A Hands-on Look at Active Probing using the IP Prespecified Timestamp Option

Walter de Donato, Pietro Marchetta, Antonio Pescapè

Department of Computer Engineering and Systems

University of Napoli "Federico II"

Agenda

- Background and motivation
- Contribution
  - Large scale active probing campaign
  - Responsiveness analysis
  - RFC-compliance analysis
  - Discussion on support of Ts_p-based techniques
- Conclusion

Agenda

- **Background and motivation**
- Contribution
  - Large scale active probing campaign
  - Responsiveness analysis
  - RFC-compliance analysis
  - Discussion on support of Ts_p-based techniques
- Conclusion

Network measurements

- **Essential for network monitoring and management**
  - **Quality of service**
  - **Network planning**
  - **Troubleshooting**
  - **. . .**
- **Active probing**
  - **Injects test packets (aka probes) into the network**
  - **Infers network characteristics from collected replies**

Active probing and IP options

- IP options are often considered useless for active probing
  - usually filtered, poorly implemented or not widely supported [1, 2]
- Recent works reconsidered the utility of IP Timestamp option
  - Reverse traceroute [3]
  - Alias resolution technique proposed by Sherry et al. [4]

[1] Fonseca et al., *Ip options are not an option*. Technical report, 2005
[2] Medina et al., *Measuring the evolution of transport protocols in the internet*. SIGCOMM Comput. Commun. Rev. 35, 2005
[3] Sherry et al., *Resolving ip aliases with prespecified timestamps*. IMC'10, New York, NY, USA, 2010
[4] Katz-Bassett et al., *Reverse traceroute*. NSDI'10, San Jose, CA, USA, 2010

Active probing and IP options

- IP options are often considered useless for active probing
  - usually filtered, poorly implemented or not widely supported [1, 2]
- Recent works reconsidered the utility of IP Timestamp option
  - Reverse traceroute [3]
  - Alias resolution technique proposed by Sherry et al. [4]

How is the IP Timestamp option managed in practice?

How much is it supported?

[1] Fonseca et al., *Ip options are not an option*. Technical report, 2005
[2] Medina et al., *Measuring the evolution of transport protocols in the internet*. SIGCOMM Comput. Commun. Rev. 35, 2005
[3] Sherry et al., *Resolving ip aliases with prespecified timestamps*. IMC'10, New York, NY, USA, 2010
[4] Katz-Bassett et al., *Reverse traceroute*. NSDI'10, San Jose, CA, USA, 2010

IP Timestamp option

- **IP option type 68**
  - **allows forwarding routers to add timestamps into the option data field**
- **RFCs 781 and 791 define it along with three variants**
  - **TS_o → each router forwarding the packet should add a timestamp, if enough space is available**
  - **TS_i → a (IP, timestamp) couple should be added**
  - **TS_p → the sender requests a timestamp for up to four "prespecified" IPs**

We focus our attention on the TS_p variant exploited in [3,4]

[3] Sherry et al., *Resolving ip aliases with prespecified timestamps*. IMC'10, New York, NY, USA, 2010
[4] Katz-Bassett et al., *Reverse traceroute*. NSDI'10, San Jose, CA, USA, 2010

IP prespecified Timestamp option

- **TSp setup on the originating host**
  - **up to four (IP, timestamp=0) records**
  - **pointer set to the first record**

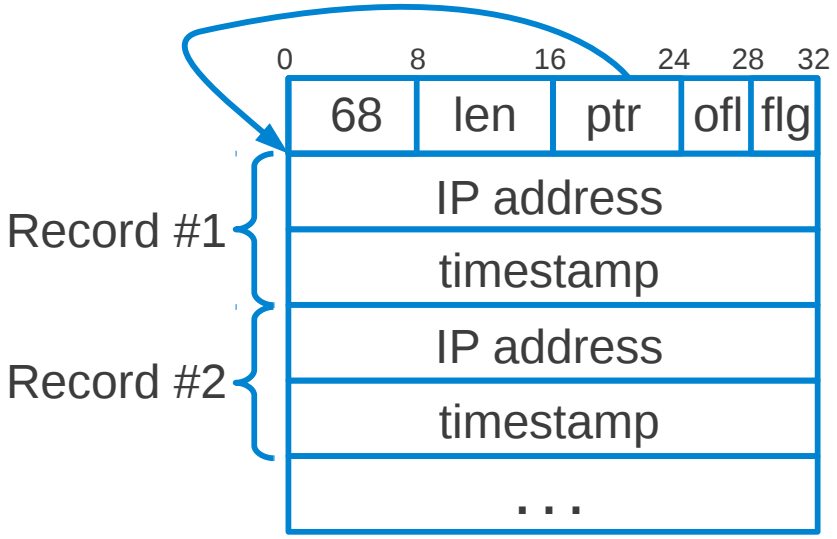

- **A forwarding router when receives a TSp probe**
  - **if the pointed record contains its own IP address**
    - **stamps the pointed record**
      - **Standard format → time in milliseconds since midnight UT**
      - **Non-standard format → most significant bit set to one**
    - **sets the pointer to the next record**
    - **if no more space is available → increments the overflow field**

Agenda

- Background and motivation
- **Contribution**
  - **Large scale active probing campaign**
  - Responsiveness analysis
  - RFC-compliance analysis
  - Discussion on support of Ts_p-based techniques
- Conclusion

Methodology (1/2)

- Target IPs extracted from a complete Archipelago [5] cycle
  - filtering non-routable addresses → 1.7 million destination IPs
- Each destination D targeted with different probes
  - ICMP ECHO_REQUEST → ICMP ECHO_REPLY
  - UDP towards unused port → ICMP PORT_UNREACHABLE
  - TCP SYN towards unassigned well-known port → TCP RESET
  - SKIP message (i.e. obsolete protocol)→ ICMP PROTOCOL_UNREACHABLE
- Each probe is sent twice with and without TS_p option
  - the option is crafted prespecifying D four times → (D|DDDD)
  - retransmissions to deal with congestion and rate limiting policies

[5] k. claffy, Y. Hyun, K. Keys, M. Fomenkov, and D. Krioukov, “Internet Mapping: from Art to Science", CATCH '09

Methodology (2/2)

- ICMP$_p$ and TCP$_p$ probes
  - the Timestamp option is extracted from the IP header of the reply
  - ICMP$_p$ and TCP$_p$ probes return the option as affected by the forward and reverse path
- UDP$_p$ and SKIP$_p$ probes
  - the Timestamp option is extracted from the IP header of the original packet encapsulated as payload in the ICMP reply
  - UDP$_p$ and SKIP$_p$ probes return the option as affected by the forward path only

Active probing campaign

- **Period:** 16 - 20 June 2011
- **Two vantage points**
  - Napoli, Italy (NA)
  - Louvain-la-Neuve, Belgium (LLN)

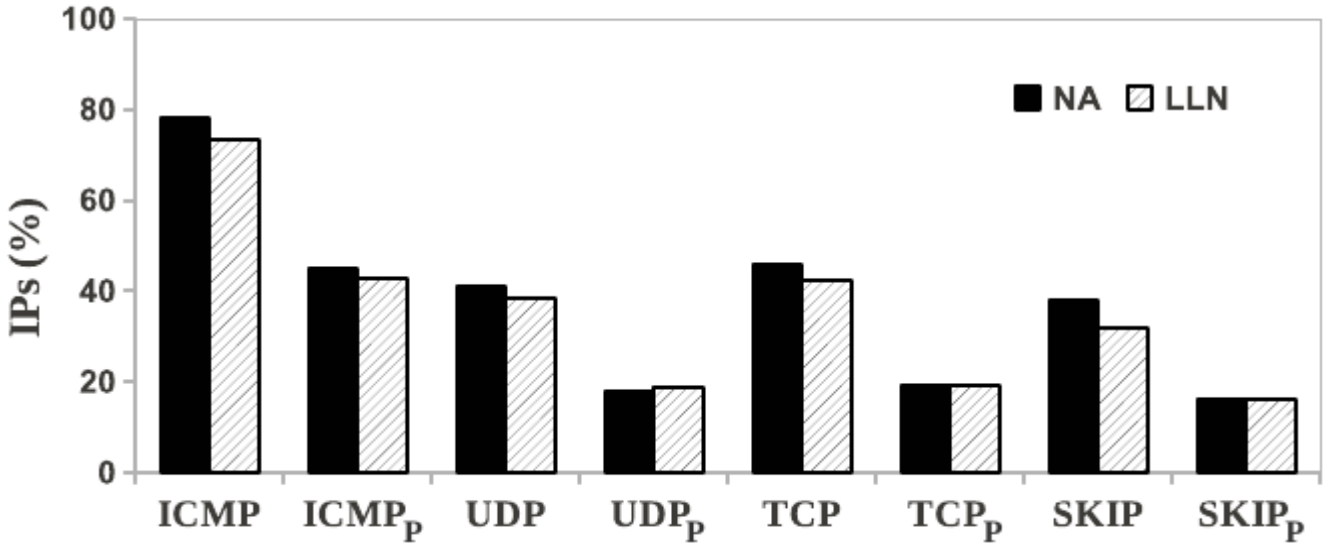

Similar results (like responsiveness profiles) were detected

Agenda

- Background and motivation
- **Contribution**
  - Large scale active probing campaign
  - **Responsiveness analysis**
  - RFC-compliance analysis
  - Discussion on support of Ts_p-based techniques
- Conclusion

Responsiveness to probes (1/2)

- **Impact of the TS_p option**
  - **A** → IPs responsive to normal probes (without TS_p option)
  - **B** → IPs responsive to probes carrying TS_p option
  - **C** → as B but preserving the option in their replies

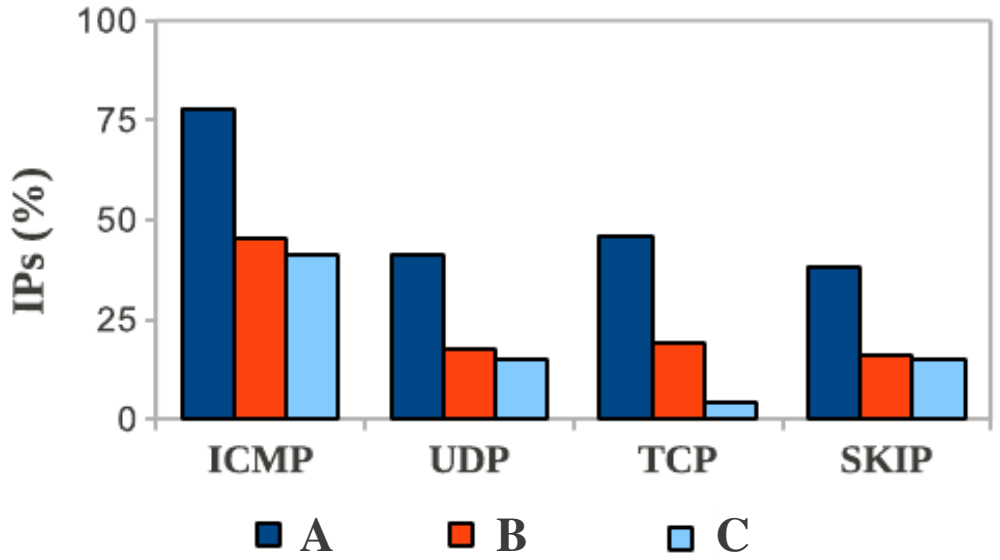

Responsiveness to probes (1/2)

- Impact of the TS_p option
  - A → IPs responsive to normal probes (without TS_p option)
  - B → IPs responsive to probes carrying TS_p option
  - C → as B but preserving the option in their replies

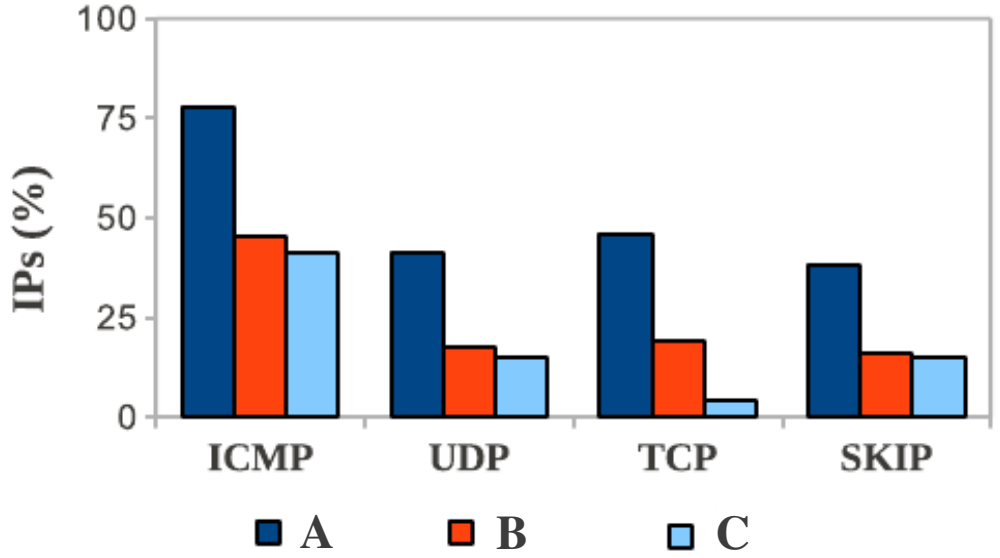

Probes	A → C
ICMP	- 37.4 %
UDP	- 26.4 %
TCP	- 42.5 %
SKIP	- 18.9 %

In the following we consider only the replies preserving the option

Responsiveness to probes (2/2)

- **Classification of destination IPs** (from the Archipelago dataset)
  - **Pathending** → appear in the dataset only as traceroute destinations
  - **Router** → otherwise

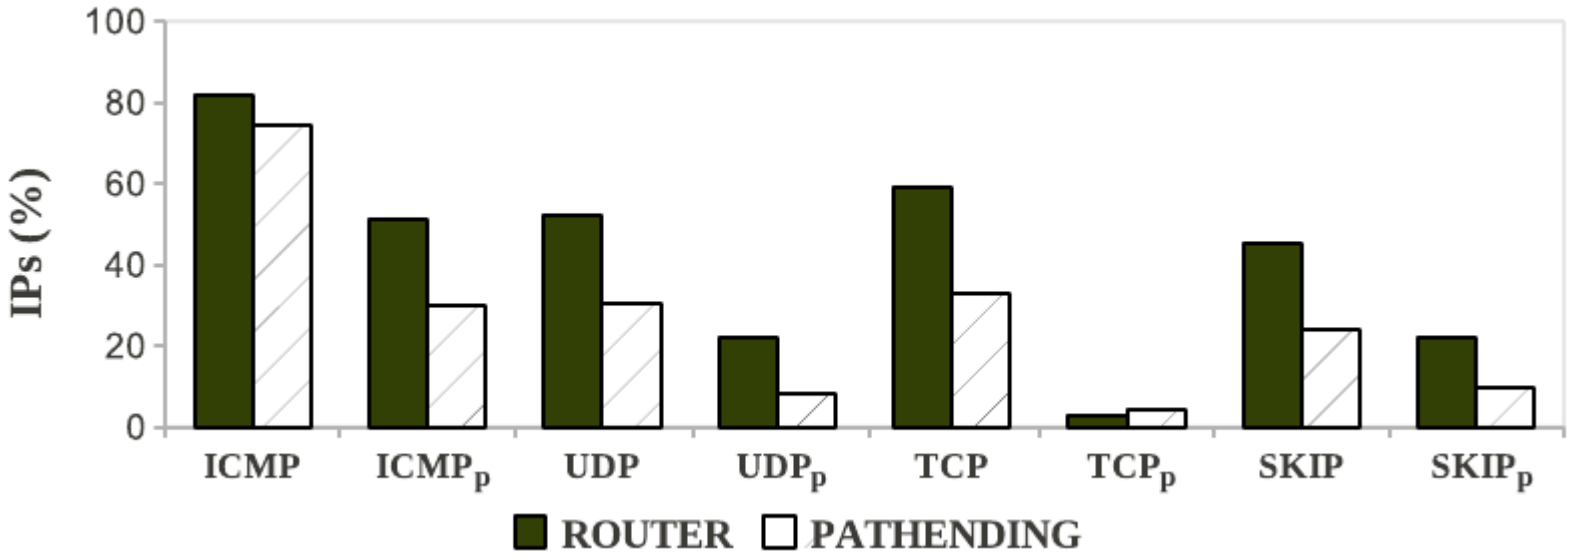

Responsiveness to probes (2/2)

- **Classification of destination IPs** (from the Archipelago dataset)
  - **Pathending** → appear in the dataset only as traceroute destinations
  - **Router** → otherwise

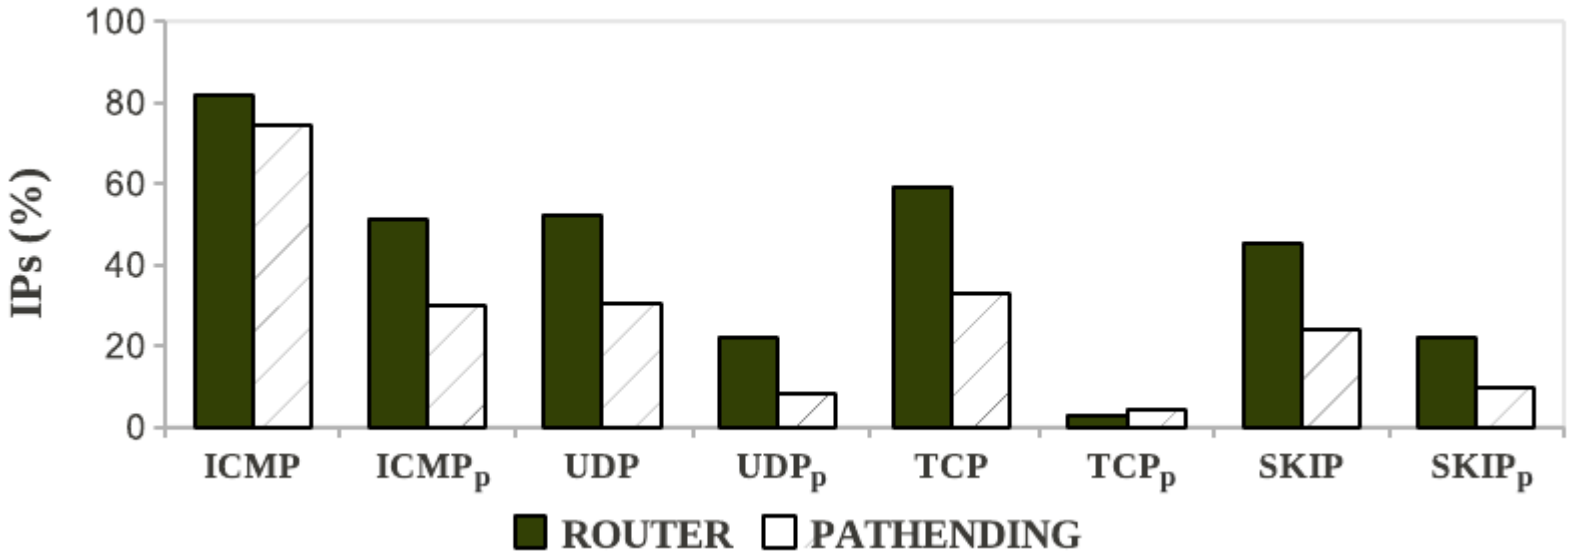

Router IPs are more responsive to most probes

Relation among different probes

- **Responsiveness Matrix**
  - **(*i, j*) → % of IPs responsive to both probes on ith row and jth column**

without TSp option

	ICMP	UDP	TCP	SKIP
ICMP	**78.1**	40.6	44.9	32.6
UDP	40.6	**41.4**	37.6	30.1
TCP	44.9	37.6	**46.1**	28.9
SKIP	32.6	30.1	28.9	**34.7**

with TSp option

	$ICMP_p$	UDP_p	TCP_p	$SKIP_p$
$ICMP_p$	**40.7**	13.2	3.5	13.5
UDP_p	13.2	**15.0**	3.2	11.6
TCP_p	3.5	3.2	**3.6**	2.6
$SKIP_p$	13.5	11.6	2.6	**15.8**

Relation among different probes

- **Responsiveness Matrix**
  - **(*i*, *j*) → % of IPs responsive to both probes on ith row and jth column**

without TSp option

	ICMP	UDP	TCP	SKIP
ICMP	**78.1**	40.6	44.9	32.6
UDP	40.6	**41.4**	37.6	30.1
TCP	44.9	37.6	**46.1**	28.9
SKIP	32.6	30.1	28.9	**34.7**

with TSp option

	$ICMP_p$	UDP_p	TCP_p	$SKIP_p$
$ICMP_p$	**40.7**	13.2	3.5	13.5
UDP_p	13.2	**15.0**	3.2	11.6
TCP_p	3.5	3.2	**3.6**	2.6
$SKIP_p$	13.5	11.6	2.6	**15.8**

- **ICMP shows the best marginal utility regardless of the option**

Relation among different probes

- **Responsiveness Matrix**
  - **(*i*, *j*) → % of IPs responsive to both probes on ith row and jth column**

without TSp option

	ICMP	UDP	TCP	SKIP
ICMP	**78.1**	40.6	44.9	32.6
UDP	40.6	**41.4**	37.6	30.1
TCP	44.9	37.6	**46.1**	28.9
SKIP	32.6	30.1	28.9	**34.7**

with TSp option

	$ICMP_p$	UDP_p	TCP_p	$SKIP_p$
$ICMP_p$	**40.7**	13.2	3.5	13.5
UDP_p	13.2	**15.0**	3.2	11.6
TCP_p	3.5	3.2	**3.6**	2.6
$SKIP_p$	13.5	11.6	2.6	**15.8**

- **ICMP shows the best marginal utility regardless of the option**
- **UDP and SKIP show a reduced but interesting marginal utility**

Timestamp rate limiting

- Some destinations not always stamp the TS_p option
  - D_j = set of destinations responding to the generic TS_p probe by stamping j records

	D_0	$D_0 \cap D_1$	$D_0 \cap D_2$	$D_0 \cap D_3$	$D_0 \cap D_4$
$ICMP_p$	98,024	2,443	299	0	0
UDP_p	54,649	643	0	0	0
TCP_p	420	0	0	0	0
$SKIP_p$	56,213	646	0	0	0

Timestamp rate limiting

- Some destinations not always stamp the TS_p option
  - D_j = set of destinations responding to the generic TS_p probe by stamping j records

	D_0	$D_0 \cap D_1$	$D_0 \cap D_2$	$D_0 \cap D_3$	$D_0 \cap D_4$
$ICMP_p$	98,024	2,443	299	0	0
UDP_p	54,649	643	0	0	0
TCP_p	420	0	0	0	0
$SKIP_p$	56,213	646	0	0	0

Mostly involves Router IPs probed with $ICMP_p$

Returned TS$_p$ option: breakdown

- Insights on how most routers manage the option in practice
  - ICMP$_p$ and TCP$_p$ probes return the option as affected by the forward and reverse path
  - UDP$_p$ and SKIP$_p$ probes return the option as affected by the forward path only

	TOT	D_0	D_1	D_2	D_3	D_4
ICMP$_p$	$723k$	13.2	26.4	54.9	$\sim$0	5.5
UDP$_p$	$267k$	20.2	74.5	0.1	0	5.1
TCP$_p$	$62k$	0.7	$\sim$0	99.3	$\sim$0	$\sim$0
SKIP$_p$	$281k$	19.8	80.1	0.1	0	$\sim$0

%

Returned TS_p option: breakdown

- Insights on how most routers manage the option in practice
  - $ICMP_p$ and TCP_p probes return the option as affected by the forward and reverse path
  - UDP_p and $SKIP_p$ probes return the option as affected by the forward path only

	TOT	D_0	D_1	D_2	D_3	D_4
$ICMP_p$	$723k$	13.2	26.4	54.9	$\sim$0	5.5
UDP_p	$267k$	20.2	74.5	0.1	0	5.1
TCP_p	$62k$	0.7	$\sim$0	99.3	$\sim$0	$\sim$0
$SKIP_p$	$281k$	19.8	80.1	0.1	0	$\sim$0

%

The option is stamped once every time the probe passes through the interface associated to the prespecified address currently pointed

Returned TS_p option: ICMP vs UDP

- Intersection between replying destinations
  - probed with UDP_p → D_j : j = 0 … 4
  - probed with $ICMP_p$ → D_i : i = 0 … 4

		UDP				
		$j=0$	$j=1$	$j=2$	$j=3$	$j=4$
	TOT	$54k$	$198.9k$	246	-	$13.7k$
ICMP $i=0$	$95.3k$	**$27.9k$**	306	-	-	-
ICMP $i=1$	$190.8k$	$12.1k$	**$32.8k$**	112	-	-
ICMP $i=2$	$397.5k$	519	$147.8k$	**92**	-	-
ICMP $i=3$	168	6	2	19	-	-
ICMP $i=4$	$39.6k$	2	2	5	-	**$13.2k$**

- IPs responsive to both $ICMP_p$ and UDP_p probes provided
  - 2 timestamps when probed with $ICMP_p$
  - 1 timestamp when probed with UDP_p

Agenda

- Background and motivation
- **Contribution**
  - Large scale active probing campaign
  - Responsiveness analysis
  - **RFC-compliance analysis**
  - Discussion on support of Ts_p-based techniques
- Conclusion

RFC compliance: timestamp format

- Replies from 660k destinations stamping at least once
  - 87.6% $\in [0, 86.4 * 10^5]$ → standard timestamp value
  - 11.3% $\in [2^{31}, 2^{32}]$ → non-standard timestamp value
  - 1.15% $\in]86.4 * 10^5, 2^{31}[$ → non RFC-compliant value

- Looking at replies carrying contiguous standard timestamps ($T_i, T_{i+1} : i \in [1, 3]$)
  - We identify three cases
    - small positive difference
    - small negative difference
    - huge absolute difference

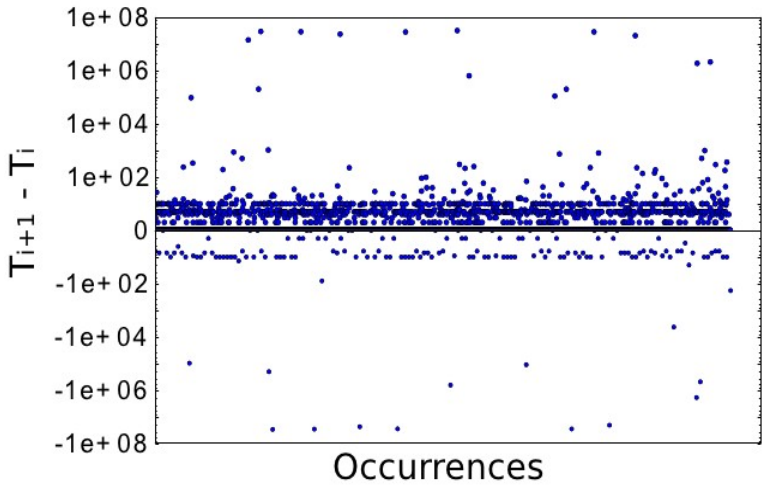

RFC compliance: anomalies (1/3)

- Several anomalies occurred in our dataset
  - ~ 40k destinations provided different non RFC-compliant replies
- Proposed taxonomy
  - **OWR:** some prespecified IP addresses are overwritten
  - **SKP:** the destination stamps the option by skipping one or more records
  - **PTR:** the pointer field is inconsistent compared to the number of timestamps
  - **CUT:** the original packet carried by the ICMP error message is truncated before the end of the option
  - **OFL:** the overflow field counts several extra−stamps, but the number of timestamps is less than four
  - **RPL:** the option data is replaced with part of the original packet header

RFC compliance: anomalies (2/3)

- Most anomalies come from Pathending destinations

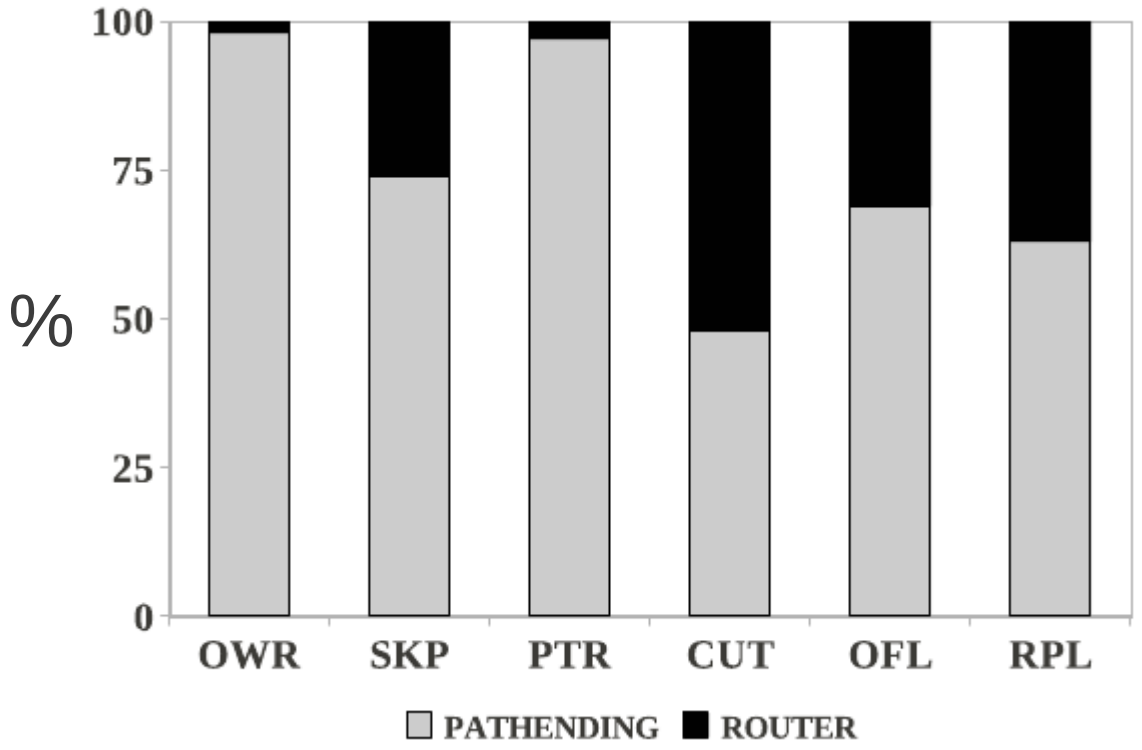

Reasonably related to end hosts having buggy TCP/IP stack implementations

RFC compliance: anomalies (3/3)

- Specific anomalies affect specific types of probe
  - CUT only affects UDP_p and $SKIP_p$ probes
  - OFL only affects $ICMP_p$ and TCP_p probes
  - RPL mostly affects UDP_p and $SKIP_p$ probes

	TOT	$ICMP_p$	UDP_p	TCP_p	$SKIP_p$
OWR	29.3k	24.8k	293	3.8k	3.7k
SKP	32	28	4	2	3
PTR	29.9k	28.5k	725	3	4.5k
CUT	6.2k	–	5.6k	–	3.5k
OFL	26	26	–	6	–
RPL	383	–	249	1	287

RFC compliance: anomalies (3/3)

- Specific anomalies affect specific types of probe
  - CUT only affects UDP_p and $SKIP_p$ probes
  - OFL only affects $ICMP_p$ and TCP_p probes
  - RPL mostly affects UDP_p and $SKIP_p$ probes

	TOT	$ICMP_p$	UDP_p	TCP_p	$SKIP_p$
OWR	29.3k	24.8k	293	3.8k	3.7k
SKP	32	28	4	2	3
PTR	29.9k	28.5k	725	3	4.5k
CUT	6.2k	–	5.6k	–	3.5k
OFL	26	26	–	6	–
RPL	383	–	249	1	287

OWR and PTR are the most common anomalies

RFC compliance: anomalies (3/3)

- Specific anomalies affect specific types of probe
  - CUT only affects UDP_p and $SKIP_p$ probes
  - OFL only affects $ICMP_p$ and TCP_p probes
  - RPL mostly affects UDP_p and $SKIP_p$ probes

	TOT	$ICMP_p$	UDP_p	TCP_p	$SKIP_p$
OWR	29.3k	24.8k	293	3.8k	3.7k
SKP	32	28	4	2	3
PTR	29.9k	28.5k	725	3	4.5k
CUT	6.2k	–	5.6k	–	3.5k
OFL	26	26	–	6	–
RPL	383	–	249	1	287

$ICMP_p$ is the most affected probe

Deepening OWR anomaly

- OWR → prespecified IPs overwritten in different ways
  - 85% → only the first prespecified IP is overwritten
    - 99.7% returned empty records
      - → identified as hosts stamping over the prespecified IP
    - 0.3% stamped at least the first record
      - → found nodes along the path treating TS_p as TS_i
  - 13% → part of the prespecified IPs is reset to zero
    - Only affected TCP_p probes
      - → generated by Windows 2000/2003 servers
  - 2% → both previous cases apply

Agenda

- Background and motivation
- **Contribution**
  - Large scale active probing campaign
  - Responsiveness analysis
  - RFC-compliance analysis
  - **Discussion on support of Ts_p-based techniques**
- Conclusion

Reverse traceroute support

- **Relies on TS_p when RecordRoute option is unable to discover the next hop on the reverse path from D to S**
  - **By using (D|DR) $ICMP_p$ probes from S**
  - **By using spoofed (D|R) $ICMP_p$ probes from a selected VP**
    - **D = last hop discovered on the reverse path**
    - **R = IP extracted from pre-collected topology information**
  - **If S receives a reply with R stamped → R is part of the reverse path**

Katz-Bassett et al., *Reverse traceroute*. NSDI'10, San Jose, CA, USA, 2010

Reverse traceroute support

- **Relies on TS_p when RecordRoute option is unable to discover the next hop on the reverse path from D to S**
  - **By using (D|DR) $ICMP_p$ probes from S**
  - **By using spoofed (D|R) $ICMP_p$ probes from a selected VP**
    - **D = last hop discovered on the reverse path**
    - **R = IP extracted from pre-collected topology information**
  - **If S receives a reply with R stamped → R is part of the reverse path**

Such approach works with 35% of IPs from our dataset

Katz-Bassett et al., *Reverse traceroute*. NSDI'10, San Jose, CA, USA, 2010

TS_p-based Alias Resolution support

- **Applied to a pair (A, B) of candidate IPs**
  - **Sends a (A|ABAB) $ICMP_p$ probe towards A**
  - **Sends a (B|BABA) $ICMP_p$ probe towards B**
    - **If both replies are stamped four times → A and B are alias**
- **It works only in some cases**
  - **A and B ∈ D_4 (i.e. provided four stamps) → 2.2% of our dataset**
  - **A and B ∈ D_1 (i.e. provided one stamp) → 10.7% of our dataset**
- **It fails in some other cases**
  - **A and B ∈ D_2 (i.e. provided two stamps) → 22.3% of our dataset**

Sherry et al., *Resolving ip aliases with prespecified timestamps*. IMC'10, New York, NY, USA, 2010

TS_p-based Alias Resolution support

- **Applied to a pair (A, B) of candidate IPs**
  - **Sends a (A|ABAB) $ICMP_p$ probe towards A**
  - **Sends a (B|BABA) $ICMP_p$ probe towards B**
    - **If both replies are stamped four times → A and B are alias**
- **It works only in some cases**
  - **A and B ∈ D_4 (i.e. provided four stamps) → 2.2% of our dataset**
  - **A and B ∈ D_1 (i.e. provided one stamp) → 10.7% of our dataset**
- **It fails in some other cases**
  - **A and B ∈ D_2 (i.e. provided two stamps) → 22.3% of our dataset**

Such approach works with 12.9% of IPs using a single VP

Sherry et al., *Resolving ip aliases with prespecified timestamps*. IMC'10, New York, NY, USA, 2010

Agenda

- Background and motivation
- Contribution
  - Large scale active probing campaign
  - Responsiveness analysis
  - RFC-compliance analysis
  - Discussion on support of Ts_p-based techniques
- **Conclusion**

Conclusion

- TS_p option heavily impacts responsiveness to probes
- The option is commoly stamped per network interface
- Unexpected behaviors may occur
  - Timestamp rate limiting policies
  - non RFC-compliant replies
  - non RFC-compliant timestamps
- Proposed TS_p-based techniques find limited support

Thank you!

Questions ?