to perform alias resolution; here we show how the same option can also be exploited to infer information about a device itself (device fingerprinting). Device fingerprinting consists in collecting information about a physical device connected to the Internet and can be adopted for troubleshooting and root-cause analysis of performance issues and outages. Fingerprinting is usually performed remotely, thus requiring no modification of the fingerprinted device, and often without its cooperation. A common approach consists in directly targeting the network device and inspecting the collected replies, exploiting implementation-specific and configuration-specific characteristics. Several solutions have already been proposed, relying on different fields and layers of the TCP/IP stack. For instance, devices can be grouped into categories by:

- Estimating which initial TTL value is set in the packets they generate [14].

On the one hand, techniques based on IP options can markedly complement the state of the art in investigating and characterizing the Internet. On the other hand, their usage often goes beyond the purpose they have been initially introduced for, revealing details not intended to be exposed by network administrators.

- Measuring the machine's clock drift sampled at the TCP layer [15].
- Grabbing device brand and model information directly from the presentation banner showed by the device.

IP-option-based approaches can complement the state of the art in grouping devices according to their brand. We have found that different devices act diversely when receiving a probe $UDP_{TS3|AAAA}$ which pre-specifies four times an address A owned by the targeted device. For instance, we verified that Juniper routers provide four timestamps in most cases while Cisco routers tend to provide no more than two timestamps.

We evaluated this fingerprinting technique using ground truth data collected with IGMP probing [1]. We inferred the router brand associated to 20.4K IP addresses belonging to 5K routers inspecting the supported version of the Distance Vector Multicast Routing Protocol (DVMRP) returned by `IGMP ASK_FOR_NEIGHBORS` message. Results are reported in Table 3. Cisco routers either ignored the TS option or recorded only one or two timestamps of the four requested (93.9 percent), while about 98.6 percent of Juniper routers recorded all the requested timestamps.

In conclusion, an IP address providing four timestamps belongs to a Juniper router in 77 percent of the cases. On the other hand, one providing no more than two timestamps belongs to a Cisco router in 95.6 percent of cases. Note that the proposed technique is less intrusive than existing *banner grabbing* ones, as it does not require an active TCP connection, and can be used simultaneously with other previously cited approaches.

Discussion and Conclusion

There is ongoing interest in using IP options for augmenting active measurements, including a rich discussion about the extent of their support and the related risks and opportunities. In this article we described six novel techniques that further unveil their potential in tasks such as identifying devices' vendor, tracing paths and calculating their length, detecting devices that implement non-standard behaviors, reducing measurement overhead, or performing alias resolution.

On the one hand, techniques based on IP options can markedly complement the state of the art in investigating and characterizing the Internet. On the other hand, their usage often goes beyond the purpose they have been initially introduced for, revealing details not intended to be exposed by network administrators. For these reasons, we believe that IP options represent both a valuable tool for network measurements and a research topic still open.

Acknowledgment

We thank the anonymous reviewers for their helpful comments on this article. This work is partially funded by art. 11 DM 593/2000 for NM2 srl (Italy).

References

[1] P. Marchetta *et al.*, "Topology Discovery at the Router Level: A New Hybrid Tool Targeting ISP Networks," *IEEE JSAC*, vol. 29, no. 9, 2011, pp. 1776–87.
[2] Y. Zhang, Z. M. Mao, and M. Zhang, "Effective Diagnosis of Routing Disruptions from End Systems," *Proc. 5th USENIX Symp. Networked Systems Design and Implementation*, NSDI'08, Berkeley, CA, USA, USENIX Association, 2008, pp. 219–32.
[3] B. Augustin, T. Friedman, and R. Teixeira," Measuring Multipath Routing in the Internet," *IEEE/ACM Trans. Net.*, vol. 19, June 2011, pp. 830–40.
[4] B. Donnet *et al.*, "Efficient Algorithms for Large-Scale Topology Discovery," *SIGMETRICS*, 2005, pp. 327–38.
[5] L. Quan, J. Heidemann, and Y. Pradkin, "Trinocular: Understanding Internet Reliability Through Adaptive Probing," *SIGCOMM Computing Commun. Rev.*, vol. 43, Aug. 2013, pp. 255–66.
[6] P. Fransson and A. Jonsson, "End-to-End Measurements on Performance Penalties of IPv4 Options," *IEEE Global Telecommun. Conf., 2004*, vol. 3, 2004, pp. 1441–47.
[7] A. Medina, M. Allman, and S. Floyd, "Measuring the Evolution of Transport Protocols in the Internet," *SIGCOMM Computer Commun. Rev.*, vol. 35, Apr. 2005, pp. 37–52.
[8] P. Marchetta *et al.*, "Experimenting with Alternative Path Tracing Solutions," *20th IEEE Symp. Computers and Commun. (IEEE ISCC 2015)*, July 6–9, 2015, pp. 81–86.
[9] P. Marchetta and A. Pescapé, "Drago: Detecting, Quantifying and Locating Hidden Routers in Traceroute IP Paths," *INFOCOM*, 2013, pp. 3237–42.
[10] P. Marchetta, V. Persico, and A. Pescapé, "Pythia: Yet Another Active Probing Technique for Alias Resolution," *Proc. 9th ACM Conf. Emerging Networking Experiments and Technologies*, *CoNEXT '13*, New York, NY, USA, ACM, 2013, pp. 229–34.
[11] G. Detal *et al.*, "Revealing Middlebox Interference with Tracebox," *Proc. 2013 Conf. Internet Measurement Conf.*, IMC '13, New York, NY, USA, ACM, 2013, pp. 1–8.
[12] T. Moors, "Streamlining Traceroute by Estimating Path Lengths," *Proc. IEEE Wksp. IP Operations and Management*, 2004, Oct 2004, pp. 123–28.
[13] G. Baltra, R. Beverly, and G. G. Xie, "Ingress Point Spreading: A New Primitive for Adaptive Active Network Mapping," *Passive and Active Measurement*, Springer, 2014, pp. 56–66.
[14] Y. Vanaubel *et al.*, "Network Fingerprinting: TTL-Based Router Signatures," *Proc. 2013 Conf. Internet Measurement Conf.*, ACM, 2013, pp. 369–76.
[15] T. Kohno, A. Broido, and K. C. Claffy, "Remote Physical Device Fingerprinting," *IEEE Trans. Dependable and Secure Computing*, vol. 2, no. 2, 2005, pp. 93–108.

Biographies

Pietro Marchetta (pietro.marchetta@unina.it) is a post doc in the Department of Electrical Engineering and Information Technology, University of Napoli Federico II. He has a Ph.D. in computer engineering from the University of Napoli Federico II. His work focuses on measurement and monitoring of network paths and topologies. He was awarded the ACM SRC at SIGCOMM 2012 and the Best Student Paper at CoNEXT 2013.

Valerio Persico [M] (valerio.persico@unina.it) is a post doc in the Department of Electrical Engineering and Information Technology, University of Napoli Federico II. He has a Ph.D. in computer engineering from the University of Napoli Federico II. His work focuses on measurement and monitoring of cloud network infrastructures. He was the recipient of the best student paper award at ACM CoNext 2013.

Giuseppe Aceto (giuseppe.aceto@unina.it) is a post doc in the Department of Electrical Engineering and Information Technology, University of Napoli Federico II. He has a Ph.D. in telecommunication engineering from the University of Napoli Federico II. His work is in measurement and monitoring of network performance and security, with a focus on censorship. He was the recipient of a best paper award at IEEE ISCC 2010.

Alessio Botta (a.botta@unina.it) is an assistant professor in the Department of Electrical Engineering and Information Technology, University of Napoli Federico II. He has a Ph.D. in computer engineering from the University of Napoli Federico II. He works on measurement and monitoring of heterogeneous networks. He was the recipient of a best paper award at IEEE ISCC 2010.

Antonio Pescapé [SM] (pescape@unina.it) is a full professor of computer engineering at the University of Napoli Federico II. His work focuses on Internet technologies, more precisely on measurement, monitoring, and analysis of the Internet. He has co-authored more than 200 conference and journal papers and is the recipient of a number of research awards.