Unintended Consequences: Effects of Submarine Cable Deployment on Internet Routing

By

Rodérick Fanou, Bradley Huffaker, Ricky Mok, KC Claffy
(CAIDA/UC San Diego)

PAM 2020

Mesh of Submarine Cables, the Backbone of the Internet

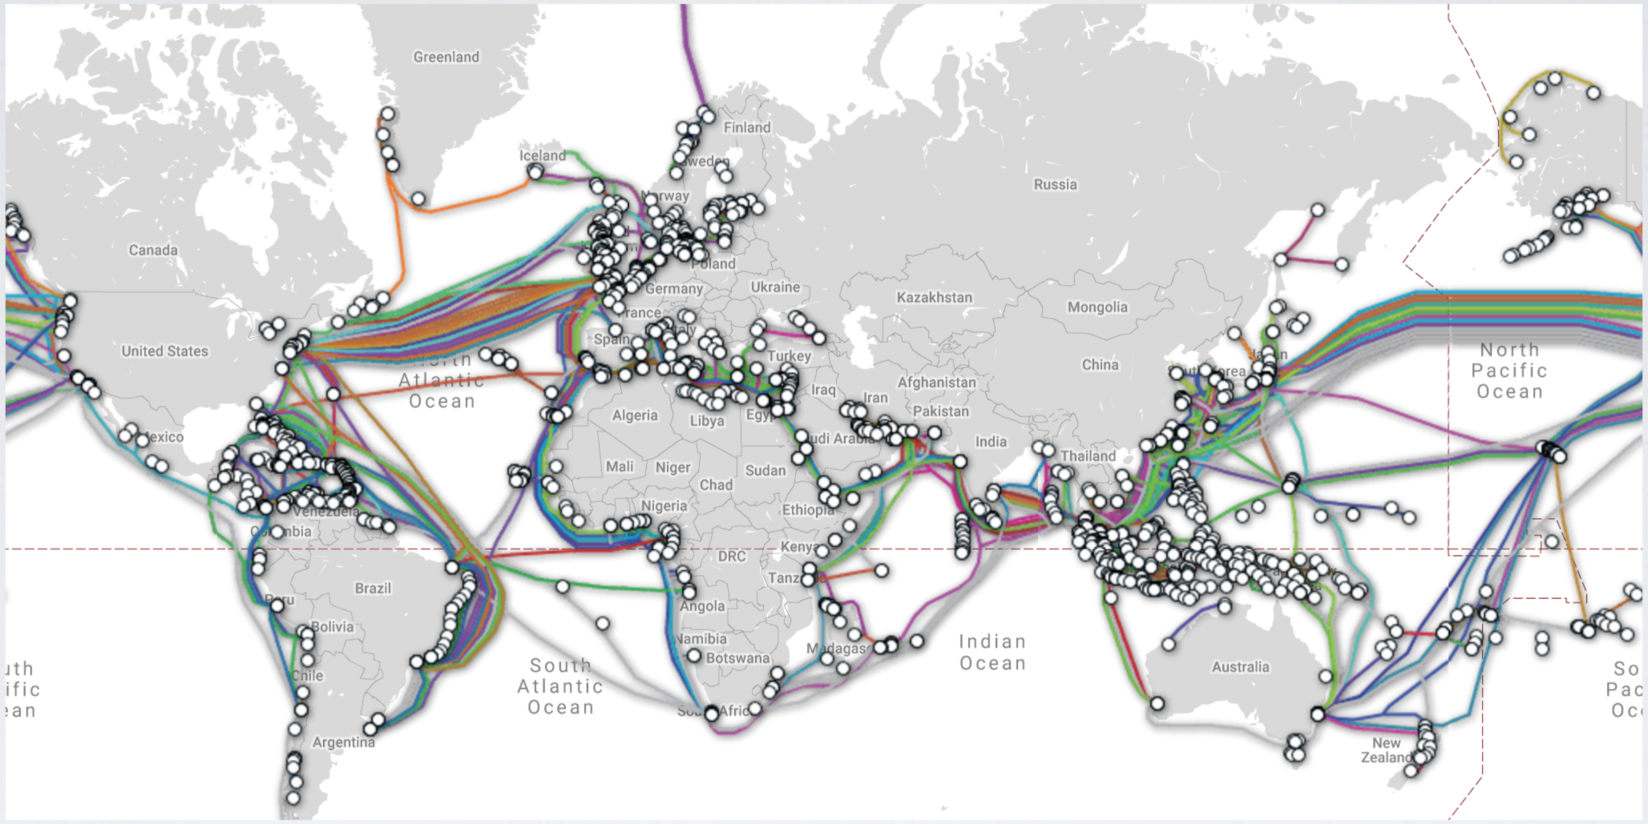

Source: Telegeography Submarine Cable Map, https://www.submarinecablemap.com, 2020

- As of early 2020, over 406 cables carry >99% of international traffic [1, 2, 3]
- Little research to isolate end-to-end performance **changes induced by their launch**

[1] Telegeography, Submarine Cable Frequently Asked Questions, 2020
[2] Bischof et al, Submarine Cables and Internet Resiliency, 2018
[3] Main D., Undersea Cables Transport 99% of International Data, 2015

Launch of SACS on Sept 18, 2018, the First South-Atlantic Cable System

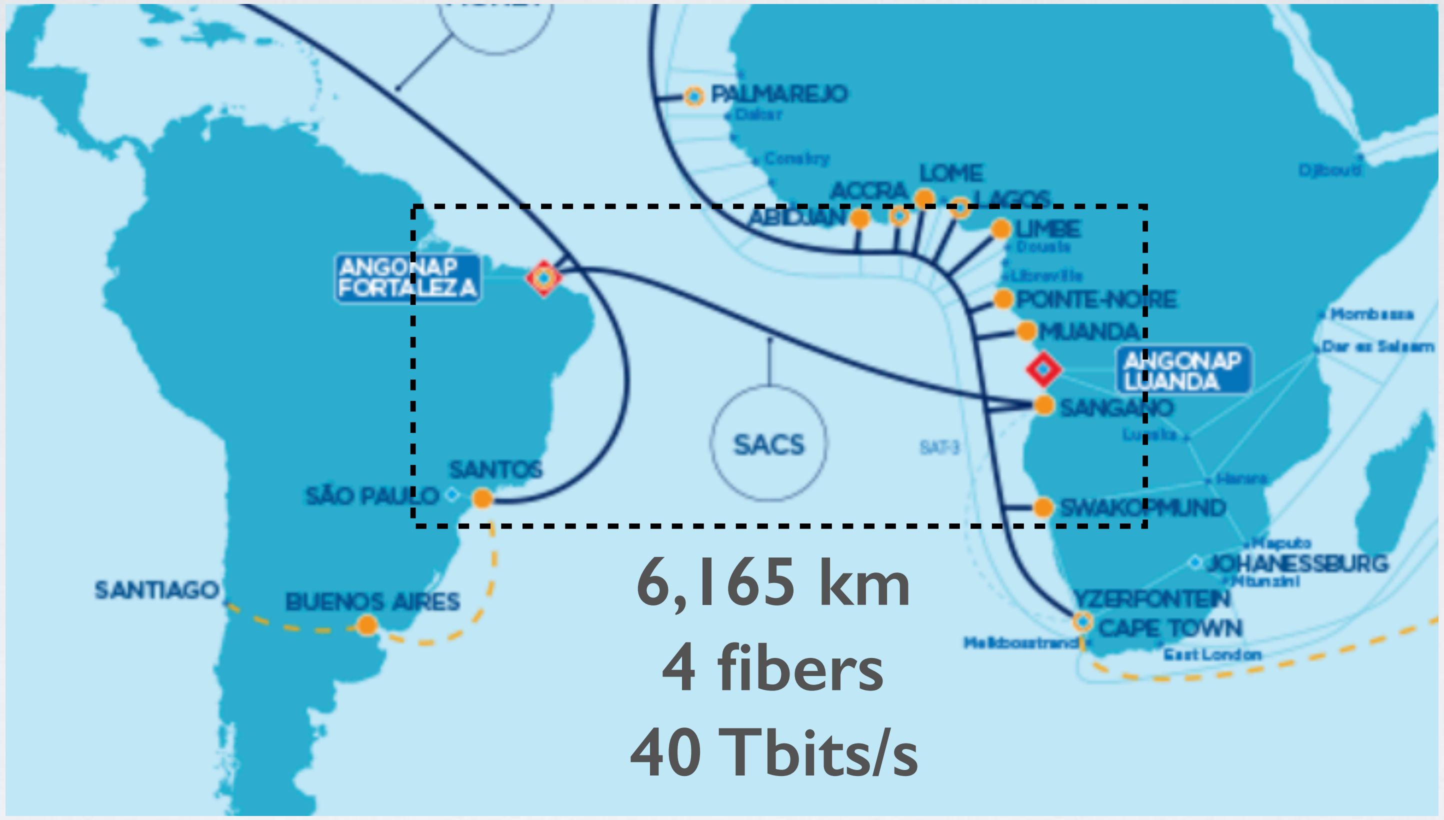

Source: Angola Cables Network, https://www.angolacables.co.ao/en/network-map/, 2019

- Angola Cables (AC, AS37468) activated the SACS cable mid-Sept 2018 [4]
- SACS links Fortaleza (**Brazil** - South America) to Sangano (**Angola** - Africa) [5]

[4] Madory et al., South-Atlantic Cable System: The Impact on the Internet, LACNIC 30 & WTR POP-BA/RNP, 2018
[5] Madory, D., First Subsea Cable Across South Atlantic Activated, Sep 2018

In the Press

First Africa–South Americas Fibre Optic Cable Opens for Commercial Traffic

ED TARGETT EDITOR
28TH SEPTEMBER 2018

+ INCREASE / DECREASE TEXT SIZE −

South Atlantic Cable System goes live – offering lower latencies to the US

Staff Writer 26 September 2018

First Subsea Cable Across South Atlantic Activated

By Doug Madory, Dyn
September 19, 2018

Angola Cables lights up world's first submarine cable linking Africa to the Americas

The South Atlantic Cable System (SACS) is ready for operation

September 28, 2018 By: Tanwen Dawn-Hiscox

7 Reasons Why SACS is a Game Changer

Winston Qiu SACS 31 October 2018

Teraco data centres will benefit from SACS cable

Bradley Prior 14 November 2018

Press Reports Expected Performance Improvements

- **Angola Cables (AC) & Oracle Dyn, 2018 [4,5]:**

First Subsea Cable Across South Atlantic Activated

By Doug Madory, Dyn
September 19, 2018

South Atlantic Cable System
The Impact on the Internet
LACNIC 30 – Lightning Talk

Doug Madory Oracle Internet Intelligence	Darwin Costa & Humberto Galiza Angola Cables

  - Latencies between servers in Brazil and Angola decreased from **over 300ms** to as low as **100ms**
  - Latencies to Angola from other locations outside Brazil (e.g. Ashburn, Tokyo, Singapore) also **experienced improvements**

[4] Madory et al., South-Atlantic Cable System: The Impact on the Internet, LACNIC 30 & WTR POP-BA/RNP, 2018
[5] Madory, D., First Subsea Cable Across South Atlantic Activated, 2018
[6] Prior, B. , Teraco Data Centres Will Benefit from SACS Cable, 2018

Press Reports Expected Performance Improvements

- Prior, 2018 [6]:

Teraco data centres will benefit from SACS cable

Bradley Prior 14 November 2018

  - SACS reduces latency to the Americas substantially including a reduction from **338ms** to **163ms** between Cape Town and Miami

[4] Madory et al., South-Atlantic Cable System: The Impact on the Internet, LACNIC 30 & WTR POP-BA/RNP, 2018
[5] Madory, D., First Subsea Cable Across South Atlantic Activated, Sep 2018
[6] Prior, B. , Teraco Data Centres Will Benefit from SACS Cable, 2018

Research Questions

- Can we scientifically study the macroscopic effects of a cable launch on AS topology and end-to-end performance?
- Does end-to-end performance improve for all regions using the new cable?
- Do AS paths connecting those regions shorten?

Challenges

Translating research questions into a set of goals/tasks

- Identify IP addresses that represent this cable using IP-layer traceroute measurements

- Figure out which (IP/AS) paths cross the cable

- Find out which effects crossing the cable had on:
  - paths and latencies between linked countries
  - paths and latencies from/destined to different continents

caida

Methodology

1. Collect candidate IP paths that could have crossed the cable
2. Identify router IP interfaces on both sides of the cable
3. Search for comparable traceroutes
4. Annotate collected paths

Datasets

- **Active traceroute (monitors and topology data)**
  - CAIDA's Ark (100+ monitors)
  - RIPE NCC's Atlas (10,000+ monitors)

Source: CAIDA Archipelago (Ark) measurement platform, https://www.caida.org/projects/ark/, 2020

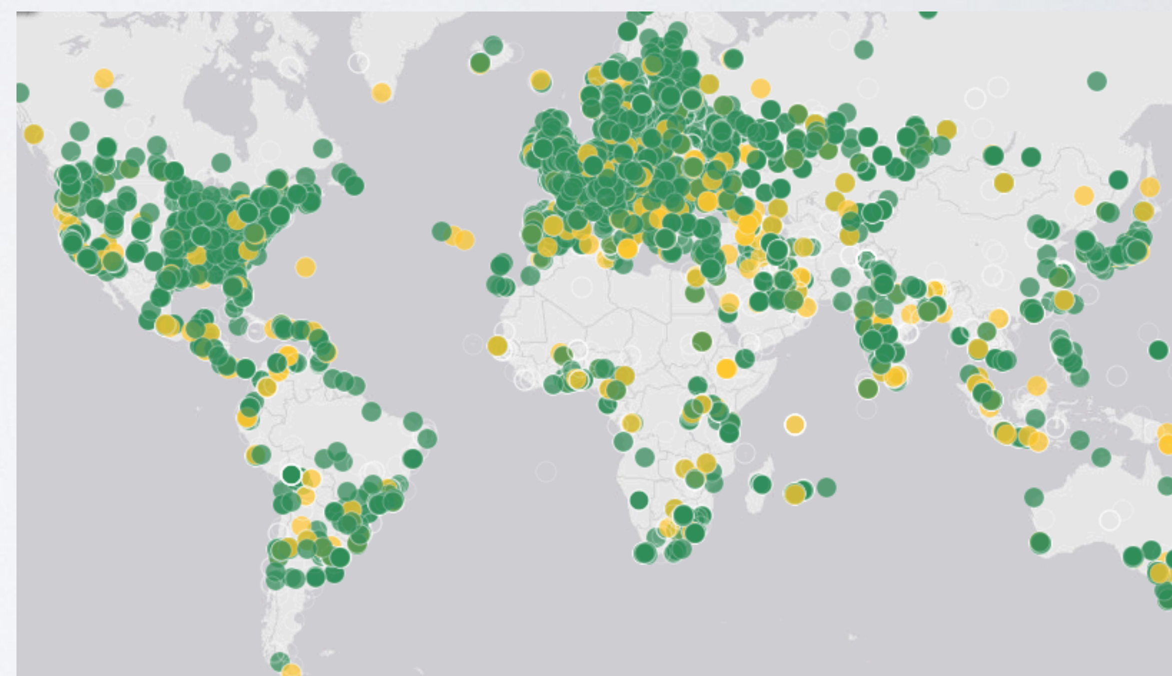

Source: RIPE NCC's Atlas measurement platform, https://www.atlas.ripe.net, 2020

Datasets

- **Active traceroute (monitors and topology data)**
  - CAIDA's Ark (100+ monitors)
  - RIPE NCC's Atlas (10,000+ monitors)
- **Geolocation**
  - Netacuity
  - Hostnames-based geolocation
- **CAIDA's MIDAR, Vela Alias, Macroscopic Internet Topology Dataset Kit (ITDK)**
- **CAIDA's IXP dataset**

Step1: Collect Candidate IP Paths crossing SACS

- Collect detailed information about SACS launch

AS topology post-SACS launch

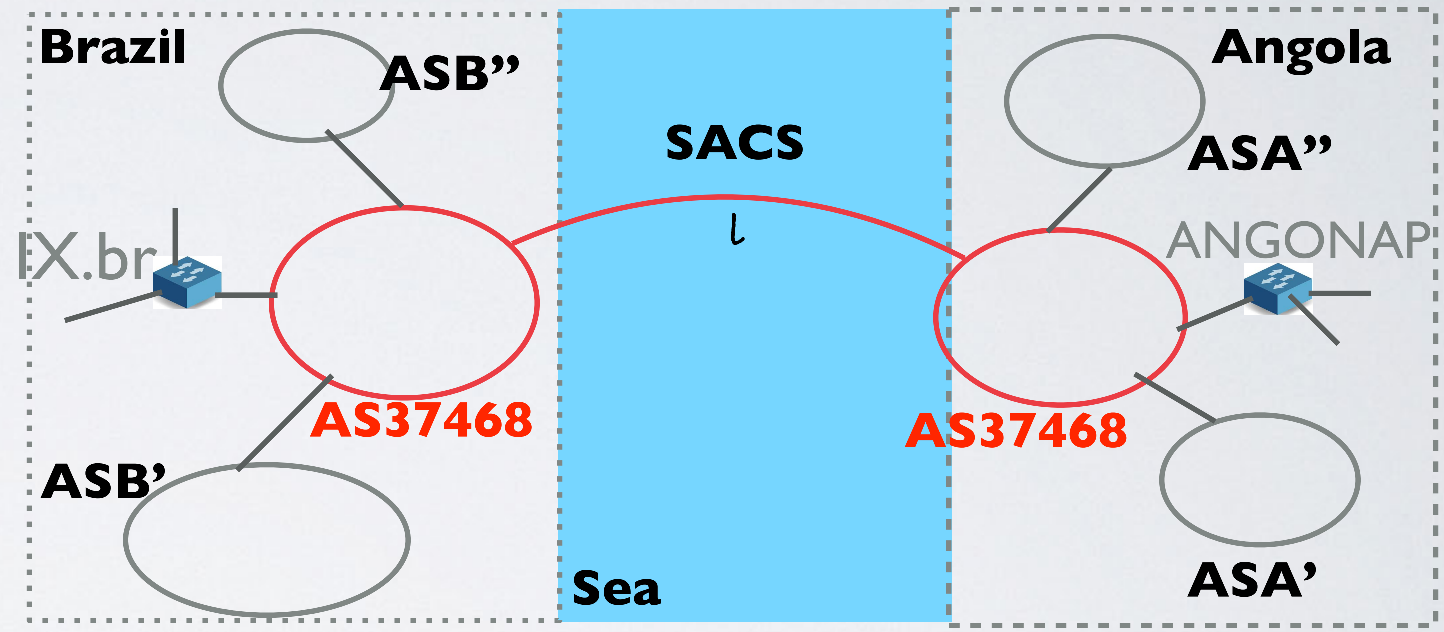

Legend

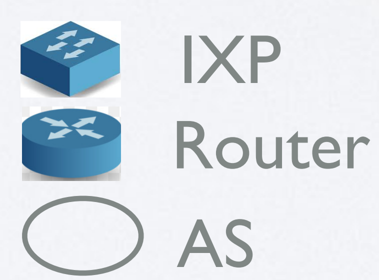

caida

Step1: Collect Candidate IP Paths crossing SACS

- Collect detailed information about SACS launch
- Select all Ark VPs in Brazil and the Angola Cable (AC) Looking Glass in Angola

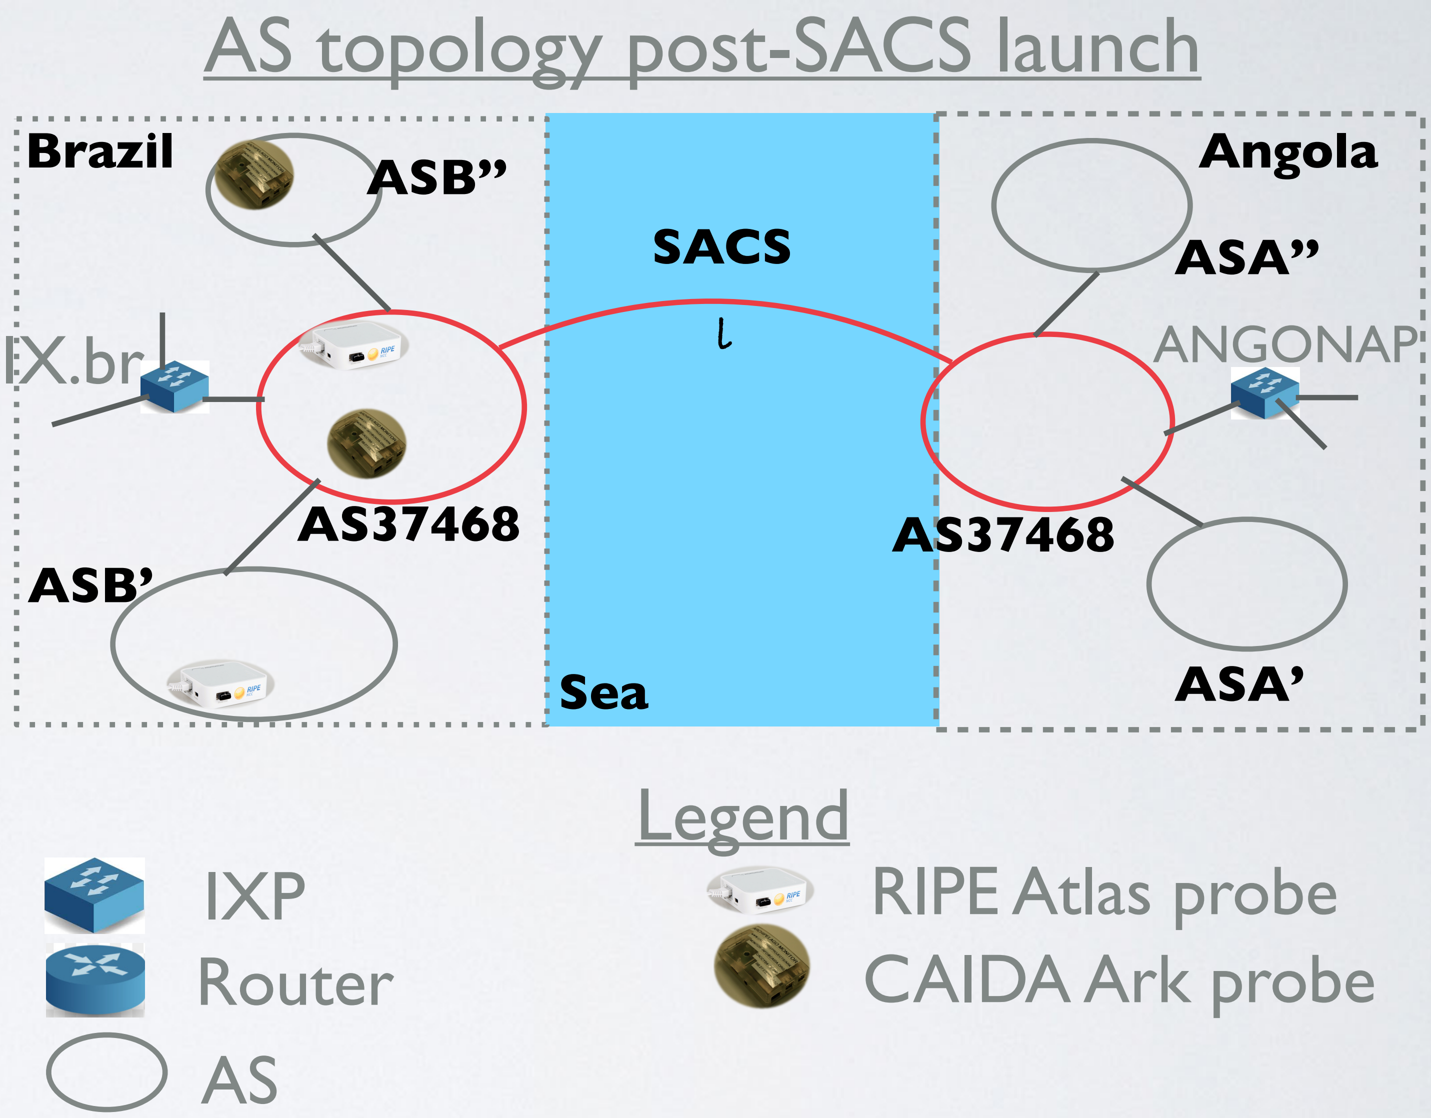

Step1: Collect Candidate IP Paths crossing SACS

- Collect detailed information about SACS launch
- Select all Ark VPs in Brazil and the Angola Cable (AC) Looking Glass in Angola
- Run traceroutes in both directions to obtain **candidate IP paths**

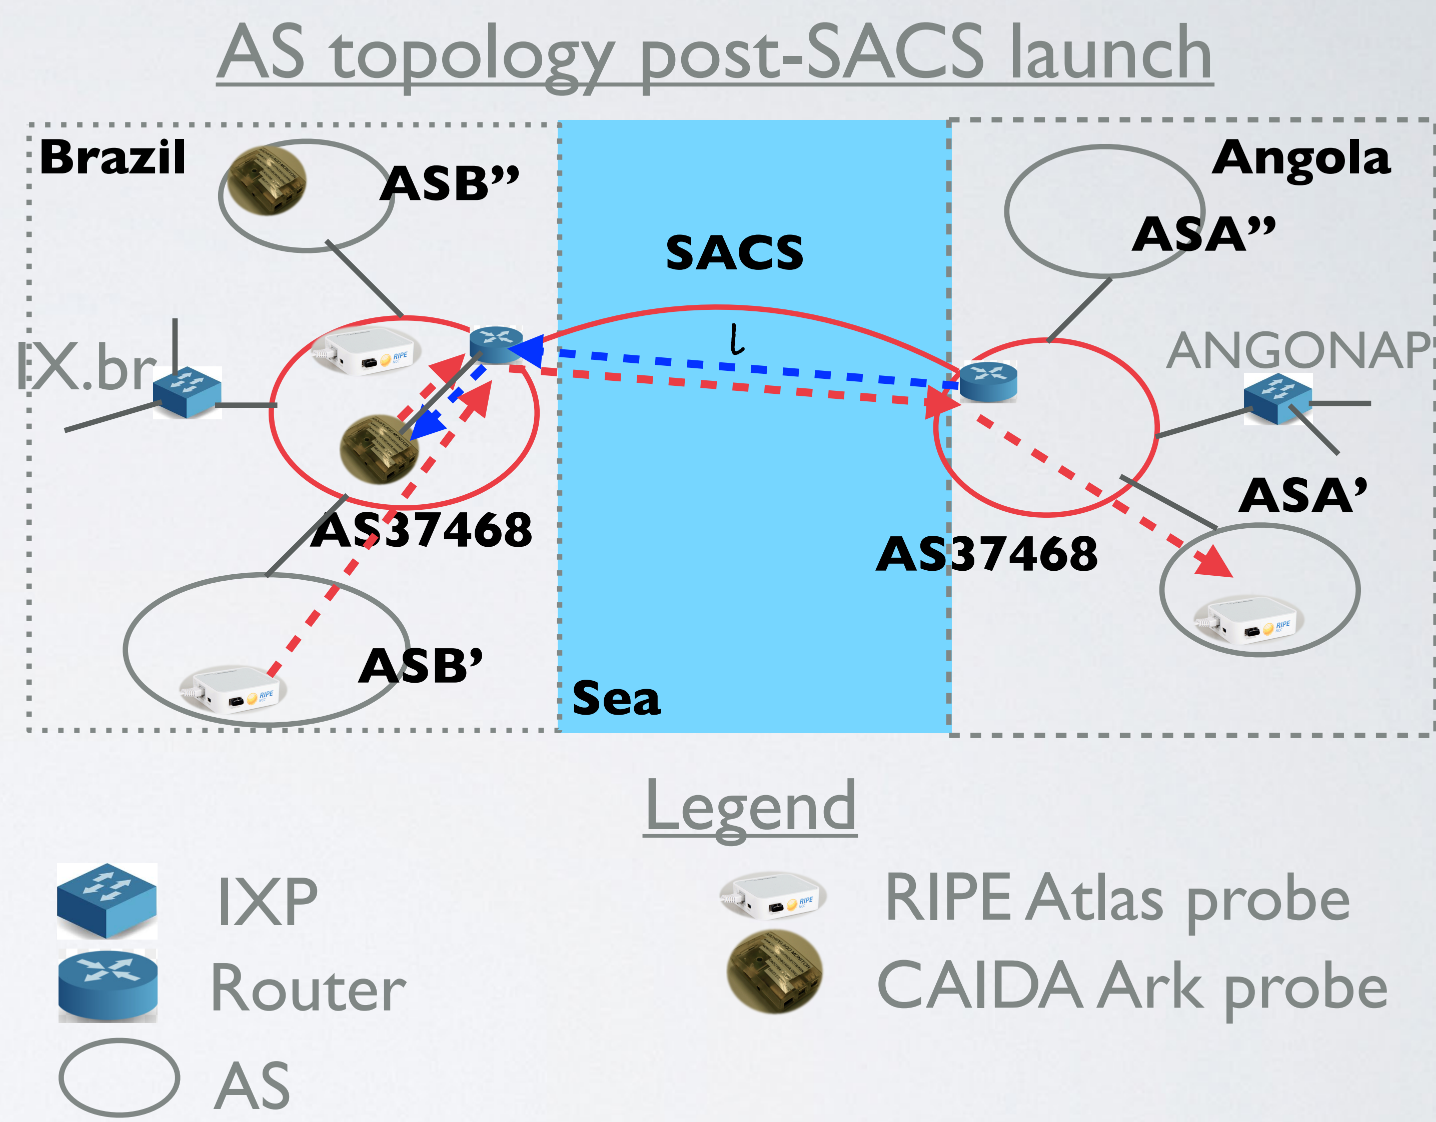

Step2 (1): Identify Link IPs

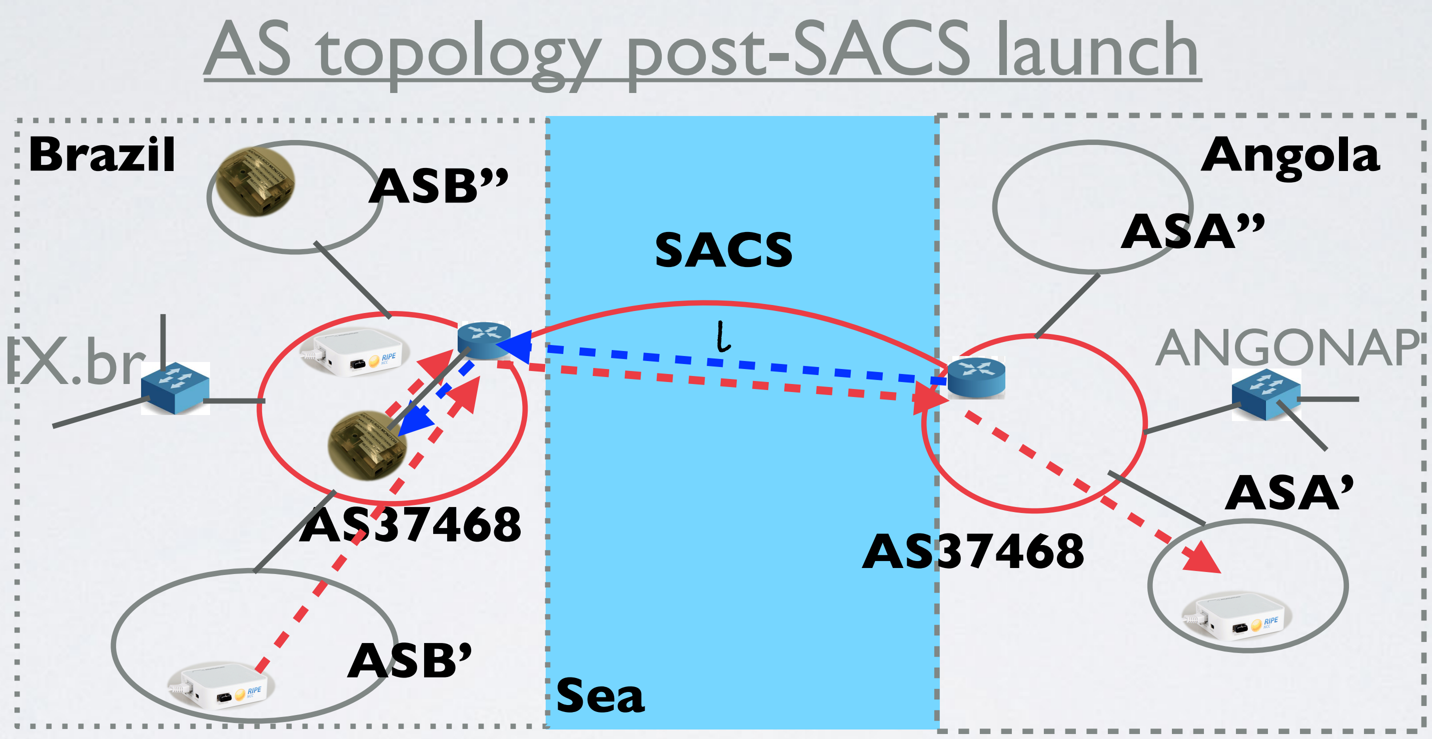

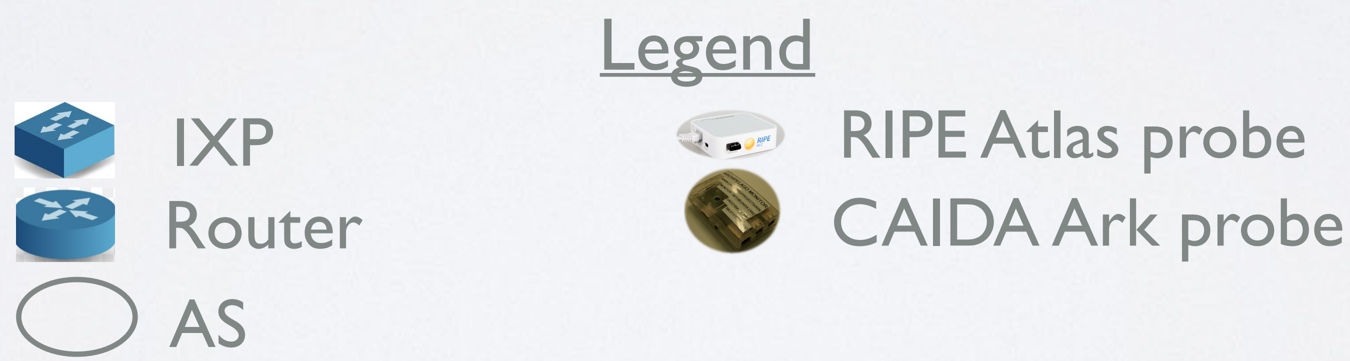

Compute $t = \frac{2L}{\frac{2c}{3}} = \frac{3L}{c}$, with $\frac{2c}{3}$ the speed of light traveling fibre optics

Step2 (2): Identify Link IPs

AS37468 Looking Glass

Router: AO-Luanda
Command: traceroute inet 189.6.232.127

$$\triangle_{RTT} \simeq 61ms \approx t = \frac{3 \times 6{,}165{,}000}{300{,}000} = 61.6ms$$

```
 1  pe2-nc026.ang.sgn.as37468.angolacables.ao (197.149.149.162)  1.432 ms  0.994 ms  1.292 ms
 2  170.238.232.145 (170.238.232.145)  61.699 ms  61.582 ms  61.808 ms
 3  pe1-nc014.br.ftz.as37468.angolacables.ao (170.238.232.97)  61.883 ms pe1-nc013.br.ftz.as37468.angolacables.ao (170.238.232.81)
 4  lag-9.ar1.rio1.gig.gblx.net (64.210.72.221)  62.455 ms  62.319 ms  62.496 ms
 5  * * *
 6  ebt-b56-uacc04.rjo.embratel.net.br (200.211.219.137)  111.980 ms  114.964 ms  111.907 ms
 7  ebt-h0-11-0-0-tcore01.rjoen.embratel.net.br (200.244.211.203)  129.229 ms ebt-h0-3-0-2-tcore01.rjo.embratel.net.br (200.244.211.
     MPLS Label=24633 CoS=0 TTL=1 S=1
 8  ebt-b1211-tcore01.ctamc.embratel.net.br (200.230.231.57)  129.993 ms  132.195 ms ebt-b10-tcore01.rjo.embratel.net.br (200.230.25
     MPLS Label=25621 CoS=0 TTL=1 S=0
     MPLS Label=24082 CoS=0 TTL=1 S=1
 9  ebt-b1451-tcore01.pae.embratel.net.br (200.230.251.217)  127.899 ms ebt-b1211-tcore01.ctamc.embratel.net.br (200.230.231.57)  13
     MPLS Label=24082 CoS=0 TTL=1 S=1
10  ebt-h0-1-0-0-uacc01.pae.embratel.net.br (200.244.213.26)  125.608 ms ebt-b1451-tcore01.pae.embratel.net.br (200.230.251.217)  13
     MPLS Label=24000 CoS=0 TTL=1 S=1
11  ebt-h0-1-0-0-uacc01.pae.embratel.net.br (200.244.213.26)  161.470 ms bd06e87f.virtua.com.br (189.6.232.127)  131.173 ms ebt-h0-1
```

Candidate IP path 1: Traceroute from AC Looking glass in Angola to Ark probe poa-br located in Brazil on 03/25/2019 (post-SACS)

Step2 (2): Identify Link IPs

hostnames hinting IPs in Brazil

AS37468 Looking Glass

hostname hinting IP in Angola

Router: AO-Luanda
Command: traceroute inet 189.6.232.127

```
 1  pe2-nc026.ang.sgn.as37468.angolacables.ao (197.149.149.162)  1.432 ms  0.994 ms  1.292 ms
 2  170.238.232.145 (170.238.232.145)  61.699 ms  61.582 ms  61.808 ms
 3  pe1-nc014.br.ftz.as37468.angolacables.ao (170.238.232.97)  61.883 ms pe1-nc013.br.ftz.as37468.angolacables.ao (170.238.232.81)
 4  lag-9.ar1.rio1.gig.gblx.net (64.210.72.221)  62.455 ms  62.319 ms  62.496 ms
 5  * * *
 6  ebt-b56-uacc04.rjo.embratel.net.br (200.211.219.137)  111.980 ms  114.964 ms  111.907 ms
 7  ebt-h0-11-0-0-tcore01.rjoen.embratel.net.br (200.244.211.203)  129.229 ms ebt-h0-3-0-2-tcore01.rjo.embratel.net.br (200.244.211.
     MPLS Label=24633 CoS=0 TTL=1 S=1
 8  ebt-b1211-tcore01.ctamc.embratel.net.br (200.230.231.57)  129.993 ms  132.195 ms ebt-b10-tcore01.rjo.embratel.net.br (200.230.25
     MPLS Label=25621 CoS=0 TTL=1 S=0
     MPLS Label=24082 CoS=0 TTL=1 S=1
 9  ebt-b1451-tcore01.pae.embratel.net.br (200.230.251.217)  127.899 ms ebt-b1211-tcore01.ctamc.embratel.net.br (200.230.231.57)  13
     MPLS Label=24082 CoS=0 TTL=1 S=1
10  ebt-h0-1-0-0-uacc01.pae.embratel.net.br (200.244.213.26)  125.608 ms ebt-b1451-tcore01.pae.embratel.net.br (200.230.251.217)  13
     MPLS Label=24000 CoS=0 TTL=1 S=1
11  ebt-h0-1-0-0-uacc01.pae.embratel.net.br (200.244.213.26)  161.470 ms bd06e87f.virtua.com.br (189.6.232.127)  131.173 ms ebt-h0-1
```

Candidate IP path 1: Traceroute from AC Looking glass in Angola to Ark probe poa-br located in Brazil on 03/25/2019 (post-SACS)

caida

Step2 (2): Identify Link IPs

AS37468 Looking Glass

Router: AO-Luanda
Command: traceroute inet 189.6.232.127

IPs on both sides of the cable

```
 1  pe2-nc026.ang.sgn.as37468.angolacables.ao (197.149.149.162)  1.432 ms  0.994 ms  1.292 ms
 2  170.238.232.145 (170.238.232.145)  61.699 ms  61.582 ms  61.808 ms
 3  pe1-nc014.br.ftz.as37468.angolacables.ao (170.238.232.97)  61.883 ms pe1-nc013.br.ftz.as37468.angolacables.ao (170.238.232.81)
 4  lag-9.ar1.rio1.gig.gblx.net (64.210.72.221)  62.455 ms  62.319 ms  62.496 ms
 5  * * *
 6  ebt-b56-uacc04.rjo.embratel.net.br (200.211.219.137)  111.980 ms  114.964 ms  111.907 ms
 7  ebt-h0-11-0-0-tcore01.rjoen.embratel.net.br (200.244.211.203)  129.229 ms ebt-h0-3-0-2-tcore01.rjo.embratel.net.br (200.244.211.
     MPLS Label=24633 CoS=0 TTL=1 S=1
 8  ebt-b1211-tcore01.ctamc.embratel.net.br (200.230.231.57)  129.993 ms  132.195 ms ebt-b10-tcore01.rjo.embratel.net.br (200.230.25
     MPLS Label=25621 CoS=0 TTL=1 S=0
     MPLS Label=24082 CoS=0 TTL=1 S=1
 9  ebt-b1451-tcore01.pae.embratel.net.br (200.230.251.217)  127.899 ms ebt-b1211-tcore01.ctamc.embratel.net.br (200.230.231.57)  13
     MPLS Label=24082 CoS=0 TTL=1 S=1
10  ebt-h0-1-0-0-uacc01.pae.embratel.net.br (200.244.213.26)  125.608 ms ebt-b1451-tcore01.pae.embratel.net.br (200.230.251.217)  13
     MPLS Label=24000 CoS=0 TTL=1 S=1
11  ebt-h0-1-0-0-uacc01.pae.embratel.net.br (200.244.213.26)  161.470 ms bd06e87f.virtua.com.br (189.6.232.127)  131.173 ms ebt-h0-1
```

Candidate IP path 1: Traceroute from AC Looking glass in Angola to Ark probe poa-br located in Brazil on 03/25/2019 (post-SACS)

Step2 (3): Identify Link IPs

Other Views: USA | South America | Europe | China | Japan

traceroute to 197.149.149.29 (AS37468) from **poa-br** using UDP

Hop	Address	Prefix	AS	Location	RTT (ms)
1	192.168.0.1				3.3
2	10.112.0.1				6.8
3	**c915c064.virtua.com.br** 201.21.192.100	201.21.192.0/18	28573	sao paulo bra	13.3
4	**c915c001.virtua.com.br** 201.21.192.1	201.21.192.0/18	28573	sao paulo bra	10.4
5	**embratel-T0-0-1-2-4004-uacc02.ctamr.embratel.net.br** 201.64.24.13	201.64.0.0/16	4230	rio de janeiro bra	27.7
6	**ligacaao-H0-14-0-0-tcore01.ctamr.embratel.net.br** 200.244.211.119	200.244.0.0/16	4230	rio de janeiro bra	35.1
7	**ebt-B1421-tcore01.spo.embratel.net.br** 200.230.231.62	200.230.0.0/16	4230	sao paulo bra	37.3
8	**ebt-H0-0-1-1-agg01.spo.embratel.net.br** 200.230.242.46	200.230.0.0/16	4230	sao paulo bra	28.0
9	**peer-T0-0-1-0-9-agg01.spo.embratel.net.br** 200.211.219.62	200.211.0.0/16	4230	sao paulo bra	66.3
10	**187-100-163-13.dsl.telesp.net.br** 187.100.163.13	187.100.0.0/16	27699	sao paulo bra	29.1
11	**187-92-6-14.customer.tdatabrasil.net.br** 187.92.6.14	187.92.0.0/16	10429	barueri bra	28.6
12	**pe2-nce024.br.spa.as37468.angolacables.ao** 170.238.232.62	170.238.232.0/22	37468	rio de janeiro bra	34.5
13	**pe1-nc011.br.ftz.as37468.angolacables.ao** 170.238.232.66	170.238.232.0/22	37468	rio de janeiro bra	74.3
14	**pe2-nc023.br.ftz.as37468.angolacables.ao** 170.238.232.82	170.238.232.0/22	37468	rio de janeiro bra	83.2
15	170.238.232.146	170.238.232.0/22	37468	rio de janeiro bra	138.6
16	**pe2-nc024.ang.lua.as37468.angolacables.ao** 197.149.149.153	197.149.148.0/22	37468	luanda ago	141.6
17	**pe1-ne011.ang.lua.as37468.angolacables.ao** 197.149.151.3	197.149.148.0/22	37468	luanda ago	142.4

STOPPED: gap limit exceeded

$\triangle_{RTT} \approx t = 61.65ms$

Candidate IP path 2: Traceroute output from poa-br to AC Looking glass on 03/25/2019 (post-SACS) using CAIDA Vela

Step2 (3): Identify Link IPs

Other Views: USA | South America | Europe | China | Japan

traceroute to 197.149.149.29 (AS37468) from **poa-br** using UDP

Hop	Address	Prefix	AS	Location	RTT (ms)
1	192.168.0.1				3.3
2	10.112.0.1				6.8
3	**c915c064.virtua.com.br** 201.21.192.100	201.21.192.0/18	28573	sao paulo bra	13.3
4	**c915c001.virtua.com.br** 201.21.192.1	201.21.192.0/18	28573	sao paulo bra	10.4
5	**embratel-T0-0-1-2-4004-uacc02.ctamr.embratel.net.br** 201.64.24.13	201.64.0.0/16	4230	rio de janeiro bra	27.7
6	**ligacaao-H0-14-0-0-tcore01.ctamr.embratel.net.br** 200.244.211.119	200.244.0.0/16	4230	rio de janeiro bra	35.1
7	**ebt-B1421-tcore01.spo.embratel.net.br** 200.230.231.62	200.230.0.0/16	4230	sao paulo bra	37.3
8	**ebt-H0-0-1-1-agg01.spo.embratel.net.br** 200.230.242.46	200.230.0.0/16	4230	sao paulo bra	28.0
9	**peer-T0-0-1-0-9-agg01.spo.embratel.net.br** 200.211.219.62	200.211.0.0/16	4230	sao paulo bra	66.3
10	**187-100-163-13.dsl.telesp.net.br** 187.100.163.13	187.100.0.0/16	27699	sao paulo bra	29.1
11	**187-92-6-14.customer.tdatabrasil.net.br** 187.92.6.14	187.92.0.0/16	10429	barueri bra	28.6
12	**pe2-nce024.br.spa.as37468.angolacables.ao** 170.238.232.62	170.238.232.0/22	37468	rio de janeiro bra	34.5
13	**pe1-nc011.br.ftz.as37468.angolacables.ao** 170.238.232.66	170.238.232.0/22	37468	rio de janeiro bra	74.3
14	**pe2-nc023.br.ftz.as37468.angolacables.ao** 170.238.232.82	170.238.232.0/22	37468	rio de janeiro bra	83.2
15	170.238.232.146	170.238.232.0/22	37468	rio de janeiro bra	138.6
16	**pe2-nc024.ang.lua.as37468.angolacables.ao** 197.149.149.153	197.149.148.0/22	37468	luanda ago	141.6
17	**pe1-ne011.ang.lua.as37468.angolacables.ao** 197.149.151.3	197.149.148.0/22	37468	luanda ago	142.4

STOPPED: gap limit exceeded

hostname hinting IPs in Brazil

hostnames hinting IPs in Angola

Candidate IP path 2: Traceroute output from poa-br to AC Looking glass on 03/25/2019 (post-SACS) using CAIDA Vela

caida

Step2 (3): Identify Link IPs

Other Views: USA | South America | Europe | China | Japan

traceroute to 197.149.149.29 (AS37468) from **poa-br** using UDP

Hop	Address	Prefix	AS	Location	RTT (ms)
1	192.168.0.1				3.3
2	10.112.0.1				6.8
3	**c915c064.virtua.com.br** 201.21.192.100	201.21.192.0/18	28573	sao paulo bra	13.3
4	**c915c001.virtua.com.br** 201.21.192.1	201.21.192.0/18	28573	sao paulo bra	10.4
5	**embratel-T0-0-1-2-4004-uacc02.ctamr.embratel.net.br** 201.64.24.13	201.64.0.0/16	4230	rio de janeiro bra	27.7
6	**ligacaao-H0-14-0-0-tcore01.ctamr.embratel.net.br** 200.244.211.119	200.244.0.0/16	4230	rio de janeiro bra	35.1
7	**ebt-B1421-tcore01.spo.embratel.net.br** 200.230.231.62	200.230.0.0/16	4230	sao paulo bra	37.3
8	**ebt-H0-0-1-1-agg01.spo.embratel.net.br** 200.230.242.46	200.230.0.0/16	4230	sao paulo bra	28.0
9	**peer-T0-0-1-0-9-agg01.spo.embratel.net.br** 200.211.219.62	200.211.0.0/16	4230	sao paulo bra	66.3
10	**187-100-163-13.dsl.telesp.net.br** 187.100.163.13	187.100.0.0/16	27699	sao paulo bra	29.1
11	**187-92-6-14.customer.tdatabrasil.net.br** 187.92.6.14	187.92.0.0/16	10429	barueri bra	28.6
12	**pe2-nce024.br.spa.as37468.angolacables.ao** 170.238.232.62	170.238.232.0/22	37468	rio de janeiro bra	34.5
13	**pe1-nc011.br.ftz.as37468.angolacables.ao** 170.238.232.66	170.238.232.0/22	37468	rio de janeiro bra	74.3
14	**pe2-nc023.br.ftz.as37468.angolacables.ao** 170.238.232.82	170.238.232.0/22	37468	rio de janeiro bra	83.2
15	170.238.232.146	170.238.232.0/22	37468	rio de janeiro bra	138.6
16	**pe2-nc024.ang.lua.as37468.angolacables.ao** 197.149.149.153	197.149.148.0/22	37468	luanda ago	141.6
17	**pe1-ne011.ang.lua.as37468.angolacables.ao** 197.149.151.3	197.149.148.0/22	37468	luanda ago	142.4

STOPPED: gap limit exceeded

IPs on both sides of the cable

Candidate IP path 2: Traceroute output from poa-br to AC Looking glass on 03/25/2019 (post-SACS) using CAIDA Vela

Step2 (3): Identify Link IPs

- **Validate IP geolocation**
  - Combine *Netacuity* and *hostname-based geolocation* to locate IPs on both sides of the RTT bump
  - Run RTT measurements to check if min RTT to each IP<10ms from VPs in inferred country
- **Find Link IPs:** Resolve selected IPs' router aliases using CAIDA's MIDAR, Vela alias, and ITDK: obtain 2 lists of IPs R_{Angola} and R_{Brazil} called **link IPs**

Step 3: Fetch Matching Traceroutes

- Use CAIDA's Ark and RIPE Atlas topology data
- Fetch traces with an IP address in $R_{Angola} \rightarrow R_{Brazil}$ or $R_{Brazil} \rightarrow R_{Angola}$ post-SACS (separated by $\Delta RTT \geq t$)
- Collect traces between the same <s,d> pairs pre-SACS, splitting these historical data into 2 sets:
  - Before SACS: ARK-BEFORE/RIPE-BEFORE
  - After SACS: ARK-AFTER/RIPE-AFTER

caida

Step 4: Add Supplementary Datasets

- **Traceroutes annotation:** For every IP hop, add
  - ASN
  - router hostname
  - geolocation information
  - whether or not it belong to an IXP & IXP name

- **Group $T_{<s,d>}$ & $T'_{<s,d>}$ and their annotation per week using their timestamp of execution**

Comparing RTT before & after SACS

- **Challenge:** Finding common IP hops in traces before and after SACS

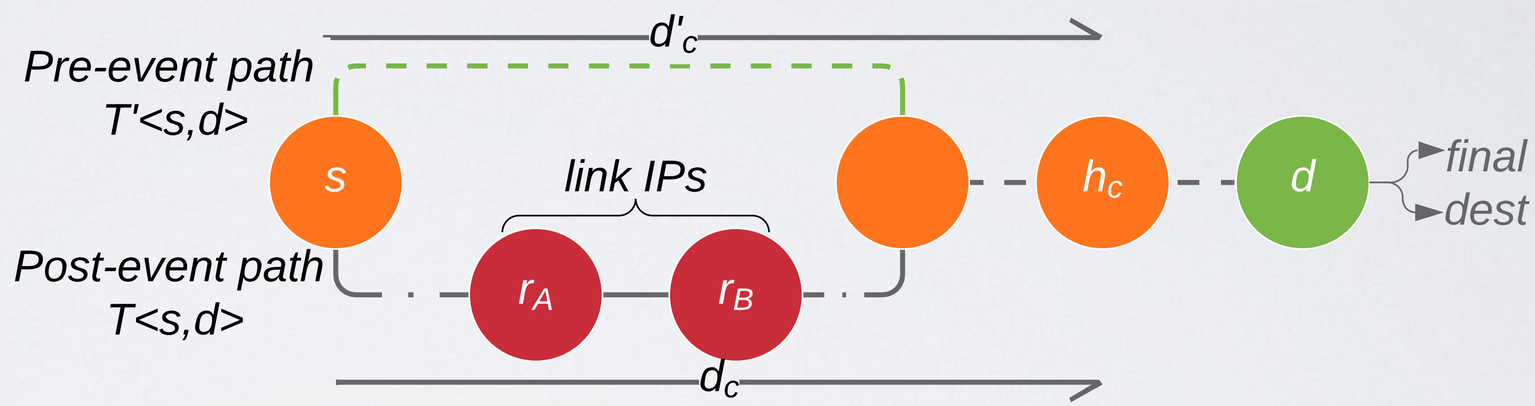

- **Method:**
  - Among all $T_{<s,d>}$ for a given <s,d>, we locate the common IP hop h_c, closest to the destination IP
  - We extract RTT from s to h_c *in* $T_{<s,d>}$ *and* $T'_{<s,d>}$, while ignoring all cases in which $h_c = \emptyset$

Surprising Results

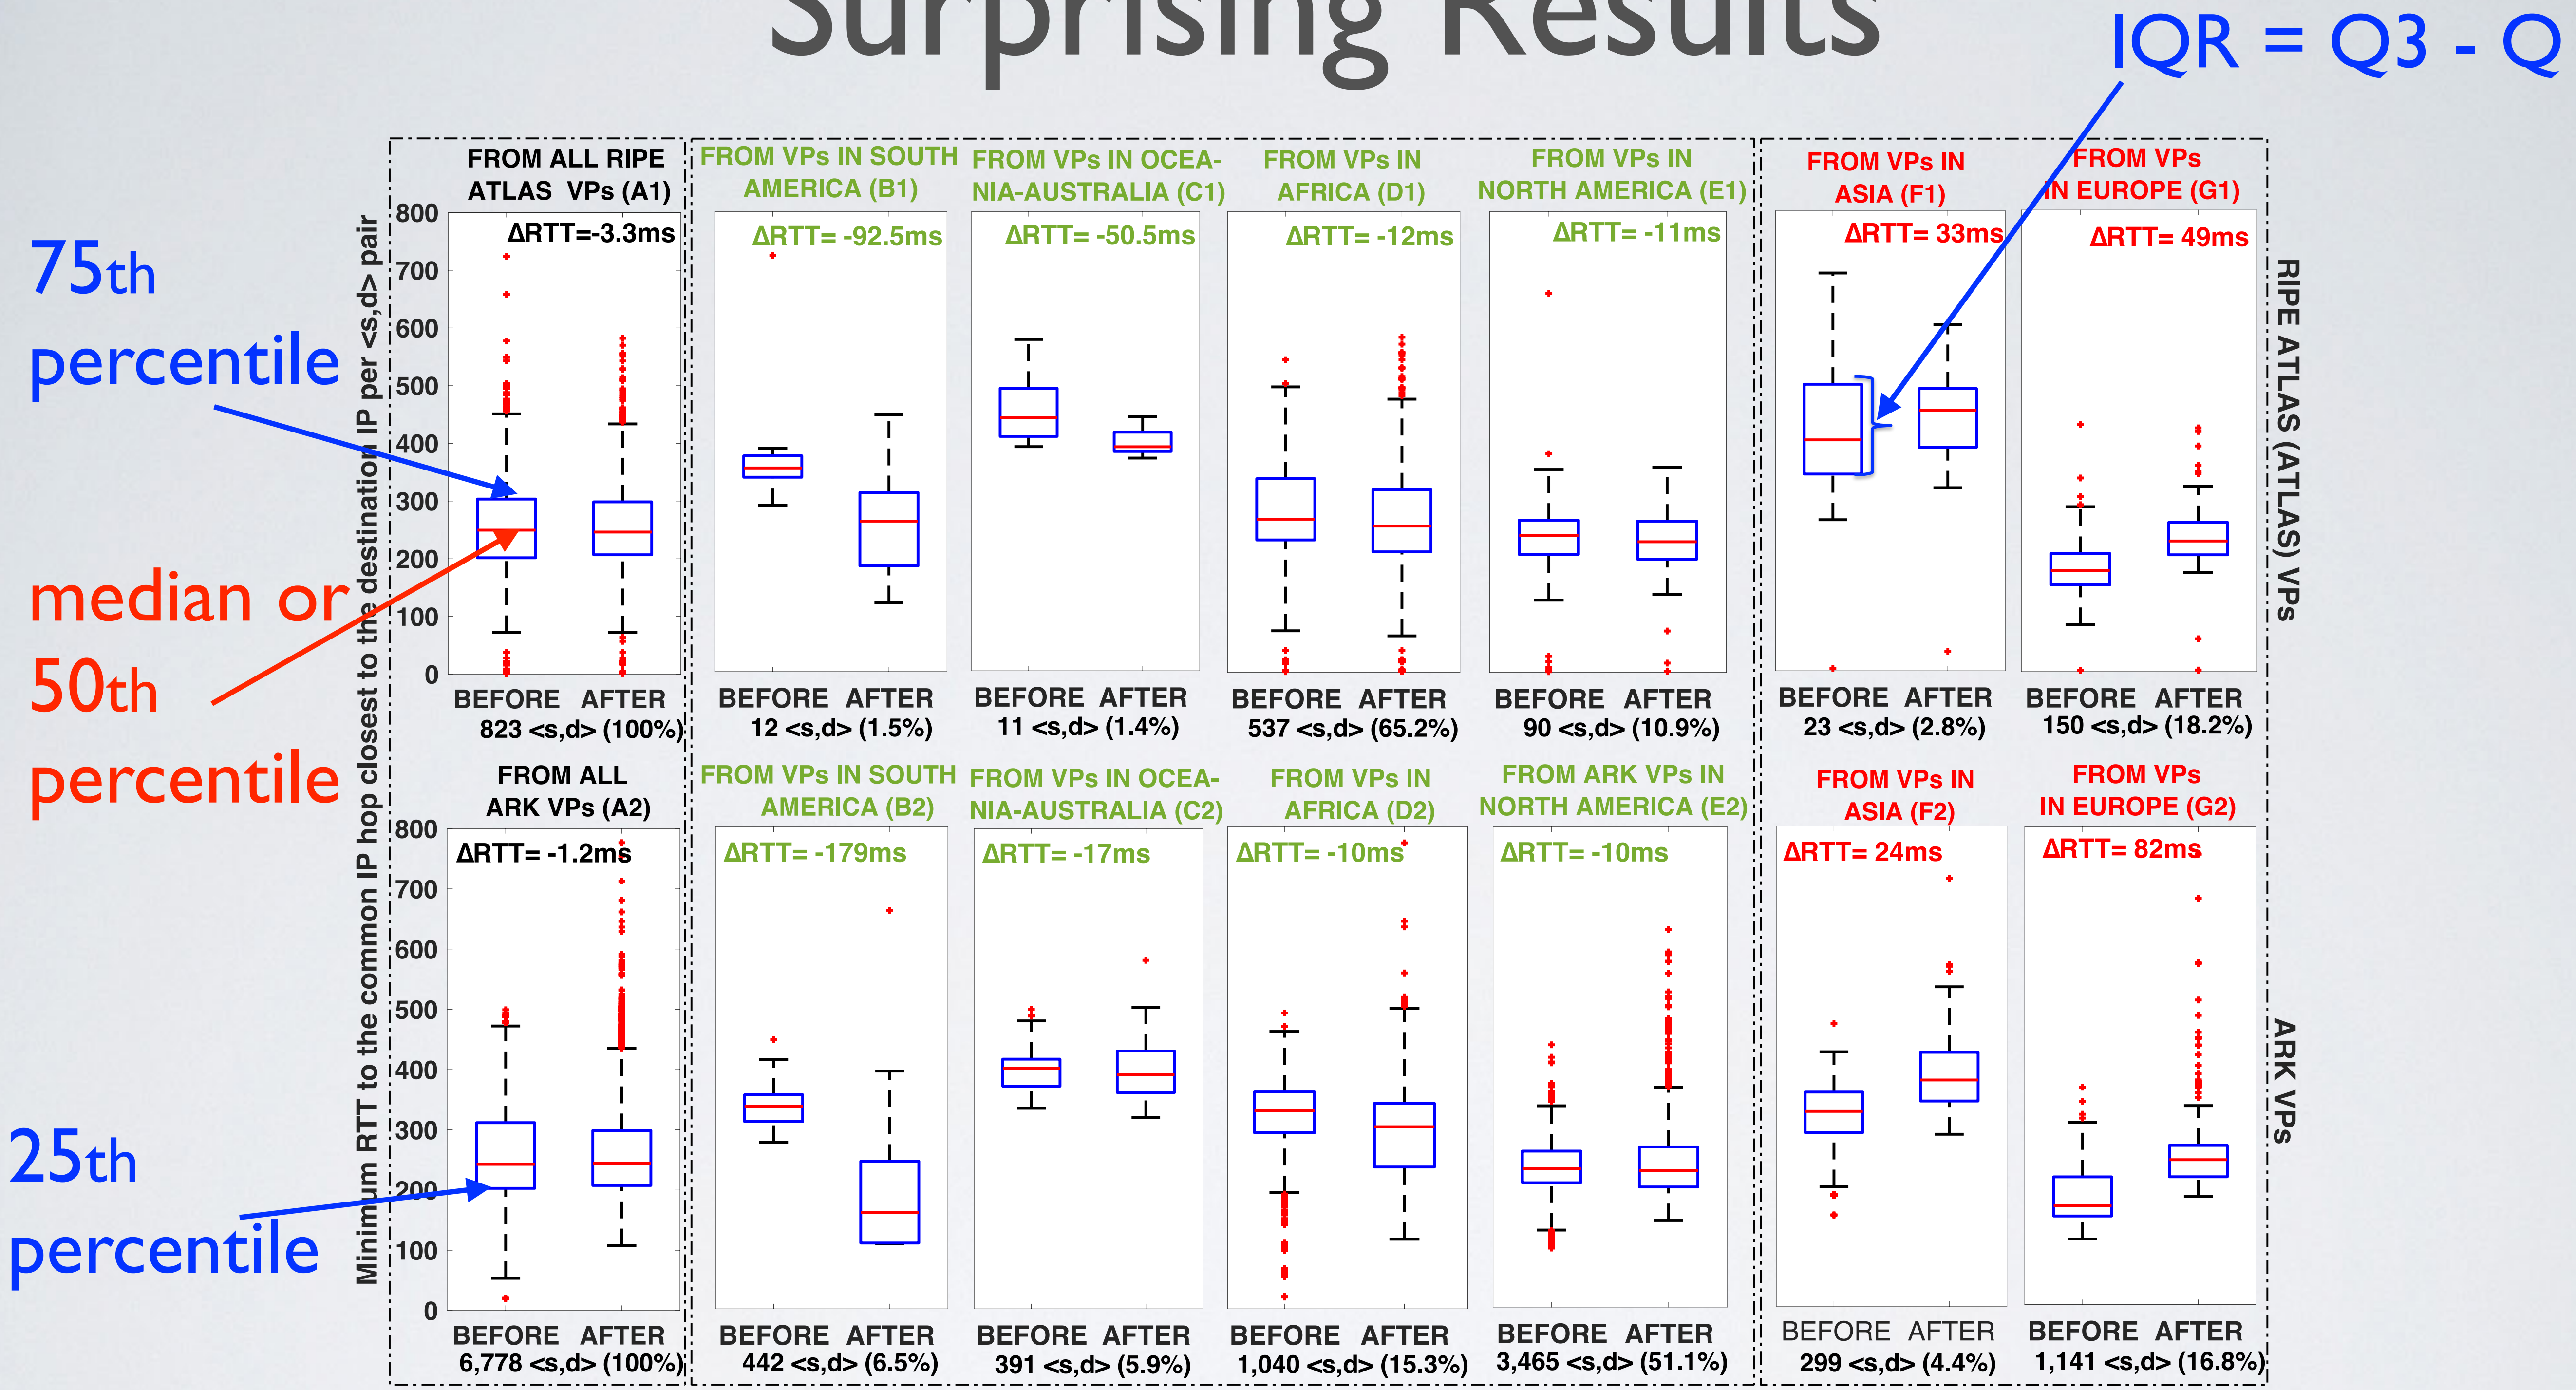

Boxplots of minimum RTTs from Ark & Atlas VPs to the common IP hops closest to the destination IPs.

Surprising Results

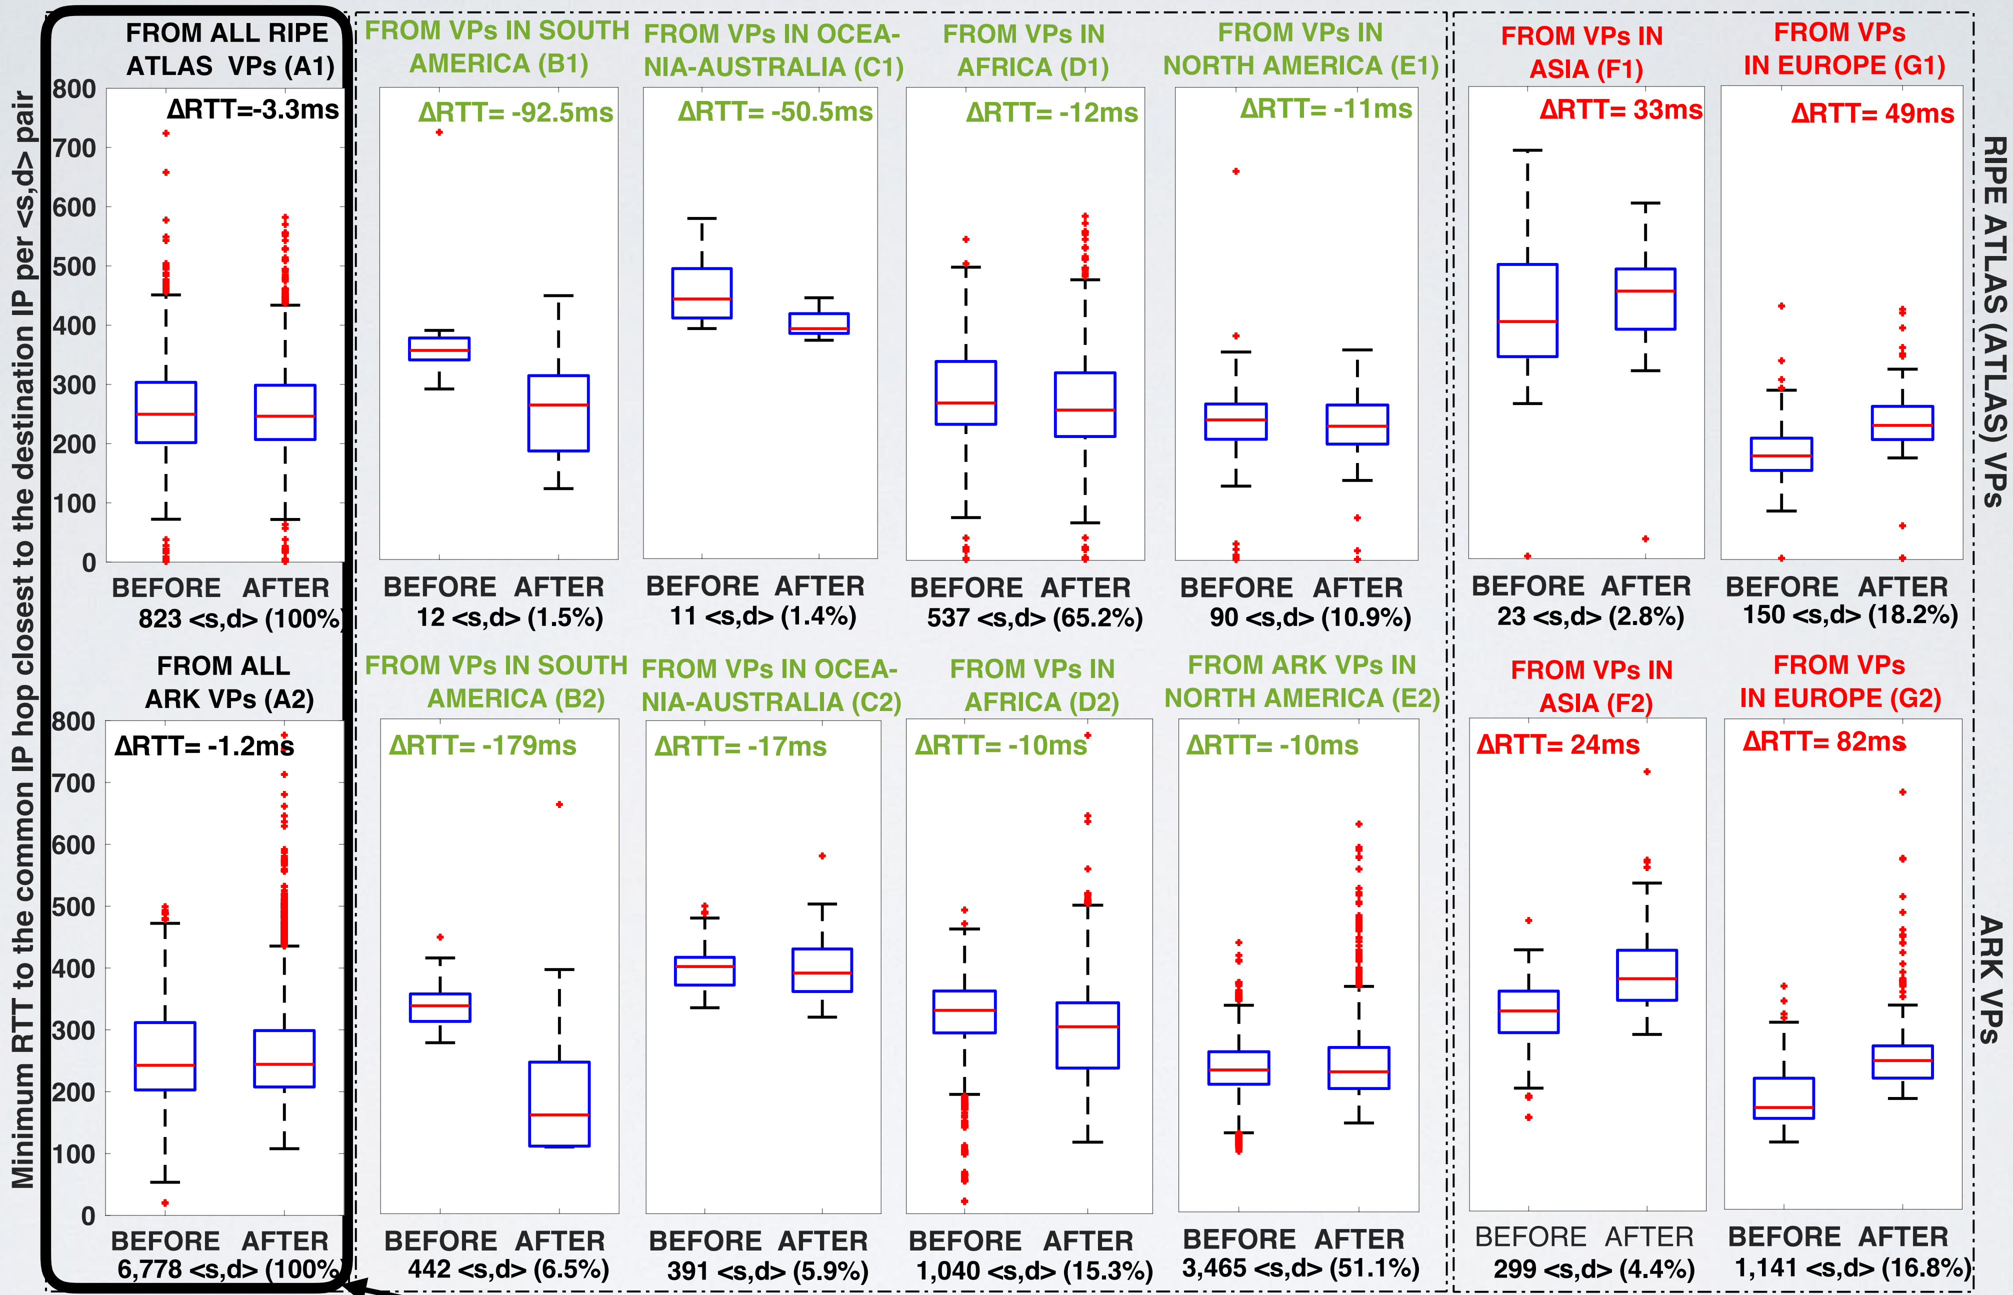

Overall, SACS had **little impact** on latencies from **all VPs to all destinations prefixes**

Surprising Results

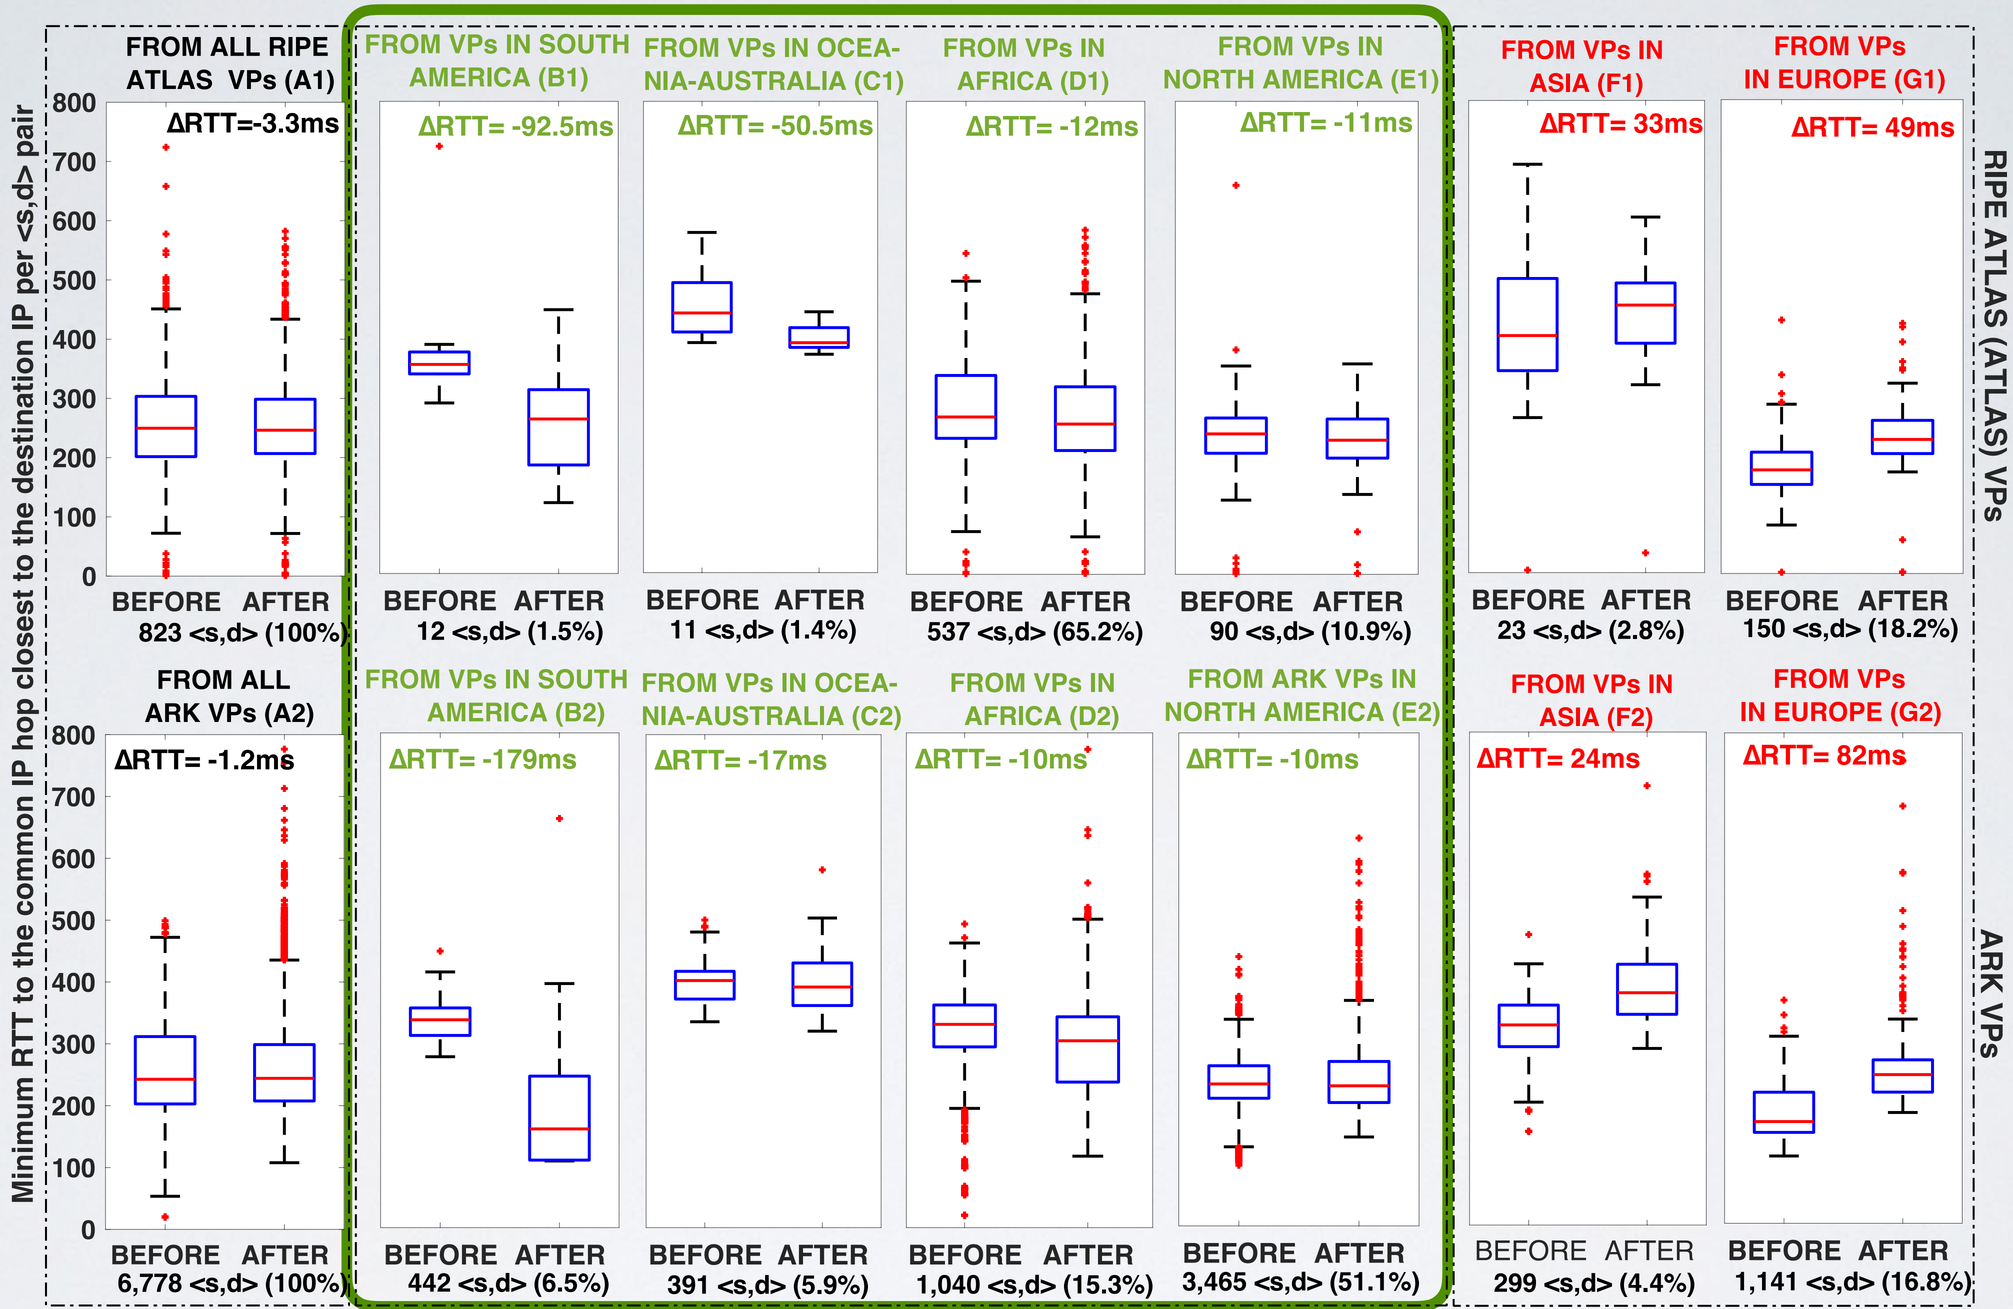

Latencies from VPs in South America significantly dropped,

Surprising Results

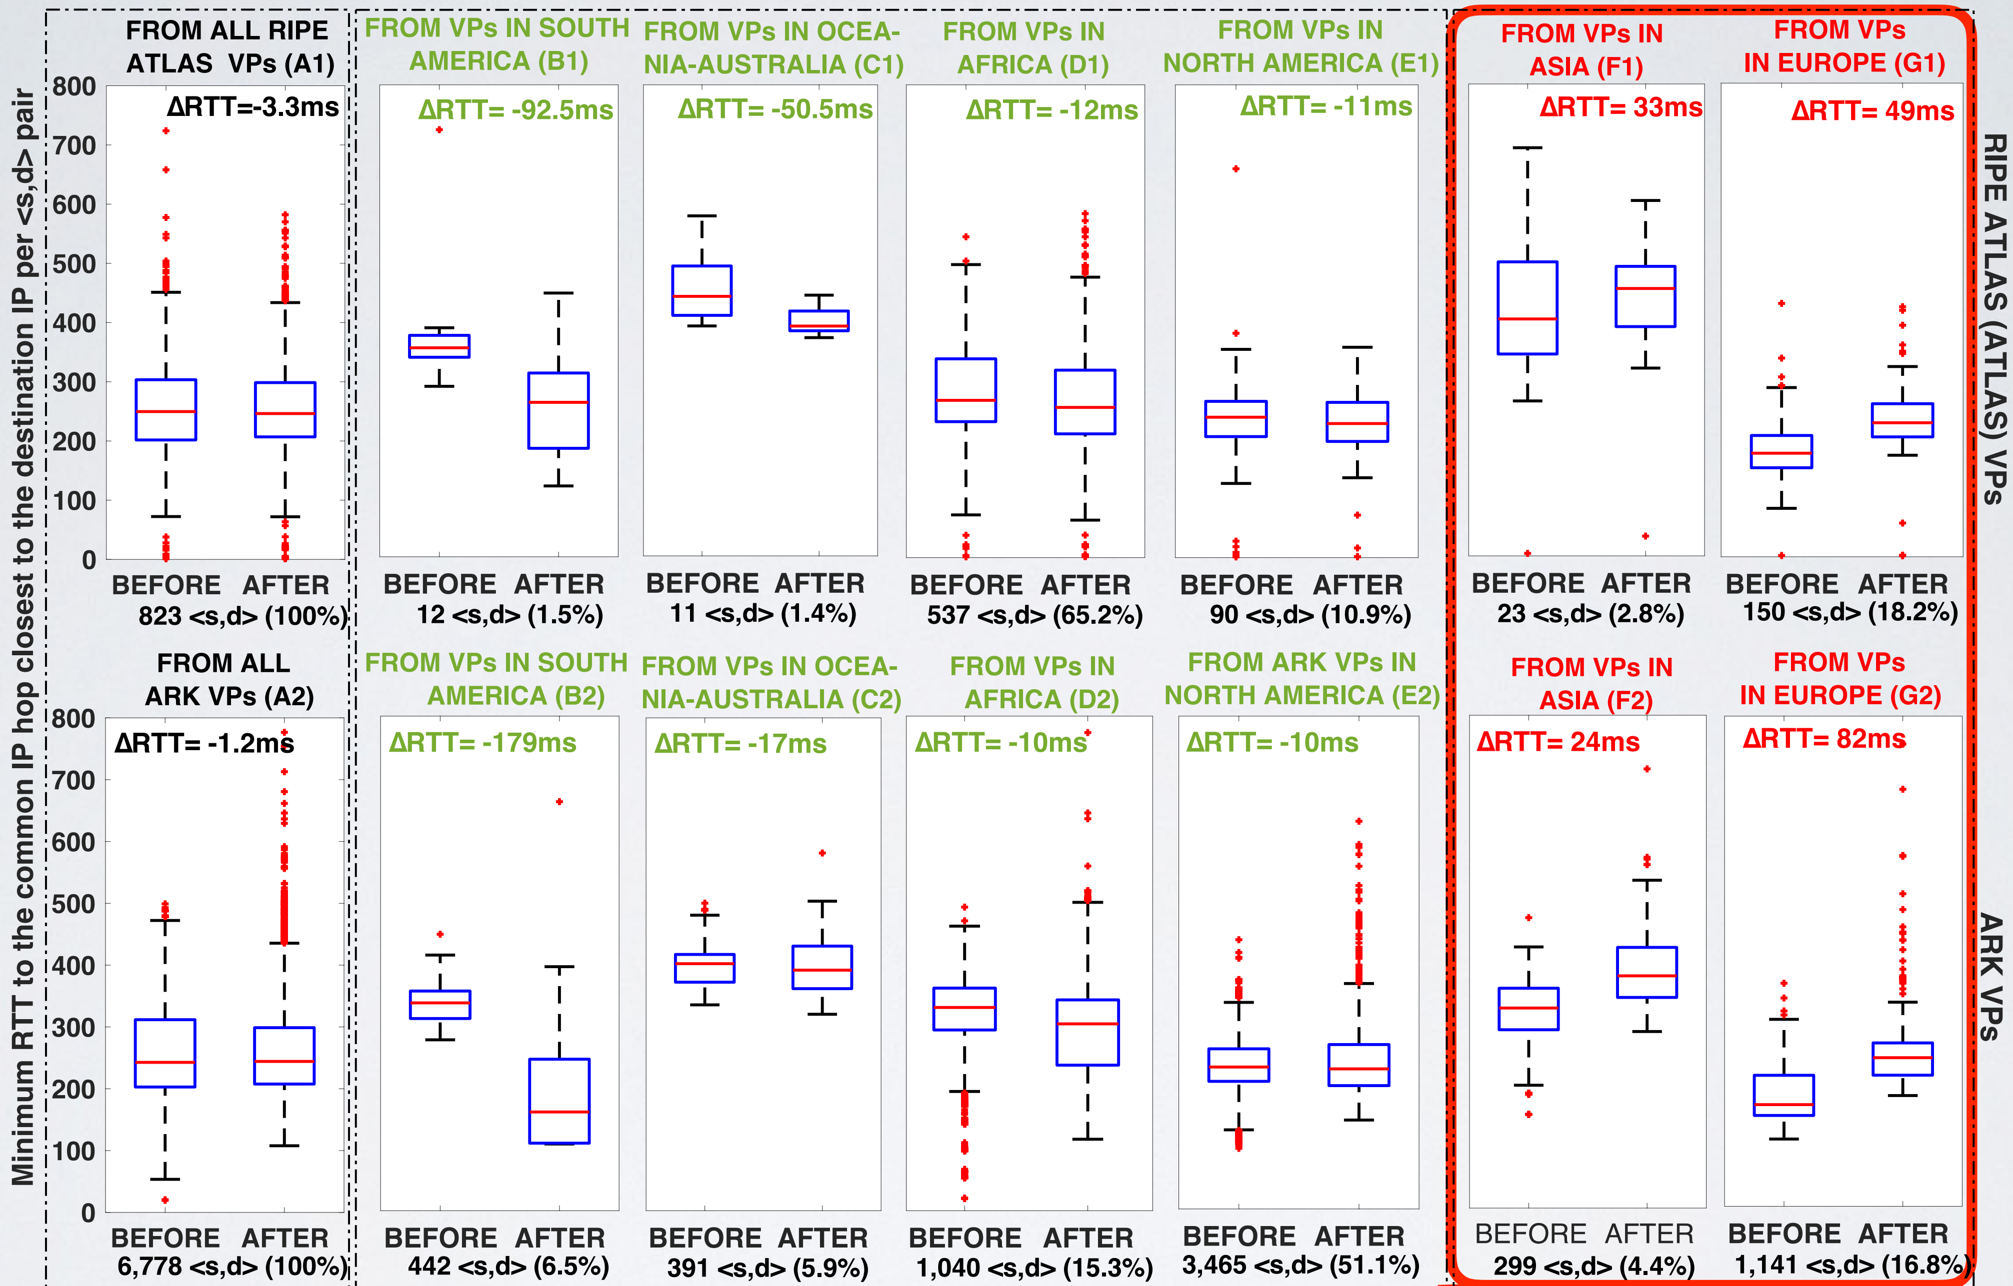

…while latencies from VPs in Europe and Asia significantly increased.

RTT changes at the Country level

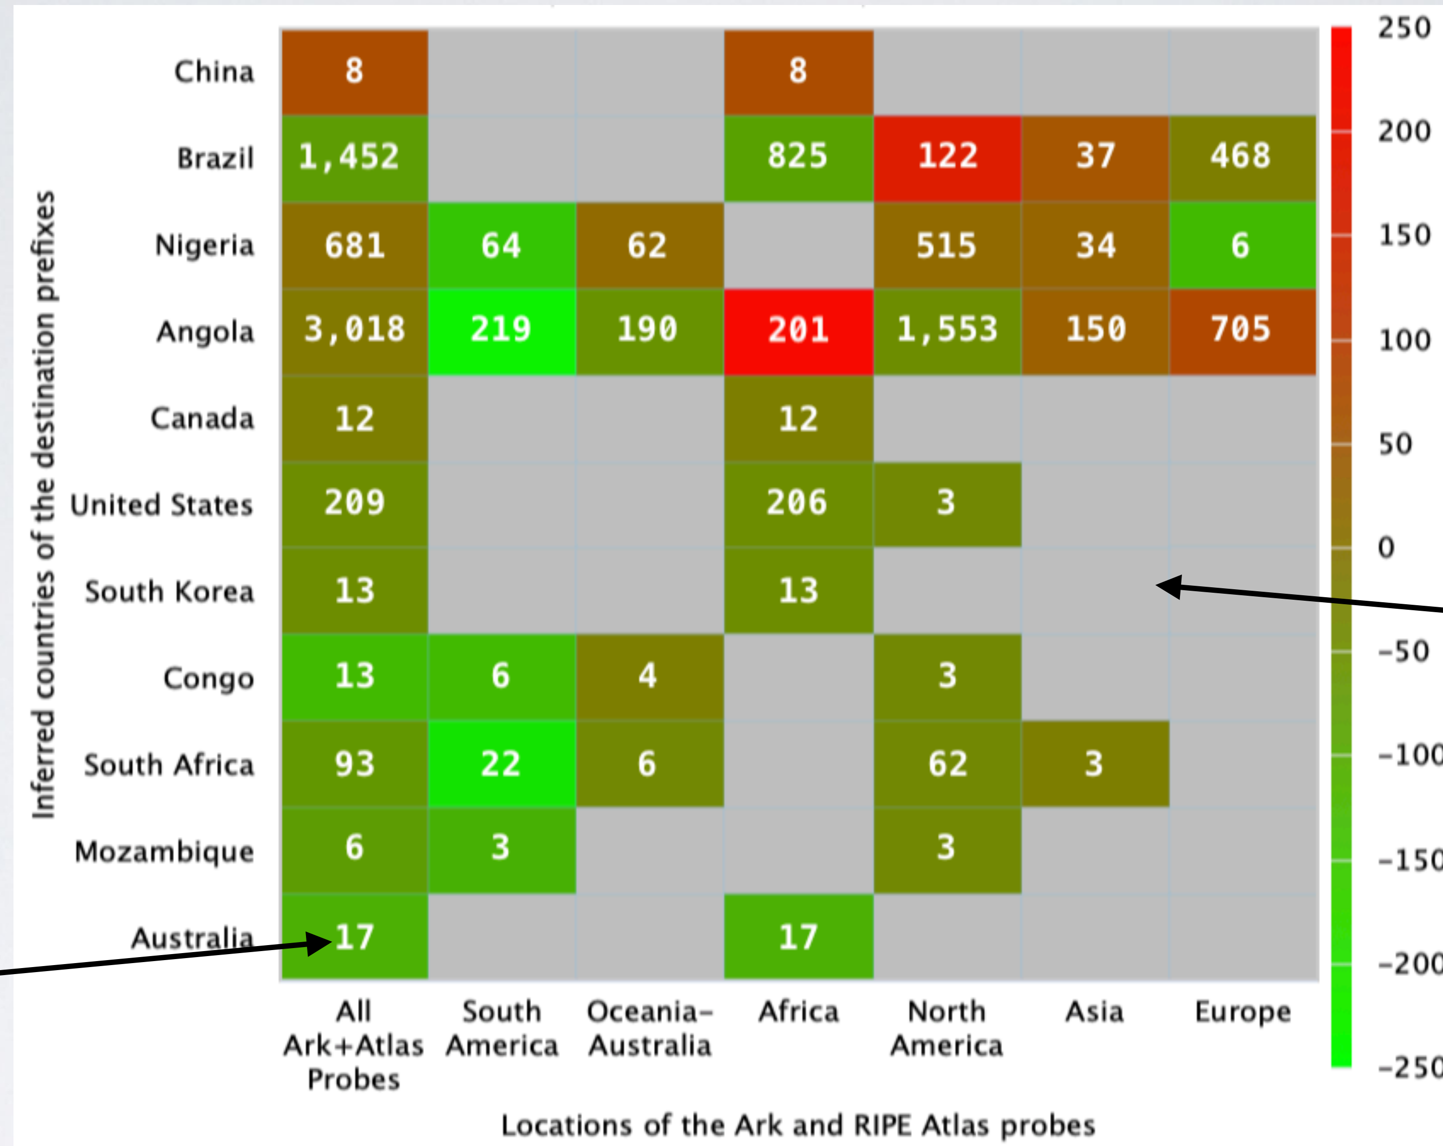

no data available

number of observed <s,d> pairs

ΔRTT AFTER−BEFORE of the medians of minimum RTTs per week pre & post SACS for observed <s,d> pairs.

RTT changes at the Country level

Predictable performance improvement

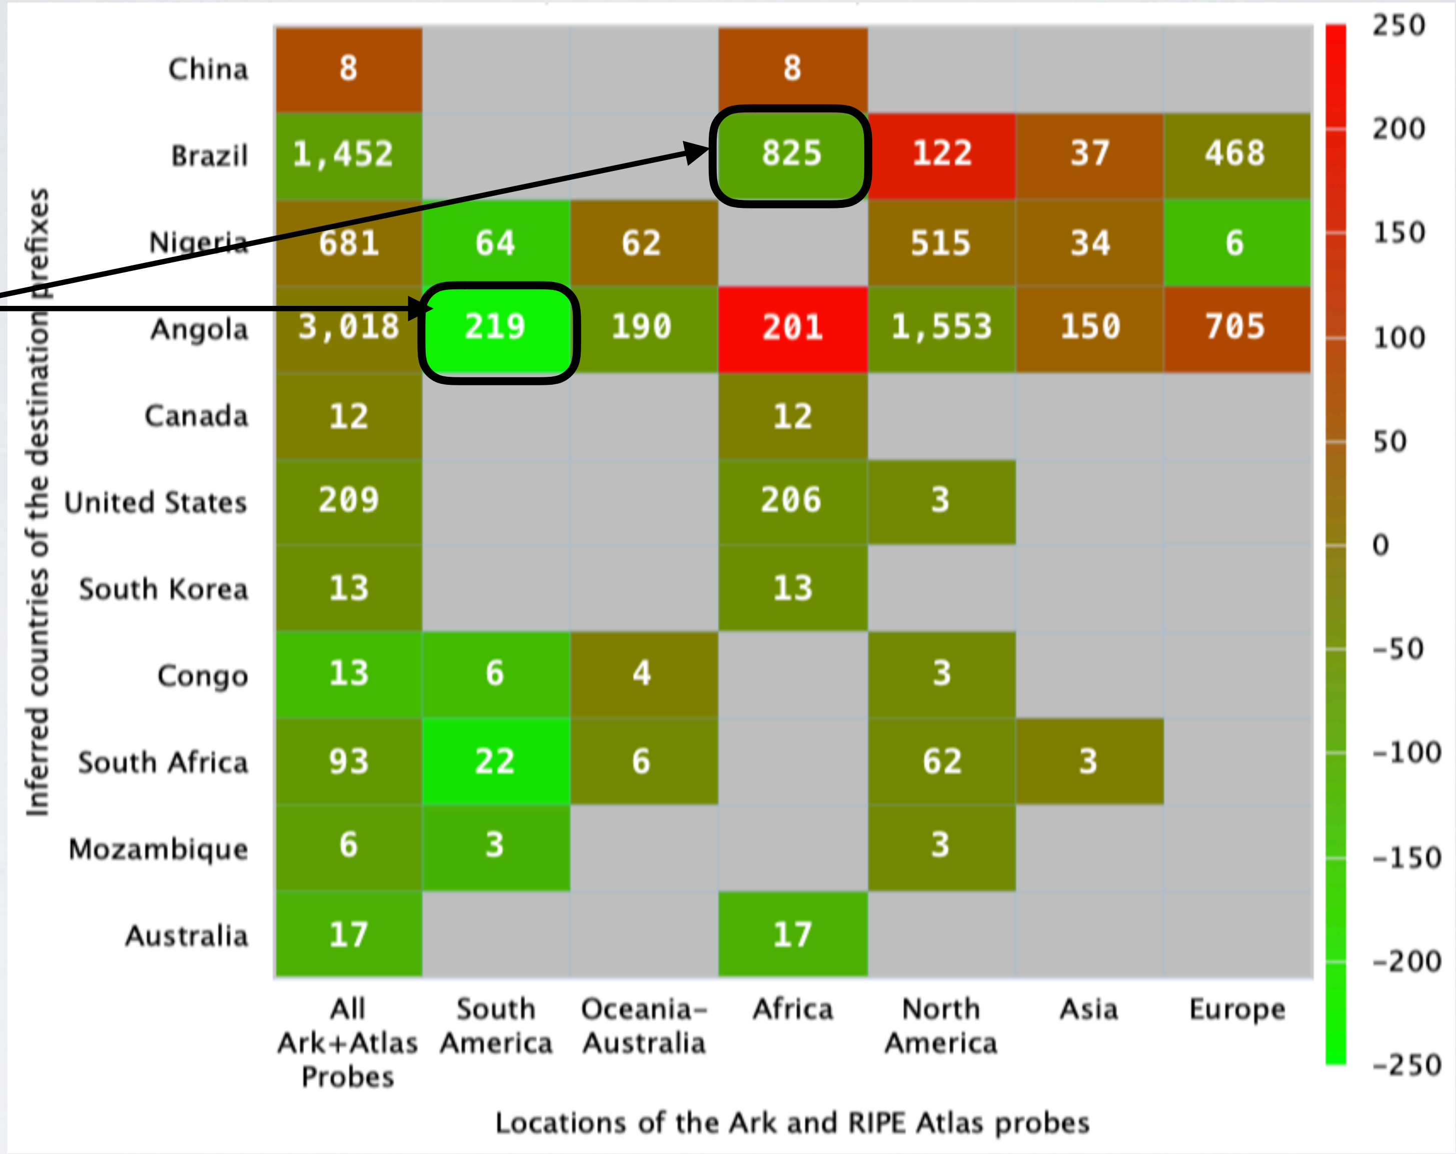

Asymmetrical RTT reduction: the decrease of the median RTT from Africa to Brazil is a third of that from South America to Angola (226ms)

caida

RTT changes at the Country level

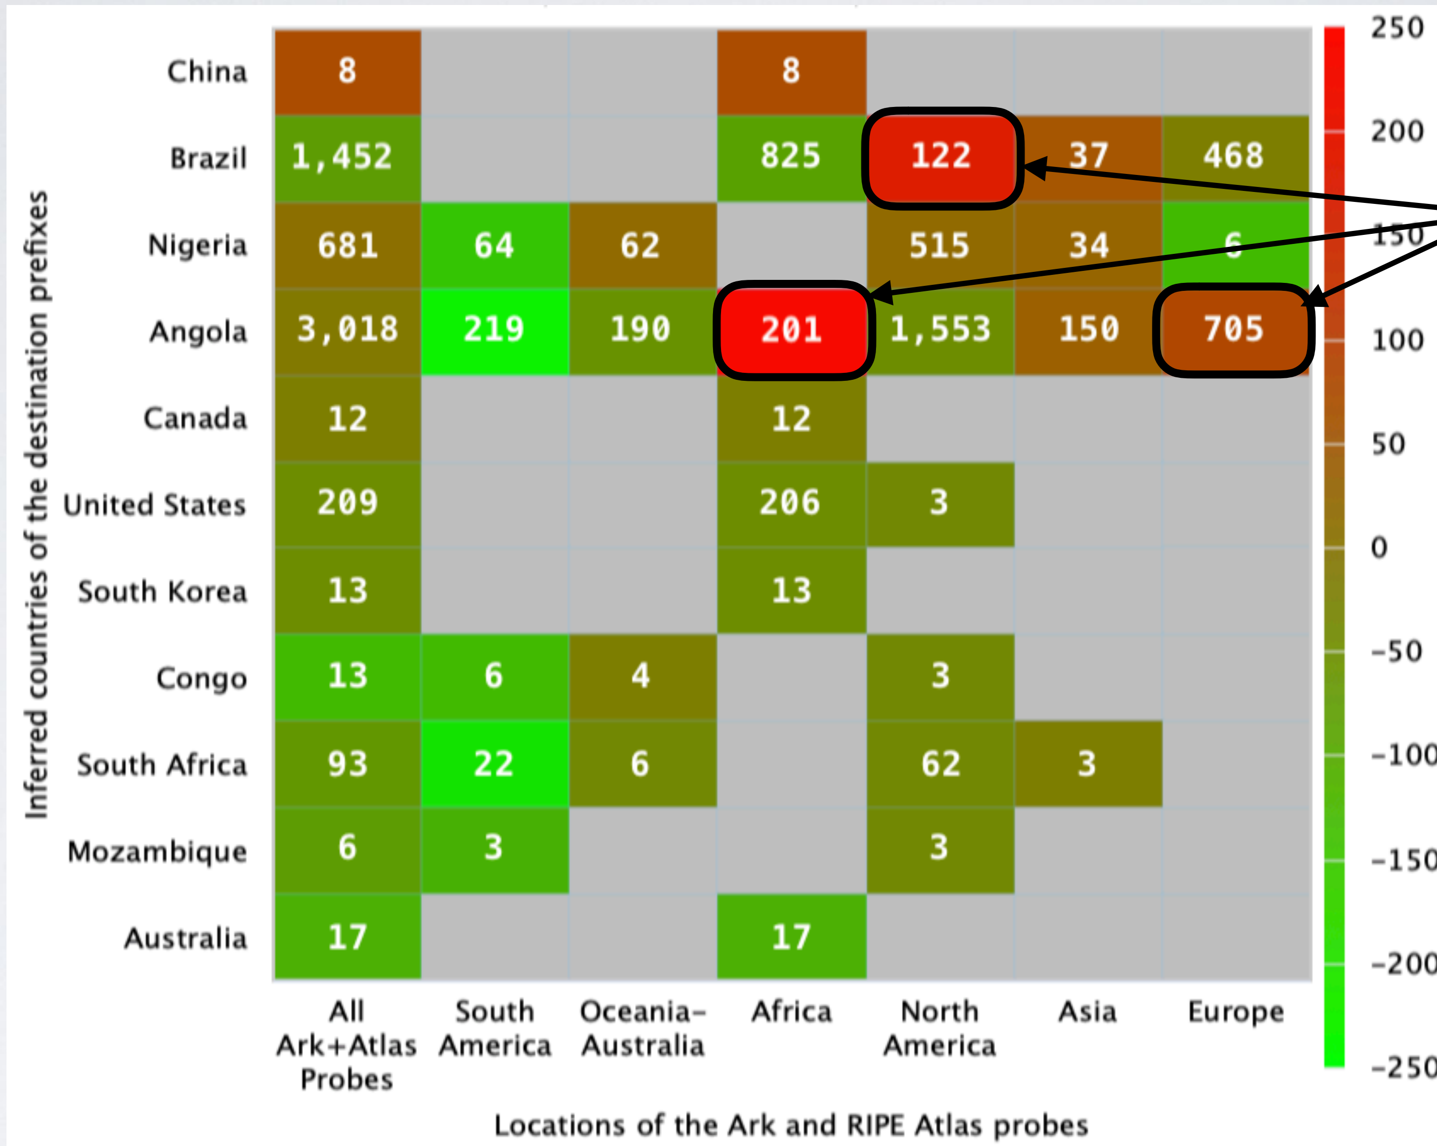

Unpredicted and unreported performance degradation

Packets routed through SACS for pairs from Africa to Angola or Europe to Angola, leading to latency increase.

Comparing Transit Structure

- **A higher centrality of an AS post-event indicates increased transit importance.**

- **Method**:

  - Use bdrmapIT [7] to infer AS paths from IP paths

  - AS's centrality [8]: percentage of observed <s,d> pairs for which the AS path with the minimum observed RTT contains the considered AS.

[7] Marder et al., Pushing the Boundaries with bdrmapIT: Mapping Router Ownership at Internet Scale. In ACM IMC, 2018
[8] Fanou et al, On the Diversity of Interdomain Routing in Africa, In PAM, 2015

Effects on Transit Structure (1)

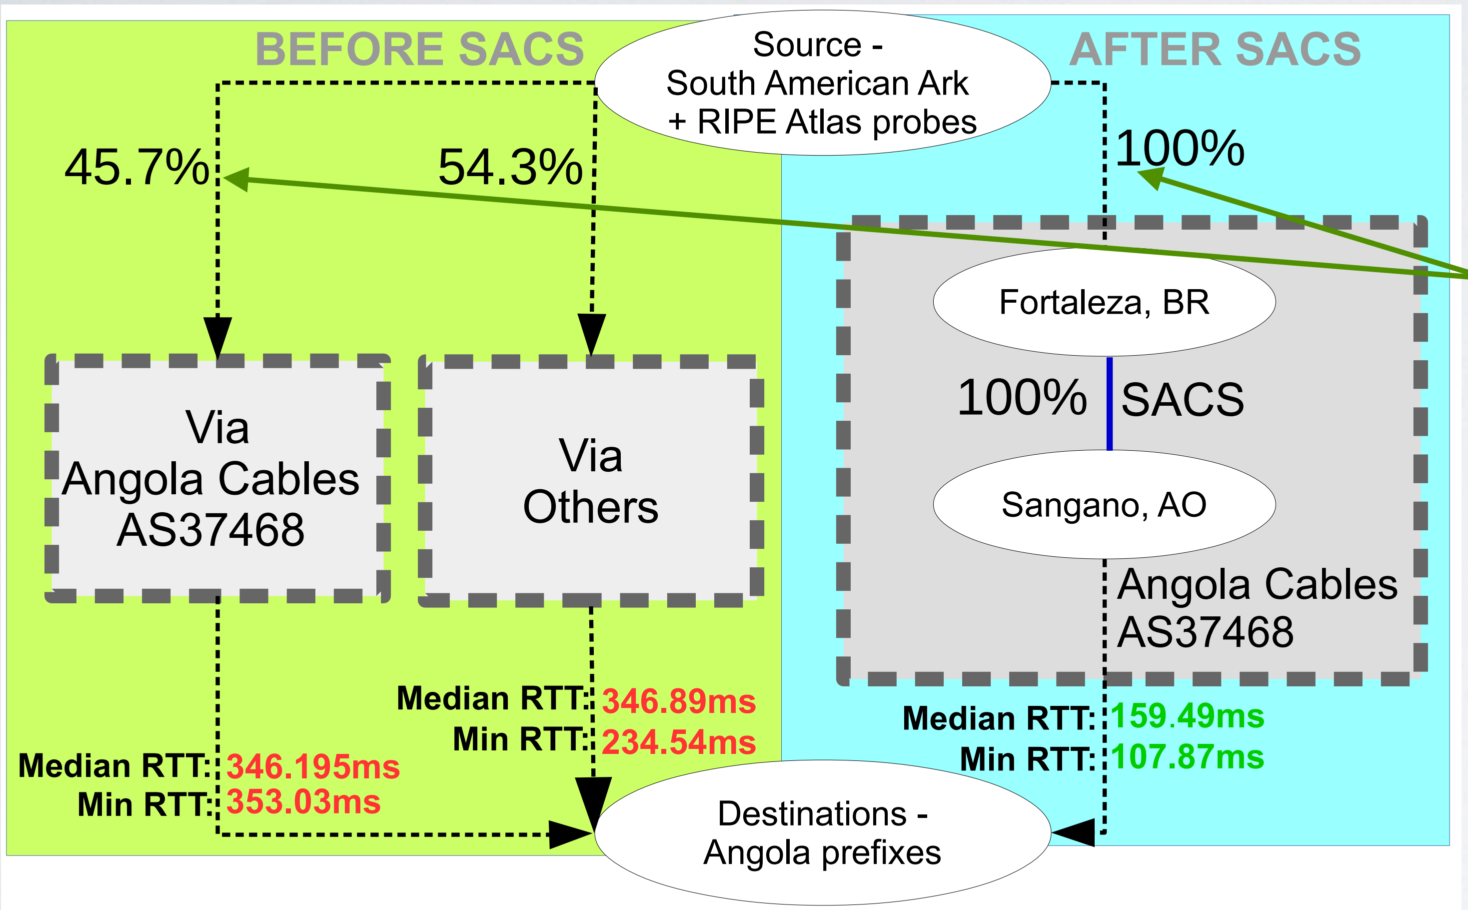

	Before	After
AS-centrality (AS37468)	46%	100%
Median RTT of IP paths via AS37468	346ms	159ms
Min RTT of IP paths via AS37468	353ms	108ms

Partial AS paths **from South America to Angola** (Observed RTT improvement)

AS-centrality of AS37468 increased from 46% to 100%

caida

Effects on Transit Structure (1)

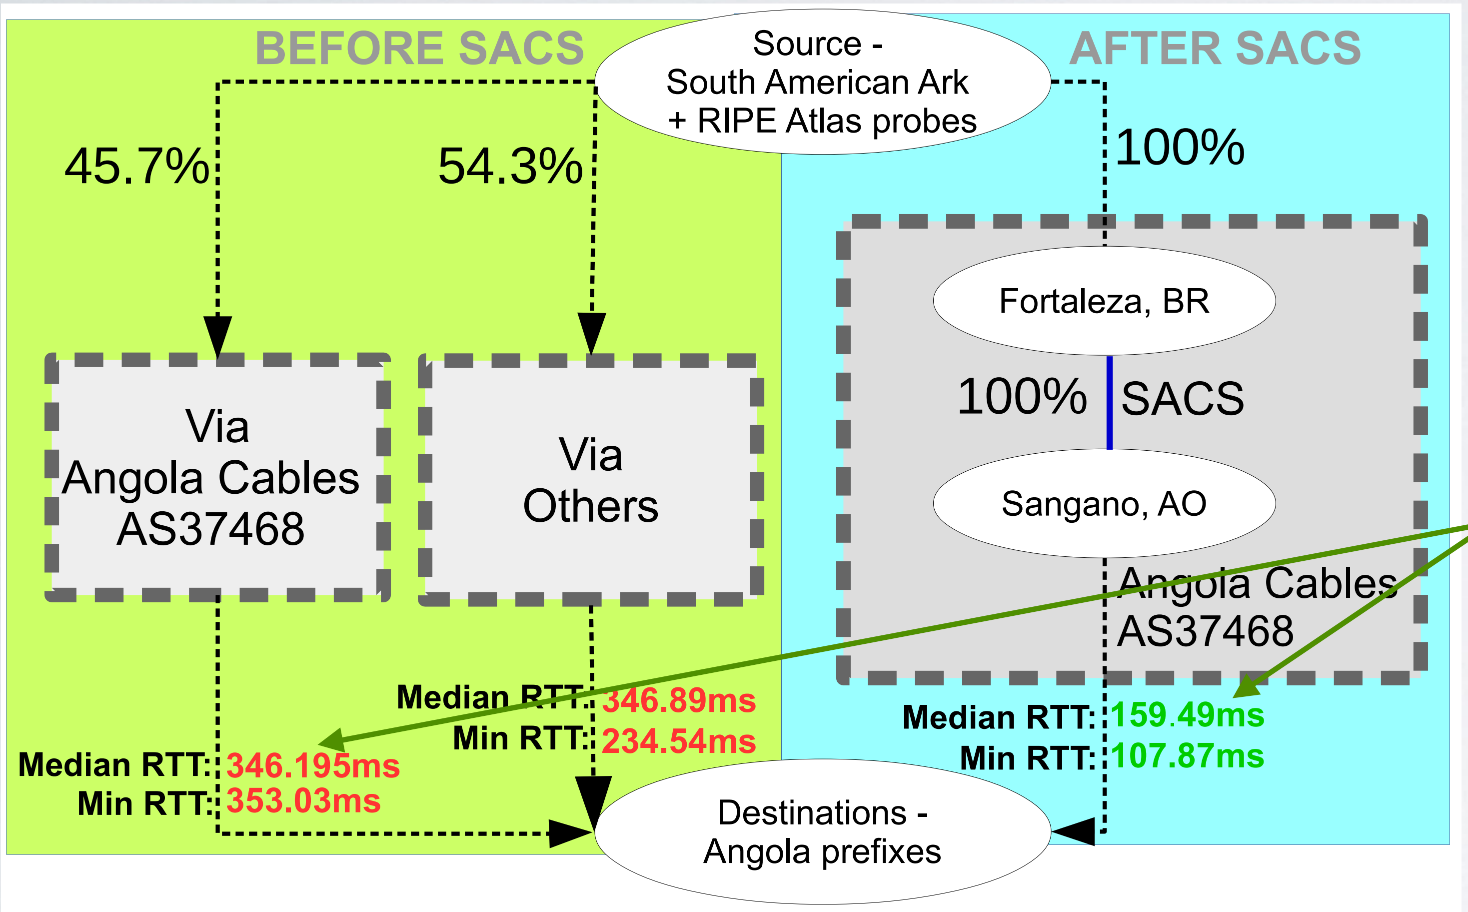

Partial AS paths **from South America to Angola** (Observed RTT improvement)

	Before	After
AS-centrality (AS37468)	46%	100%
Median RTT of IP paths via AS37468	346ms	159ms
Min RTT of IP paths via AS37468	353ms	108ms

Median RTT on IP paths crossing AS37468 decreased from 346ms to 159ms

Effects on Transit Structure (2)

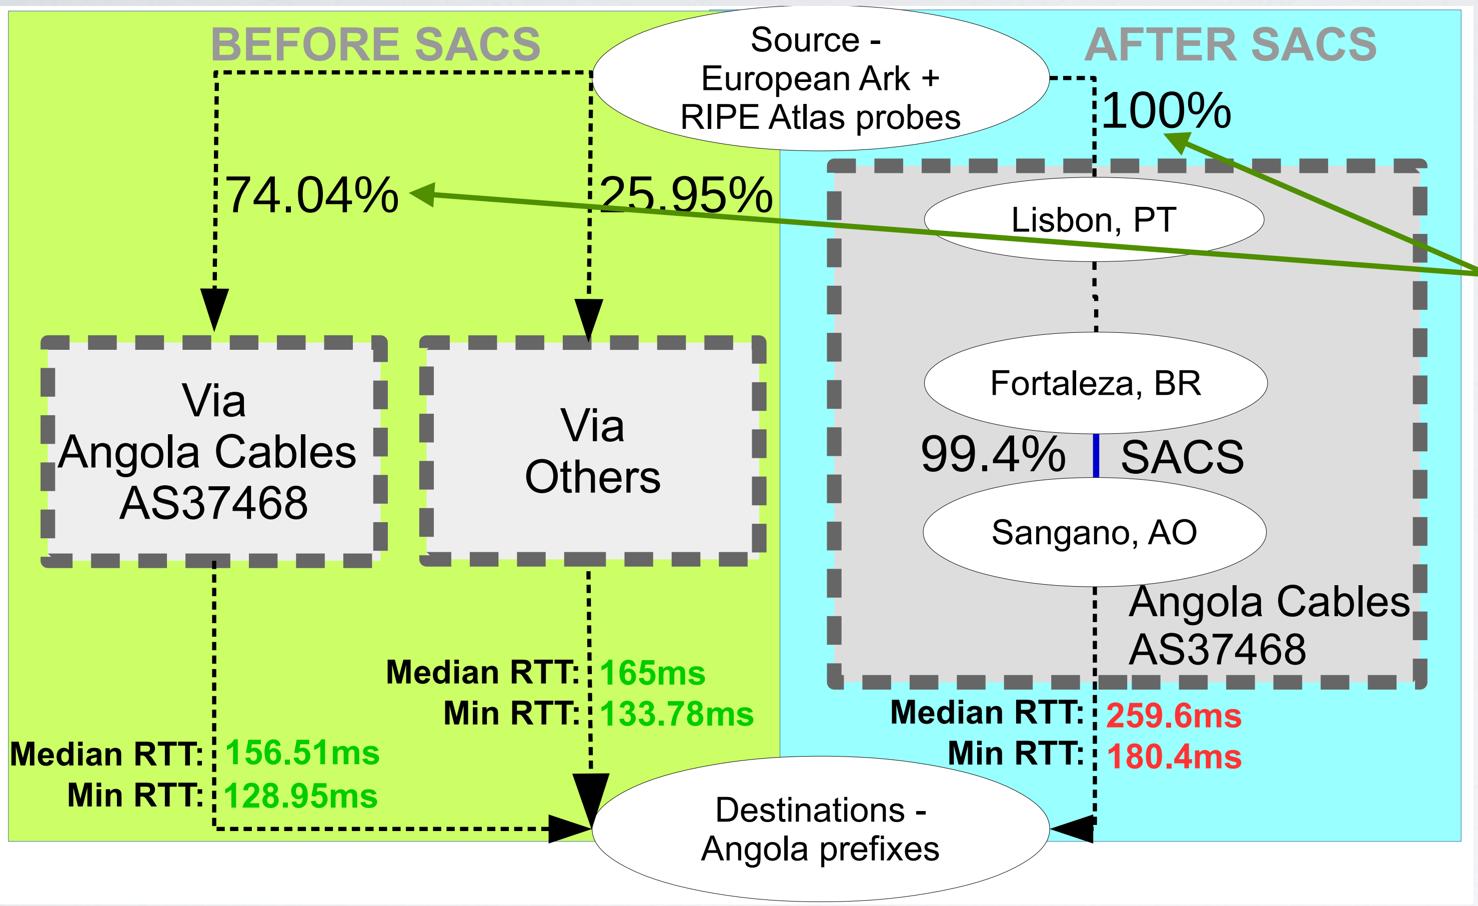

	Before	After
AS-centrality (AS37468)	74%	100%
Median RTT of IP paths via AS37468	156ms	259ms
Min RTT of IP paths via AS37468	128ms	180ms

Partial AS paths **from Europe to Angola** (Observed RTT degradation)

AS-centrality of AS37468 increased from 74% to 100%

Effects on Transit Structure (2)

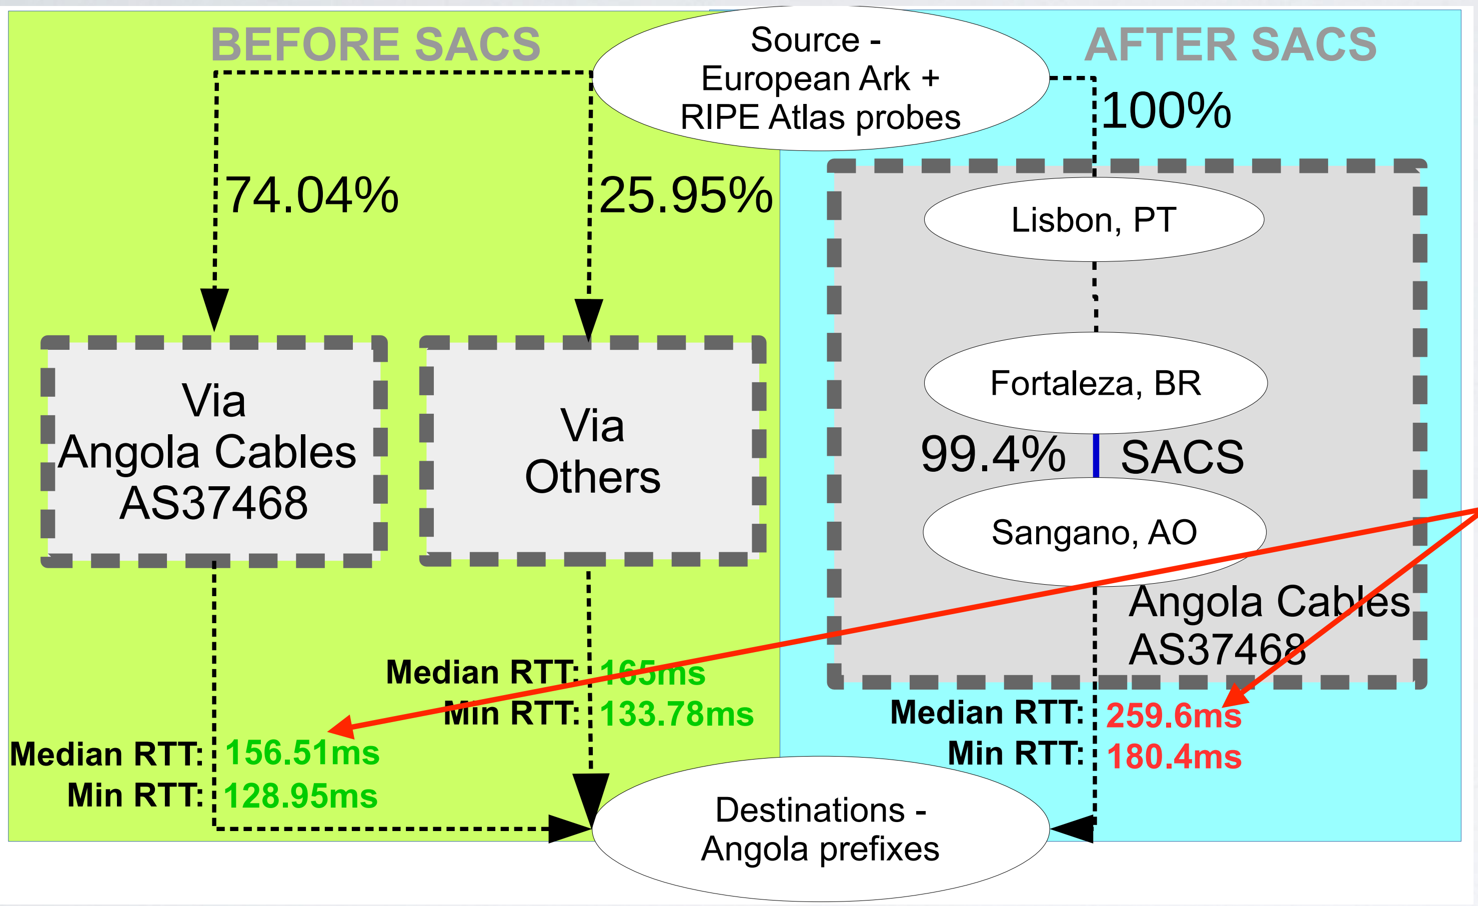

Partial AS paths **from Europe to Angola** (Observed RTT degradation)

	Before	After
AS-centrality (AS37468)	74%	100%
Median RTT of IP paths via AS37468	156ms	259ms
Min RTT of IP paths via AS37468	128ms	180ms

Median RTT on IP paths crossing AS37468 increased from 156ms to 259ms

Examples of Paths Changes

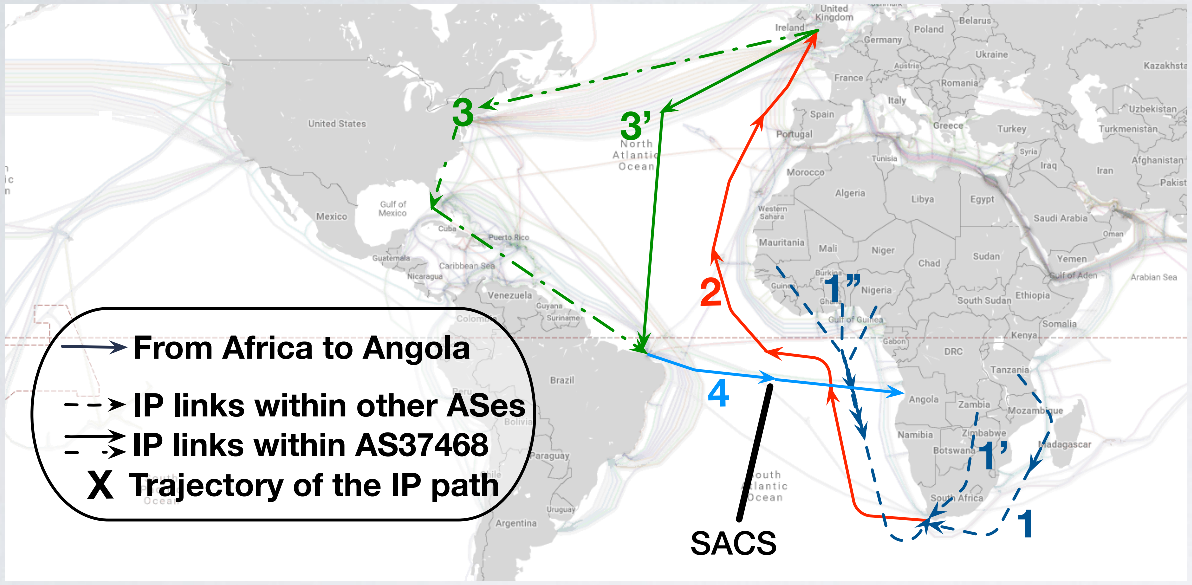

(A) Africa to Angola (at least 55% <s,d>)

(B) North America to Brazil (25% <s,d>)

(C) Europe to Angola (99.3% <s,d>)

(D) South America to Angola (100%)

Most IP paths from Africa to Angola crossing SACS went through South Africa, Europe, North America, Fortaleza before reaching their destination

Examples of Paths Changes

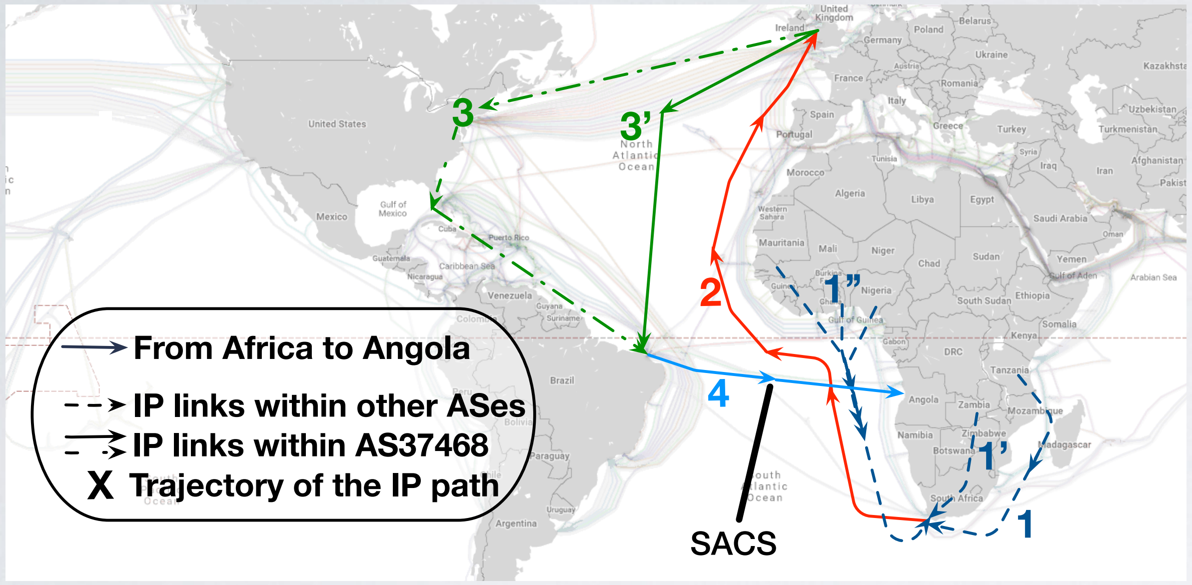

(A) Africa to Angola (at least 55% <s,d>)

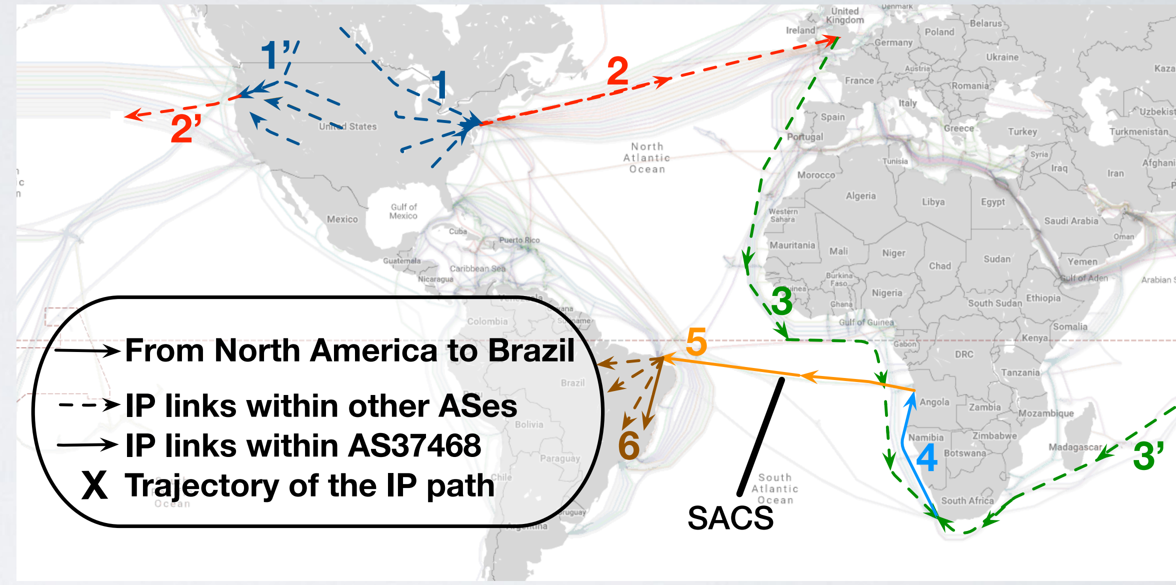

(B) North America to Brazil (25% <s,d>)

(C) Europe to Angola (99.3% <s,d>)

(D) South America to Angola (100%)

Most IP paths from North America to Brazil crossing SACS went through Europe/Asia, South Africa, Angola before reaching Fortaleza

Examples of Paths Changes

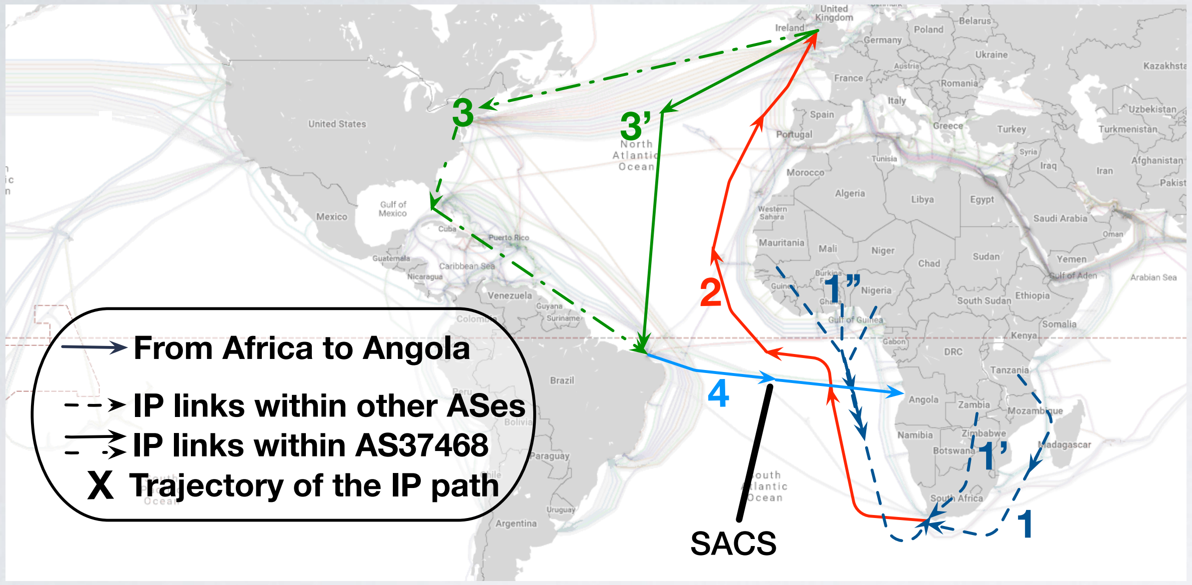

(A) Africa to Angola (at least 55% <s,d>)

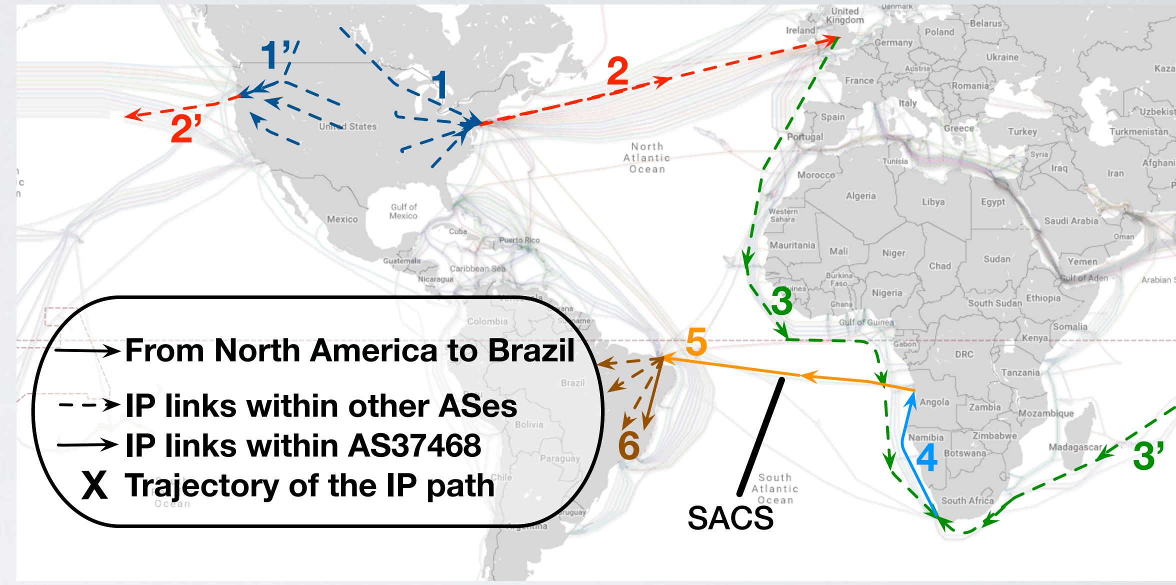

(B) North America to Brazil (25% <s,d>)

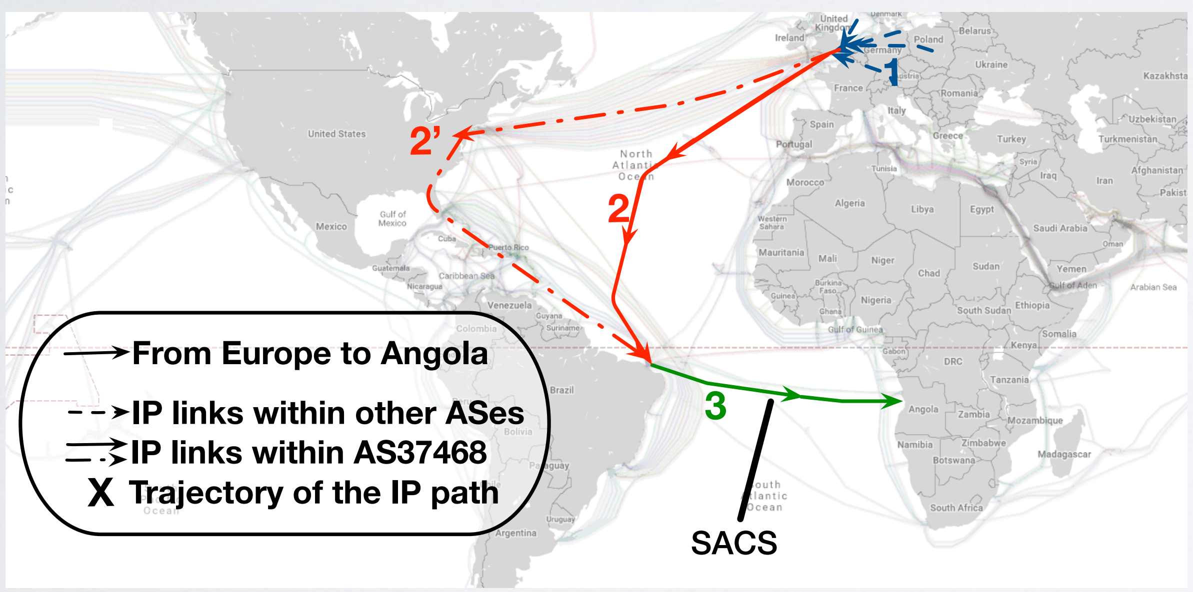

(C) Europe to Angola (99.3% <s,d>)

(D) South America to Angola (100%)

Most observed IP paths from Europe to Angola crossing SACS went through North America before reaching their destination

Examples of Paths Changes

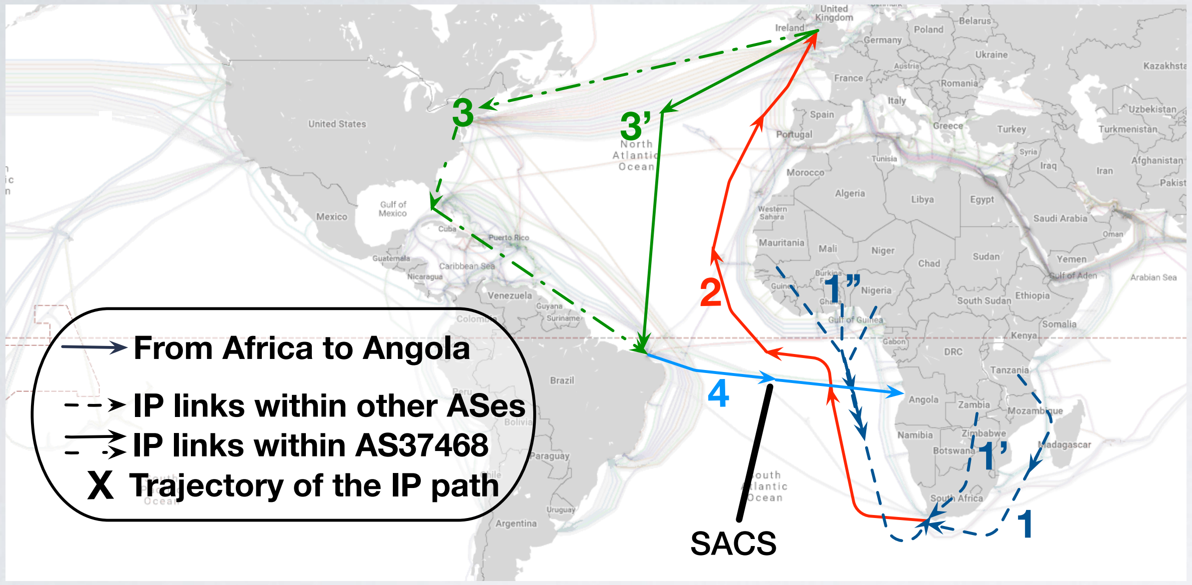

(A) Africa to Angola (at least 55% <s,d>)

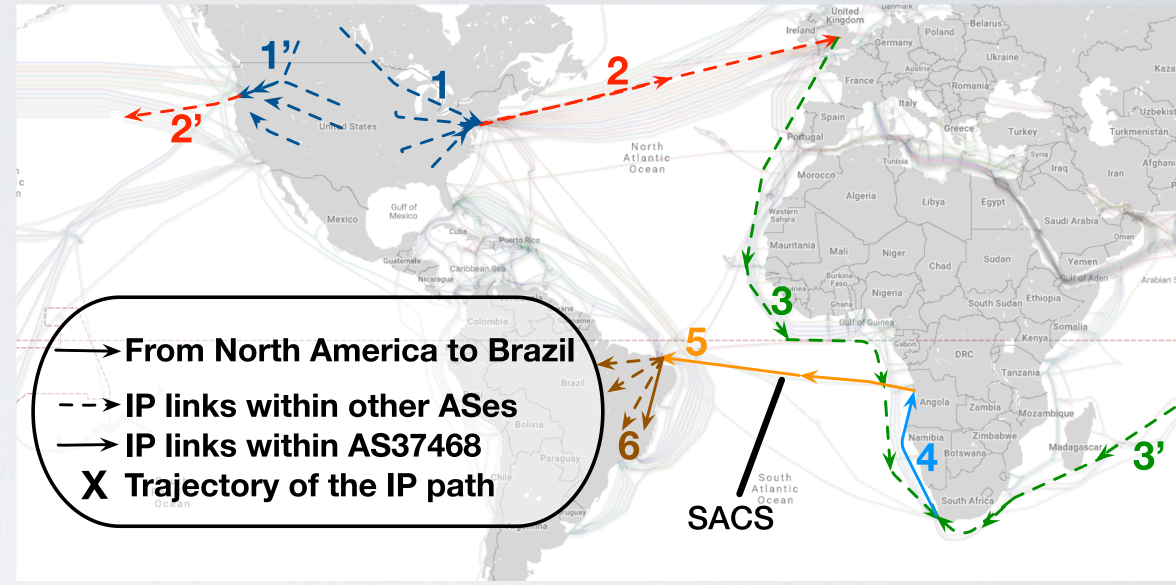

(B) North America to Brazil (25% <s,d>)

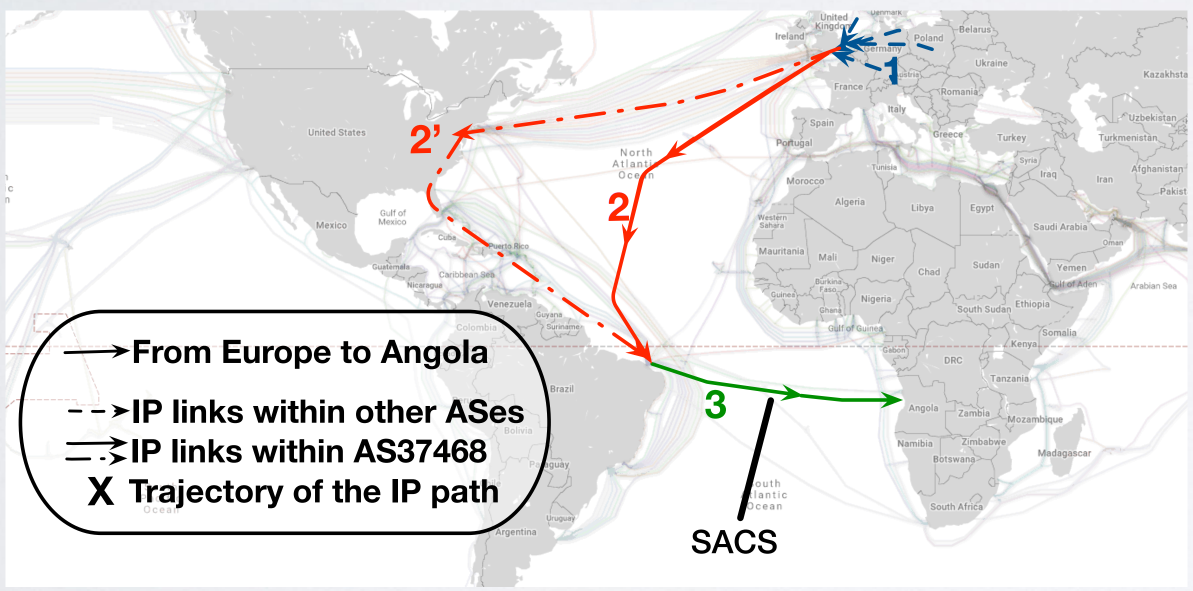

(C) Europe to Angola (99.3% <s,d>)

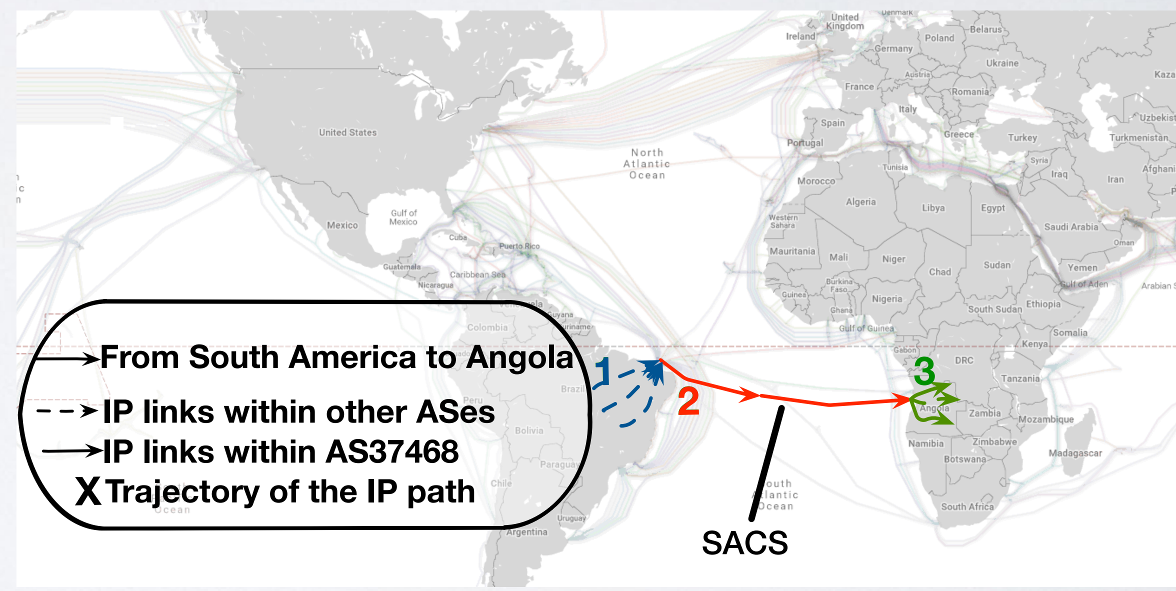

(D) South America to Angola (100%)

All observed IP paths from South America to Angola
directly crossed SACS to reach their destination

Comparing AS path lengths

- **Analyze length of AS paths between source AS/destination prefix pairs crossing cable operator's network in BGP pre & post-event**
- **Method**:
  - Consider AS paths collected during the first 5 days of the month pre & post-event
  - Filter out all $ASpaths_{<s,d>}$ crossing AS37468 operating the cable post-event
  - Extract all $ASpaths'_{<s,d>}$ between the same <s,d> pairs pre-event

Increased AS paths lengths

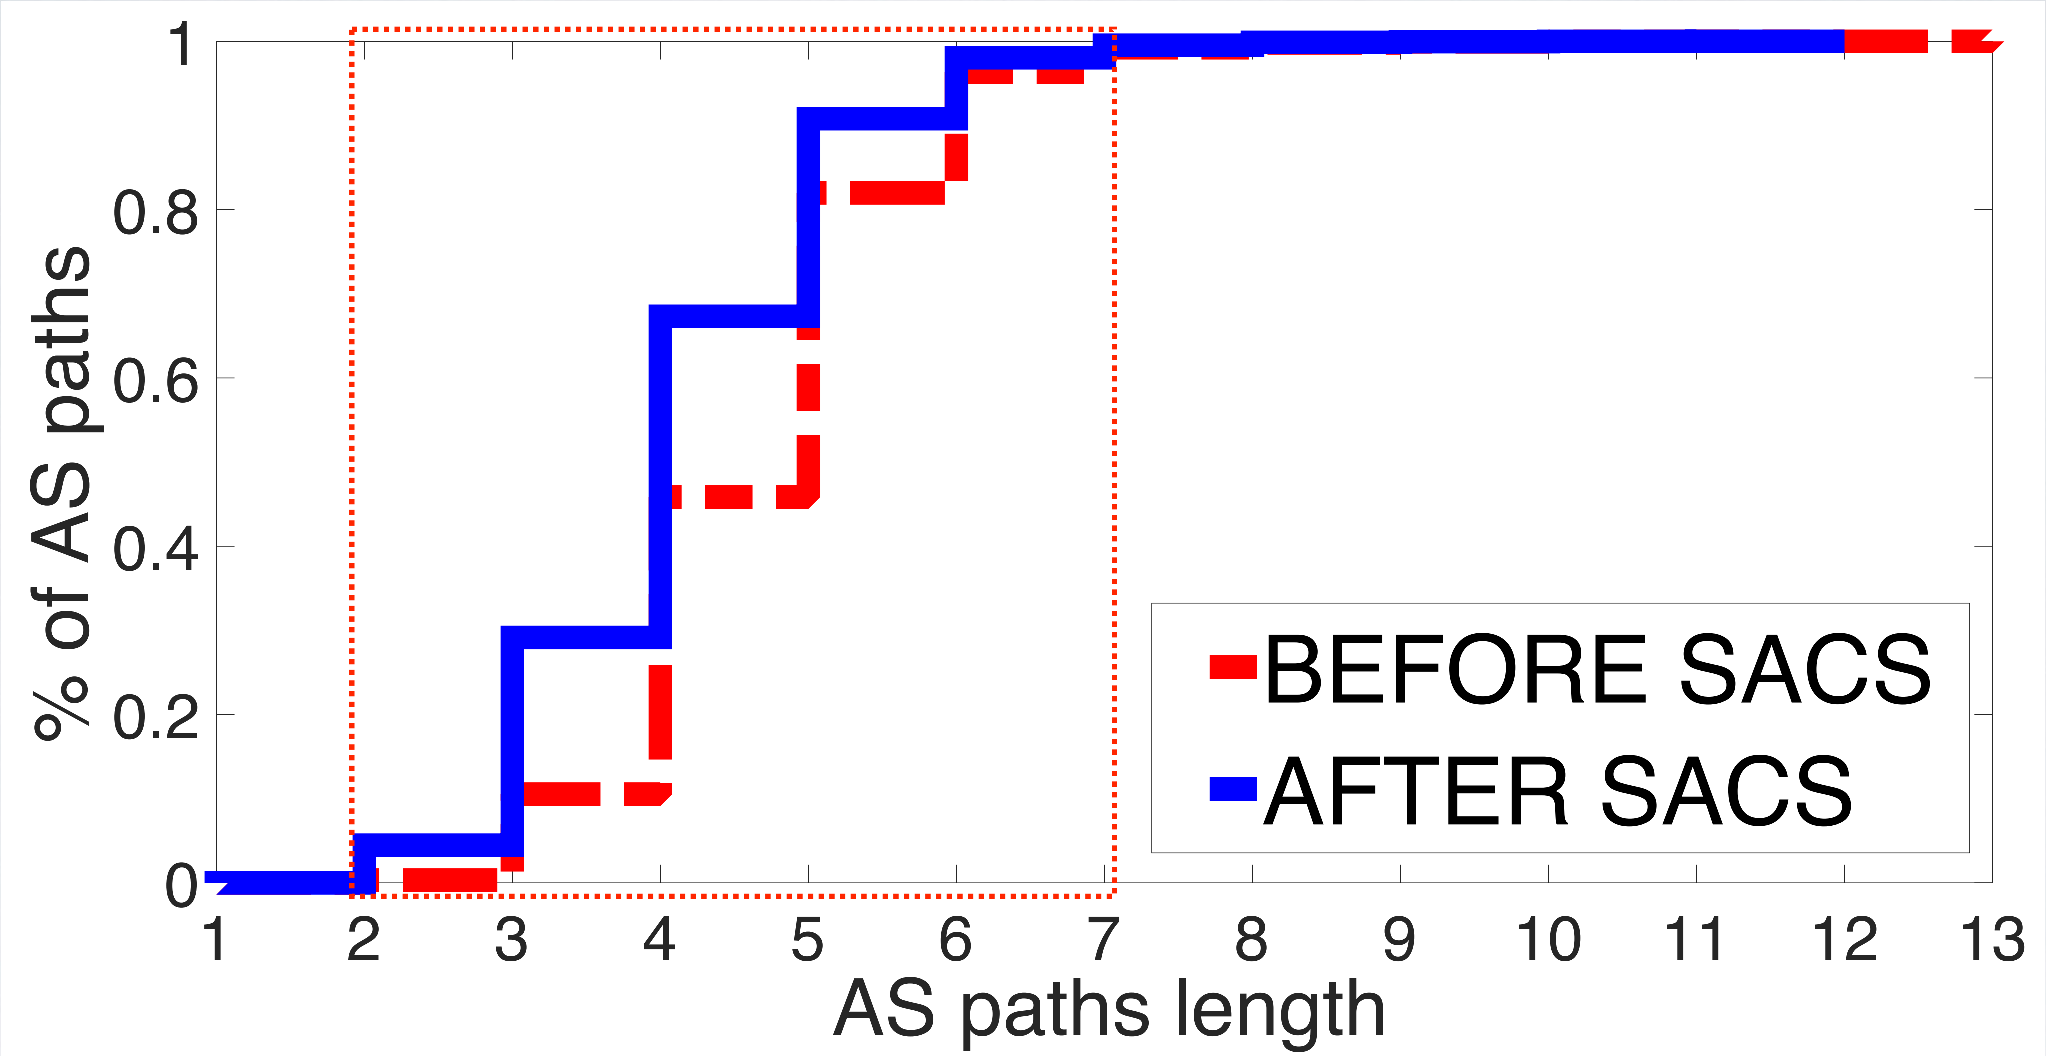

Distribution of the length of AS paths between same source AS/destination prefix pairs served via AC (AS37468) pre&post SACS, showing the increase of paths of length 2-7

Validation with the ISP

- AC identified cases where the suboptimality happened **outside of their network**
- AC highlighted that **internal link failures could account for the performance degradations**
- AC did not validate suboptimal cases within their network, but **most observed IP paths switched to optimal ones** after our conversation

Potential Causes of Suboptimal Routing

- IGP/EGP misconfigurations (typos, errors, etc.)
- slow IGP or EGP convergence
- lack of traffic engineering after event in neighbouring AS
- persistent lack of peering among local ASes in Africa
- frequent use of default routes via international transit providers in developing regions

Caveat

- AC noted most traffic crossing SACS through its network goes from **South-America to Angola** or **South Africa to Brazil,** the regions pairs that experienced substantial performance improvements

- No complaints from customers, hinting low amount of traffic carried, but **suboptimal routes lasted 3.5 months**

Contributions

- **Reproducible method** to investigate impact of a cable deployment on macroscopic topology and performance

- **Application of our methodology to the case of SACS**, the first South-Atlantic cable from South America to Africa:

  - Discovered RTT decrease from **Africa to Brazil** was roughly **a third** of that from **South America to Angola**

  - Discovered surprising performance degradations to/from some regions

Further Contributions

- **Suggestions to avoid suboptimal routing post-activation of cables in the future:**
  - Inform BGP neighbours to allow time for changes
  - Ensure optimal iBGP configs post-activation
  - Use measurements platforms to verify path optimality
- **Access to Code & Data:**
  - https://github.com/CAIDA/submarine-cable-impact-analysis-public

caida

Help build the Ark measurement network by hosting a VP!

We are always looking for volunteers to host VPs!

Contact us:

ark-info@caida.org

or

roderick@caida.org

Thank you. Questions?

Backup Slides

Related Work

- **[4,5] investigated effects of disruptions to existing cables**
  - In 2010, [4] observed significant congestion on paths from Hong Kong to EU websites caused by submarine cable fault
  - In 2011, [5] studied the impact of SEA-ME-WE 4 submarine cable fault on end-to-end path quality in terms of the delay and packet loss rates

- **[6] underlined the need to undertake research to characterize the global submarine cable network**

- No scientific study of the macroscopic effects of a cable launch on AS topology and end-to-end performance

[4] Chang et al, Could Ash Cloud or Deep-Sea Current Overwhelm the Internet? In USENIX HotDep, 2010
[5] Chan et al. Non-cooperative diagnosis of submarine cable faults. In PAM, 2011
[6] Bischof et al, Submarine Cables and Internet Resiliency, 2018

Example of AS Topology Pre-cable Launch

- AS8 is yet to activate a cable between countries A and B
- AS8 is present at an IXP in both countries (not necessary)
- Red line: AS3 => AS4
- Blue line: AS4 => AS3

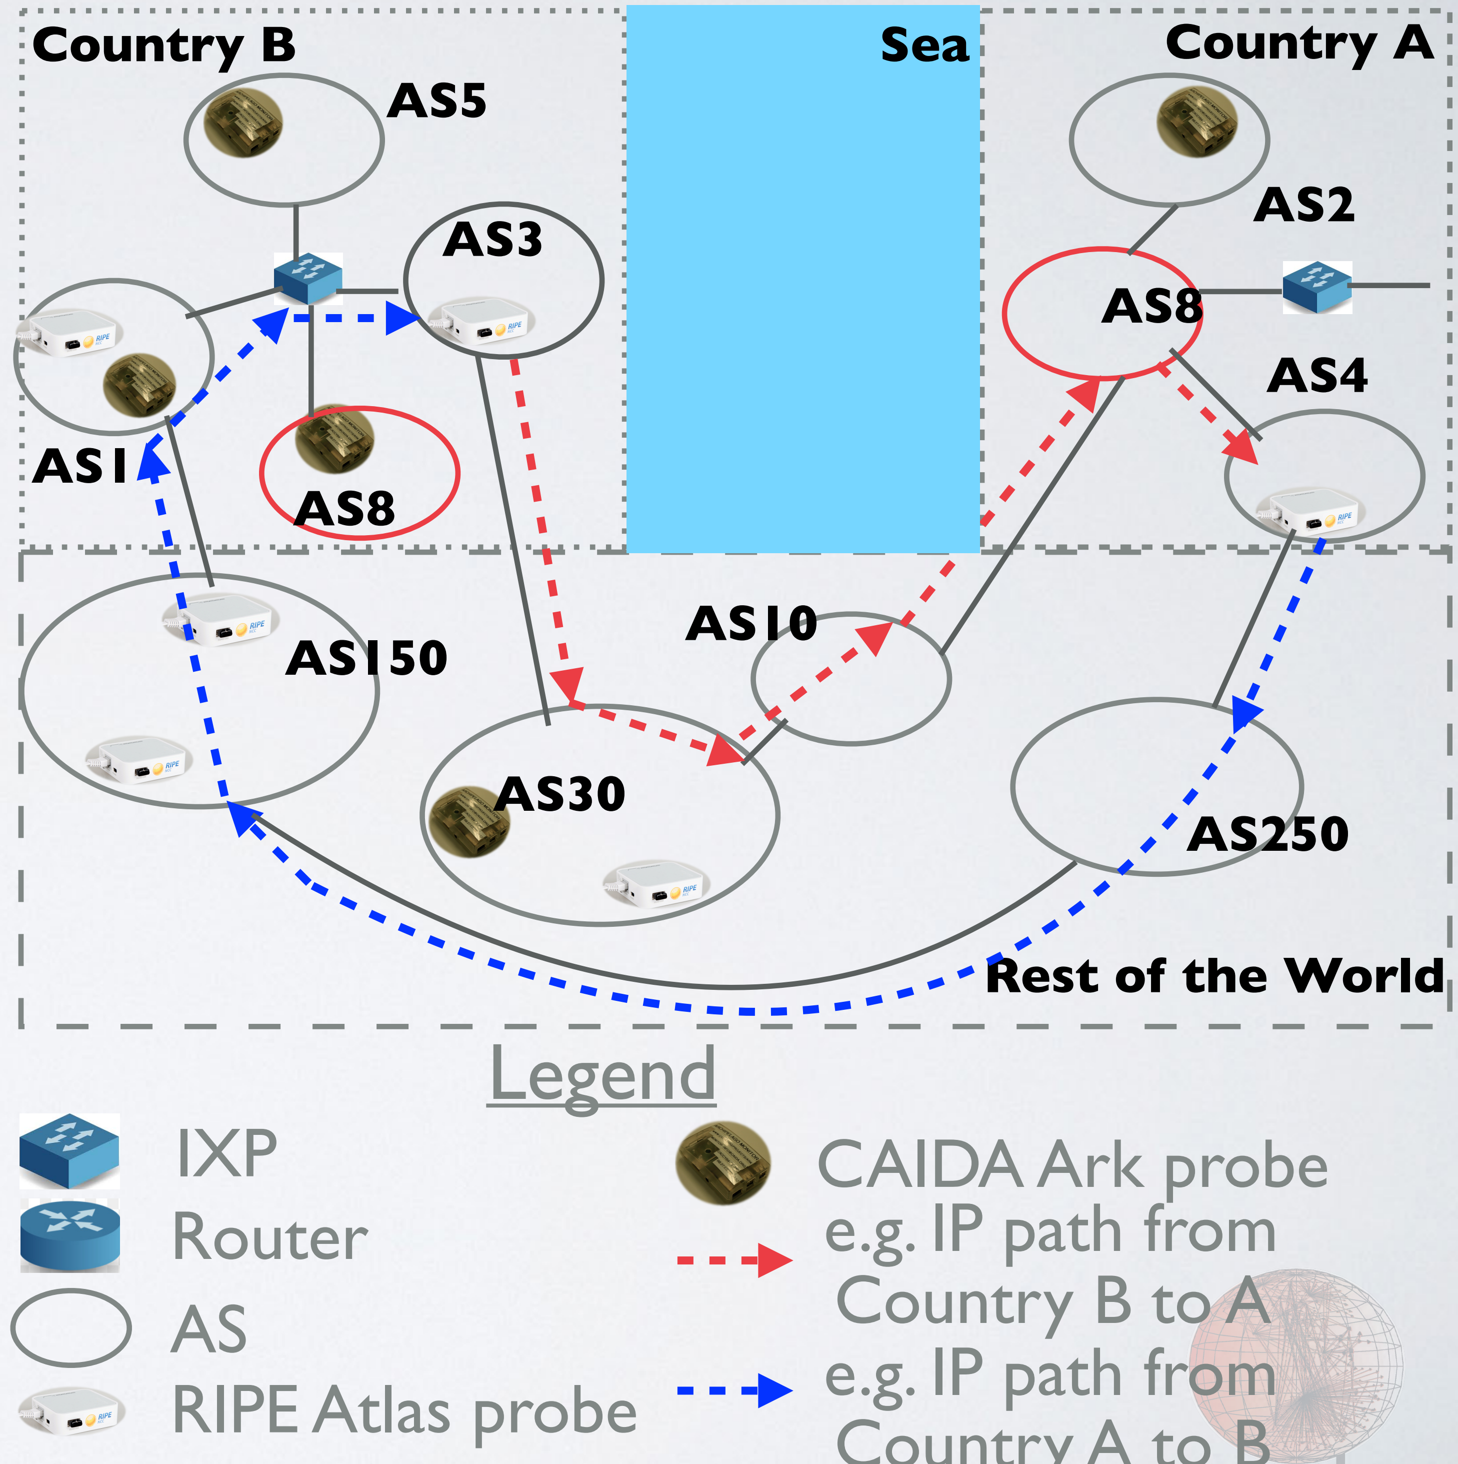

Step1. Collect candidate IP paths that could have crossed the cable

- Collect detailed information about the cable launch
- Select VPs in connected countries and near the 2 cables endpoints (eg: VPs in **AS3, AS8, *AS2* & *AS4***)
- Conduct traceroutes in both directions to obtain **candidate IP paths**

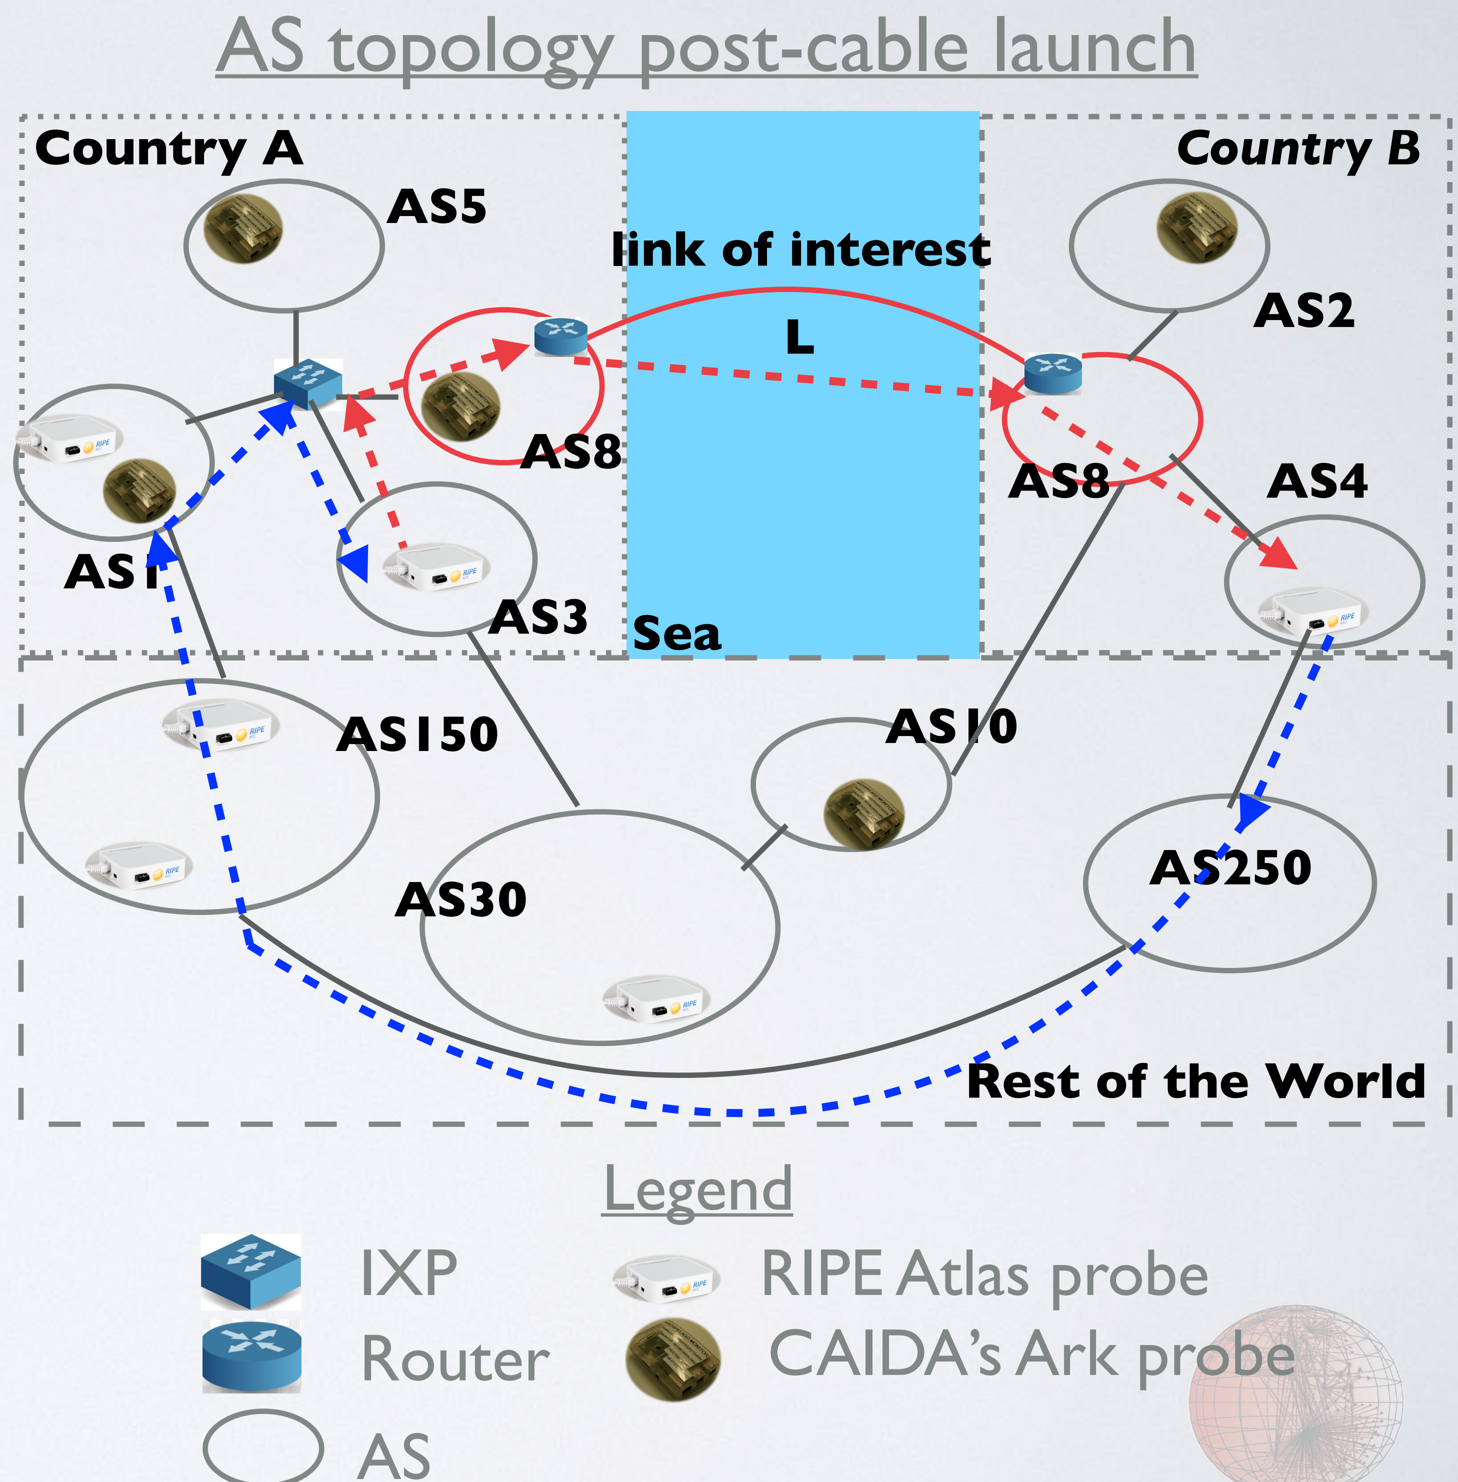

caida

Step 2. Identify Routers Interfaces on both Sides of the Cable (1)

- In our example, AS8 is the AS operating the new cable long of L km between countries A & B.

- **Combine 2 approaches to identify traces crossing the link of interest and interfaces on both sides (1):**

  A. **RTT-threshold based on speed-of-light constraints**

  - Use bdrmapIT to map IP in **candidate IP paths** to ASes

  - Look for an RTT difference gap between consecutive hops within AS8 of $t = \frac{2L}{\frac{2c}{3}} = \frac{3L}{c}$, with $\frac{2c}{3}$ the speed of light traveling fibre optics

  - **Those IP hops likely match the cable landing sites.**

caida

Step 2. Identify Routers Interfaces on both Sides of the Cable (2)

- **Combine 2 approaches to identify traces crossing the link of interest and interfaces on both sides (2)**:

  B. IP geolocation:

  - Use IP geolocation databases and hostname-based geolocation to locate IPs on both sides of the threshold RTTs
  - Validate with RTT measurements from VPs in inferred countries to those IPs and their adjacent IPs.
  - If inferred geolocations are countries A and B, resolve the router aliases of the selected IPs using CAIDA's MIDAR, Vela alias, and ITDK (the obtained lists Ra and Rb are called **link IPs**)

Step 3. Search for Comparable Historical Traceroutes

- Collect traceroutes post-cable launch traversing link IPs:
  - **Ark** and **Atlas** randomly probe within BGP prefixes: We denote s each source IP address and d the longest prefix match for each traceroute destination IP in BGP
  - We fetch all traceroutes $T_{<s,d>}$ post-cable containing $Ra \rightarrow Rb$ or $Rb \rightarrow Ra$ (separated by RTT difference t)
- Search for pre-event traceroutes $T'_{<s,d>}$ for the same source IPs/destination prefix <s,d> pairs for comparison

Step 3: Fetch Matching Traceroutes

- Use CAIDA's Ark and RIPE Atlas topology data
- Fetch traces with $Ra \rightarrow Rb$ or $Rb \rightarrow Ra$ post-SACS
- Collect traces between the same <s,d> pairs pre-SACS, splitting these historical data into 2 sets:
  - mid-Sept 2018 - late Jan 2019: ARK-AFTER/RIPE-AFTER
  - Jan 2018-mid - Sept 2018: ARK-BEFORE/RIPE-BEFORE
- **Cleaned up dataset:** Atlas traceroutes (823 <s,d> pairs) & Ark traceroutes (6,778 <s,d> pairs)

Step 4. Annotate Collected Paths

- **For every hop in $T_{<s,d>}$ & $T'_{<s,d>}$:**
  - We resolve hostname (zdns & qr) & ASN (bdrmapIT[7])
  - We compute difference in RTT with the previous hop
  - We check if it belongs to an IXP prefix
- **We group $T_{<s,d>}$ & $T'_{<s,d>}$ and their annotation per week using their timestamp of execution**

[7] Marder et al., Pushing the Boundaries with bdrmapIT: Mapping Router Ownership at Internet Scale. In ACM IMC, 2018

AS-centrality Computation

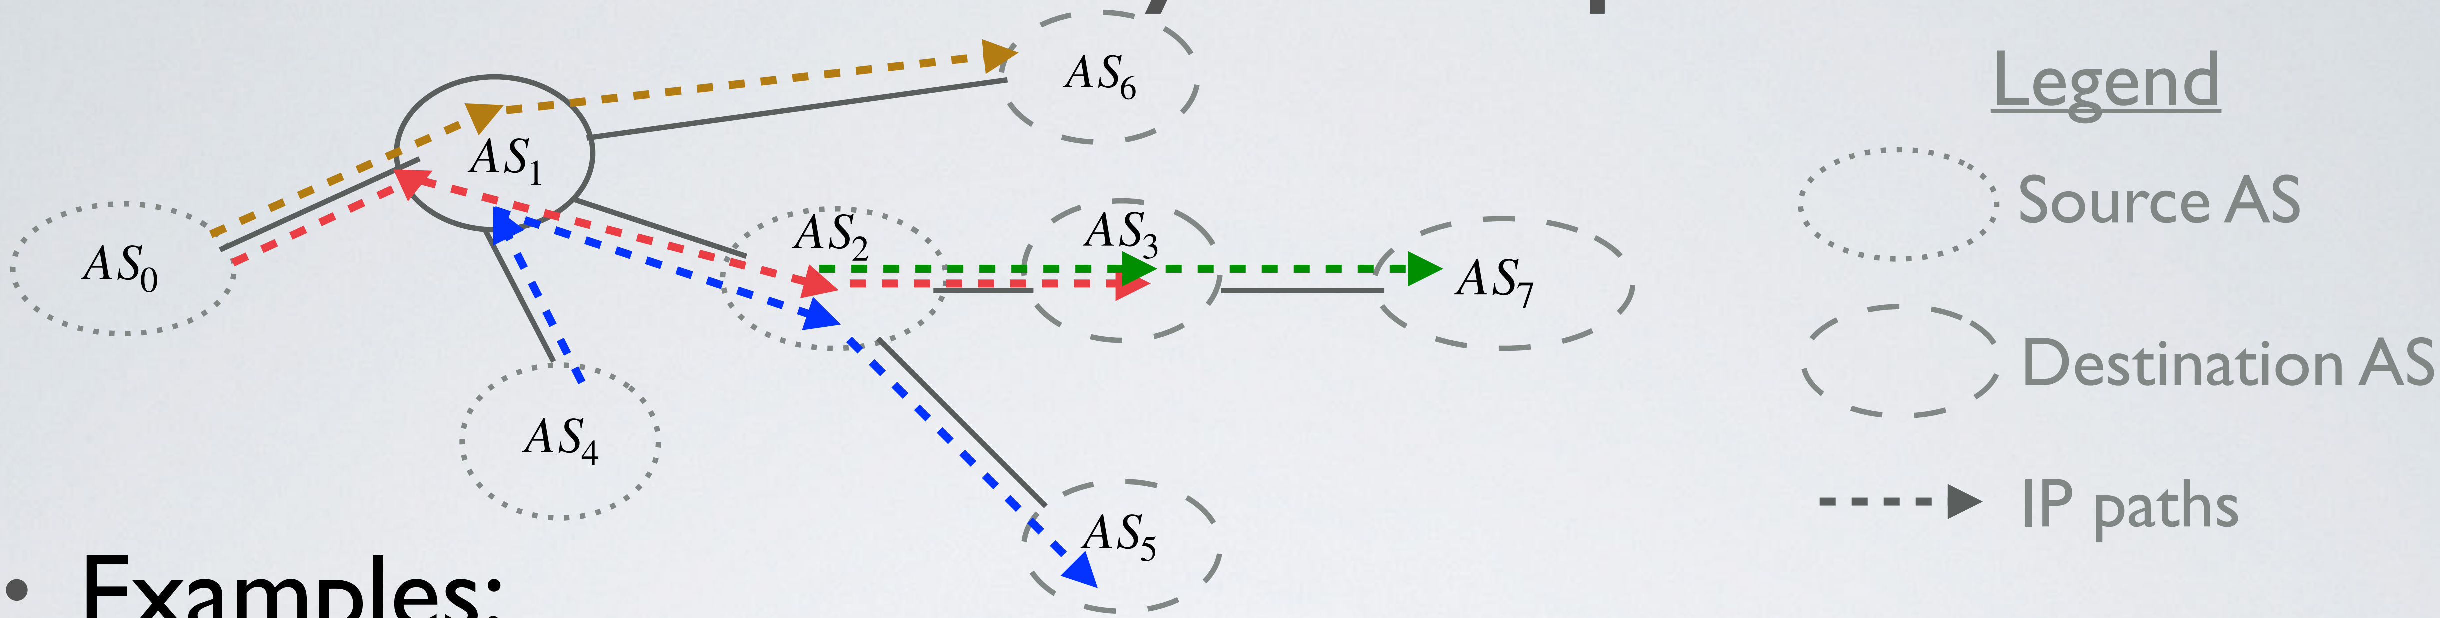

- **Examples:**
  - Suppose 4 <s,d> pairs were measured from region A to country B and that the IP paths with the min RTTs were crossing the ASes above
  - after mapping IPs to ASes, we obtain the AS paths between: $AS_0 - AS_6, AS_4 - AS_5, AS_2 - AS_7, AS_0 - AS_3$
  - AS-c(AS_1) $= \frac{3}{4}$, AS-c(AS_2) $= \frac{1}{4}$, AS-c(AS_3) $= \frac{1}{4}$, AS-c(AS_6) $= 0$

Effects on Transit Structure

Top 3 transit ASes serving <s, d> pairs from continents to destination countries. The categories in which we noticed Suboptimal routing are highlighted in italic.

<table>
<tr><th rowspan="2">Category (# $< s, d >$)</th><th colspan="3">Before</th><th colspan="3">After</th></tr>
<tr><th>CC</th><th>AS-centrality</th><th colspan="2">Transit AS</th><th>AS-centrality</th><th>CC</th></tr>
<tr><td rowspan="4">From Africa to Angola (201)</td><td>AO</td><td>66.7%</td><td colspan="2">Angola Cables (AS37468)</td><td>90.1%</td><td>AO</td></tr>
<tr><td>ZA</td><td>32.3%</td><td colspan="2">Internet Solutions (AS3741)</td><td>22.4%</td><td>ZA</td></tr>
<tr><td>BG</td><td>20.9%</td><td>Sofia Connect (AS47872)</td><td>WIOCC-AS (AS37662)</td><td>16.4%</td><td>MU</td></tr>
<tr><td></td><td></td><td></td><td>IPPLANET (AS12491)</td><td>16.4%</td><td>IL</td></tr>
<tr><td rowspan="3">From North America to Brazil (122)</td><td>US</td><td>44.4%</td><td>ATT-Internet4 (AS7018)</td><td>Angola Cables (AS37468)</td><td>100%</td><td>AO</td></tr>
<tr><td>BR</td><td>30.1%</td><td>NipBr (AS27693)</td><td>Chinanet-B. (AS4134)</td><td>60.2%</td><td>CN</td></tr>
<tr><td>US</td><td>23%</td><td>Nitel (AS53828)</td><td>Abilene (AS11537)</td><td>58.3%</td><td>US</td></tr>
<tr><td rowspan="3">From Europe to Angola (705)</td><td>AO</td><td>62.9%</td><td colspan="2">Angola Cables (AS37468)</td><td>78.1%</td><td>AO</td></tr>
<tr><td>BG</td><td>18.6%</td><td>Sofia-Connect (AS47872)</td><td>Telianet (AS1299)</td><td>17.6%</td><td>EU</td></tr>
<tr><td>EU</td><td>14.2%</td><td>Telianet (AS1299)</td><td>TWTC (AS4323)</td><td>9.9%</td><td>US</td></tr>
<tr><td rowspan="3">From Asia to Brazil (141)</td><td>AO</td><td>50.3%</td><td colspan="2">Angola Cables (AS37468)</td><td>90.1%</td><td>AO</td></tr>
<tr><td>US</td><td>28.4%</td><td>TATA (AS6453)</td><td>TWTC (AS4323)</td><td>31.9%</td><td>US</td></tr>
<tr><td>JP</td><td>24.1%</td><td colspan="2">KDDI (AS2516)</td><td>26.2%</td><td>JP</td></tr>
<tr><td rowspan="3">From South America to Angola (212)</td><td>AO</td><td>45.7%</td><td colspan="2">Angola Cables (AS37468)</td><td>96.2%</td><td>AO</td></tr>
<tr><td>BR</td><td>36.8%</td><td>Terremark do Brasil (AS28625)</td><td>Cilnet (AS28580)</td><td>18.4%</td><td>BR</td></tr>
<tr><td>US</td><td>36.3%</td><td>Cogent (AS174)</td><td>CO.PA.CO. (AS27768)</td><td>11.8%</td><td>PY</td></tr>
</table>

After SACS activation, AC became the top transit AS for observed paths found to suffer or not from suboptimal routing.

Comparison of Round Trip Time (RTTs) or Delays (RTDs) of IP paths crossing SACS

Sept 2018

Origin	Destination	RTT (ms) Before SACS	RTT (ms) After SACS	Gain RTT (ms)
Miami	**Luanda**	**278**	**128**	**-150**
Fortaleza	London (UK)	172	150	-22
Fortaleza	Lisbon (PT)	193	162	-31
Fortaleza	Madrid (ES)	208	173	-35
Fortaleza	Cape Town (ZA)	337	112	-225
Fortaleza	**Luanda**	**350**	**63**	**-287**
São Paulo	London (UK)	212	190	-22
São Paulo	Frankfurt (DE)	222	200	-22
São Paulo	Johannesburg (ZA)	384	130	-254
São Paulo	Cape Town (ZA)	377	152	-225
São Paulo	**Luanda**	**380**	**109**	**-271**

Variations of RTT Gains on IP paths crossing SACS hint (internal) routing changes

LOW LATENCY IN INTERNATIONAL ROUTES (RTD)

March 2020

Fortaleza - São Paulo	~45ms
Fortaleza - Luanda	~63ms
Fortaleza - Miami	~65ms
Fortaleza - Johannesburgo	~125ms
Fortaleza - Lagos	~115ms
Fortaleza - Lisbon	~163ms
Miami - São Paulo	~104ms
Miami - Luanda	~125ms
São Paulo - Cape Town	~140ms
Miami - Cape Town	~163ms
Miami - Johannesburgo	~180ms
New York - Cape Town	~190ms
Luanda - Cape Town	~37ms
Luanda - Johannesburgo	~54ms
Luanda - Lisbon	~102ms
Luanda - São Paulo	~106ms
Luanda - Miami	~125ms
Luanda - London	~128ms

Source: Angola Cables, https://www.angolacables.co.ao, March 2020