IIJ INNOVATION INSTITUTE

Pinpointing Anomalies in Large-Scale Traceroute Measurements

Romain Fontugne & Kenjiro Cho

November 10, 2016

This presentation

On going research work conducted at IIJ-II

In collaboration with:

- Emile Aben (RIPE NCC)
- Cristel Pelsser (University of Strasbourg)
- Randy Bush (IIJ)

Agenda

Background:

- Understanding Internet health
- Challenges

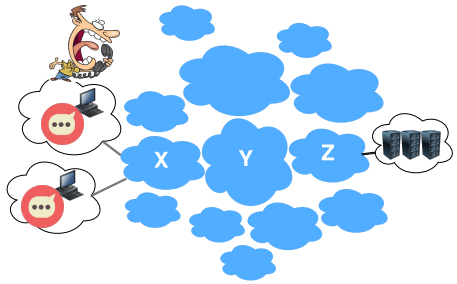

Detect and locate Internet congestion:

- Analysis of traceroutes from RIPE Atlas
- Differential RTT and robust statistics

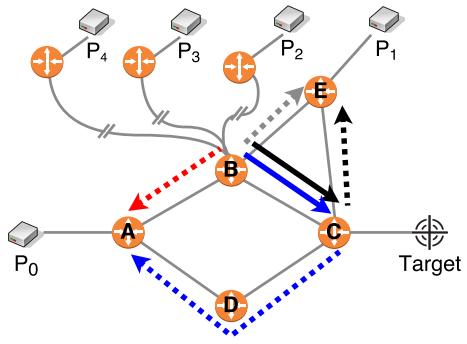

Results:

- Study cases: DDoS attack and BGP leak

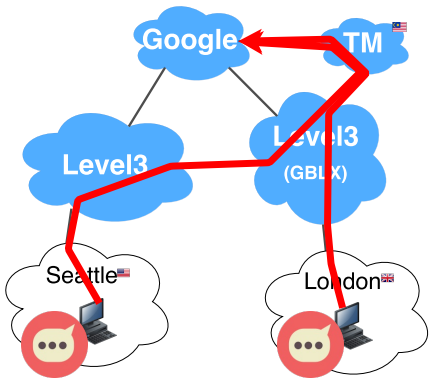

Understanding Internet health?

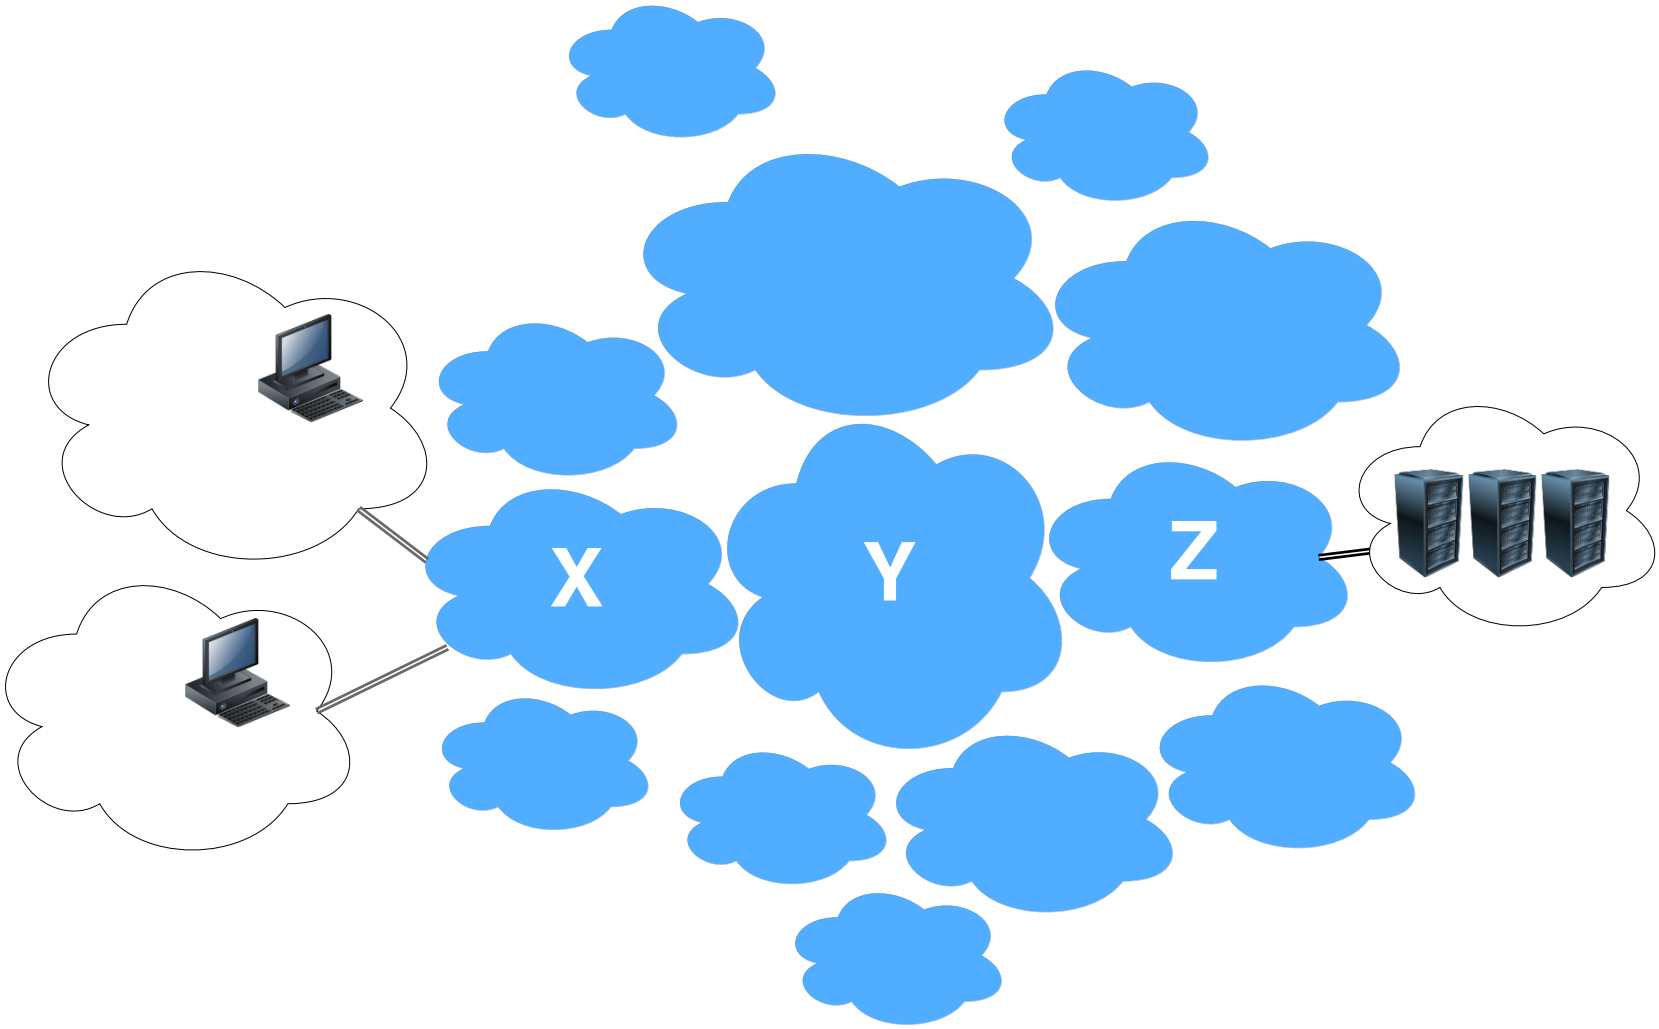

Understanding Internet health?

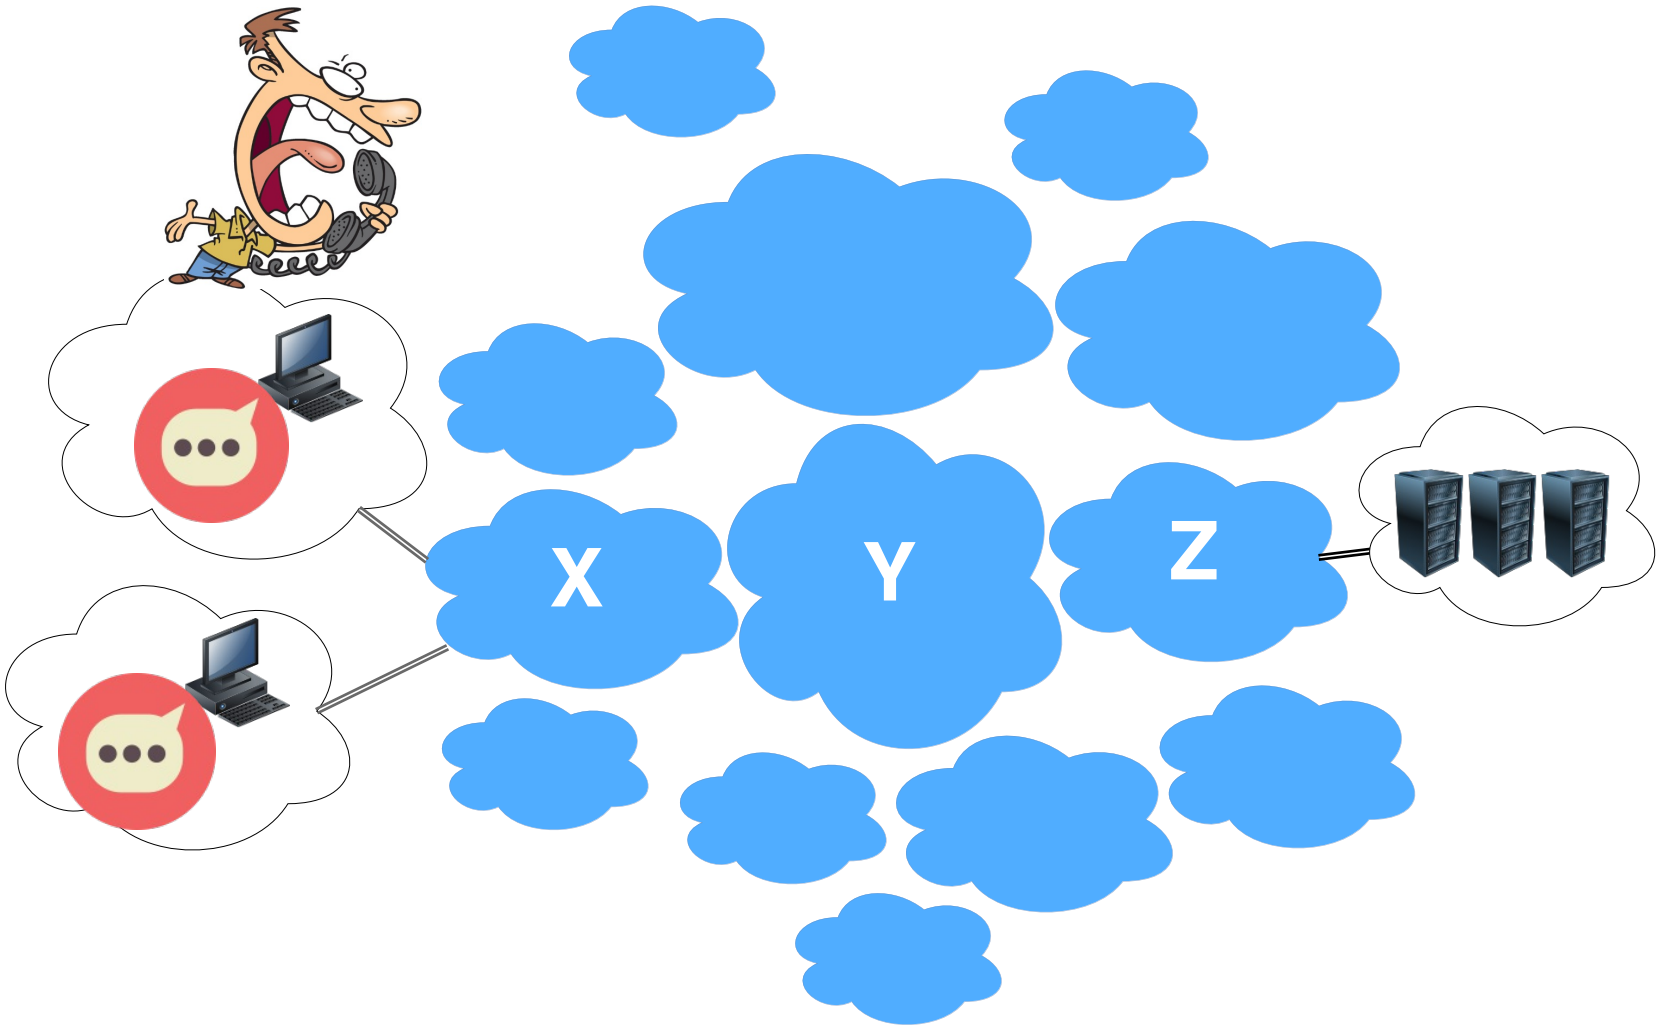

Understanding Internet health?

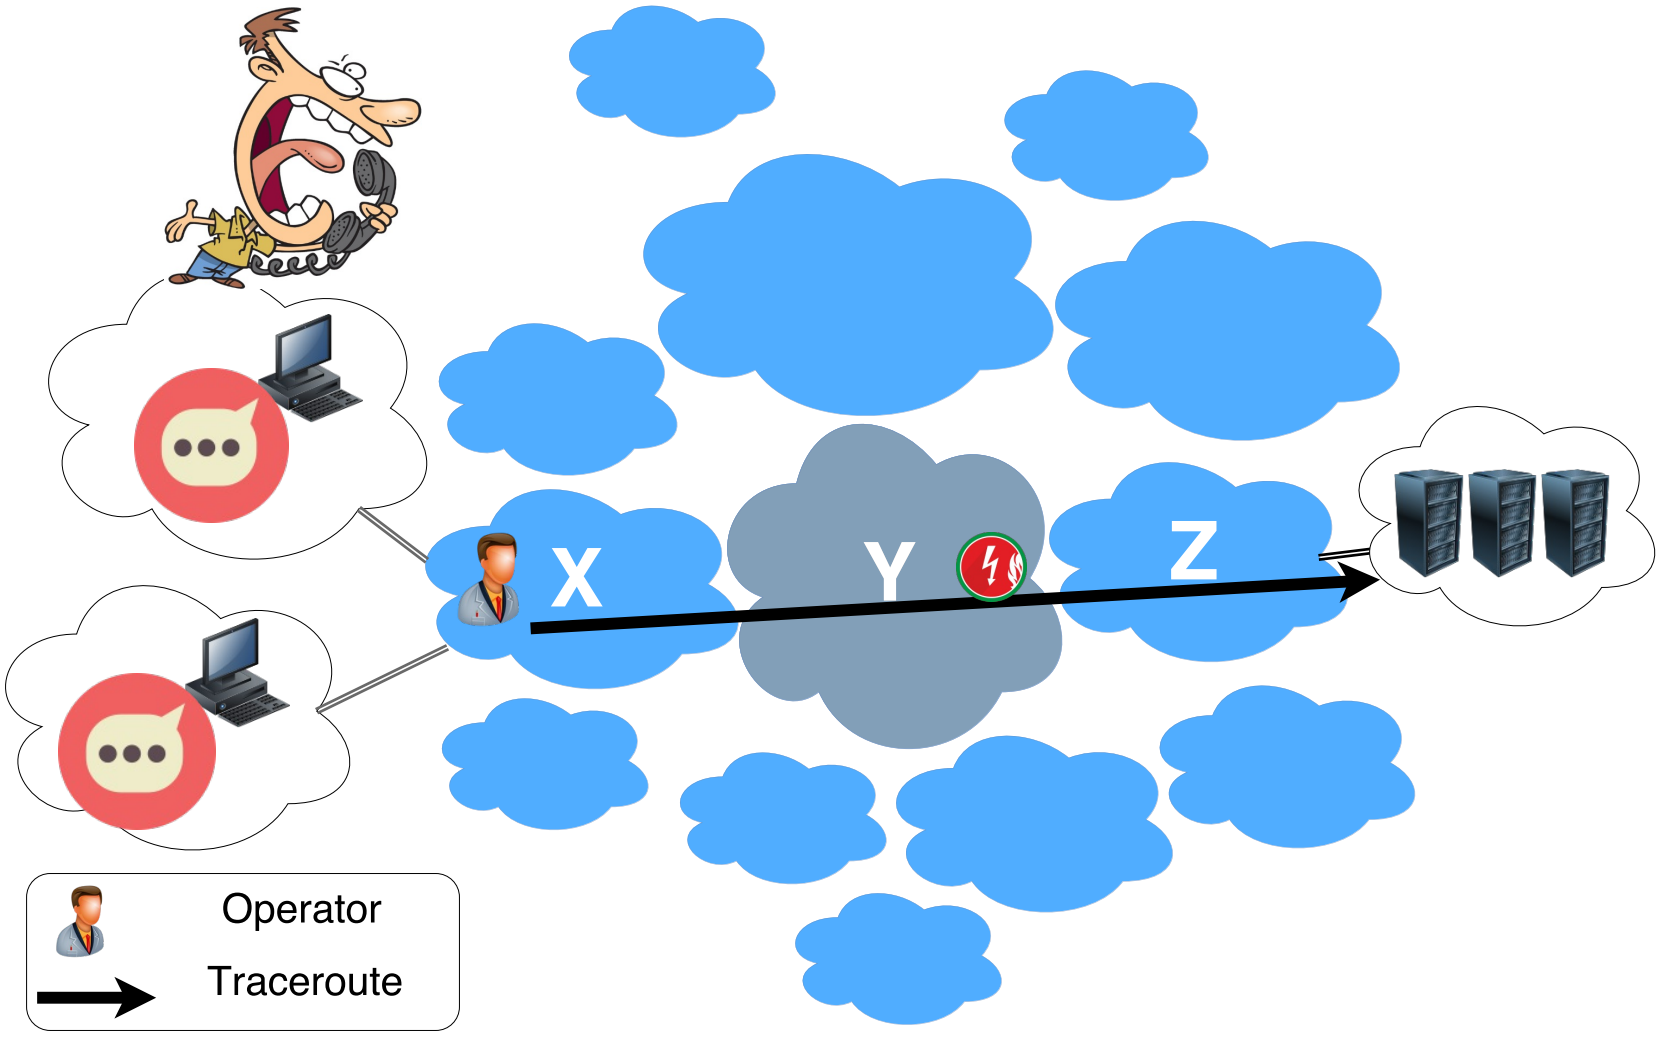

Understanding Internet health? (Problems)

Manual observations and operations

- Traceroute / Ping / Operators' group mailing lists
- Time consuming
- Slow process
- Small visibility

→ Our goal: Pinpointing network disruptions (i.e. congestion and packet loss)

Silly solution: frequent traceroutes to the whole Internet!

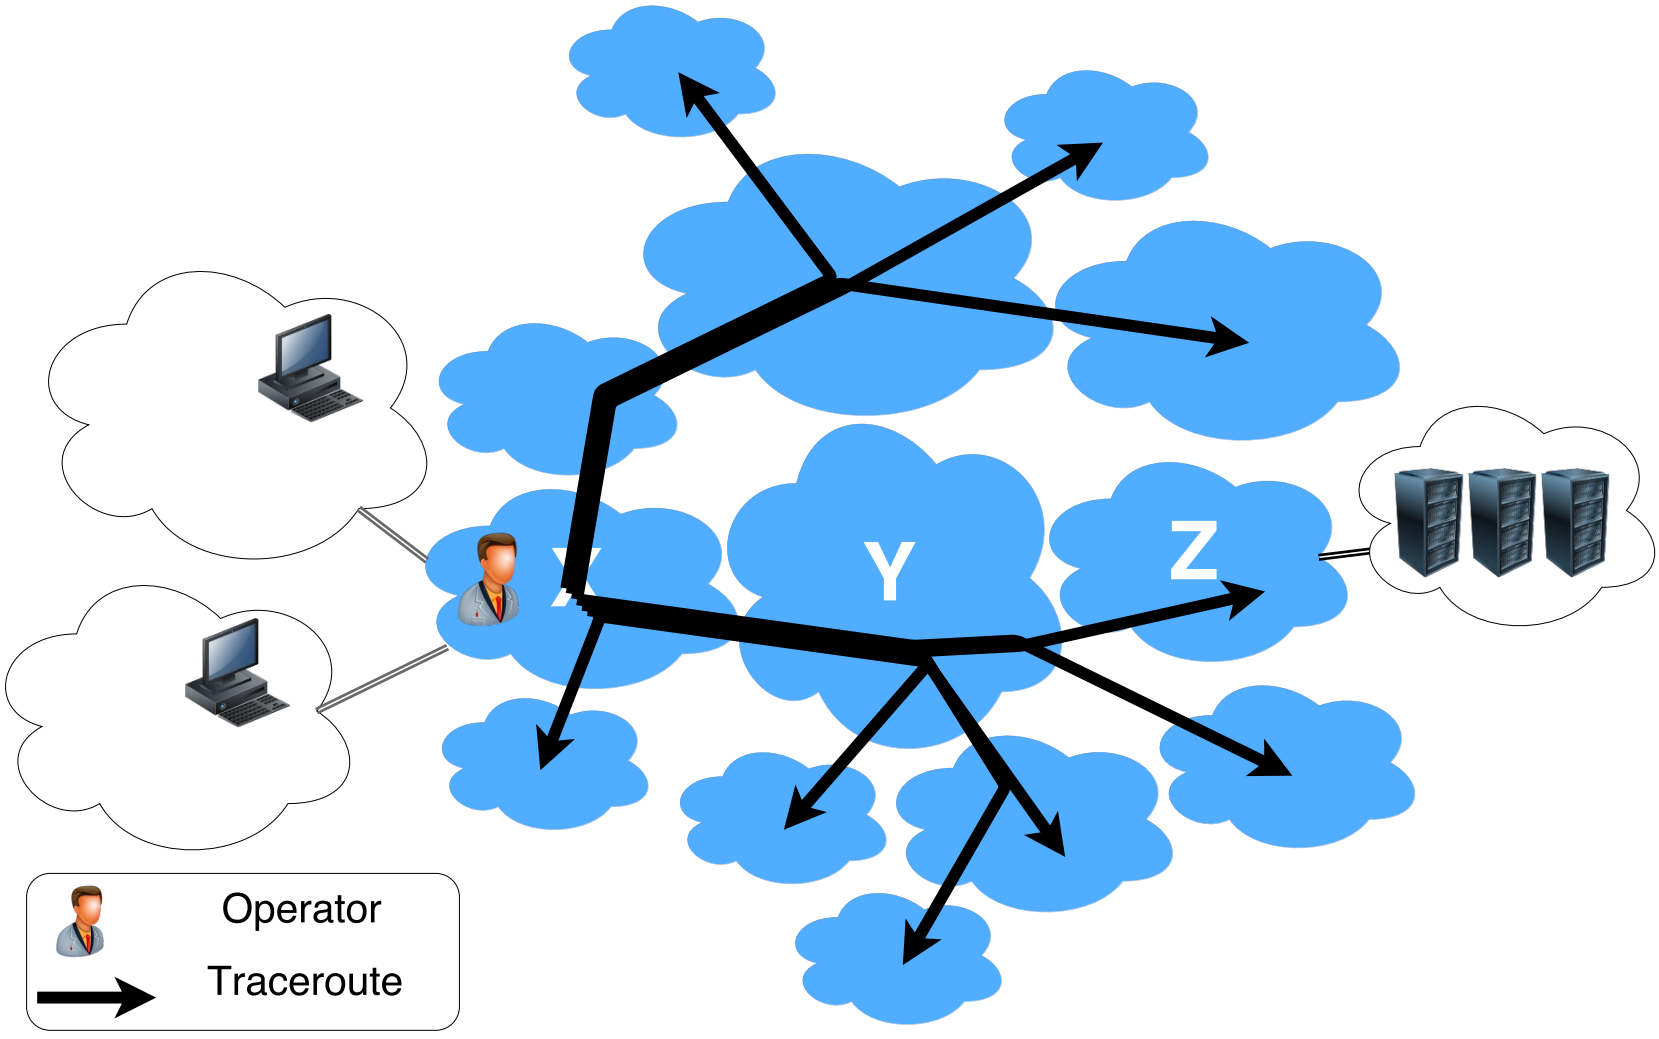

→ Doesn't scale

→ Overload the network

Better solution: mine results from deployed platforms

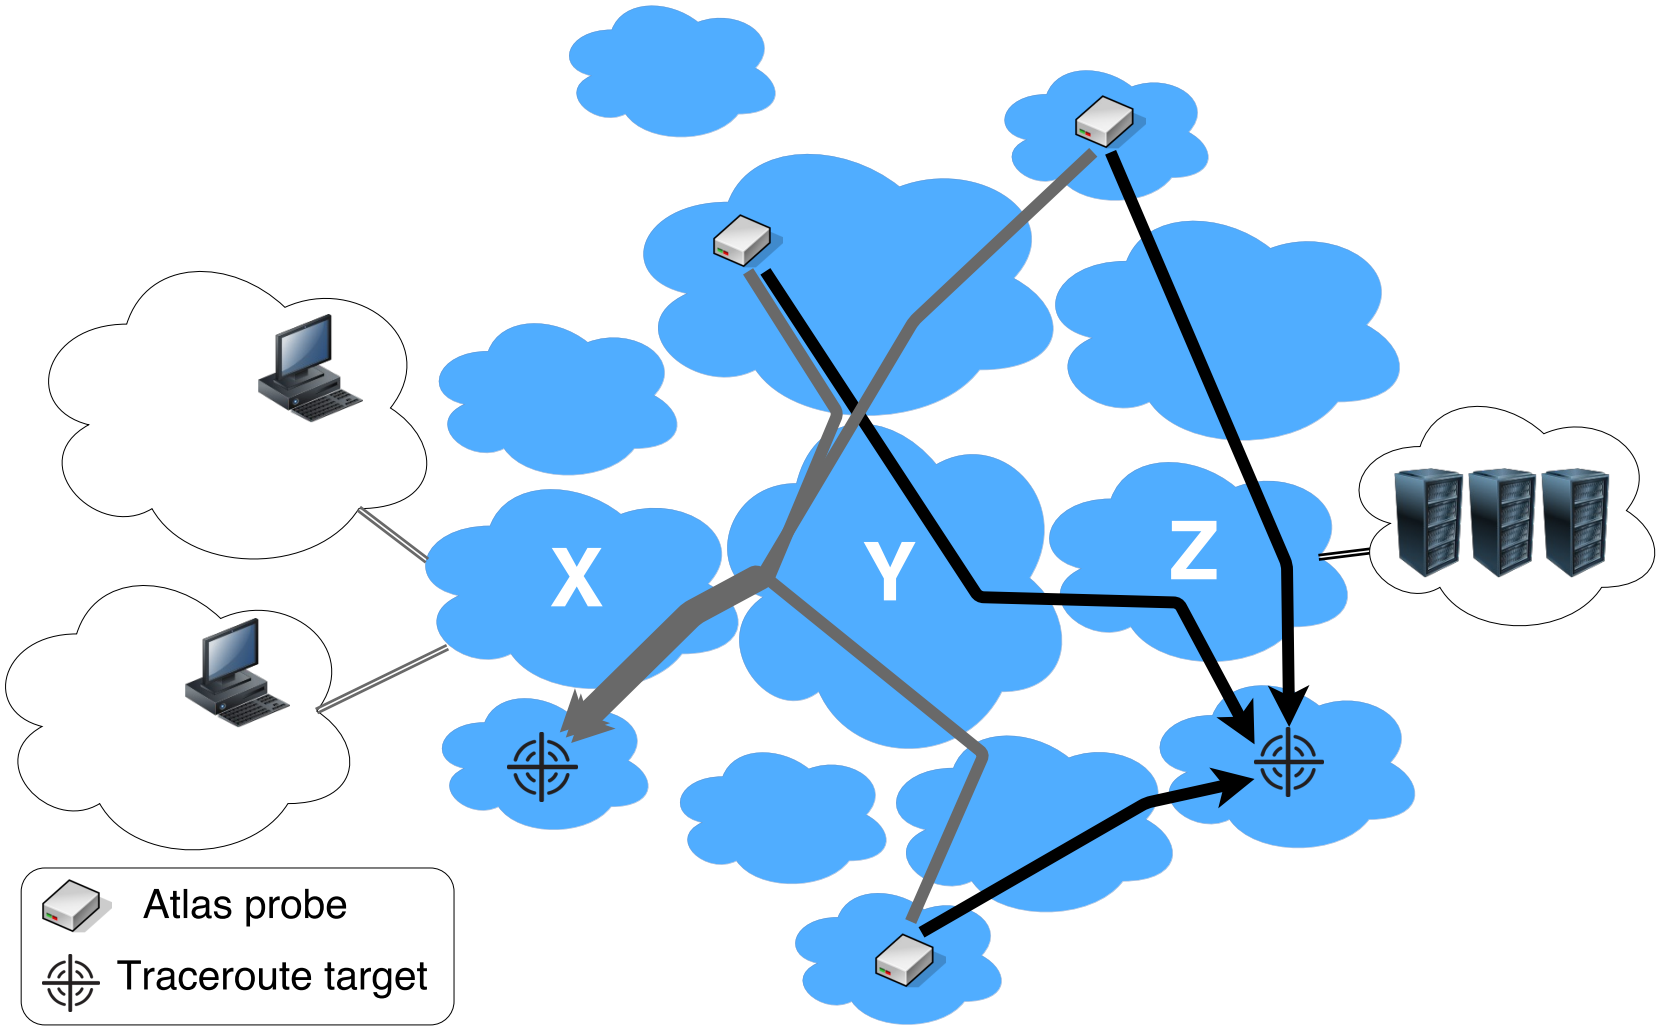

→ Cooperative and distributed approach

→ Using existing data, no added burden to the network

RIPE Atlas

Actively measures Internet connectivity

- Ethernet port
- Automatically perform active measurements: ping, **traceroute**, DNS, SSL, NTP and HTTP
- All results are collected by RIPE NCC

RIPE Atlas: coverage

9300+ active probes!

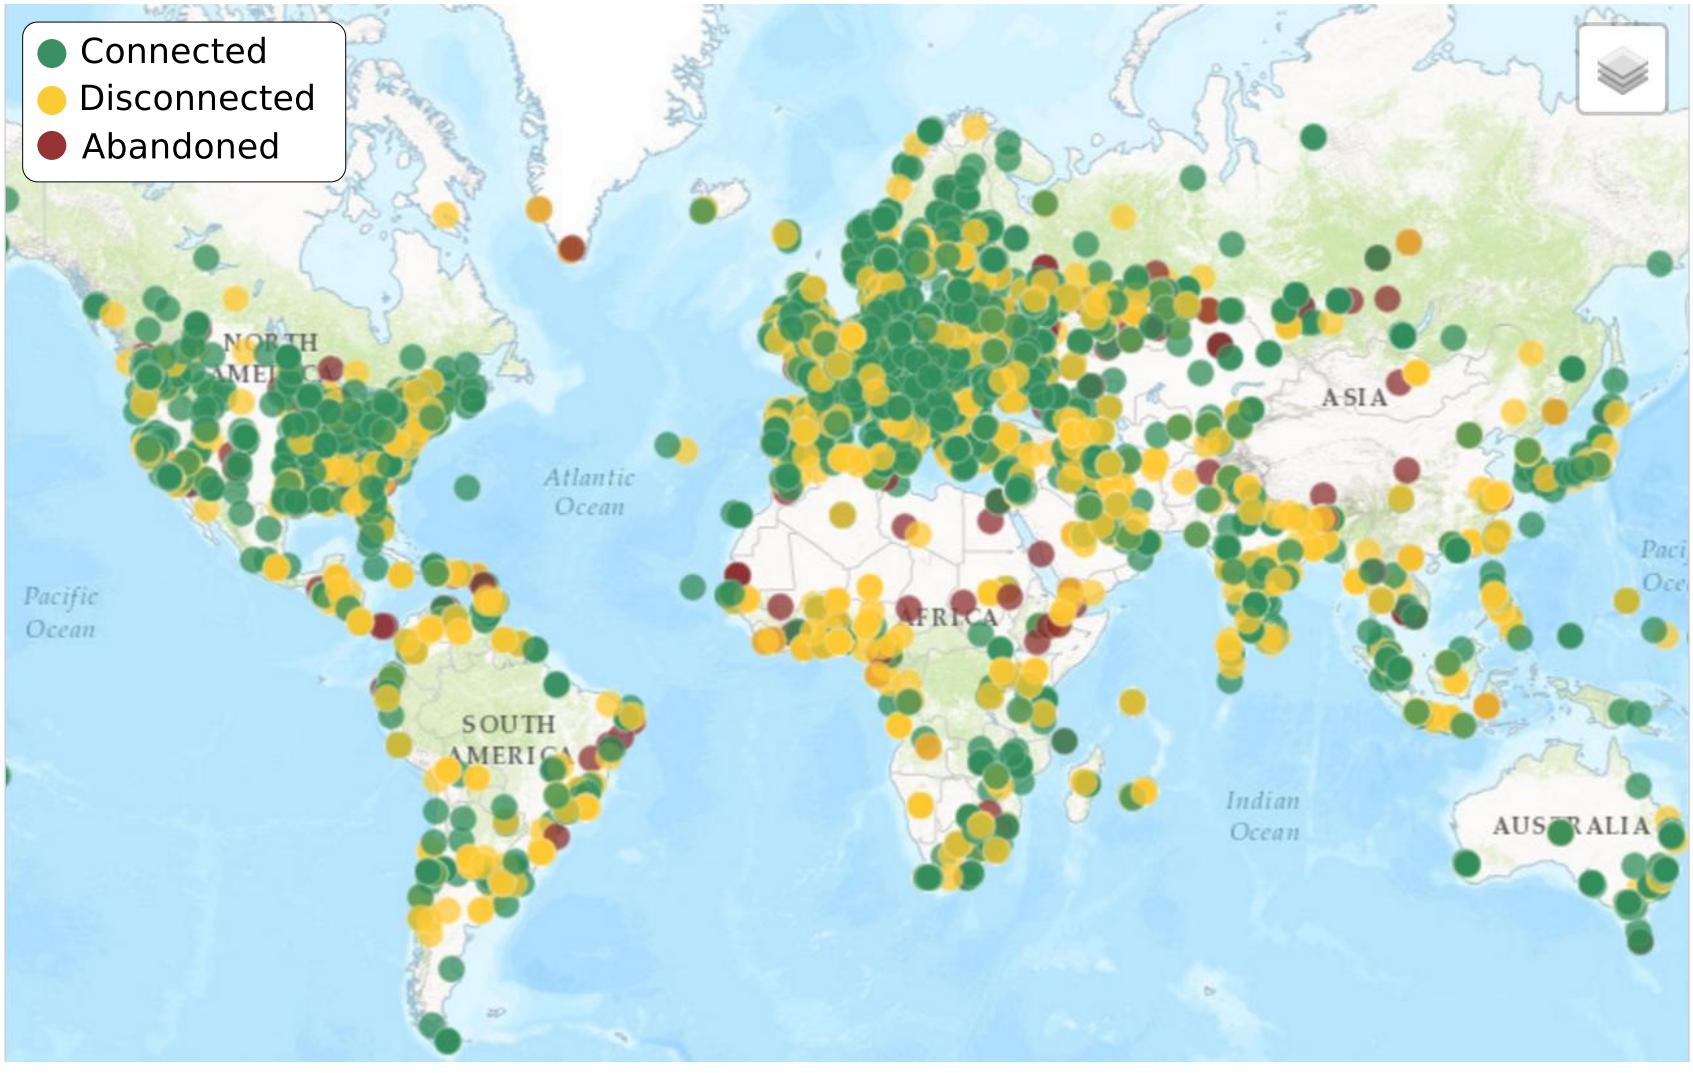

RIPE Atlas: traceroutes

Two repetitive large-scale measurements

- *Builtin*: traceroute every 30 minutes to all DNS root servers ($\approx$ 500 server instances)
- *Anchoring*: traceroute every 15 minutes to 189 collaborative servers

Analyzed dataset

- May to December 2015
- 2.8 billion IPv4 traceroutes
- 1.2 billion IPv6 traceroutes

Monitor delays with traceroute?

Traceroute to “www.target.com”

```
~$ traceroute www.target.com
traceroute to target, 30 hops max, 60 byte packets
 1  A       0.775 ms  0.779 ms  0.874 ms
 2  B       0.351 ms  0.365 ms  0.364 ms
 3  C       2.833 ms  3.201 ms  3.546 ms
 4  Target  3.447 ms  3.863 ms  3.872 ms
```

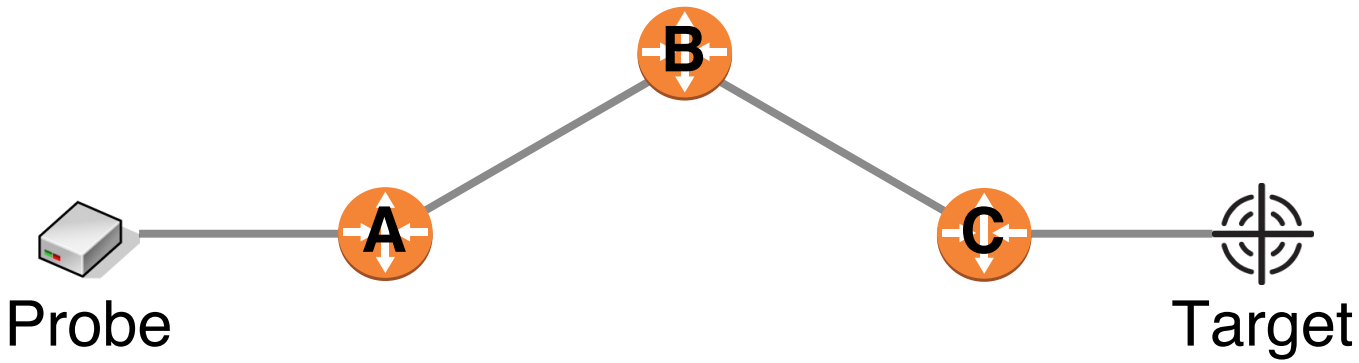

Round Trip Time (RTT) between B and C?

Report abnormal RTT between B and C?

Monitor delays with traceroute?

Challenges:

- Noisy data

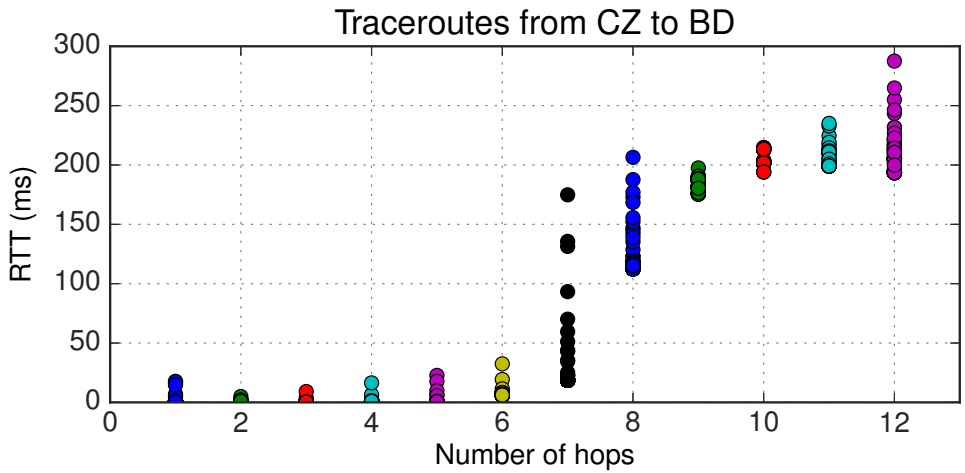

Monitor delays with traceroute?

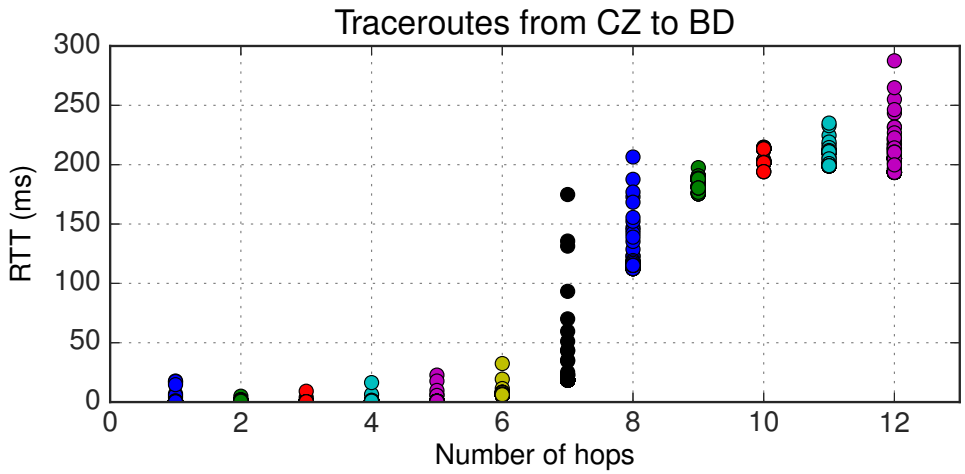

Challenges:

- Noisy data
- Traffic asymmetry
- Packet loss

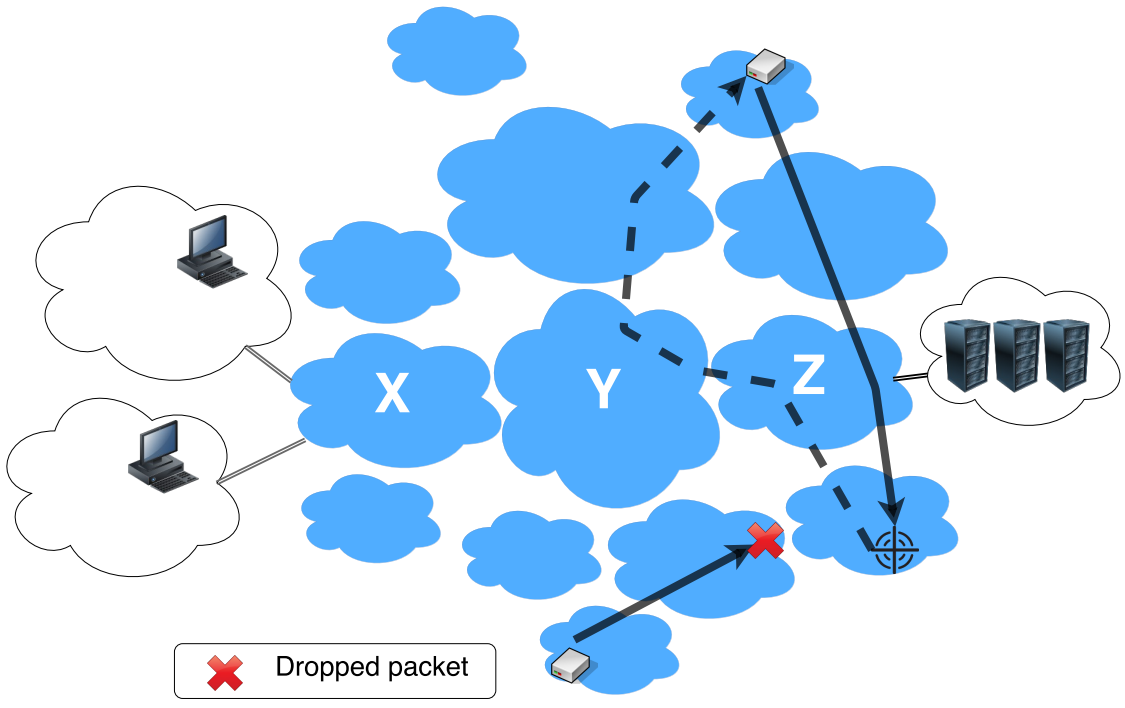

What is the RTT between B and C?

```
~$ traceroute www.target.com
traceroute to target, 30 hops max, 60 byte packets
 1  A       0.775 ms  0.779 ms  0.874 ms
 2  B       0.351 ms  0.365 ms  0.364 ms
 3  C       2.833 ms  3.201 ms  3.546 ms
 4  Target  3.447 ms  3.863 ms  3.872 ms
```

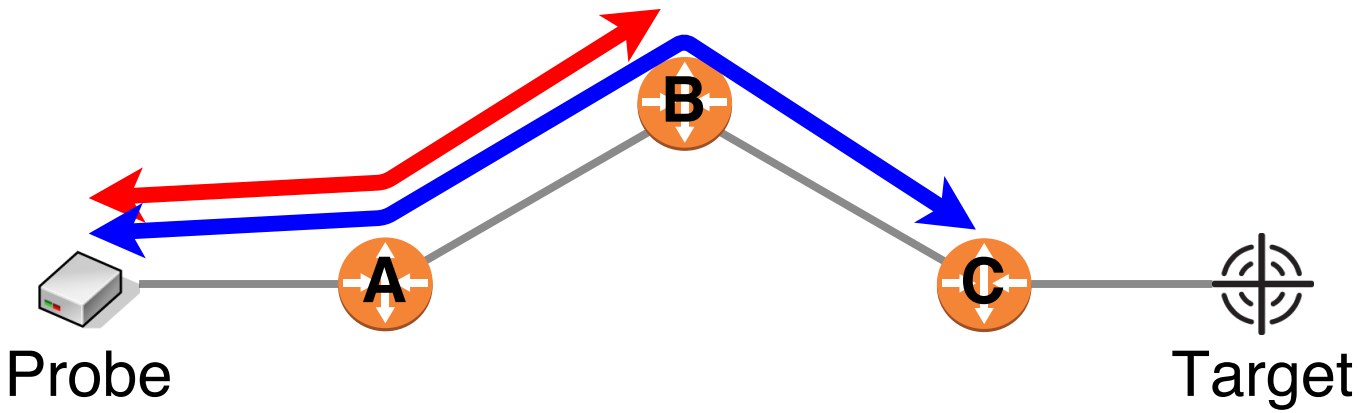

$$RTT_C - RTT_B = RTT_{CB}?$$

What is the RTT between B and C?

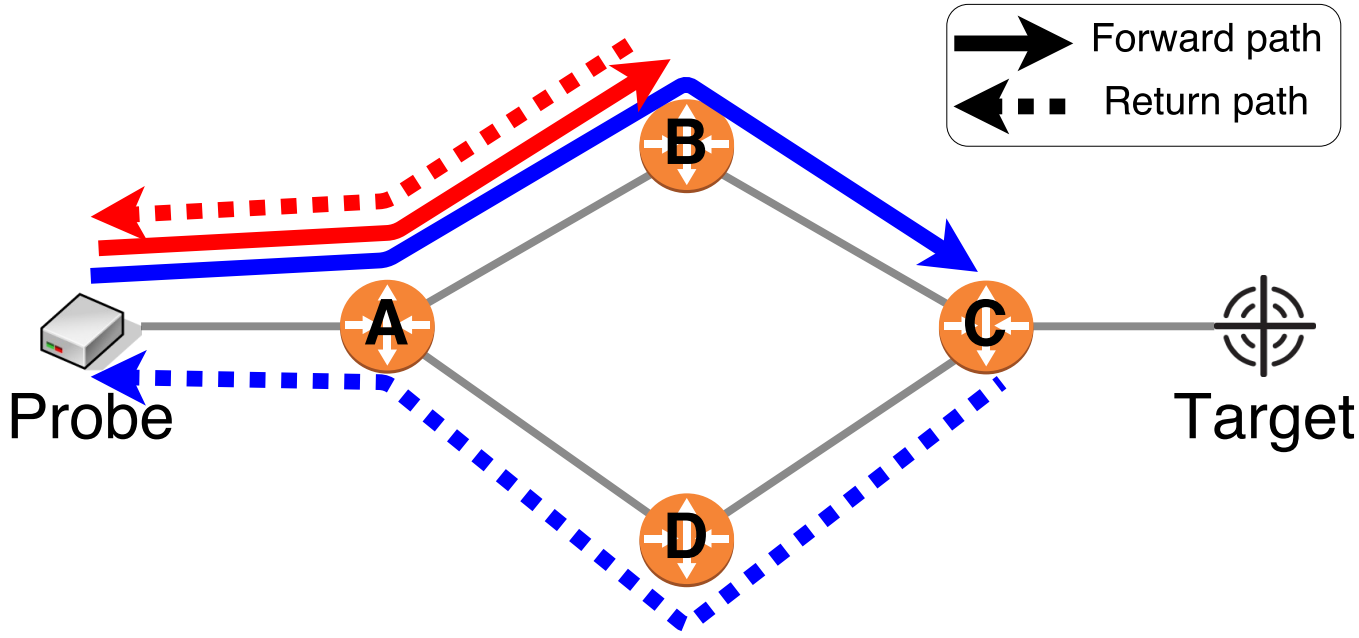

RTT_C - $RTT_B = RTT_{CB}$?

- No!
- Traffic is asymmetric
- RTT_B and RTT_C take **different return paths!**

What is the RTT between B and C?

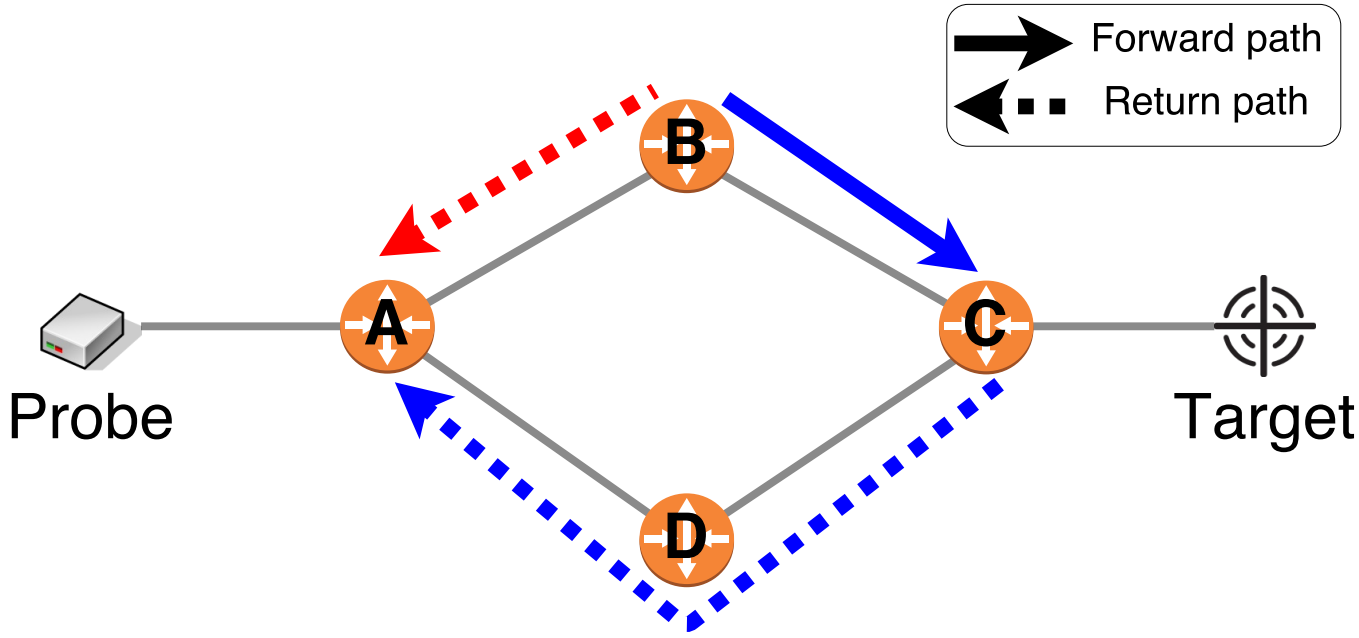

RTT_C - $RTT_B = RTT_{CB}$?

- No!
- Traffic is asymmetric
- RTT_B and RTT_C take **different return paths!**
- **Differential RTT**: $\Delta_{CB} = RTT_C - RTT_B = d_{BC} + e_p$

Problem with differential RTT

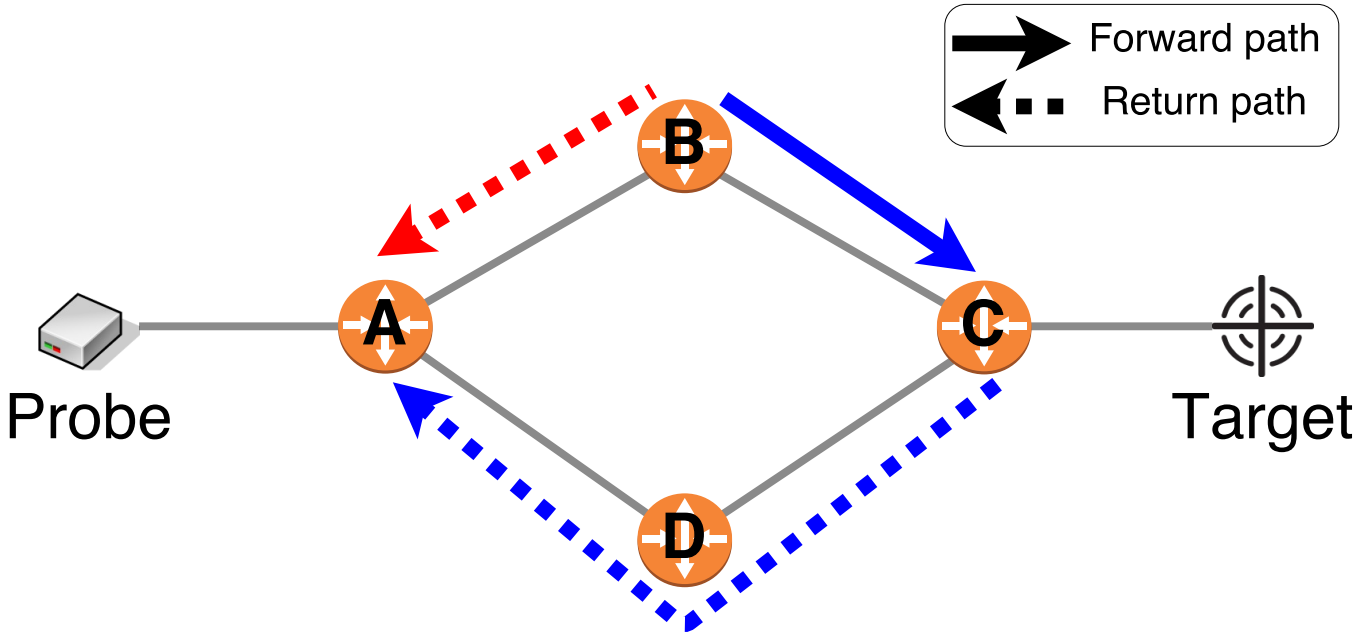

Monitoring Δ_{CB} over time:

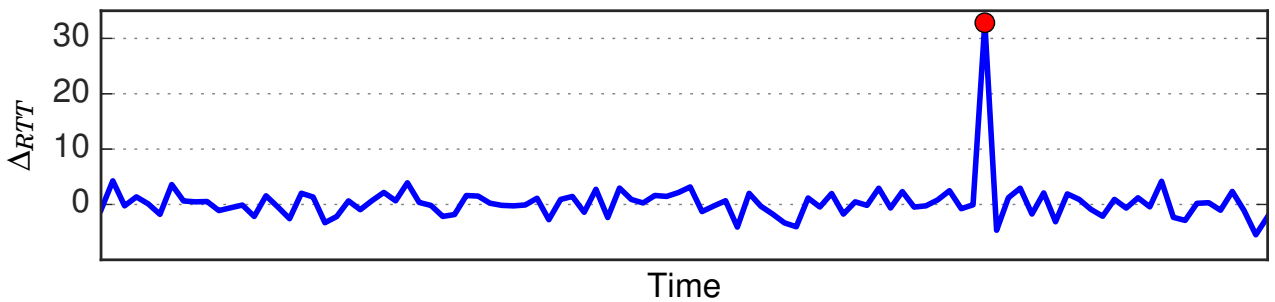

$\rightarrow$ Delay change on BC? CD? DA? BA???

Proposed Approach: Use probes with different return paths

Differential RTT: $\Delta_{CB} = x_0$

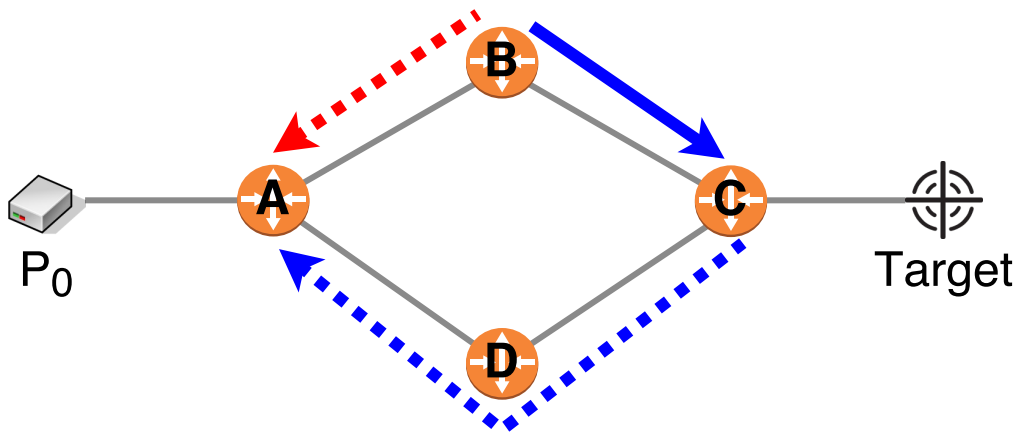

Proposed Approach: Use probes with different return paths

Differential RTT: $\Delta_{CB} = \{x_0, x_1\}$

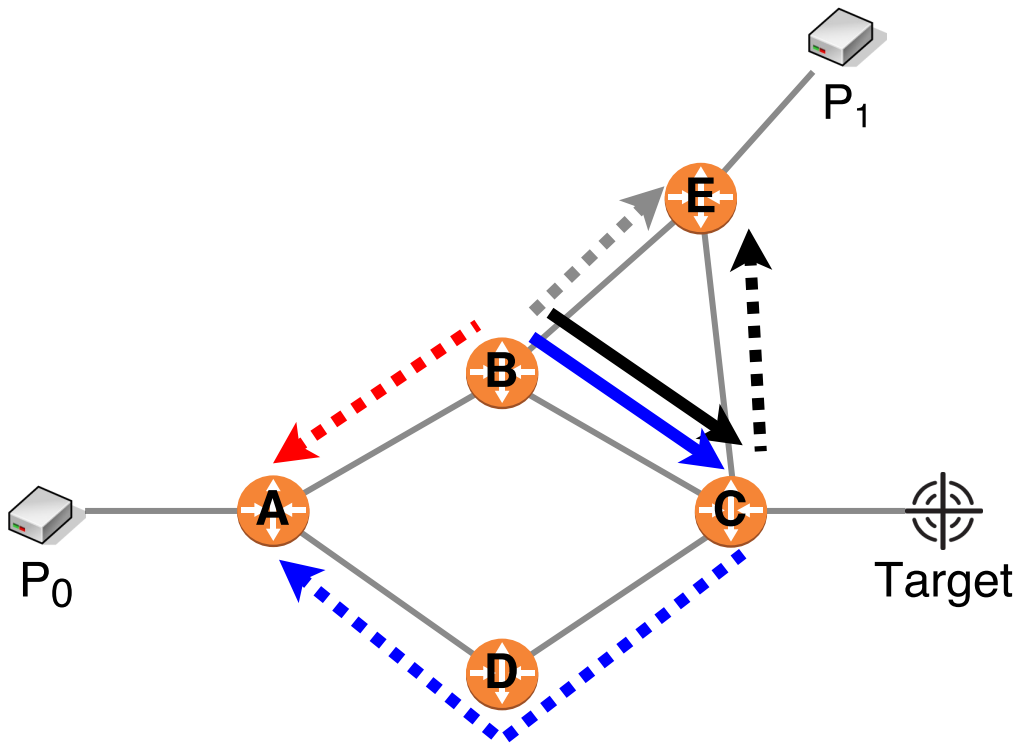

Proposed Approach: Use probes with different return paths

Differential RTT: $\Delta_{CB} = \{x_0, x_1, x_2, x_3, x_4\}$

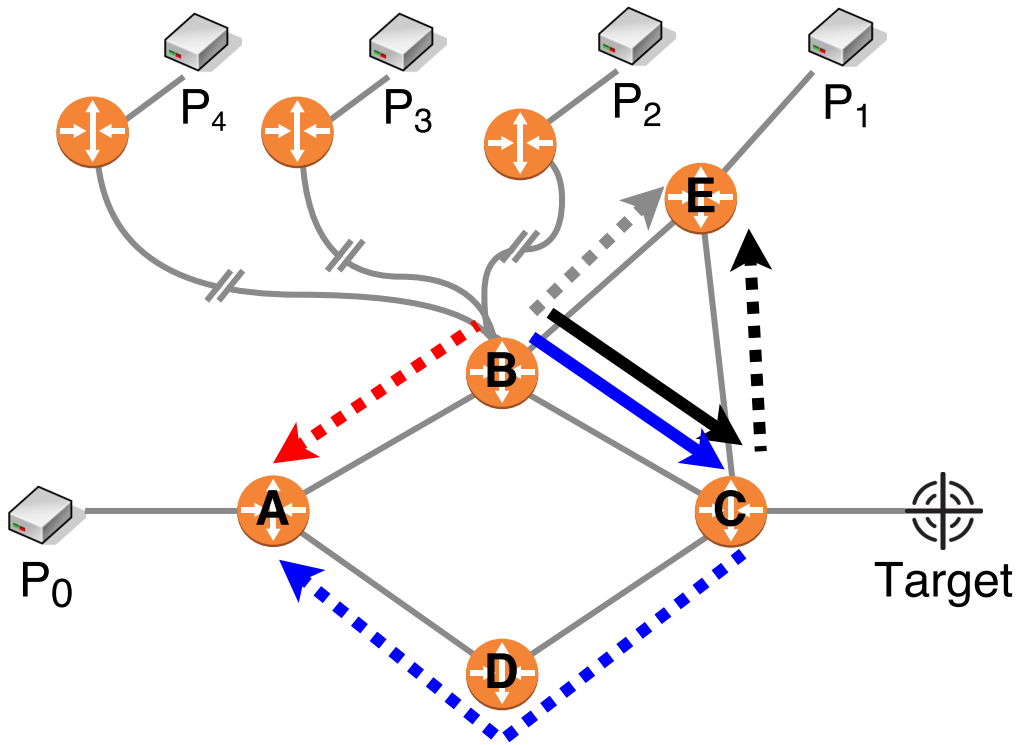

Proposed Approach: Use probes with different return paths

Differential RTT: $\Delta_{CB} = \{x_0, x_1, x_2, x_3, x_4\}$

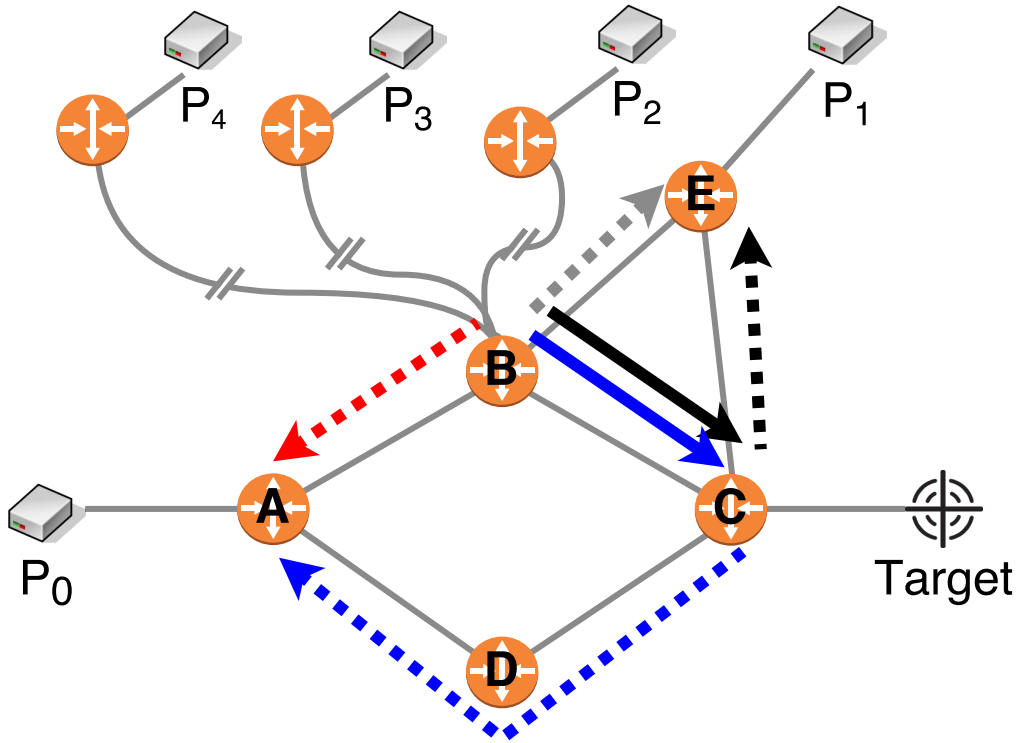

Median Δ_{CB}:

- Stable if a few return paths delay change
- Fluctuate if delay on BC changes

Median Diff. RTT: Example

Tier1 link, 2 weeks of data, 95 probes:

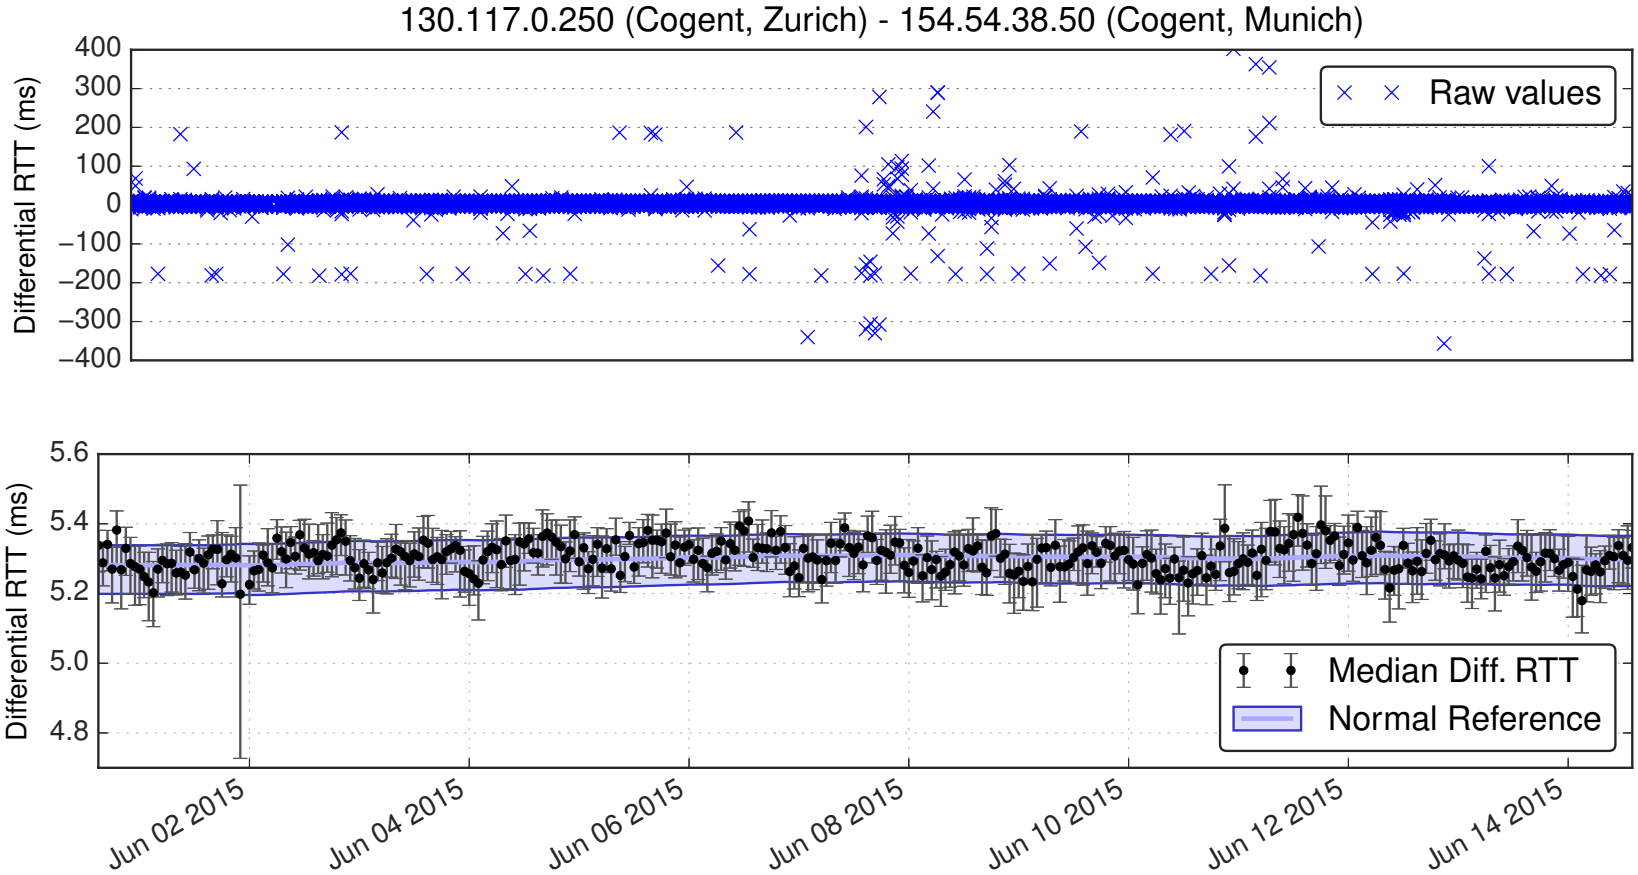

- **Stable** despite noisy RTTs (not true for average)
- Normally distributed

Detecting congestion

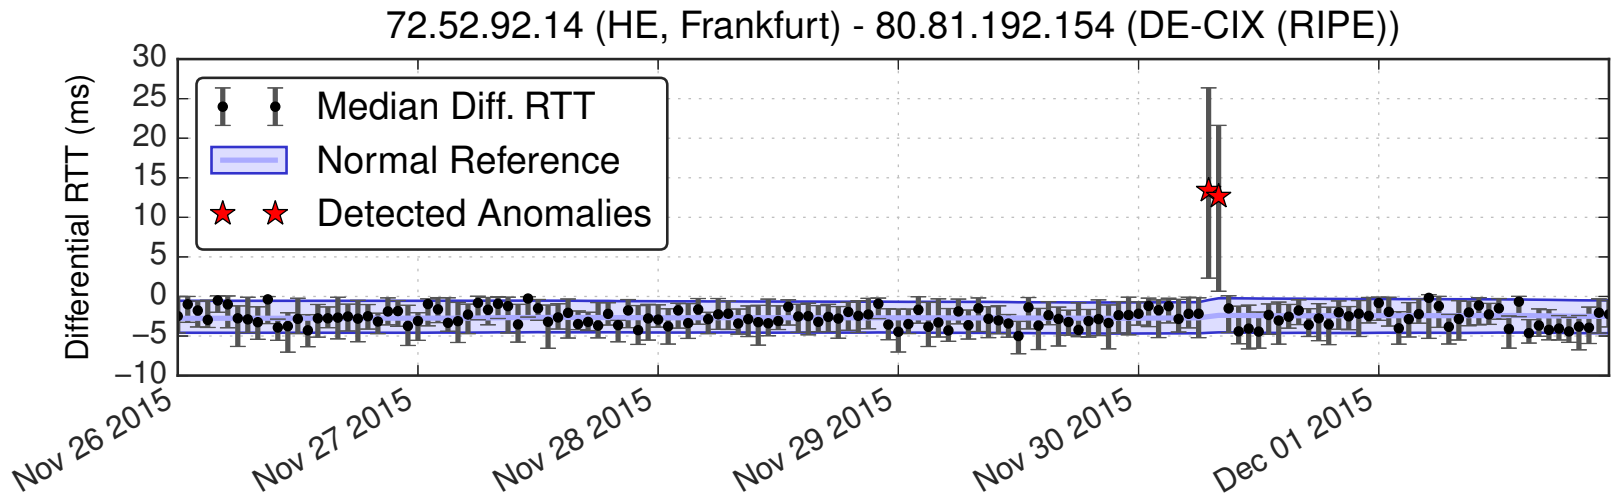

Significant RTT changes:

Confidence interval not overlapping with the normal reference

Packet loss

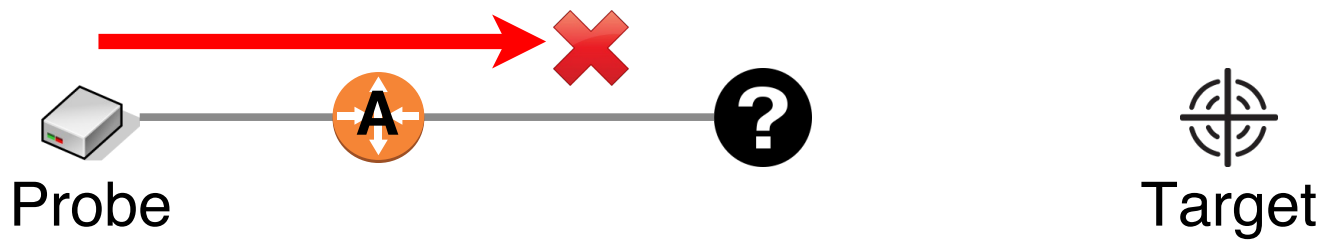

Worst case: router is not responding

- Cannot obtain RTT values
- Need to identify the faulty link

Packet forwarding model

Learn usual paths from past traceroutes:

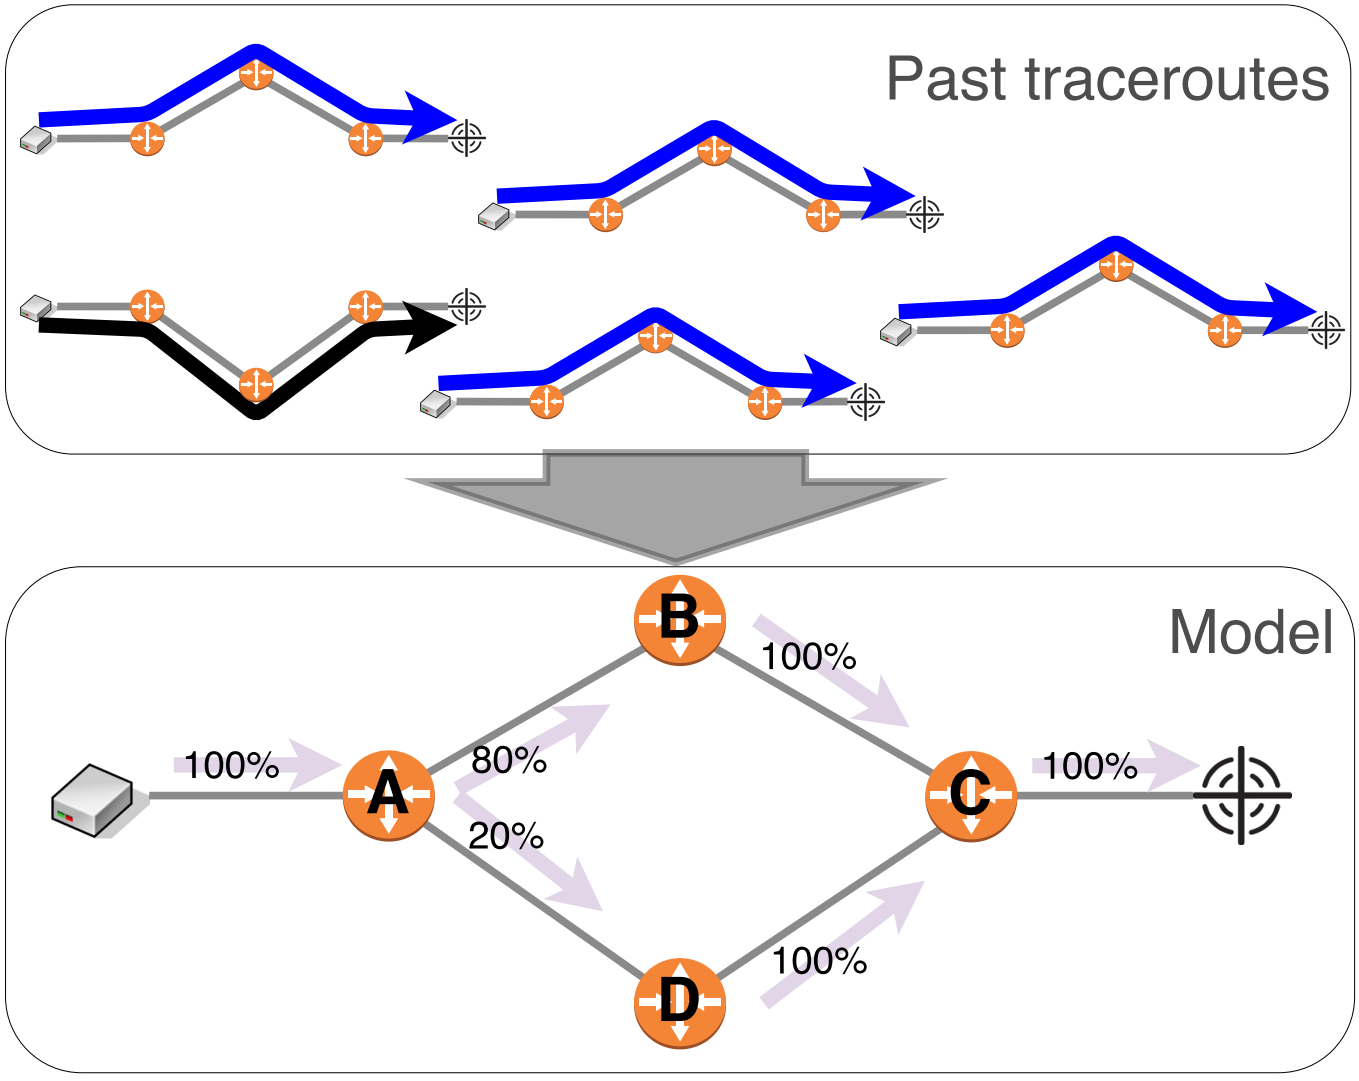

Identifying faulty links

In case of packet loss:

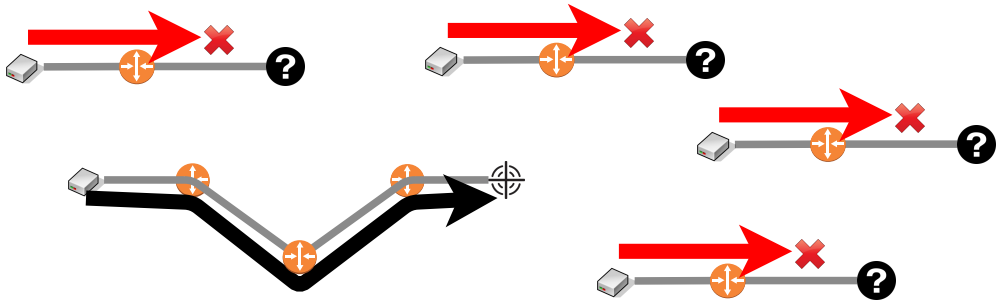

Query the model for the expected next hop

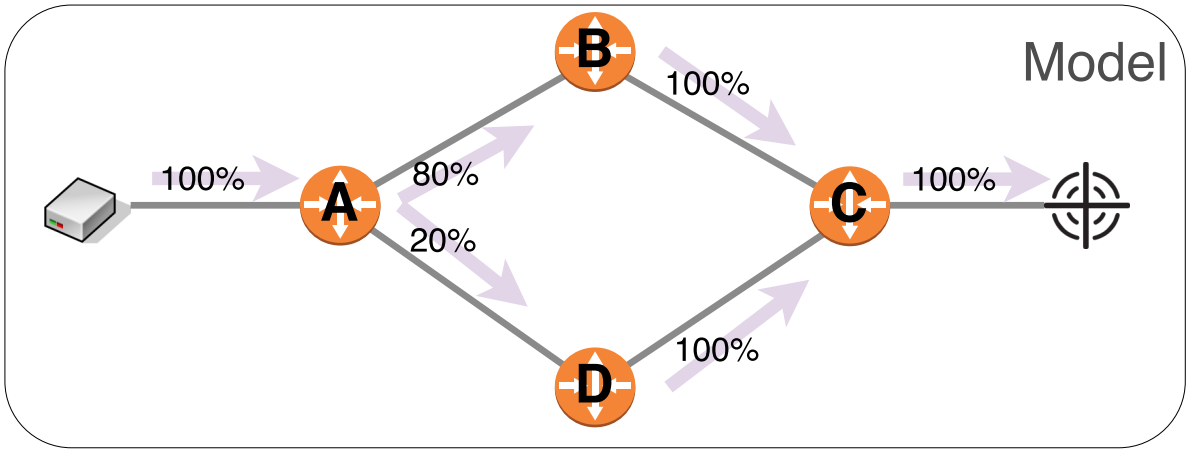

$\rightarrow$ Link AB is dropping packets!

Results

Analyzed dataset

- Atlas *builtin*/*anchoring* measurements
- From May to Dec. 2015
- Observed 262k IPv4 and 42k IPv6 links

We found a lot of congested links!
Let's see only two significant examples

Study case: DDoS on DNS root servers

Two attacks:

- Nov. 30th 2015
- Dec. 1st 2015

Almost all server are anycast

- Congestion at the 531 sites?
- Found 129 instances altered by the attacks

The Register®
Biting the hand that feeds IT

DATA CENTRE SOFTWARE NETWORKS SECURITY INFRASTRUCTURE DEVOPS BUSINESS HARDWARE

Networks

Internet's root servers take hit in DDoS attack

Who's testing the limits of the DNS system?

8 Dec 2015 at 23:10, Kieren McCarthy

The internet's root servers came under a c
effectively knocked three of the 13 critical pi

The attack came just days before the Janet

According to a first analysis of the root serv
attack occurred on November 30, 2015 bet

Many, but not all, of the root servers receive
flood network connections and cause timec
messages for a single domain name; the s

Ultimately, the operators affected by the atta
proper analysis is now underway to discove

Of perhaps most concern is the fact that ev
deal with such an attack, a number of the s

The root servers themselves make up the p
as a sort of global directory for all the other

Due to the internet's design, the servers the
you compare it to what companies like Goo
immediate problems for the wider internet,
thousands of other servers.

That said, any attack on the DNS' infrastruc
longer than a day, it would start causing sig

≤ 66% > 99% Data resolution: 10 minutes

Home Hacking Tech Cyber Attacks Vulnerabilities Malware Spying

The Hacker News™
Security in a serious way

Someone Just Tried to Take Down Internet's Backbone with 5 Million Queries/Sec

Wednesday, December 09, 2015 Swati Khandelwal

G+1 192 Like 5.9K Share 6538 Tweet 1049 Share 58 share 8417

The Internet's Backbone

DNS Root Servers Hit by a Massive Cyber Attack

Someone just DDoSed one of the most critical organs of the Internet anatomy – **The Internet's DNS Root Servers**.

Early last week, a flood of as many as 5 Million queries per second hit many of the Internet's DNS (*Domain Name System*) Root Servers that act as the authoritative reference for mapping domain names to IP addresses and are a total of 13 in numbers.

The attack, commonly known as **Distributed Denial of Service** (DDoS) attack, took place on two separate occasions.

The first DDoS attack to the Internet's backbone root servers launched on *November 30* that lasted 160 minutes (*almost 3 hours*), and the second one started on *December 1* that lasted almost an hour.

Massive Attacks Knocked Many of the 13 Root Servers Offline

Observed congestion

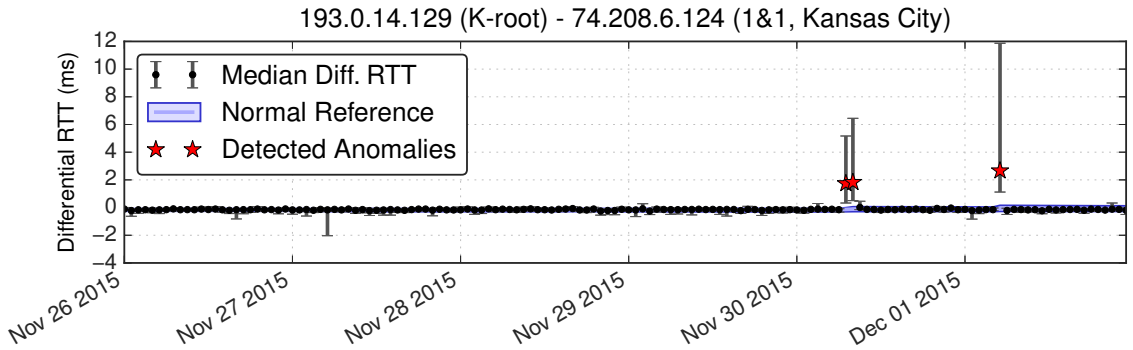

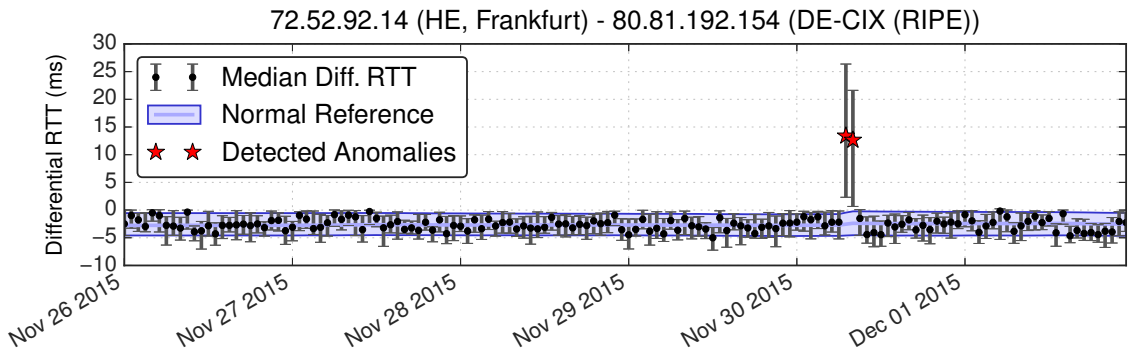

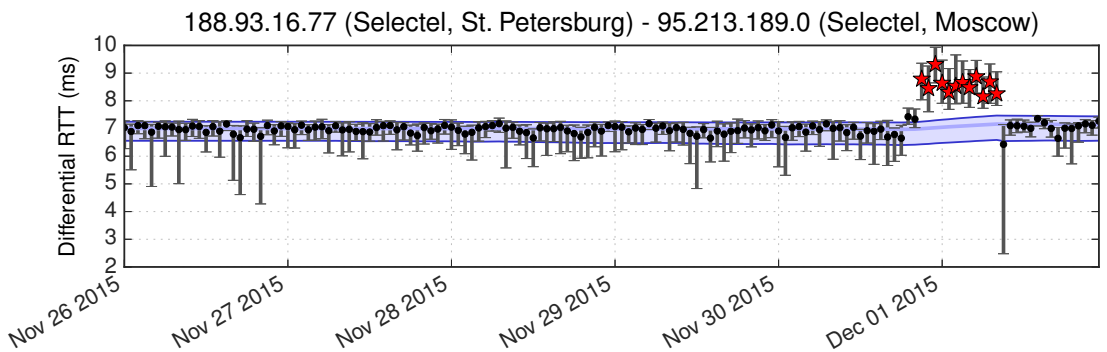

- Certain servers are affected only by one attack
- Continuous attack in Russia

Unaffected root servers

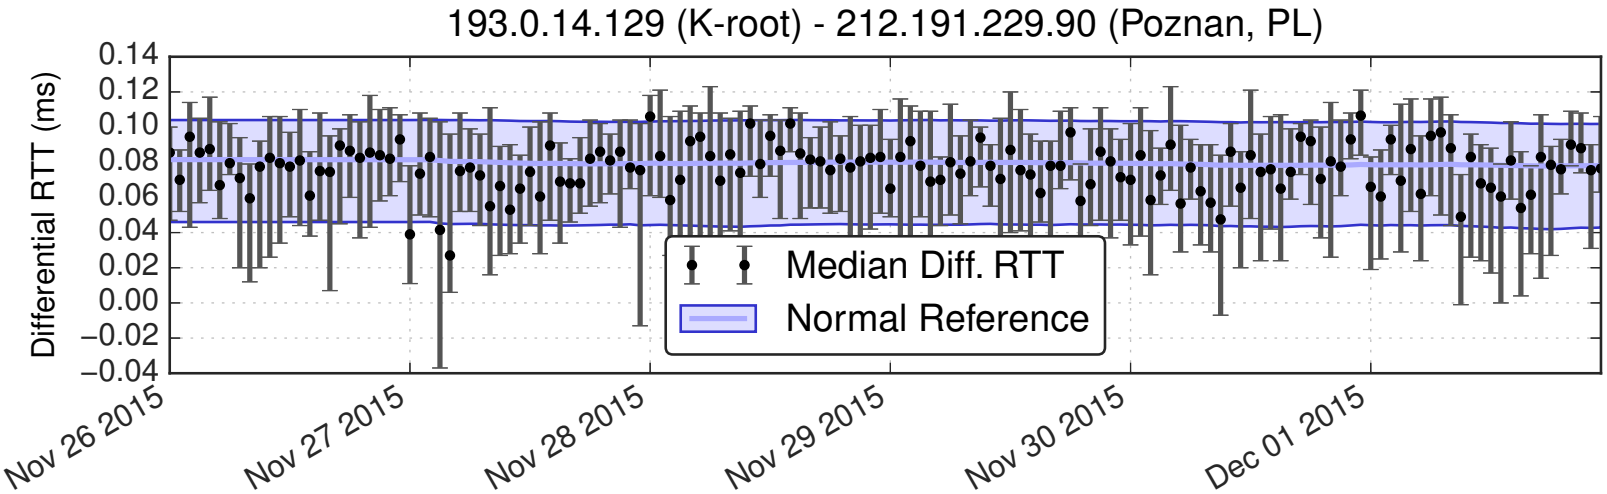

Very stable delay during the attacks

- Thanks to anycast!
- Far from the attackers

Congested links for servers F, I, and K

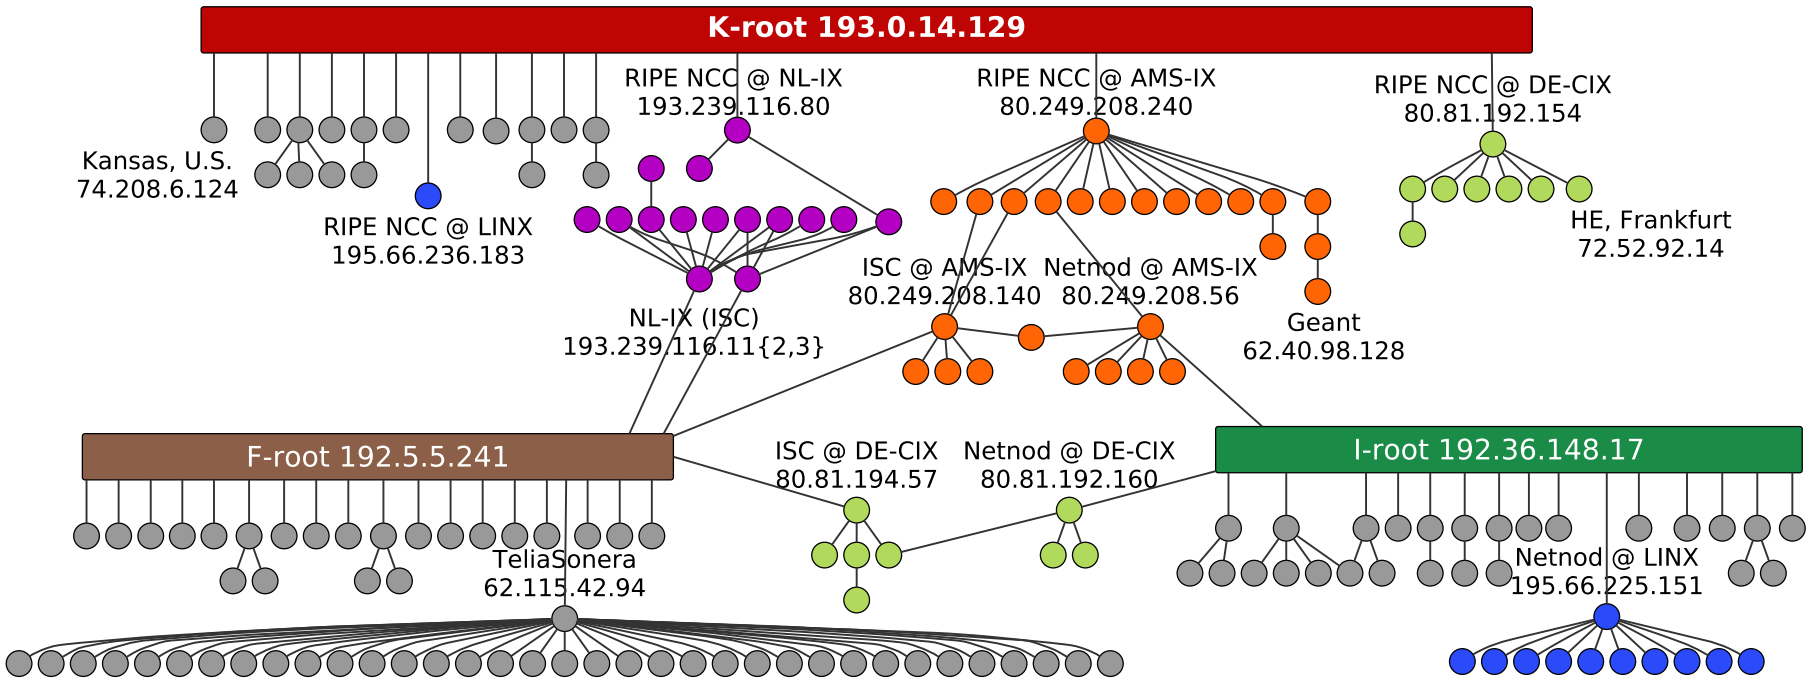

→ **Concentration of malicious traffic in IXPs**

Study case: Telekom Malaysia BGP leak

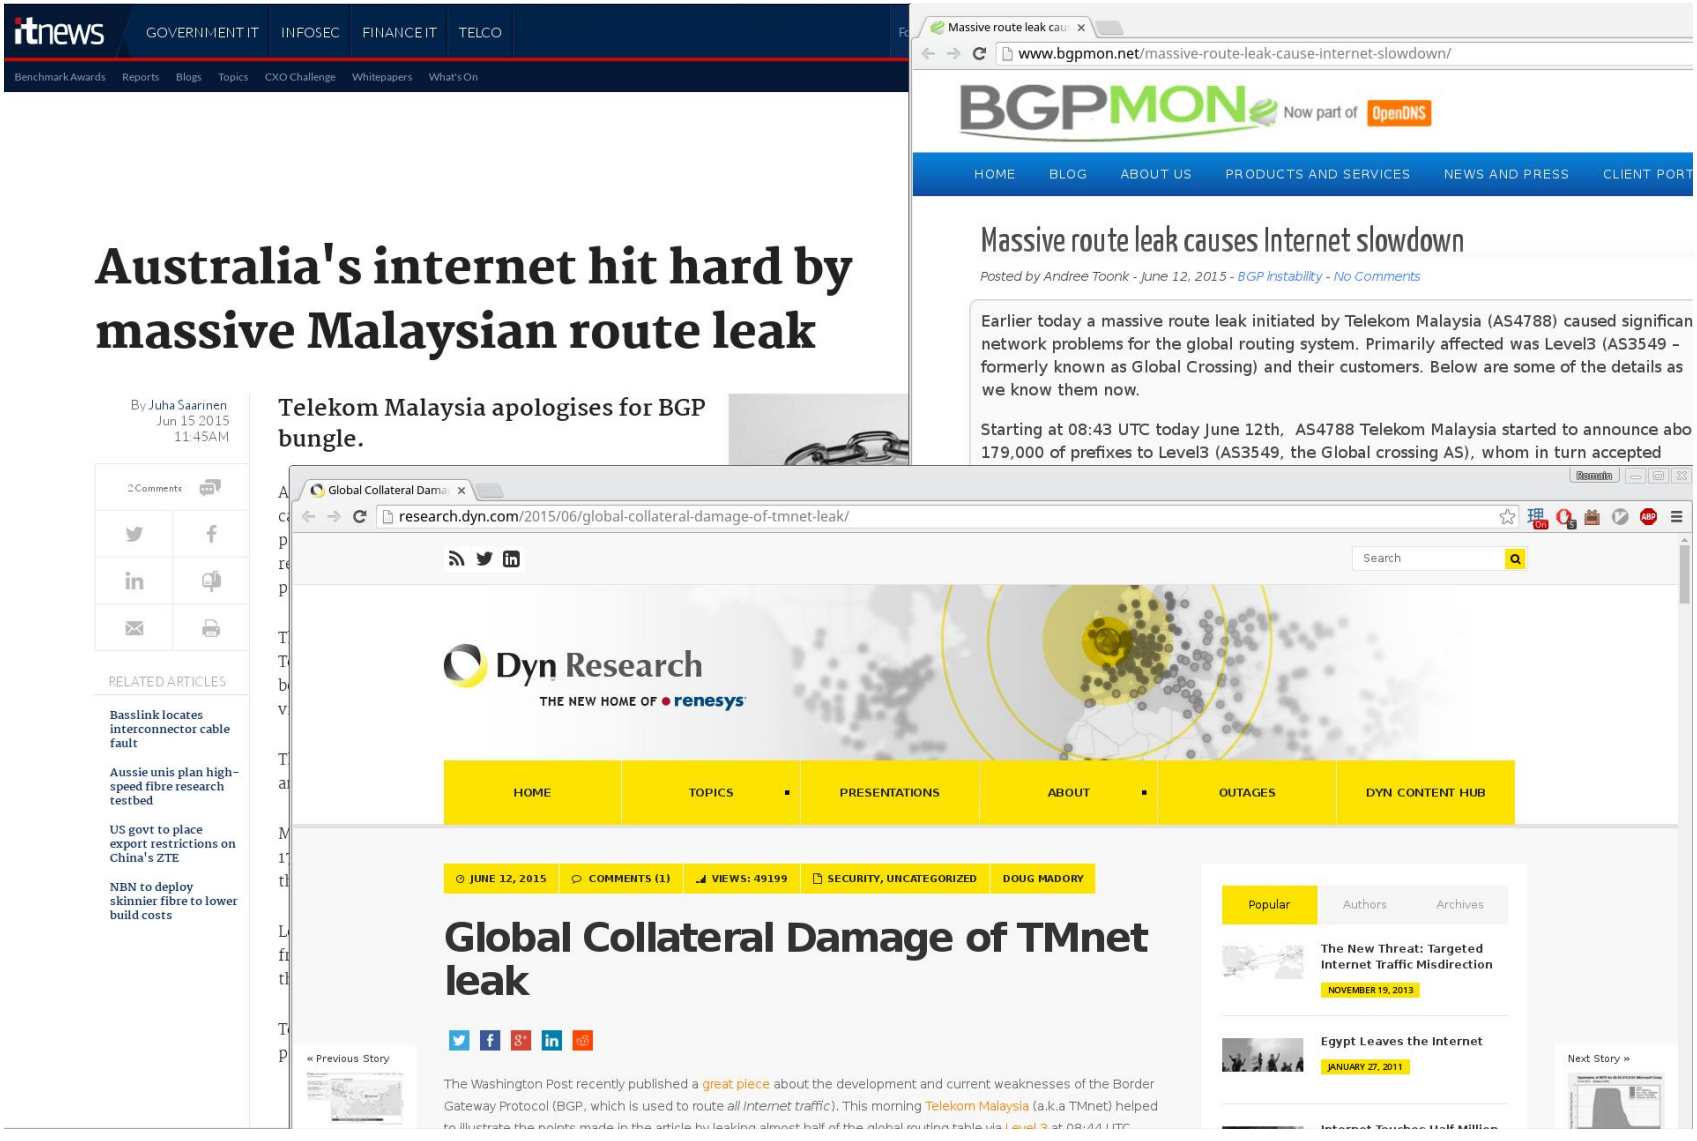

Study case: Telekom Malaysia BGP leak

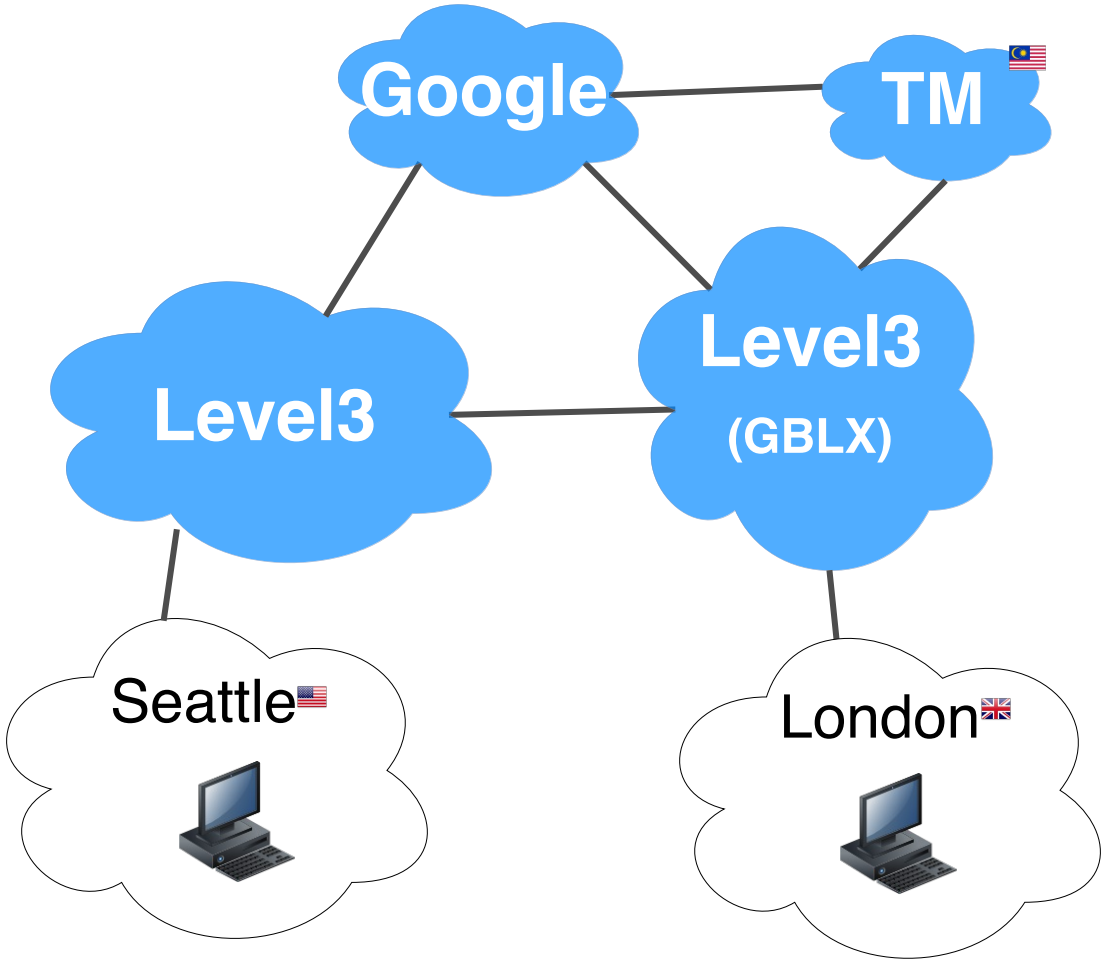

Study case: Telekom Malaysia BGP leak

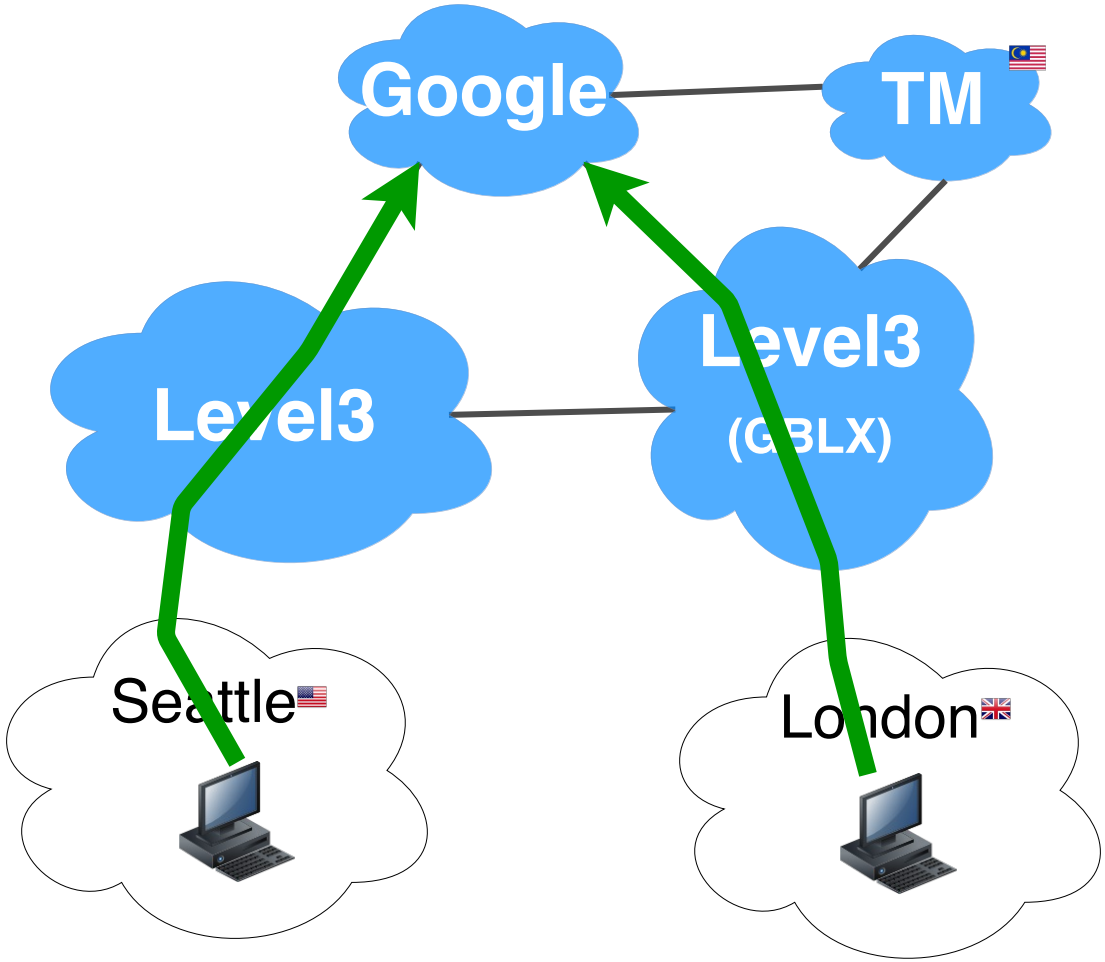

Study case: Telekom Malaysia BGP leak

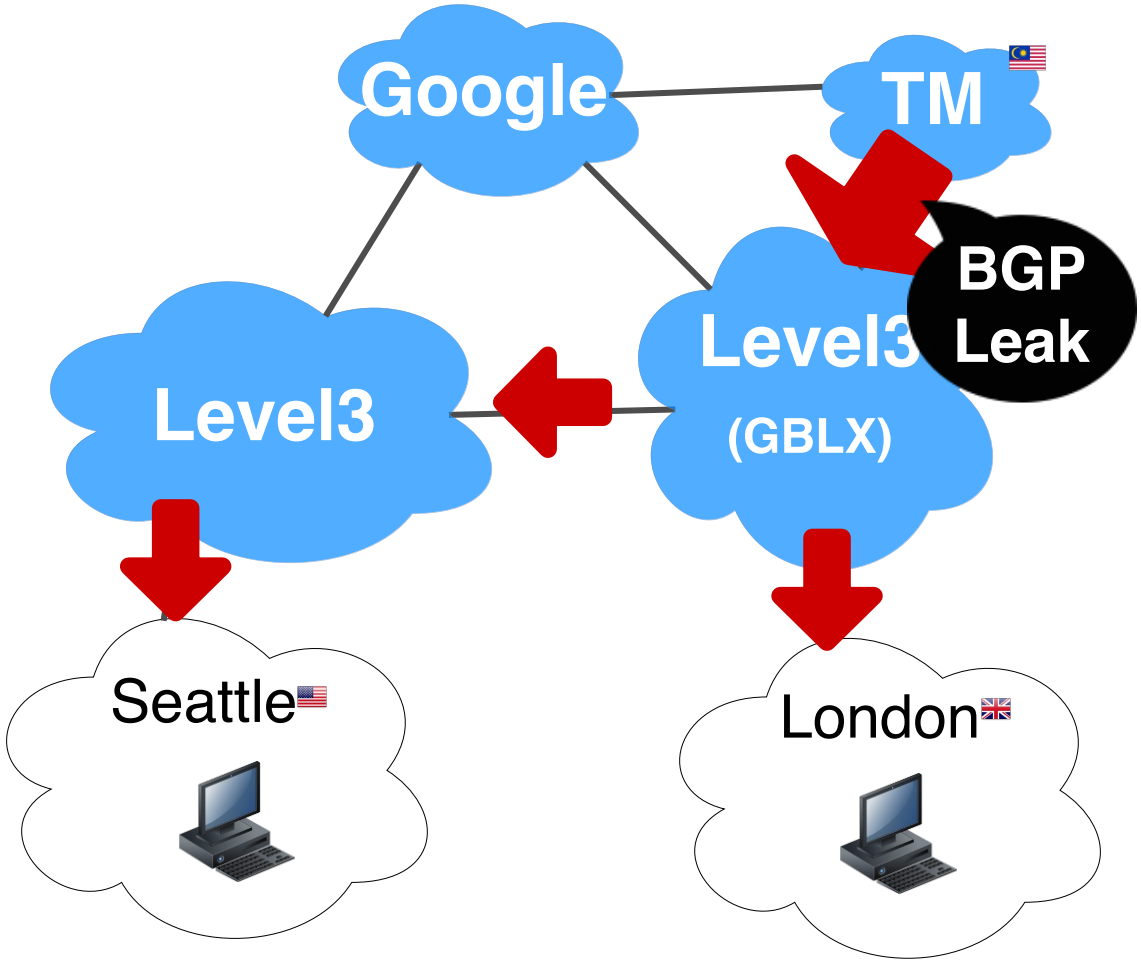

Study case: Telekom Malaysia BGP leak

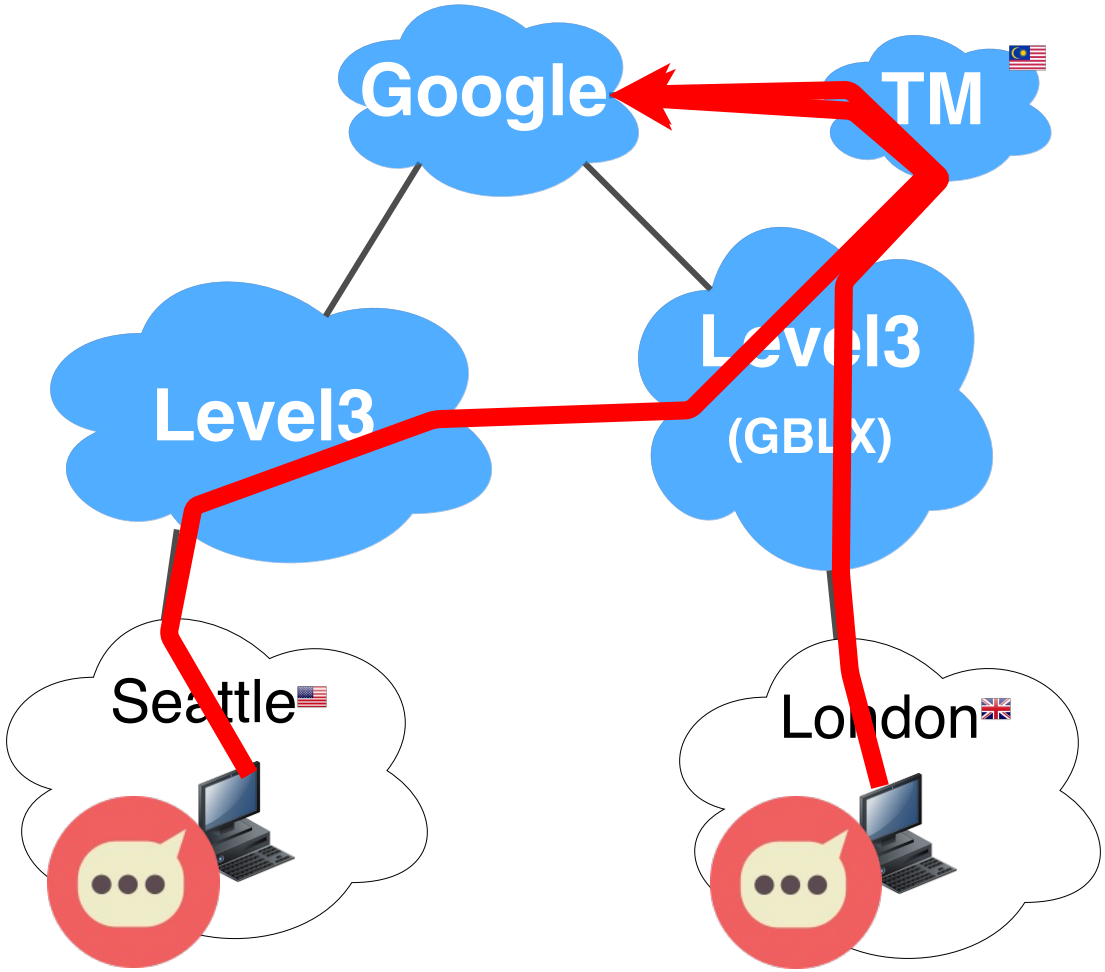

Not only with Google... but about **170k prefixes!**

Congestion in Level3

Rerouted traffic has congested Level3 (120 reported links)

- Example: 229ms increase between two routers in London!

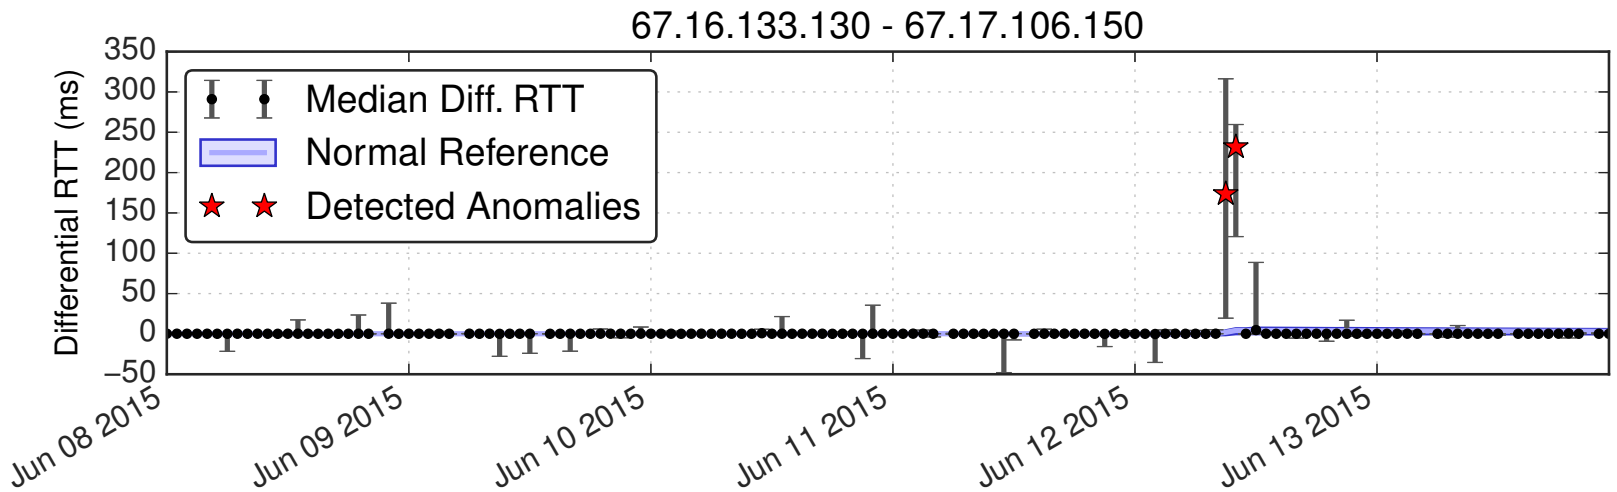

Congestion in Level3

Reported links in London:

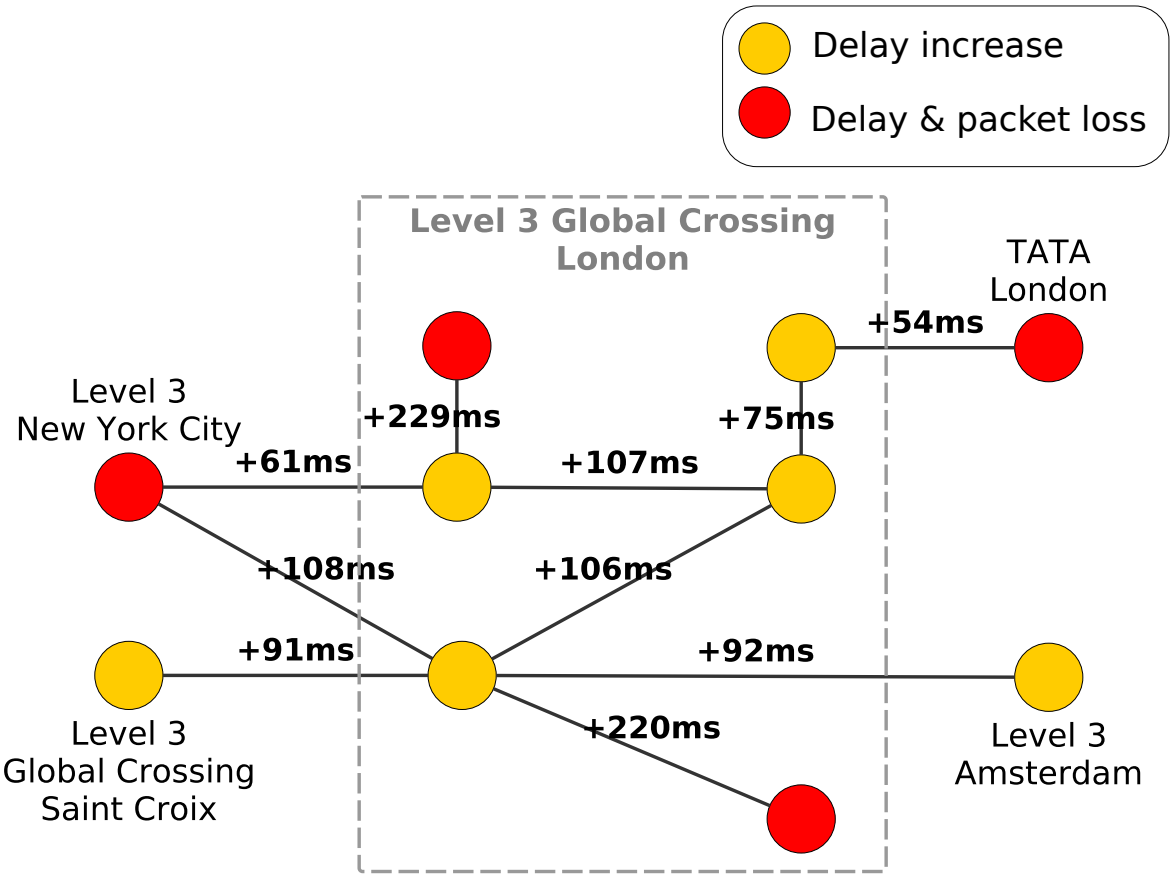

→ Traffic staying within UK/Europe may also be altered

Summary

Monitor delays with the Atlas platform

- Billions of (noisy) traceroutes

Detect and locate Internet congestion

- Robust statistical analysis
- Diverse root causes: remote attacks, routing anomalies, etc...
- Give a lot of new insights on reported events

On going work with RIPE NCC:

- Online detection and reports for network operators

References: `http://romain.iijlab.net/ihr/`