UNIVERSITÄT DES SAARLANDES

Yarrpbox: Detecting Middleboxes at Internet-Scale

Fahad Hilal, Oliver Gasser

MPI-INF

CoNEXT'23

max planck institut informatik

Introduction

- ***End-to-end principle***
  - simplicity in middle, intelligence at ends

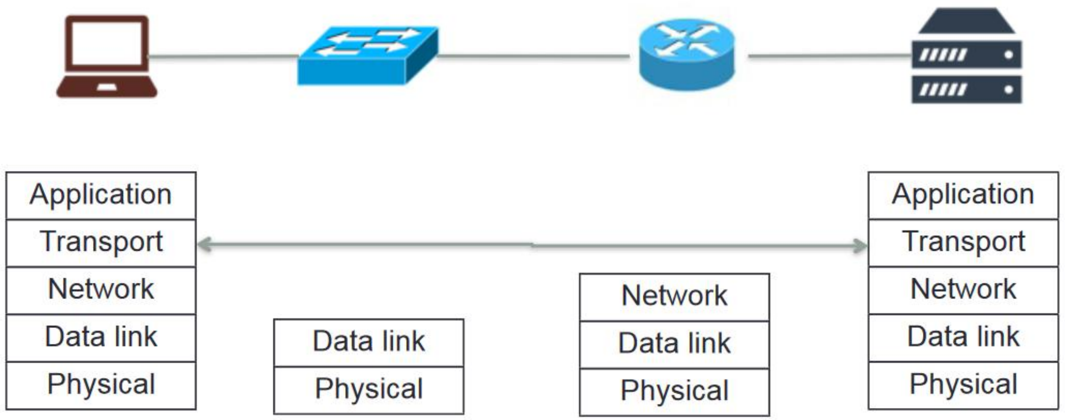

https://www.ietf.org/proceedings/interim-2013-nmrg-01/slides/slides-interim-2013-nmrg-1-9.pdf

Introduction

- ***Paradigm shift***
  - Internet as ***deployed***, no longer as ***designed***

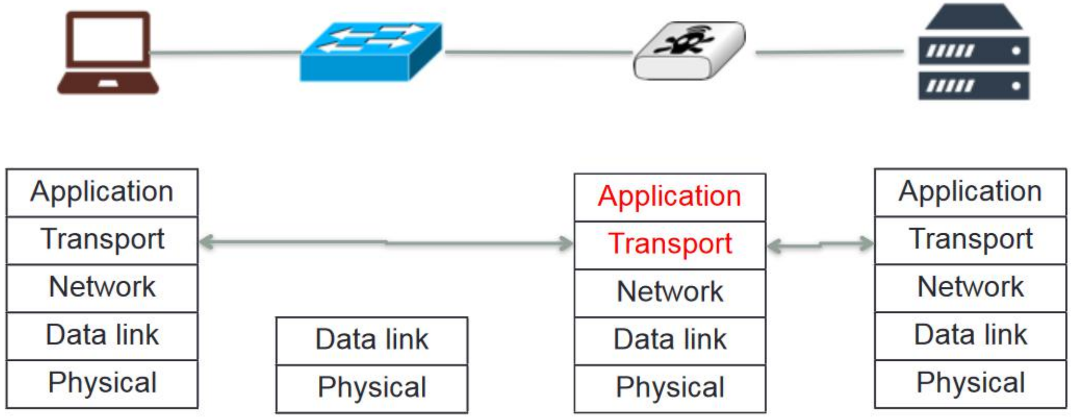

https://www.ietf.org/proceedings/interim-2013-nmrg-01/slides/slides-interim-2013-nmrg-1-9.pdf

Introduction

UNIVERSITÄT DES SAARLANDES

- ***Paradigm shift***
  - Internet as ***deployed***, no longer as ***designed***
  - ***invasion*** of ***middleboxes***

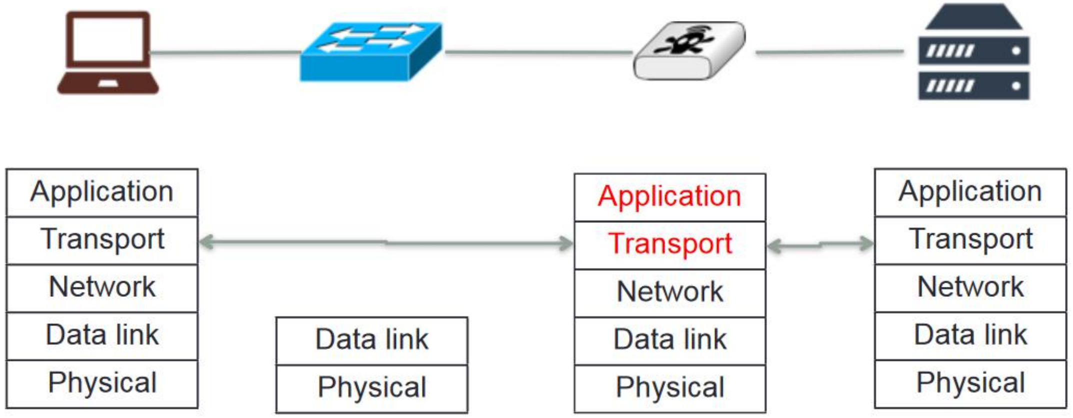

https://www.ietf.org/proceedings/interim-2013-nmrg-01/slides/slides-interim-2013-nmrg-1-9.pdf

Introduction

- ***Paradigm shift***
  - Internet as ***deployed***, no longer as ***designed***
  - ***invasion*** of ***middleboxes***
    - inspect, filter, modify

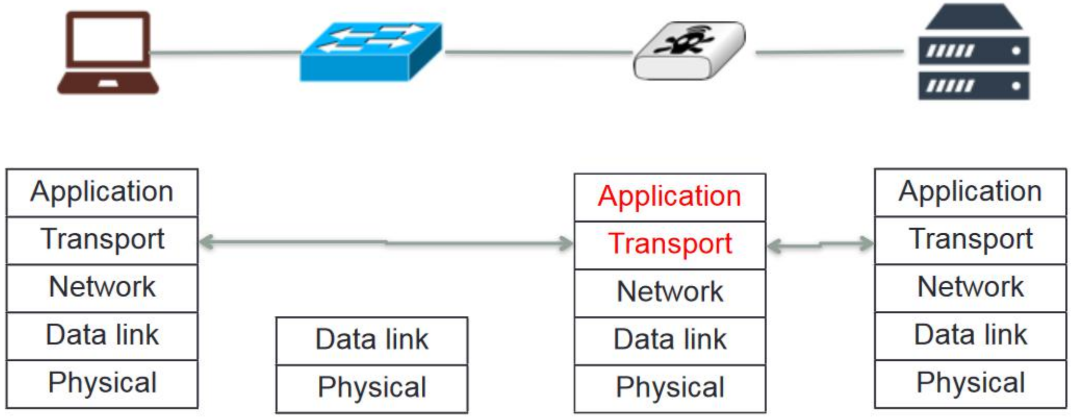

https://www.ietf.org/proceedings/interim-2013-nmrg-01/slides/slides-interim-2013-nmrg-1-9.pdf

Introduction

- ***Paradigm shift***
  - Internet as ***deployed***, no longer as ***designed***
  - ***invasion*** of ***middleboxes***
    - inspect, filter, modify
    - thwart attacks, expand address space, balance resources

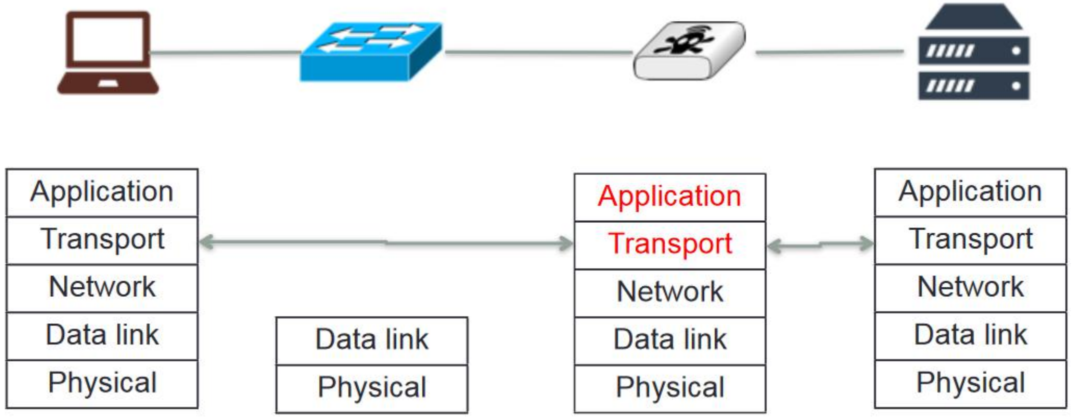

https://www.ietf.org/proceedings/interim-2013-nmrg-01/slides/slides-interim-2013-nmrg-1-9.pdf

Introduction

- ***Paradigm shift***
  - Internet as ***deployed***, no longer as ***designed***
  - ***invasion*** of ***middleboxes***
    - inspect, filter, modify
    - thwart attacks, expand address space, balance resources
    - ossification

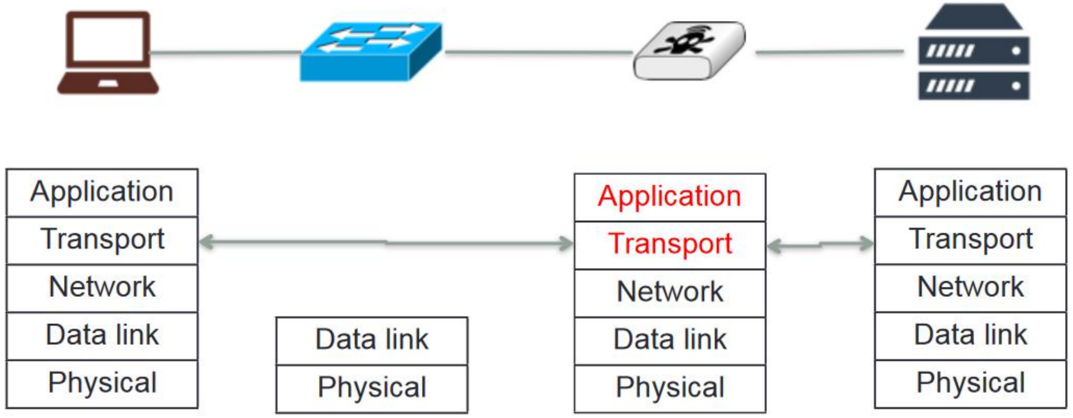

https://www.ietf.org/proceedings/interim-2013-nmrg-01/slides/slides-interim-2013-nmrg-1-9.pdf

Goals

- Detect **packet-rewriting** middlebox interferences
  - IP and transport layer

Goals

- Detect **packet-rewriting** middlebox interferences
  - IP and transport layer
  - legacy, misconfigs., previously benign

Goals

- Detect **packet-rewriting** middlebox interferences
  - IP and transport layer
  - legacy, misconfigs., previously benign
- Detect middlebox **location**

Goals

- Detect **packet-rewriting** middlebox interferences
  - IP and transport layer
  - legacy, misconfigs., previously benign
- Detect middlebox **location**
- Improve upon existing tools
  - make **Internet-scale** middlebox detection **feasible**

Motivation

- ***Side effects***
  - negative impact on ***evolvability***

Motivation

- ***Side effects***
  - negative impact on ***evolvability***
    - DCCP, SCTP standardized,
      - failed to be deployed at large scale

Motivation

- ***Side effects***
  - negative impact on ***evolvability***
    - DCCP, SCTP standardized,
      - failed to be deployed at large scale
  - ***middlebox-proof*** solutions needed

Motivation

- ***Side effects***
  - negative impact on ***evolvability***
    - DCCP, SCTP standardized,
      - failed to be deployed at large scale
  - ***middlebox-proof*** solutions needed
- **Transient dynamics**
  - ***~50%*** display ***dynamic*** behavior [1]

[1] *Edeline, Korian, and Benoit Donnet. "A first look at the prevalence and persistence of middleboxes in the wild." 2017 29th International Teletraffic Congress (ITC 29).*

Motivation

- ***Side effects***
  - negative impact on ***evolvability***
    - DCCP, SCTP standardized,
      - failed to be deployed at large scale
  - ***middlebox-proof*** solutions needed
- **Transient dynamics**
  - ***~50%*** display ***dynamic*** behavior [1]
  - *timely* detection needed

[1] *Edeline, Korian, and Benoit Donnet. "A first look at the prevalence and persistence of middleboxes in the wild." 2017 29th International Teletraffic Congress (ITC 29).*

Background-Traceroute-Style Detection

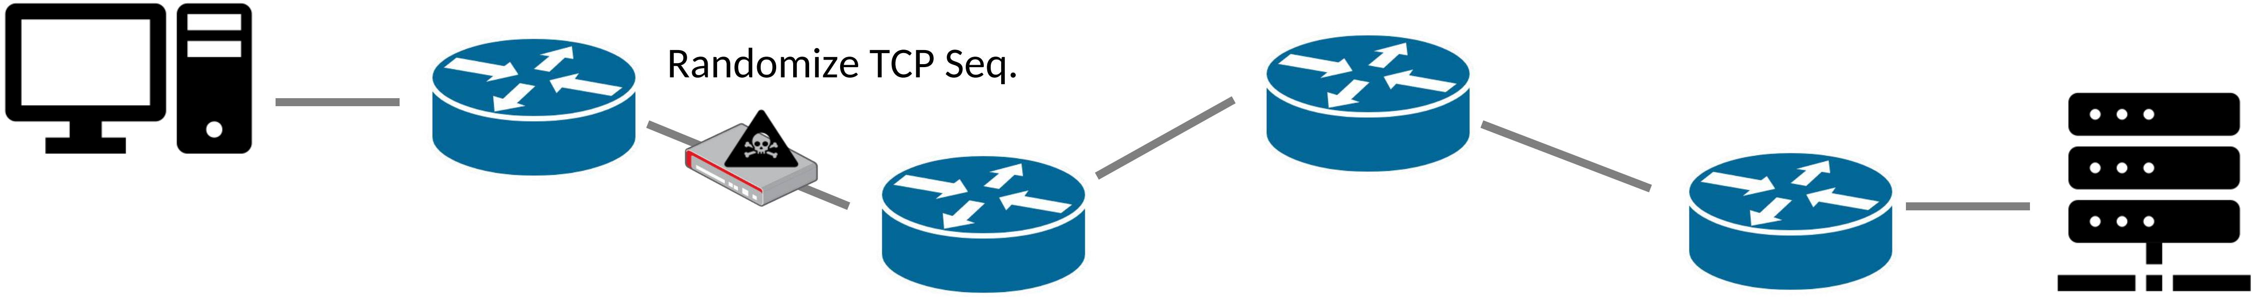

Background-Traceroute-Style Detection

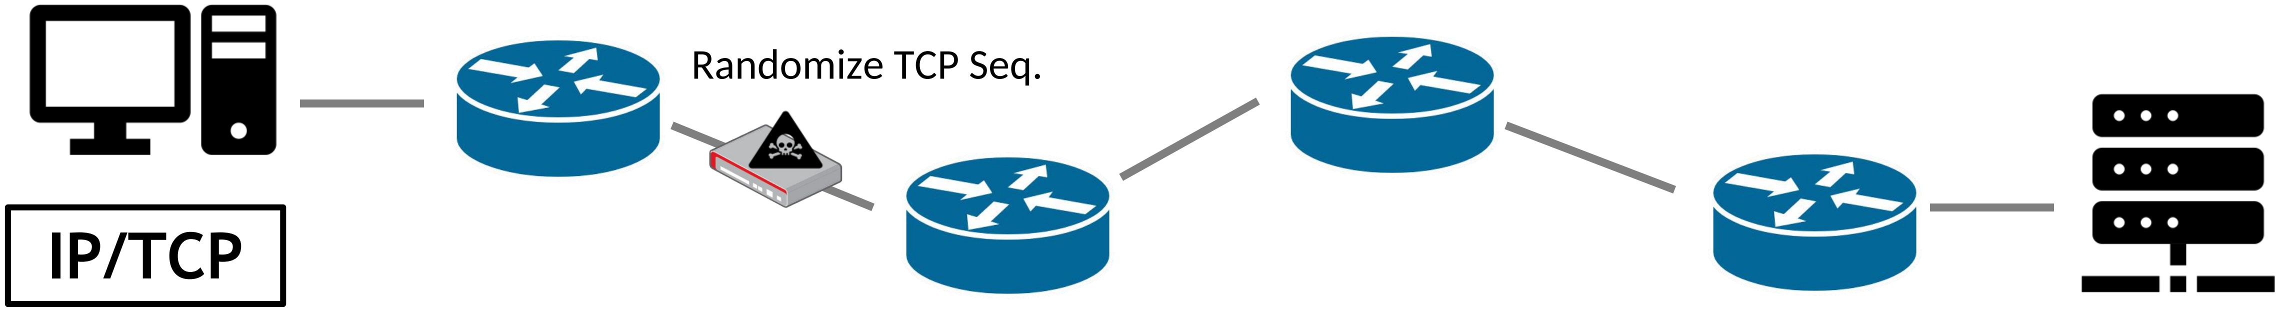

Background-Traceroute-Style Detection

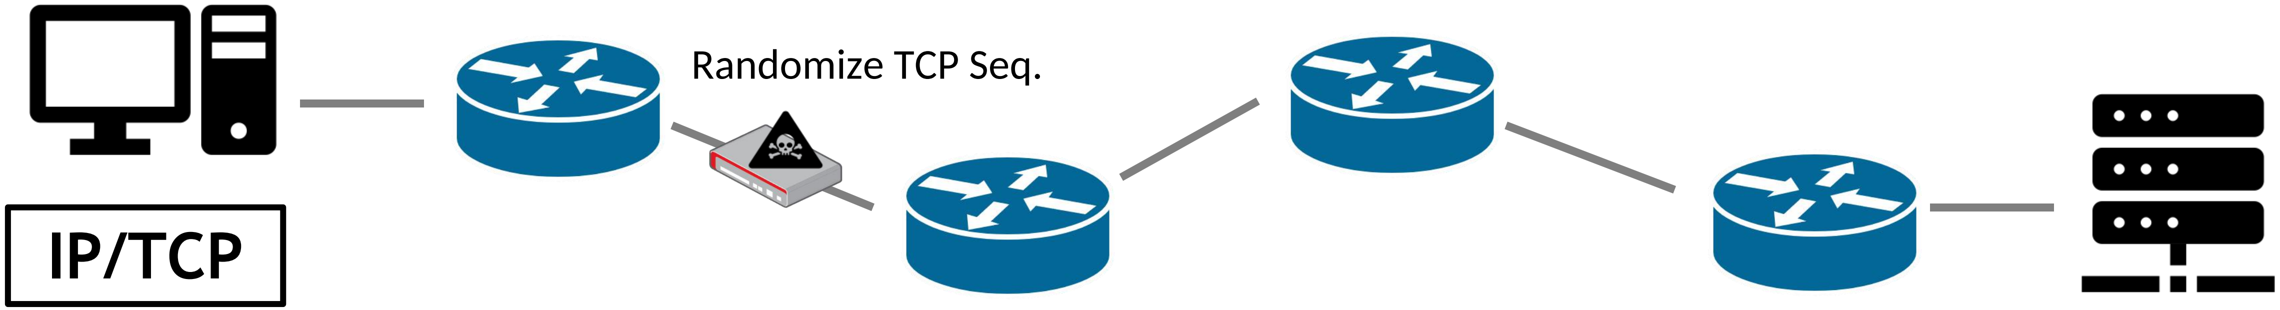

<table>
<tr><td>Ver</td><td>IHL</td><td colspan="2">ToS</td><td colspan="2">Total length</td></tr>
<tr><td colspan="4">Identification</td><td>Flags</td><td>Frag. Offset</td></tr>
<tr><td colspan="2">TTL</td><td colspan="2">Protocol</td><td colspan="2">Checksum</td></tr>
<tr><td colspan="6">Source IP address</td></tr>
<tr><td colspan="6">Destination IP address</td></tr>
<tr><td colspan="4">Source port</td><td colspan="2">Destination port</td></tr>
<tr><td colspan="6">Sequence number</td></tr>
<tr><td colspan="6">Acknowledgment number</td></tr>
<tr><td>THL</td><td colspan="2">Reserved</td><td>Flags</td><td colspan="2">Window</td></tr>
<tr><td colspan="4">Checksum</td><td colspan="2">Urgent pointer</td></tr>
</table>

Background-Traceroute-Style Detection

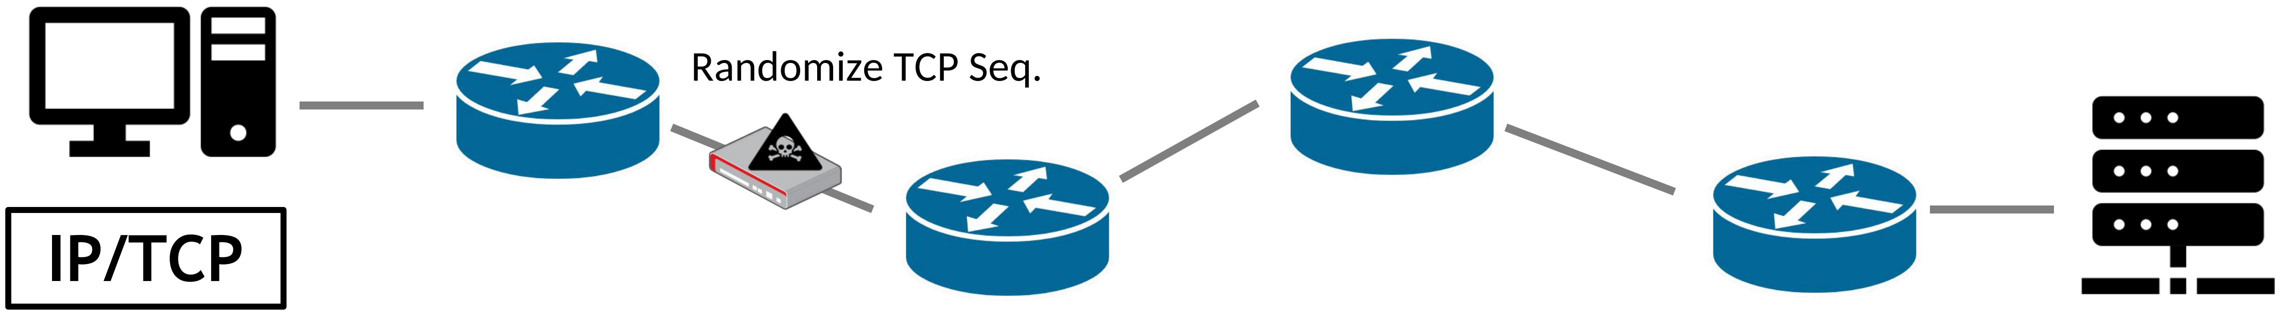

Ver	IHL	ToS	Total length
Identification		Flags	Frag. Offset
TTL = 1	Protocol	Checksum	
Source IP address			
Destination IP address			
Source port		Destination port	
Sequence number			
Acknowledgment number			
THL	Reserved	Flags	Window
Checksum		Urgent pointer	

Background-Traceroute-Style Detection

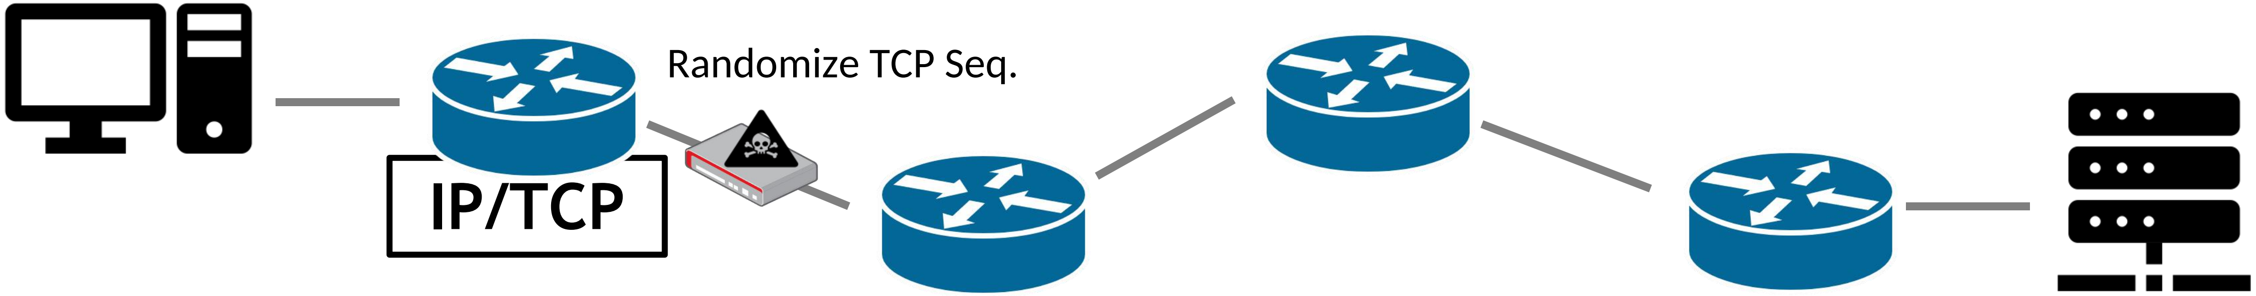

<table>
<tr><td>Ver</td><td>IHL</td><td>ToS</td><td colspan="2">Total length</td></tr>
<tr><td colspan="3">Identification</td><td>Flags</td><td>Frag. Offset</td></tr>
<tr><td colspan="2">TTL = 0</td><td>Protocol</td><td colspan="2">Checksum</td></tr>
<tr><td colspan="5">Source IP address</td></tr>
<tr><td colspan="5">Destination IP address</td></tr>
<tr><td colspan="3">Source port</td><td colspan="2">Destination port</td></tr>
<tr><td colspan="5">Sequence number</td></tr>
<tr><td colspan="5">Acknowledgment number</td></tr>
<tr><td>THL</td><td>Reserved</td><td>Flags</td><td colspan="2">Window</td></tr>
<tr><td colspan="3">Checksum</td><td colspan="2">Urgent pointer</td></tr>
</table>

Background-Traceroute-Style Detection

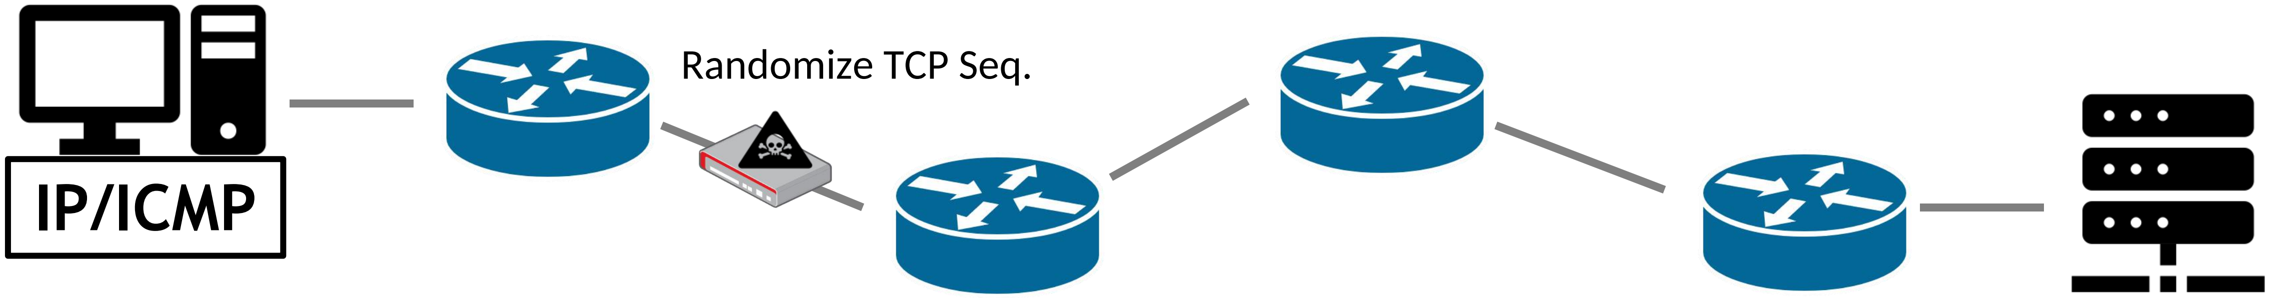

Background-Traceroute-Style Detection

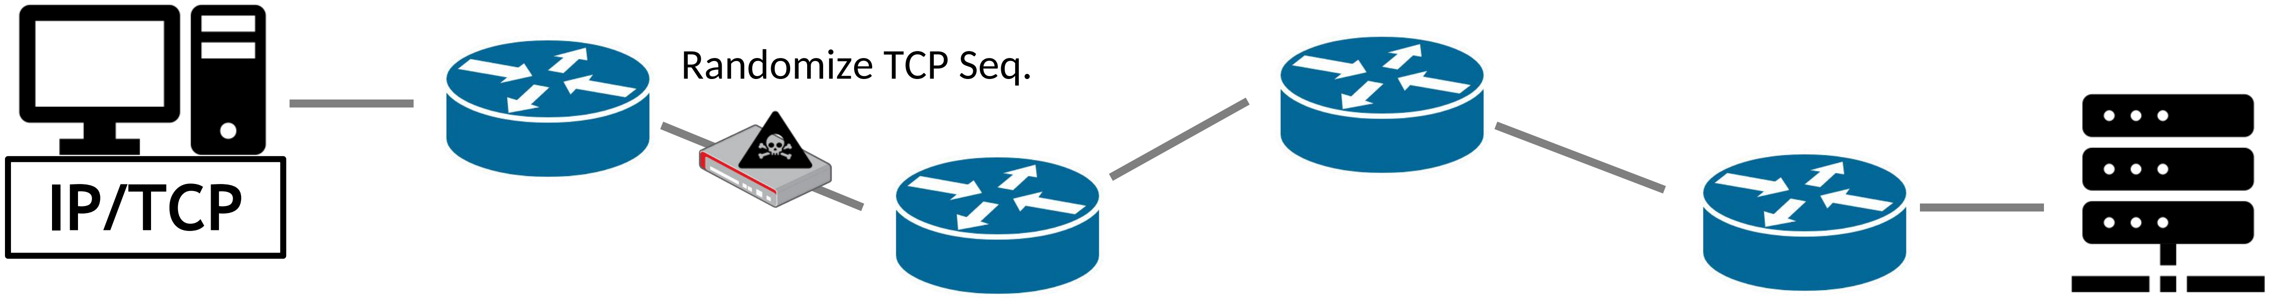

Ver	IHL	ToS	Total length
Identification		Flags	Frag. Offset
TTL = 2	Protocol	Checksum	
Source IP address			
Destination IP address			
Source port		Destination port	
Sequence number			
Acknowledgment number			
THL	Reserved	Flags	Window
Checksum		Urgent pointer	

Background-Traceroute-Style Detection

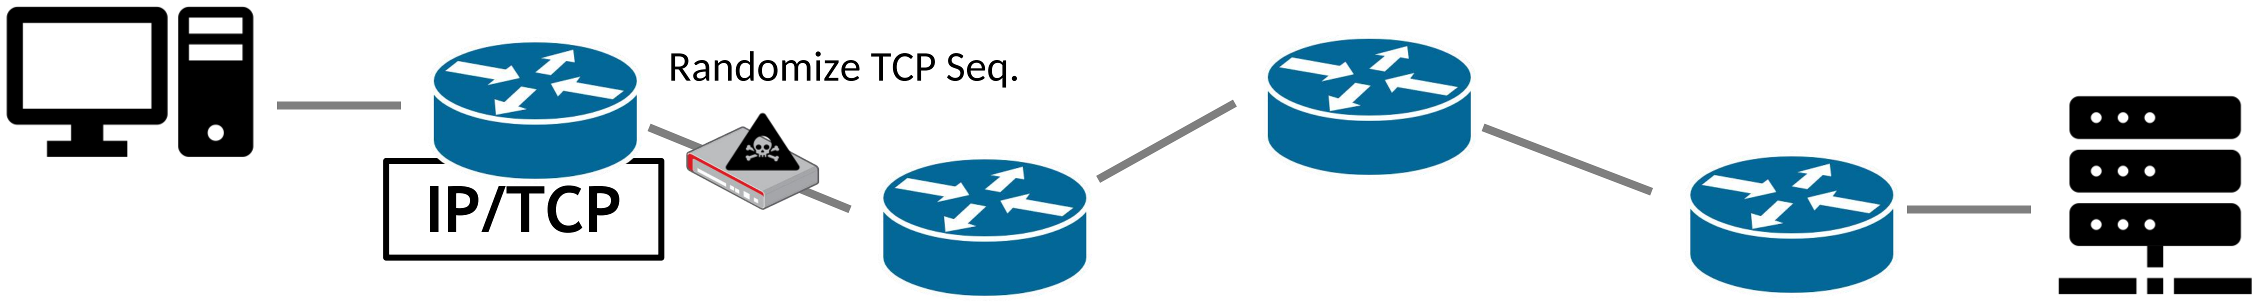

Ver	IHL	ToS	Total length
Identification		Flags	Frag. Offset
TTL = 1	Protocol	Checksum	
Source IP address			
Destination IP address			
Source port		Destination port	
Sequence number			
Acknowledgment number			
THL	Reserved	Flags	Window
Checksum		Urgent pointer	

Background-Traceroute-Style Detection

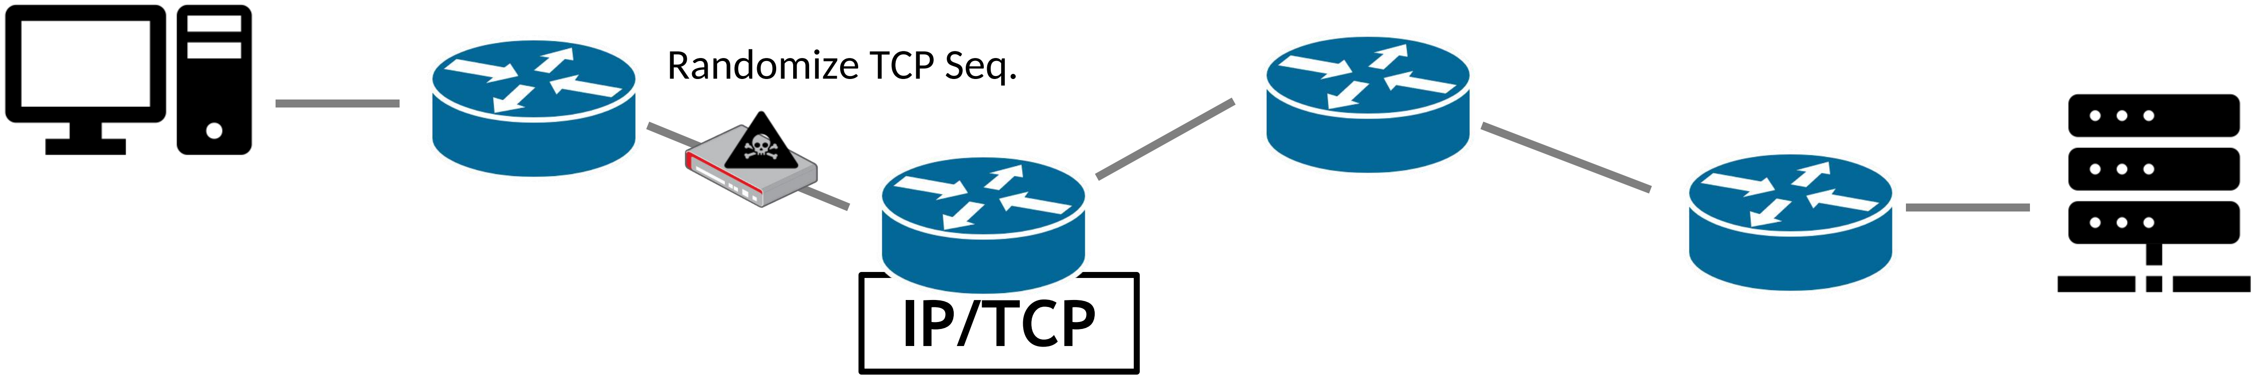

<table>
<tr><td>Ver</td><td>IHL</td><td>ToS</td><td colspan="2">Total length</td></tr>
<tr><td colspan="3">Identification</td><td>Flags</td><td>Frag. Offset</td></tr>
<tr><td colspan="2">TTL = 0</td><td>Protocol</td><td colspan="2">Checksum</td></tr>
<tr><td colspan="5">Source IP address</td></tr>
<tr><td colspan="5">Destination IP address</td></tr>
<tr><td colspan="3">Source port</td><td colspan="2">Destination port</td></tr>
<tr><td colspan="5"><i>Sequence number</i></td></tr>
<tr><td colspan="5">Acknowledgment number</td></tr>
<tr><td>THL</td><td>Reserved</td><td>Flags</td><td colspan="2">Window</td></tr>
<tr><td colspan="3">Checksum</td><td colspan="2">Urgent pointer</td></tr>
</table>

Background-Traceroute-Style Detection

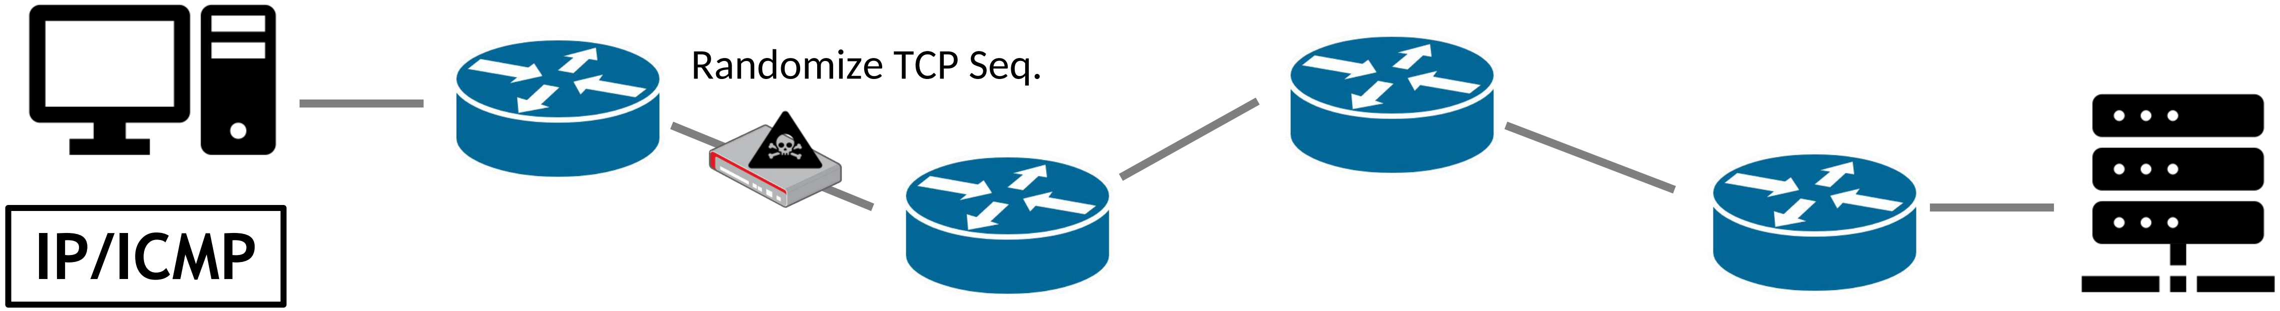

IP/ICMP

Randomize TCP Seq.

Background-Traceroute-Style Detection

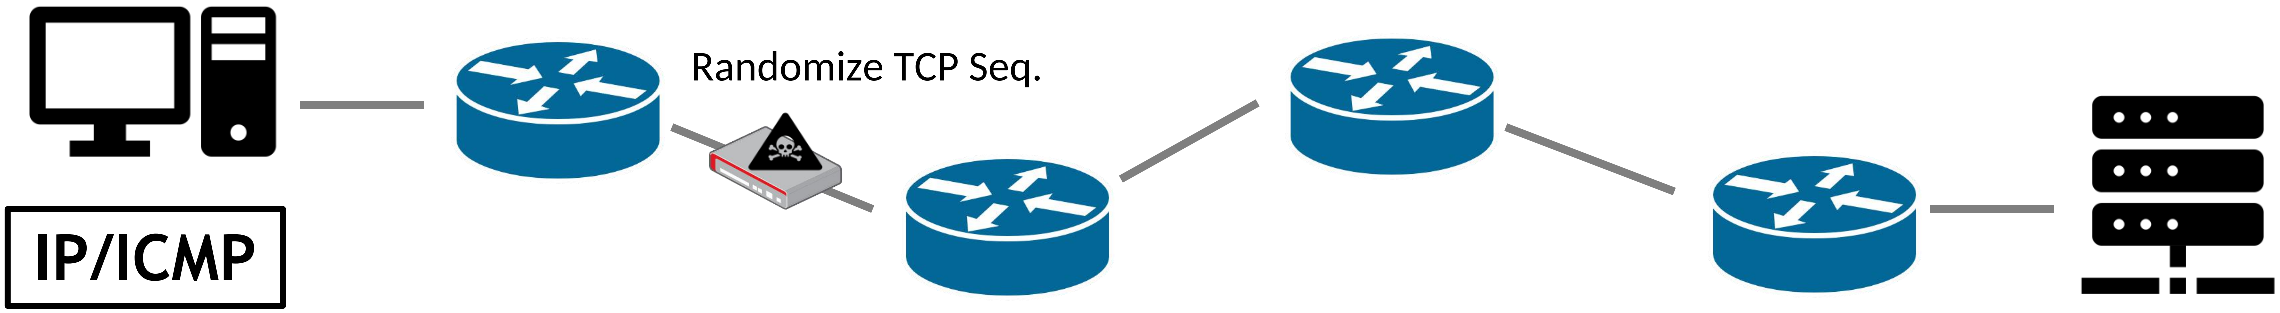

Randomize TCP Seq.

IP/ICMP

IP		
type = 11	code = 0	*checksum*
0 (unused)		

Background-Traceroute-Style Detection

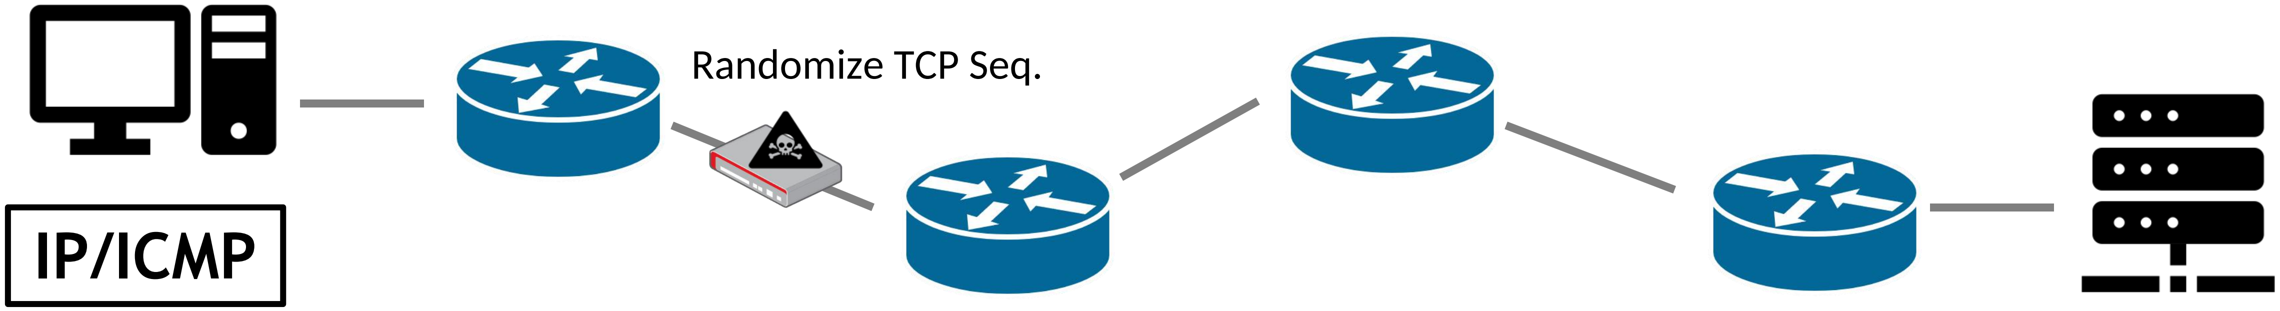

IP/ICMP

Randomize TCP Seq.

IP			
type = 11	code = 0	*checksum*	
0 (unused)			
Ver	IHL	ToS	Total length
Identification		Flags	Frag. Offset
TTL = 0	Protocol	Checksum	
Source IP address			
Destination IP address			
Source port		Destination port	
Sequence number			
Acknowledgment number			
THL	Reserved	Flags	Window
Checksum		Urgent pointer	

Background-Traceroute-Style Detection

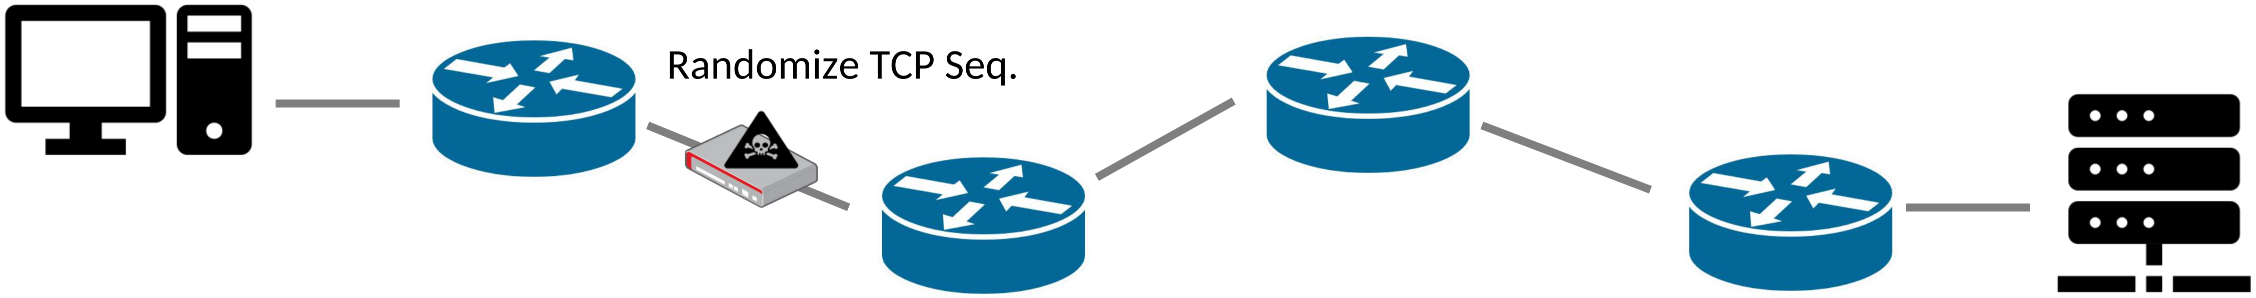

Snapshot at Router 2:

<table>
<tr><td>Ver</td><td>IHL</td><td>ToS</td><td colspan="2">Total length</td></tr>
<tr><td colspan="3">Identification</td><td>Flags</td><td>Frag. Offset</td></tr>
<tr><td colspan="2">1</td><td>Protocol</td><td colspan="2">Checksum</td></tr>
<tr><td colspan="5">Source IP address</td></tr>
<tr><td colspan="5">Destination IP address</td></tr>
<tr><td colspan="3">Source port</td><td colspan="2">Destination port</td></tr>
<tr><td colspan="5"><i>Sequence number</i></td></tr>
<tr><td colspan="5">Acknowledgment number</td></tr>
<tr><td>THL</td><td>Reserved</td><td>Flags</td><td colspan="2">Window</td></tr>
<tr><td colspan="3">Checksum</td><td colspan="2">Urgent pointer</td></tr>
</table>

Background-Traceroute-Style Detection

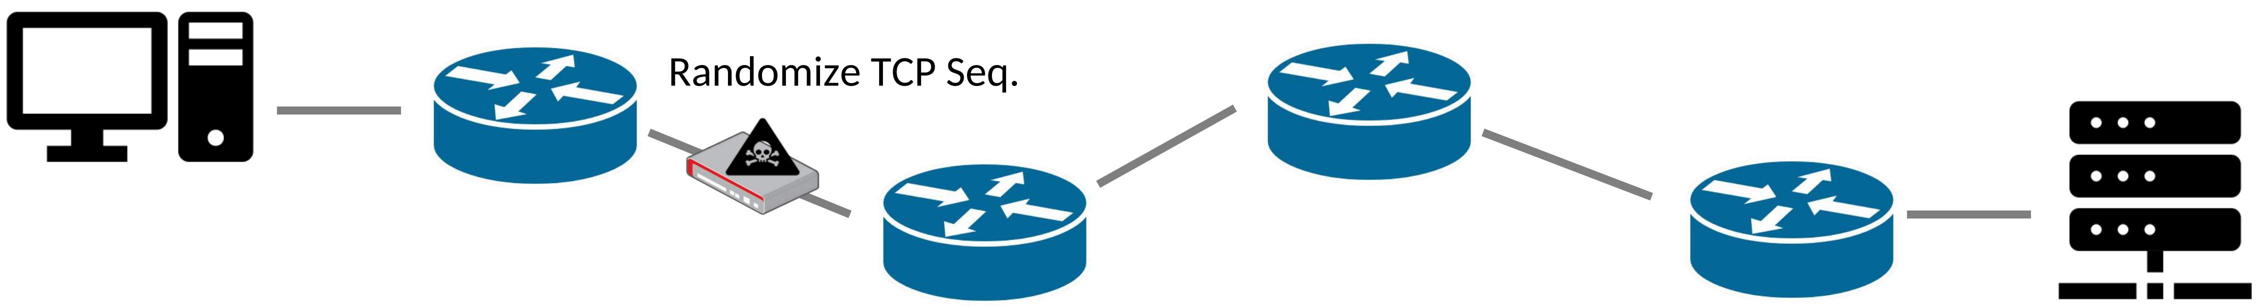

Snapshot at Router 2:

Ver	IHL	ToS	Total length
Identification		Flags	Frag. Offset
1		Protocol	Checksum
Source IP address			
Destination IP address			
Source port		Destination port	
Sequence number			
Acknowledgment number			
THL	Reserved	Flags	Window
Checksum		Urgent pointer	

Sent Probe:

Ver	IHL	ToS	Total length
Identification		Flags	Frag. Offset
2		Protocol	Checksum
Source IP address			
Destination IP address			
Source port		Destination port	
Sequence number			
Acknowledgment number			
THL	Reserved	Flags	Window
Checksum		Urgent pointer	

Background-Traceroute-Style Detection

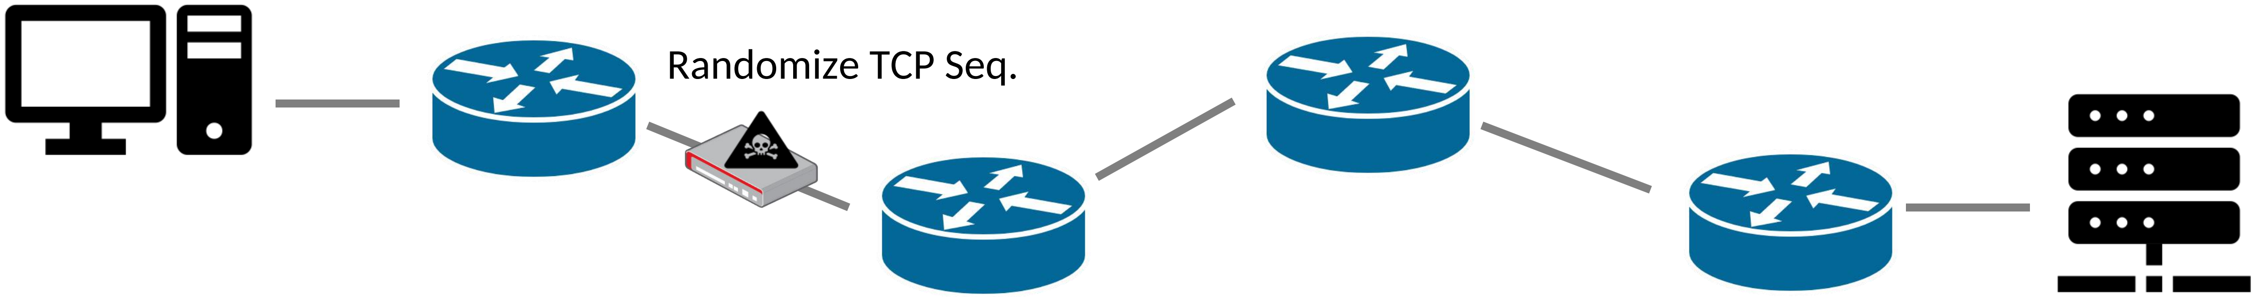

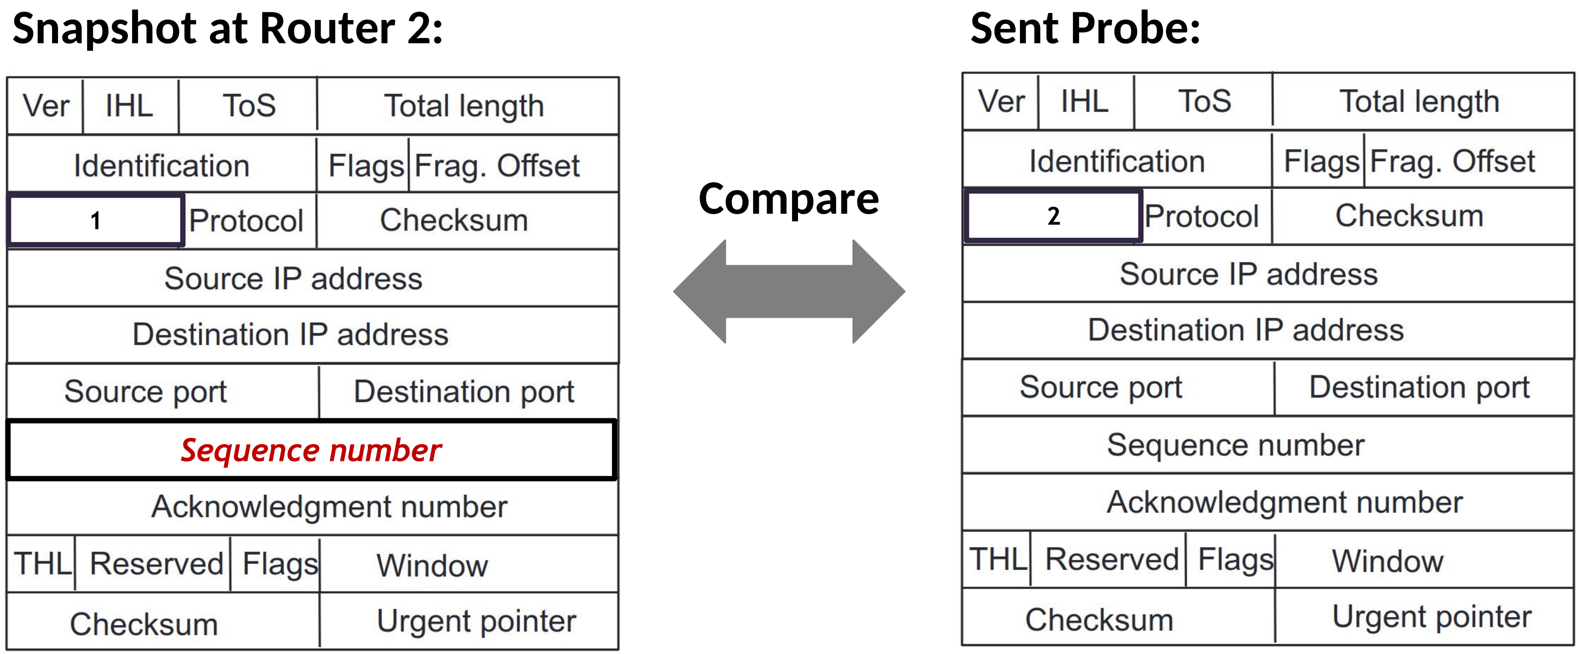

Snapshot at Router 2:

Ver | IHL | ToS | Total length
Identification | Flags | Frag. Offset
1 | Protocol | Checksum
Source IP address
Destination IP address
Source port | Destination port
Sequence number
Acknowledgment number
THL | Reserved | Flags | Window
Checksum | Urgent pointer

Compare

Sent Probe:

Ver | IHL | ToS | Total length
Identification | Flags | Frag. Offset
2 | Protocol | Checksum
Source IP address
Destination IP address
Source port | Destination port
Sequence number
Acknowledgment number
THL | Reserved | Flags | Window
Checksum | Urgent pointer

Related Work

- ***Tracebox*** [1]
  - ***limitations***
    - ***state*** over outstanding probes
      - slow

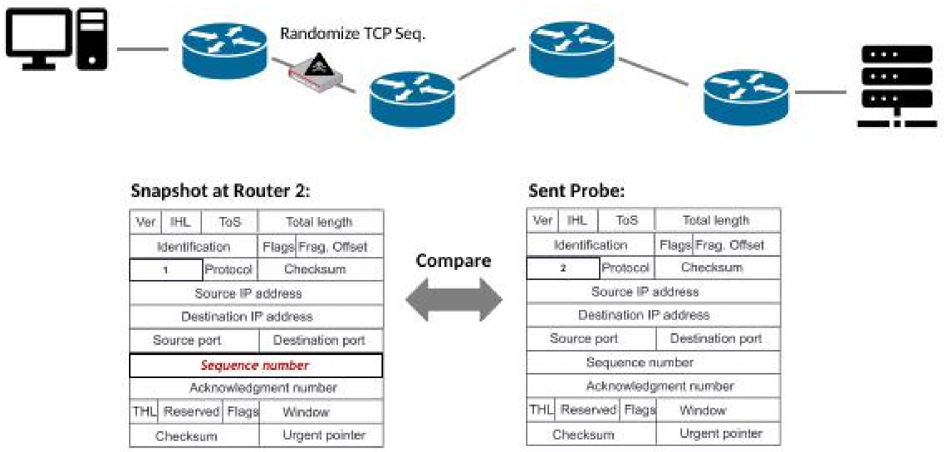

Snapshot at Router 2:

Ver	IHL	ToS	Total length
Identification		Flags	Frag. Offset
1	Protocol	Checksum	
Source IP address			
Destination IP address			
Source port		Destination port	
Sequence number			
Acknowledgment number			
THL	Reserved	Flags	Window
Checksum		Urgent pointer	

Sent Probe:

Ver	IHL	ToS	Total length
Identification		Flags	Frag. Offset
2	Protocol	Checksum	
Source IP address			
Destination IP address			
Source port		Destination port	
Sequence number			
Acknowledgment number			
THL	Reserved	Flags	Window
Checksum		Urgent pointer	

[1] Detal, Gregory, et al. "Revealing middlebox interference with Tracebox." Proceedings of the 2013 conference on Internet measurement conference. 2013.

Related Work

- ***Tracebox*** [1]
  - ***limitations***
    - ***state*** over outstanding probes
      - **slow**
    - sequential, probing hops in ***sequence***
      - ***rate-limiting***

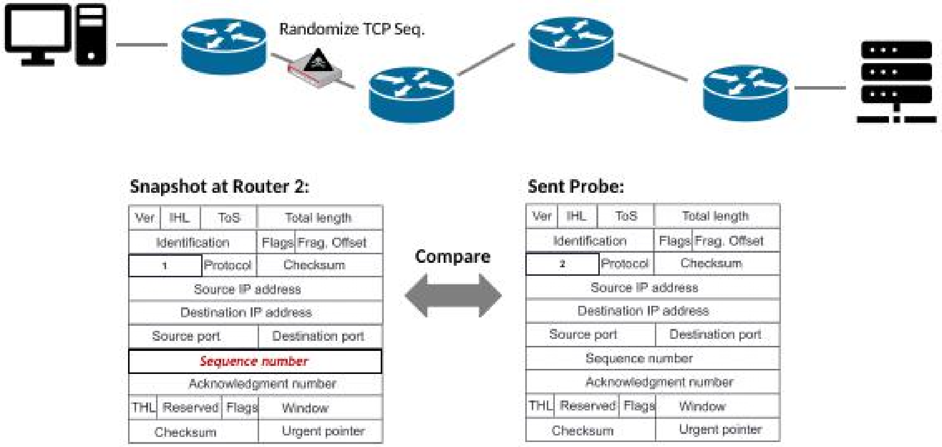

[1] Detal, Gregory, et al. "Revealing middlebox interference with Tracebox." Proceedings of the 2013 conference on Internet measurement conference. 2013.

Related Work

- ***Tracebox*** [1]
  - ***limitations***
    - ***state*** over outstanding probes
      - **slow**
    - sequential, probing hops in ***sequence***
      - ***rate-limiting***

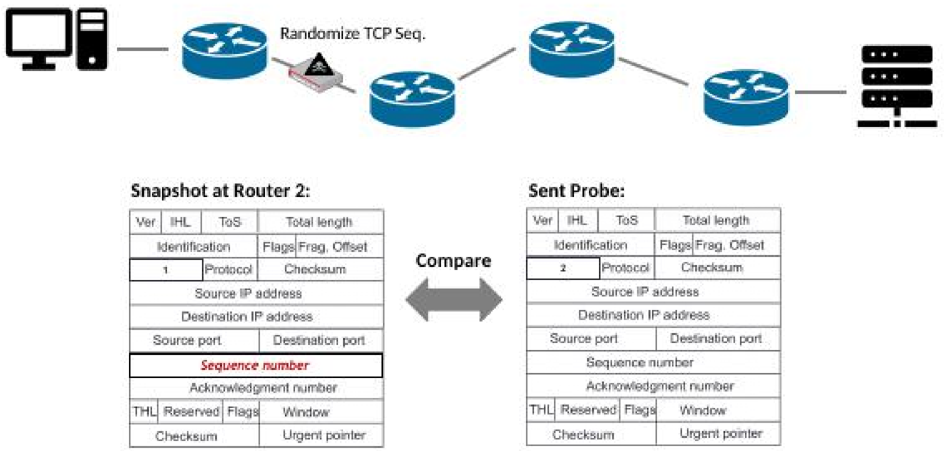

Infeasible for Internet-scale studies!

[1] *Detal, Gregory, et al. "Revealing middlebox interference with Tracebox." Proceedings of the 2013 conference on Internet measurement conference. 2013.*

Yarrpbox

- ***Based on yarrp*** [1]
  - network topology discovery
  - probing over 100Kpps

[1] Beverly, Robert. "Yarrp'ing the Internet: Randomized high-speed active topology discovery." Proceedings of the 2016 Internet Measurement Conference. 2016.

Yarrpbox

- ***Based on yarrp*** [1]
  - network topology discovery
  - probing over 100Kpps

- ***Improvements over Tracebox***
  - ***randomized probing***

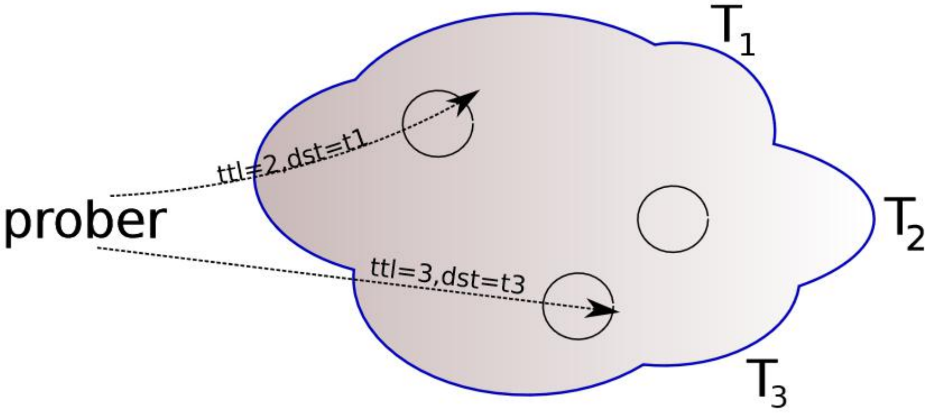

https://www.caida.org/workshops/aims/1602/slides/aims1602_rbeverly.pdf

[1] Beverly, Robert. "Yarrp'ing the Internet: Randomized high-speed active topology discovery." Proceedings of the 2016 Internet Measurement Conference. 2016.

Yarrpbox

- ***Based on yarrp*** [1]
  - network topology discovery
  - probing over 100Kpps

- ***Improvements over Tracebox***
  - ***randomized probing***
    - IP target and TTL domain
    - minimize ICMP rate-limiting

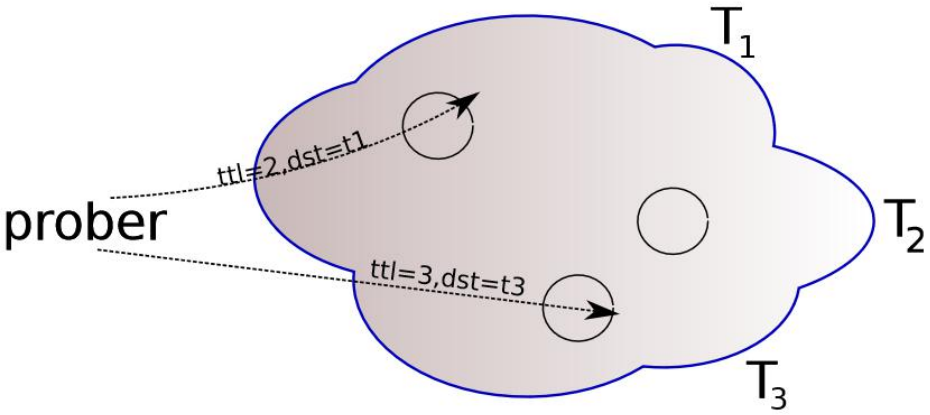

https://www.caida.org/workshops/aims/1602/slides/aims1602_rbeverly.pdf

[1] *Beverly, Robert. "Yarrp'ing the Internet: Randomized high-speed active topology discovery." Proceedings of the 2016 Internet Measurement Conference. 2016.*

Yarrpbox

- ***Based on yarrp*** [1]
  - network topology discovery
  - probing over 100Kpps
- ***Improvements over Tracebox***
  - ***randomized probing***
    - IP target and TTL domain
    - minimize ICMP rate-limiting
  - ***stateless operation***
    - fire and forget
    - reconstitutes necessary info. from ***replies***

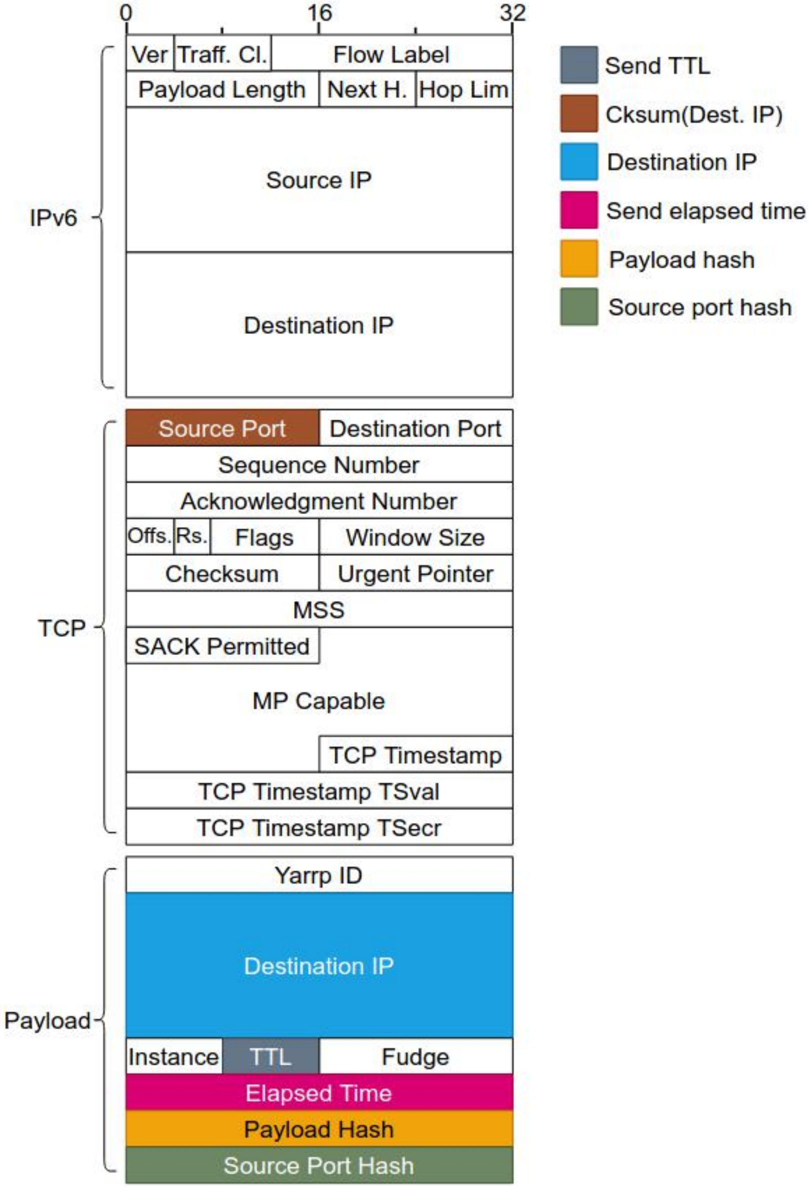

0 16 32

IPv6

Ver | Traff. Cl. | Flow Label

Payload Length | Next H. | Hop Lim

Source IP

Destination IP

TCP

Source Port | Destination Port

Sequence Number

Acknowledgment Number

Offs. | Rs. | Flags | Window Size

Checksum | Urgent Pointer

MSS

SACK Permitted

MP Capable

TCP Timestamp

TCP Timestamp TSval

TCP Timestamp TSecr

Payload

Yarrp ID

Destination IP

Instance | TTL | Fudge

Elapsed Time

Payload Hash

Source Port Hash

⬛ Send TTL

🟫 Cksum(Dest. IP)

🟦 Destination IP

🟪 Send elapsed time

🟧 Payload hash

🟩 Source port hash

[1] *Beverly, Robert. "Yarrp'ing the Internet: Randomized high-speed active topology discovery." Proceedings of the 2016 Internet Measurement Conference. 2016.*

Yarrpbox

- ***Features***
  - builds middlebox detection into yarrp
    - packet crafting, response parsing, file writing
    - result - ***stateless, high speed, load distributive middlebox detection***

Yarrpbox

- ***Features***
  - builds middlebox detection into yarrp
    - packet crafting, response parsing, file writing
    - result - ***stateless, high speed, load distributive middlebox detection***
  - detects
    - interferences ***without*** support from target
    - ***approx. location***

Yarrpbox

- ***Limitations***
  - not all modifications can be detected
    - state storage in packet, ***stateless*** operation

Yarrpbox

- ***Limitations***
  - not all modifications can be detected
    - state storage in packet, ***stateless*** operation
  - middlebox location can only be approx
    - due to ***missing responses***

Results

Results

- ***Contributions***
  - **first Internet-scale** IPv4 middlebox study
    - ***random IP*** from each ***/24***

Results

- ***Contributions***
  - **first Internet-scale** IPv4 middlebox study
    - ***random IP*** from each ***/24***
  - **first IPv6** middlebox study till date
    - BGP - ***100*** IPs per prefix

Results

- ***Contributions***
  - **first Internet-scale** IPv4 middlebox study
    - ***random IP*** from each ***/24***
  - **first IPv6** middlebox study till date
    - BGP - *100* IPs per prefix
  - ***geo-dist.*** scans from 8 VPs (6 continents)

Results

- ***Contributions***
  - **first Internet-scale** IPv4 middlebox study
    - ***random IP*** from each ***/24***
  - **first IPv6** middlebox study till date
    - BGP - ***100*** IPs per prefix
  - ***geo-dist.*** scans from 8 VPs (6 continents)
  - major speed-up
    - approx. **58 years** with **Tracebox**

Results

- ***Contributions***
  - **first Internet-scale** IPv4 middlebox study
    - ***random IP*** from each ***/24***
  - **first IPv6** middlebox study till date
    - BGP - *100* IPs per prefix
  - ***geo-dist.*** scans from 8 VPs (6 continents)
  - major speed-up
    - approx. **58 years** with **Tracebox**
    - under ***10 hrs*** with **Yarrpbox!**

Results-Traces

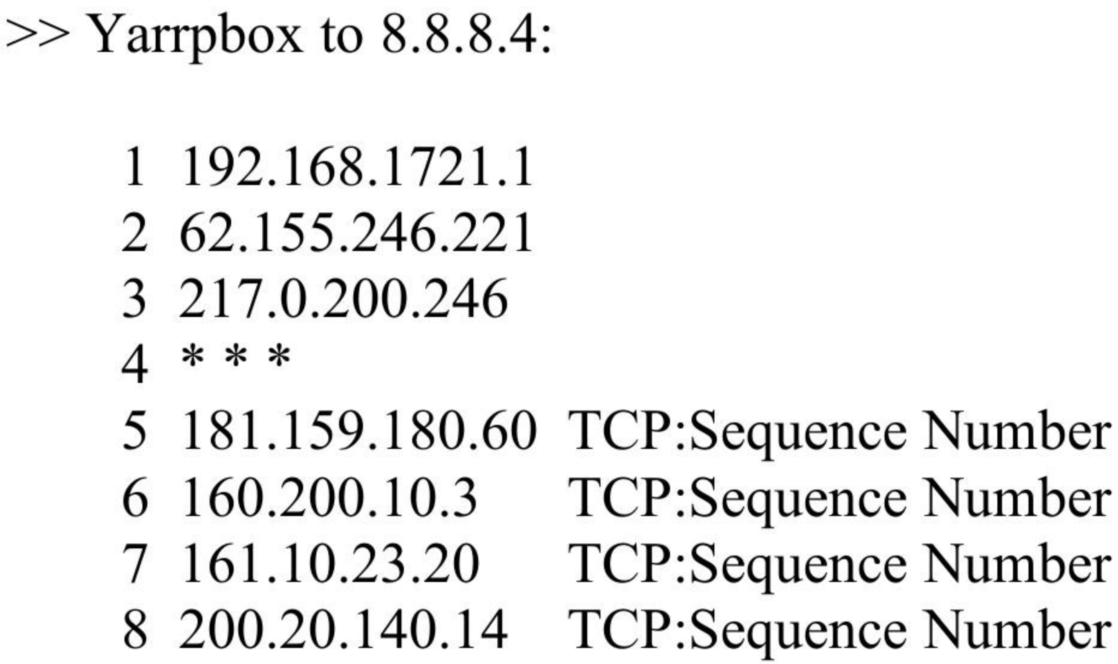

```
>> Yarrpbox to 8.8.8.4:

    1  192.168.1721.1
    2  62.155.246.221
    3  217.0.200.246
    4  * * *
    5  181.159.180.60  TCP:Sequence Number
    6  160.200.10.3    TCP:Sequence Number
    7  161.10.23.20    TCP:Sequence Number
    8  200.20.140.14   TCP:Sequence Number
```

Results-Trace Gaps

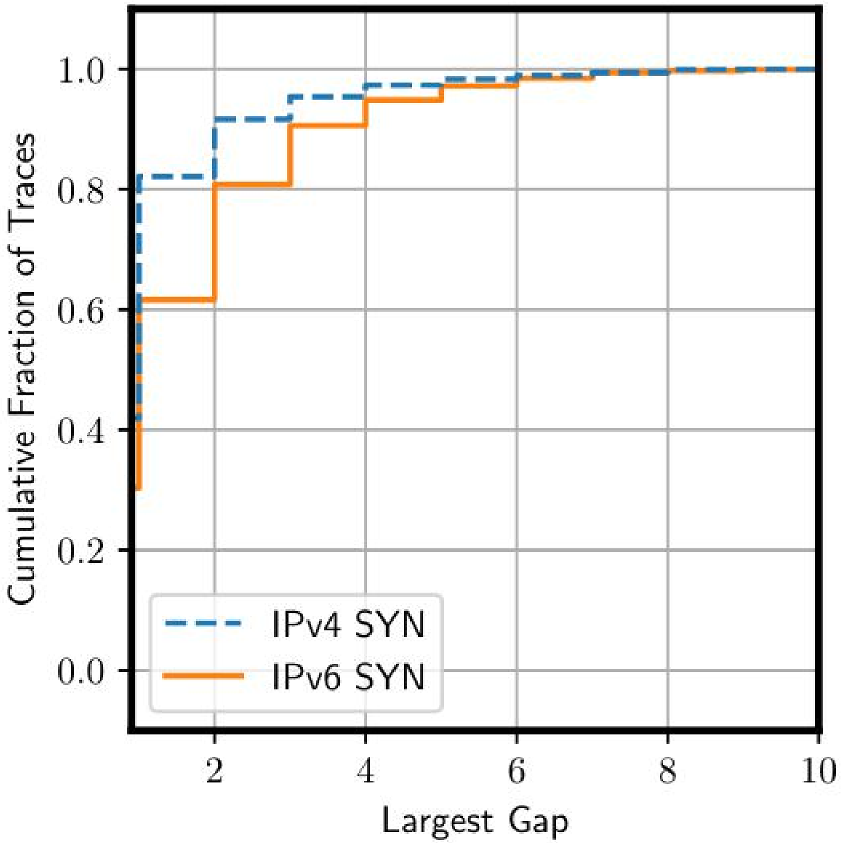

Results-Trace Gaps

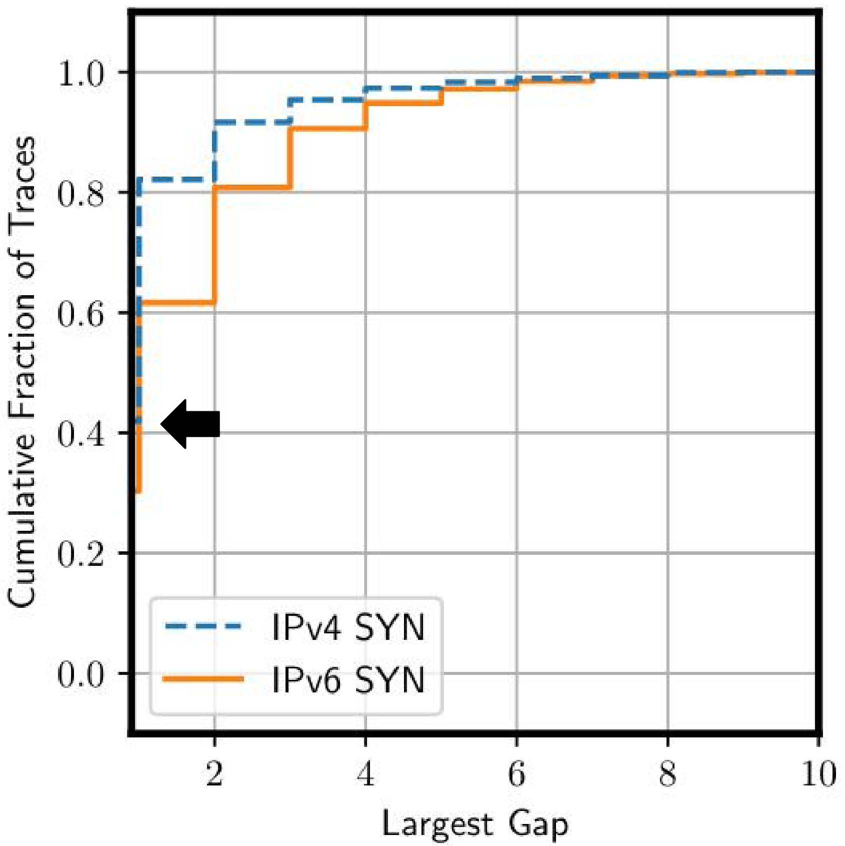

Results-Trace Gaps

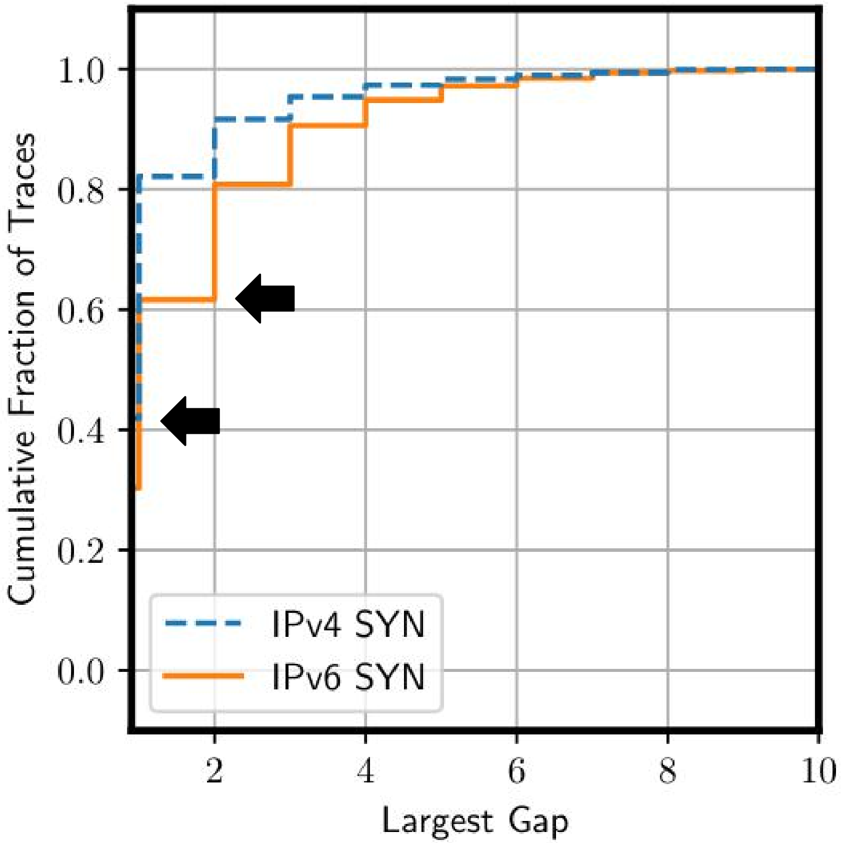

Results-Trace Gaps

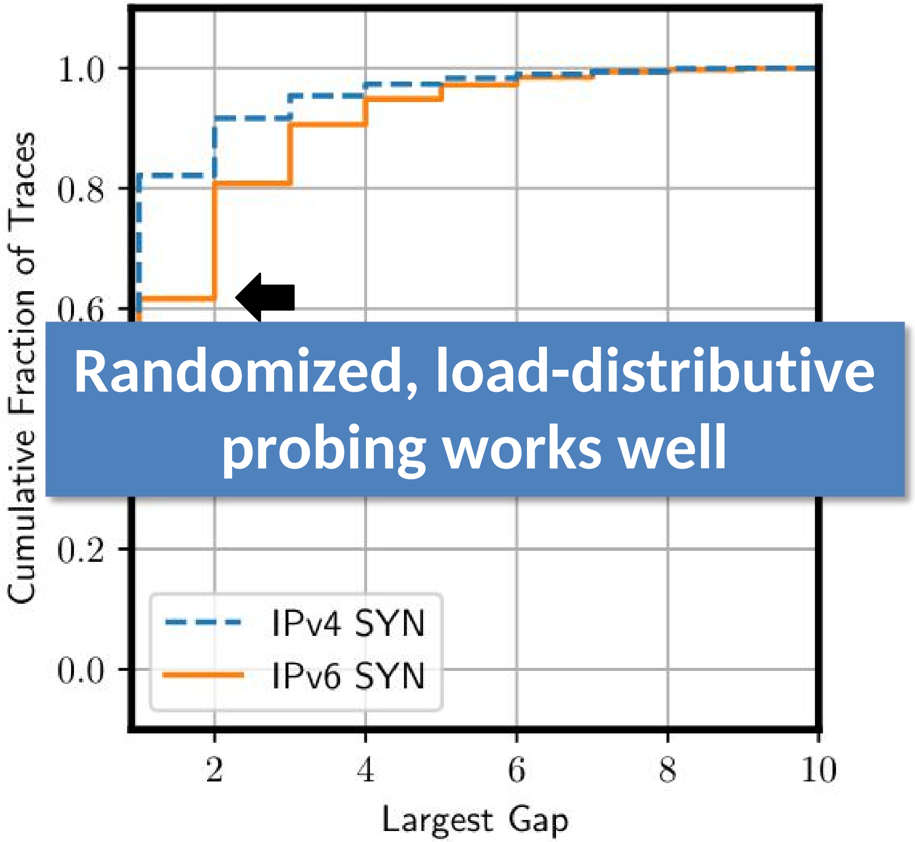

Randomized, load-distributive probing works well

Results-Middlebox Location

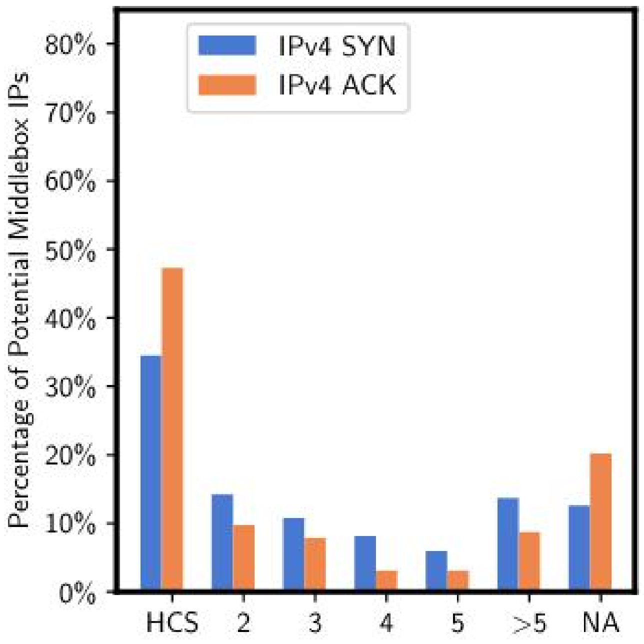

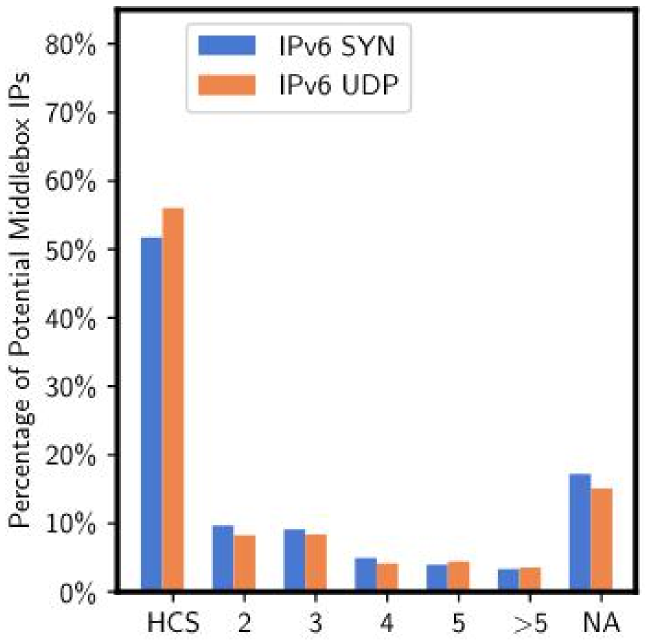

Results-Middlebox Location

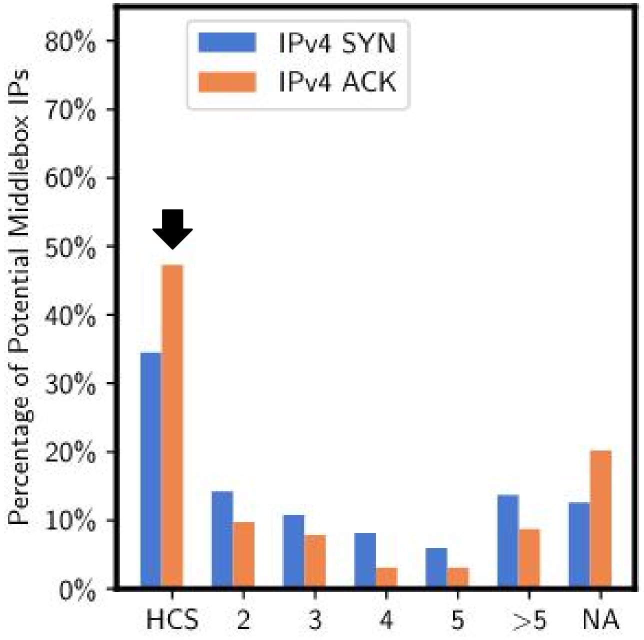

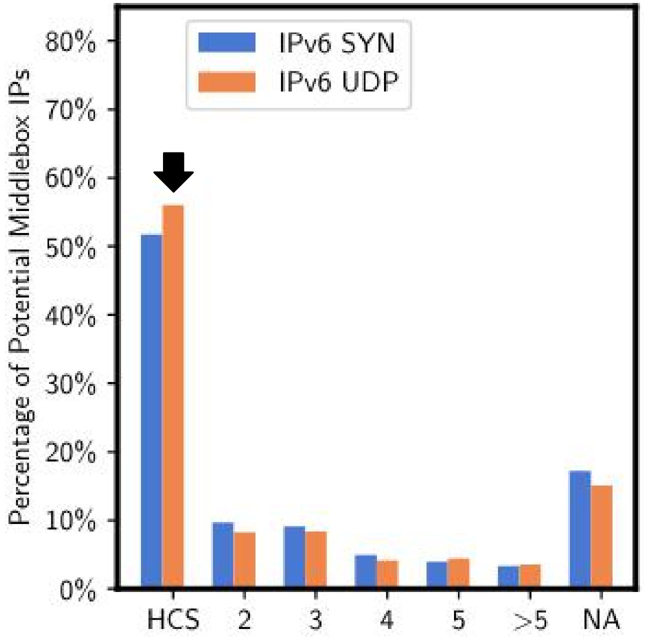

Results-Middlebox Location

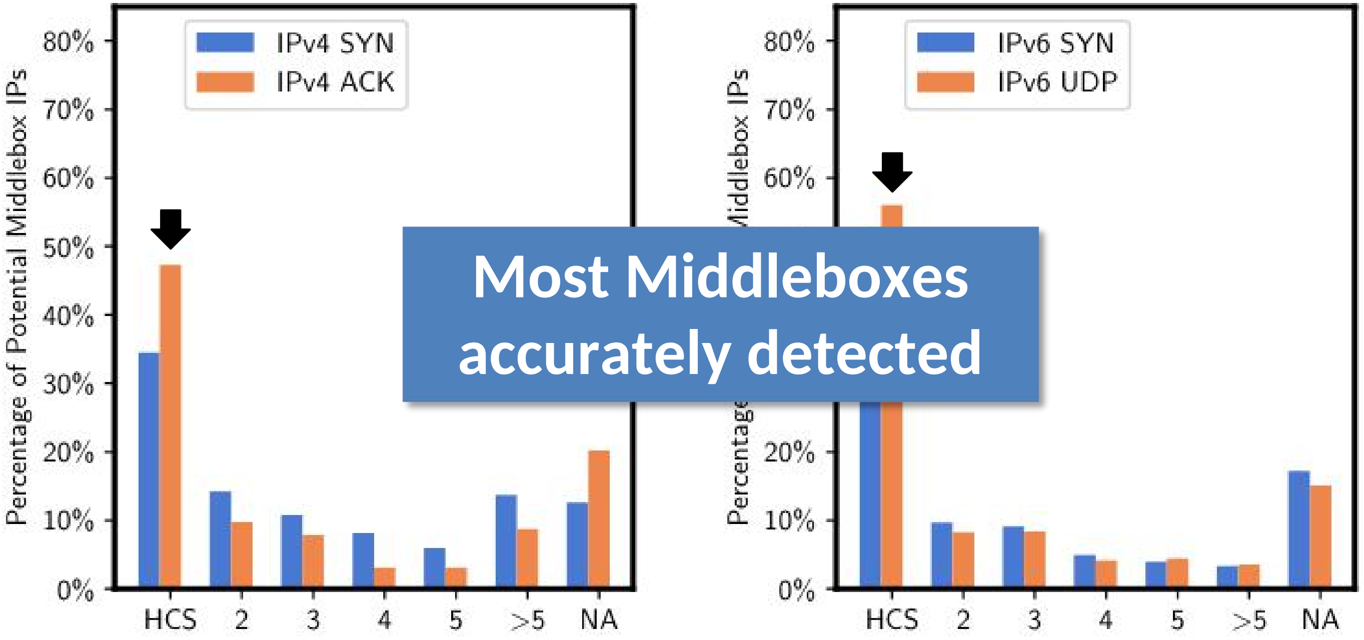

Most Middleboxes accurately detected

Results-Middlebox Interference (HC ASes)

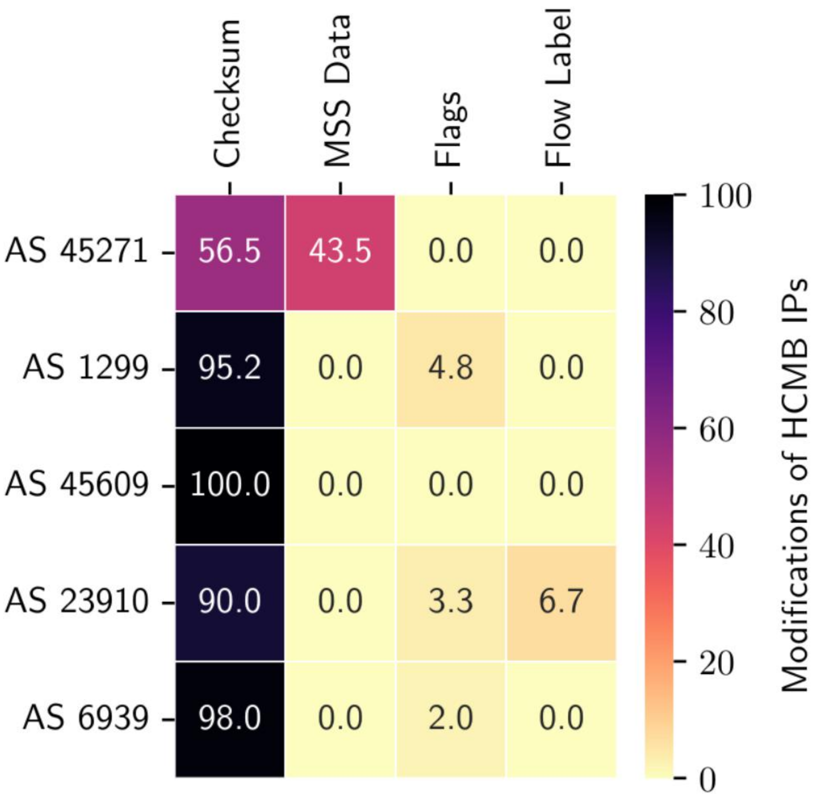

- Highest Confidence IPs (HC) -> ASNs

Results-Middlebox Interference (HC ASes)

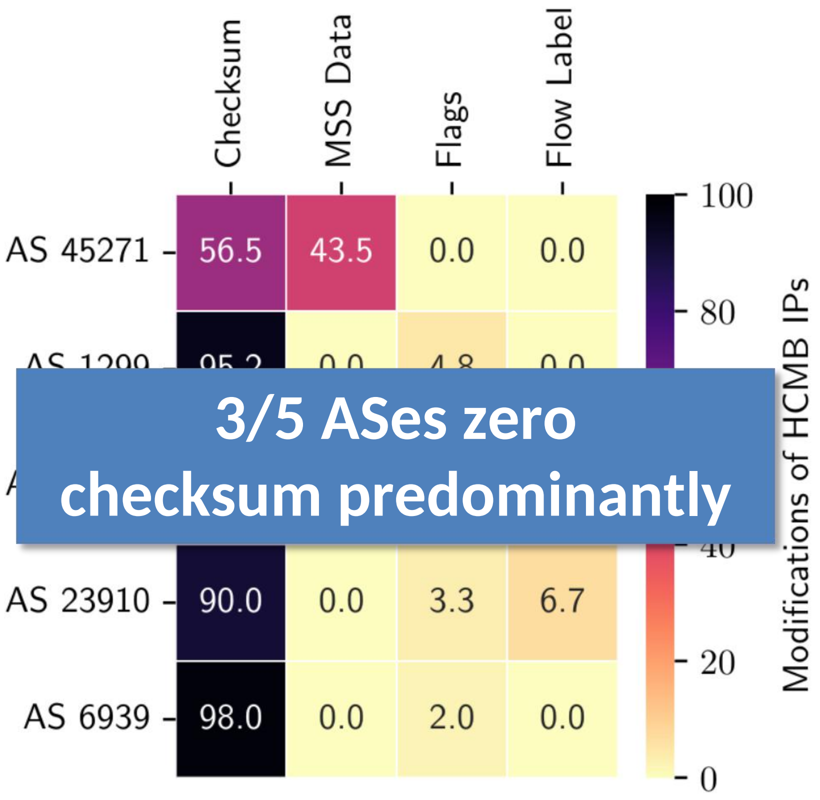

3/5 ASes zero checksum predominantly

- Highest Confidence IPs (HC) -> ASNs

Results-Middlebox Interference

- ***Observations***

Interference (IPv6)	Percentage
TCP Checksum	29.3%
NOP Addition	18.7%
MP CAPABLE Removal	**13.9%**
TCP Sequence Number	12%
TCP Timestamp Removal	**11%**
Sack Permitted Removal	**10.9%**
IP Flow Label	1.7%
MSS Data	1.1%
IP Payload Length	1.1%

Results-Middlebox Interference

- ***Observations***
  - resistance to extensions
    - option removals (**~40%** of all Interferences)

Interference (IPv6)	Percentage
TCP Checksum	29.3%
NOP Addition	18.7%
MP CAPABLE Removal	**13.9%**
TCP Sequence Number	12%
TCP Timestamp Removal	**11%**
Sack Permitted Removal	**10.9%**
IP Flow Label	1.7%
MSS Data	1.1%
IP Payload Length	1.1%

Results-Middlebox Interference

- ***Observations***
  - resistance to extensions
    - option removals (**~40%** of all Interferences)
      - ***~13%*** in ***Tier-1s***

Interference (IPv6)	Percentage
TCP Checksum	29.3%
NOP Addition	18.7%
MP CAPABLE Removal	**13.9%**
TCP Sequence Number	12%
TCP Timestamp Removal	**11%**
Sack Permitted Removal	**10.9%**
IP Flow Label	1.7%
MSS Data	1.1%
IP Payload Length	1.1%

Results-Middlebox Interference

- ***Observations***
    - resistance to extensions
        - option removals (**~40%** of all Interferences)
            - ***~13%*** in ***Tier-1s***
                - ***~28%*** in Tier-1s ***(IPv4)***

Interference (IPv6)	Percentage
TCP Checksum	29.3%
NOP Addition	18.7%
MP CAPABLE Removal	**13.9%**
TCP Sequence Number	12%
TCP Timestamp Removal	**11%**
Sack Permitted Removal	**10.9%**
IP Flow Label	1.7%
MSS Data	1.1%
IP Payload Length	1.1%

Results-Middlebox Interference

- ***Observations***
  - resistance to extensions
    - option removals (**~40%** of all Interferences)
      - ***~13%*** in ***Tier-1s***
        - ***~28%*** in Tier-1s ***(IPv4)***
      - ***< 1%*** MB-affected paths

Interference (IPv6)	Percentage
TCP Checksum	29.3%
NOP Addition	18.7%
MP CAPABLE Removal	**13.9%**
TCP Sequence Number	12%
TCP Timestamp Removal	**11%**
Sack Permitted Removal	**10.9%**
IP Flow Label	1.7%
MSS Data	1.1%
IP Payload Length	1.1%

Results-Middlebox Interference

- ***Observations***
  - resistance to extensions
    - option removals (**~40%** of all Interferences)
      - ***~13%*** in ***Tier-1s***
        - ***~28%*** in Tier-1s ***(IPv4)***
      - ***< 1%*** MB-affected paths
        - ***~14%*** MB-affected paths (***IPv4***)

Interference (IPv6)	Percentage
TCP Checksum	29.3%
NOP Addition	18.7%
MP CAPABLE Removal	**13.9%**
TCP Sequence Number	12%
TCP Timestamp Removal	**11%**
Sack Permitted Removal	**10.9%**
IP Flow Label	1.7%
MSS Data	1.1%
IP Payload Length	1.1%

Results-Middlebox Interference

- ***Observations***
    - resistance to extensions
        - option removals (***~40%*** of all Interferences)
            - ***~13%*** in ***Tier-1s***
                - ***~28%*** i
            - ***< 1%*** MB-affecte
                - ***~14%*** MB-affected paths (***IPv4***)

Lower path-brokenness in IPv6!

Interference (IPv6)	Percentage
TCP Checksum	29.3%
NOP Addition	18.7%
MP CAPABLE Removal	**13.9%**
TCP Sequence Number	12%
TCP Timestamp Removal	**11%**
Sack Permitted Removal	**10.9%**
IP Flow Label	1.7%
MSS Data	1.1%
IP Payload Length	1.1%

Results-Alias Resolution

HC IPs	Map to NS Sets	NS MBs
2.4k (11.4%)	2.1k (87.5%)	132

- SNMPv3 dataset
    - IPv4 + IPv6 alias sets

Results-Alias Resolution

HC IPs	Map to NS Sets	NS MBs
2.4k (11.4%)	2.1k (87.5%)	132

- SNMPv3 dataset
  - IPv4 + IPv6 alias sets
- Multiple-alias MBs (***NS MBs***)
  - **~20%** dual-stack

Results-Alias Resolution

HC IPs	Map to NS Sets	NS MBs
2.4k (11.4%)	2.1k (87.5%)	132

- SNMPv3 dataset
  - IPv4 + IPv6 alias sets
- Multiple-alias MBs (***NS MBs***)
  - ***~20%*** dual-stack
- Vendors
  - ***~54%*** (singleton + non-singleton) MBs -> ***Cisco*** devices
  - ~20% -> Juniper

Results-Alias Resolution

HC IPs	Map to NS Sets	NS MBs
2.4k (11.4%)	2.1k (87.5%)	132

MBs from lower concentration of vendors!

- SNMPv3 dataset
  - IPv4 + IPv6 alias sets
- Multiple-alias MBs (***NS MBs***)
  - ***~20%*** dual-stack
- Vendors
  - ***~54%*** (singleton + non-singleton) MBs -> ***Cisco*** devices
  - ~20% -> Juniper

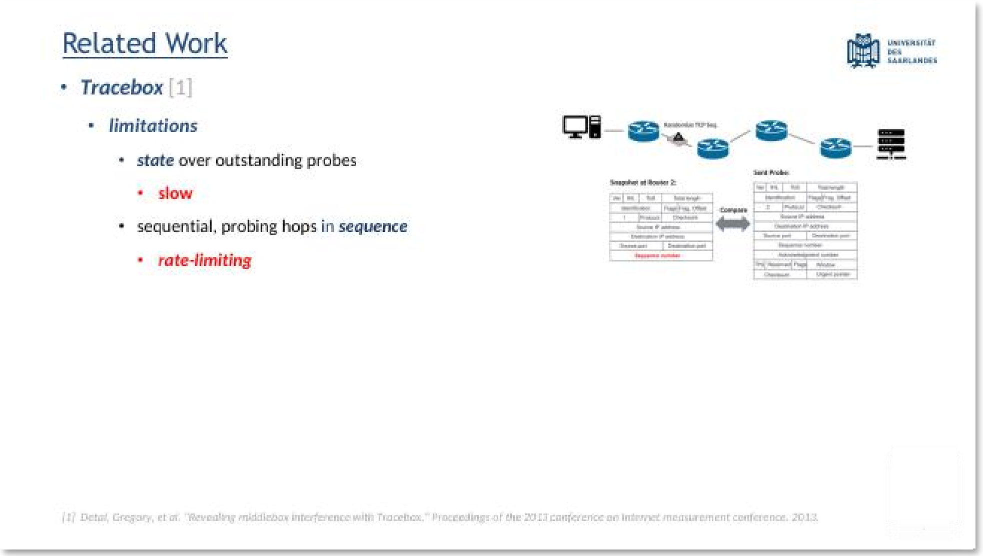

Related Work

- ***Tracebox*** [1]
  - ***limitations***
    - ***state*** over outstanding probes
      - slow
    - sequential, probing hops in ***sequence***
      - ***rate-limiting***

Randomize TCP Seq.
Snapshot at Router 2:
Ver | IHL | ToS | Total length
Identification | Flags | Frag. Offset
1 | Protocol | Checksum
Source IP address
Destination IP address
Source port | Destination port
Sequence number
Compare
Send Probe:
Ver | IHL | ToS | Total length
Identification | Flags | Frag. Offset
2 | Protocol | Checksum
Source IP address
Destination IP address
Source port | Destination port
Sequence number
Acknowledgement number
THL | Reserved | Flags | Window
Checksum | Urgent pointer

[1] Detal, Gregory, et al. "Revealing middlebox interference with Tracebox." Proceedings of the 2013 conference on Internet measurement conference. 2013.

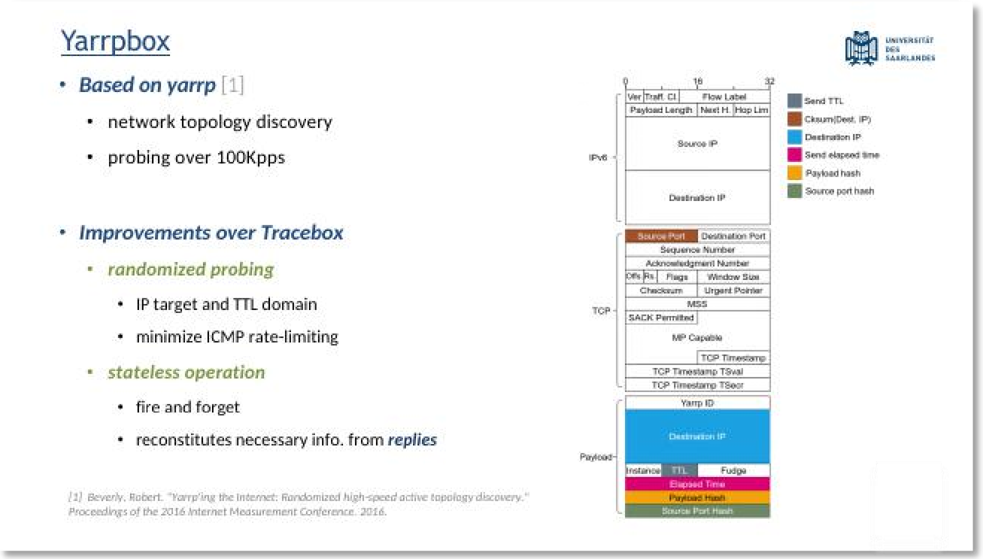

Yarrpbox

- ***Based on yarrp*** [1]
  - network topology discovery
  - probing over 100Kpps
- ***Improvements over Tracebox***
  - ***randomized probing***
    - IP target and TTL domain
    - minimize ICMP rate-limiting
  - ***stateless operation***
    - fire and forget
    - reconstitutes necessary info. from ***replies***

[1] Beverly, Robert. "Yarrp'ing the Internet: Randomized high-speed active topology discovery." Proceedings of the 2016 Internet Measurement Conference. 2016.

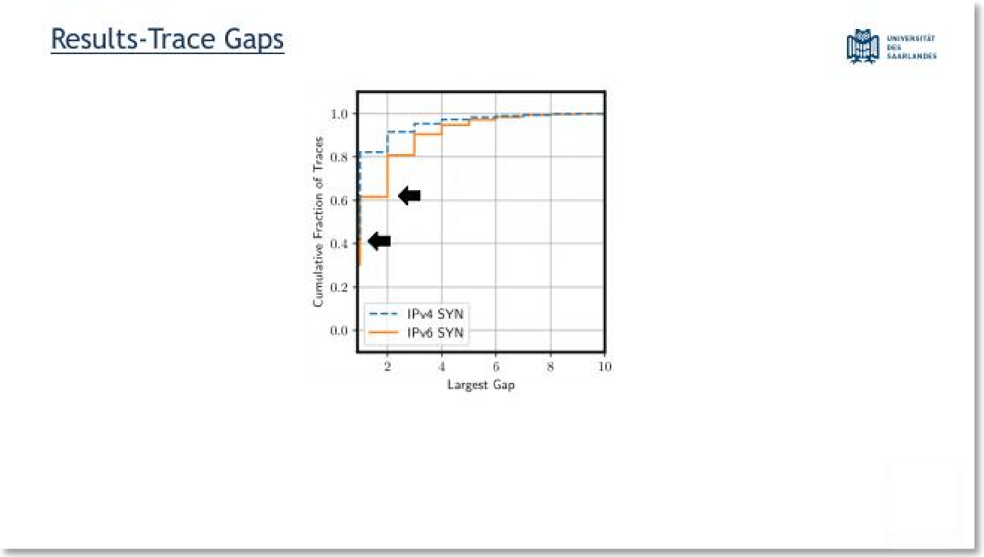

Results-Trace Gaps

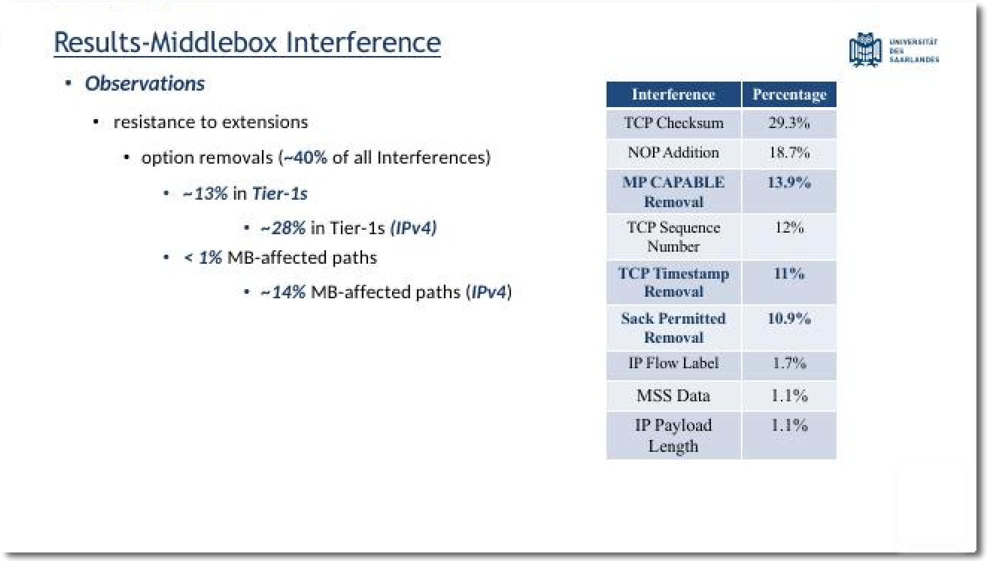

Results-Middlebox Interference

- ***Observations***
  - resistance to extensions
    - option removals (***~40%*** of all Interferences)
      - ***~13%*** in ***Tier-1s***
        - ***~28%*** in Tier-1s ***(IPv4)***
      - ***< 1%*** MB-affected paths
        - ***~14%*** MB-affected paths (***IPv4***)

Interference	Percentage
TCP Checksum	29.3%
NOP Addition	18.7%
MP CAPABLE Removal	**13.9%**
TCP Sequence Number	12%
TCP Timestamp Removal	**11%**
Sack Permitted Removal	**10.9%**
IP Flow Label	1.7%
MSS Data	1.1%
IP Payload Length	1.1%

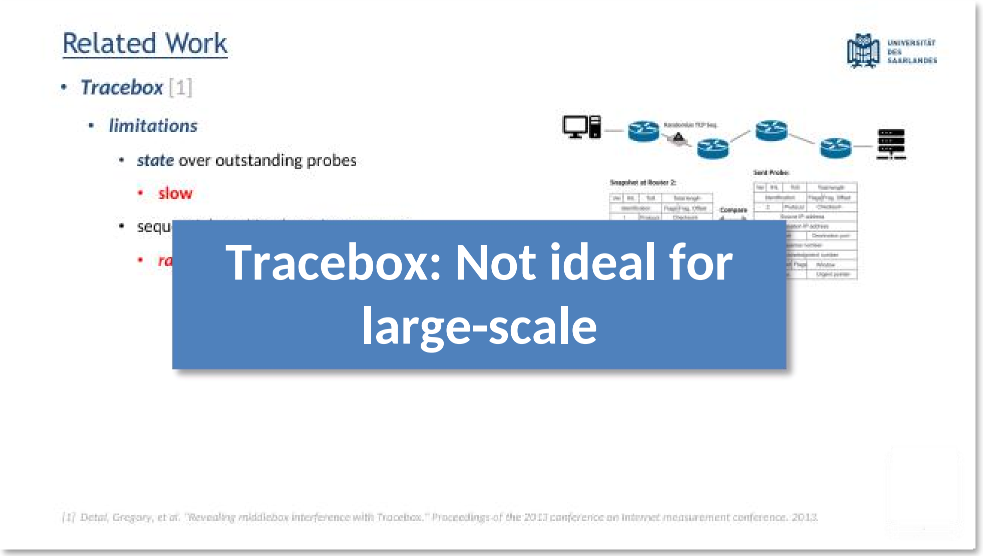

Related Work

- ***Tracebox*** [1]
  - *limitations*
    - *state* over outstanding probes
      - slow
    - sequ
      - ra

Tracebox: Not ideal for large-scale

[1] Detal, Gregory, et al. "Revealing middlebox interference with Tracebox." Proceedings of the 2013 conference on Internet measurement conference. 2013.

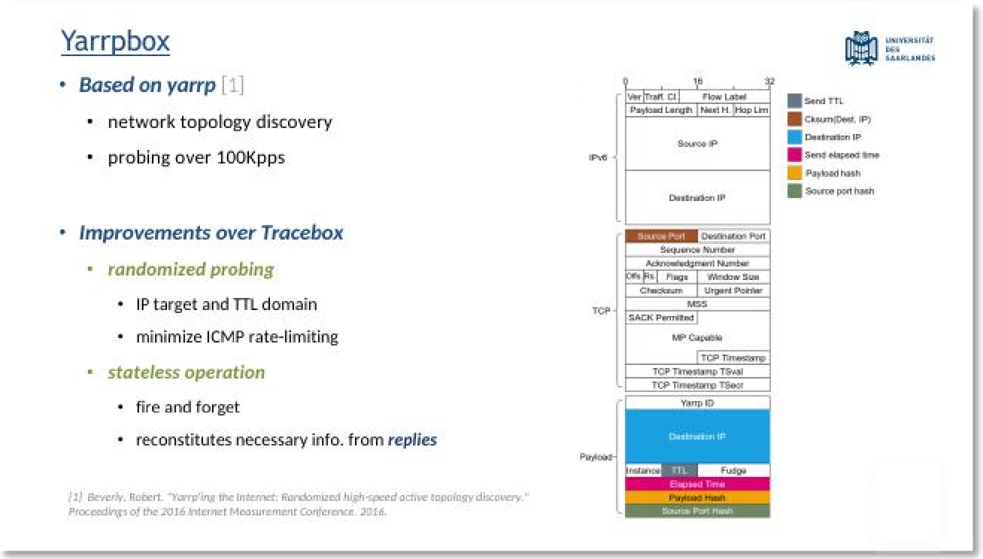

Yarrpbox

- ***Based on yarrp*** [1]
  - network topology discovery
  - probing over 100Kpps
- ***Improvements over Tracebox***
  - *randomized probing*
    - IP target and TTL domain
    - minimize ICMP rate-limiting
  - *stateless operation*
    - fire and forget
    - reconstitutes necessary info. from ***replies***

[1] Beverly, Robert. "Yarrp'ing the Internet: Randomized high-speed active topology discovery." Proceedings of the 2016 Internet Measurement Conference. 2016.

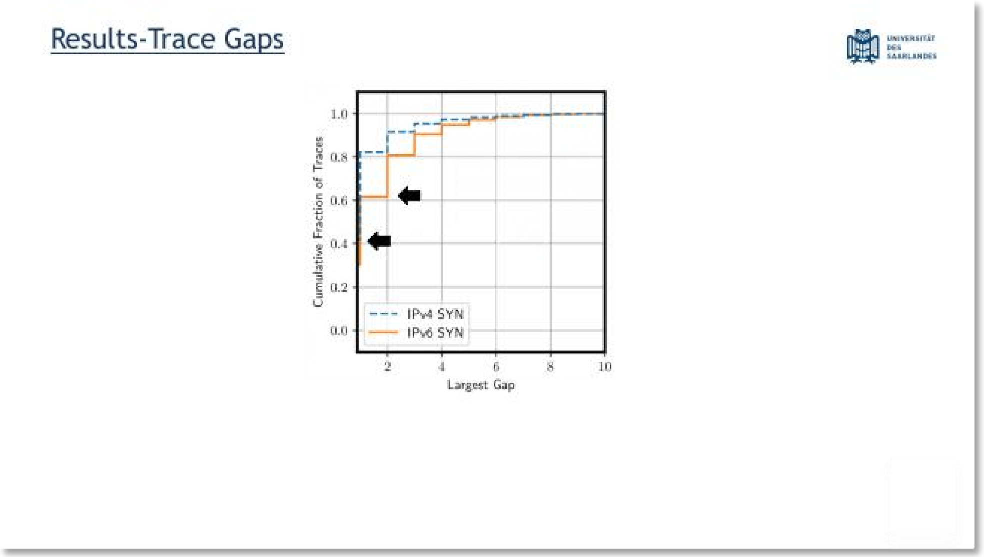

Results-Trace Gaps

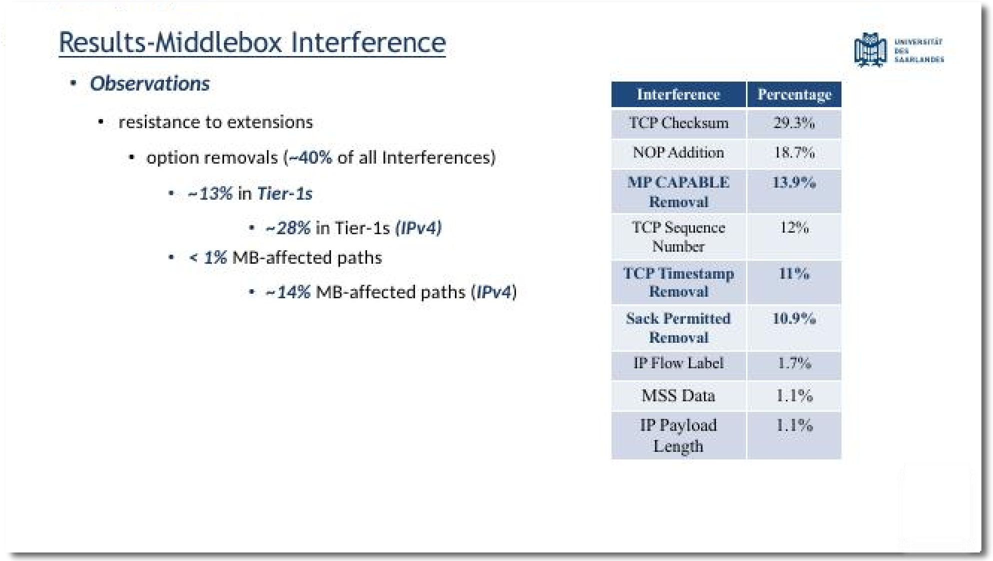

Results-Middlebox Interference

- ***Observations***
  - resistance to extensions
    - option removals (*~40%* of all Interferences)
      - *~13%* in ***Tier-1s***
        - *~28%* in Tier-1s *(IPv4)*
      - *< 1%* MB-affected paths
        - *~14%* MB-affected paths *(IPv4)*

Interference	Percentage
TCP Checksum	29.3%
NOP Addition	18.7%
MP CAPABLE Removal	**13.9%**
TCP Sequence Number	12%
TCP Timestamp Removal	**11%**
Sack Permitted Removal	**10.9%**
IP Flow Label	1.7%
MSS Data	1.1%
IP Payload Length	1.1%

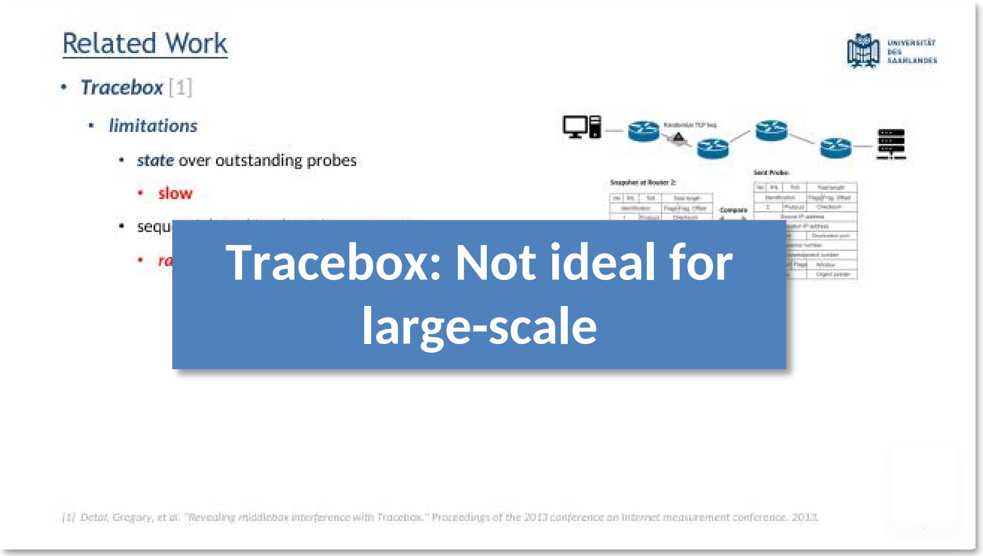

Related Work

- *Tracebox* [1]
  - *limitations*
    - *state* over outstanding probes
      - slow
    - sequ
      - ra

Tracebox: Not ideal for large-scale

[1] Detal, Gregory, et al. "Revealing middlebox interference with Tracebox." Proceedings of the 2013 conference on Internet measurement conference. 2013.

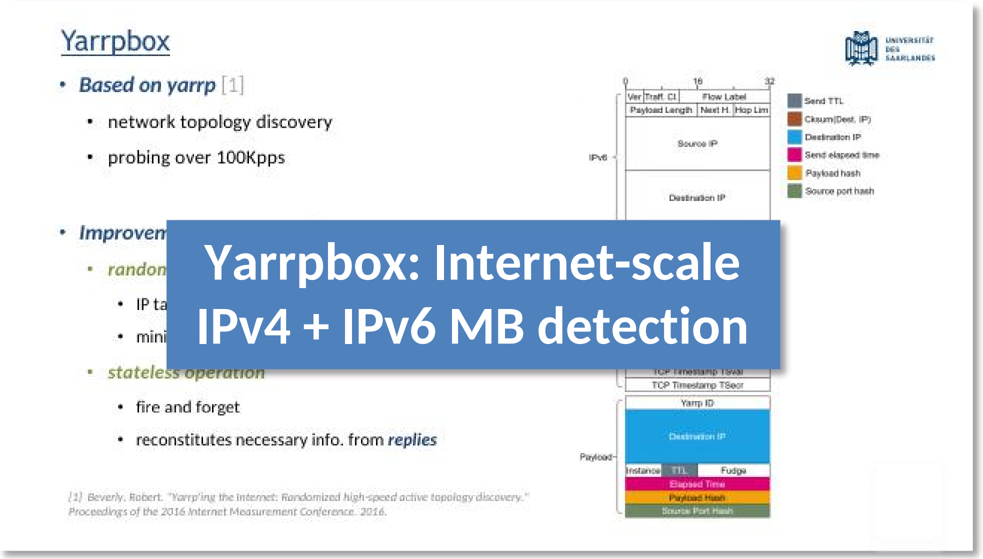

Yarrpbox

- *Based on yarrp* [1]
  - network topology discovery
  - probing over 100Kpps
- *Improvem*
  - *randon*
    - IP ta
    - mini
  - *stateless operation*
    - fire and forget
    - reconstitutes necessary info. from *replies*

Yarrpbox: Internet-scale IPv4 + IPv6 MB detection

[1] Beverly, Robert. "Yarrp'ing the Internet: Randomized high-speed active topology discovery." Proceedings of the 2016 Internet Measurement Conference. 2016.

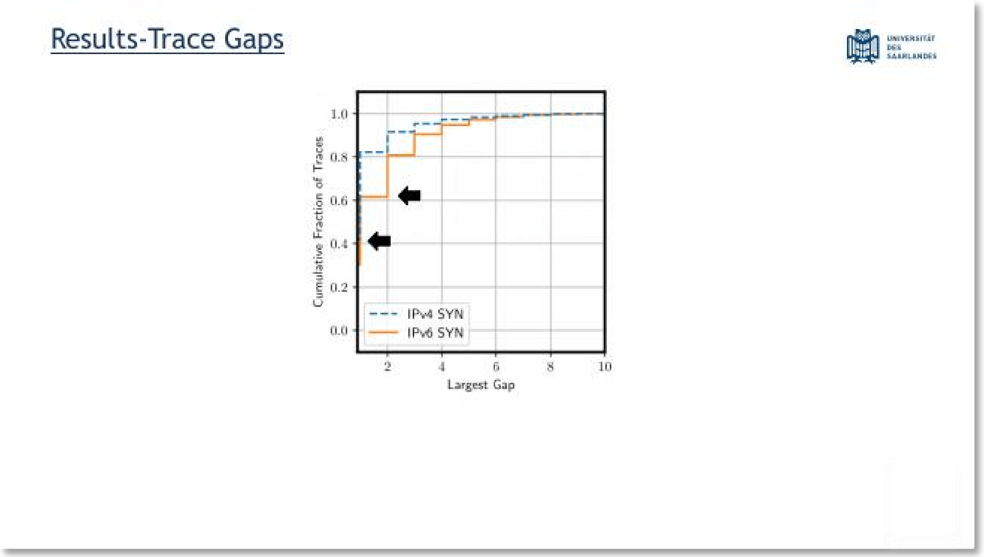

Results-Trace Gaps

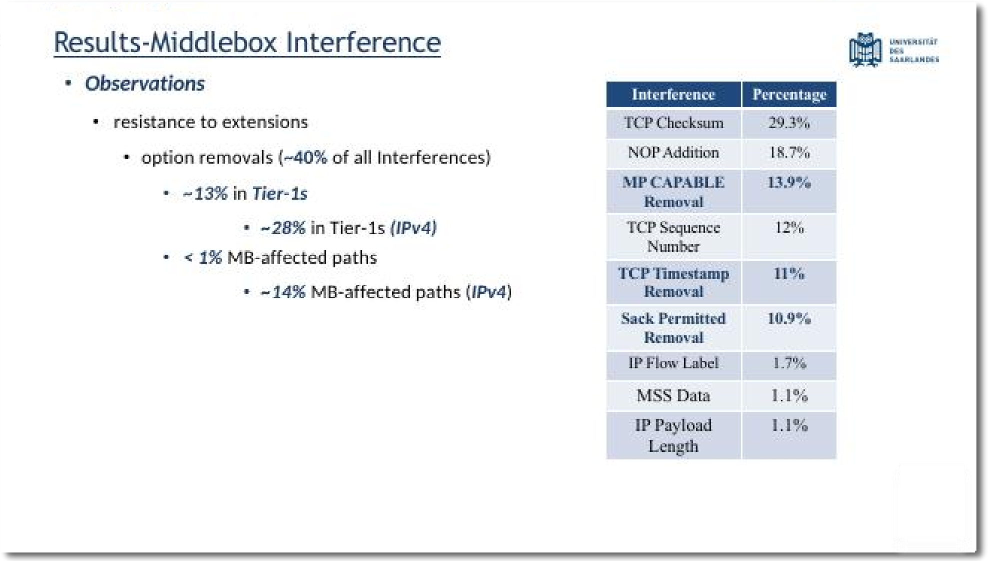

Results-Middlebox Interference

- *Observations*
  - resistance to extensions
    - option removals (*~40%* of all Interferences)
      - *~13%* in *Tier-1s*
        - *~28%* in Tier-1s *(IPv4)*
      - *< 1%* MB-affected paths
        - *~14%* MB-affected paths (*IPv4*)

Interference	Percentage
TCP Checksum	29.3%
NOP Addition	18.7%
MP CAPABLE Removal	13.9%
TCP Sequence Number	12%
TCP Timestamp Removal	11%
Sack Permitted Removal	10.9%
IP Flow Label	1.7%
MSS Data	1.1%
IP Payload Length	1.1%

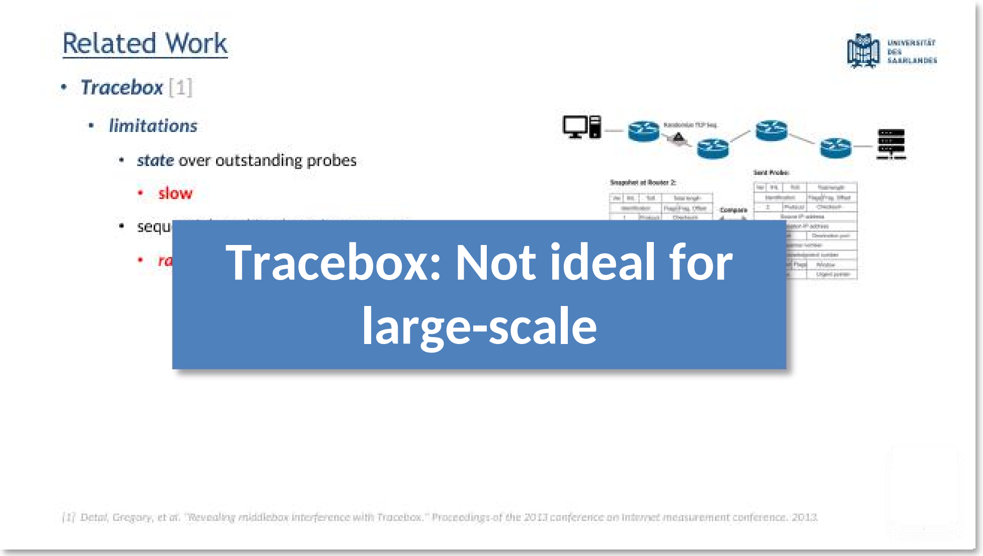

Related Work

- *Tracebox* [1]
  - *limitations*
    - *state* over outstanding probes
      - slow
    - sequ
      - ra

Tracebox: Not ideal for large-scale

[1] Detal, Gregory, et al. "Revealing middlebox interference with Tracebox." Proceedings of the 2013 conference on Internet measurement conference. 2013.

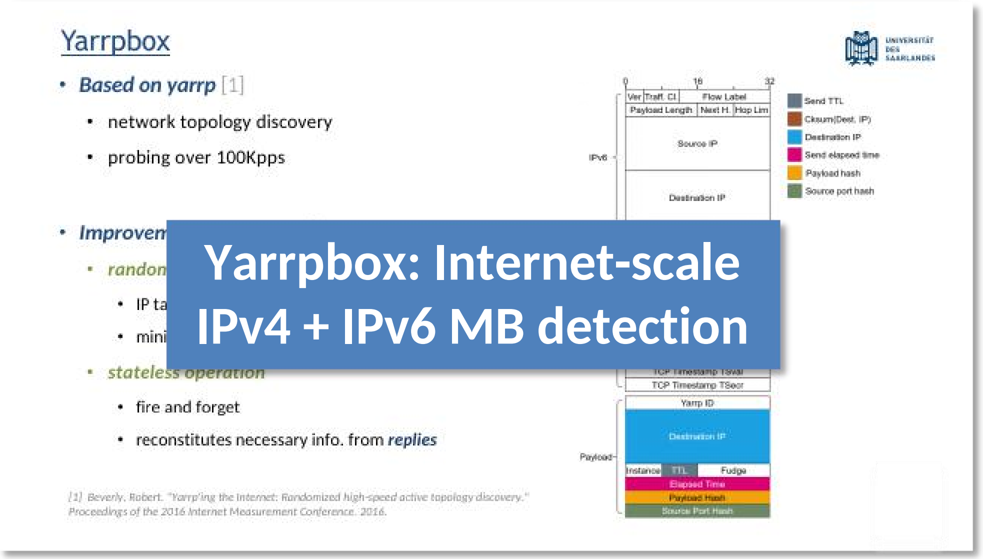

Yarrpbox

- *Based on yarrp* [1]
  - network topology discovery
  - probing over 100Kpps
- *Improvem*
  - *randon*
    - IP ta
    - mini
  - *stateless operation*
    - fire and forget
    - reconstitutes necessary info. from *replies*

Yarrpbox: Internet-scale IPv4 + IPv6 MB detection

[1] Beverly, Robert. "Yarrp'ing the Internet: Randomized high-speed active topology discovery." Proceedings of the 2016 Internet Measurement Conference. 2016.

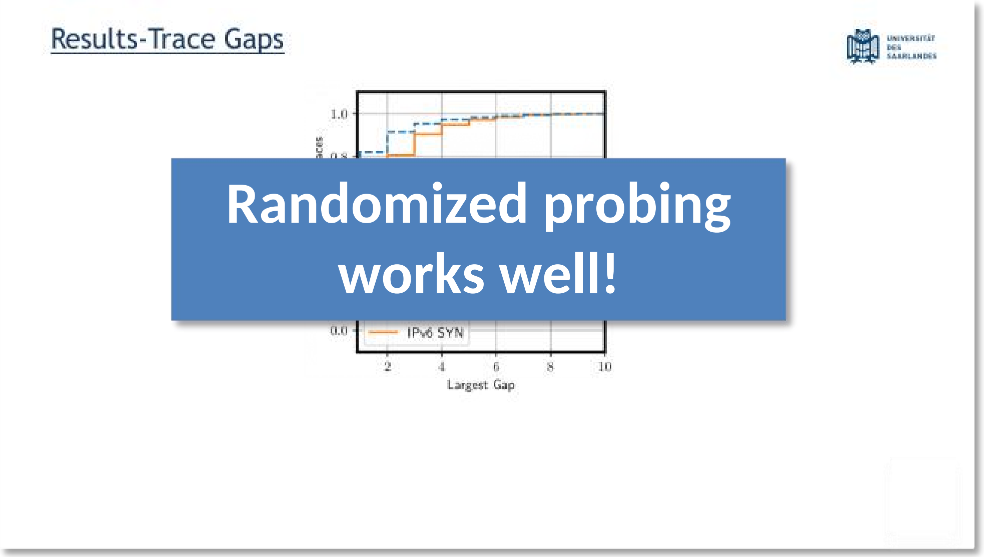

Results-Trace Gaps

Randomized probing works well!

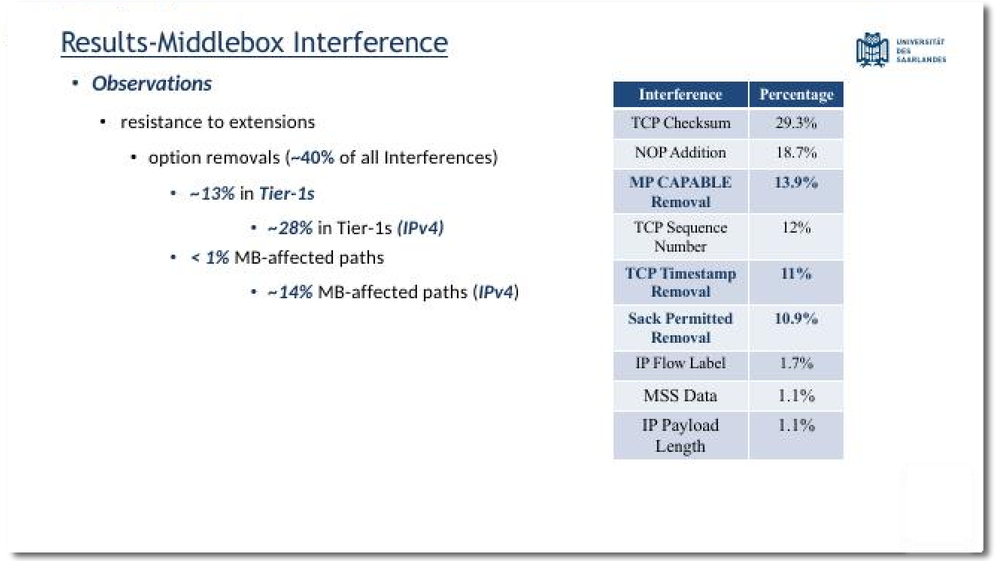

Results-Middlebox Interference

- *Observations*
  - resistance to extensions
    - option removals (*~40%* of all Interferences)
      - *~13%* in *Tier-1s*
        - *~28%* in Tier-1s *(IPv4)*
      - *< 1%* MB-affected paths
        - *~14%* MB-affected paths (*IPv4*)

Interference	Percentage
TCP Checksum	29.3%
NOP Addition	18.7%
MP CAPABLE Removal	13.9%
TCP Sequence Number	12%
TCP Timestamp Removal	11%
Sack Permitted Removal	10.9%
IP Flow Label	1.7%
MSS Data	1.1%
IP Payload Length	1.1%

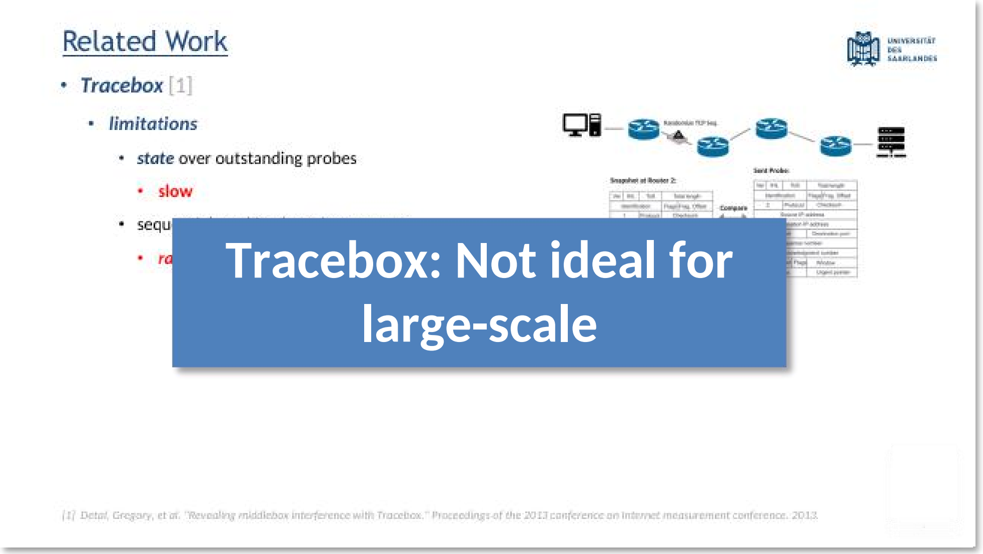

Related Work

- *Tracebox* [1]
  - *limitations*
    - *state* over outstanding probes
      - slow
    - sequ
      - ra

Tracebox: Not ideal for large-scale

[1] Detal, Gregory, et al. "Revealing middlebox interference with Tracebox." Proceedings of the 2013 conference on Internet measurement conference. 2013.

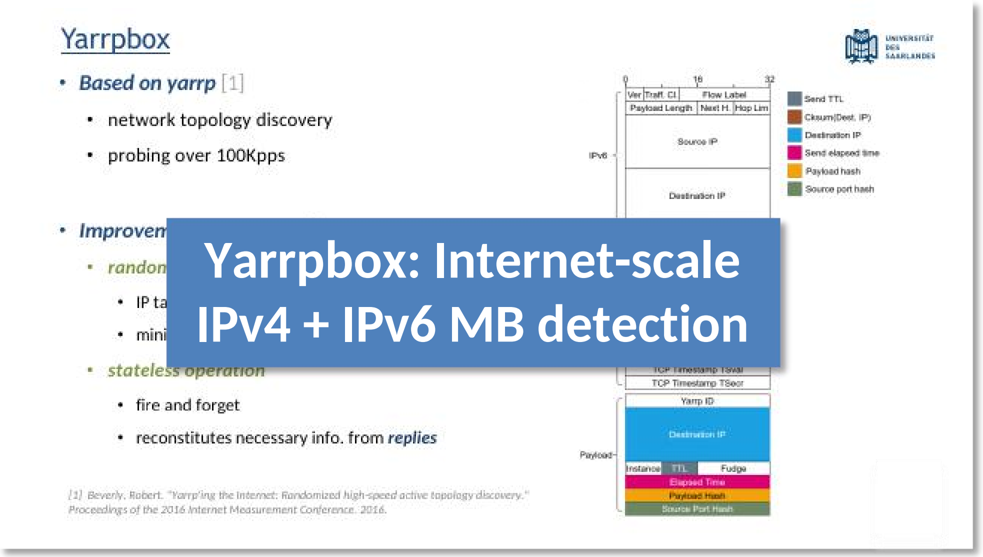

Yarrpbox

- *Based on yarrp* [1]
  - network topology discovery
  - probing over 100Kpps
- *Improvem*
  - *randon*
    - IP ta
    - mini
  - *stateless operation*
    - fire and forget
    - reconstitutes necessary info. from *replies*

Yarrpbox: Internet-scale IPv4 + IPv6 MB detection

[1] Beverly, Robert. "Yarrp'ing the Internet: Randomized high-speed active topology discovery." Proceedings of the 2016 Internet Measurement Conference. 2016.

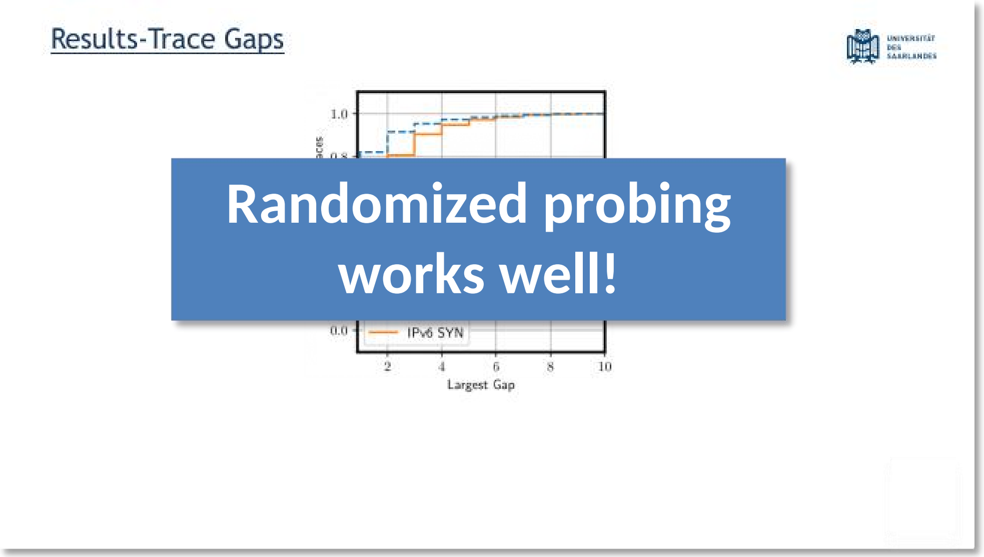

Results-Trace Gaps

Randomized probing works well!

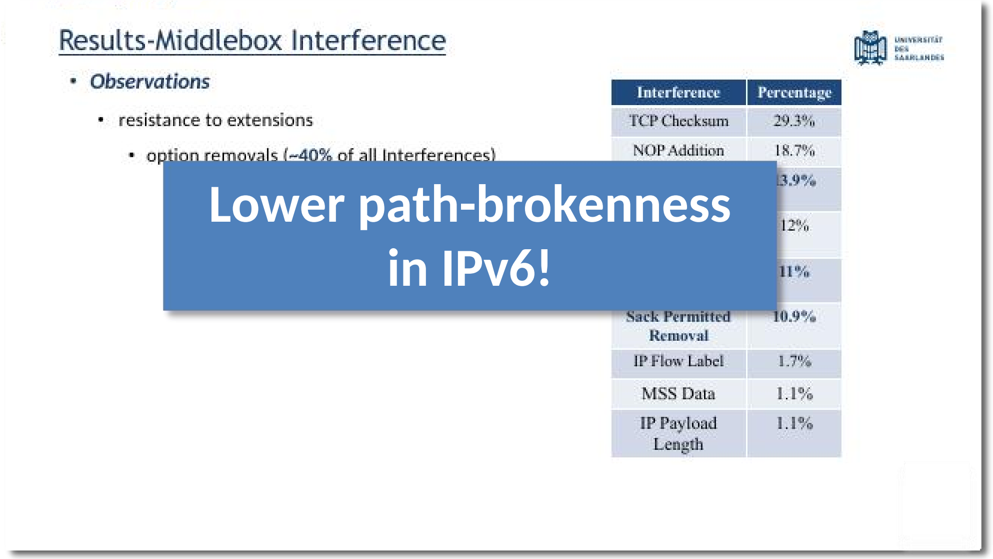

Results-Middlebox Interference

- *Observations*
  - resistance to extensions
    - option removals (~40% of all Interferences)

Interference	Percentage
TCP Checksum	29.3%
NOP Addition	18.7%
[illegible]	13.9%
[illegible]	12%
[illegible]	11%
Sack Permitted Removal	10.9%
IP Flow Label	1.7%
MSS Data	1.1%
IP Payload Length	1.1%

Lower path-brokenness in IPv6!

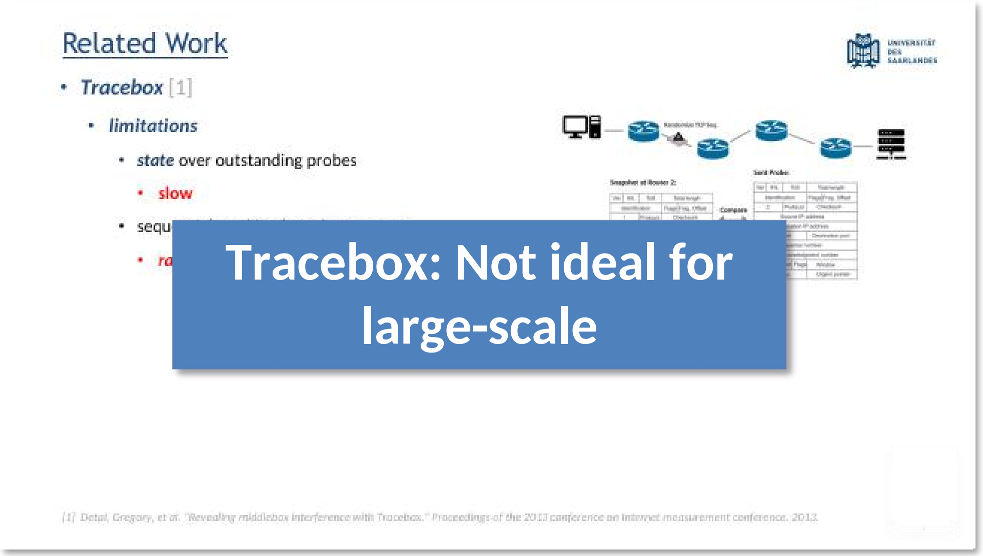

Related Work

- *Tracebox* [1]
  - *limitations*
    - *state* over outstanding probes
      - slow
    - sequ
      - ra

[1] Detal, Gregory, et al. "Revealing middlebox interference with Tracebox." Proceedings of the 2013 conference on Internet measurement conference. 2013.

Tracebox: Not ideal for large-scale

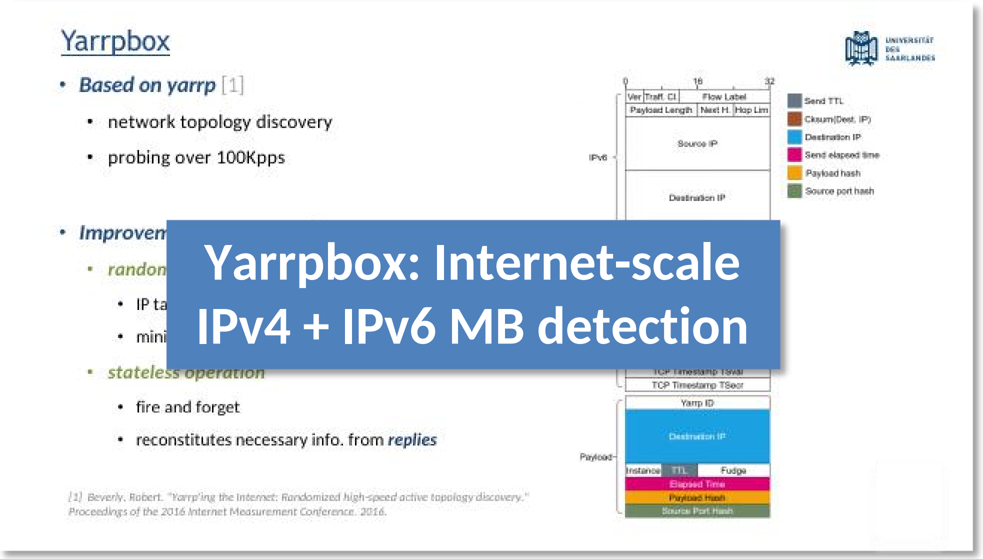

Yarrpbox

- *Based on yarrp* [1]
  - network topology discovery
  - probing over 100Kpps
- *Improvem*
  - *randon*
    - IP ta
    - mini
  - *stateless operation*
    - fire and forget
    - reconstitutes necessary info. from *replies*

[1] Beverly, Robert. "Yarrp'ing the Internet: Randomized high-speed active topology discovery." Proceedings of the 2016 Internet Measurement Conference. 2016.

Yarrpbox: Internet-scale IPv4 + IPv6 MB detection

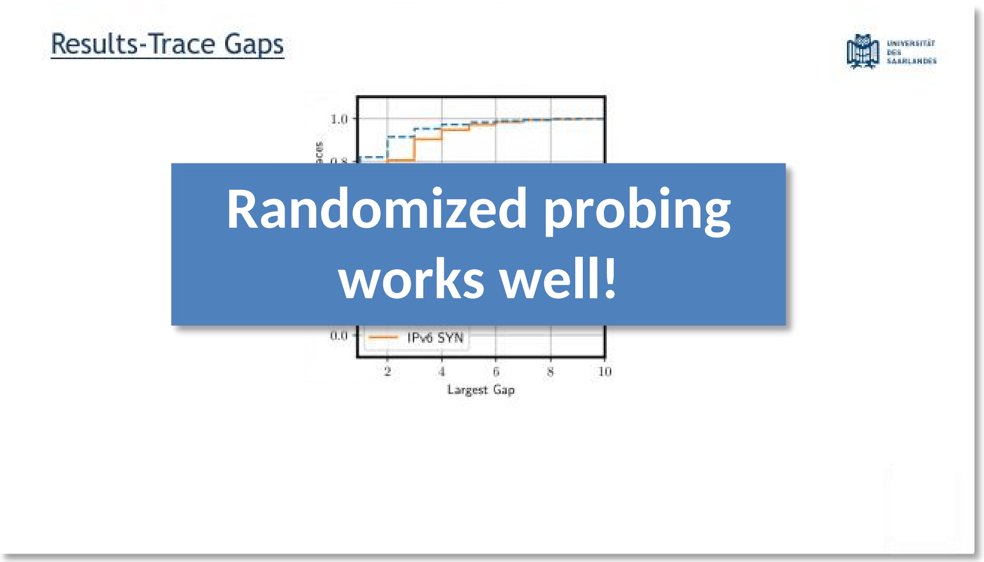

Results-Trace Gaps

Randomized probing works well!

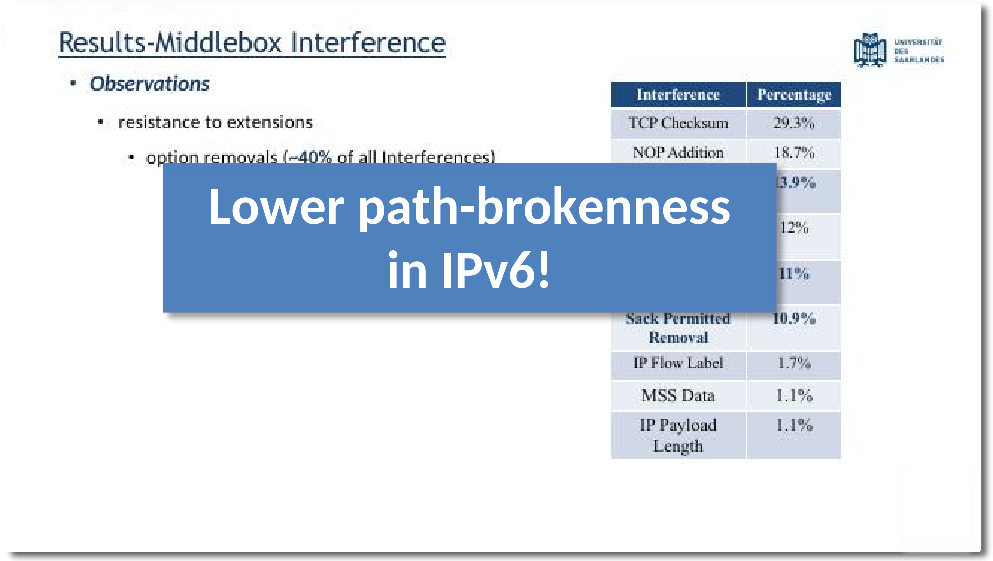

Results-Middlebox Interference

- *Observations*
  - resistance to extensions
    - option removals (~40% of all Interferences)

Interference	Percentage
TCP Checksum	29.3%
NOP Addition	18.7%
[illegible]	13.9%
[illegible]	12%
[illegible]	11%
Sack Permitted Removal	10.9%
IP Flow Label	1.7%
MSS Data	1.1%
IP Payload Length	1.1%

Lower path-brokenness in IPv6!

fhilal@mpi-inf.mpg.de

UNIVERSITÄT DES SAARLANDES

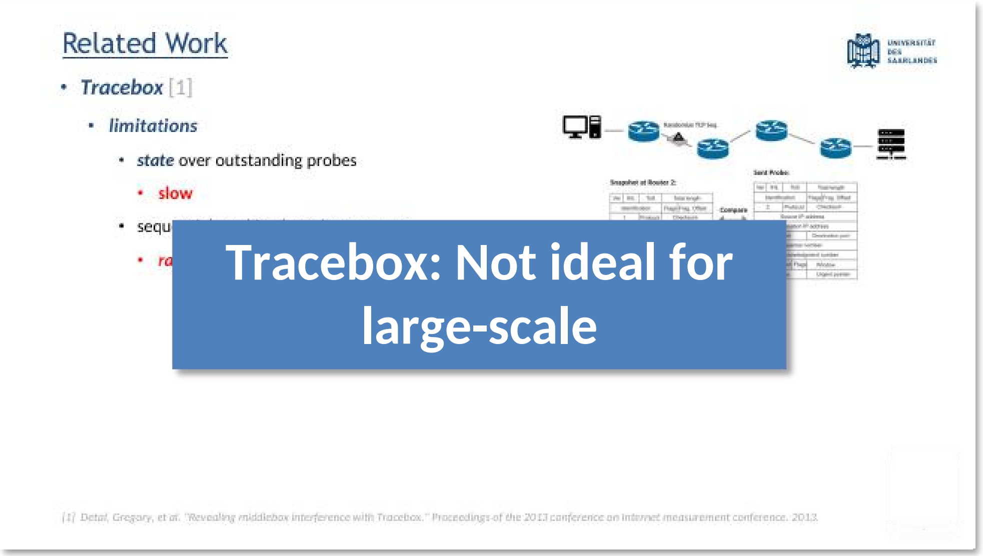

Related Work

- *Tracebox* [1]
  - *limitations*
    - *state* over outstanding probes
      - slow
    - sequ
      - ra

[1] Detal, Gregory, et al. "Revealing middlebox interference with Tracebox." Proceedings of the 2013 conference on Internet measurement conference. 2013.

Tracebox: Not ideal for large-scale

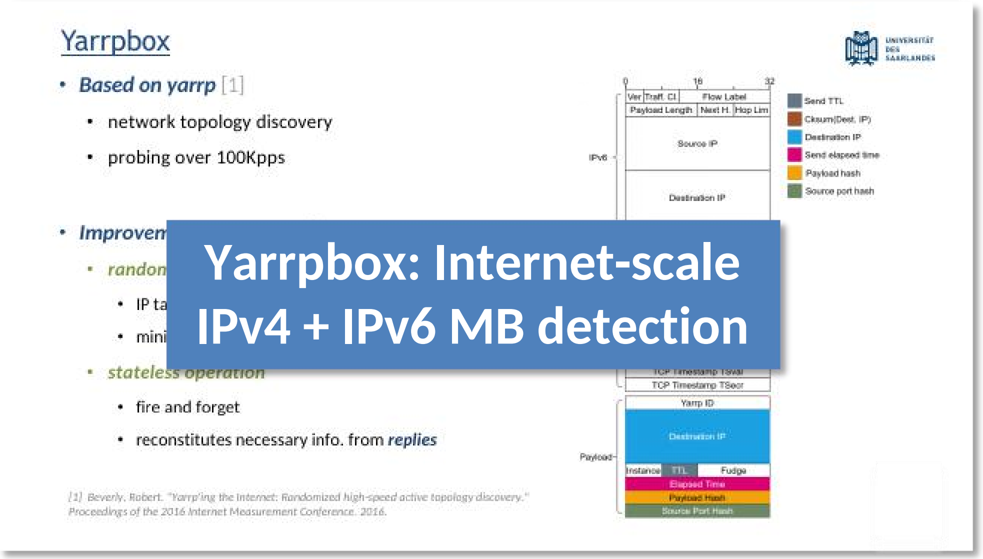

Yarrpbox

- *Based on yarrp* [1]
  - network topology discovery
  - probing over 100Kpps
- *Improvem*
  - *randon*
    - IP ta
    - mini
  - *stateless operation*
    - fire and forget
    - reconstitutes necessary info. from *replies*

[1] Beverly, Robert. "Yarrp'ing the Internet: Randomized high-speed active topology discovery." Proceedings of the 2016 Internet Measurement Conference. 2016.

Yarrpbox: Internet-scale IPv4 + IPv6 MB detection

Questions?

Artifacts Available acm

Artifacts Evaluated Reusable acm

Results Reproduced acm

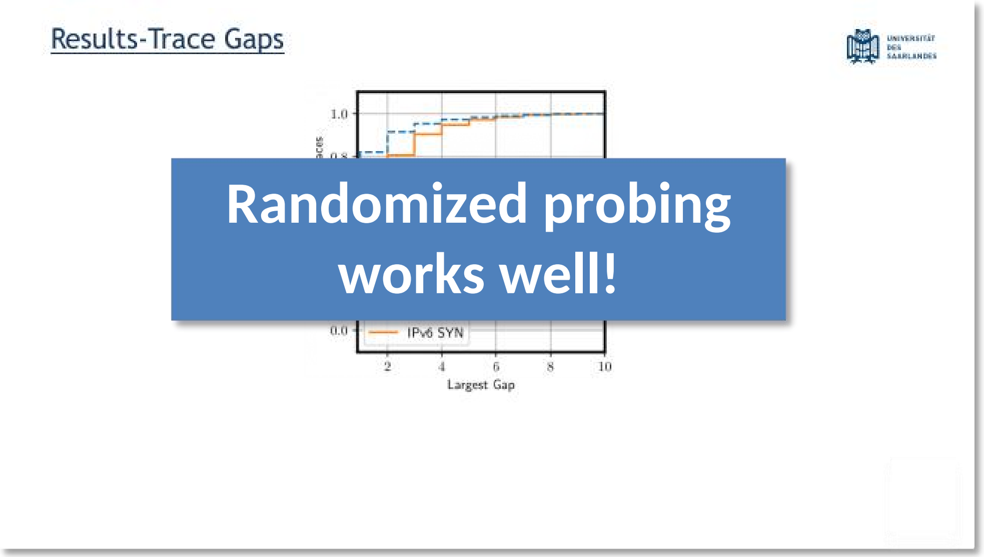

Results-Trace Gaps

Randomized probing works well!

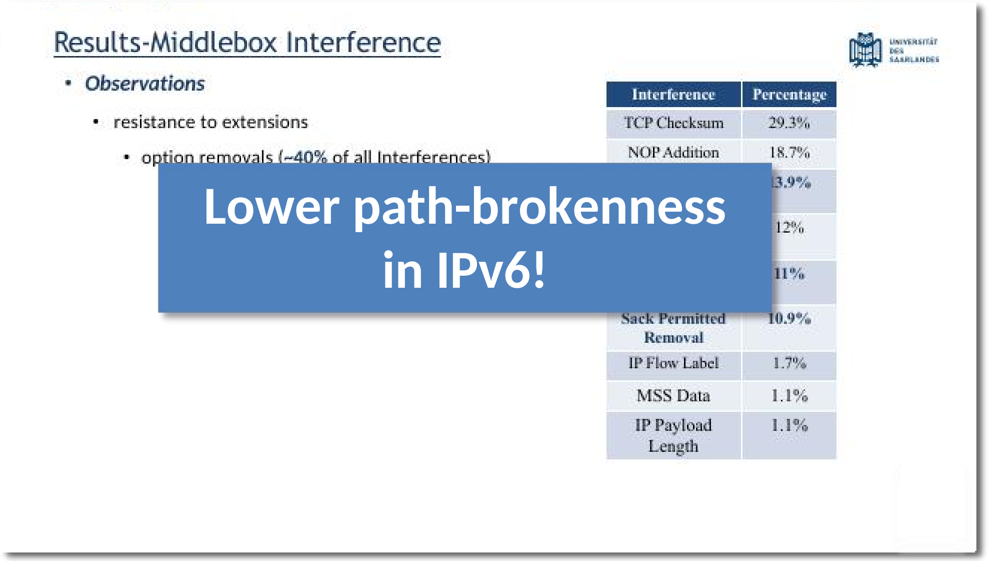

Results-Middlebox Interference

- *Observations*
  - resistance to extensions
    - option removals (~40% of all Interferences)

Interference	Percentage
TCP Checksum	29.3%
NOP Addition	18.7%
[illegible]	13.9%
[illegible]	12%
[illegible]	11%
Sack Permitted Removal	10.9%
IP Flow Label	1.7%
MSS Data	1.1%
IP Payload Length	1.1%

Lower path-brokenness in IPv6!

Additional Slides

Yarrpbox

- ***Methodology***
  - store ***state in probe packet***

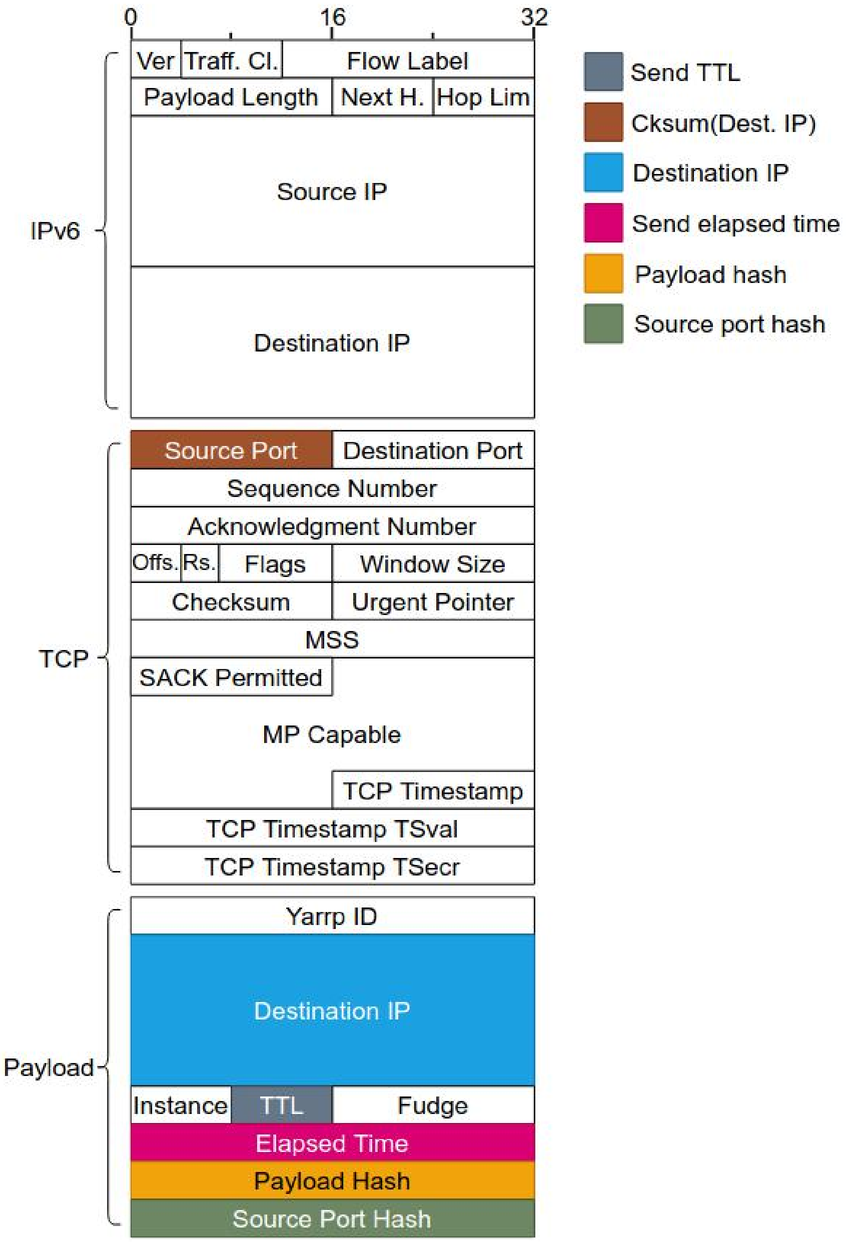

Yarrpbox

- ***Methodology***
  - store ***state in probe packet***
  - assign ***fixed values*** to other fields

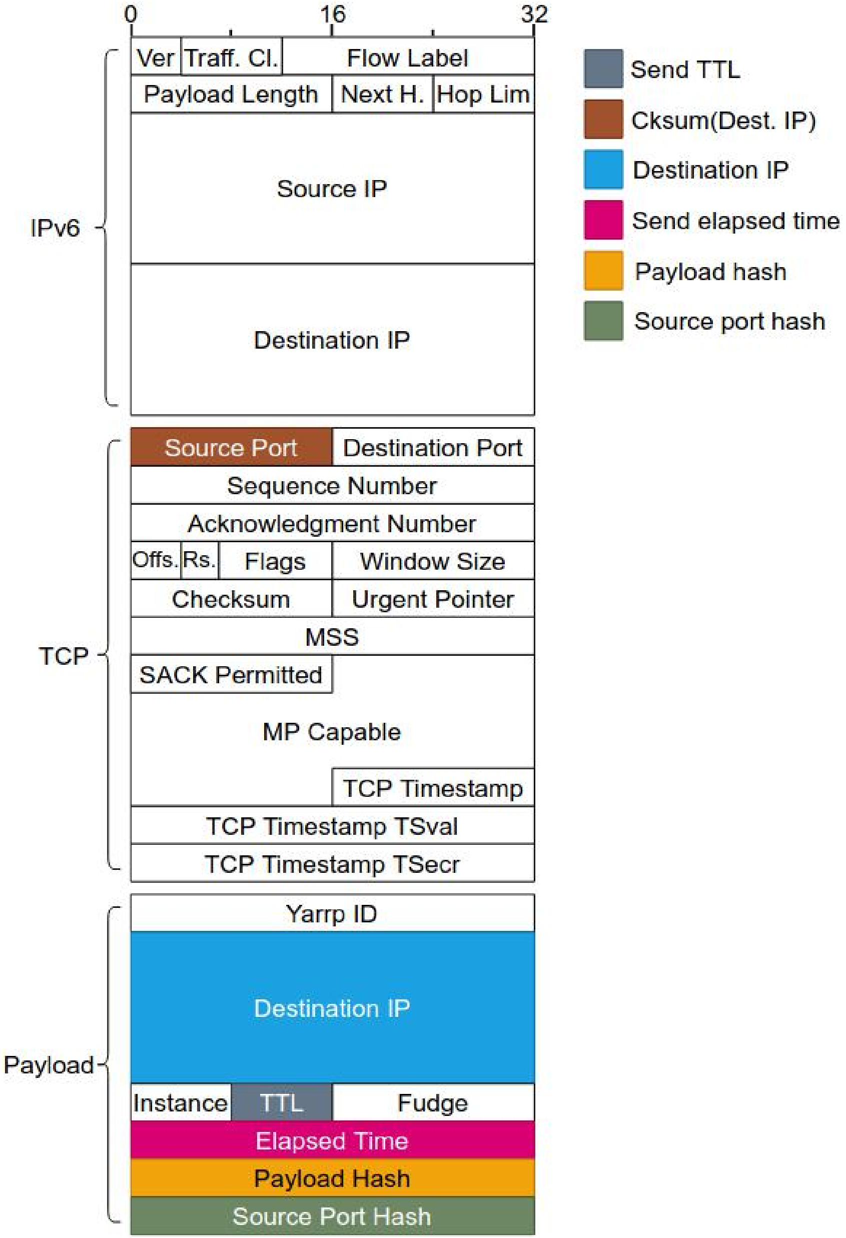

0 16 32

IPv6

Ver | Traff. Cl. | Flow Label

Payload Length | Next H. | Hop Lim

Source IP

Destination IP

TCP

Source Port | Destination Port

Sequence Number

Acknowledgment Number

Offs. | Rs. | Flags | Window Size

Checksum | Urgent Pointer

MSS

SACK Permitted

MP Capable

TCP Timestamp

TCP Timestamp TSval

TCP Timestamp TSecr

Payload

Yarrp ID

Destination IP

Instance | TTL | Fudge

Elapsed Time

Payload Hash

Source Port Hash

⬛ Send TTL

🟫 Cksum(Dest. IP)

🟦 Destination IP

🟪 Send elapsed time

🟧 Payload hash

🟩 Source port hash

Yarrpbox

- ***Methodology***
  - store ***state in probe packet***
  - assign ***fixed values*** to other fields
  - send TTL limited probe

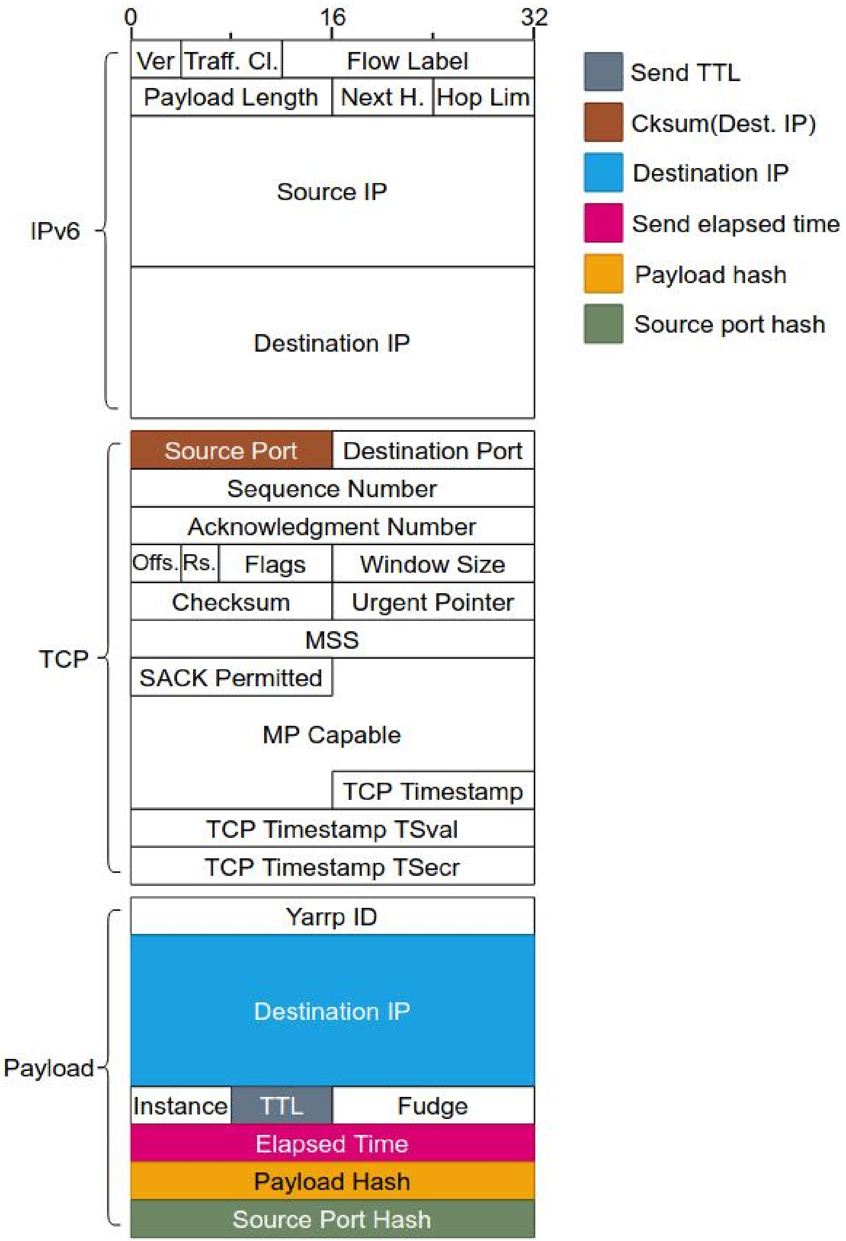

Yarrpbox

- ***Methodology***
  - store ***state in probe packet***
  - assign ***fixed values*** to other fields
  - send TTL limited probe
  - retrieve state from quoted packet in ICMP time
    - ***identify target, originating ttl***

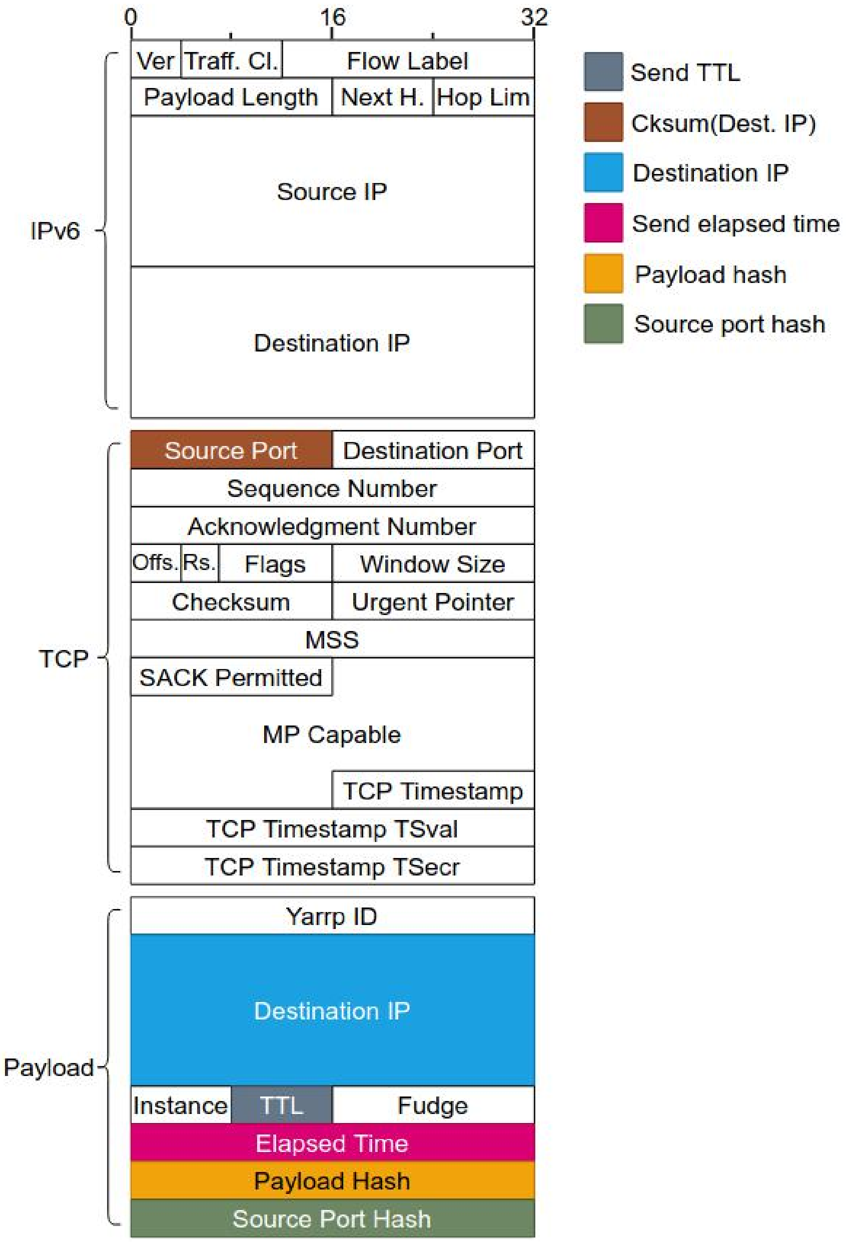

Yarrpbox

- ***Methodology***
  - store ***state in probe packet***
  - assign ***fixed values*** to other fields
  - send TTL limited probe
  - retrieve state from quoted packet in ICMP time
    - ***identify target, originating ttl***
  - inspect if header field values in quote ***match*** ex
  - if mismatch, middlebox along path

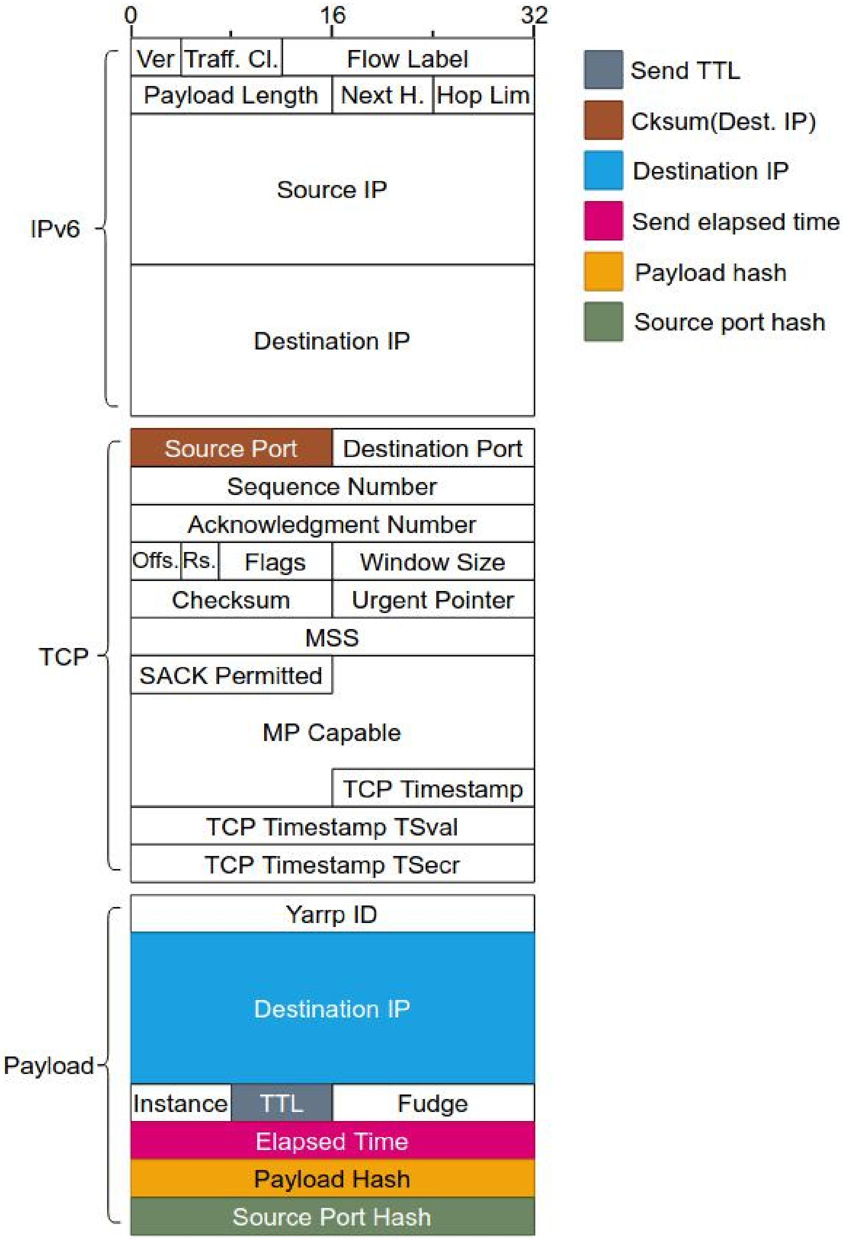

Yarrpbox-State Encoding

- **IPv4**
  - ***state*** encoded into ***IP and TCP header***, payload not used

Yarrpbox-State Encoding

- **IPv4**
  - ***state*** encoded into ***IP and TCP header***, payload not used
  - RFC792 (ICMPv4): "Internet Header + 64 bits of Data Datagram"

Yarrpbox-State Encoding

- **IPv4**
  - ***state*** encoded into ***IP and TCP header***, payload not used
  - RFC792 (ICMPv4): "Internet Header + 64 bits of Data Datagram"
  - RFC1812 (ICMPv4): "the ICMP datagram SHOULD contain ***as much*** of the original datagram as possible without the length of the ICMP datagram exceeding 576 bytes."

Yarrpbox-State Encoding

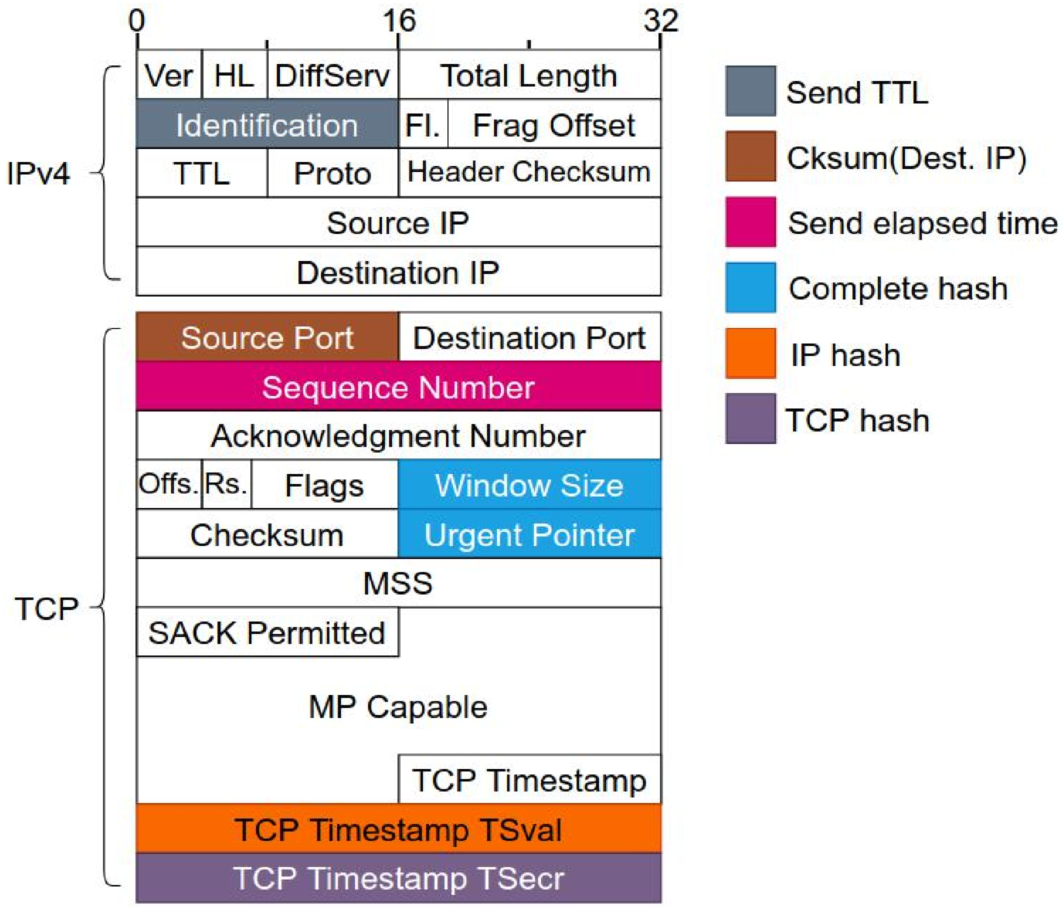

0 16 32

IPv4

Ver | HL | DiffServ | Total Length

Identification | Fl. | Frag Offset

TTL | Proto | Header Checksum

Source IP

Destination IP

TCP

Source Port | Destination Port

Sequence Number

Acknowledgment Number

Offs. | Rs. | Flags | Window Size

Checksum | Urgent Pointer

MSS

SACK Permitted

MP Capable

TCP Timestamp

TCP Timestamp TSval

TCP Timestamp TSecr

■ Send TTL

🟫 Cksum(Dest. IP)

🟥 Send elapsed time

🟦 Complete hash

🟧 IP hash

🟪 TCP hash

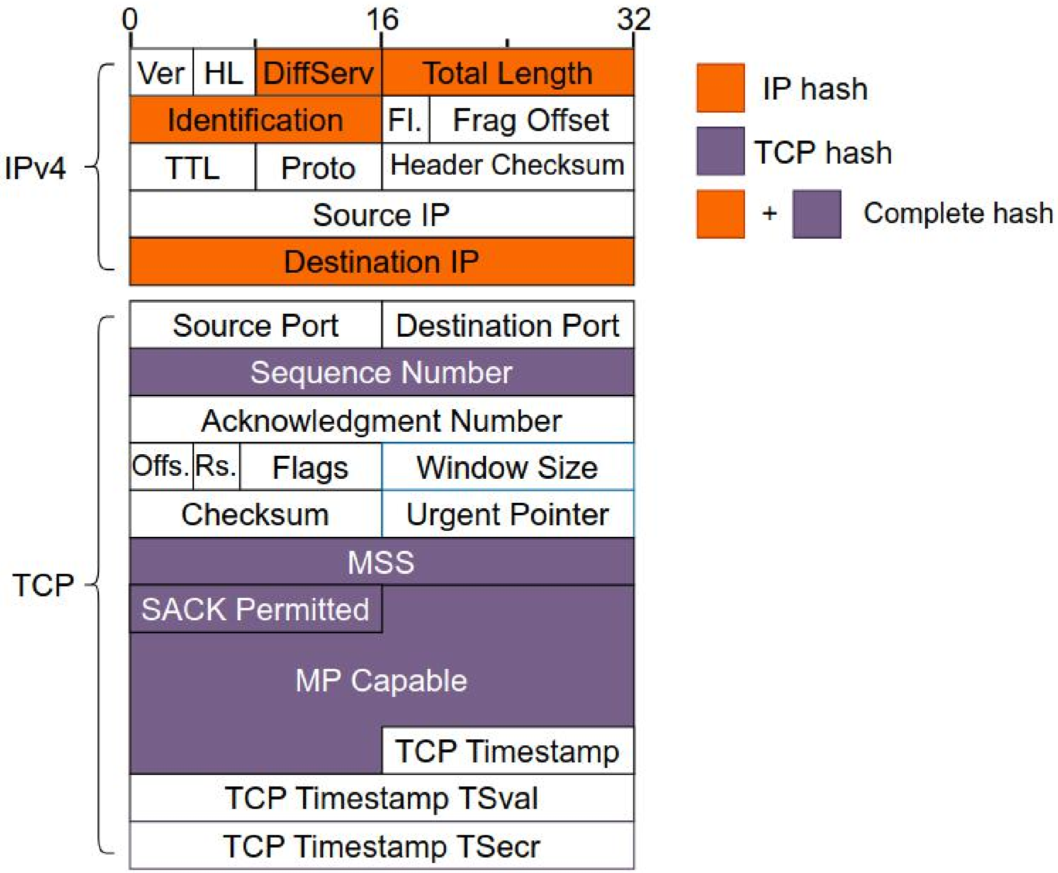

0 16 32

IPv4

Ver | HL | DiffServ | Total Length

Identification | Fl. | Frag Offset

TTL | Proto | Header Checksum

Source IP

Destination IP

TCP

Source Port | Destination Port

Sequence Number

Acknowledgment Number

Offs. | Rs. | Flags | Window Size

Checksum | Urgent Pointer

MSS

SACK Permitted

MP Capable

TCP Timestamp

TCP Timestamp TSval

TCP Timestamp TSecr

🟧 IP hash

🟪 TCP hash

🟧 + 🟪 Complete hash

Yarrpbox-State Encoding

- **IPv6**
  - ***state*** within ***payload***
  - RFC4443 (ICMPv6): "***As much*** of invoking packet as possible without the ICMPv6 packet exceeding the minimum IPv6 MTU”
  - more modifications detected

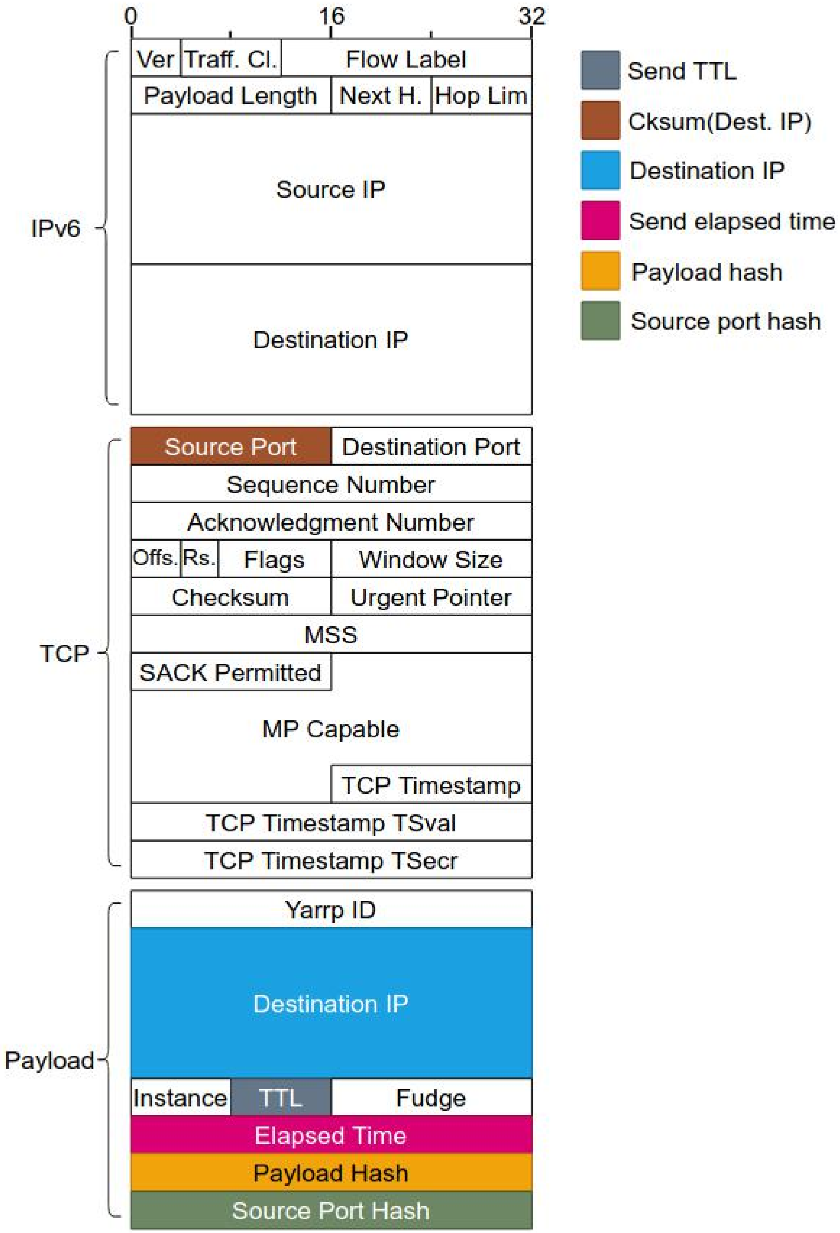

0 16 32

IPv6

Ver	Traff. Cl.	Flow Label

Payload Length	Next H.	Hop Lim

Source IP

Destination IP

TCP

Source Port	Destination Port

Sequence Number

Acknowledgment Number

Offs.	Rs.	Flags	Window Size

Checksum	Urgent Pointer

MSS

SACK Permitted

MP Capable

TCP Timestamp

TCP Timestamp TSval

TCP Timestamp TSecr

Payload

Yarrp ID

Destination IP

Instance	TTL	Fudge

Elapsed Time

Payload Hash

Source Port Hash

■	Send TTL
■	Cksum(Dest. IP)
■	Destination IP
■	Send elapsed time
■	Payload hash
■	Source port hash

Results-Overview

Scan	Replies	Interferences	MB IPs
IPv4 SYN	89.4M	**759.5k**	**16.8k**
IPv4 ACK	88.8M	636.6k	8.9k
IPv6 SYN	92.4M	197.9k	8.2k
IPv6 ACK	94.6M	25.1k	7.6k
IPv6 UDP	94.1M	50.7k	10.8k
ICMPv6	98.7M	**4**	**2**

- ***Vantage Point***
  - MPI
- ***Targets***
  - IPv4: ***random IP*** from each ***/24***
  - IPv6: BGP announced prefixes, ***100*** IPs per prefix

Results-Overview

Scan	Replies	Interferences	MB IPs
IPv4 SYN	89.4M	**759.5k**	**16.8k**
IPv4 ACK	88.8M	636.6k	8.9k
IPv6 SYN	92.4M	197.9k	8.2k
IPv6 ACK	94.6M	25.1k	7.6k
IPv6 UDP	94.1M	50.7k	10.8k
ICMPv6	98.7M	**4**	**2**

Scan with multiple protocols!

- ***Vantage Point***
  - MPI
- ***Targets***
  - IPv4: ***random IP*** from each ***/24***
  - IPv6: BGP announced prefixes, ***100*** IPs per prefix

Results-Middlebox Interference

Interference	Percentage
IP ID/**TSval**/RW+UP	**76.5%**
TCP Tmsp. TSval	**6.6%**
NOP Addition	**5.4%**
MP CAP. Removal	**5%**
TSecr/RW+UP	2.2%
TCP UP+RW	2%
IP ID	0.7%
TCP Seq. Number	**0.5%**
IP Total Length	0.5%
TCP Tmsp. Removal	**0.2%**
SACK Perm. Removal	**0.1%**
MSS Data	0.1%

Results-HC MB ASes

Scan	#1 AS	#2 AS	#3 AS
IPv4 SYN	**7018 (9%)**	5617 (4.3%)	**1299 (3.3%)**
IPv6 SYN	45271 (5.8%)	**1299 (4.7%)**	23910 (3.25)

- ***IPv4:*** 1.6k, ***IPv6:*** 1.3k
- IPv4: AS 7018 (AT&T, US), IPv6: AS 45271 (Idea Cellular Limited, IN)
- ISPs and ***Tier-1s***

Results-Middlebox Location

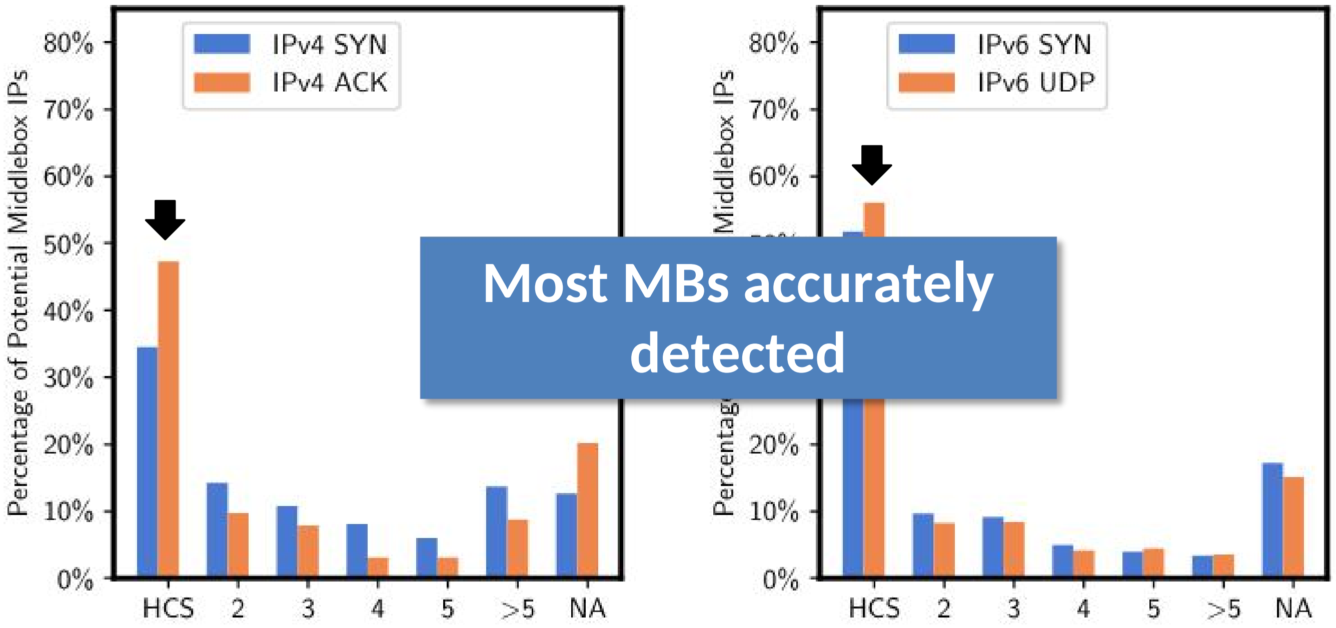

Most MBs accurately detected

- ~ 71% IPv4 replies -> full quotes -> more accuracy
  - < 30% full quoters
- IPv6 full quotes only

Results-Middlebox Interference (HC ASes)

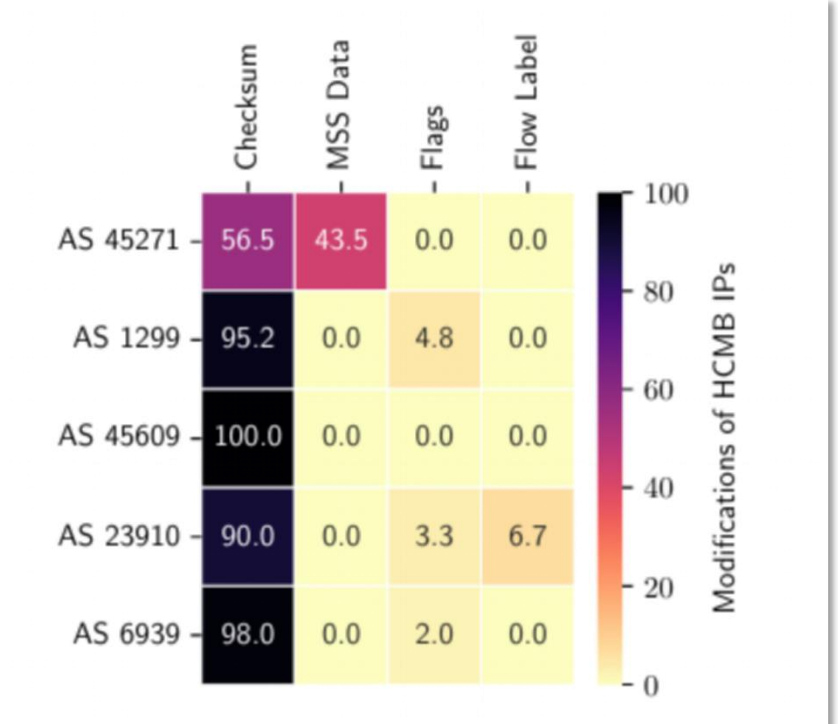

3/5 ASes zero the checksum predominantly

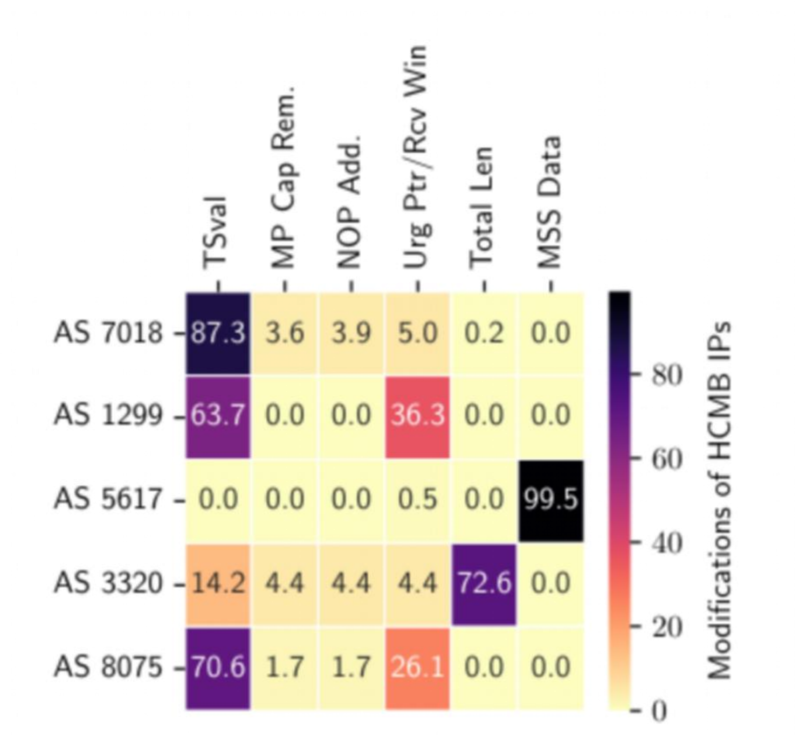

More diverse interferences

Future Work

- Longitudinal measurements
- Target popular servers
- Transient dynamics
- Traffic drops
- Port-based scans
- TCP checksum zeroings
- Proxies

Results-Geo Distributed Measurements

- ***Influence of probing location***
  - IPv4 MB affected paths: **6% (US East)** to **29% (Australia)**

Results-Geo Distributed Measurements

- ***Influence of probing location***
  - IPv4 MB affected paths: **6% (US East)** to **29% (Australia)**
  - IPv6 MB affected paths: **0.05% (Australia)** to **0.15% (MPI VP)**

Results-Geo Distributed Measurements

- ***Influence of probing location***
  - IPv4 MB affected paths: **6% (US East)** to **29% (Australia)**
  - IPv6 MB affected paths: **0.05% (Australia)** to **0.15% (MPI VP)**

Scan from multiple locations!

Results-Validation

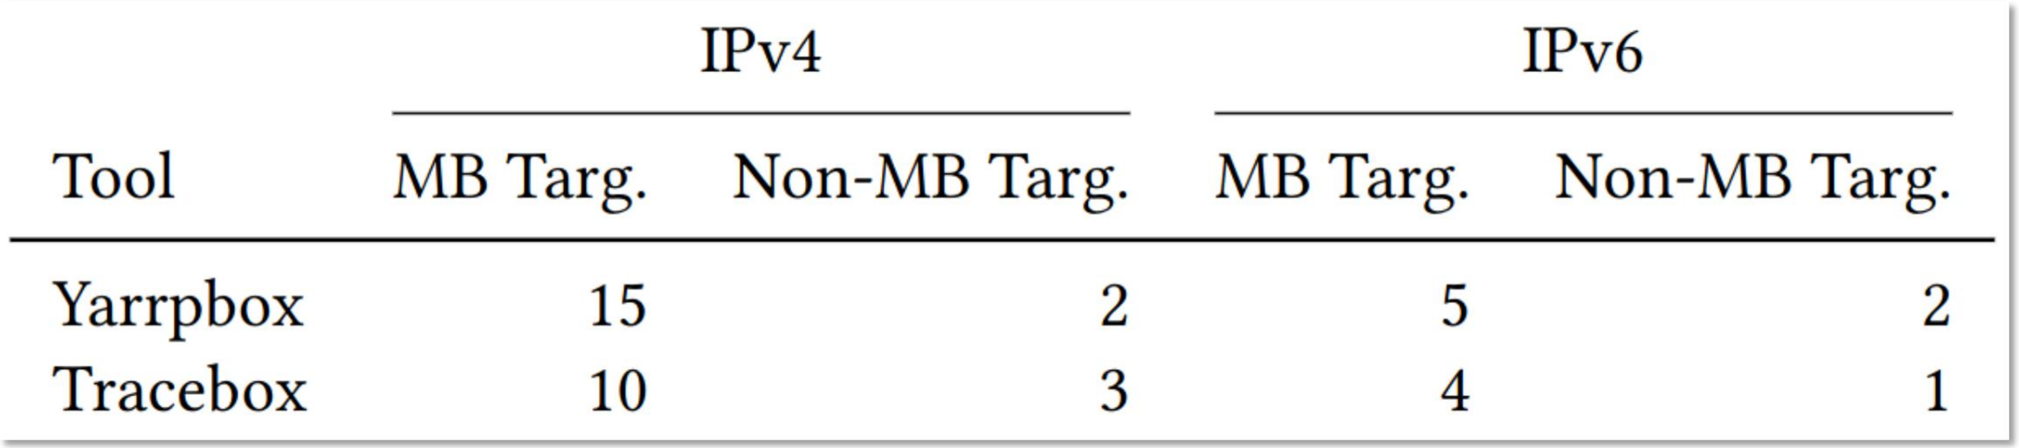

	IPv4		IPv6	
Tool	MB Targ.	Non-MB Targ.	MB Targ.	Non-MB Targ.
Yarrpbox	15	2	5	2
Tracebox	10	3	4	1

Results-ICMP Quoting Behaviour

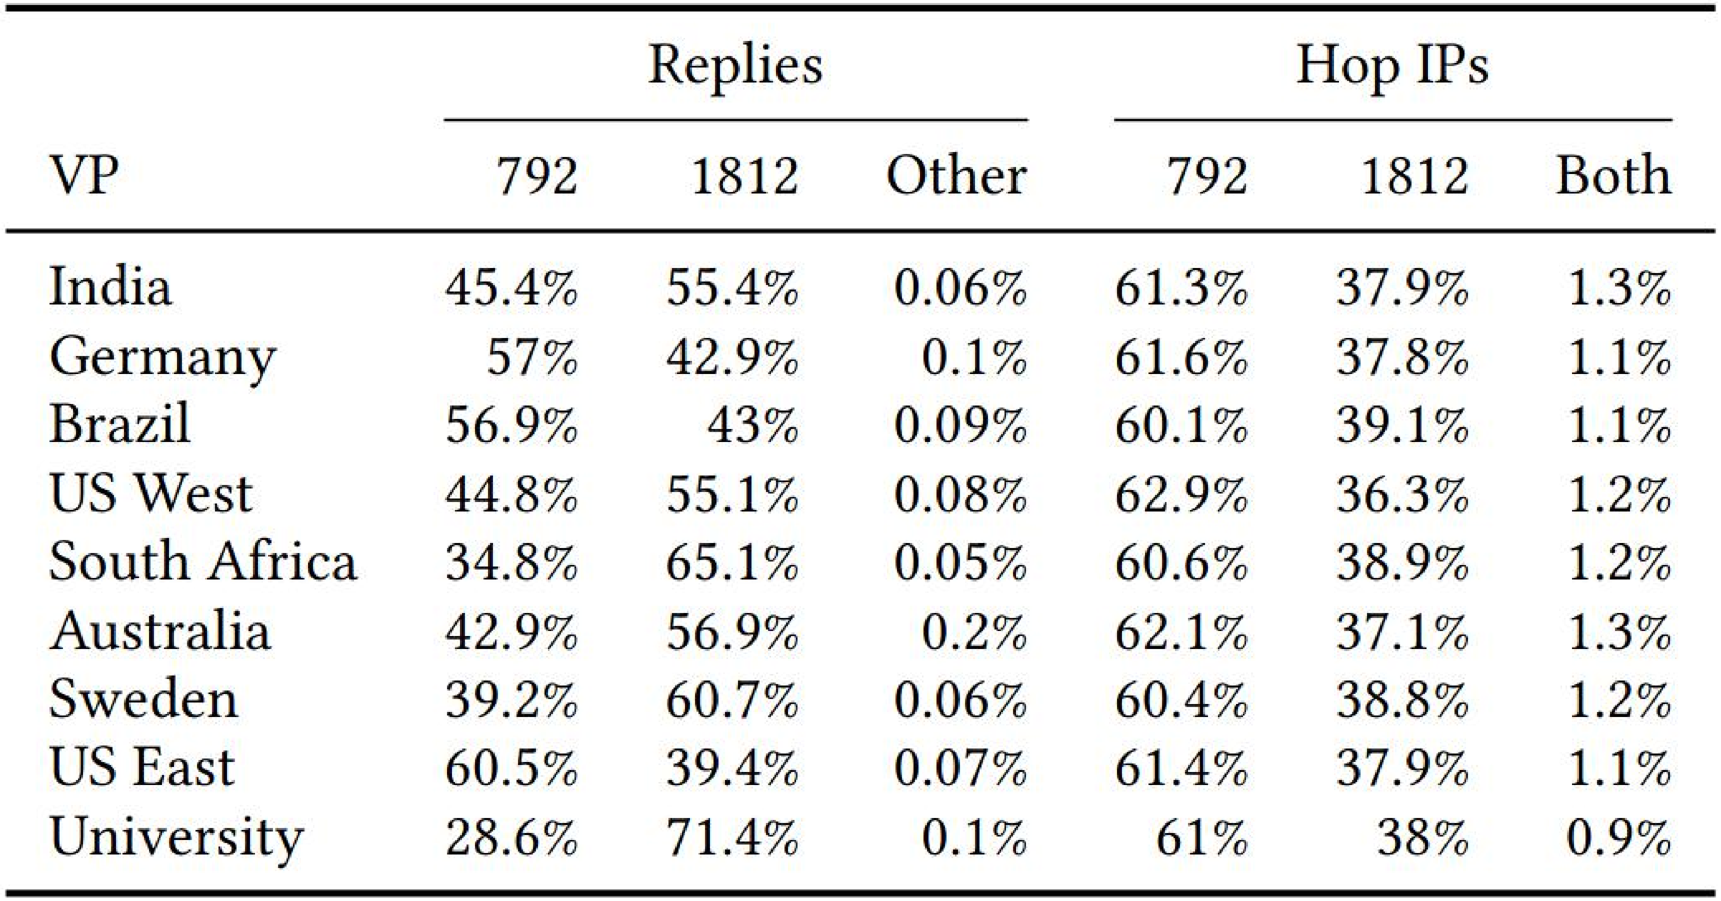

	Replies			Hop IPs		
VP	792	1812	Other	792	1812	Both
India	45.4%	55.4%	0.06%	61.3%	37.9%	1.3%
Germany	57%	42.9%	0.1%	61.6%	37.8%	1.1%
Brazil	56.9%	43%	0.09%	60.1%	39.1%	1.1%
US West	44.8%	55.1%	0.08%	62.9%	36.3%	1.2%
South Africa	34.8%	65.1%	0.05%	60.6%	38.9%	1.2%
Australia	42.9%	56.9%	0.2%	62.1%	37.1%	1.3%
Sweden	39.2%	60.7%	0.06%	60.4%	38.8%	1.2%
US East	60.5%	39.4%	0.07%	61.4%	37.9%	1.1%
University	28.6%	71.4%	0.1%	61%	38%	0.9%

Background-MB Behaviour

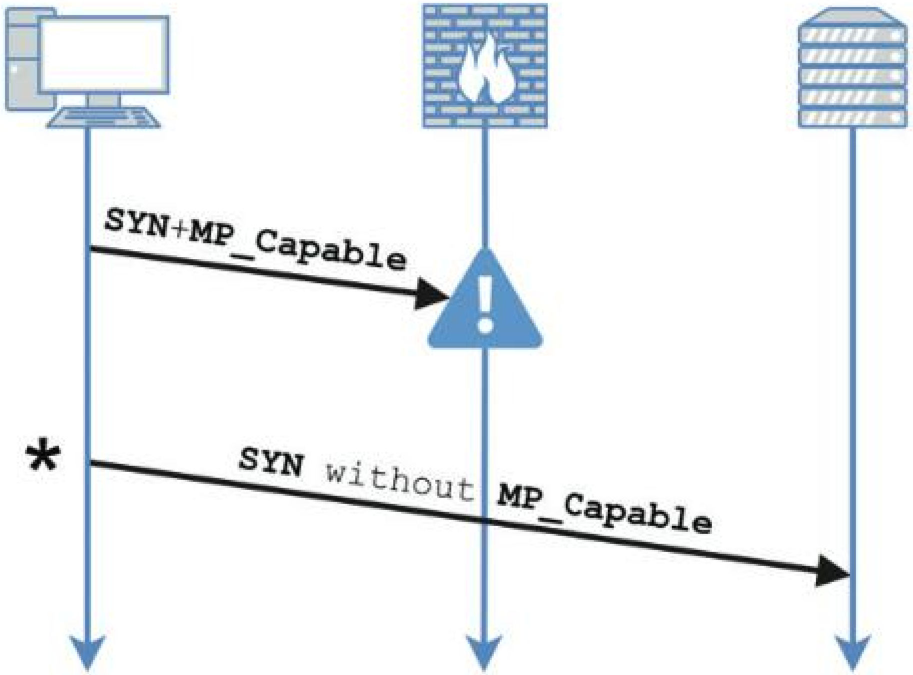

Background-MB Behaviour

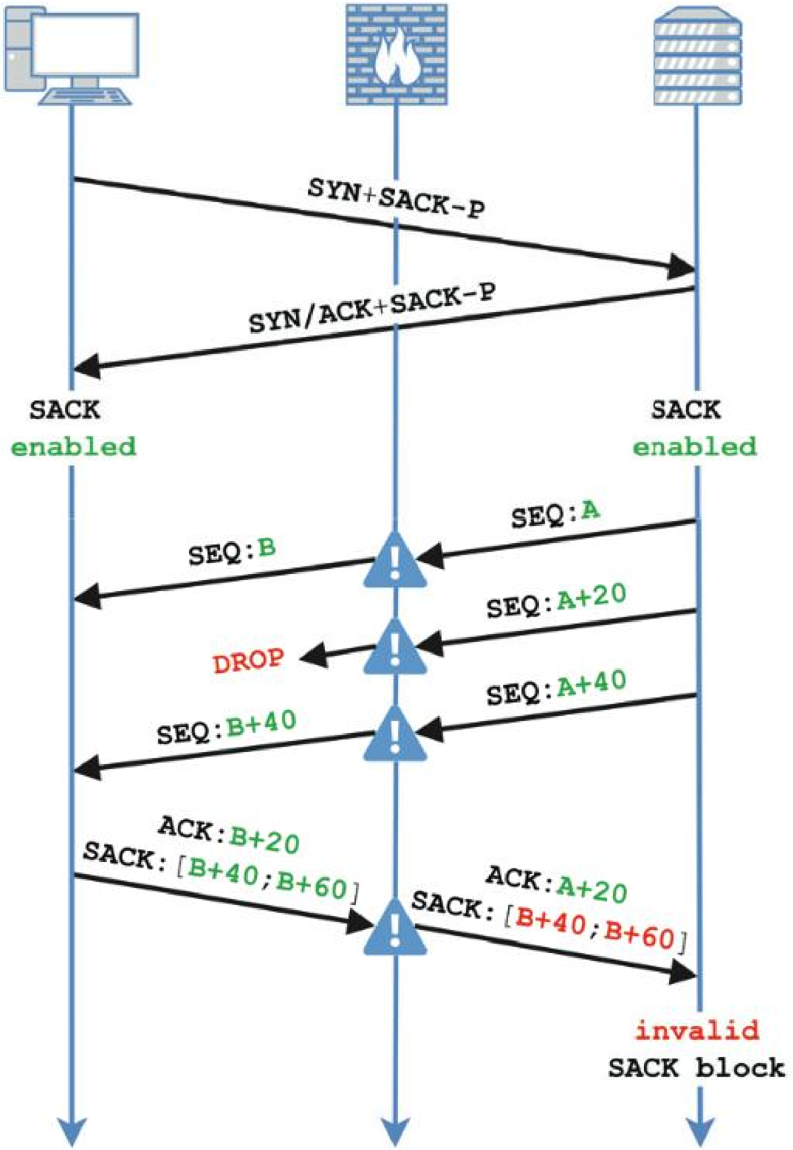

Background-MB Behaviour

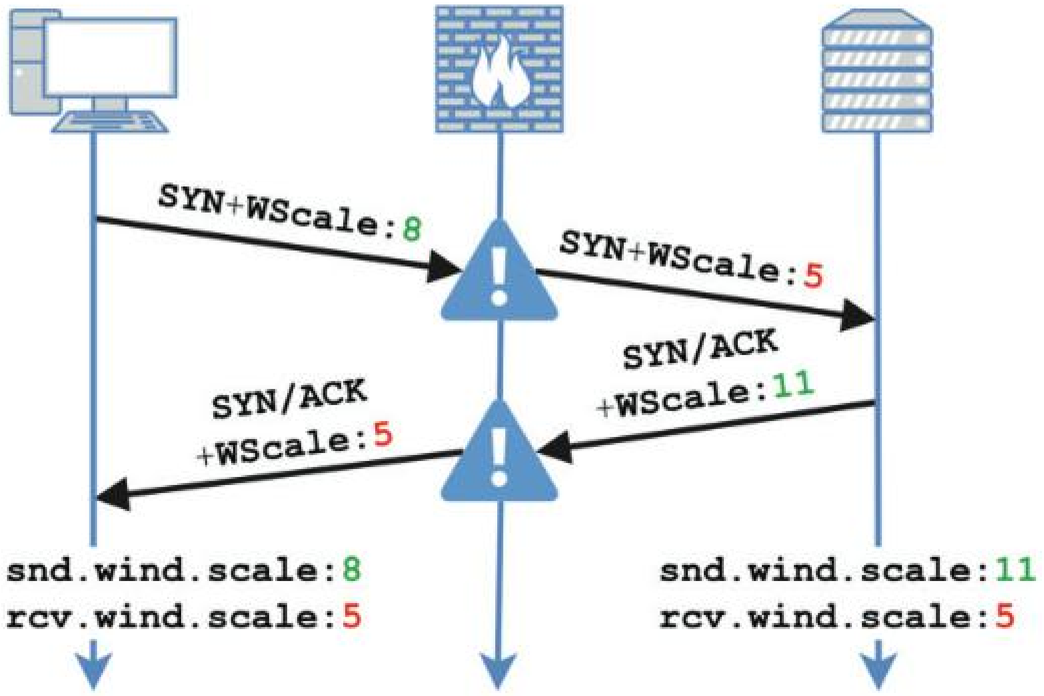

Hash Storage (IPv4 only)

- **Most modified fields**
  - DSCP, IP total length, IPID (IP Hash)
  - Sequence number, MSS option, MP_CAPABLE and SACK-Permitted (TCP Hash)

Hash Storage (IPv4 only)

- **Most modified fields**
  - DSCP, IP total length, IPID (IP Hash)
  - Sequence number, MSS option, MP_CAPABLE and SACK-Permitted (TCP Hash)
- **Hash storage (4 byte hash)**
  - Urgent pointer (2 byte) + rcv window (2 byte) (Complete Hash)
  - Timestamp Option
    - TSval (4 byte) (IP Hash)
    - TSecr (4 byte) (TCP Hash)

TCP/IPv4-Modification of IPID

- ***IPID/ TSval and Receiver Window or Urgent Pointer***
  - arises when only IP hash and Complete hash modified, nothing else
  - only IPID modif. from IP hash can not be pinpointed

TCP/IPv4-Modification of IPID

- ***IPID/ TSval and Receiver Window or Urgent Pointer***
  - arises when only IP hash and Complete hash modified, nothing else
  - only IPID modif. from IP hash can not be pinpointed
  - only IP hash and Complete hash modified,
    - might be down to modif. of IPID or hash storages (TSval and Urg + Rcv)

TCP/IPv4-Modification of IPID

- ***IPID/ TSval and Receiver Window or Urgent Pointer***
  - arises when only IP hash and Complete hash modified, nothing else
  - only IPID modif. from IP hash can not be pinpointed
  - only IP hash and Complete hash modified,
    - might be down to modif. of IPID or hash storages (TSval and Urg + Rcv)
  - only IP hash modified and nothing else (not even complete hash)
    - IP hash storage (Tsval) modified

Fields Most Likely to be Modified

- ***Based on Edeline and Donnet*** *[1]*
    - active probing using Tracebox from 89 PlanetLab nodes located in different continents
    - Aimed at 594,241 popular HTTP servers (extracted from Alexa 1M)

[1] Edeline, K., & Donnet, B. (2019, June). A bottom-up investigation of the transport-layer ossification. In 2019 Network Traffic Measurement and Analysis Conference (TMA) (pp. 169-176). IEEE

Fields Most Likely to be Modified

- ***Based on Edeline and Donnet*** [1]
  - active probing using Tracebox from 89 PlanetLab nodes located in different continents
  - Aimed at 594,241 popular HTTP servers (extracted from Alexa 1M)
  - resulting dataset has 232 million observations attributed to 18,667 middleboxes

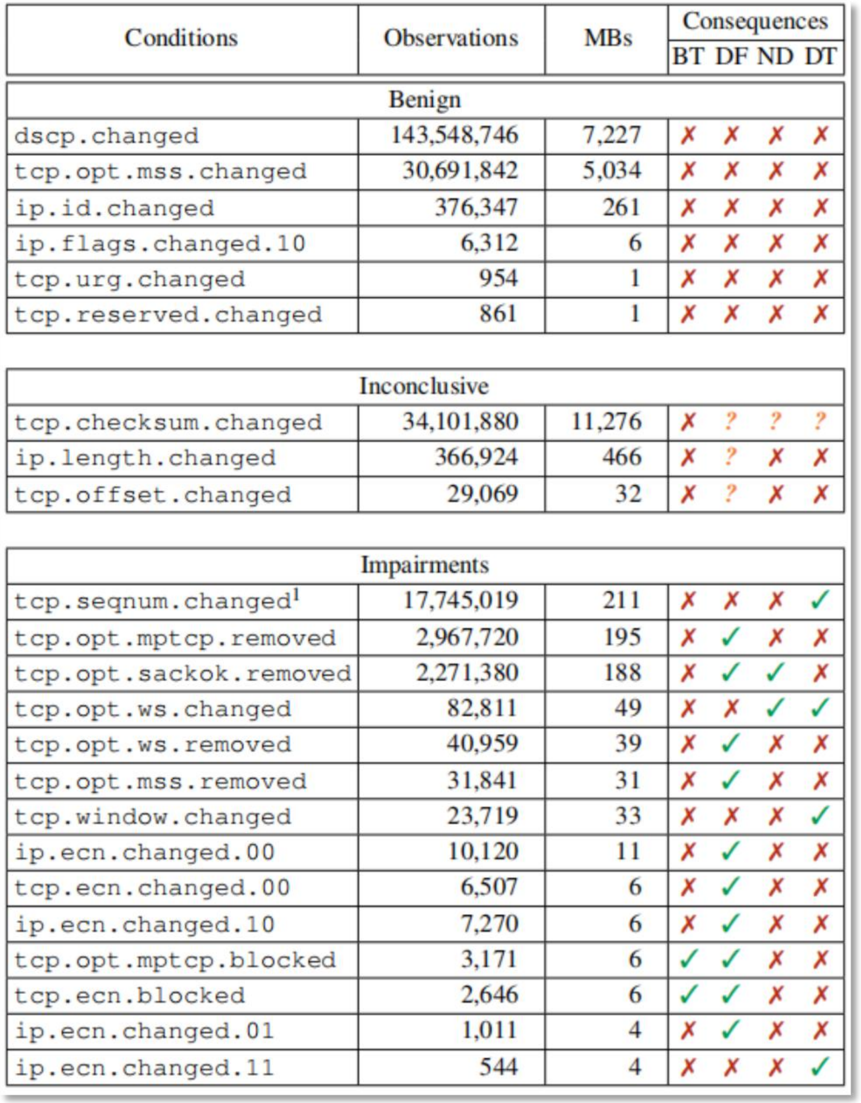

Conditions	Observations	MBs	Consequences			
			BT	DF	ND	DT
Benign						
`dscp.changed`	143,548,746	7,227	✗	✗	✗	✗
`tcp.opt.mss.changed`	30,691,842	5,034	✗	✗	✗	✗
`ip.id.changed`	376,347	261	✗	✗	✗	✗
`ip.flags.changed.10`	6,312	6	✗	✗	✗	✗
`tcp.urg.changed`	954	1	✗	✗	✗	✗
`tcp.reserved.changed`	861	1	✗	✗	✗	✗
Inconclusive						
`tcp.checksum.changed`	34,101,880	11,276	✗	?	?	?
`ip.length.changed`	366,924	466	✗	?	✗	✗
`tcp.offset.changed`	29,069	32	✗	?	✗	✗
Impairments						
`tcp.seqnum.changed`[1]	17,745,019	211	✗	✗	✗	✓
`tcp.opt.mptcp.removed`	2,967,720	195	✗	✓	✗	✗
`tcp.opt.sackok.removed`	2,271,380	188	✗	✓	✓	✗
`tcp.opt.ws.changed`	82,811	49	✗	✗	✓	✓
`tcp.opt.ws.removed`	40,959	39	✗	✓	✗	✗
`tcp.opt.mss.removed`	31,841	31	✗	✓	✗	✗
`tcp.window.changed`	23,719	33	✗	✗	✗	✓
`ip.ecn.changed.00`	10,120	11	✗	✓	✗	✗
`tcp.ecn.changed.00`	6,507	6	✗	✓	✗	✗
`ip.ecn.changed.10`	7,270	6	✗	✓	✗	✗
`tcp.opt.mptcp.blocked`	3,171	6	✓	✓	✗	✗
`tcp.ecn.blocked`	2,646	6	✓	✓	✗	✗
`ip.ecn.changed.01`	1,011	4	✗	✓	✗	✗
`ip.ecn.changed.11`	544	4	✗	✗	✗	✓

[1] Edeline, K., & Donnet, B. (2019, June). A bottom-up investigation of the transport-layer ossification. In 2019 Network Traffic Measurement and Analysis Conference (TMA) (pp. 169-176). IEEE

Fields Most Likely to be Modified

- ***Based on Edeline and Donnet*** [1]
  - active probing using Tracebox from 89 PlanetLab nodes located in different continents
  - Aimed at 594,241 popular HTTP servers (extracted from Alexa 1M)
  - resulting dataset has 232 million observations attributed to 18,667 middleboxes
  - observed packet manipulations, classified as
    - benign middlebox modifications
    - inconclusive
    - impairments (capable of harm to TCP)

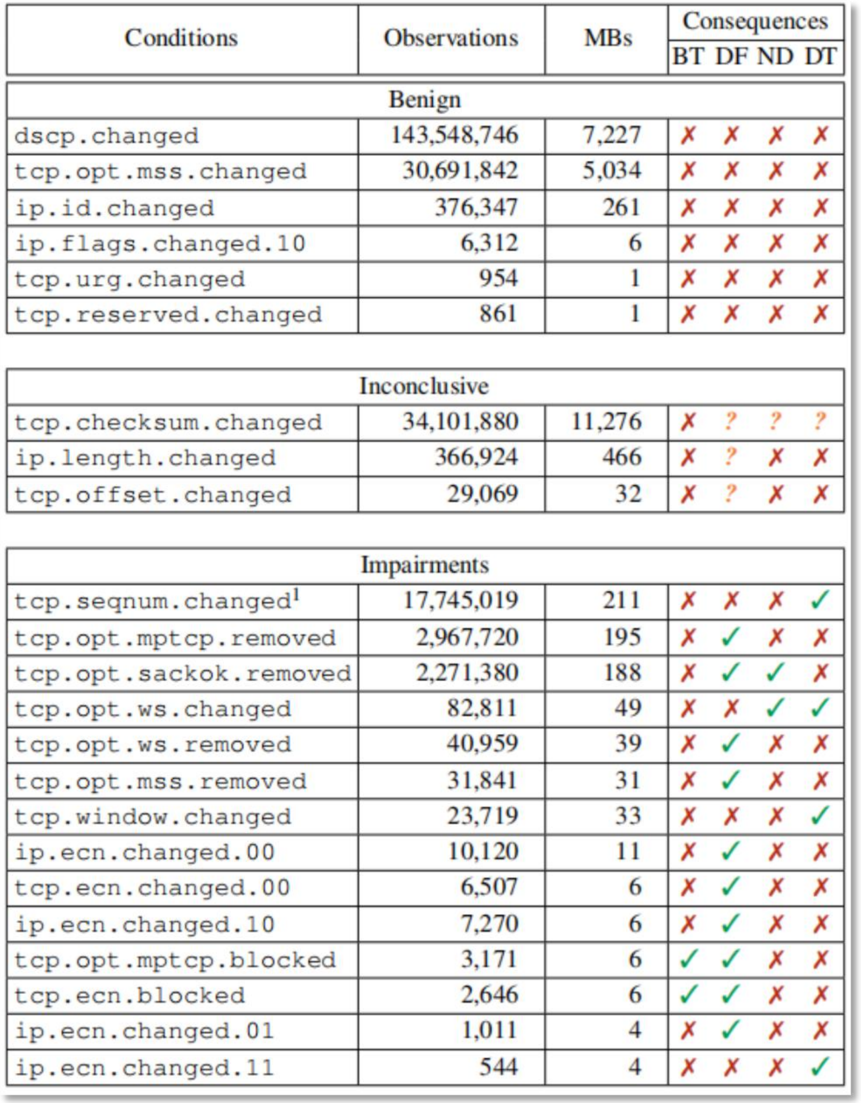

Conditions	Observations	MBs	Consequences			
			BT	DF	ND	DT
Benign						
dscp.changed	143,548,746	7,227	✗	✗	✗	✗
tcp.opt.mss.changed	30,691,842	5,034	✗	✗	✗	✗
ip.id.changed	376,347	261	✗	✗	✗	✗
ip.flags.changed.10	6,312	6	✗	✗	✗	✗
tcp.urg.changed	954	1	✗	✗	✗	✗
tcp.reserved.changed	861	1	✗	✗	✗	✗
Inconclusive						
tcp.checksum.changed	34,101,880	11,276	✗	?	?	?
ip.length.changed	366,924	466	✗	?	✗	✗
tcp.offset.changed	29,069	32	✗	?	✗	✗
Impairments						
tcp.seqnum.changed[1]	17,745,019	211	✗	✗	✗	✓
tcp.opt.mptcp.removed	2,967,720	195	✗	✓	✗	✗
tcp.opt.sackok.removed	2,271,380	188	✗	✓	✓	✗
tcp.opt.ws.changed	82,811	49	✗	✗	✓	✓
tcp.opt.ws.removed	40,959	39	✗	✓	✗	✗
tcp.opt.mss.removed	31,841	31	✗	✓	✗	✗
tcp.window.changed	23,719	33	✗	✗	✗	✓
ip.ecn.changed.00	10,120	11	✗	✓	✗	✗
tcp.ecn.changed.00	6,507	6	✗	✓	✗	✗
ip.ecn.changed.10	7,270	6	✗	✓	✗	✗
tcp.opt.mptcp.blocked	3,171	6	✓	✓	✗	✗
tcp.ecn.blocked	2,646	6	✓	✓	✗	✗
ip.ecn.changed.01	1,011	4	✗	✓	✗	✗
ip.ecn.changed.11	544	4	✗	✗	✗	✓

[1] Edeline, K., & Donnet, B. (2019, June). A bottom-up investigation of the transport-layer ossification. In 2019 Network Traffic Measurement and Analysis Conference (TMA) (pp. 169-176). IEEE

Related Work

- ***Controlling both ends***
  - detection of middleboxes on one path

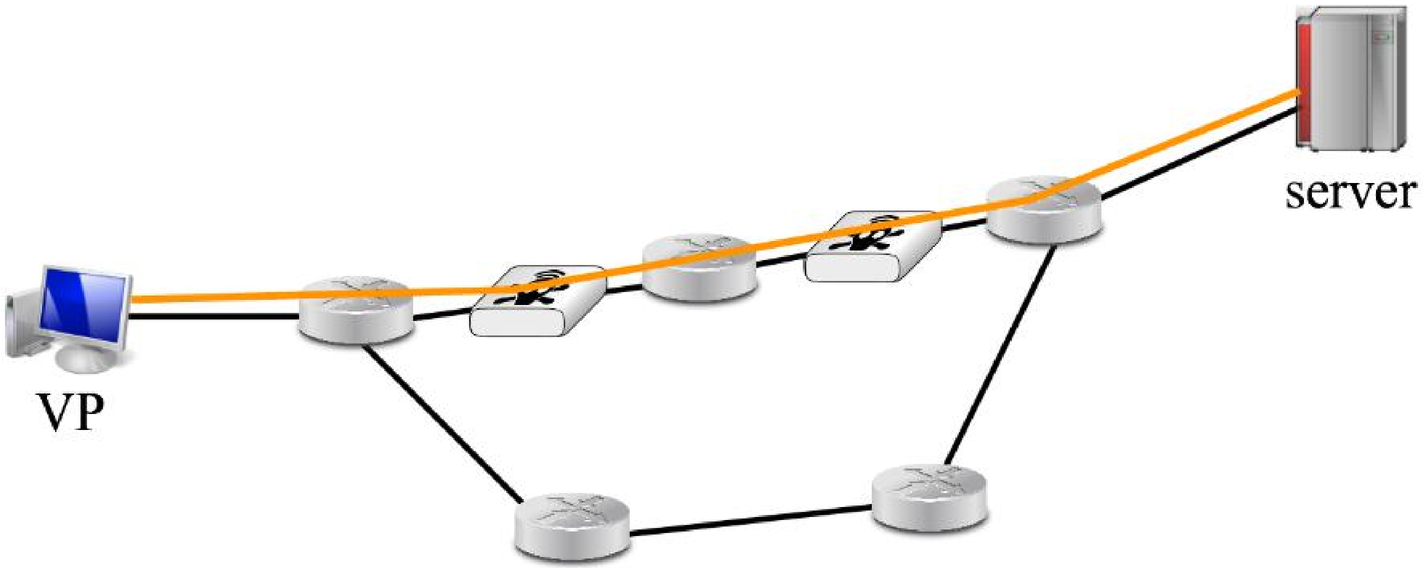

https://www.ietf.org/proceedings/93/slides/slides-93-hopsrg-3.pdf

Related Work

- ***Controlling both ends***
  - detection of middleboxes on one path
  - detects middleboxes only on paths to that server from different VPs

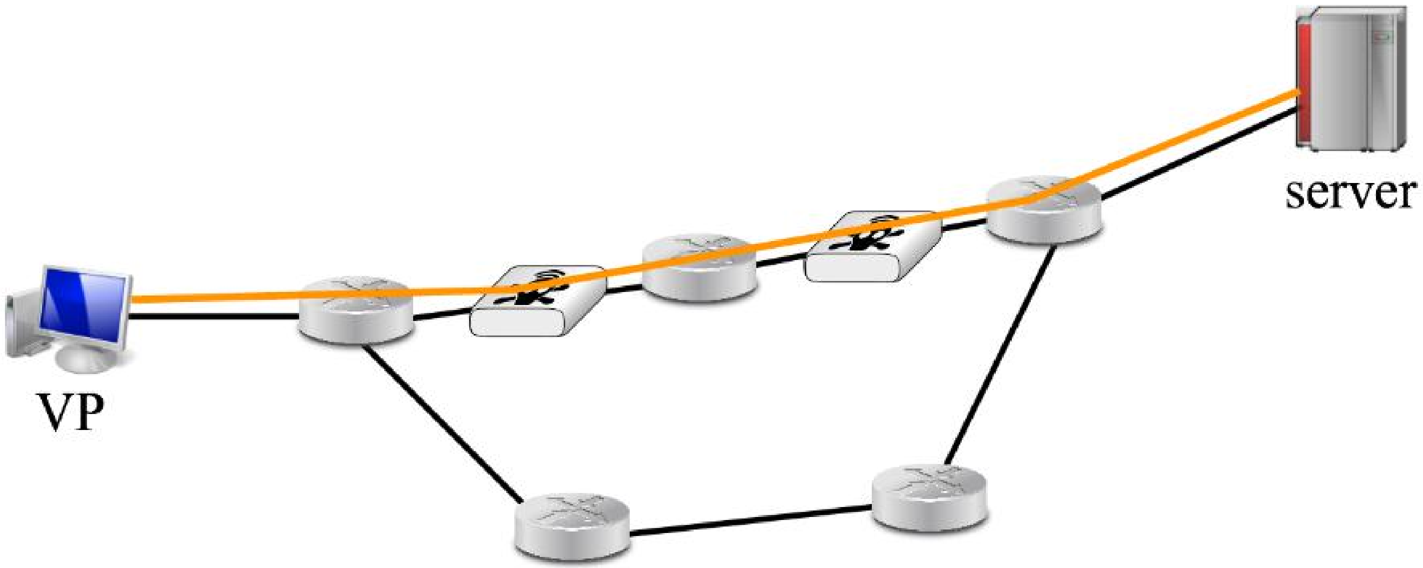

https://www.ietf.org/proceedings/93/slides/slides-93-hopsrg-3.pdf

Related Work

- ***With limited controlled servers***
  - lot of middleboxes missed

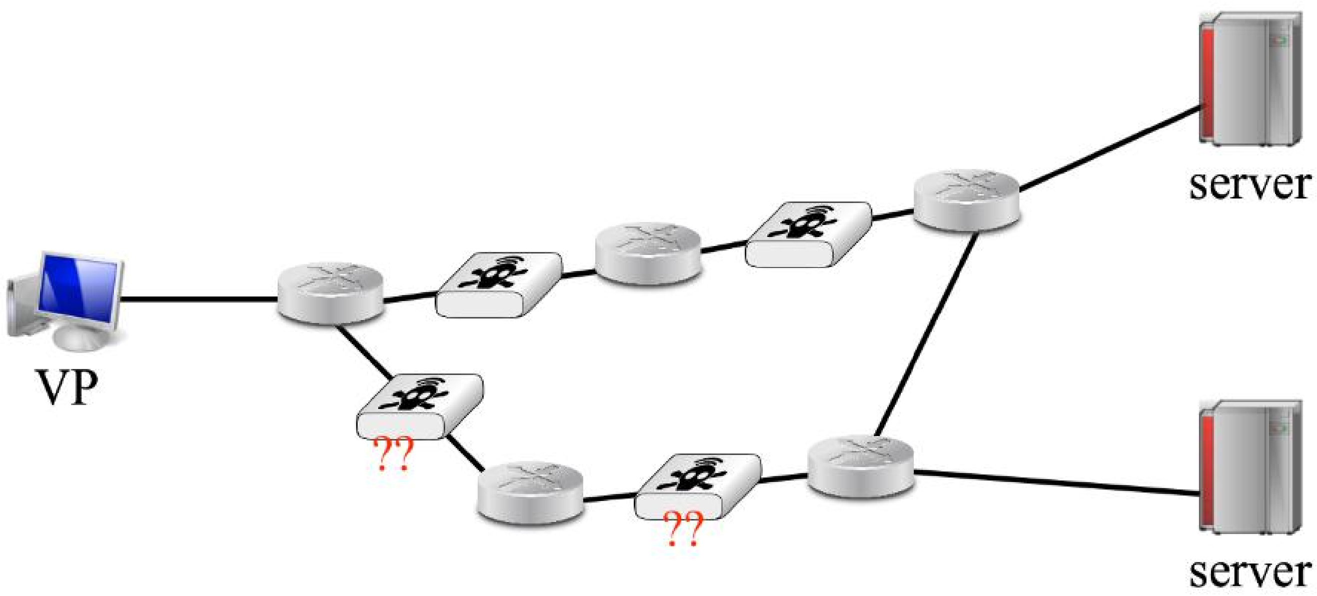

https://www.ietf.org/proceedings/93/slides/slides-93-hopsrg-3.pdf