FCC Headquarters
26 June, 2013 - Washington DC, USA

Monitoring Large-scale Internet Outages

Alberto Dainotti, Emile Aben*, Alistair King,
Karyn Benson, Young Hyun, KC Claffy
alberto@caida.org

Cooperative Association for Internet Data Analysis
University of California, San Diego
*RIPE NCC

ANALYSIS OF INTERNET OUTAGES

Combining different measurement sources

- BGP
  - BGP updates from route collectors of ***RIPE-NCC RIS*** and ***RouteViews***
- Active Traceroute Probing
  - Archipelago Measurement Infrastructure (***ARK***)
  - ***RIPE-NCC Atlas***
- Internet Background Radiation (IBR)
  - Traffic reaching the ***UCSD Network Telescope***
- *more data sources to come...*

CASE STUDIES

Different for causes/tech implications/impact

- Country-level Internet Blackouts *(BGP withdrawals, packet-filtering, satellite-signal jamming, ...)*

EGYPT, JAN 2011
GOVERNMENT ORDERS TO SHUT DOWN THE INTERNET

- Natural disasters affecting the infrastructure/population (earthquakes, hurricanes, ...)

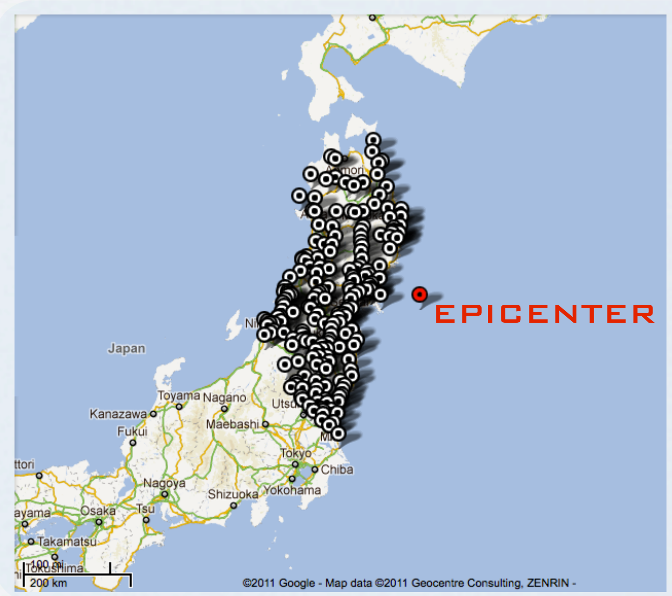

JAPAN, MAR 2011
EARTHQUAKE OF MAGNITUDE 9.0

THE EVENTS (1/3)

Internet Disruptions in North Africa

- Egypt
  - *January 25th, 2011*: protests start in the country
  - The government orders service providers to "shut down" the Internet
  - ***January 27th, around 22:34 UTC***: several sources report the withdrawal in the Internet's global routing table of almost all routes to Egyptian networks
  - The disruption lasts ***5.5 days***

- Libya
  - *February 17th, 2011*: protests start in the country
  - The government controls most of the country's communication infrastructure
  - ***February 18th (6.8 hrs), 19th (8.3 hrs), March 3rd (3.7 days):*** three different connectivity disruptions:

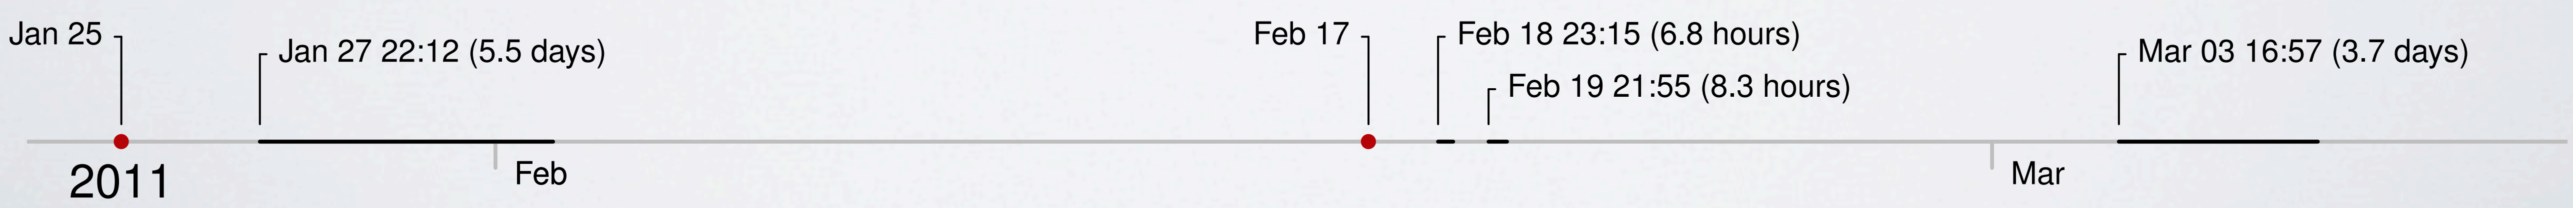

NETWORK INFO

Prefixes, ASes, Filtering

- Egypt
  - 3165 *IPv4* and 6 *IPv6* prefixes are delegated to Egypt by AfriNIC
  - They are managed by 51 Autonomous Systems
  - Filtering type: BGP only

- Libya
  - 13 *IPv4* prefixes, no *IPv6* prefixes
  - 3 Autonomous Systems operate in the country
  - Filtering type: mix of BGP, packet filtering, satellite signal jamming

A. Dainotti, C. Squarcella, E. Aben, K. C. Claffy, M. Chiesa, M. Russo, A. Pescapè, "Analysis of Country-wide Internet Outages Caused by Censorship" ACM SIGCOMM Internet Measurement Conference 2011

BGP

prefix reachability

- We reconstruct prefixes losing and regaining reachability
  - we build the routing history of every collector's peer for each collector
  - using both RIBs and UPDATES
  - we mark a prefix as disappeared if it is withdrawn in each routing history

Egyptian disconnection and reconnection <u>NOTE</u>: IPv6 routes stayed up!

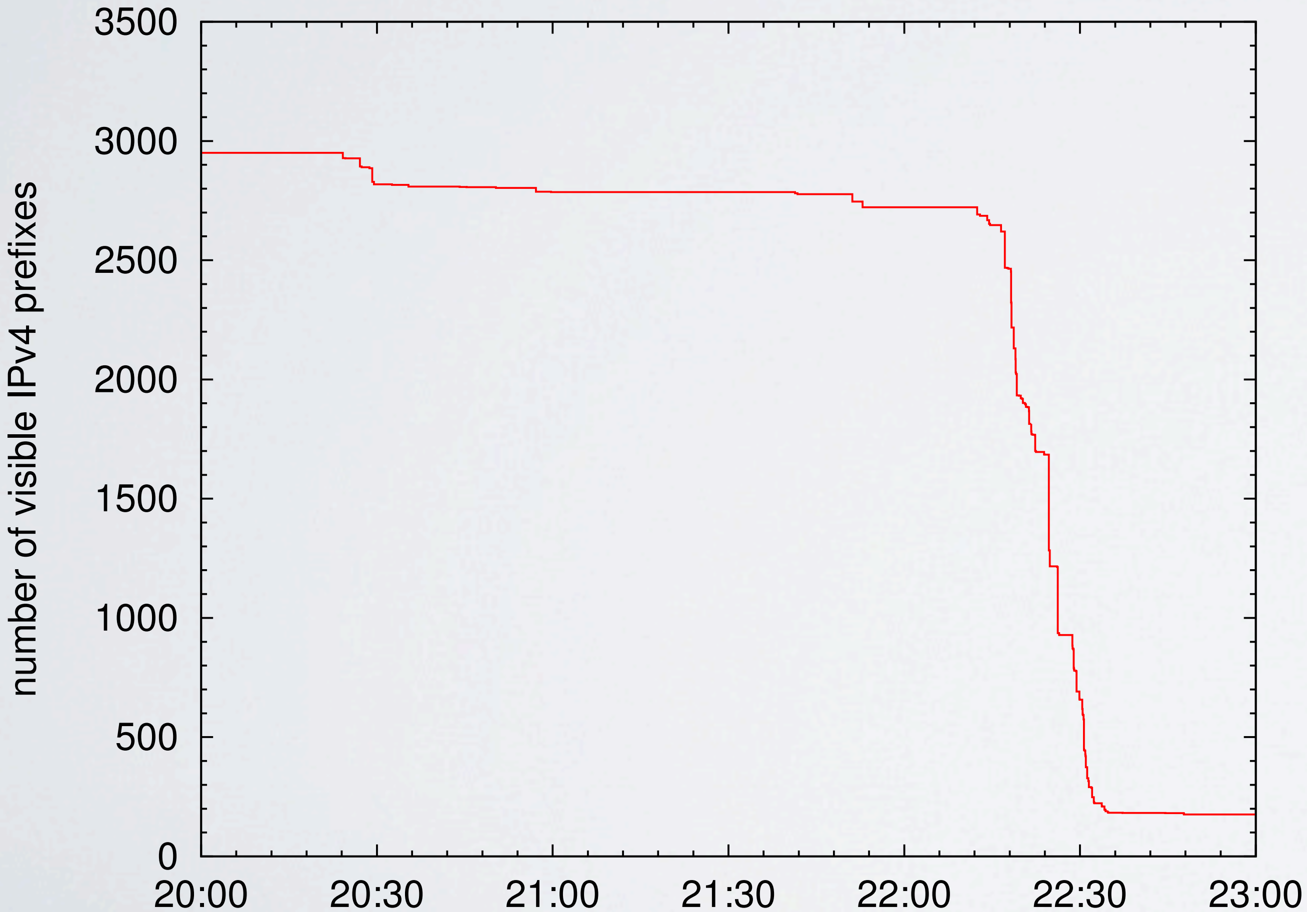

January 27th

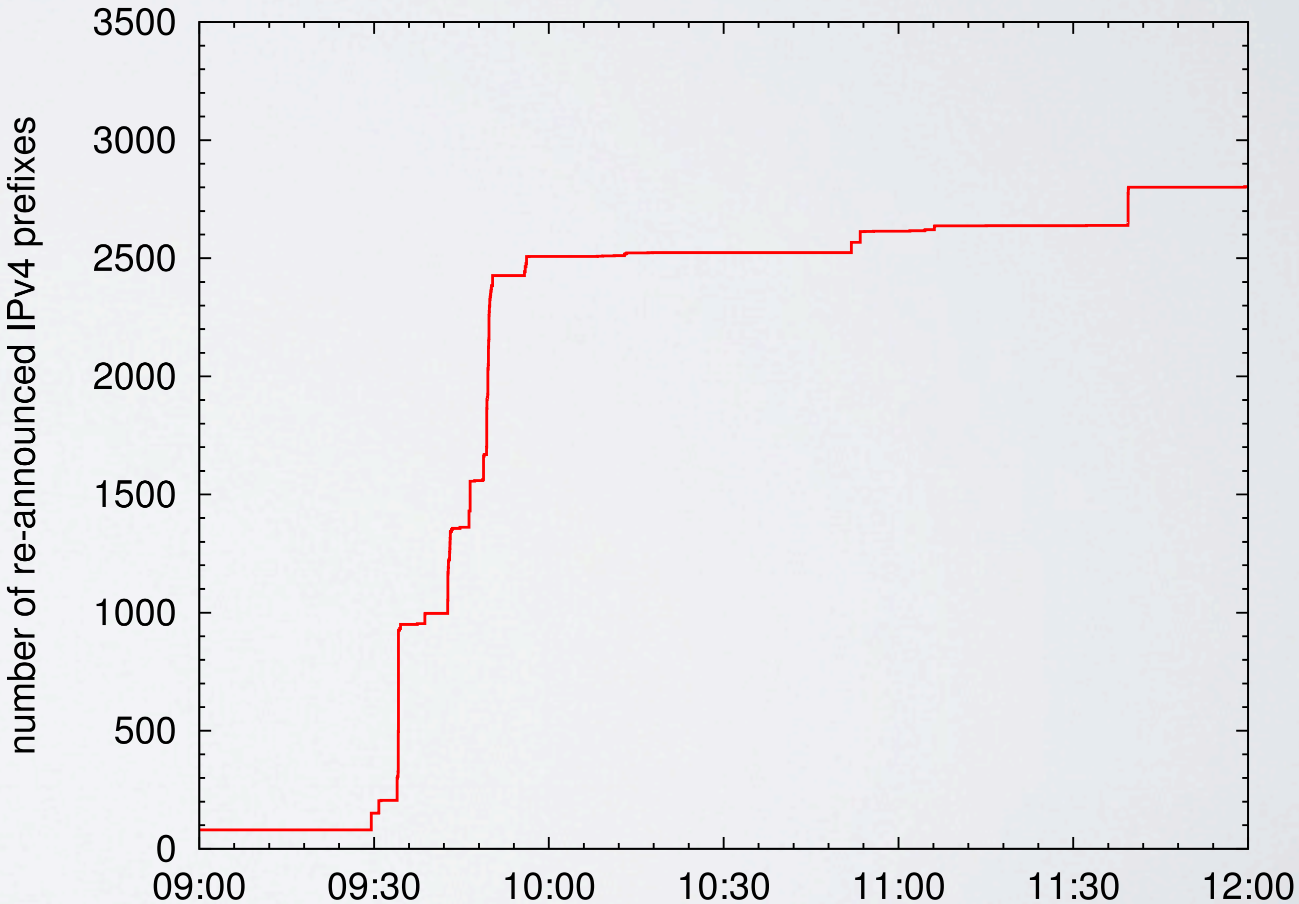

February 2nd

BGP

per-AS analysis

- A detailed analysis shows there is synchronization among ASes

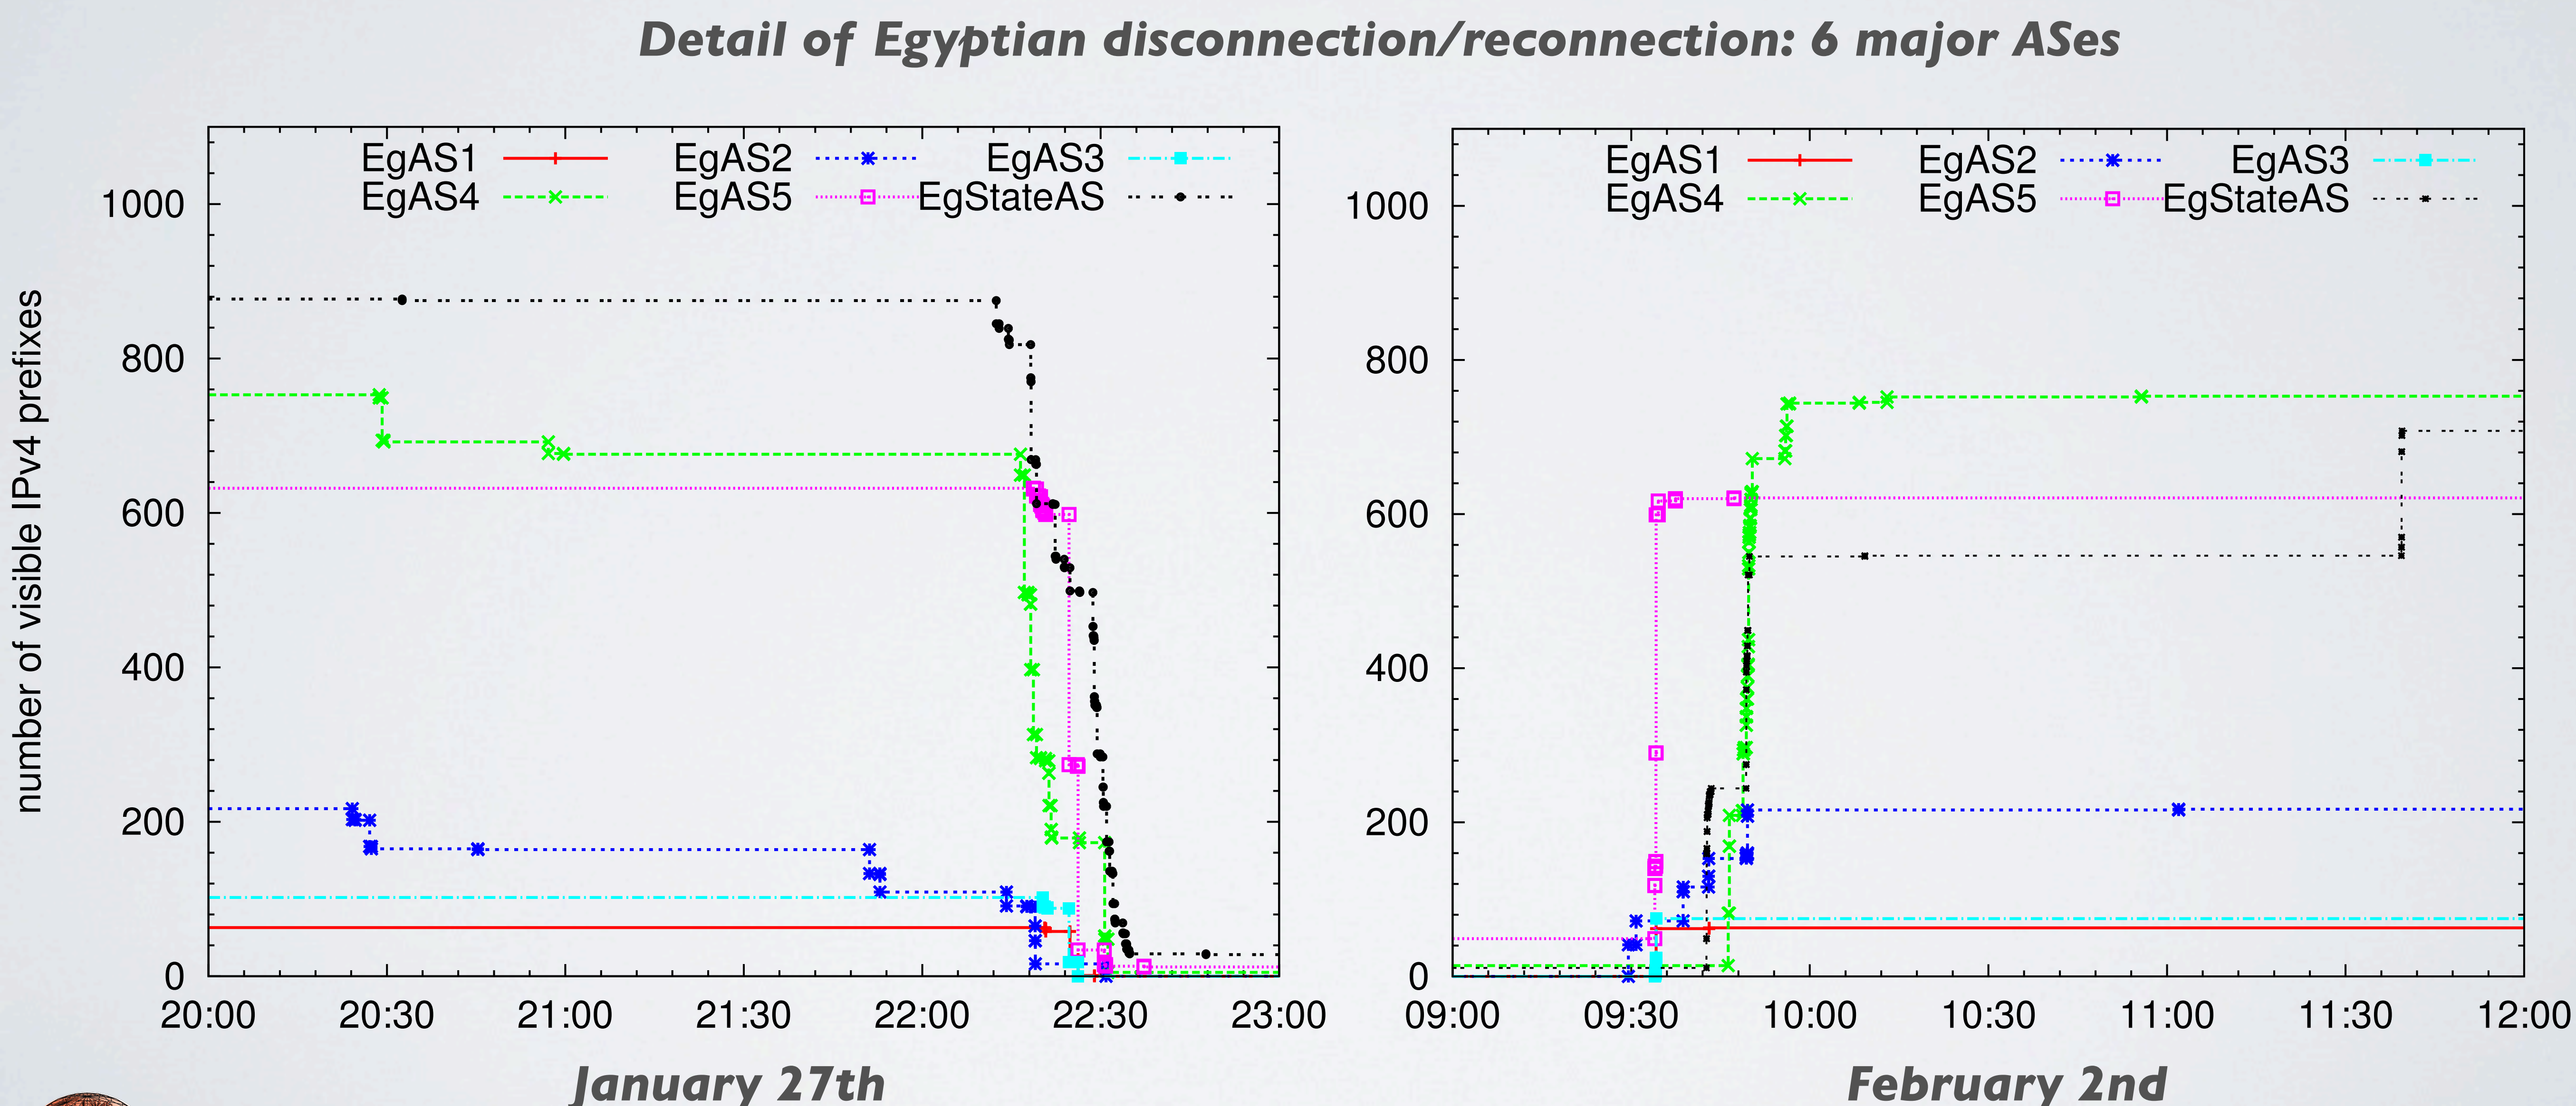

IBR

"Extracting benefit from harm.."

- Use *Internet Background Radiation (IBR) generated by malware-infected hosts as a "signal"*

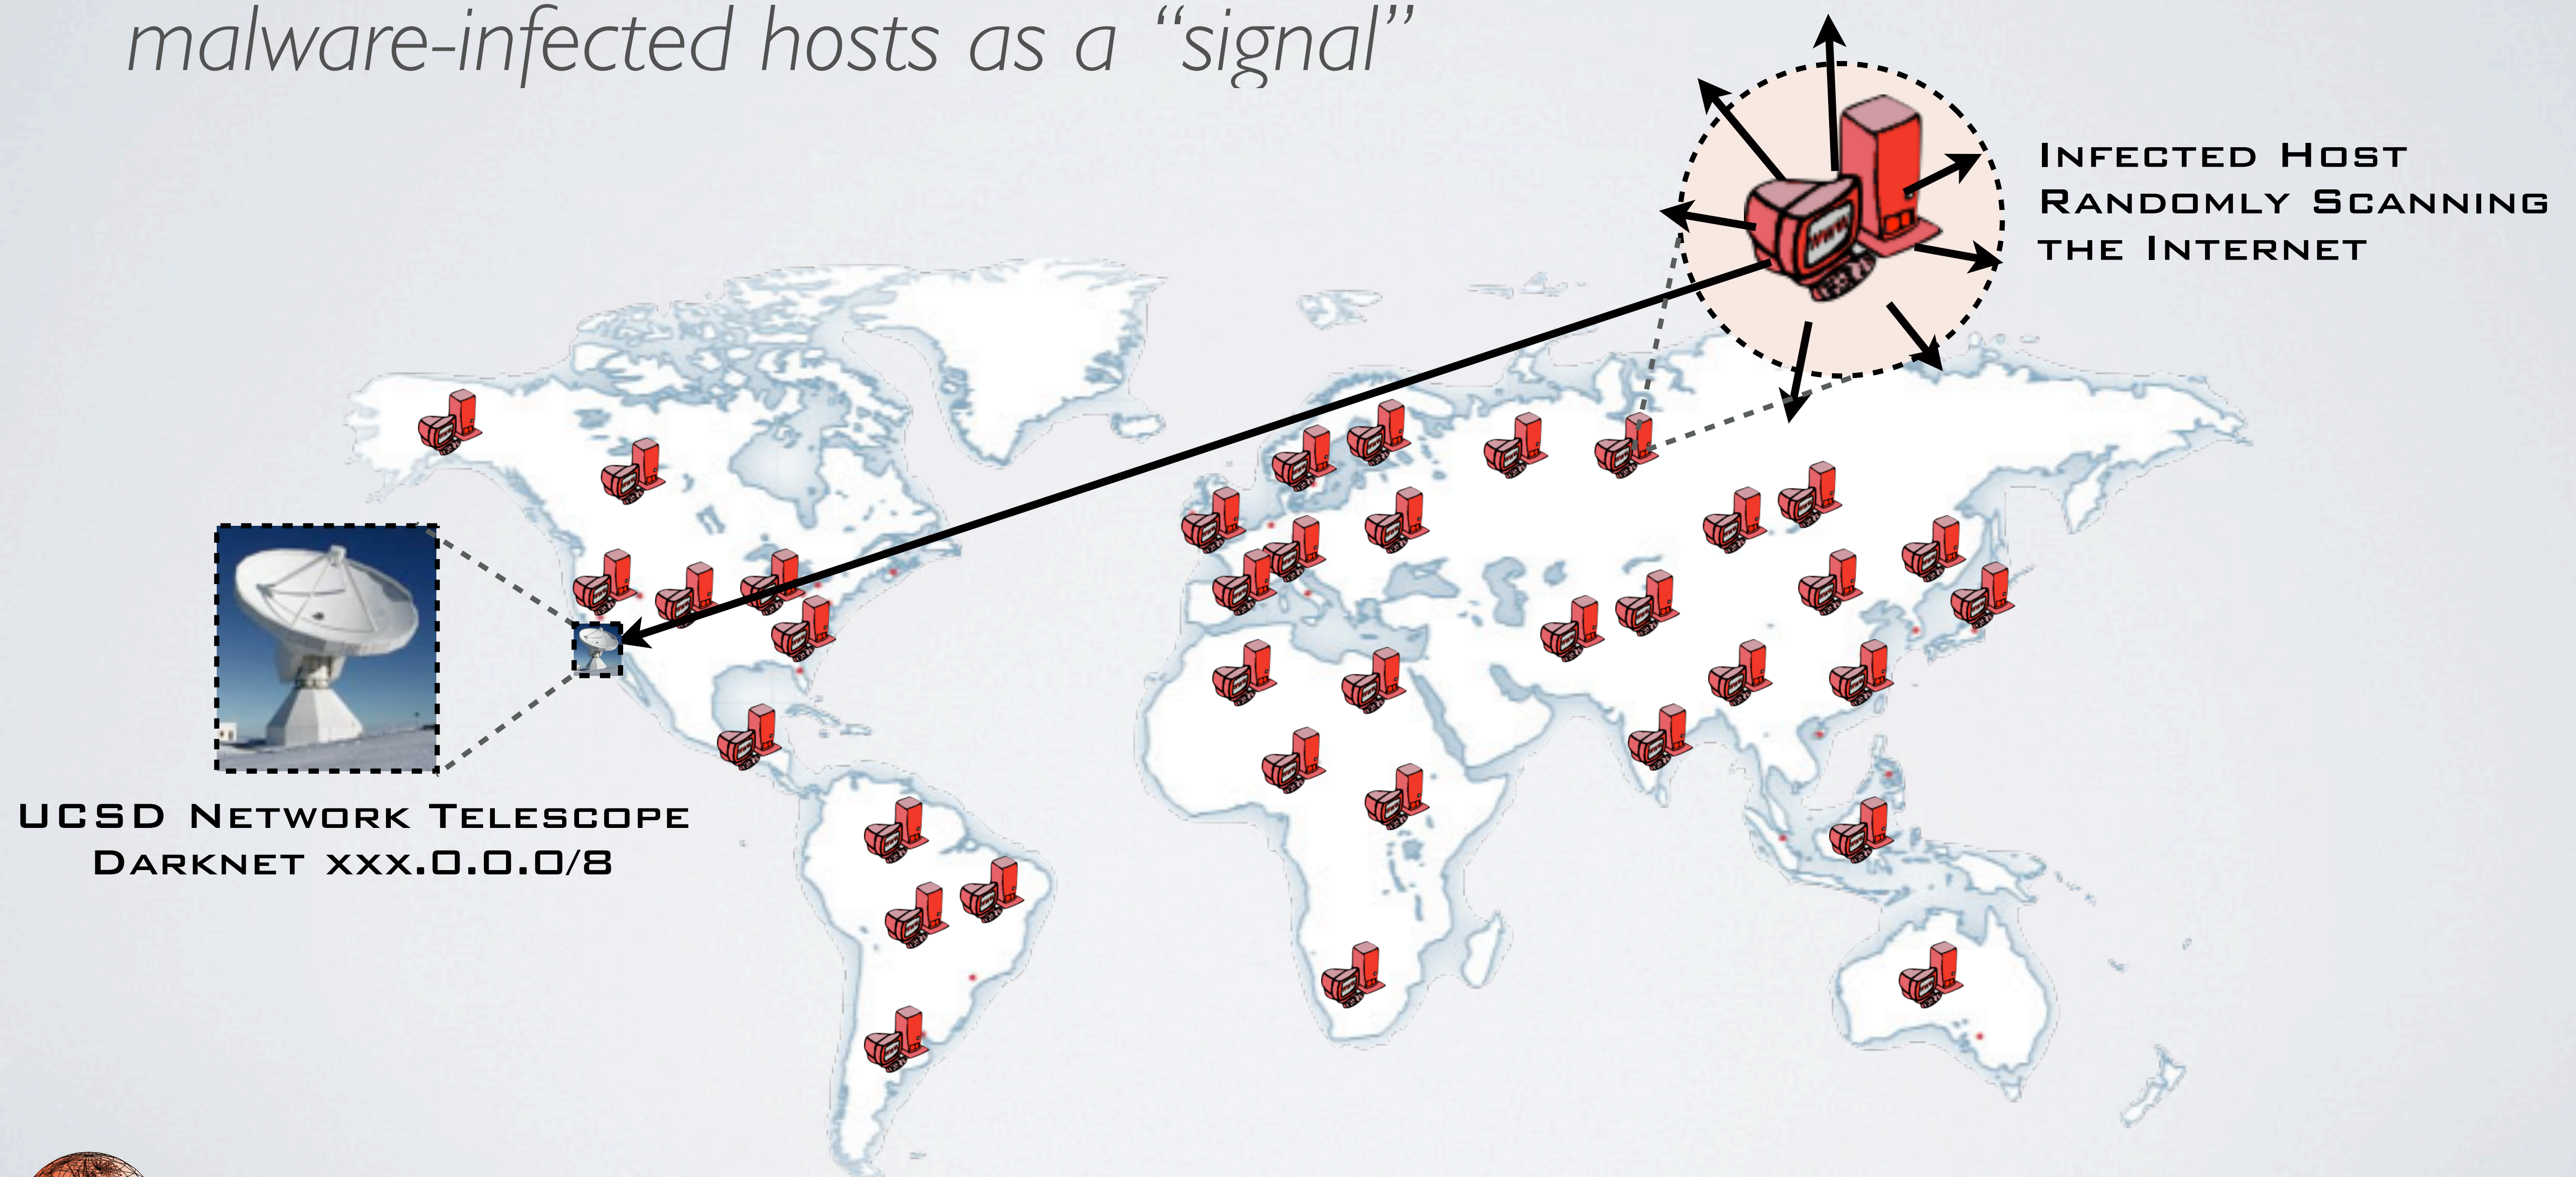

UCSD TELESCOPE

when malware helps..

- Unsolicited traffic, *a.k.a. Internet Background Radiation* - e.g. scanning from conficker-infected hosts - from the observed country reveals several aspects of these outages!

Egypt

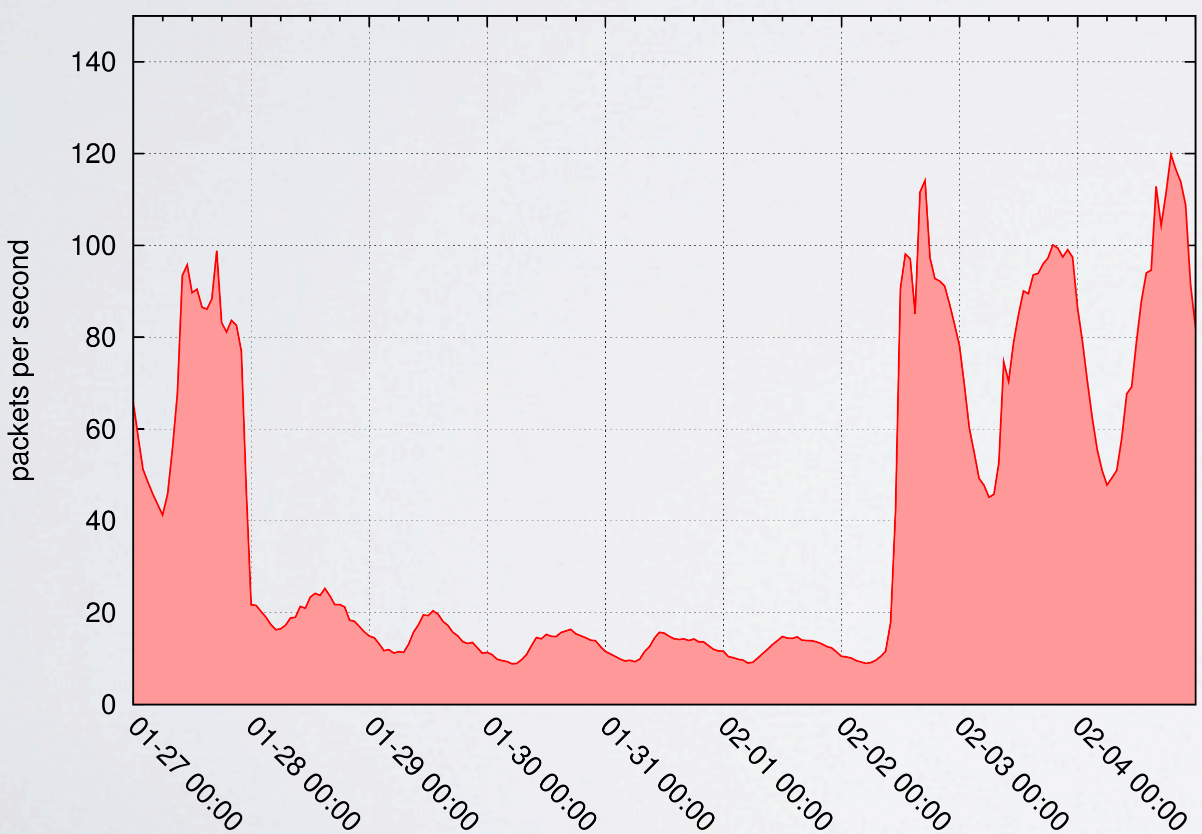

Libya

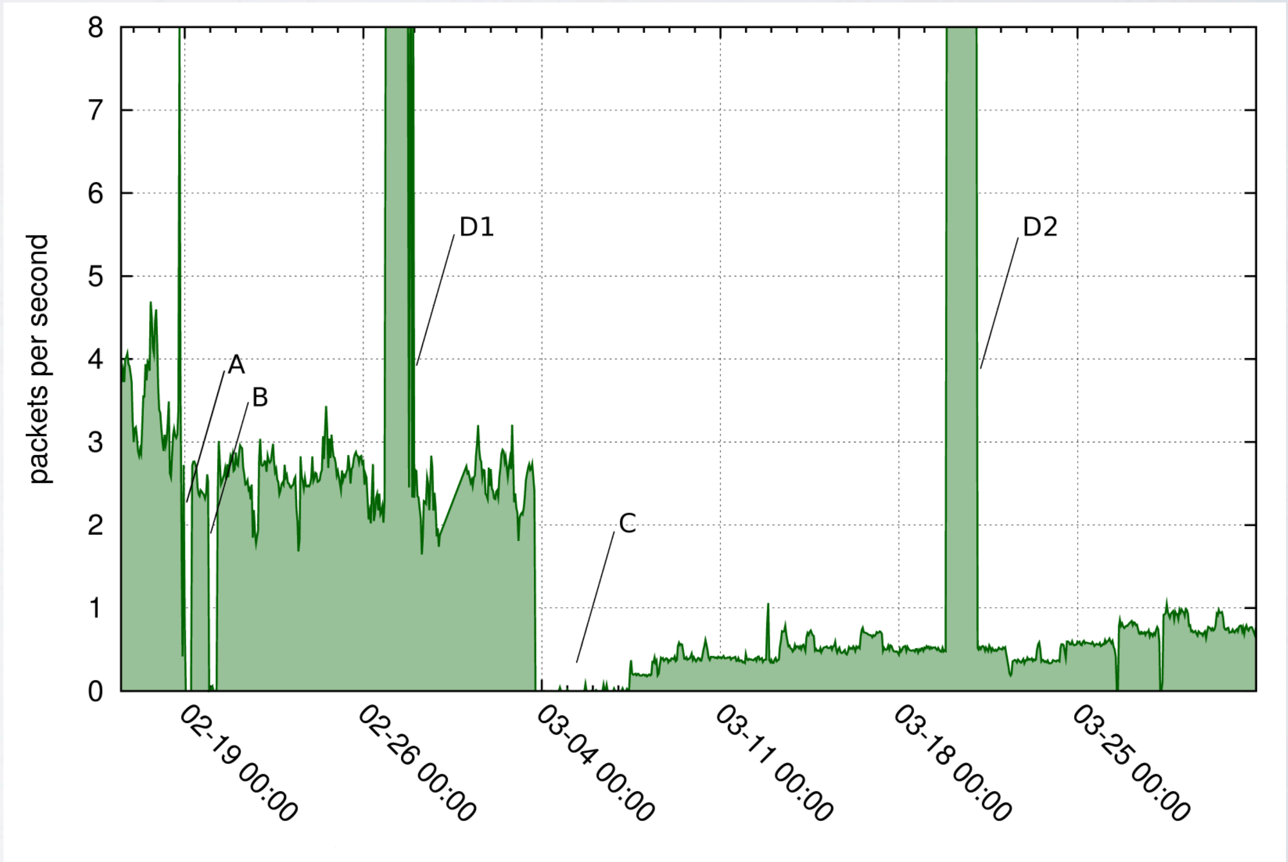

A,B,C: Outages
D1, D2: Denial of Service attacks

RANDOM PROBING

E.g., Conficker

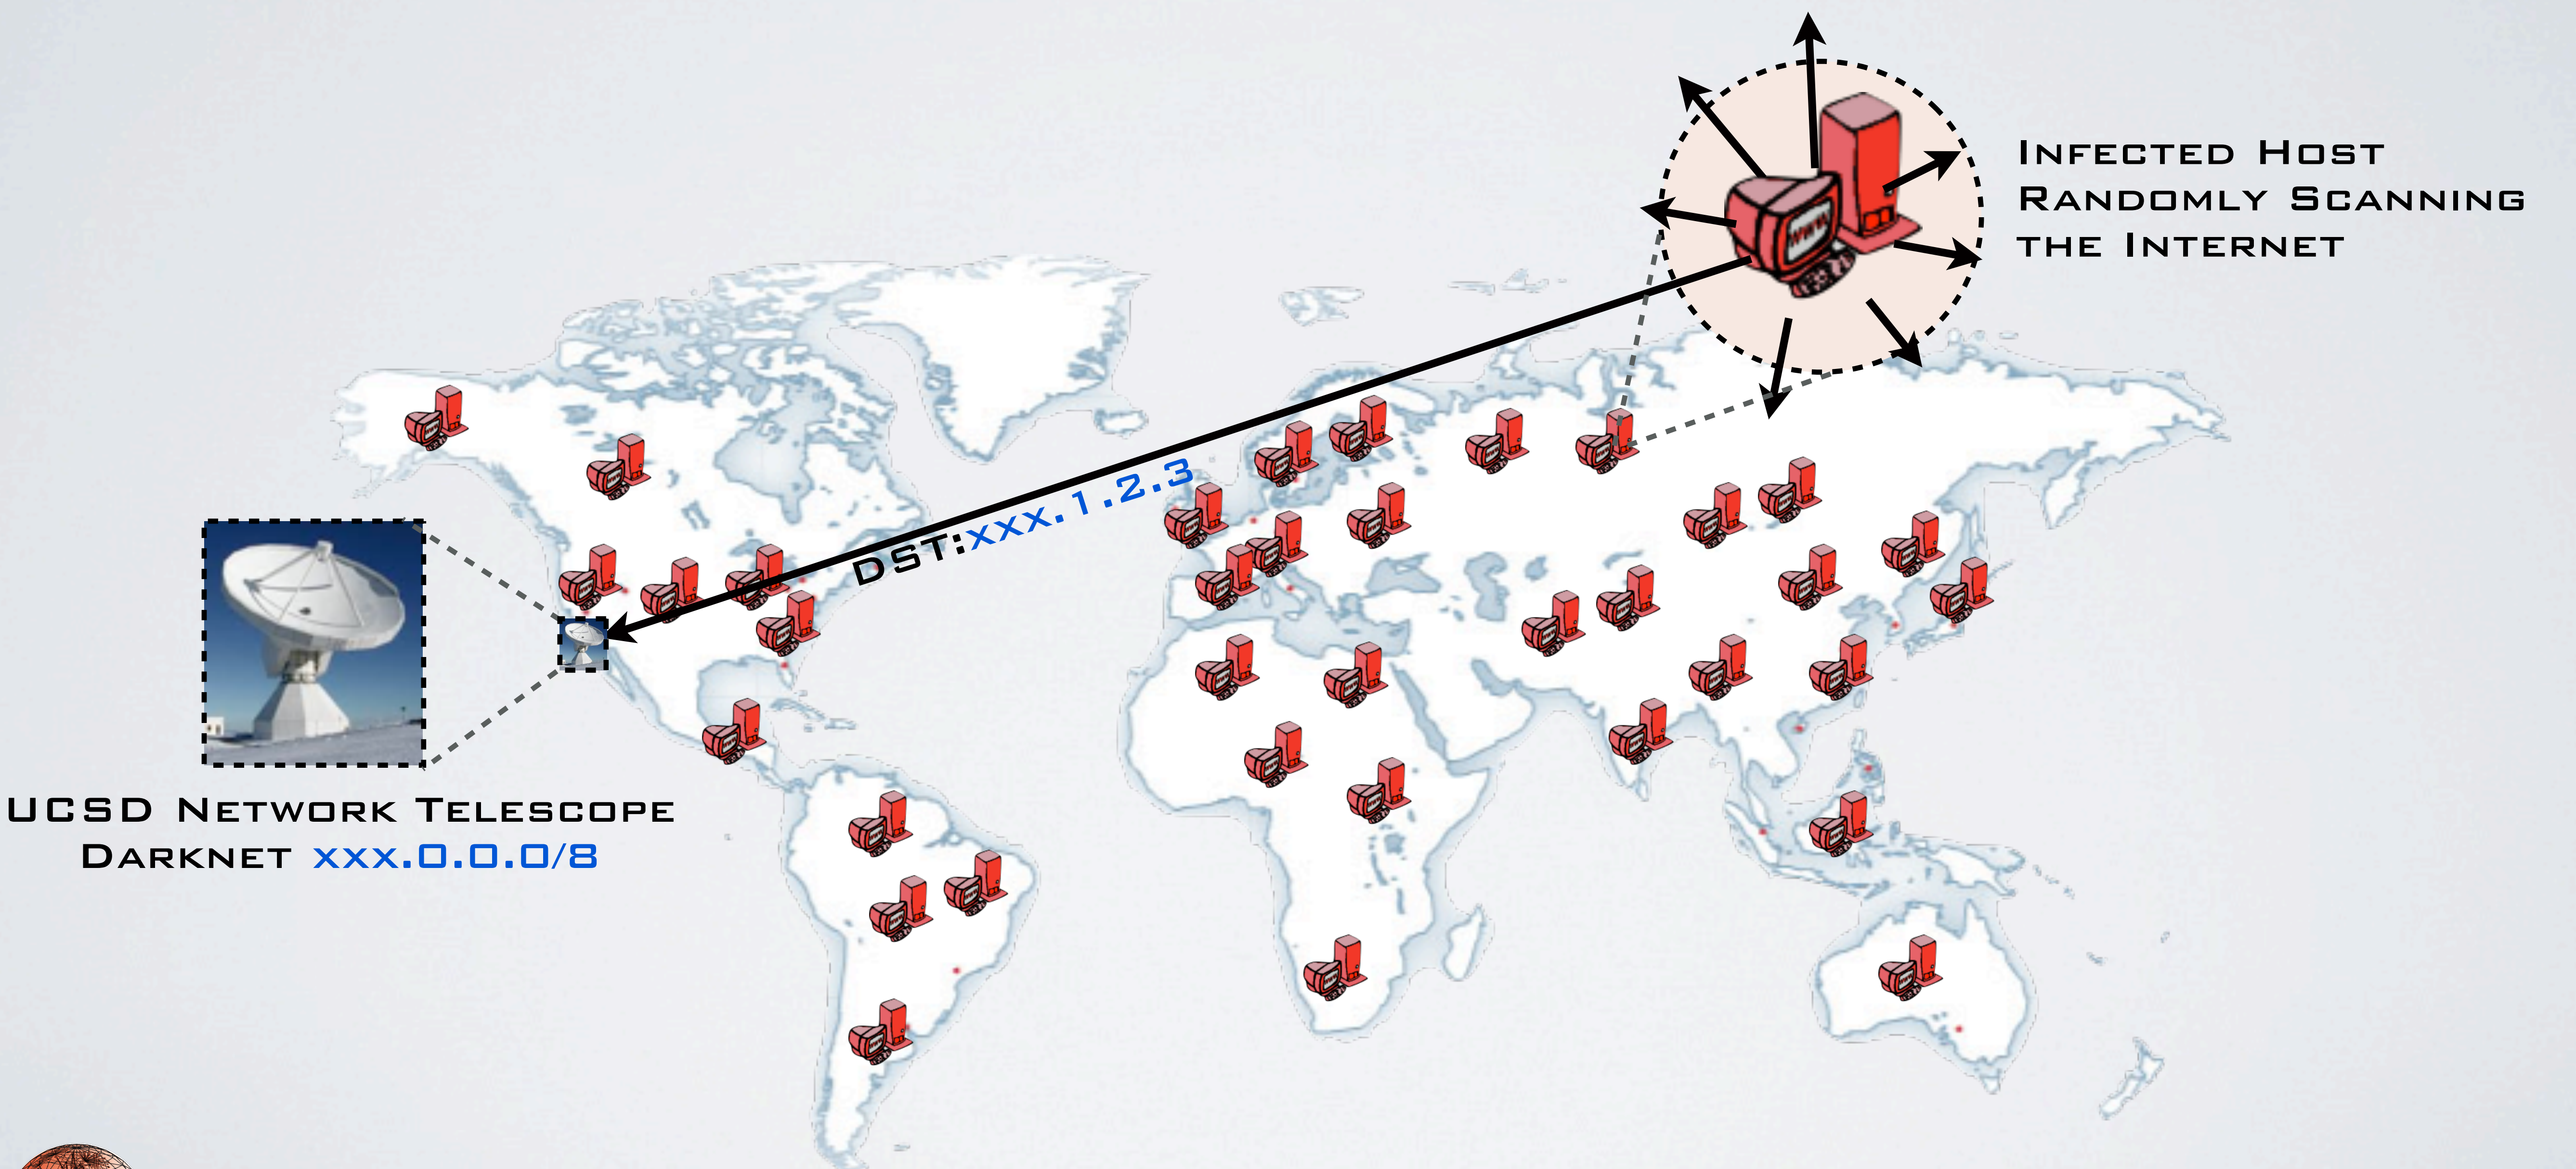

BACKSCATTER

e.g., SYN+ACK replies to spoofed SYNs

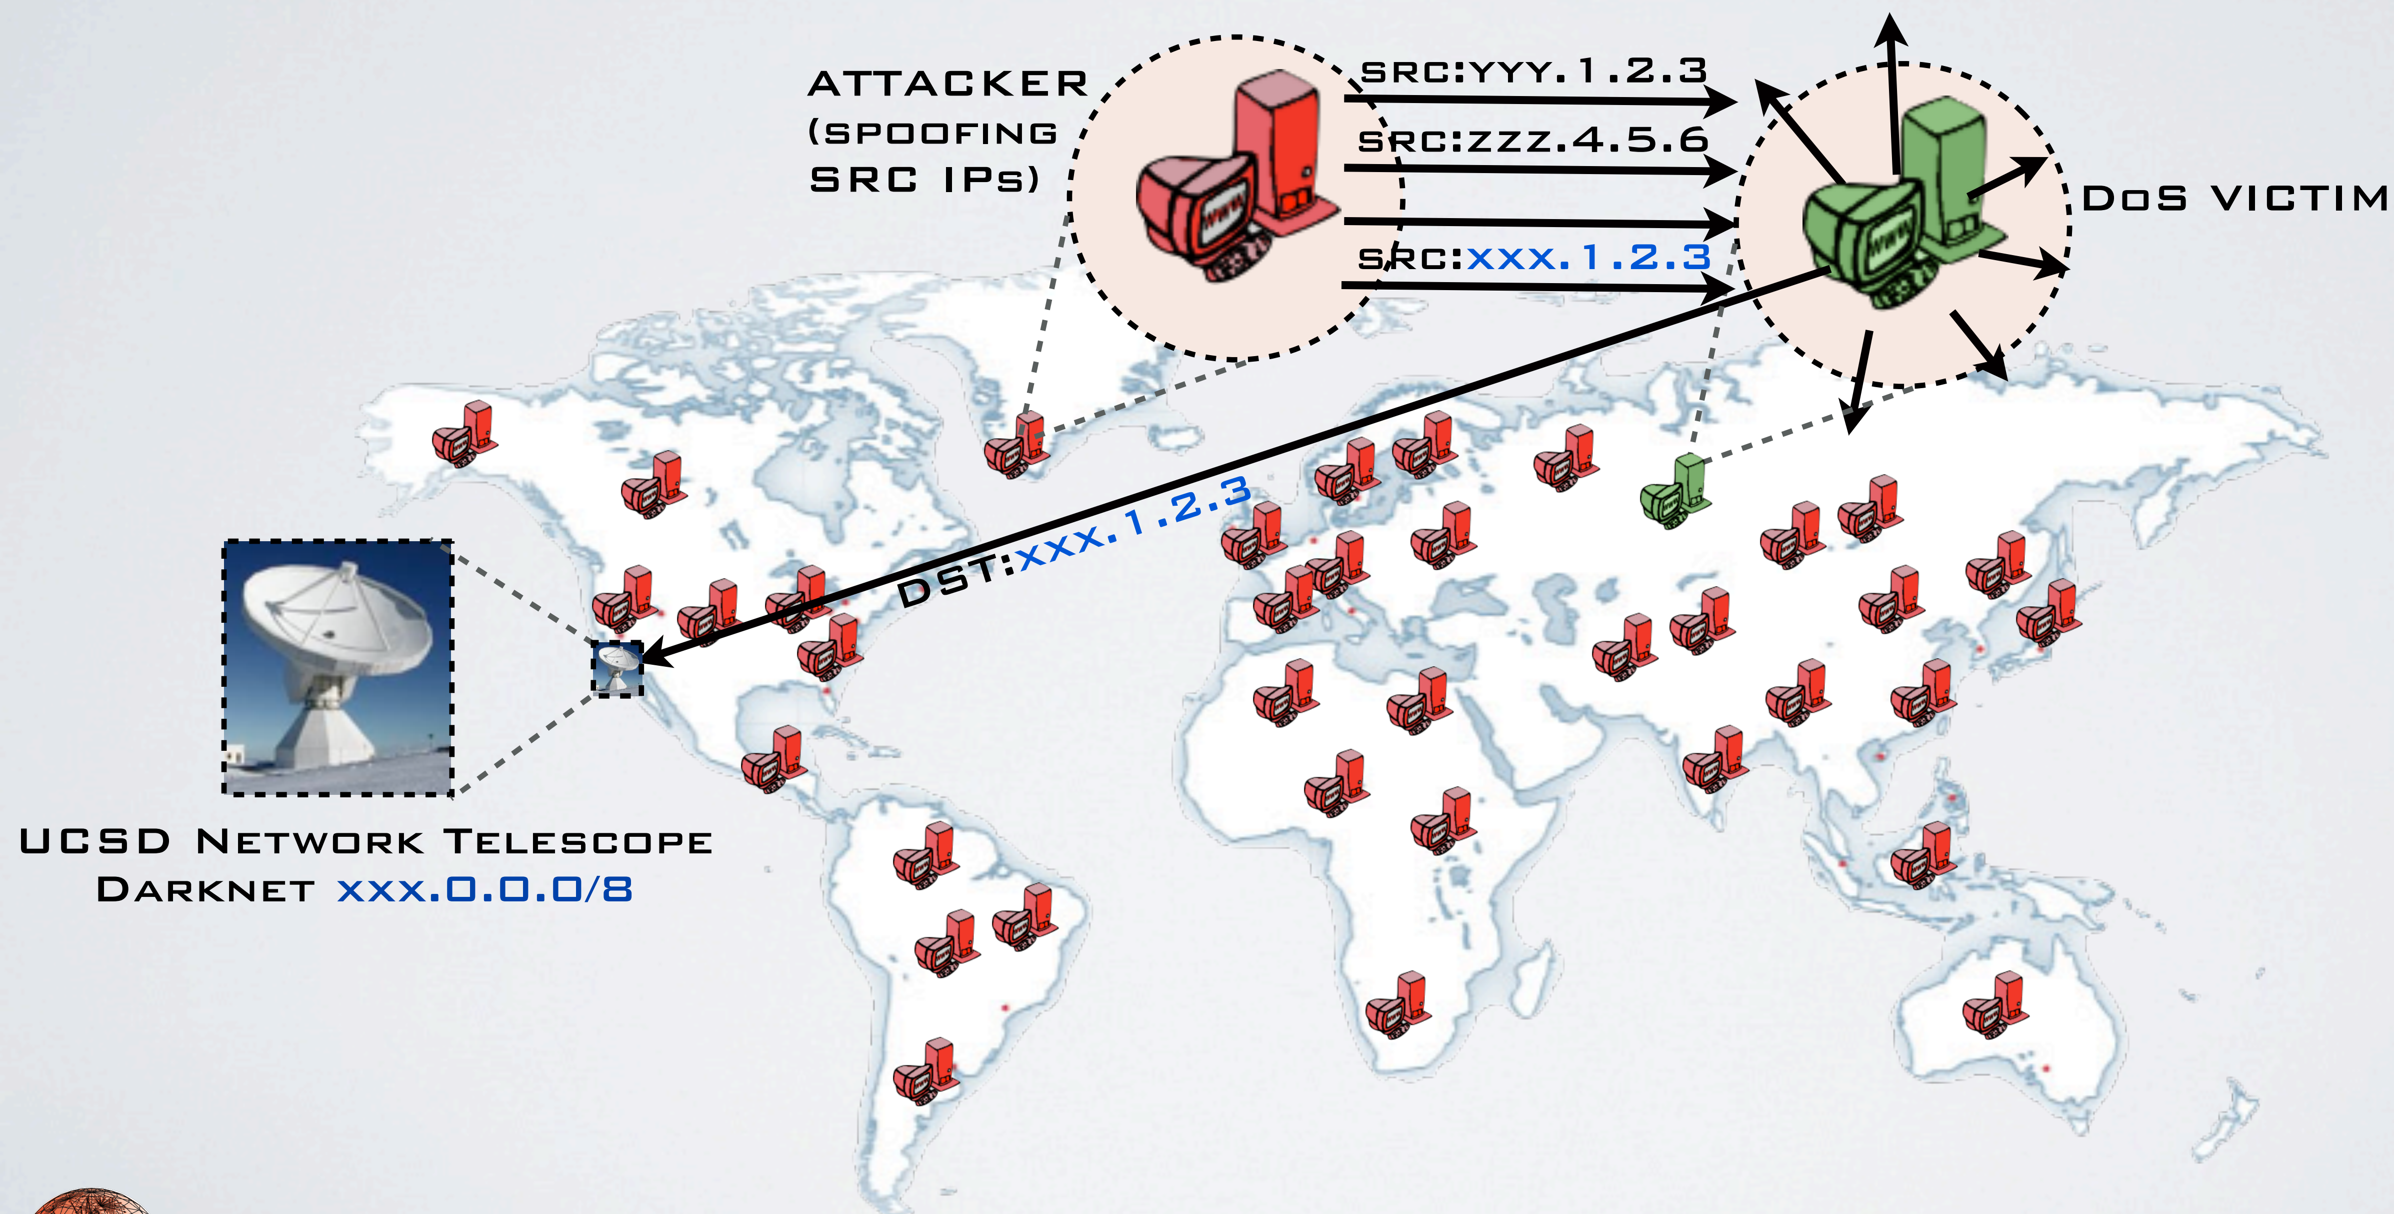

Cooperative Association for Internet Data Analysis
University of California San Diego

EGYPT

IBR: dissecting it

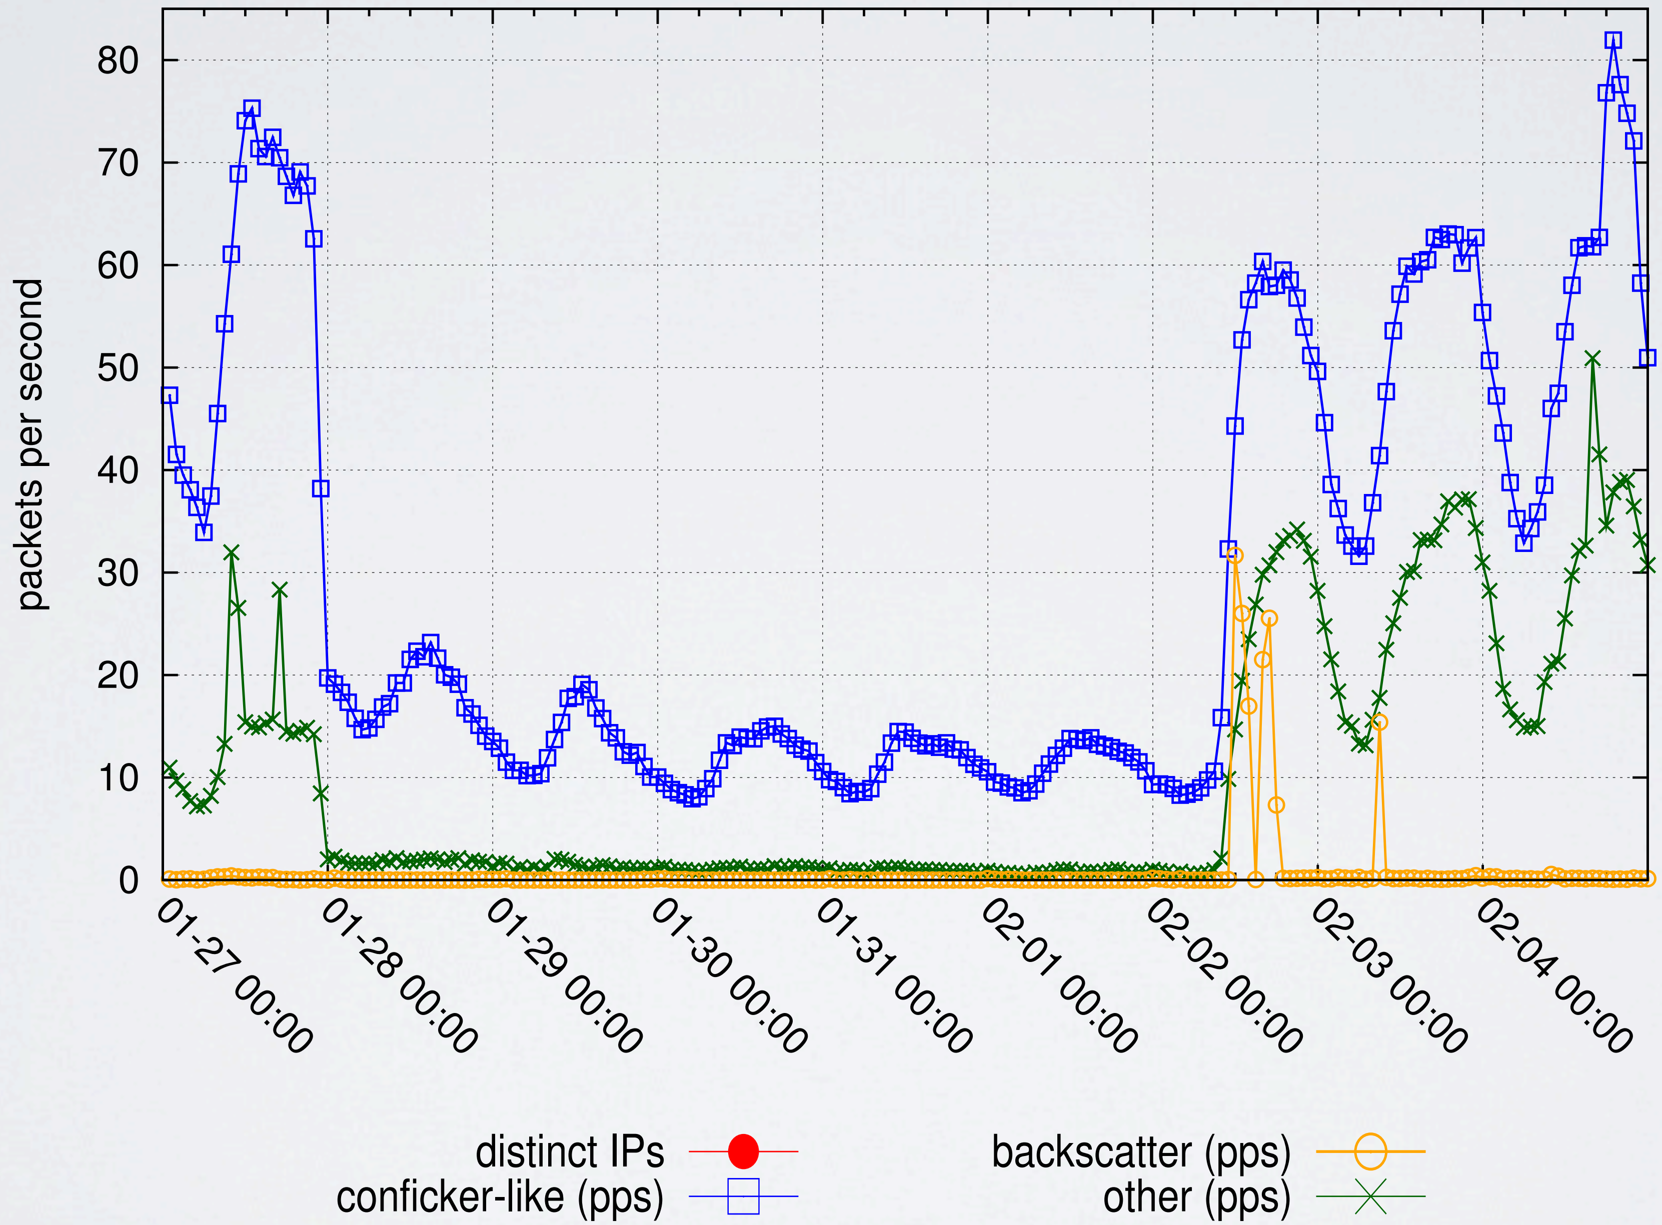

EGYPT

IBR: rate of distinct src IPs vs packet rate

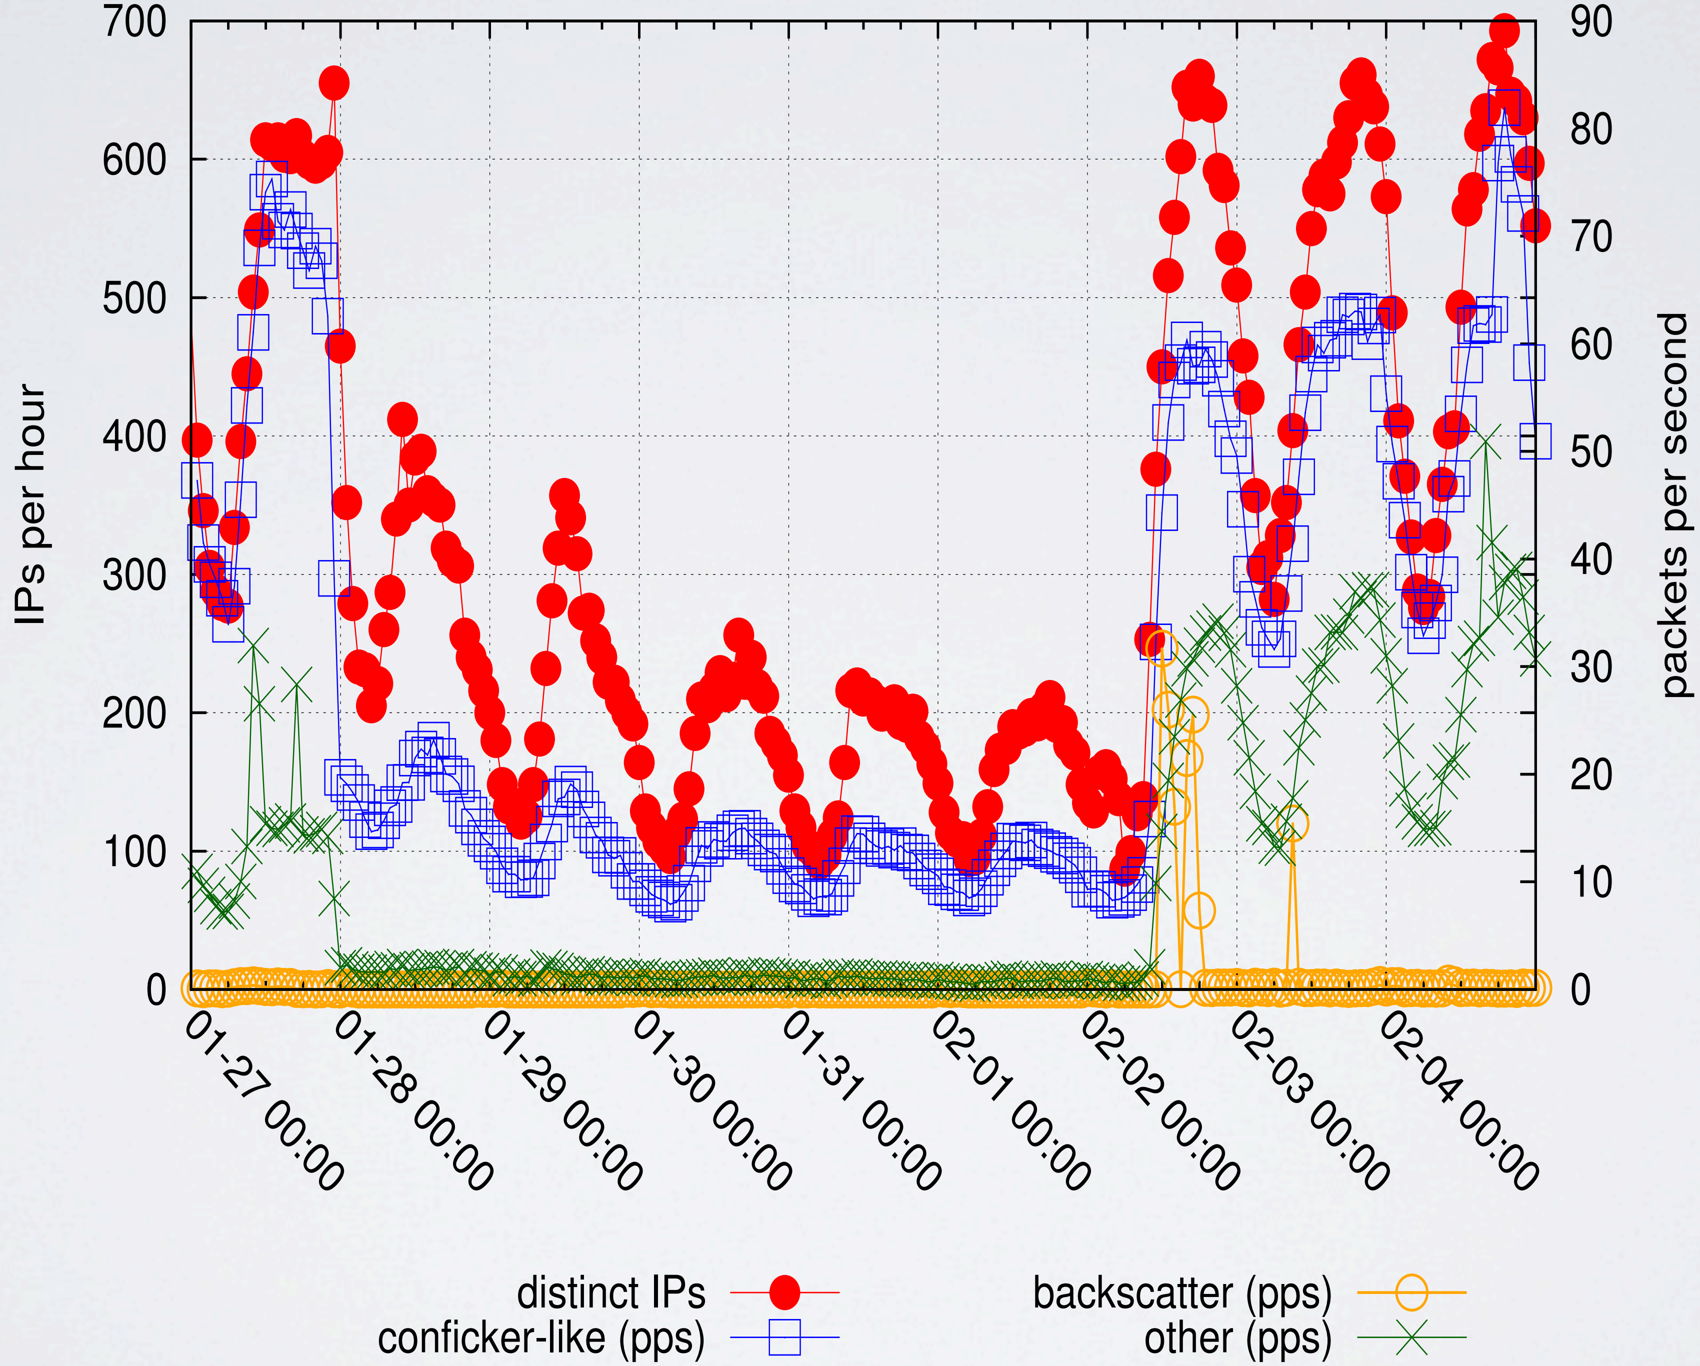

Cooperative Association for Internet Data Analysis
University of California San Diego

TELESCOPE *VS* BGP

Consistency

- The sample case of *EgAS7*shows the consistency between telescope traffic and BGP measurements

Egypt: disconnection of EgAS7

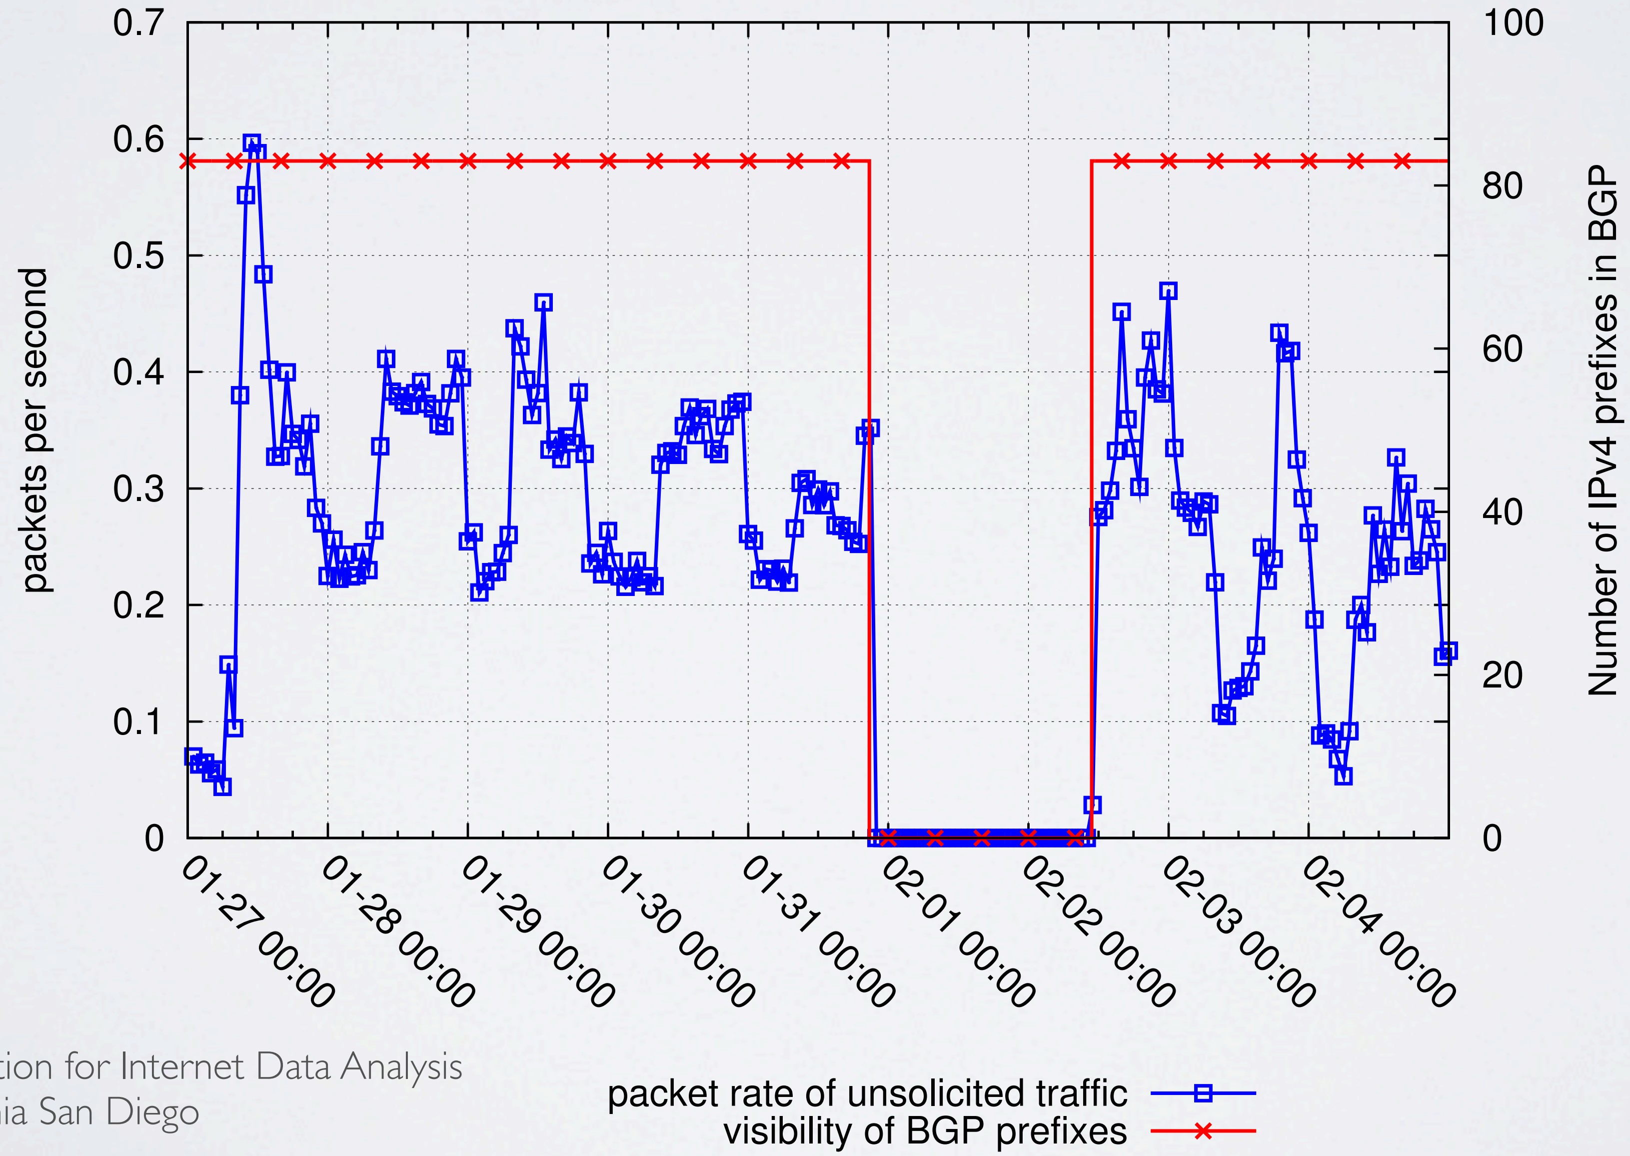

TELESCOPE vs BGP

Complementarity

- Contrasting telescope traffic with BGP measurements revealed a mix of blocking techniques that was not publicized by others
- The second Libyan outage involved overlapping of ***BGP withdrawals*** and ***packet filtering***

ACTIVE MEASUREMENTS

ARK + ATLAS

- CAIDA ARCHIPELAGO (*ARK*)

  \- Coordinate traceroute-based topology measurement probing the full routed IPv4 address space

  [***http://www.caida.org/projects/ark/***](http://www.caida.org/projects/ark/)

- RIPE ATLAS

  \- traceroutes/pings to fixed destinations

  \- user-defined measurements (a community-oriented tool)

  [***https://atlas.ripe.net/***](https://atlas.ripe.net/)

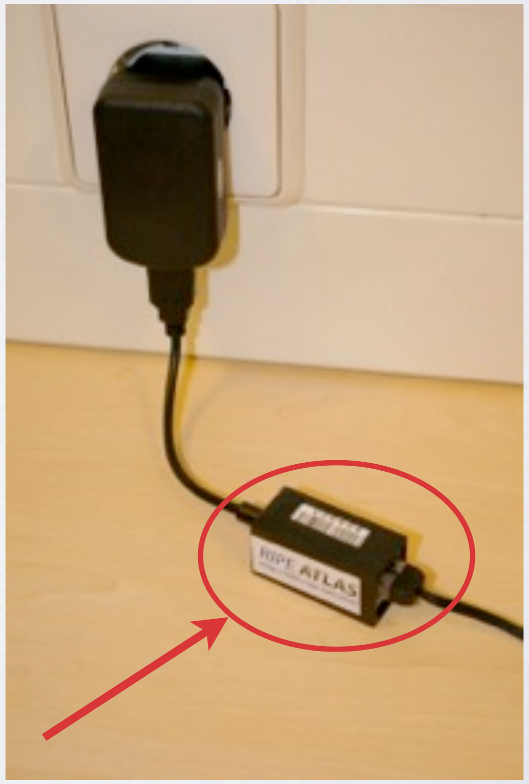

ARK

active measurements

- ARK active measurements are consistent with other sources
  - limitation due to frequency of probes and because they target random addresses
  - the first two Libyan outages are not visible
  - we used them only to test *reachability*, not to analyze topology

Egypt

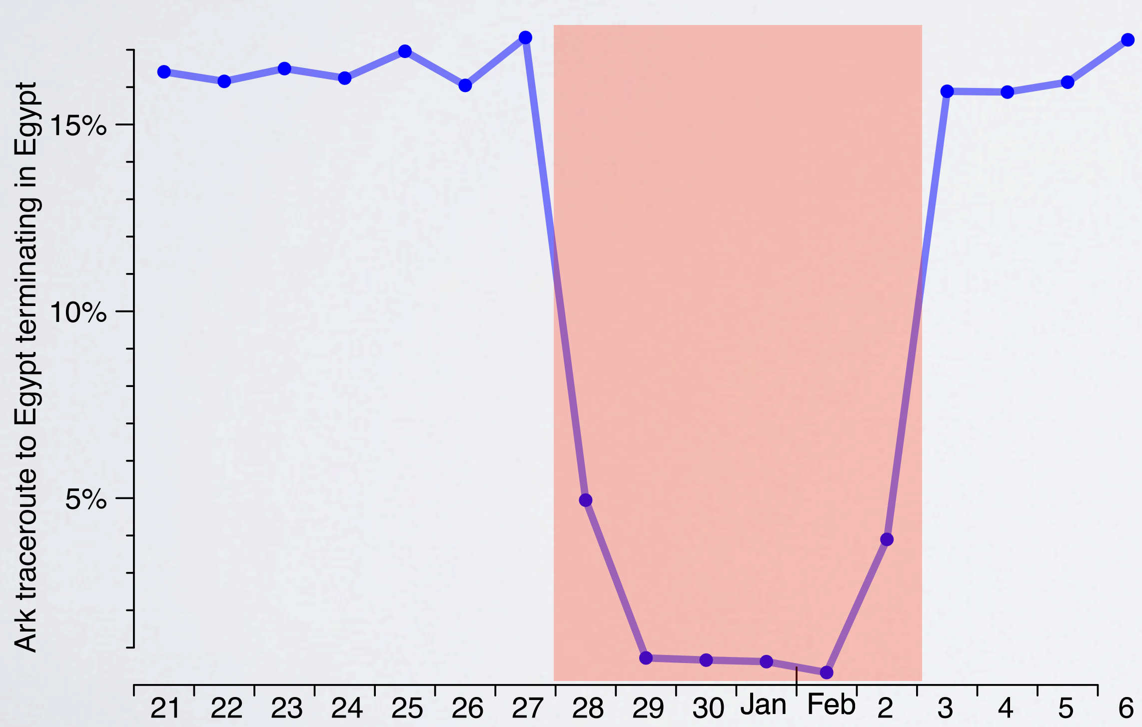

Libya

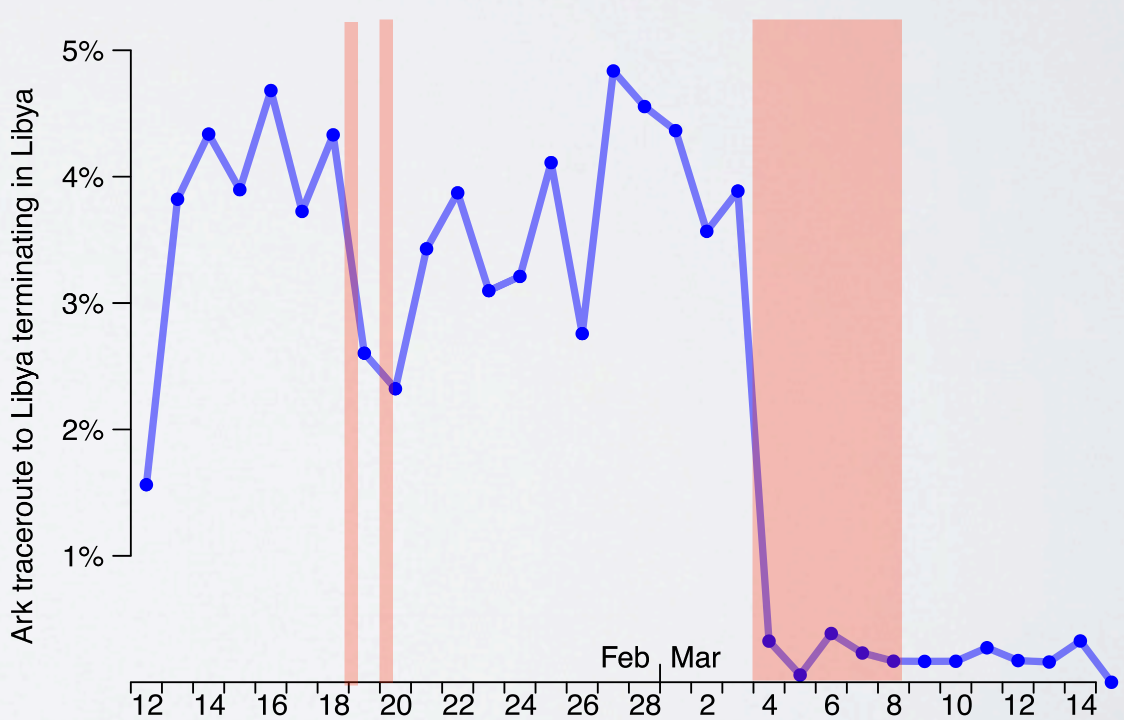

ARK

confirming telescope's findings

- Third Libyan outage: while BGP reachability was up, most of Libya was disconnected
  - ARK measurements confirmed the finding from the telescope
    1) disconnection
    2) identification of some reachable networks

  suggesting the use of packet filtering by the censors

Libya seen by ARK

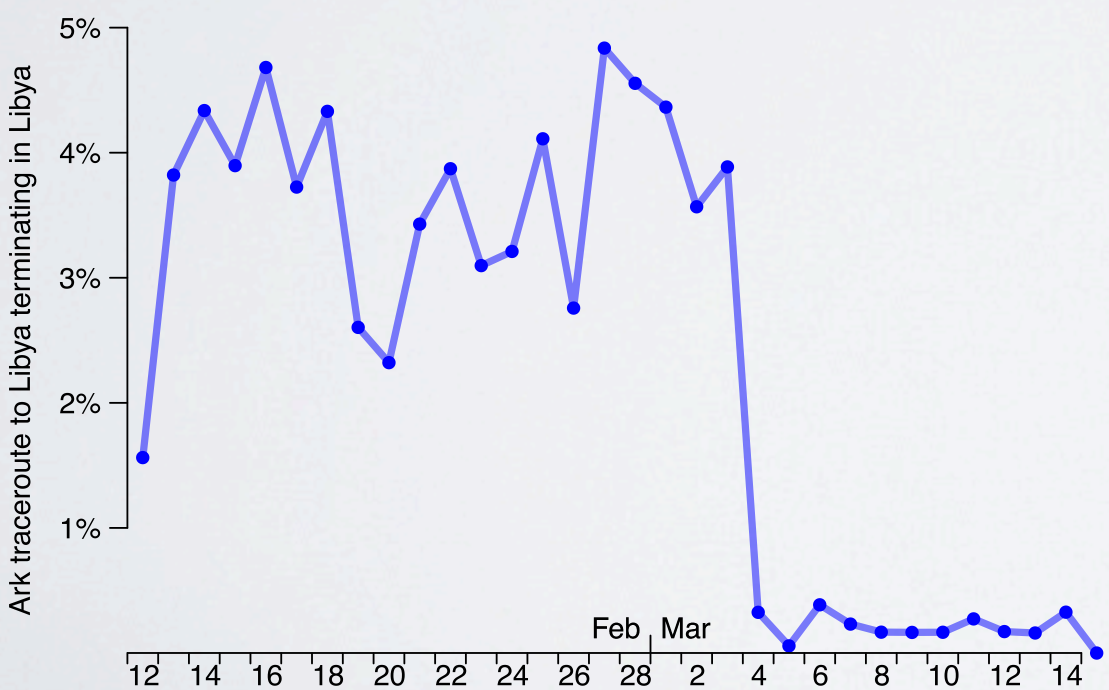

Libya seen by the Telescope

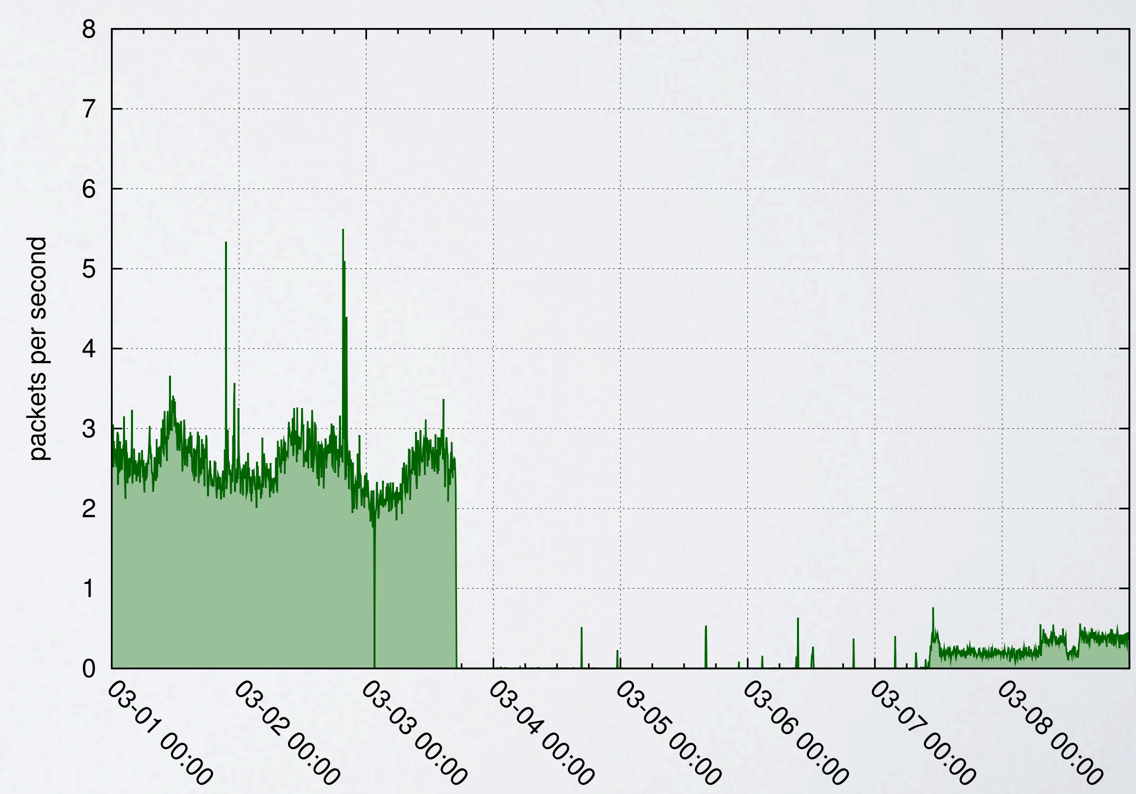

SATELLITE CONNECTIVITY

probable signal jamming

- Third Libyan outage
  - A Libyan IPv4 prefix managed by SatAS1 was BGP-reachable
  - Only a small amount of traffic from that prefix reaches the telescope during the outage

Libya: Telescope traffic from national operator and satellite-based ISP

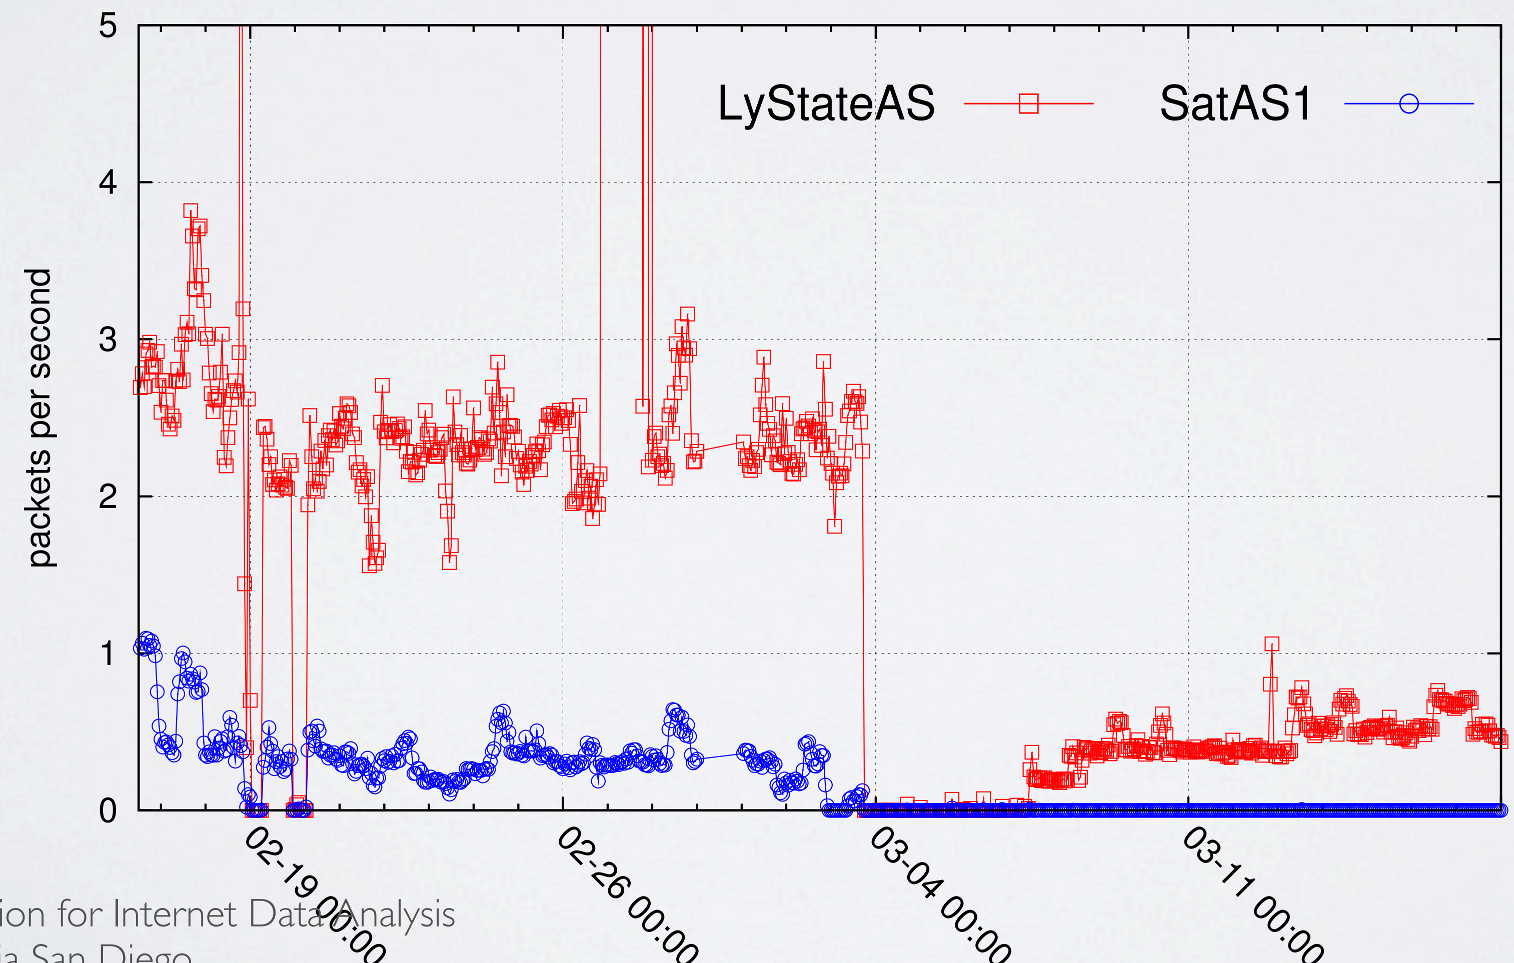

THE EVENTS (2/3)

Earthquakes

- Christchurch - NZ
  - *February 21st, 2011* 23:51:42 UTC
  - Local time 22nd, 12:51:42 PM
  - Magnitude: 6.1
- Tohoku - JP
  - *March 11th, 2011* 05:46:23 UTC
  - Local time 02:46:23 PM
  - Magnitude: 9.0

	Christchurch - NZ		Tohoku - JP	
Distance (Km)	Networks	IP Addresses	Networks	IP Addresses
< 5	1	255	0	0
< 10	283	662,665	0	0
< 20	292	732,032	0	0
< 40	299	734,488	0	0
< 80	309	738,062	5	91
< 100	310	738,317	58	42,734
< 200	348	769,936	1,352	1,691,560
< 300	425	828,315	3,953	4,266,264
< 400	1,531	3,918,964	16,182	63,637,753
< 500	1,721	4,171,527	41,522	155,093,650

We used MaxMind GeoLite City DB to compute distance from a given network to the epicenters

A. Dainotti, R. Amman, E. Aben, K. C. Claffy, "Extracting Benefit from Harm: Using Malware Pollution to Analyze the Impact of Political and Geophysical Events on the Internet" ACM SIGCOMM Computer Communication Review, Jan 2012

A SIMPLE METRIC

to evaluate impact and extension

- $I_{\Delta t_i}$ number of distinct source IP addresses seen by the telescope over the interval **Δ**ti,
- $\Delta t_1, ..., \Delta t_n$ 1-hour time slots ***following*** the event
- $\Delta t_{-1}, ..., \Delta t_{-n}$ 1-hour time slots ***preceding*** the event

$$\theta = \frac{\sum_{i=-1}^{-24} I_{\Delta t_i}}{\sum_{j=1}^{24} I_{\Delta t_j}}$$

RADIUS OF IMPACT

rough estimate based on θ

- We compute θ for address ranges geolocated at different distances from the epicenter of the earthquake *(0 to 500km in bins of 1km each)*
- θ around 1 indicates no substantial change in the number of unique IP addresses observed in IBR before and after the event.

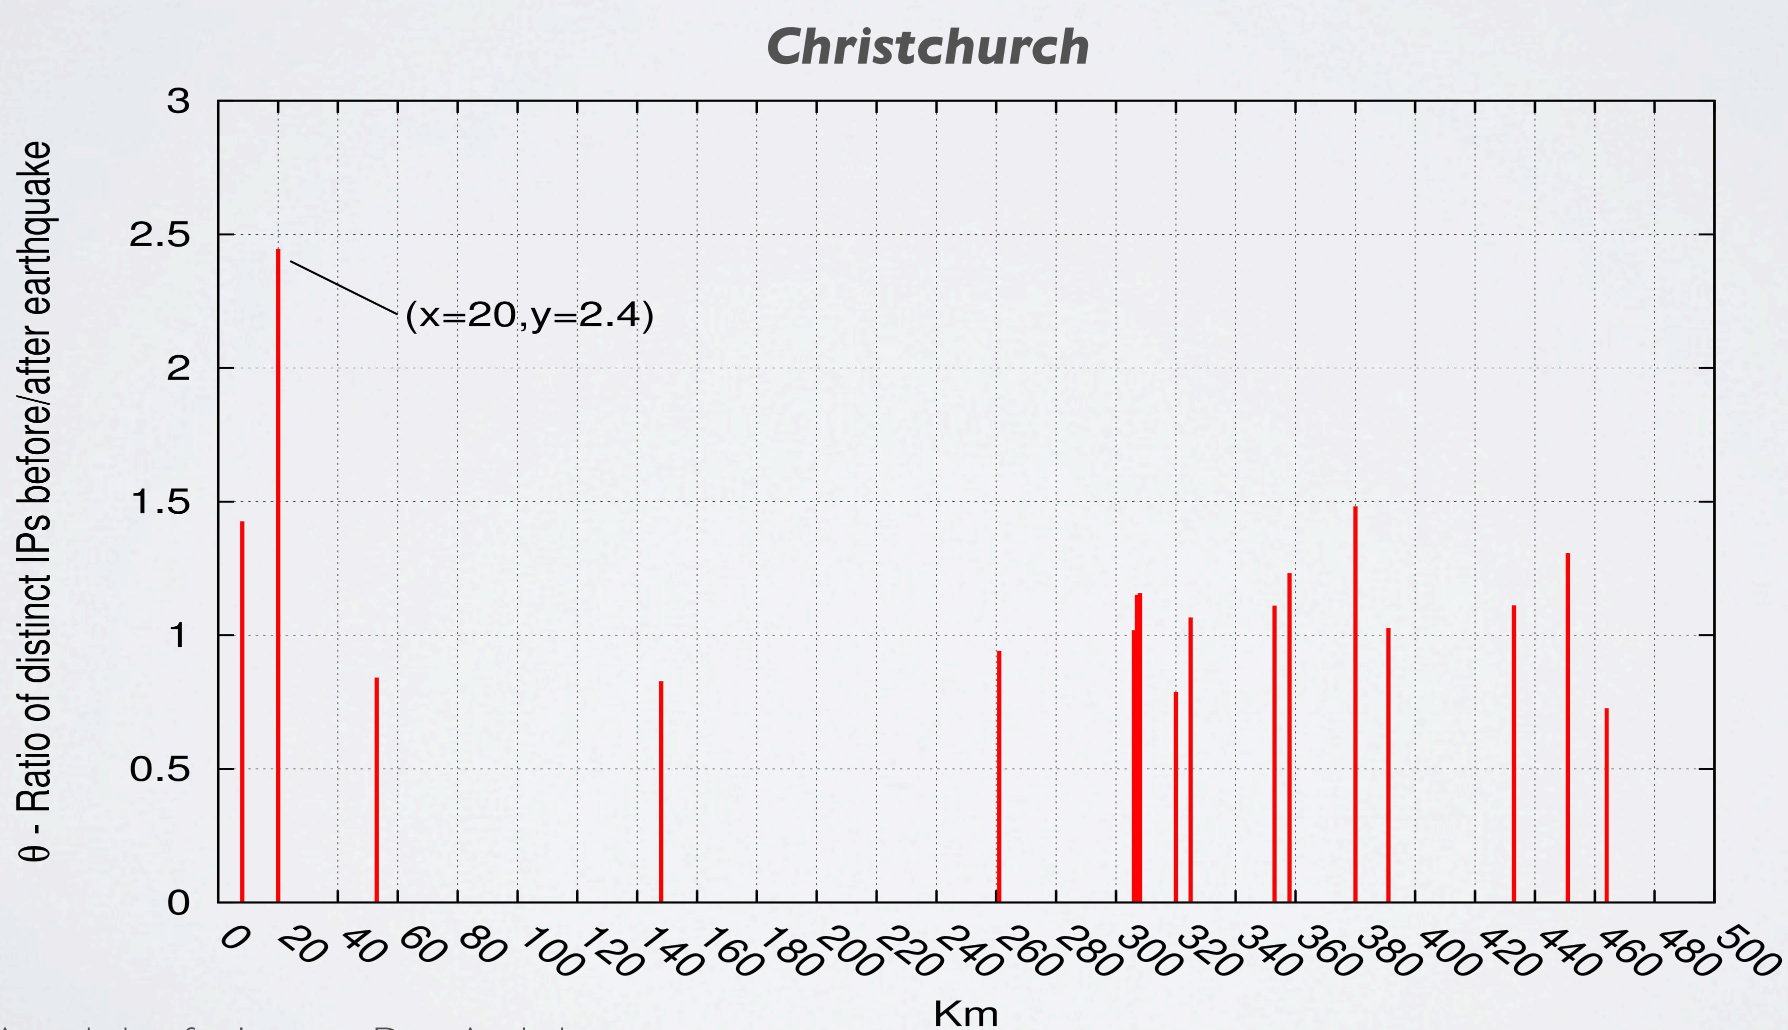

Cooperative Association for Internet Data Analysis
University of California San Diego

RADIUS OF IMPACT

rough estimate based on θ

We call ρ_{max} the maximum distance at which we observe a value of θ significantly > 1

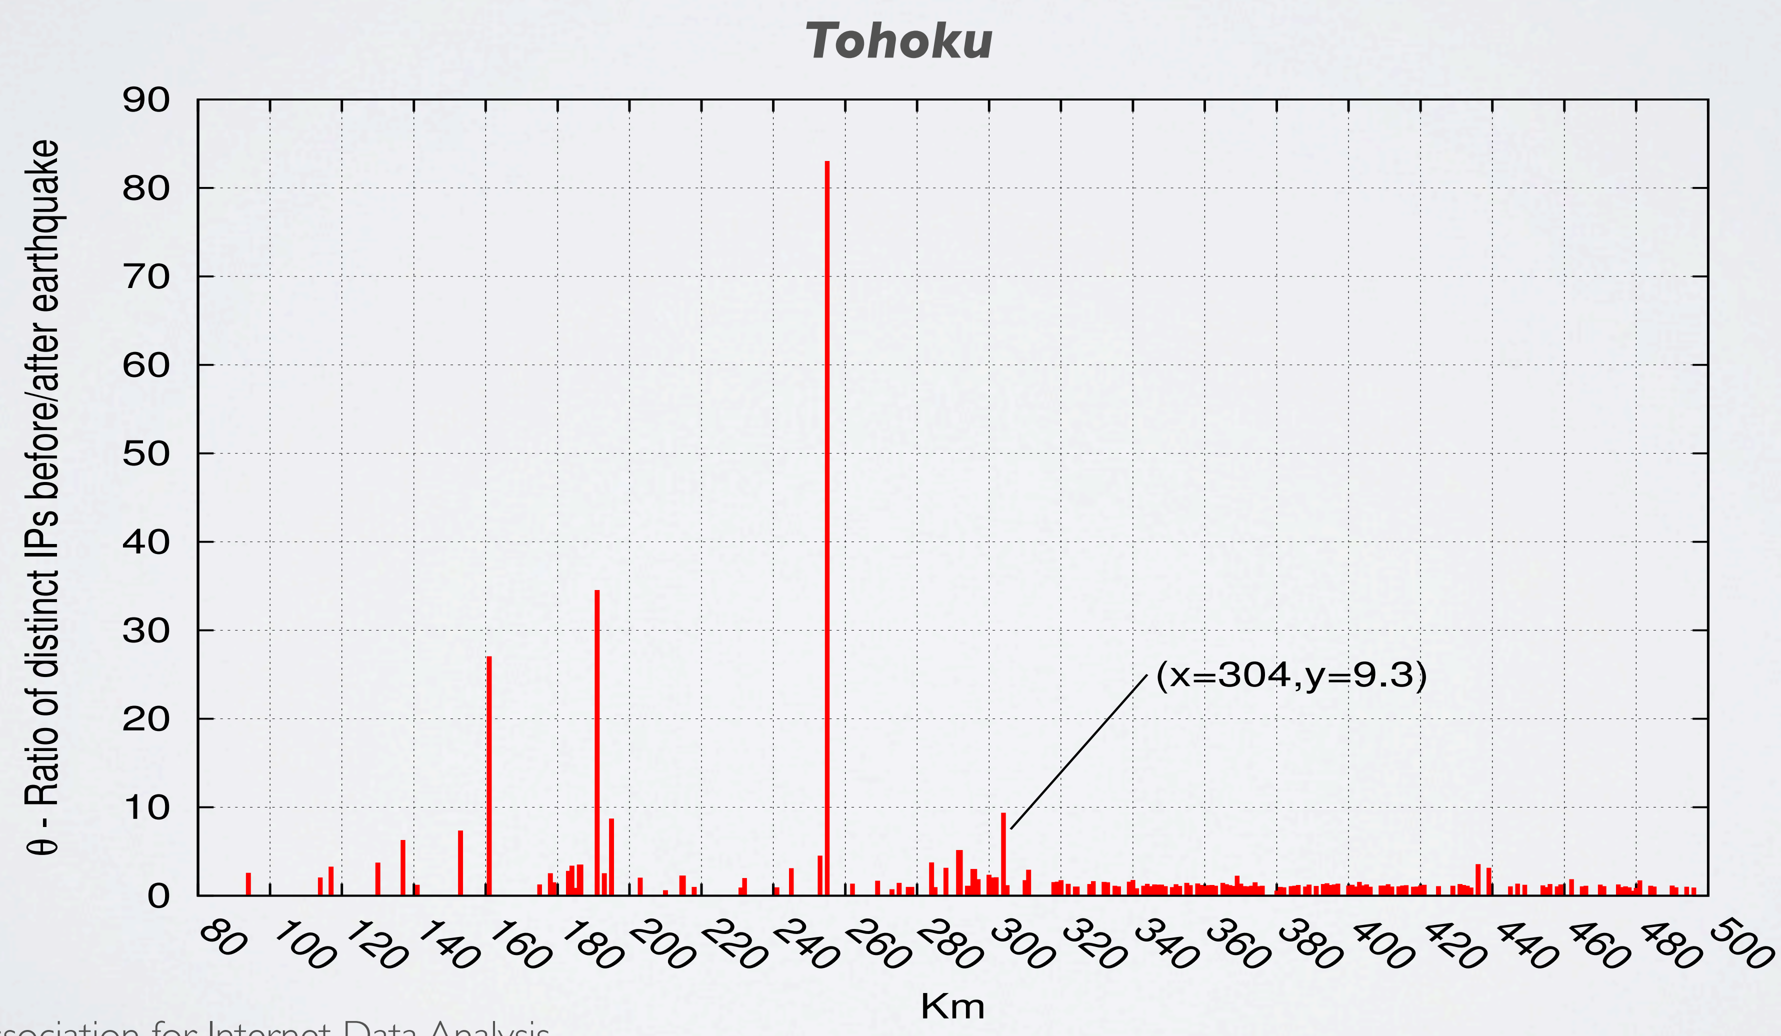

EXTENSION OF IMPACT

geo coordinates of most affected networks

Networks within each respective ρ_{max}

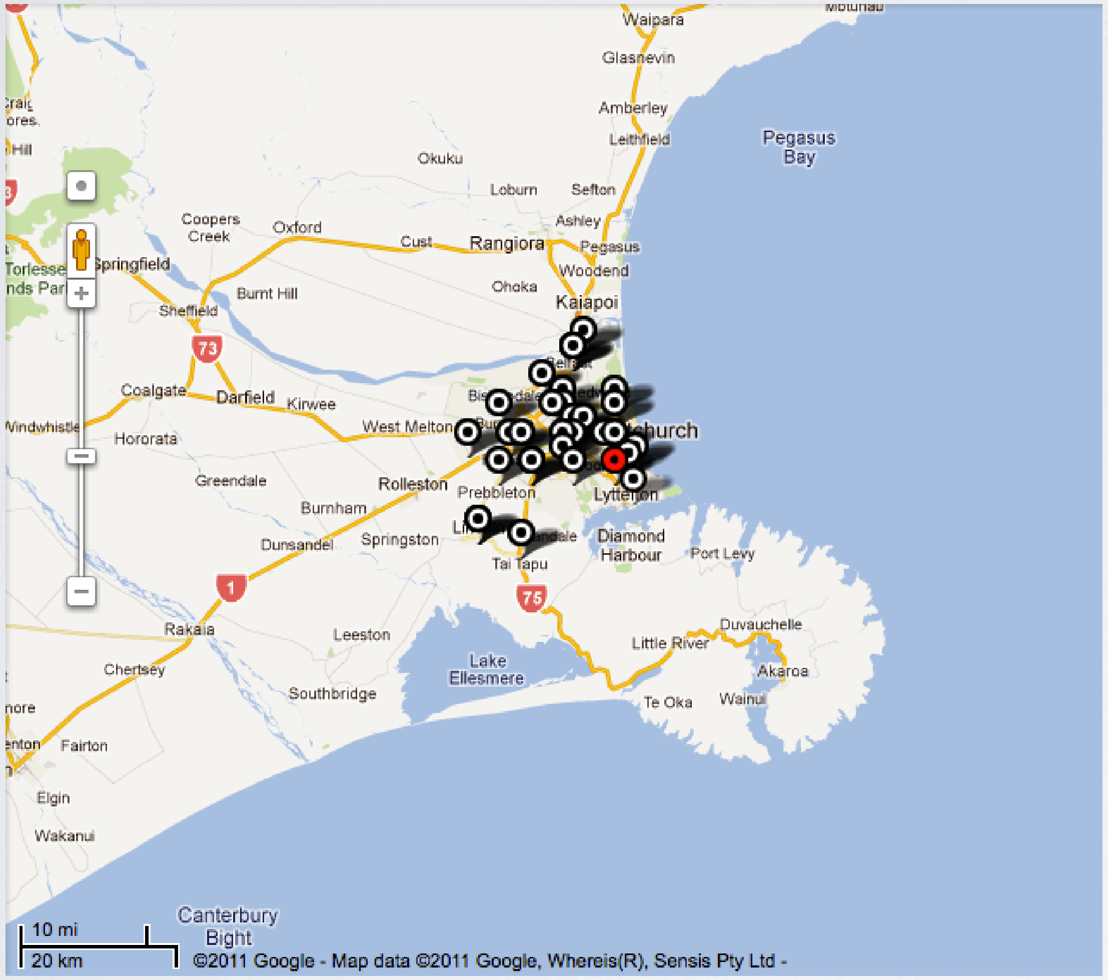

(a) Christchurch

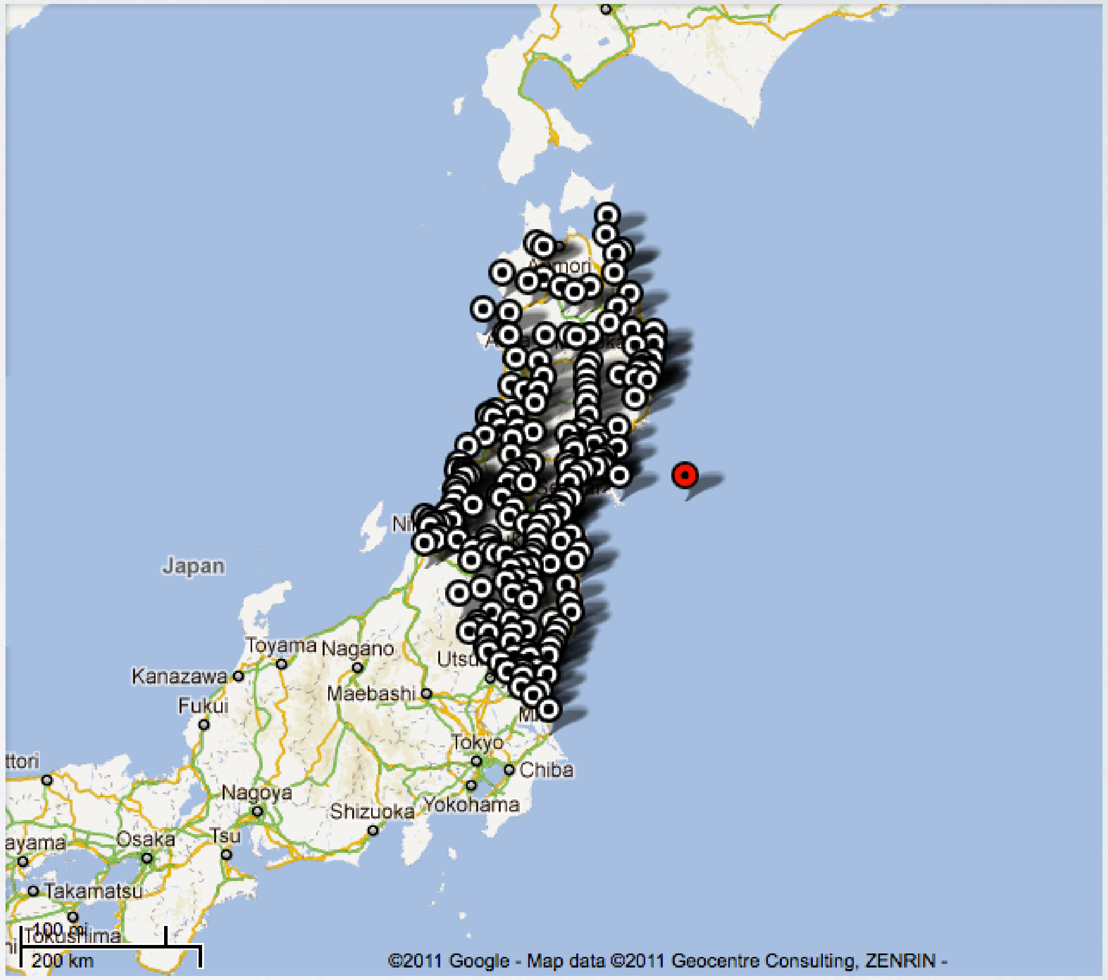

(b) Tohoku

"MAGNITUDE"

A measure of impact

- Varying the radius, we pick the highest value of θ calculated for *the whole set of* networks within the corresponding circle

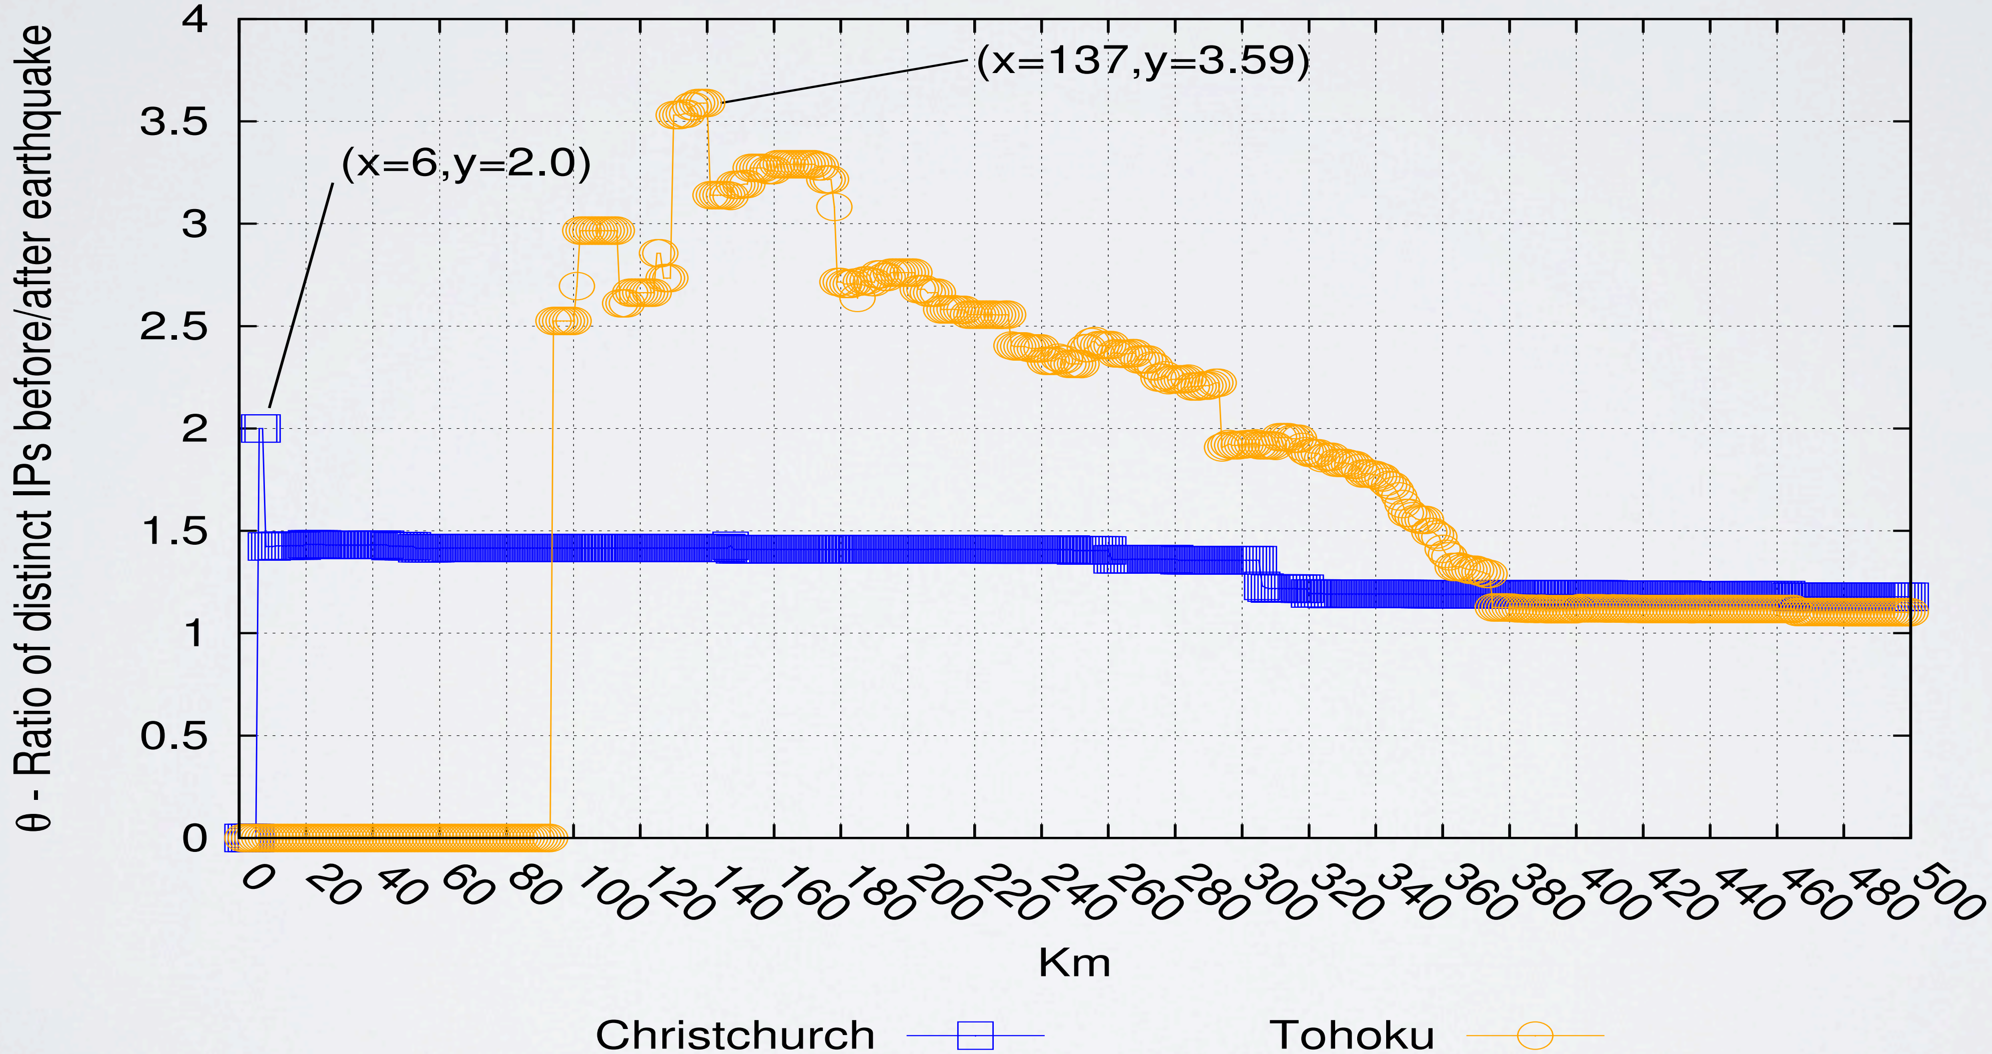

	Christchurch	Tohoku
Magnitude (θ_{max})	2 at $6km$	3.59 at $137km$
Radius (ρ_{max})	$20km$	$304km$

IP RATE IN TIME

reflects the dynamics of the event

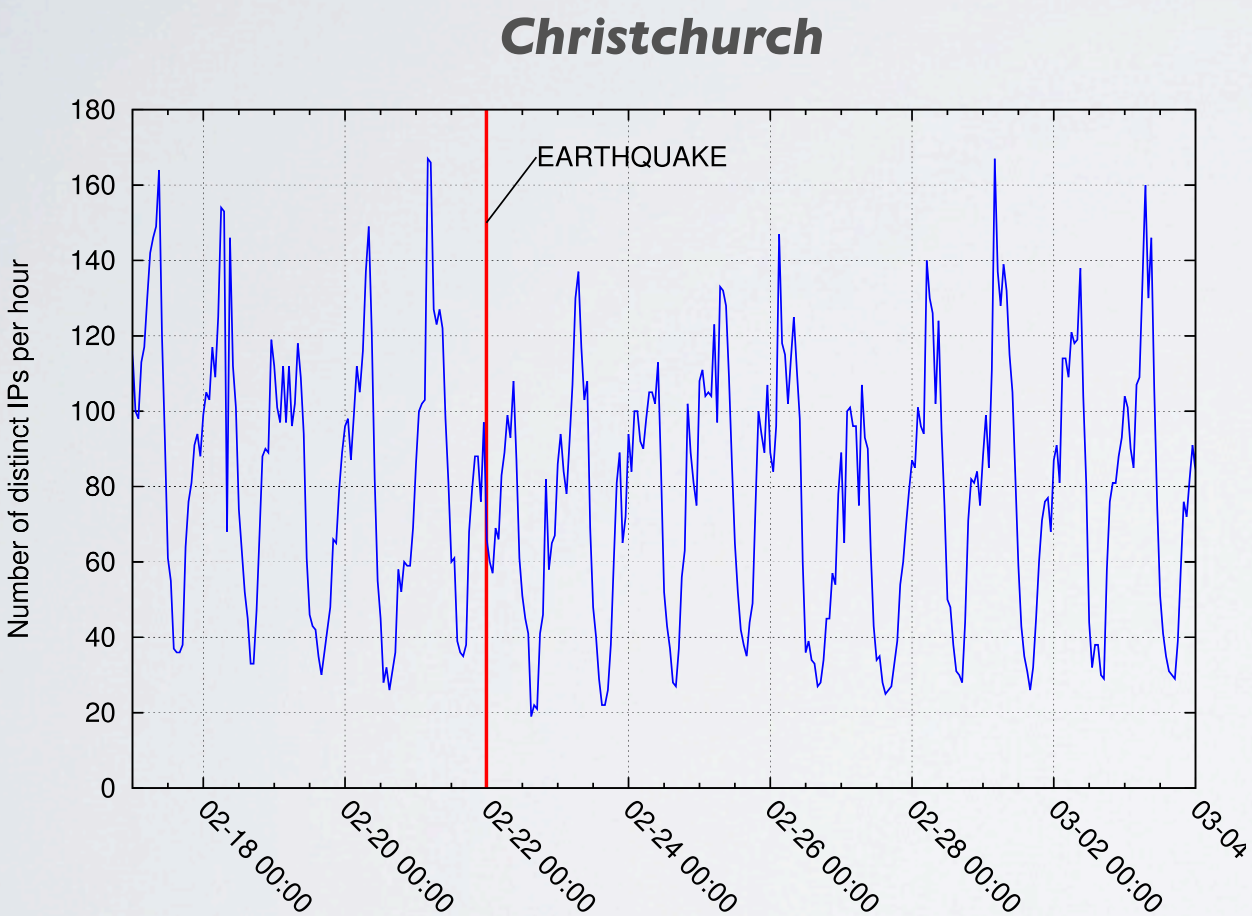

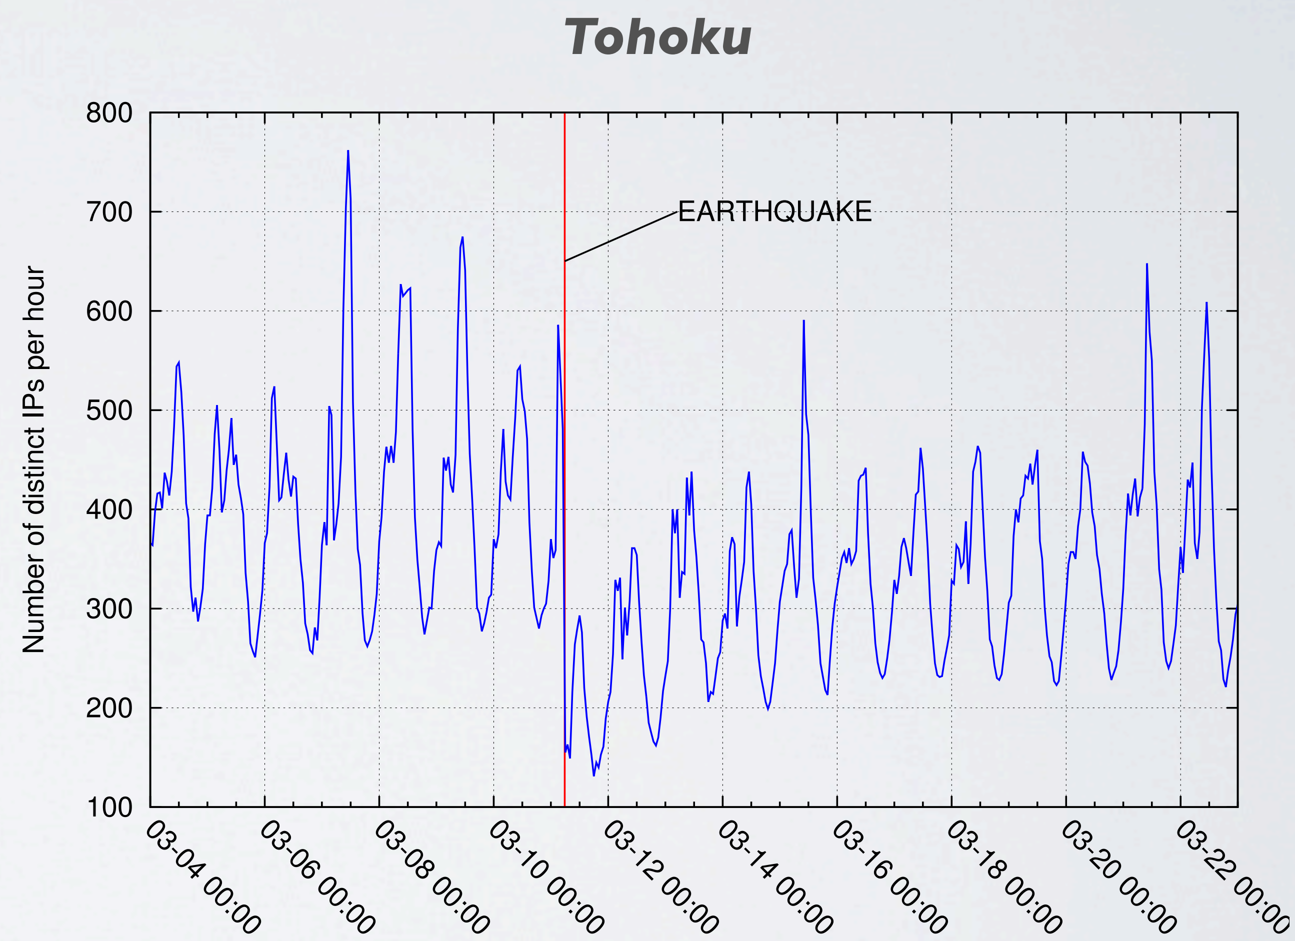

EVALUATING Θ

variations over a long time period

- 2 months period of observation
- θ normally stays within [0.7 - 1.3]

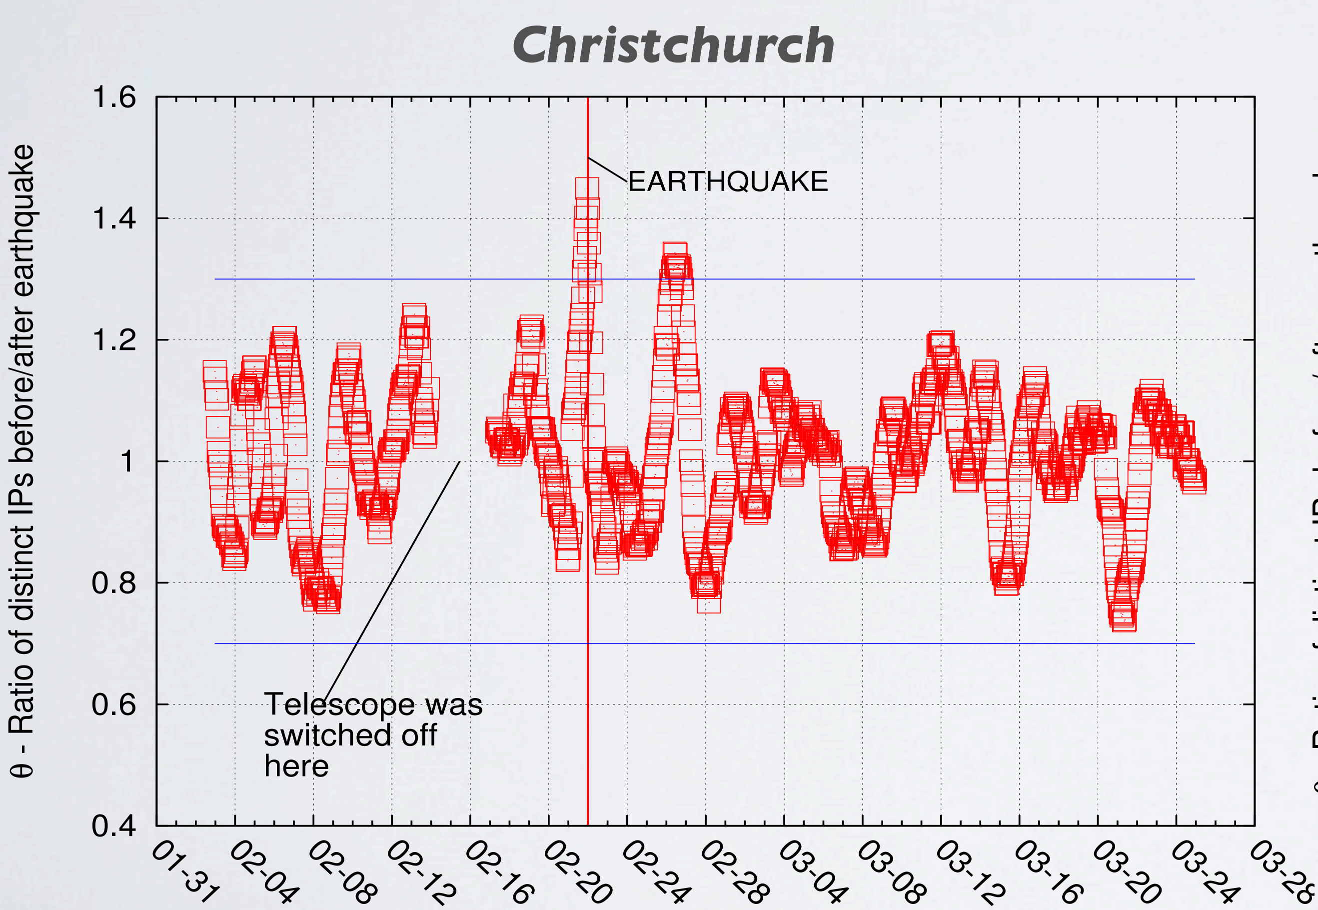

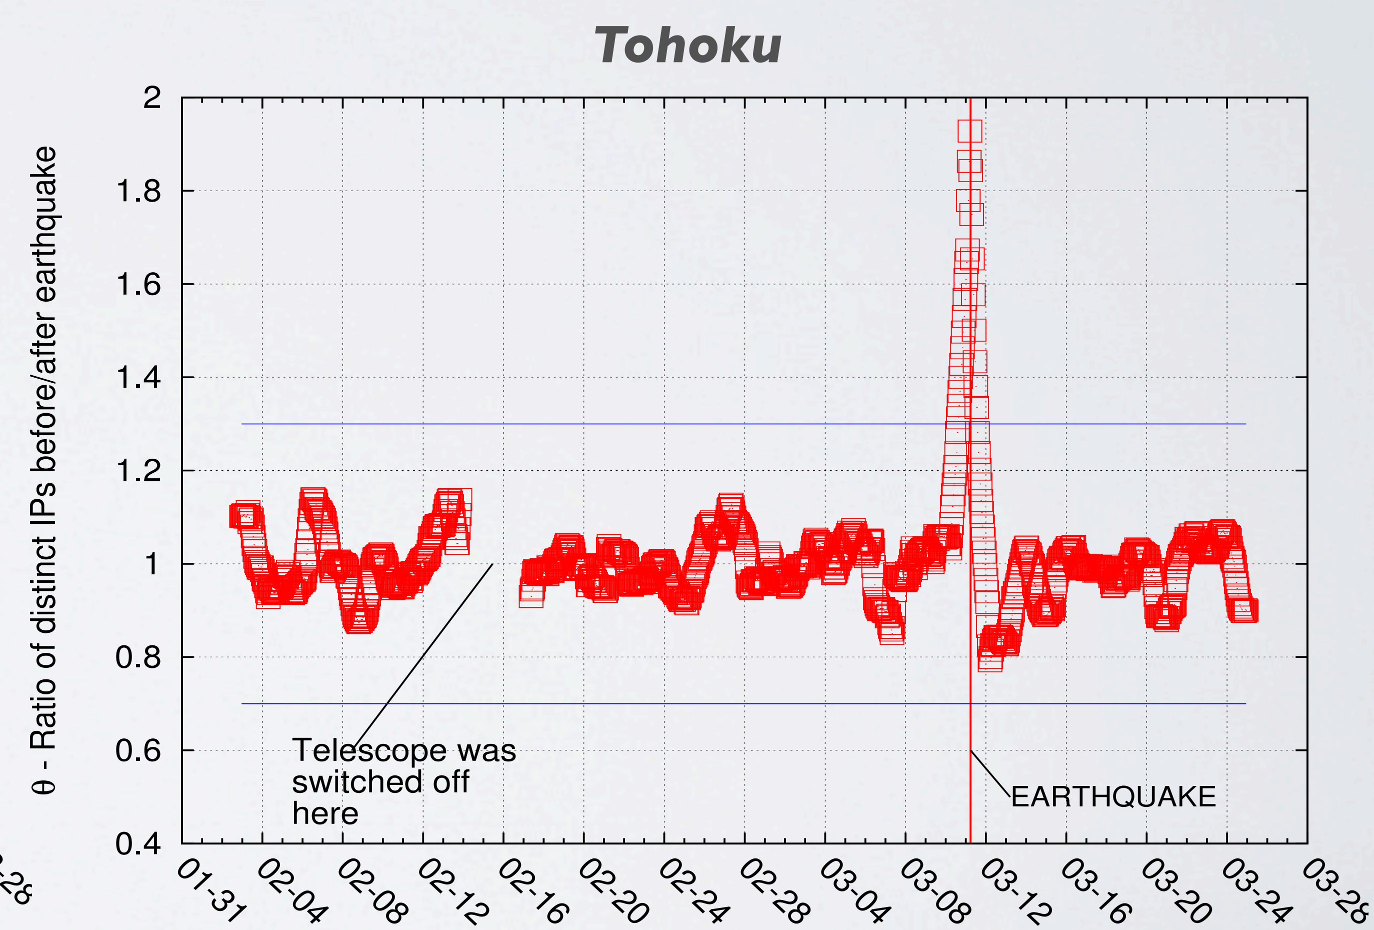

THE EVENTS (3/3)

Hurricane Sandy

- Atlantic, Caribbean, US east coast
  - October *22nd - 31st. 2012*

SANDY: IS IT DIFFERENT?

(compared to our previous case studies)

- Movement over a large area
  - with no fixed epicenter like an earthquake has
- High level of Internet penetration in the affected region, including major hubs for international Internet connectivity
- Disruption was limited to only a subset of networks/hubs in the affected region, making it harder to identify geographic areas of massive impact
- For the 1st time we tried to measure in realtime

IBR: SANDY IN NYC

Reusing the same metric based on ratio of distinct source IPs

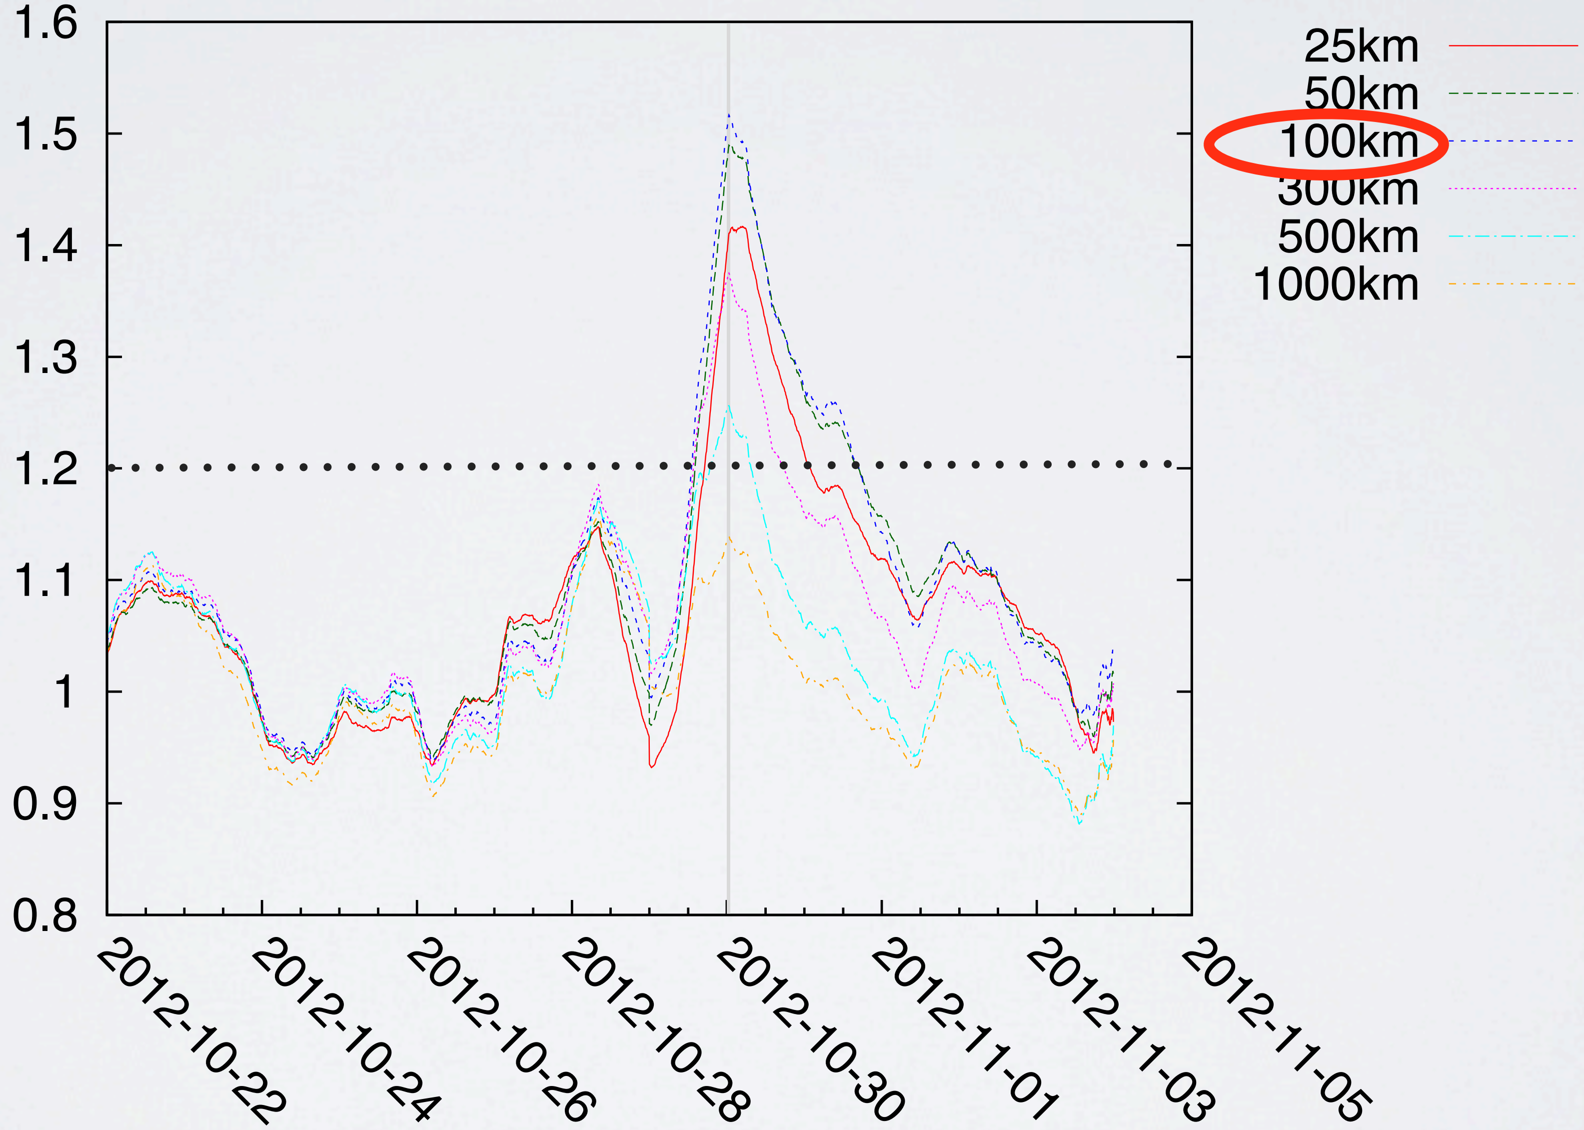

IBR: NY, HOME *vs* BUSINESS

*Different impact on home vs business users**

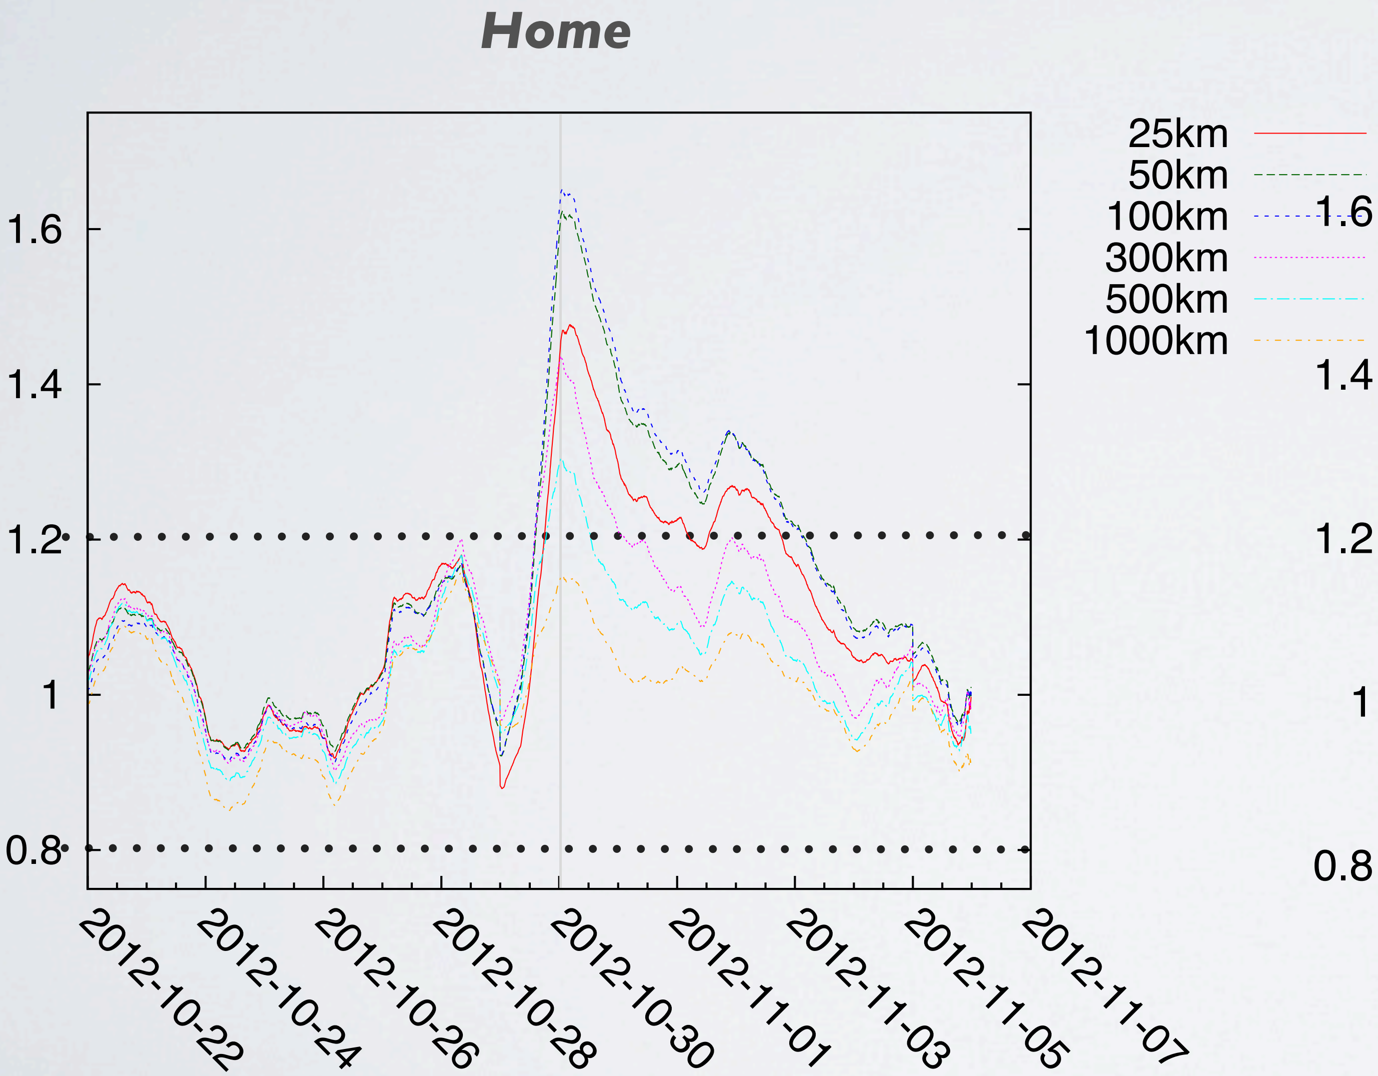

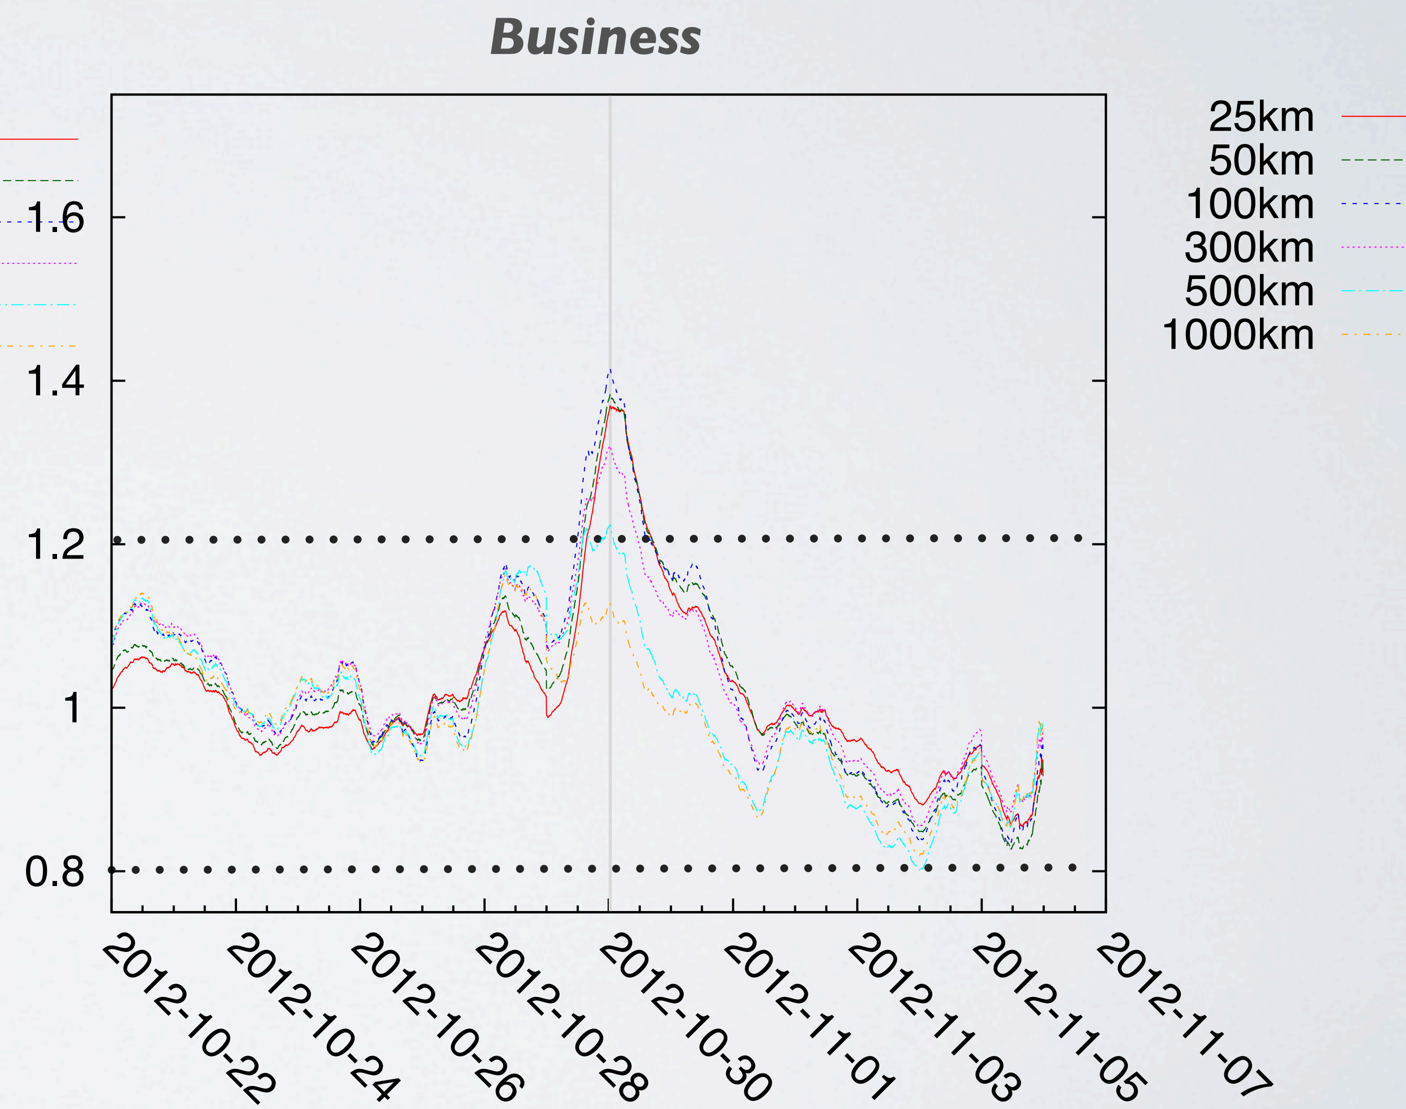

Cooperative Association for Internet Data Analysis
University of California San Diego

***** according to NetAcuity**
www.digitalelement.com/NetAcuity

ATLAS: RTT

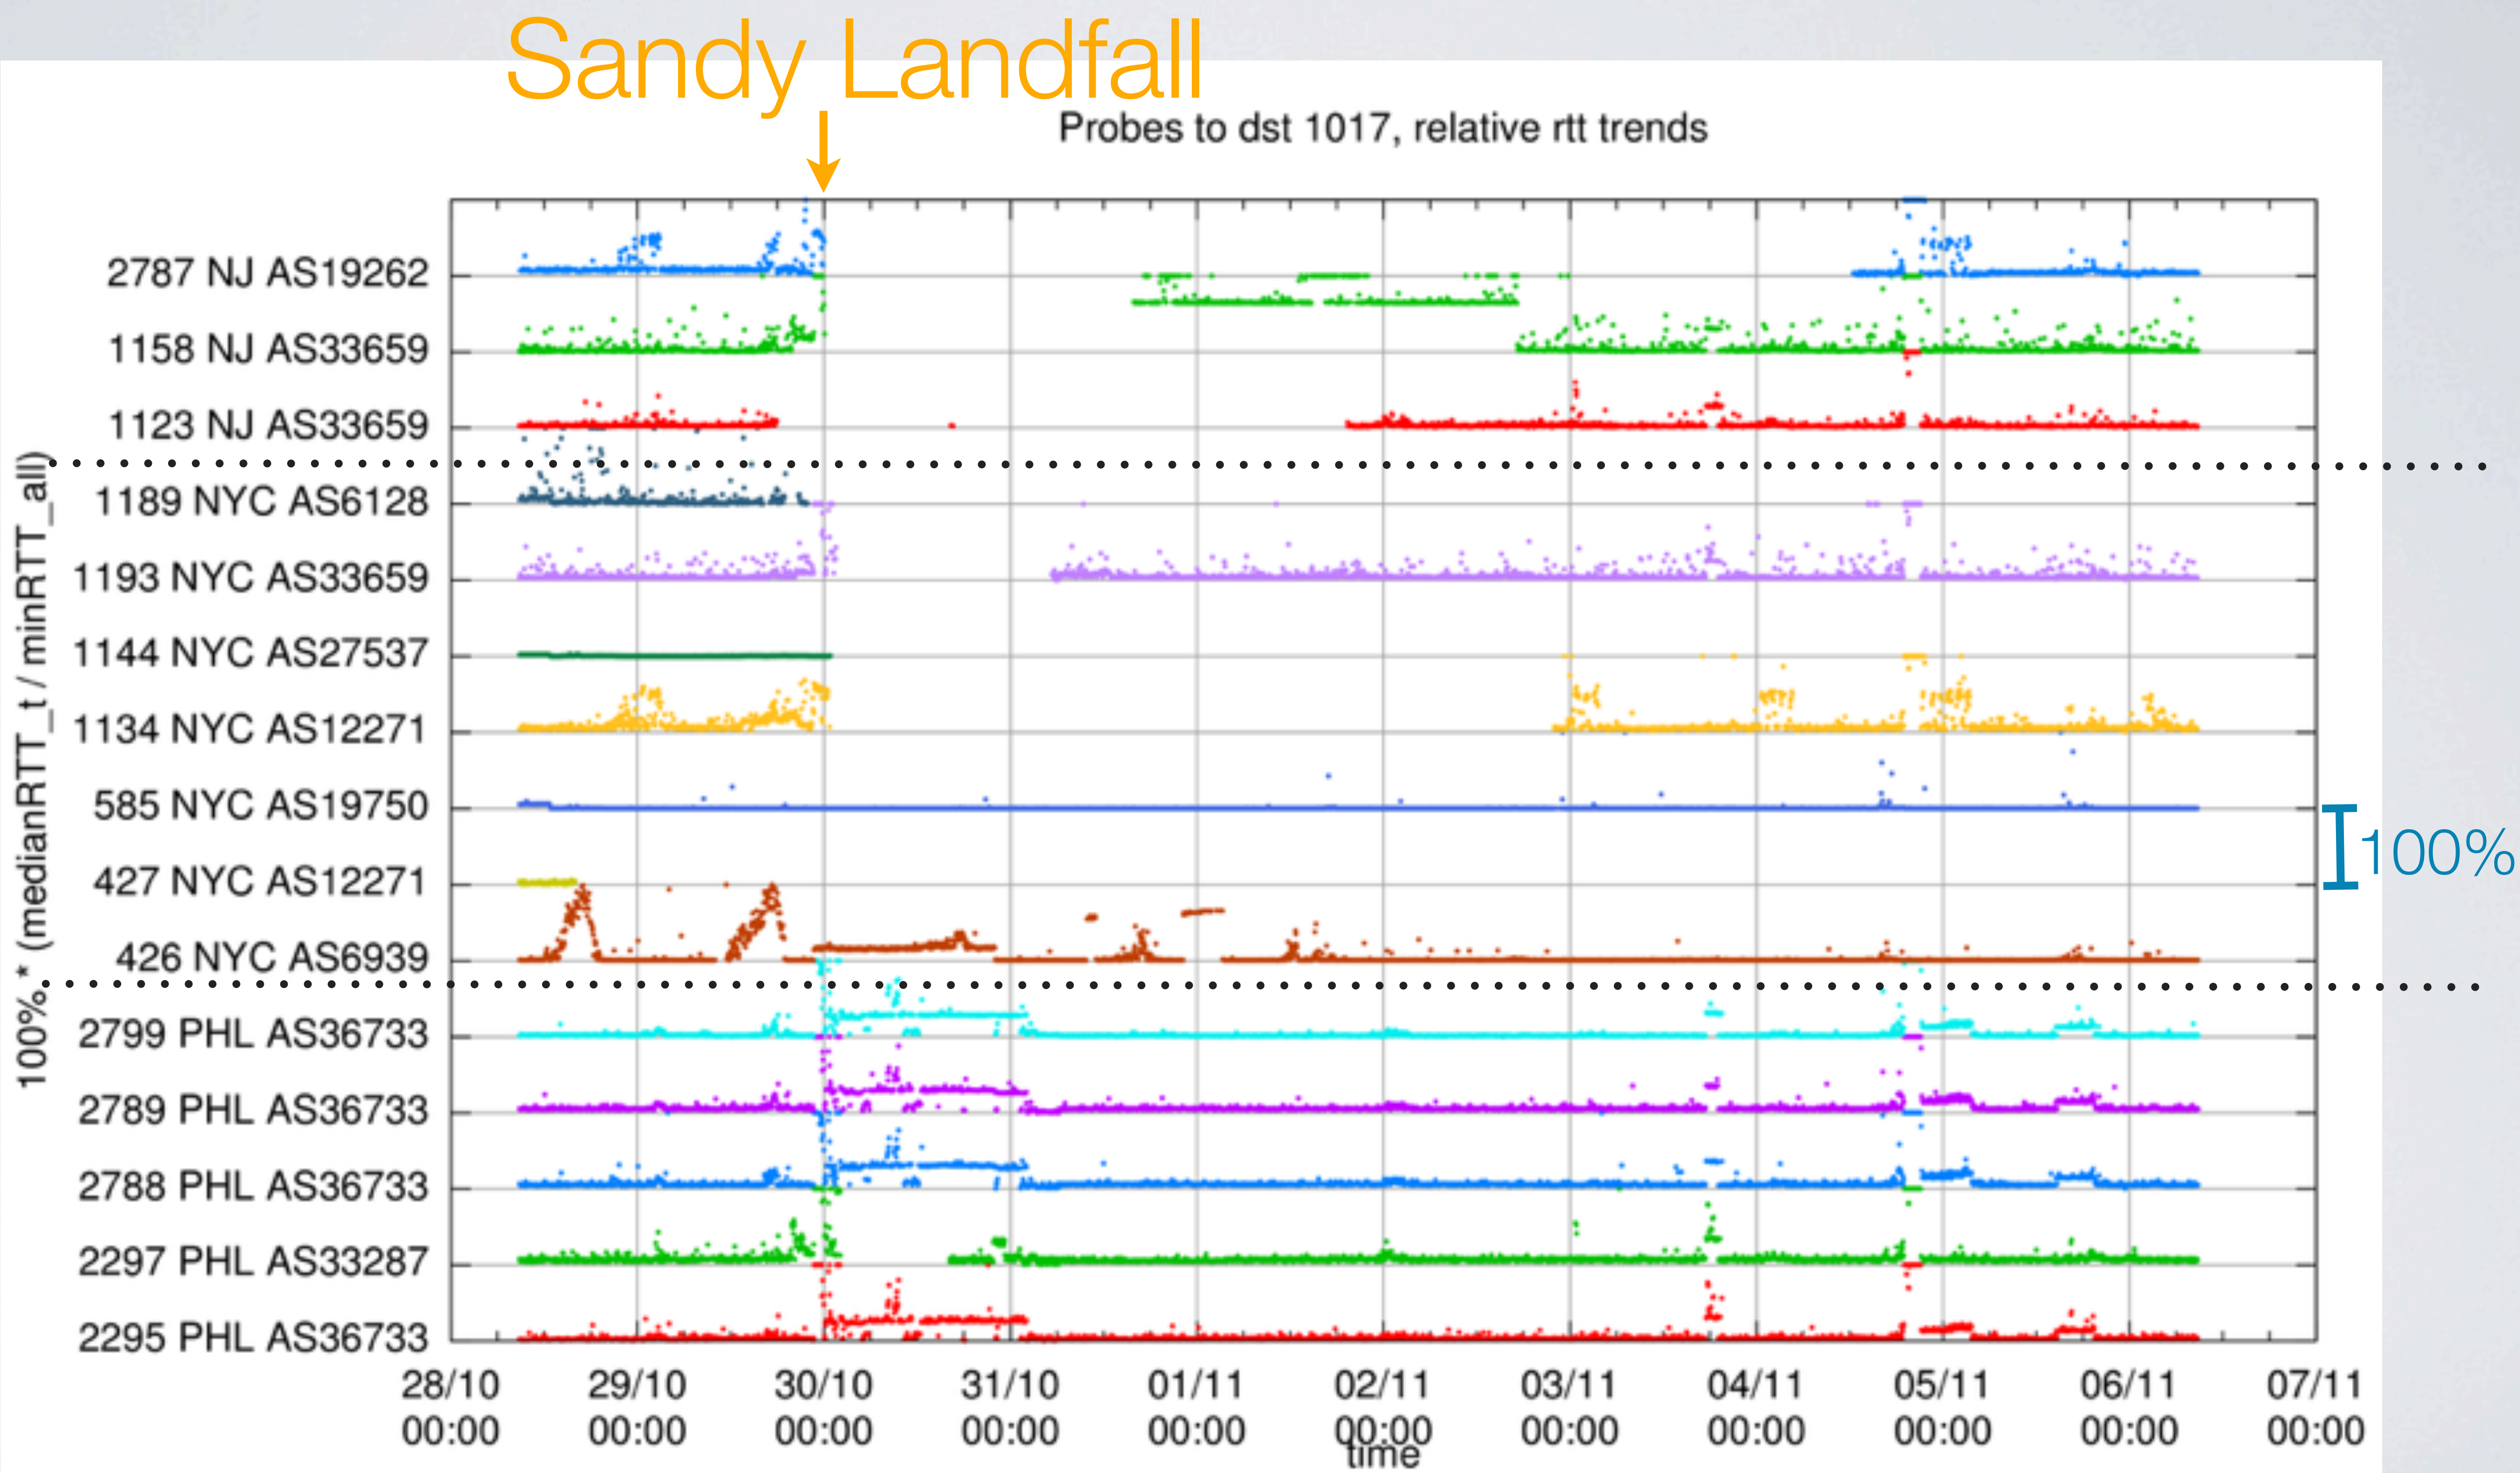

Sandy Landfall
Probes to dst 1017, relative rtt trends
2787 NJ AS19262
1158 NJ AS33659
1123 NJ AS33659
1189 NYC AS6128
1193 NYC AS33659
1144 NYC AS27537
1134 NYC AS12271
585 NYC AS19750
427 NYC AS12271
426 NYC AS6939
2799 PHL AS36733
2789 PHL AS36733
2788 PHL AS36733
2297 PHL AS33287
2295 PHL AS36733
100%.* (medianRTT_t / minRTT_all)
28/10 00:00
29/10 00:00
30/10 00:00
31/10 00:00
01/11 00:00
02/11 00:00
03/11 00:00
04/11 00:00
05/11 00:00
06/11 00:00
07/11 00:00
time
100%

ATLAS: PATH CHANGES

Looking at two major hubs

- New York City (NYC) is a major Internet connectivity hub
- Ashburn/Washington DC (ASH) is the other for US-Europe traffic

ATLAS: PATH CHANGES

dst: ns.ripe.net / AS3333 / NL

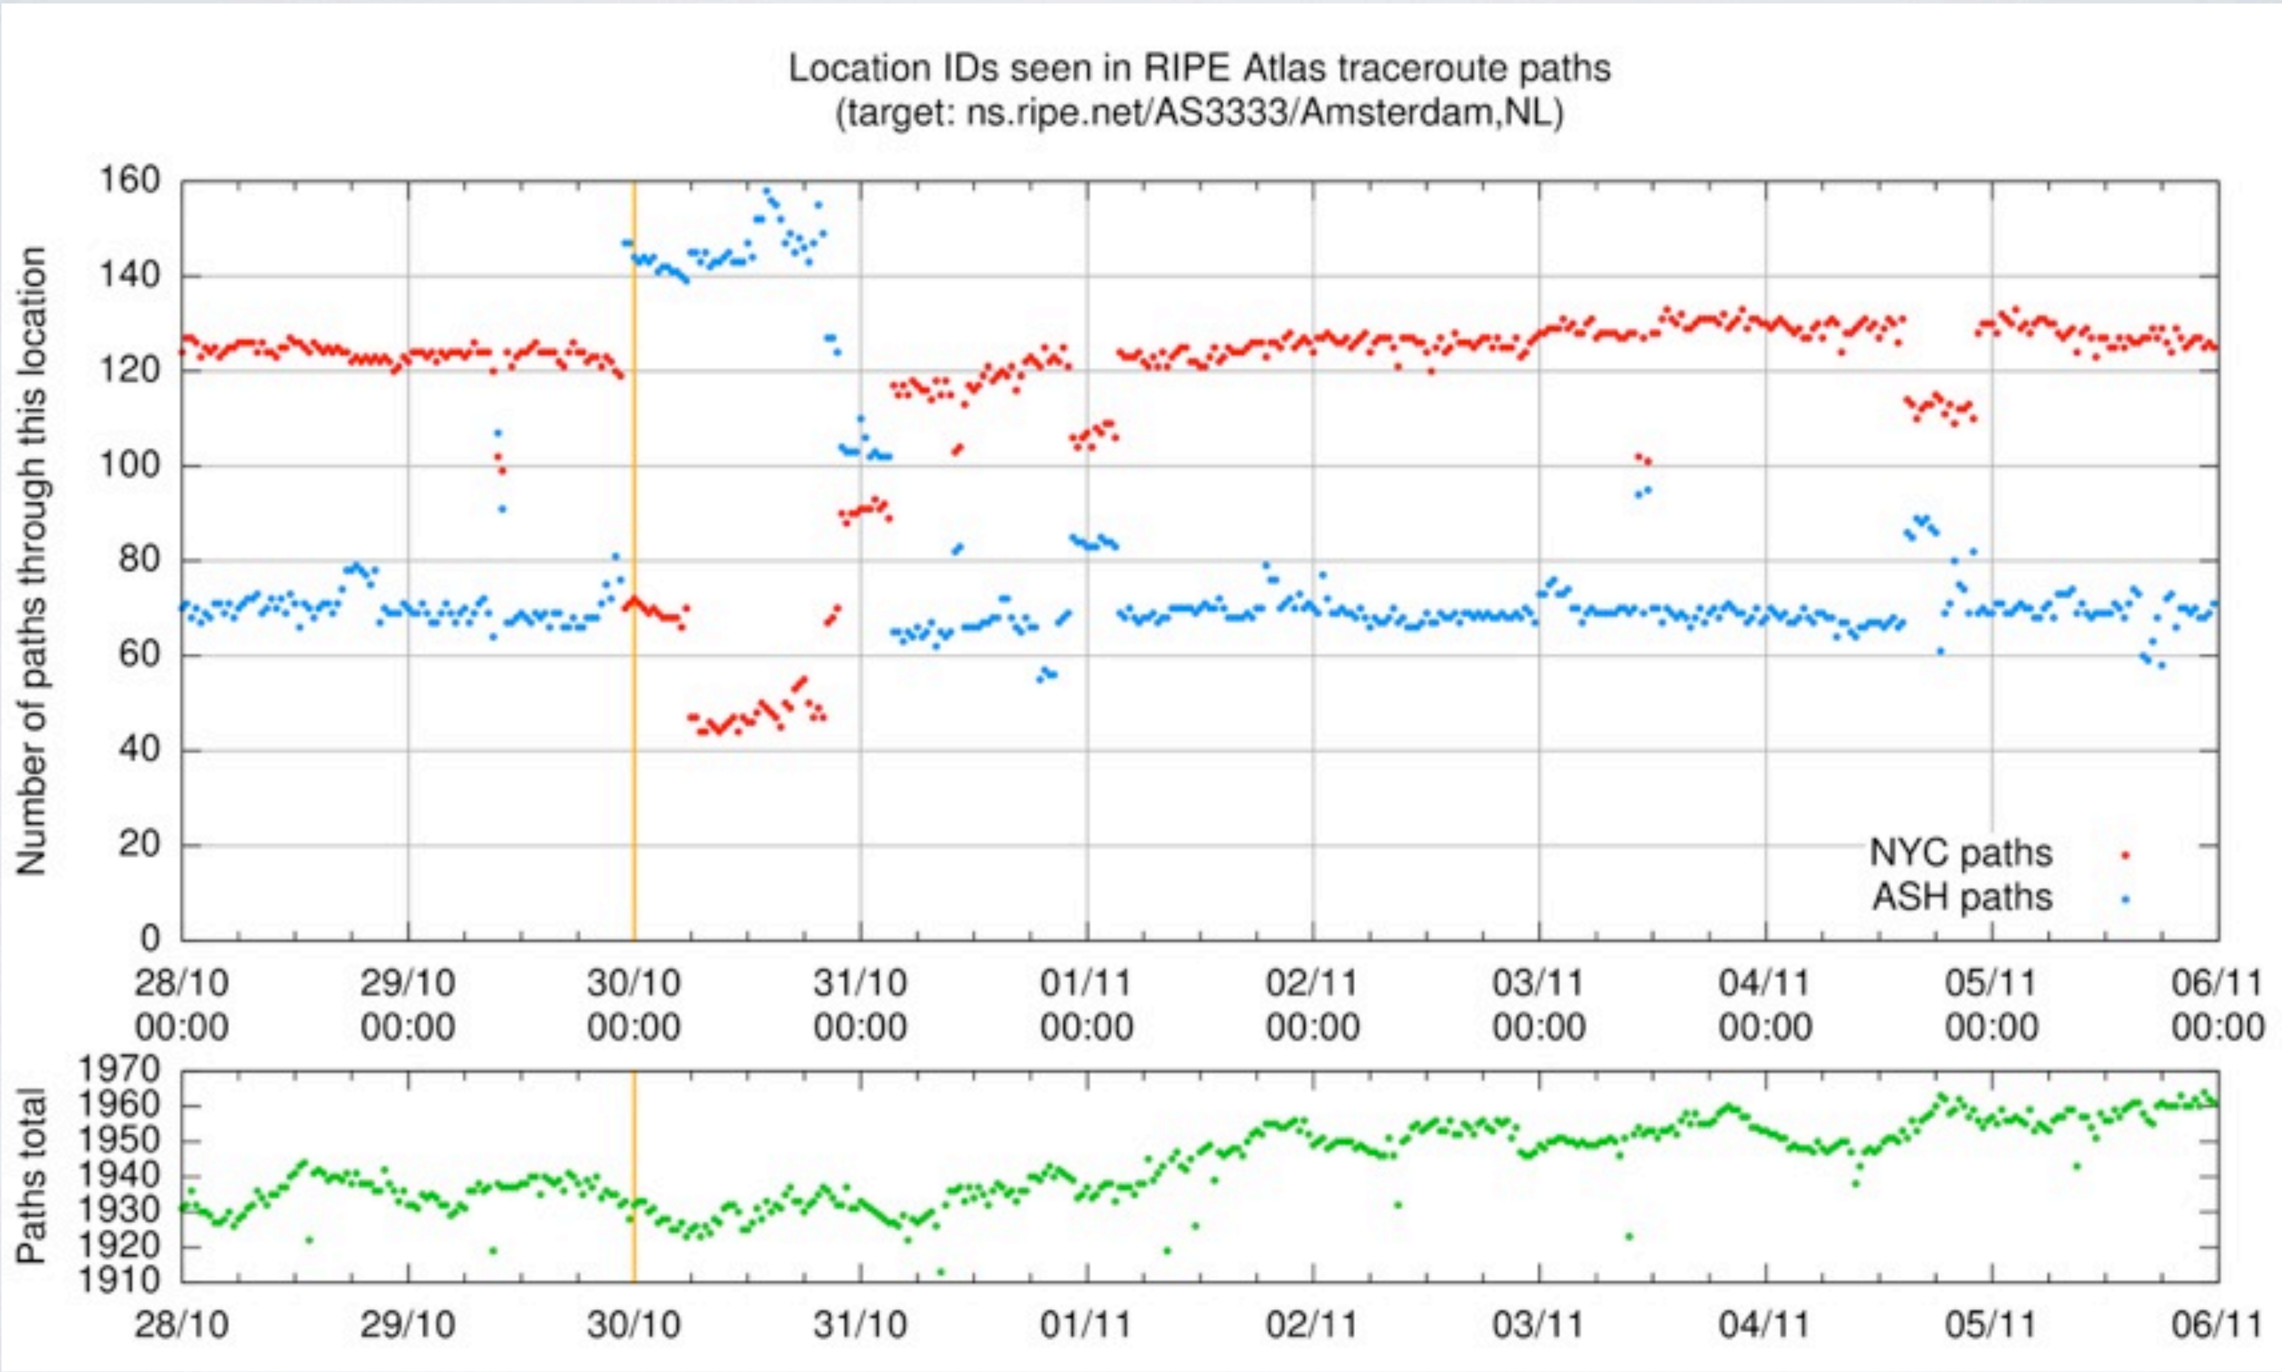

ATLAS: NYC PATH CHANGES

dst: ns.ripe.net / AS3333 / NL

pre: 22:00 UTC vs. post: 09:00 UTC

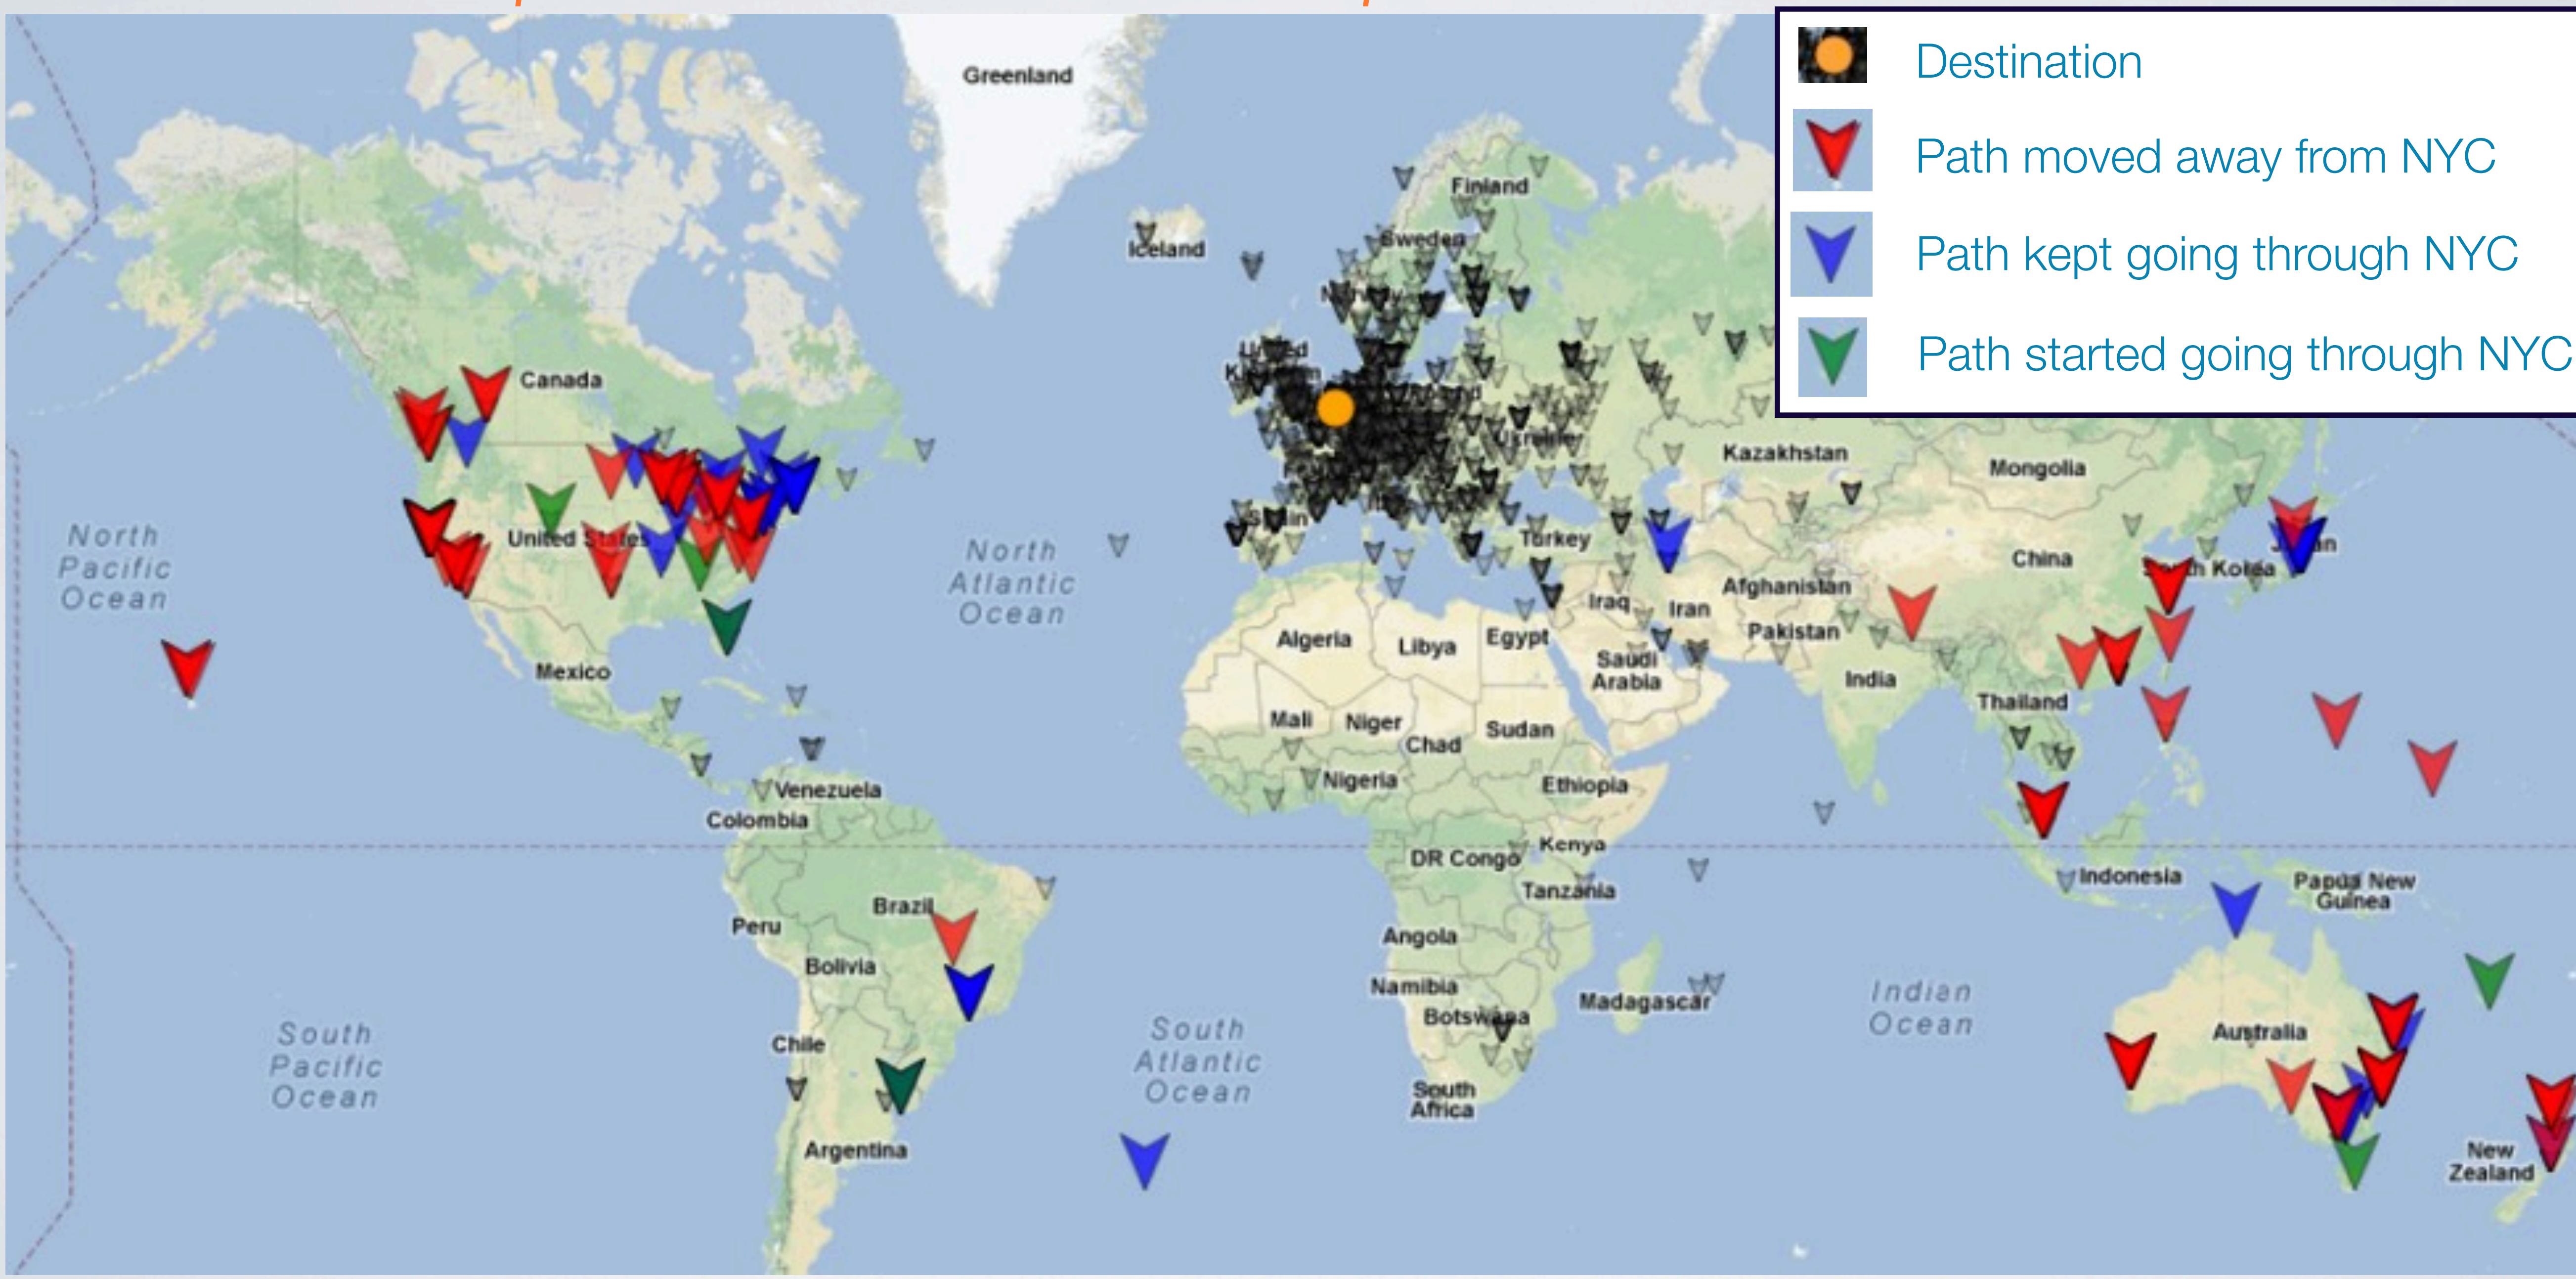

ATLAS: LATENCY

RTT US -> AS3333/NL (+20 ms)

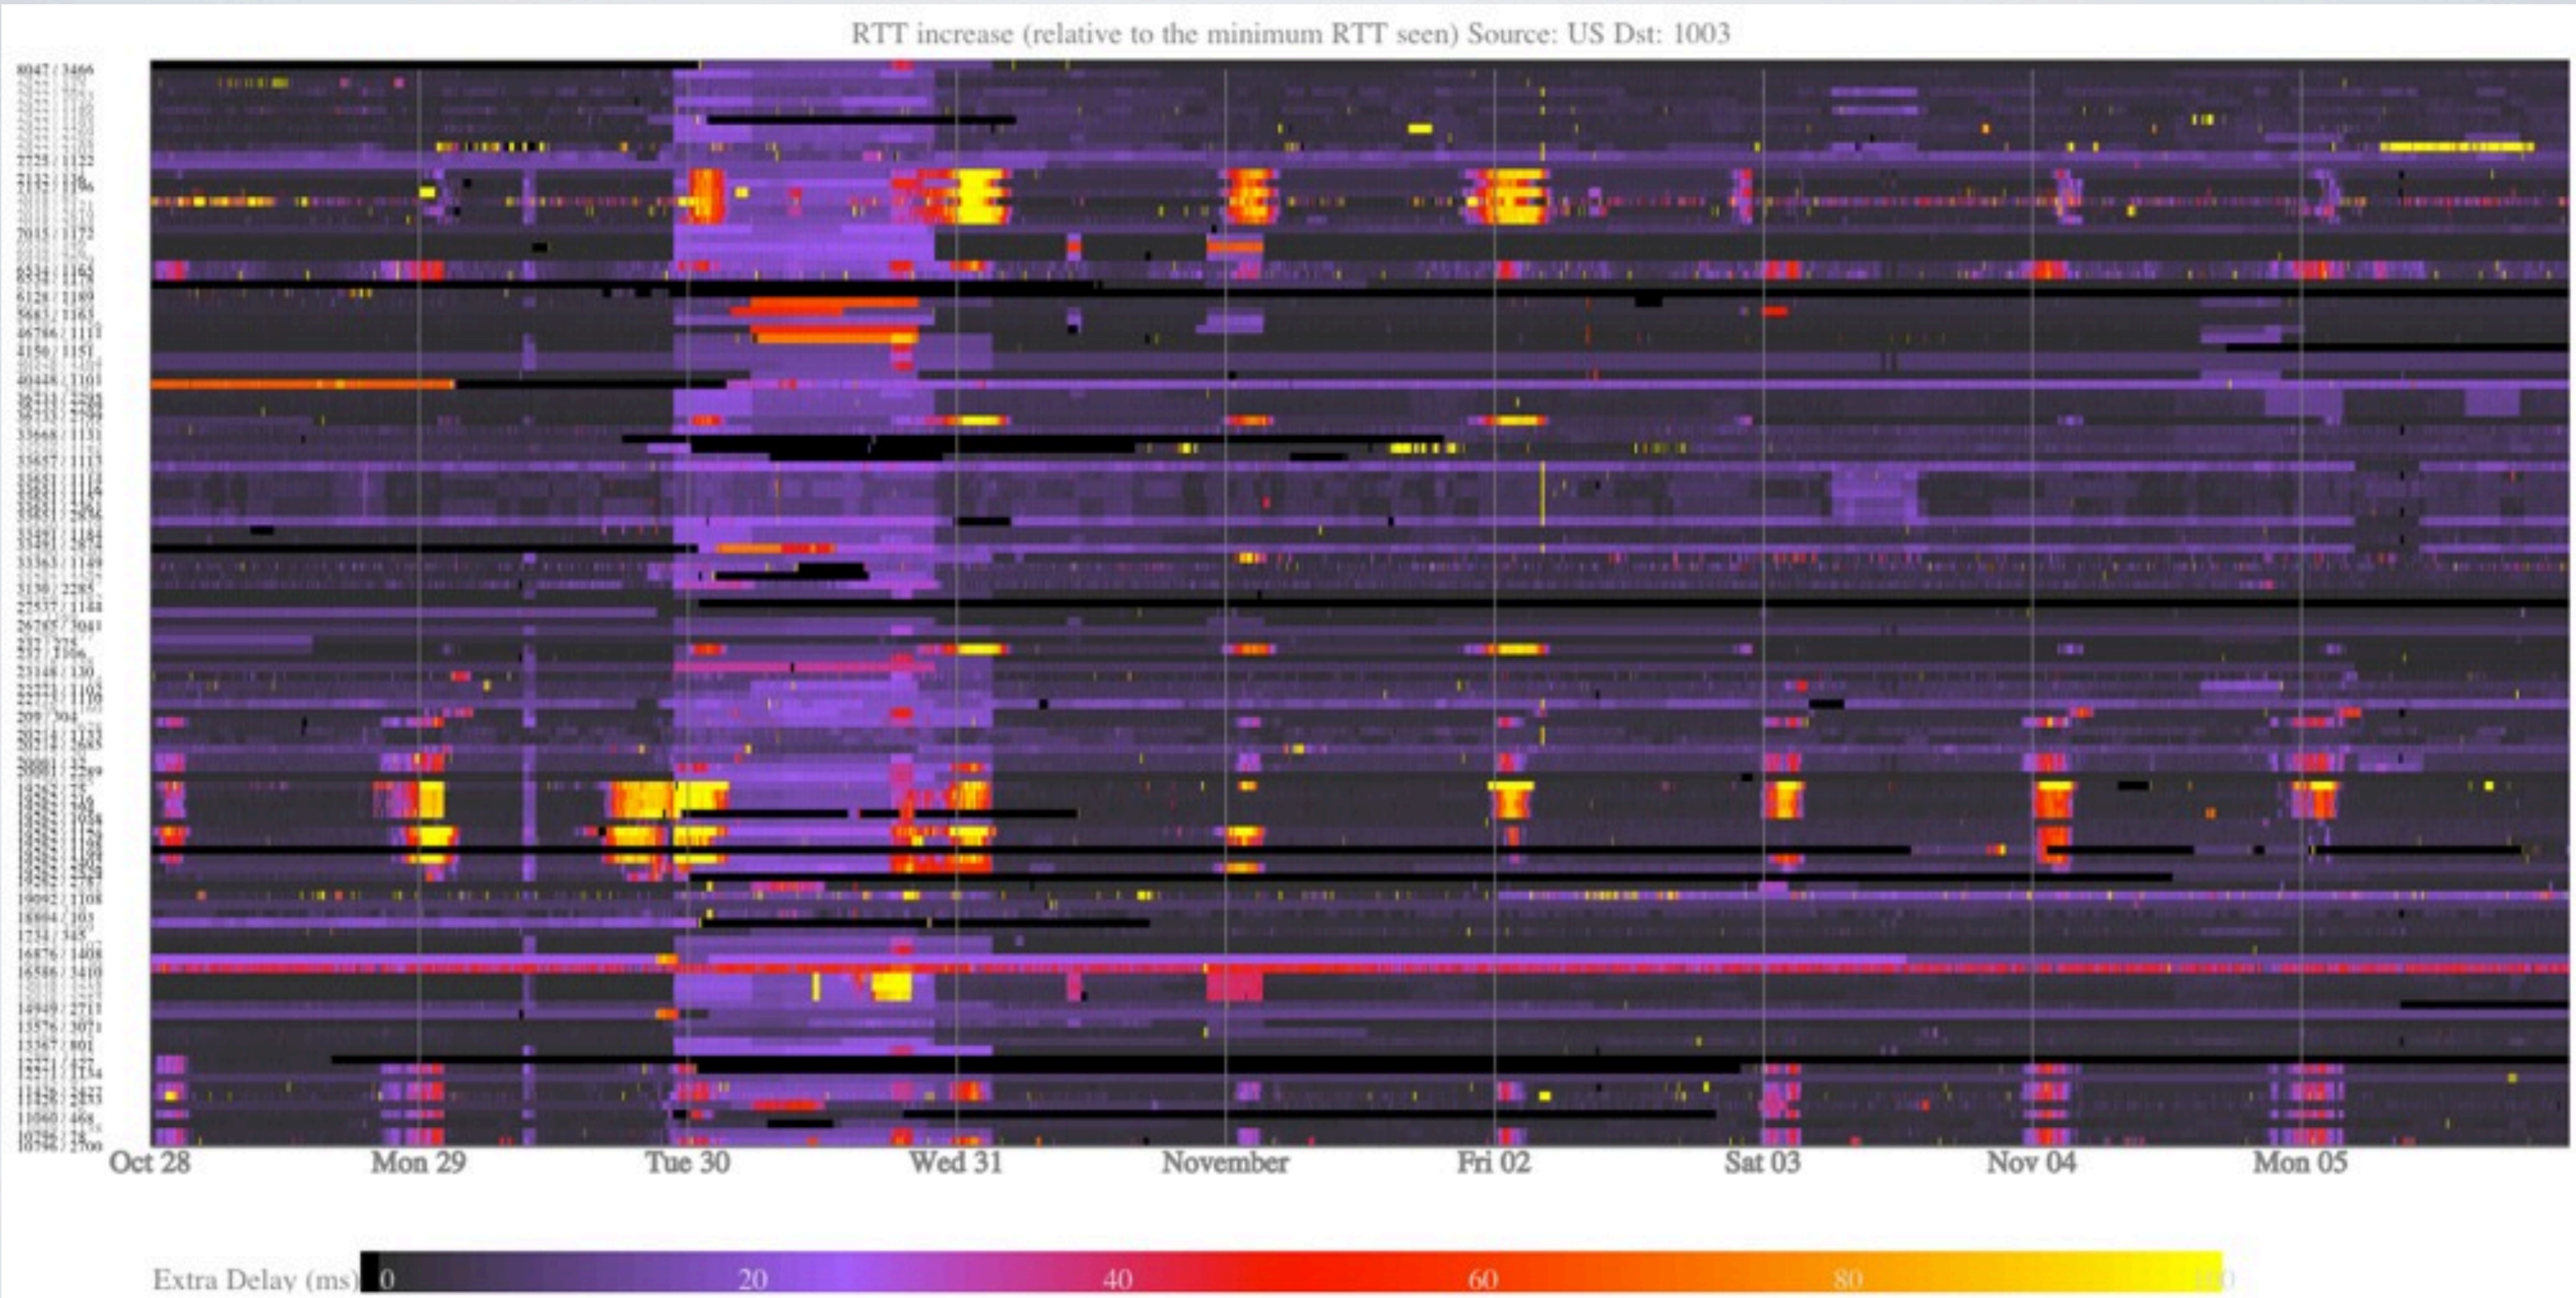

THANKS