SIGCOMM 2012
13-17 August, 2012 - Helsinki, Finland

Extracting Benefit from Harm: Using Malware Pollution to Analyze the Impact of Political and Geophysical Events on the Internet

A. Dainotti, R. Amman, E. Aben, K. C. Claffy
alberto@caida.org
CAIDA/UCSD

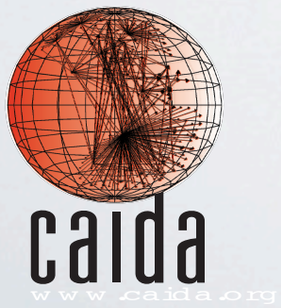

CONTEXT

Analysis of large-scale Internet Outages

- 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

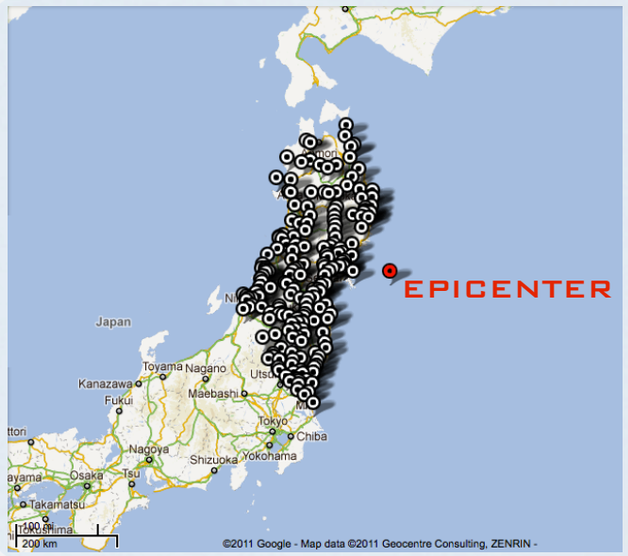

Japan, Mar 2011
Earthquake of Magnitude 9.0

IDEA

"Extracting benefit from harm.."

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

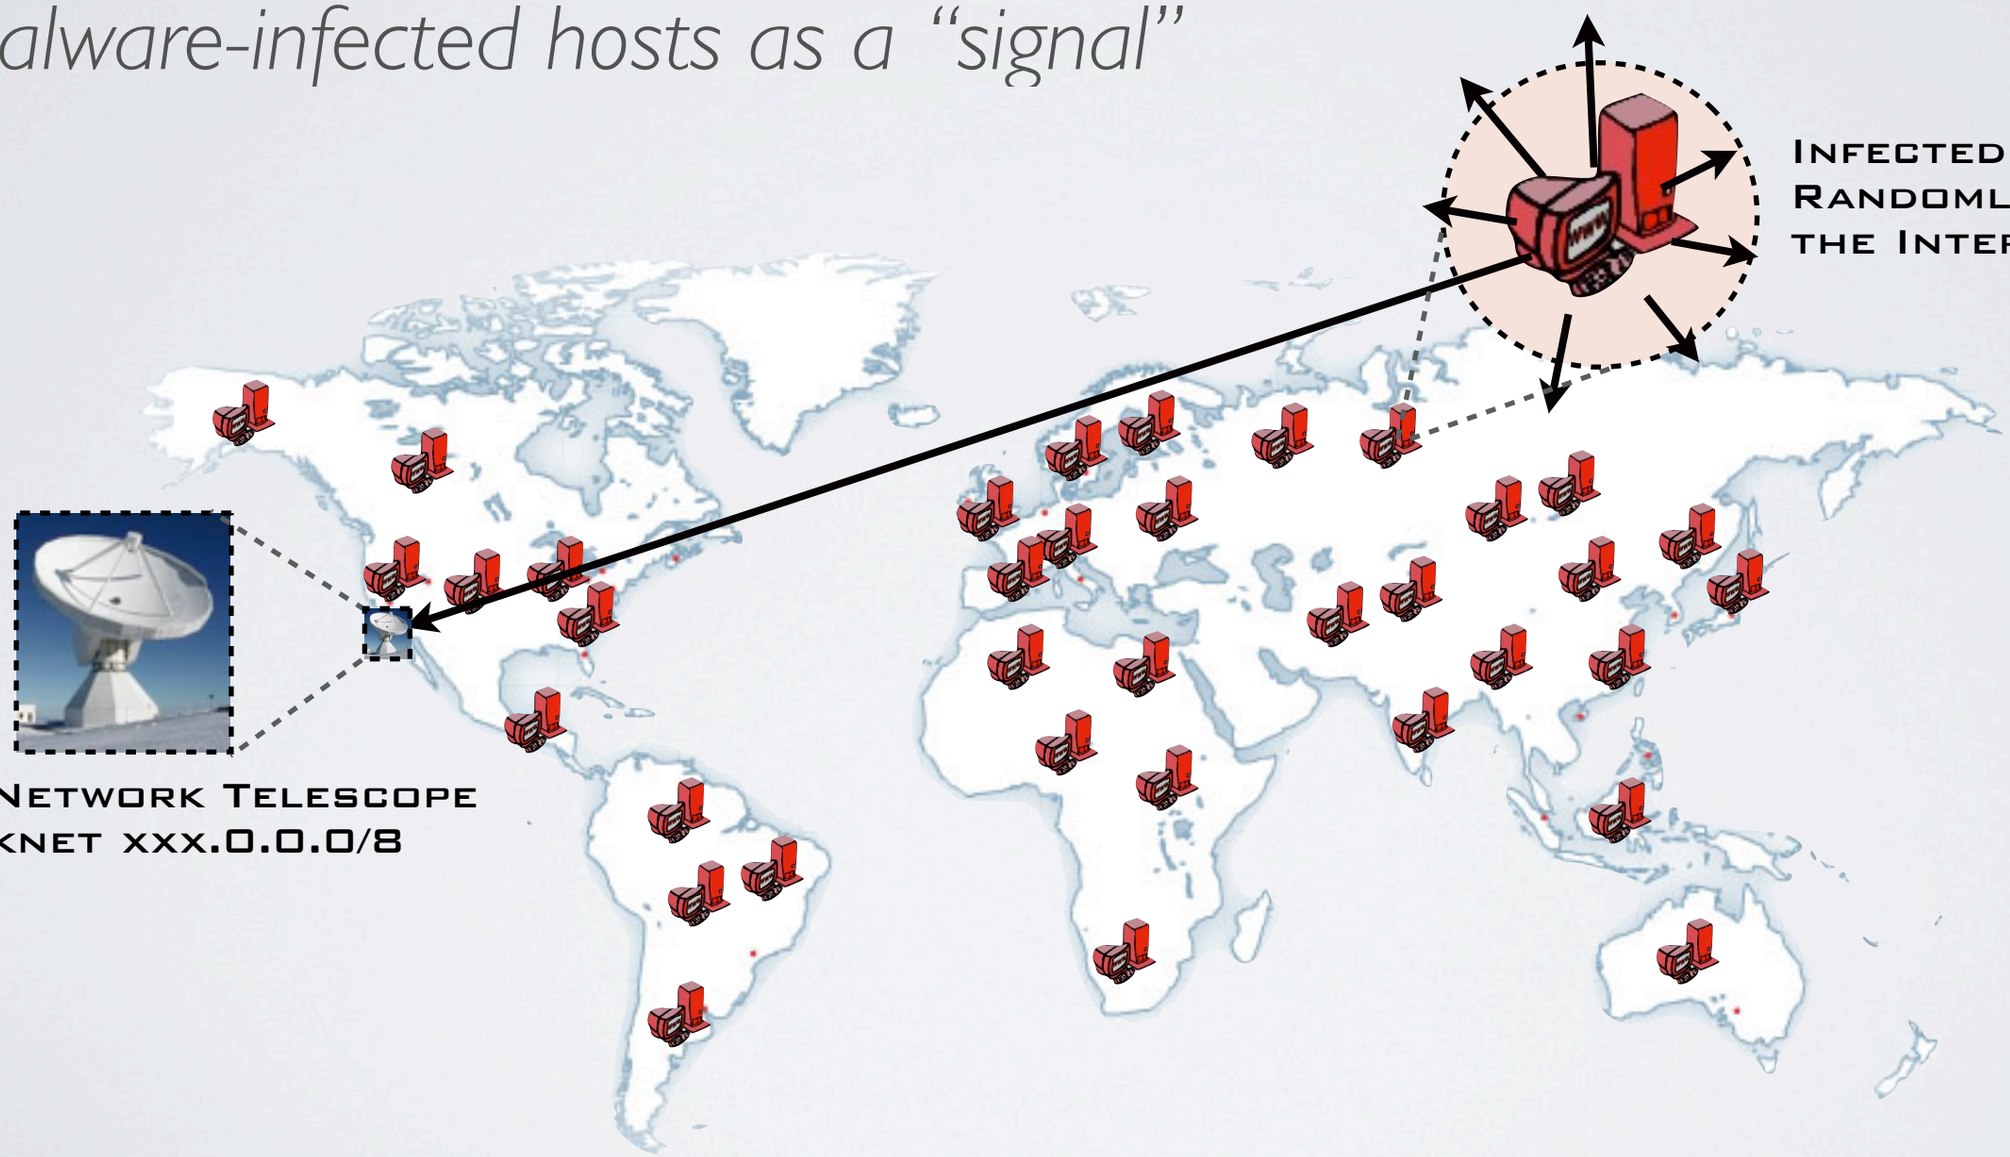

NOVELTY

Using IBR to study Internet Outages

- Revival of Network Telescopes

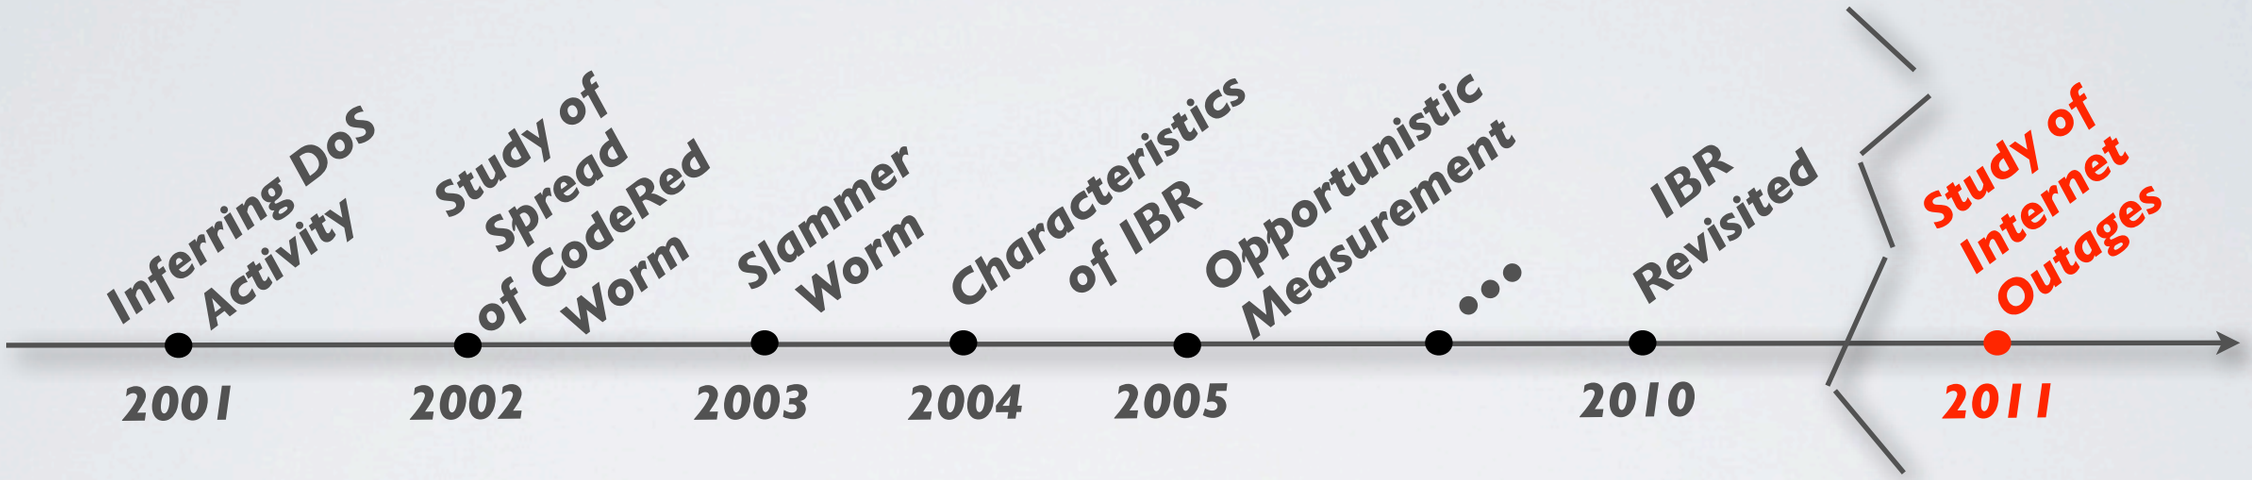

- Alternative/Complementary measurement approaches to study outages
  - *BGP* **[13][28]**
  - *Active Probing* **[20][42]**
  - *Passive Traffic* **[22][24]**
  - *Google services* **[13][14]**
  - *Peer-to-Peer traffic* **[5][6]**

THE EVENTS (1/2)

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:

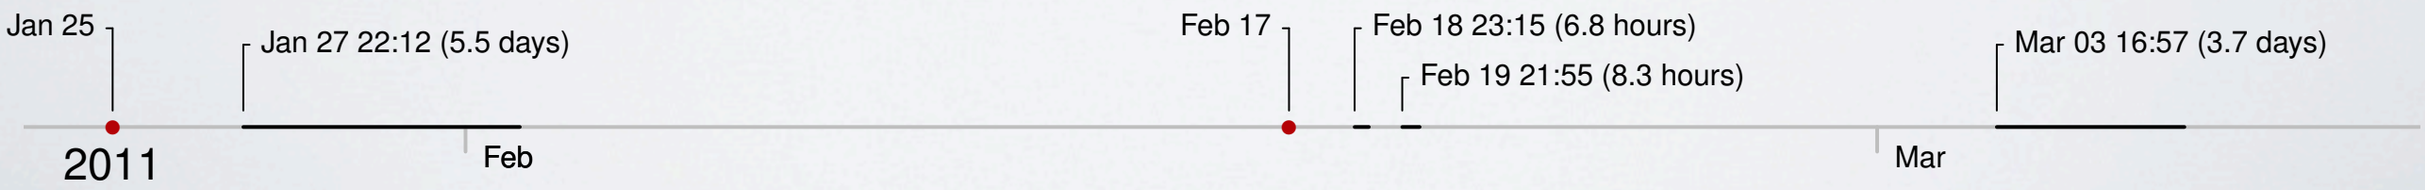

NETWORK INFO

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

EGYPT

IBR: packet rate

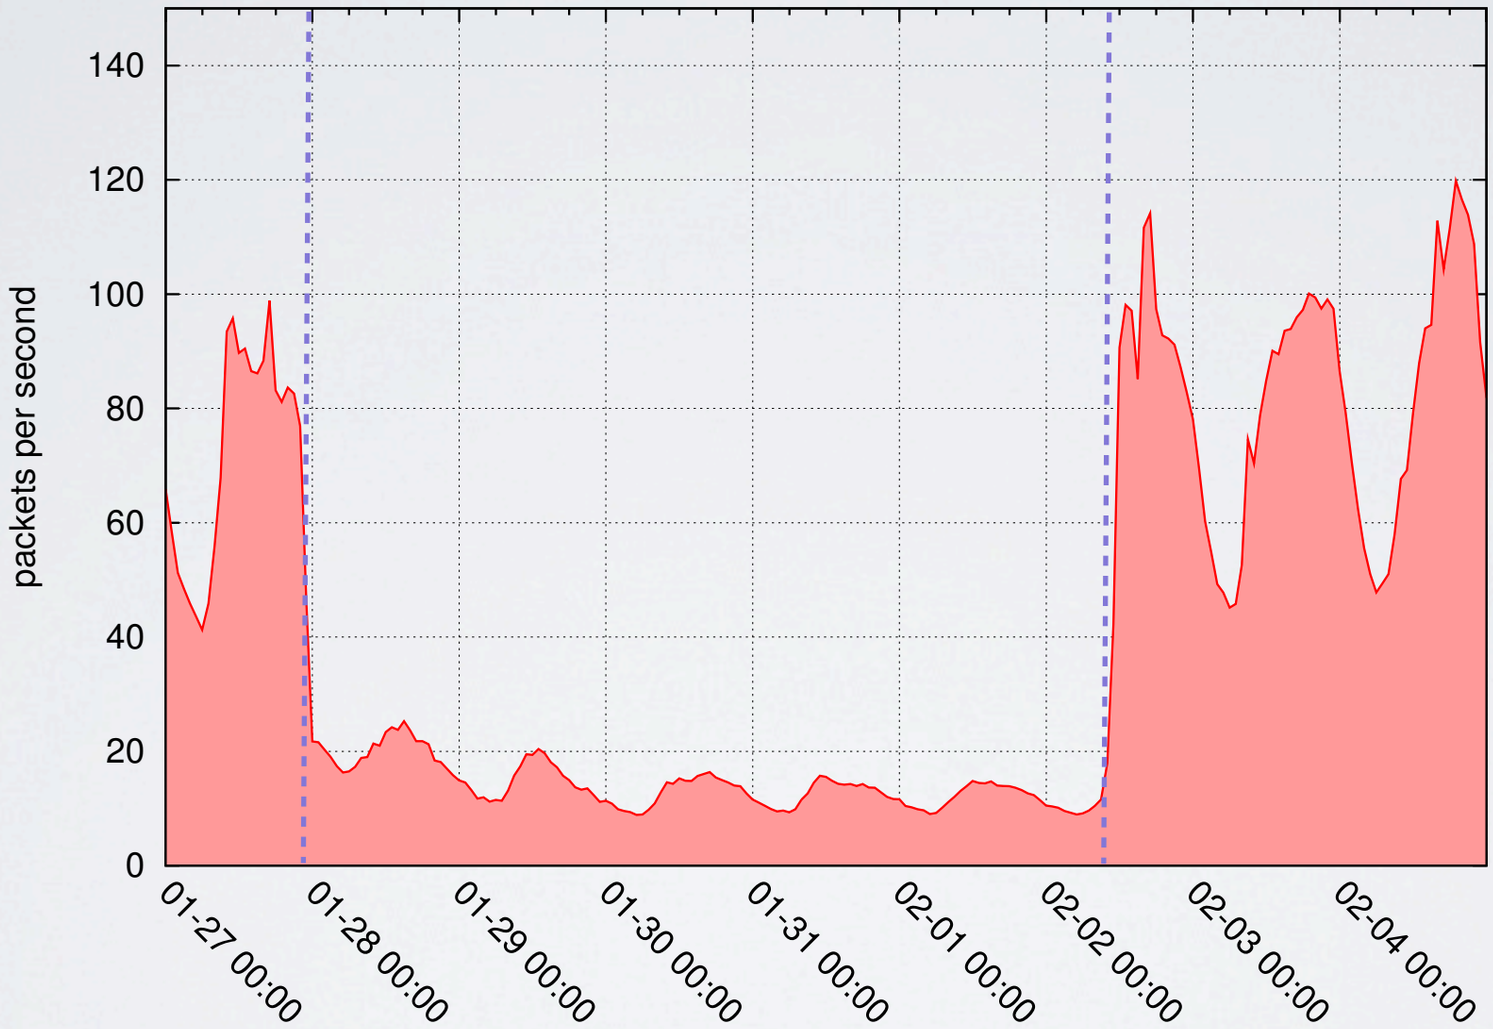

RANDOM PROBING

E.g., Conficker

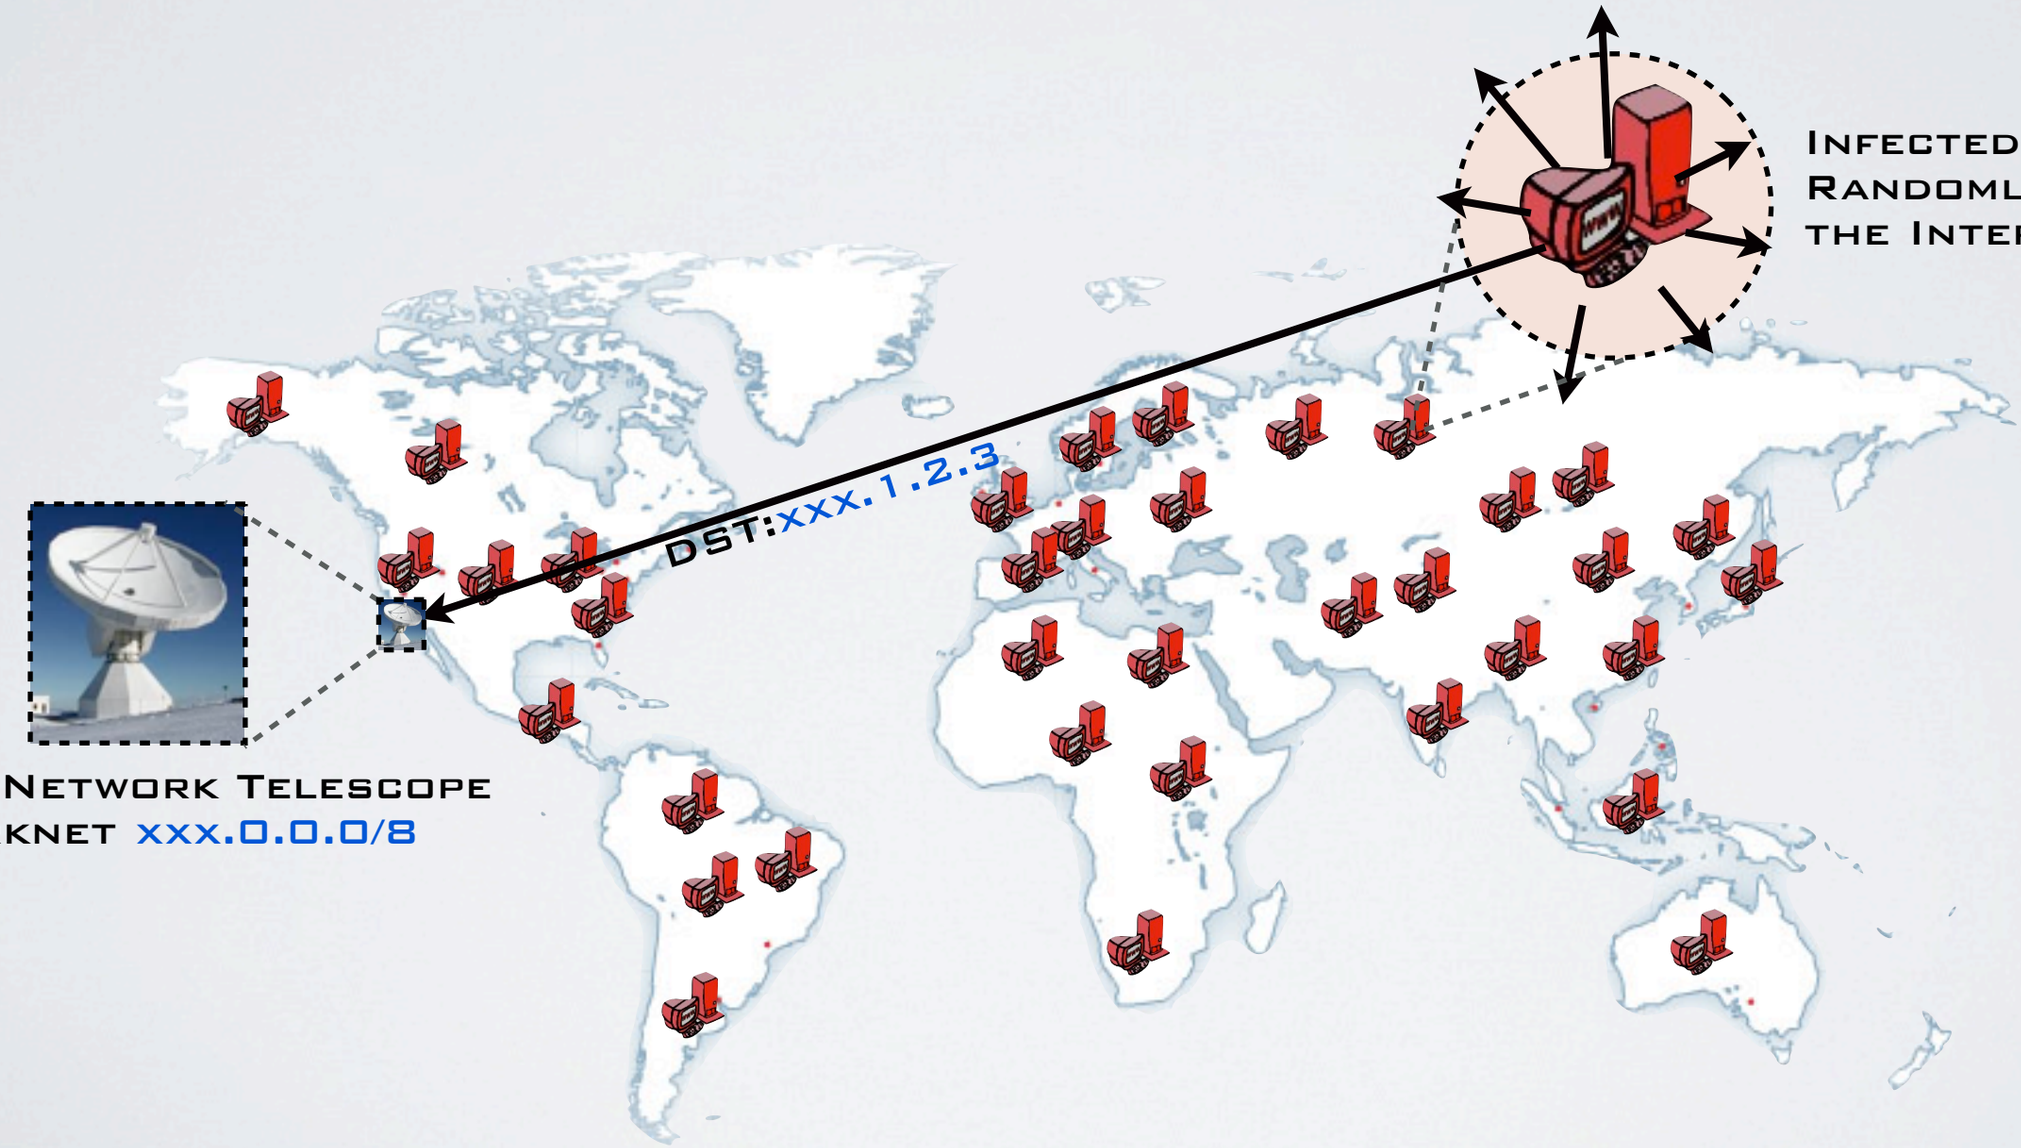

BACKSCATTER

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

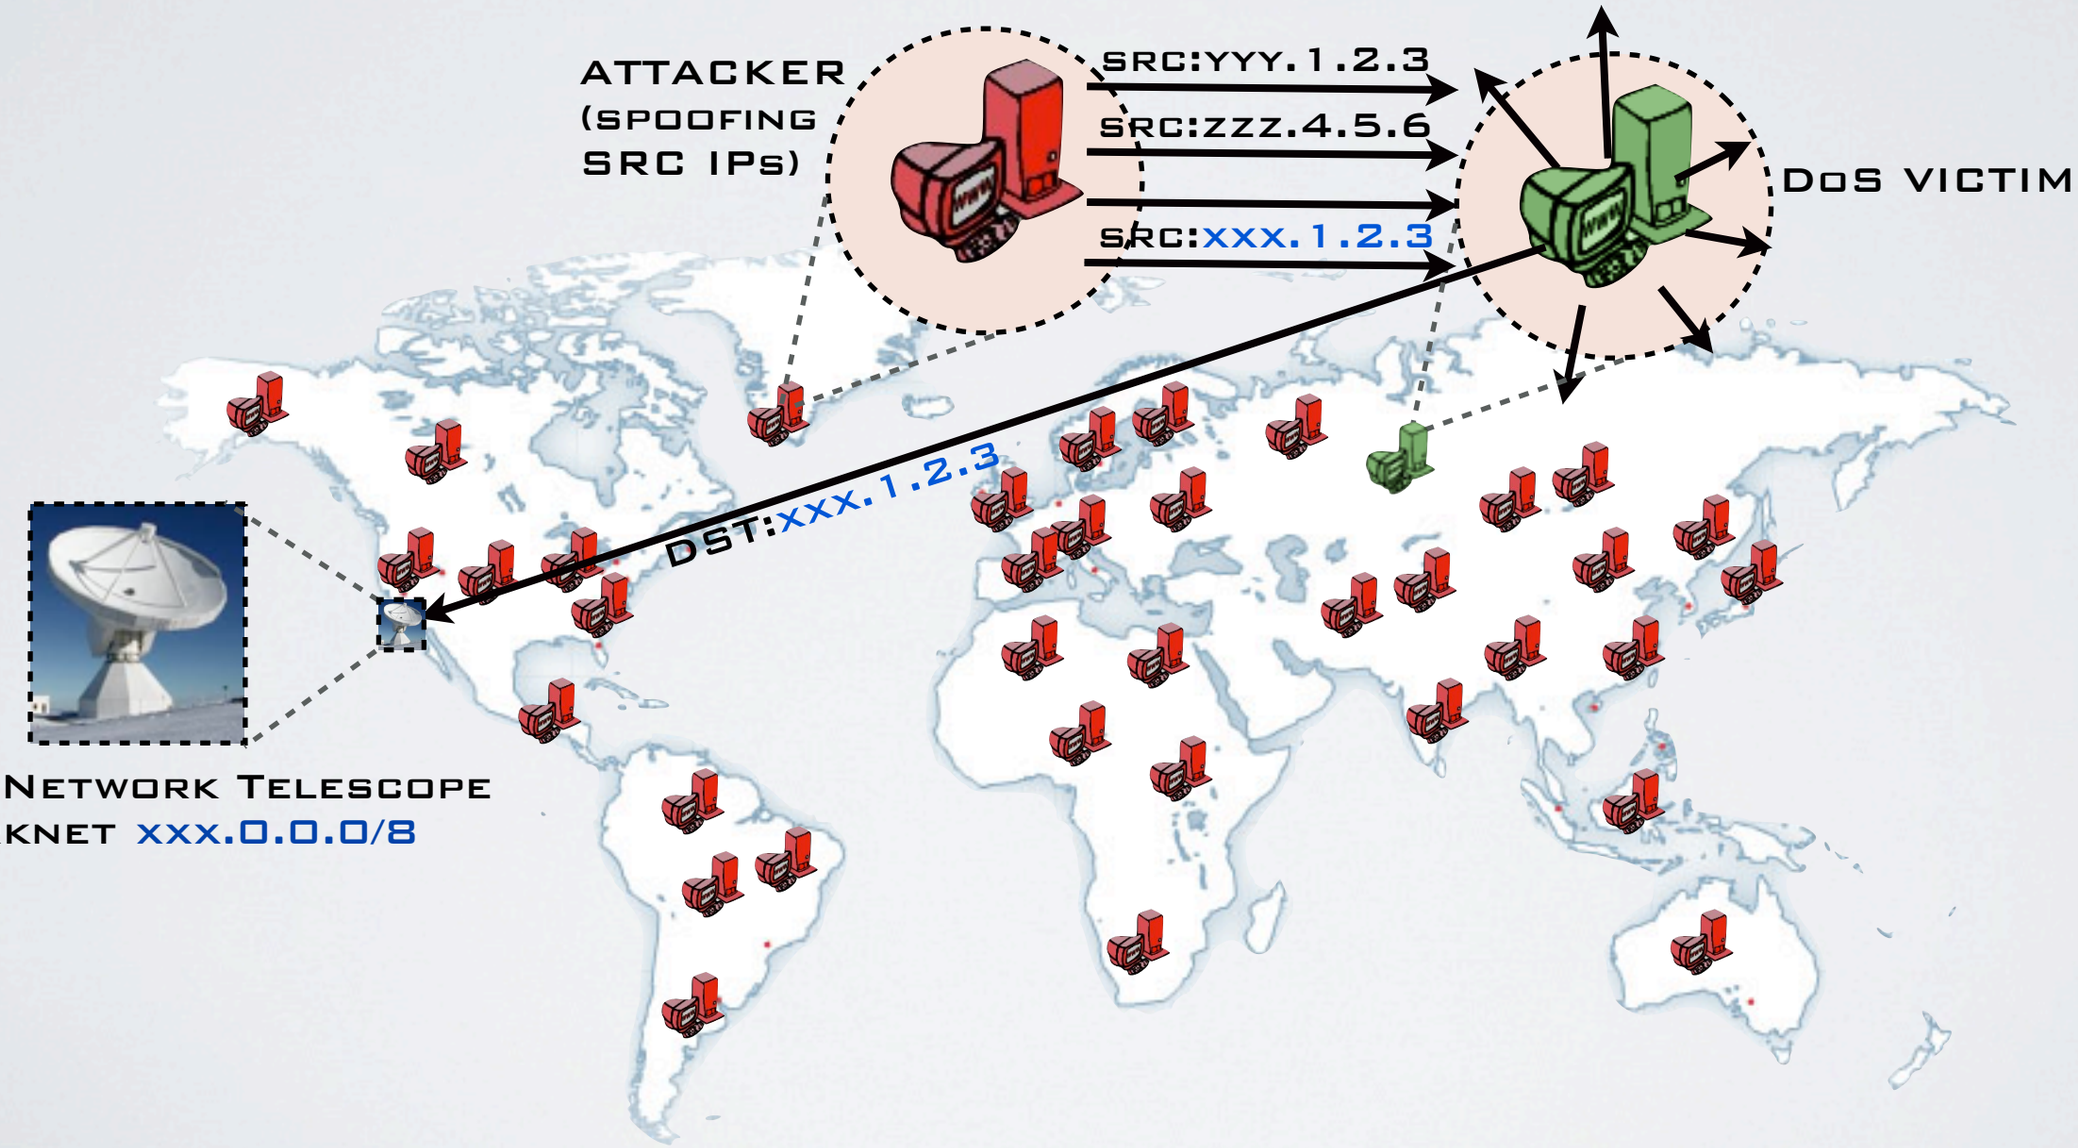

EGYPT

IBR: dissecting it

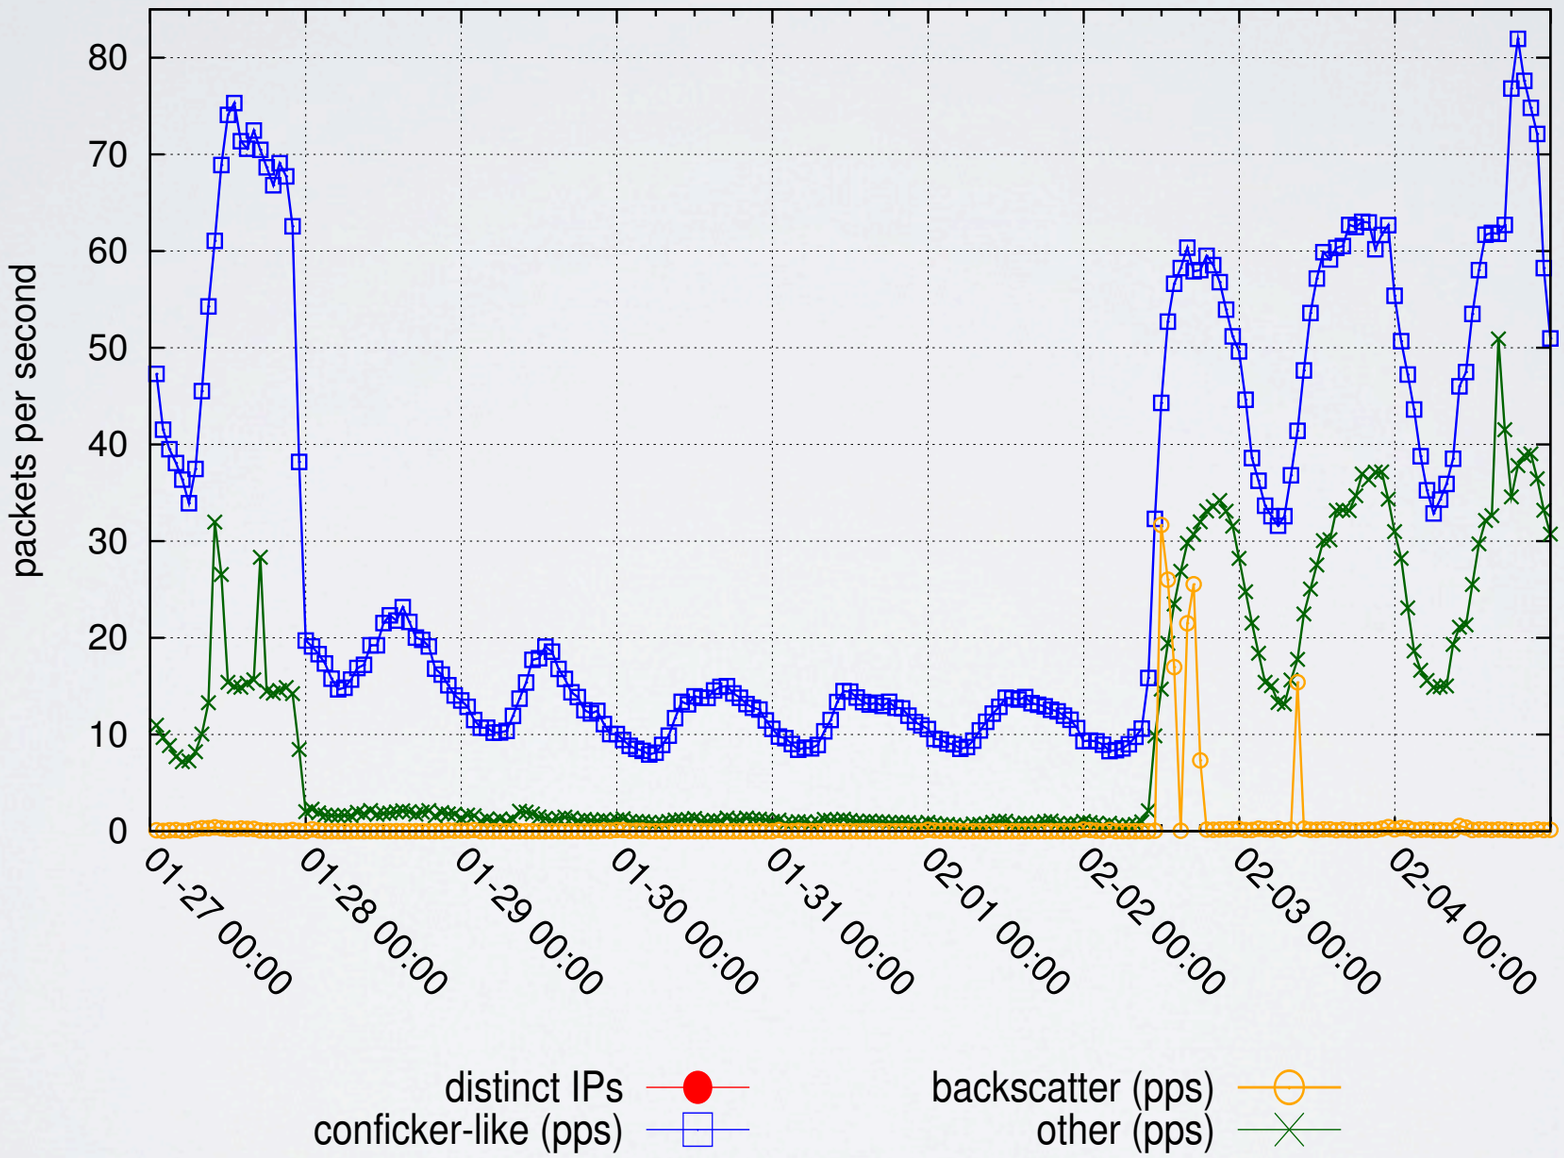

EGYPT

IBR: rate of distinct src IPs vs packet rate

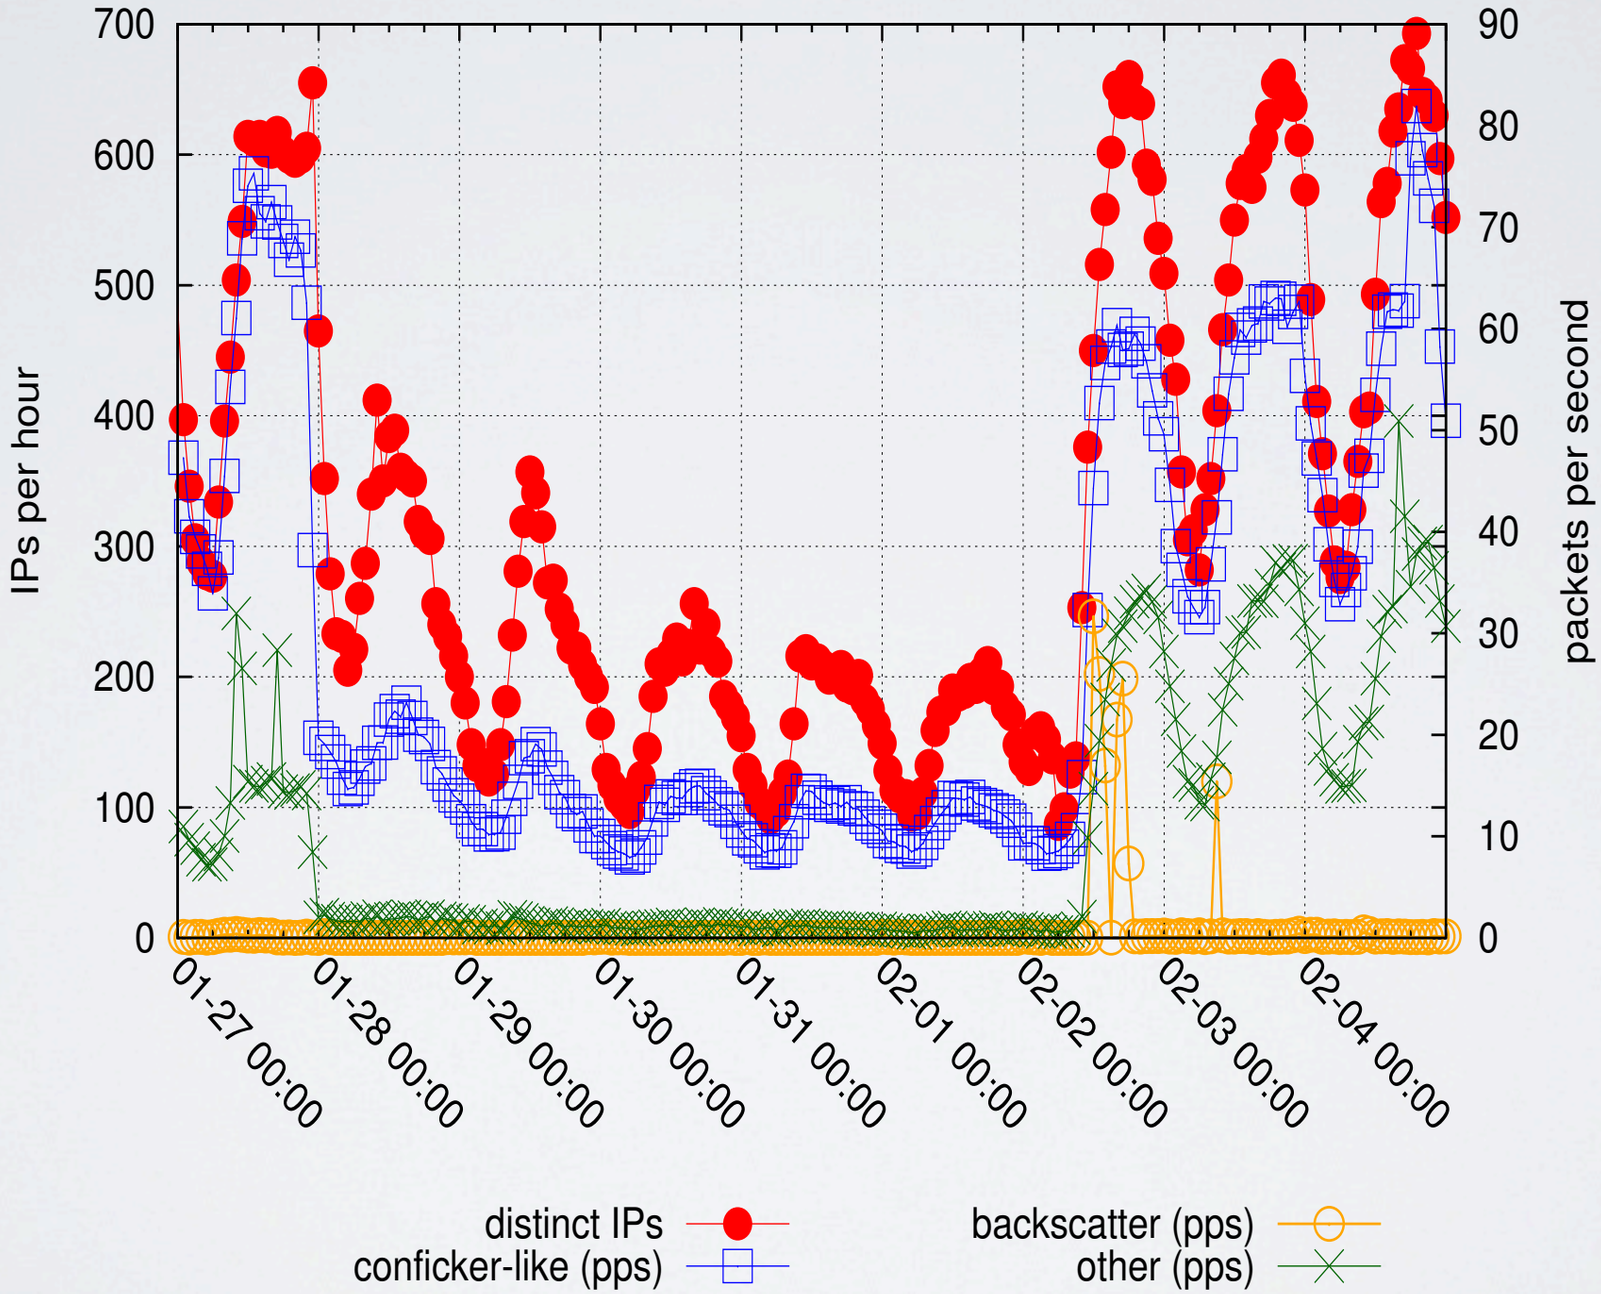

LIBYA

the first two outages

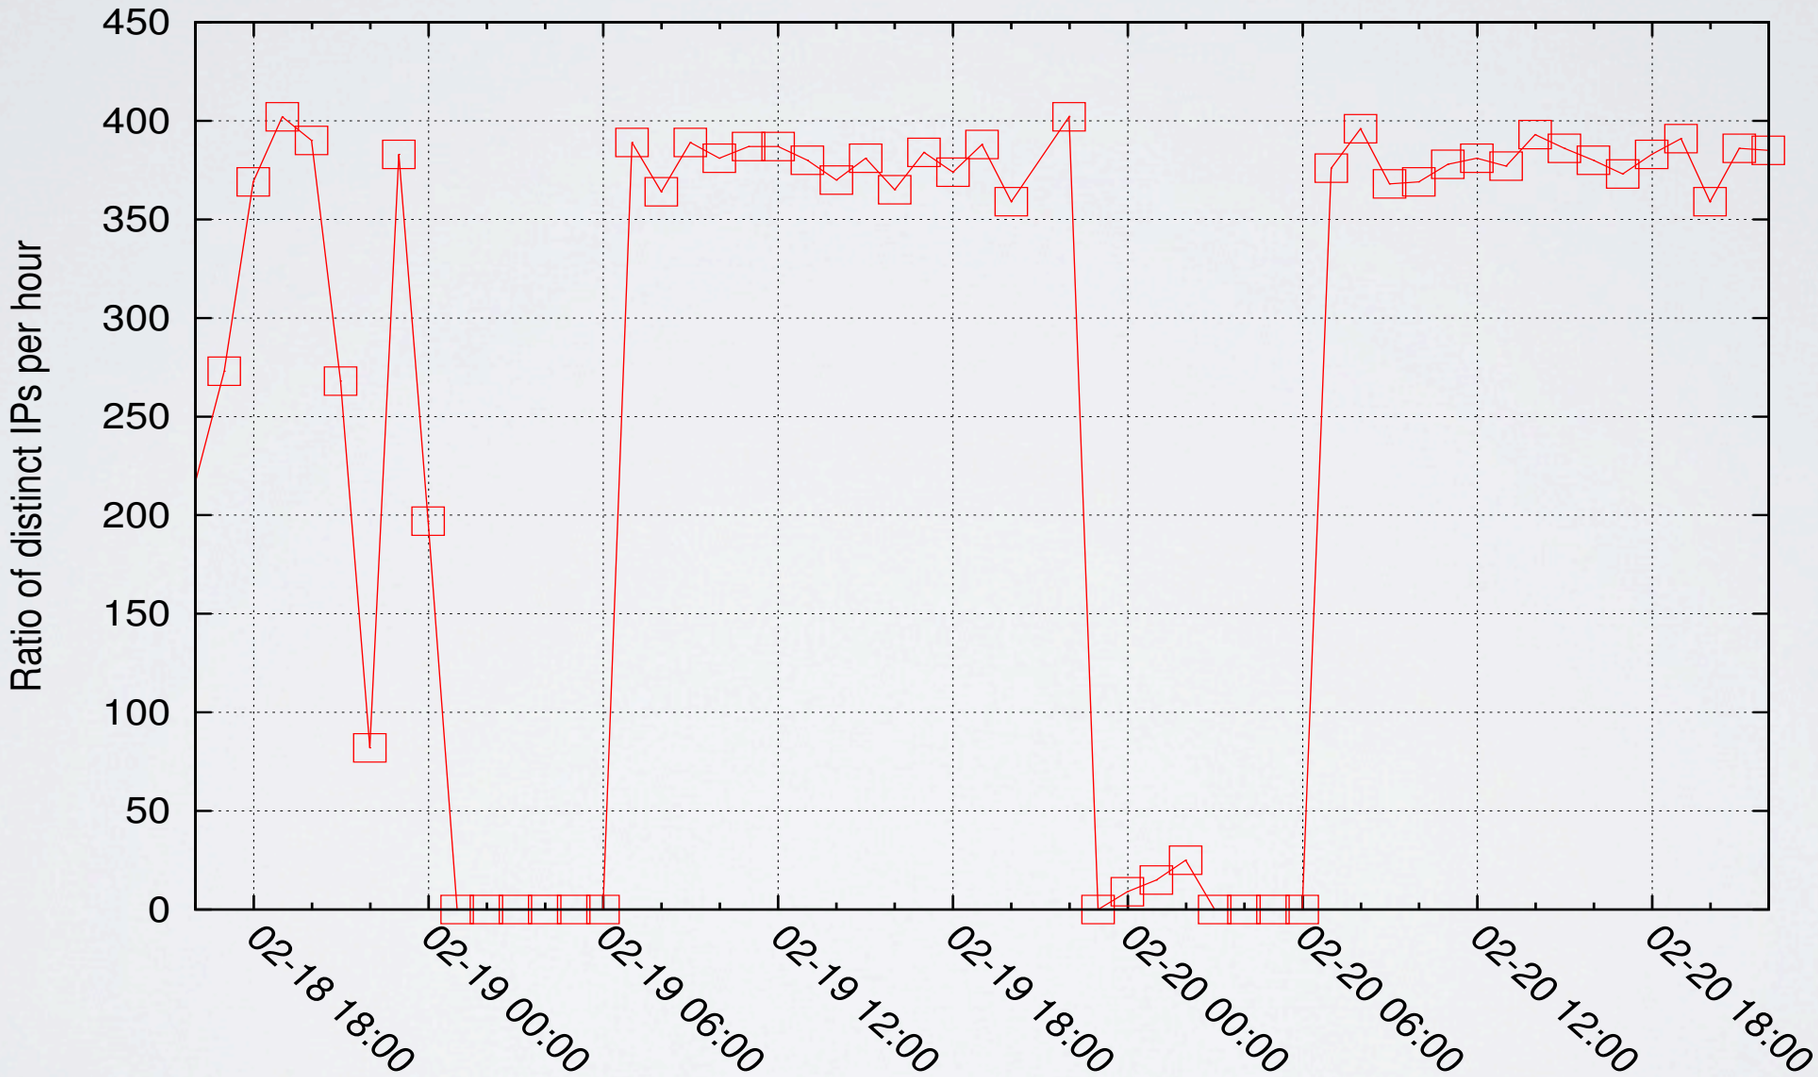

THE EVENTS (2/2)

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

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 Δti,
- $\Delta t_1, ..., \Delta t_n$ 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.

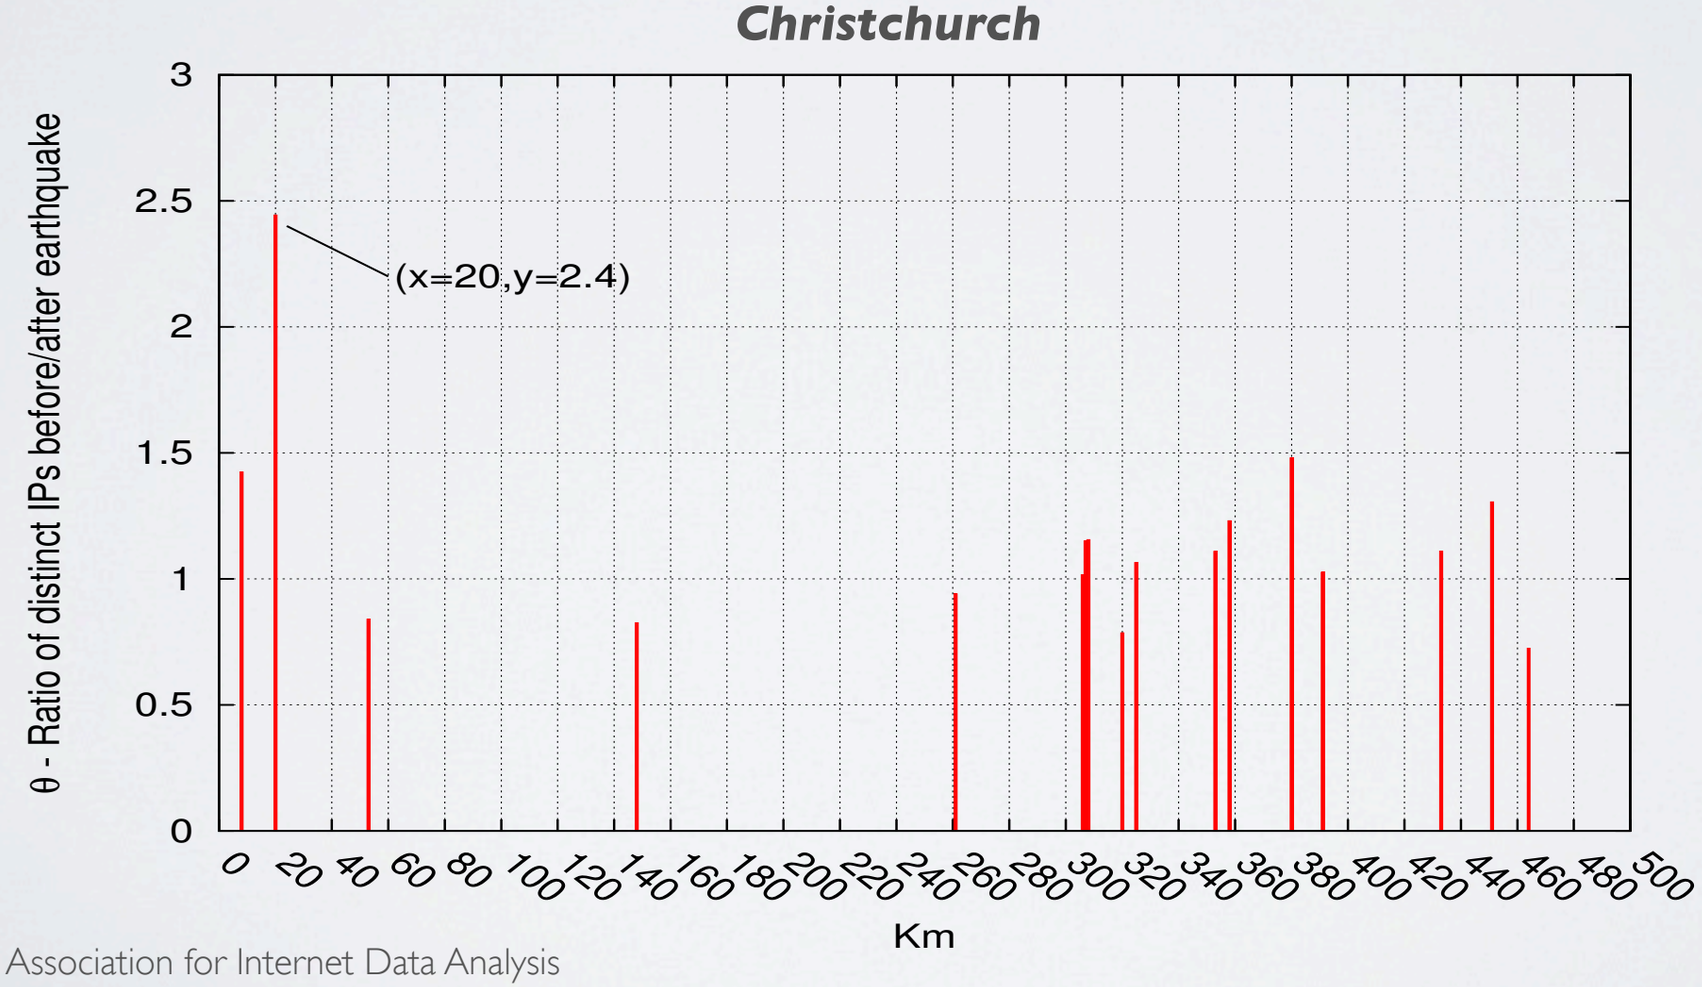

RADIUS OF IMPACT

rough estimate based on θ

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

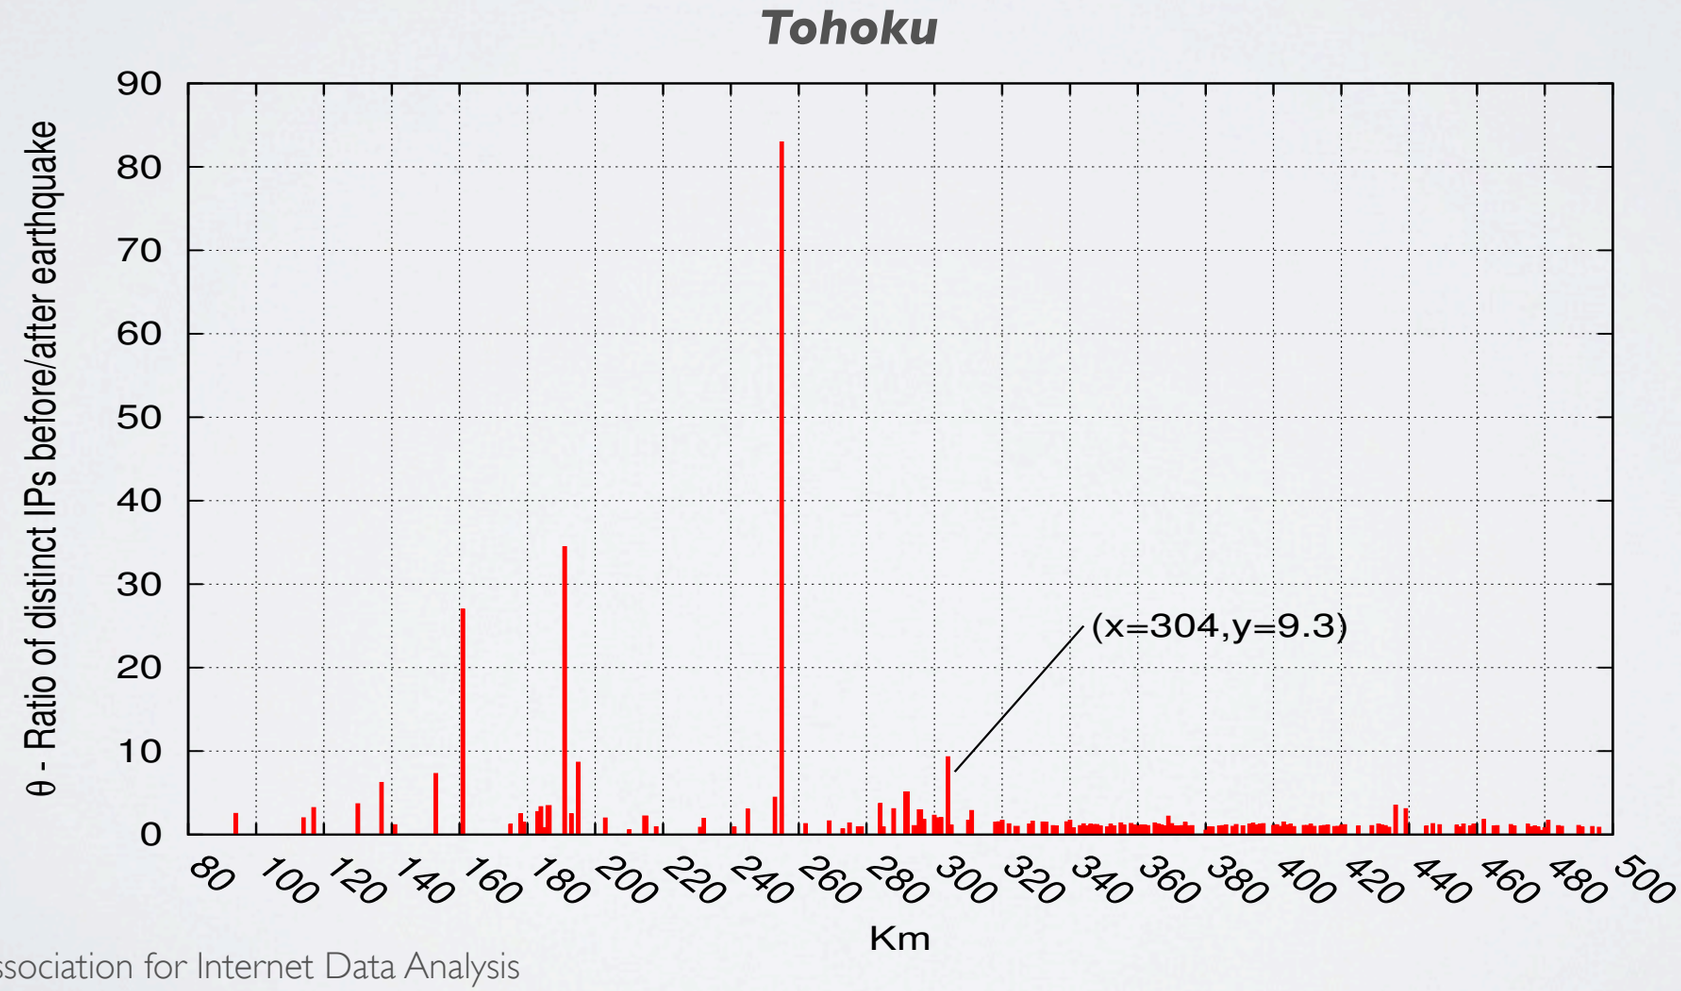

EXTENSION OF IMPACT

geo coordinates of most affected networks

Networks within each respective ρ_{max}

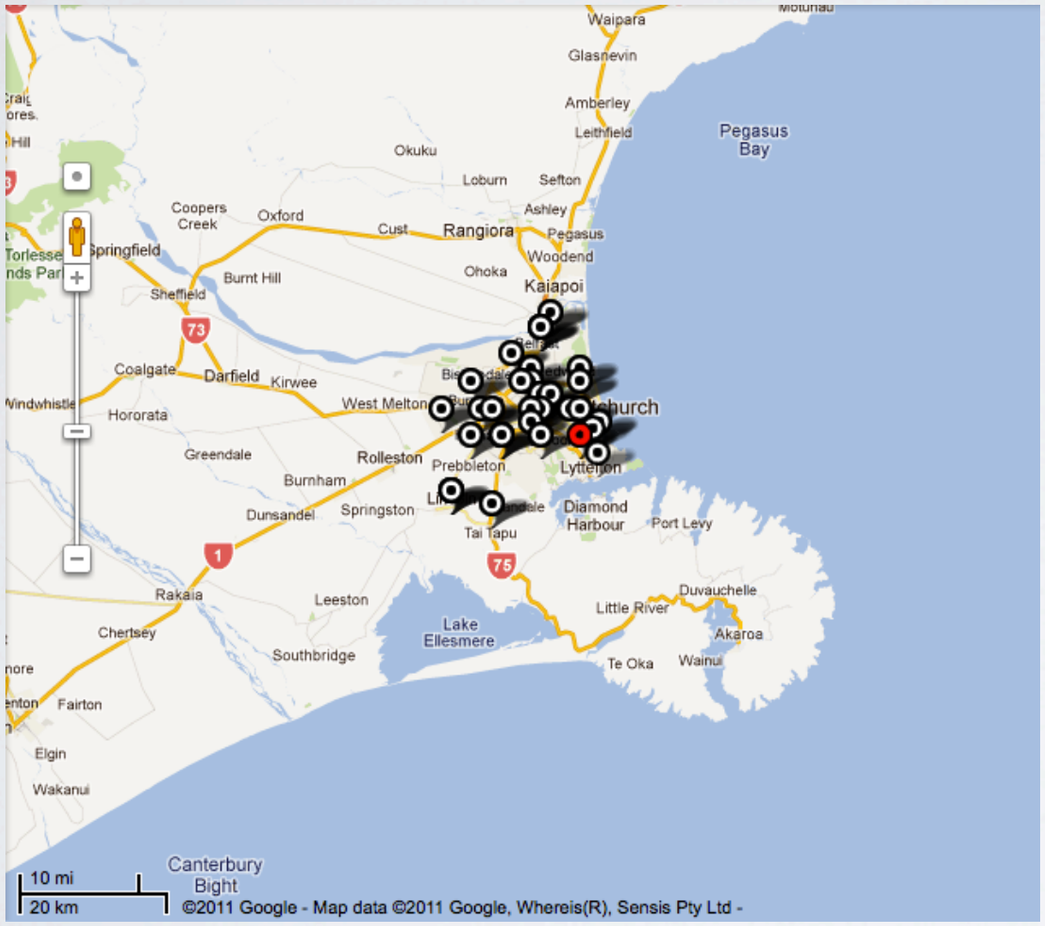

(a) Christchurch

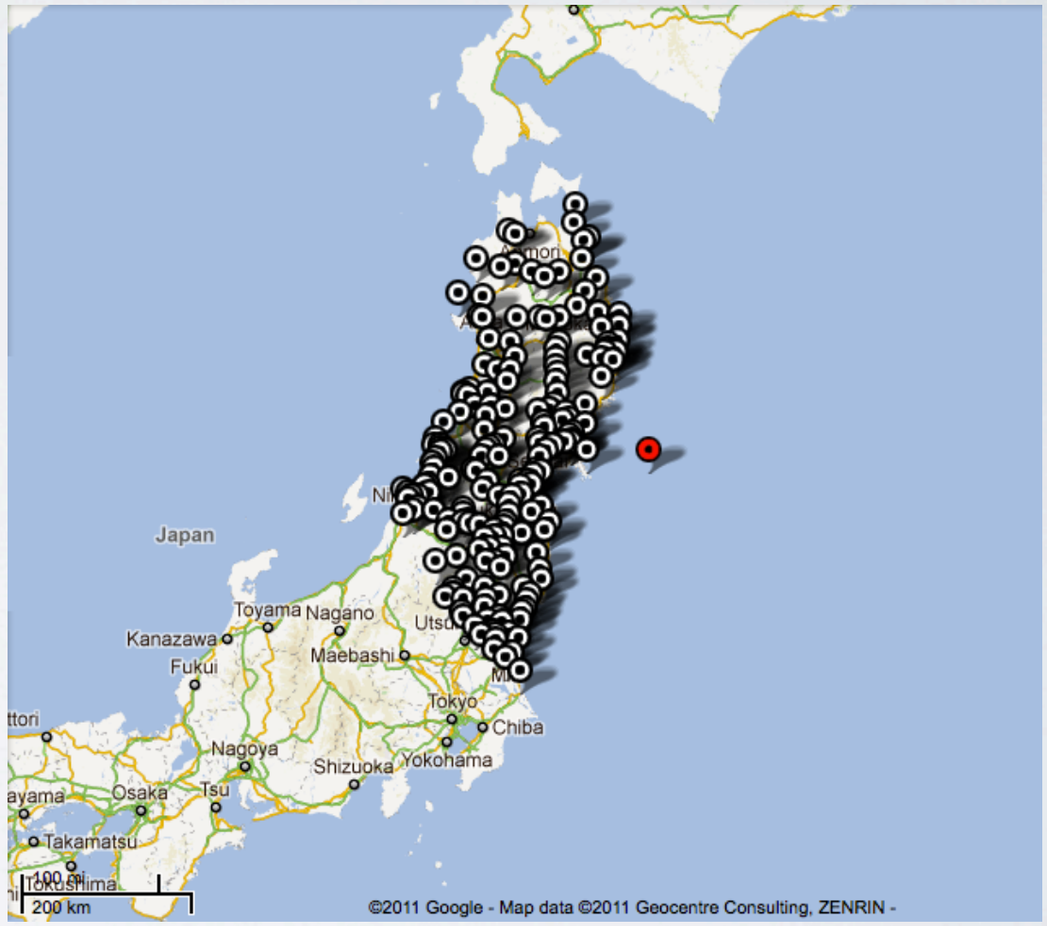

(b) Tohoku

"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

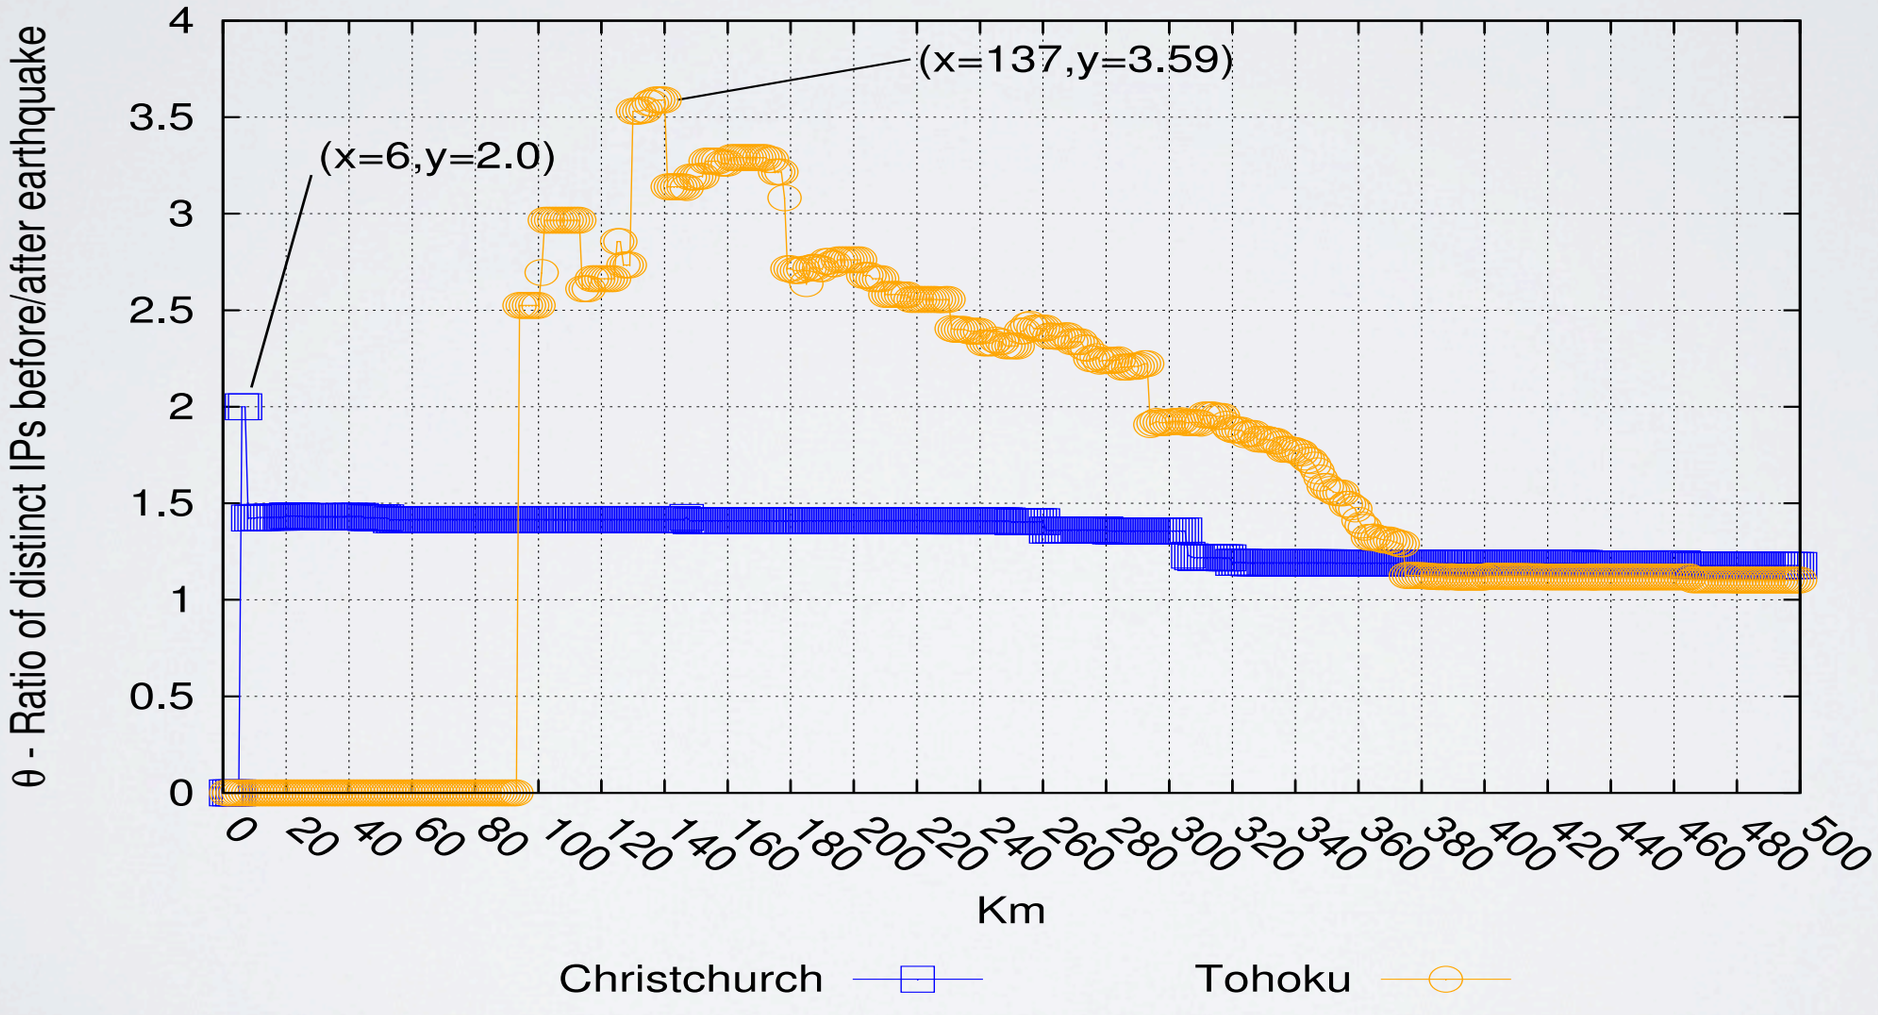

	Christchurch	Tohoku
Magnitude (θ_{max})	2 at $6km$	3.59 at $137km$
Radius (ρ_{max})	$20km$	$304km$

IP RATE IN TIME

reflects the dynamics of the event

Christchurch

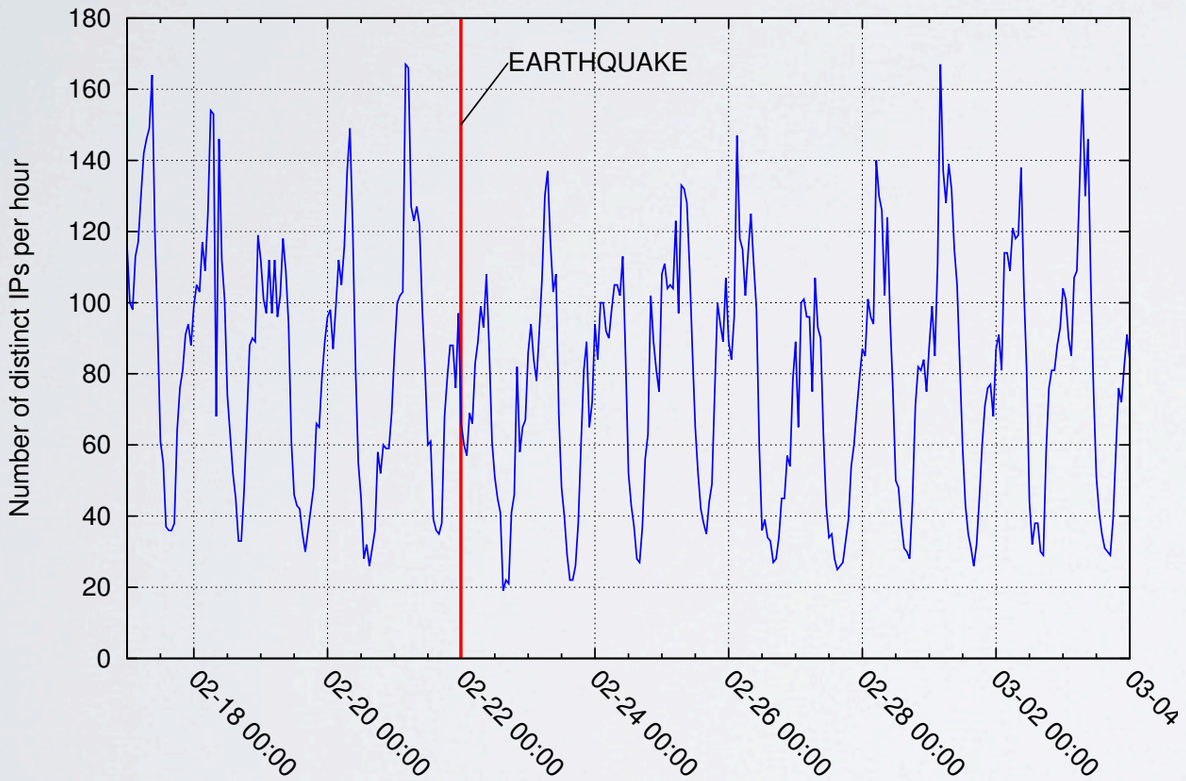

Tohoku

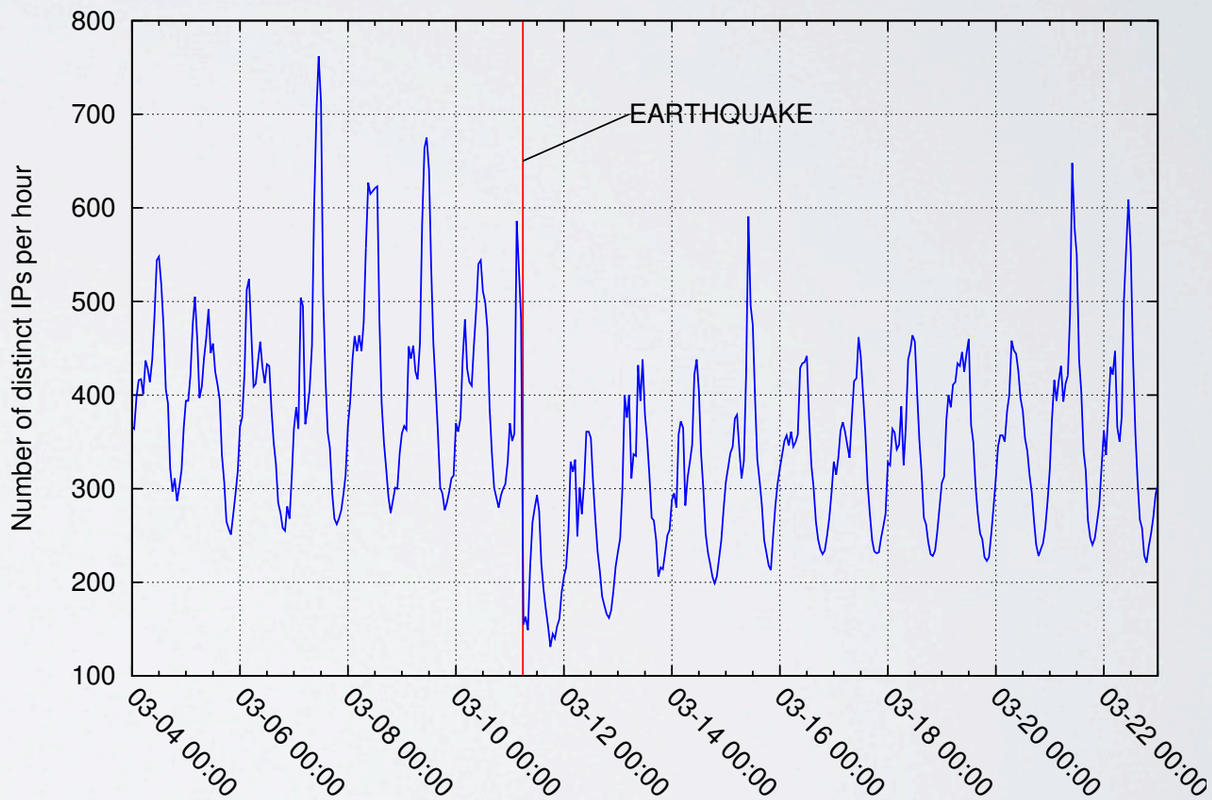

EVALUATING Θ

variations over a long time period

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

Christchurch

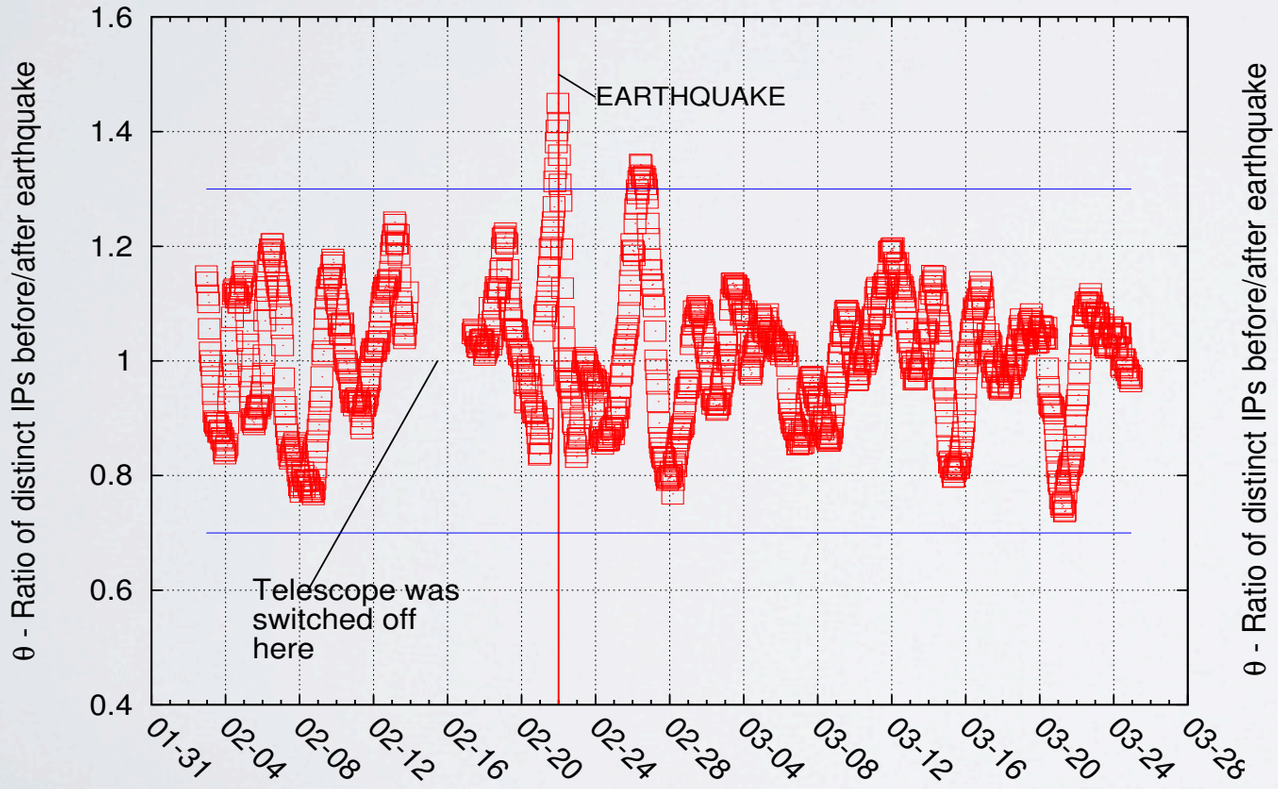

Tohoku

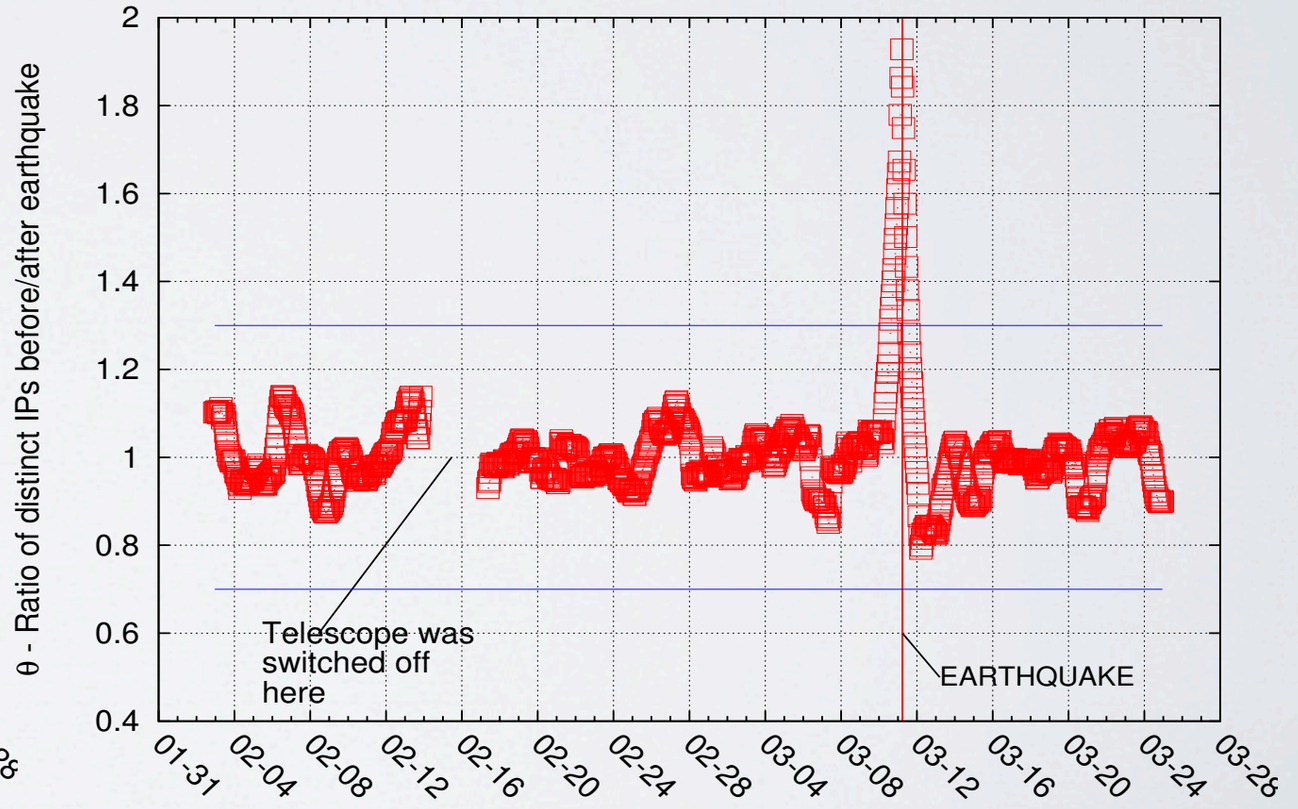

CONCLUSION

ongoing work

- IBR is an effective source of data for the analysis of network outages caused by events of different type

- Future work
  - Integrate and combine analysis of multiple data sources (BGP, IBR, active measurement, ...)
  - Analysis of AS/Link-level topology
  - Automated detection + triggered active measurements

THANKS

Cooperative Association for Internet Data Analysis
University of California San Diego