BADGERS 2012
15 October, 2012 - Raleigh, NC, USA

Analysis of Internet-wide Probing using Darknets

A. Dainotti, A.J. King, K.C. Claffy
alberto@caida.org
CAIDA - University of California, San Diego

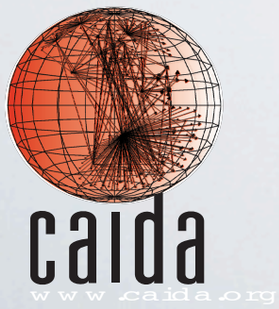

DISCLAIMER

more questions than answers!

GOAL OF THE TALK

1. Point you at our finding

2. Report on our experience in analyzing it

3. QUESTION: How can we detect similar events?

4. Half of a proposal: collaborative (large-)data sharing

Cooperative Association for Internet Data Analysis
University of California San Diego

DARKNETS

a.k.a. Network Telescopes

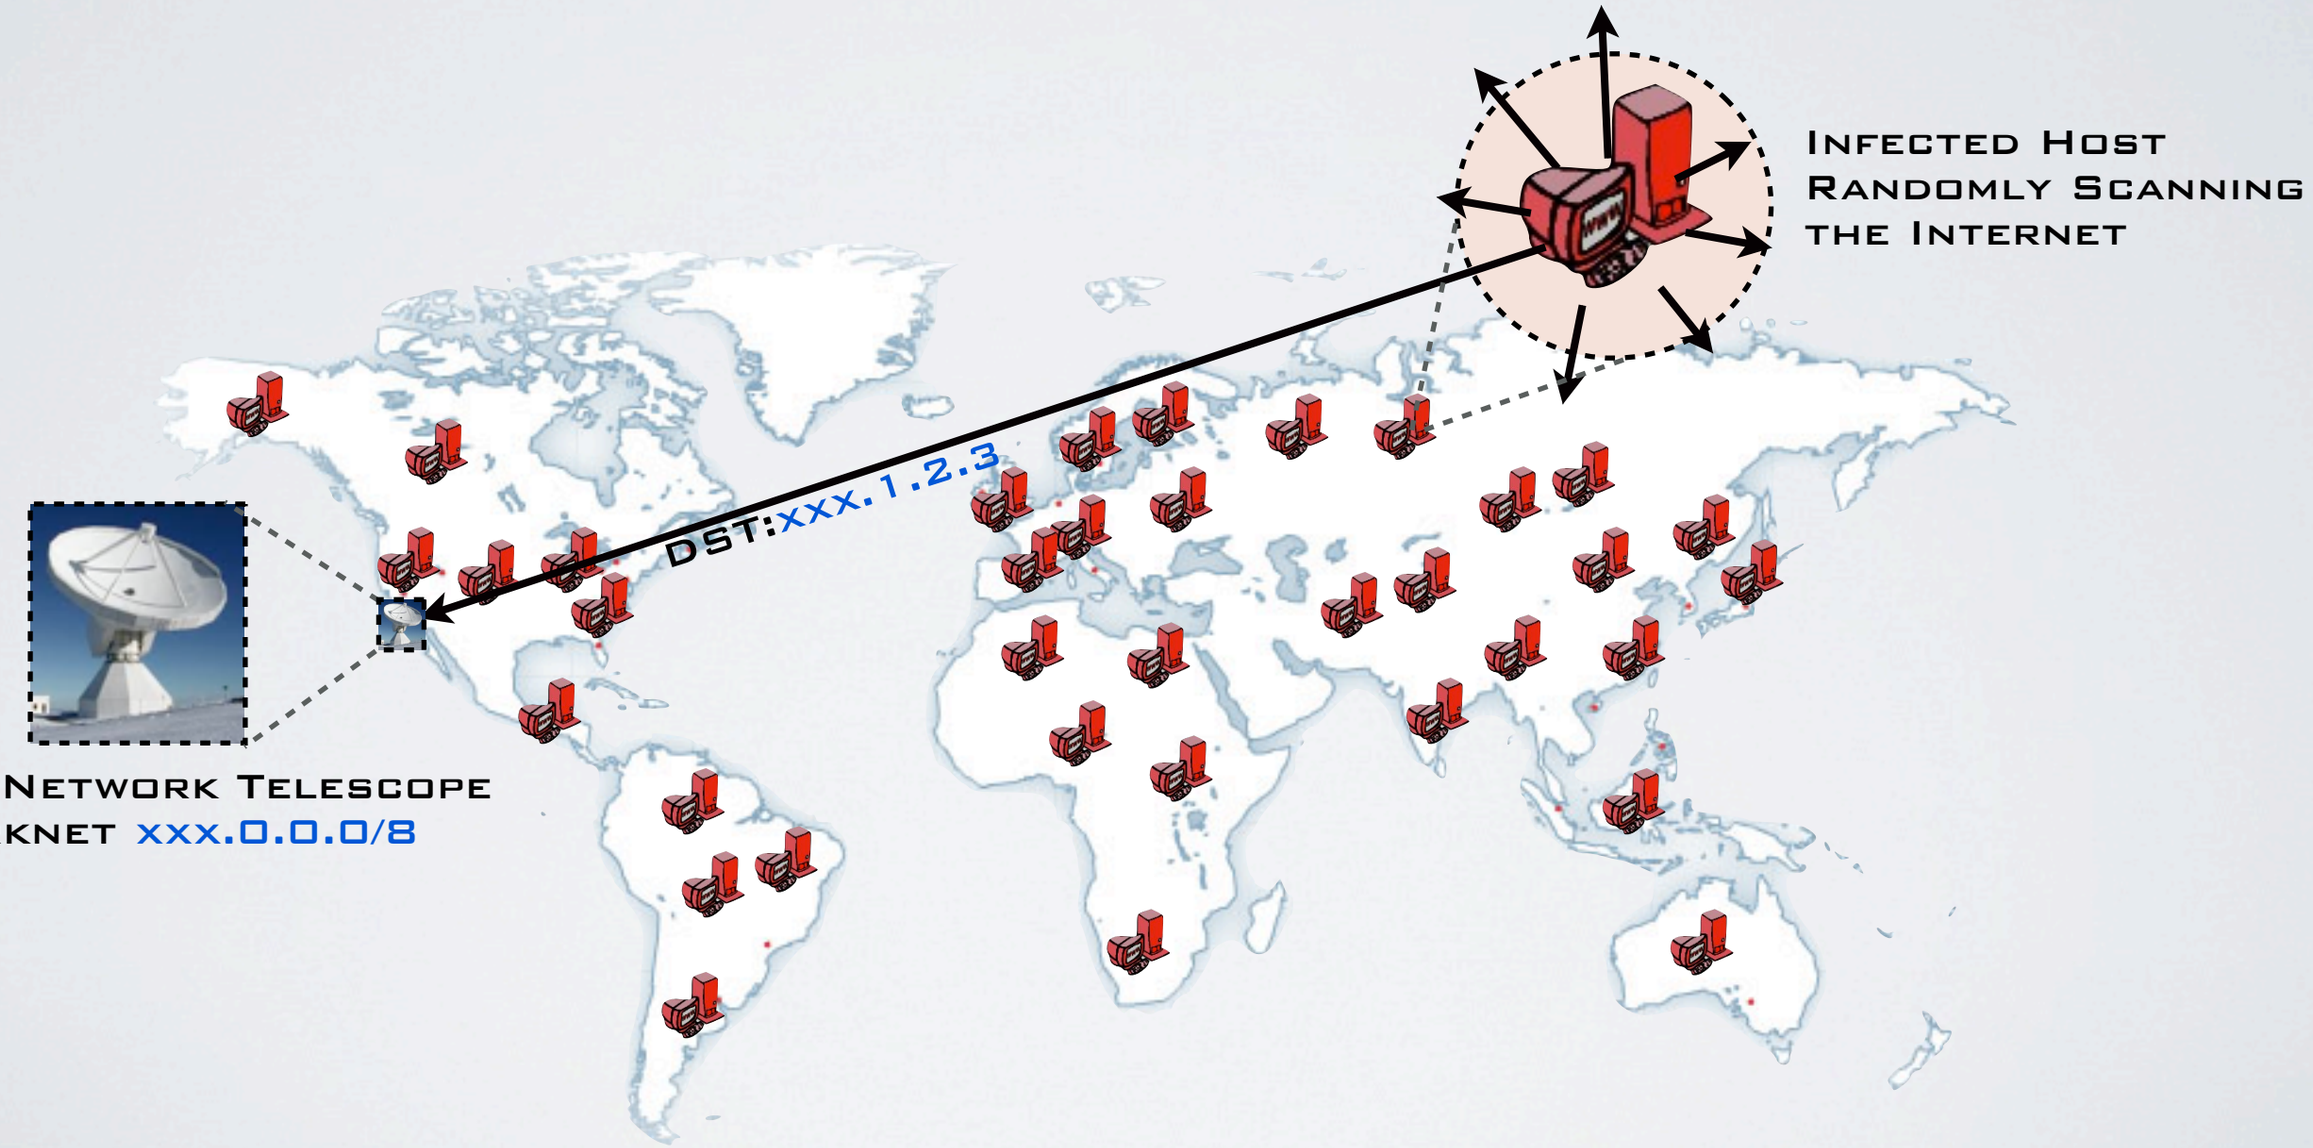

SIPSCAN

Feb 2011

- A "/0" scan from a botnet
- Observed by the UCSD telescope (a /8 darknet)
- Scanning SIP servers with a query on UDP port 5060

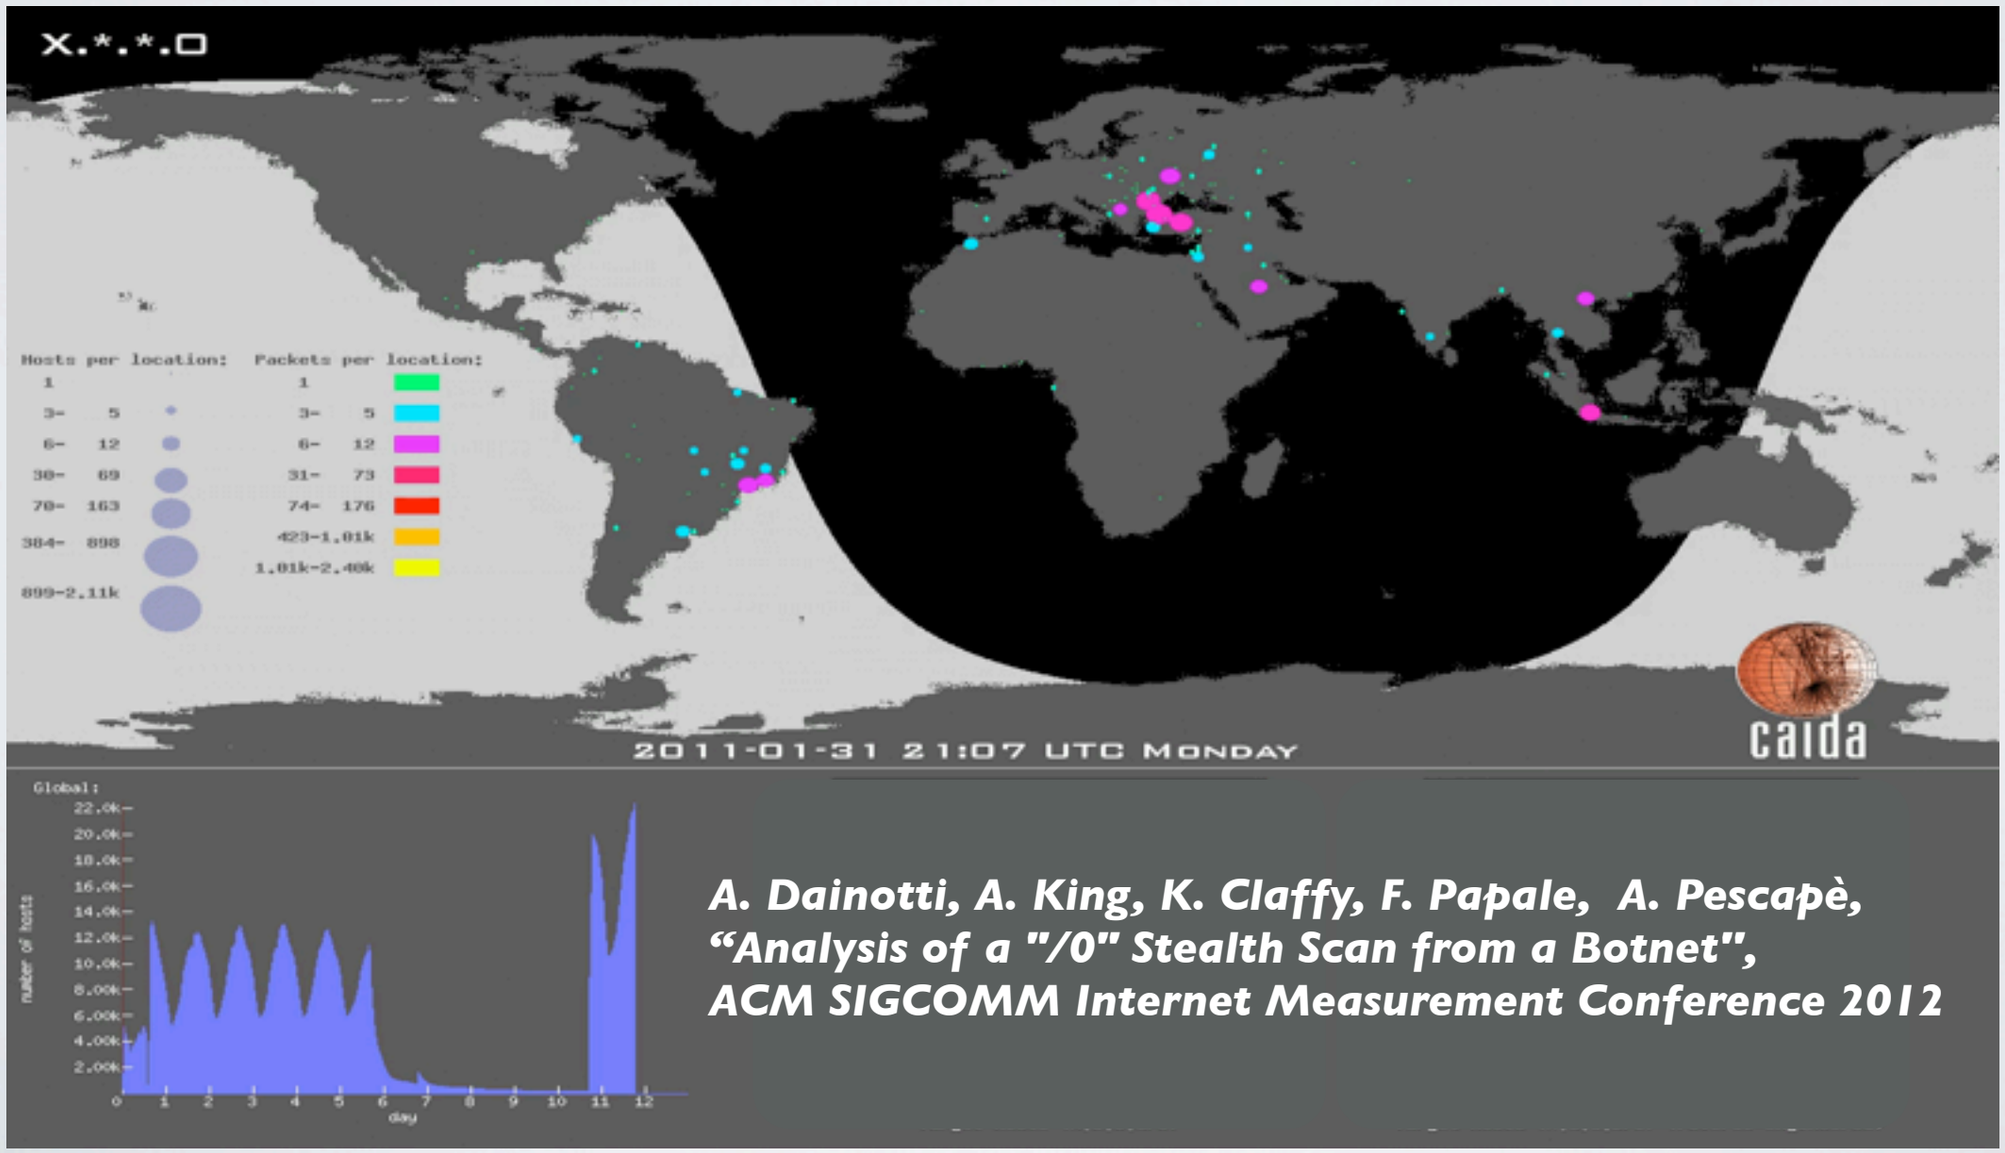

A. Dainotti, A. King, K. Claffy, F. Papale, A. Pescapè, "Analysis of a "/0" Stealth Scan from a Botnet", ACM SIGCOMM Internet Measurement Conference 2012

SIPSCAN

Why so interesting?

- Probing the entire IPv4 address space (in 12 days)
- Great coordination: small overlap with good coverage
- Stealth!
  - Large bots turnover
  - Reverse byte order in the progression of target IPs

SIPSCAN

Why so interesting?

- Probing the entire IPv4 address space (in 12 days)
- Great coordination: small overlap with good coverage
- Stealth!
  - Large bots turnover
  - Reverse byte order in the progression of target IPs

SERENDIPITY

the "Egyptian Killswitch" (Feb 2011)

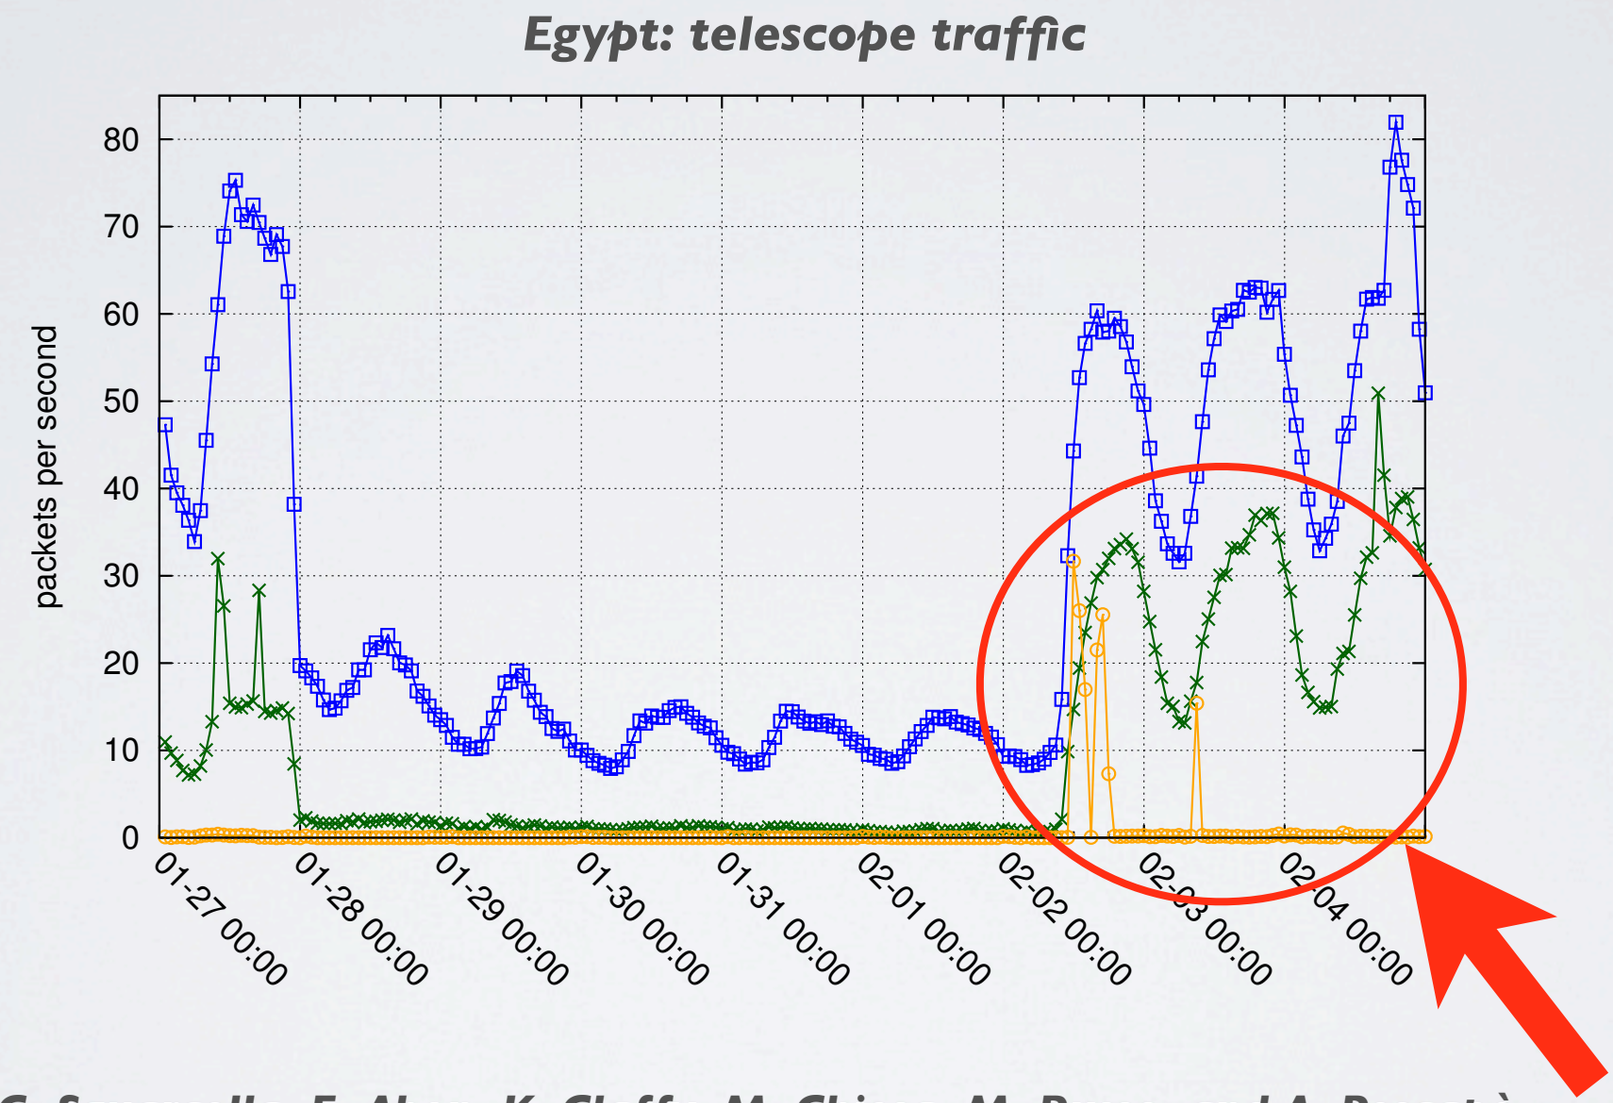

A. Dainotti, C. Squarcella, E. Aben, K. Claffy, M. Chiesa, M. Russo, and A. Pescapè, "Analysis of Country-wide Internet Outages Caused by Censorship", ACM SIGCOMM Internet Measurement Conference 2011

SIPSCAN

isolating the "SipScan"

- Thanks to the unique payload fingerprint we could isolate it without inferences

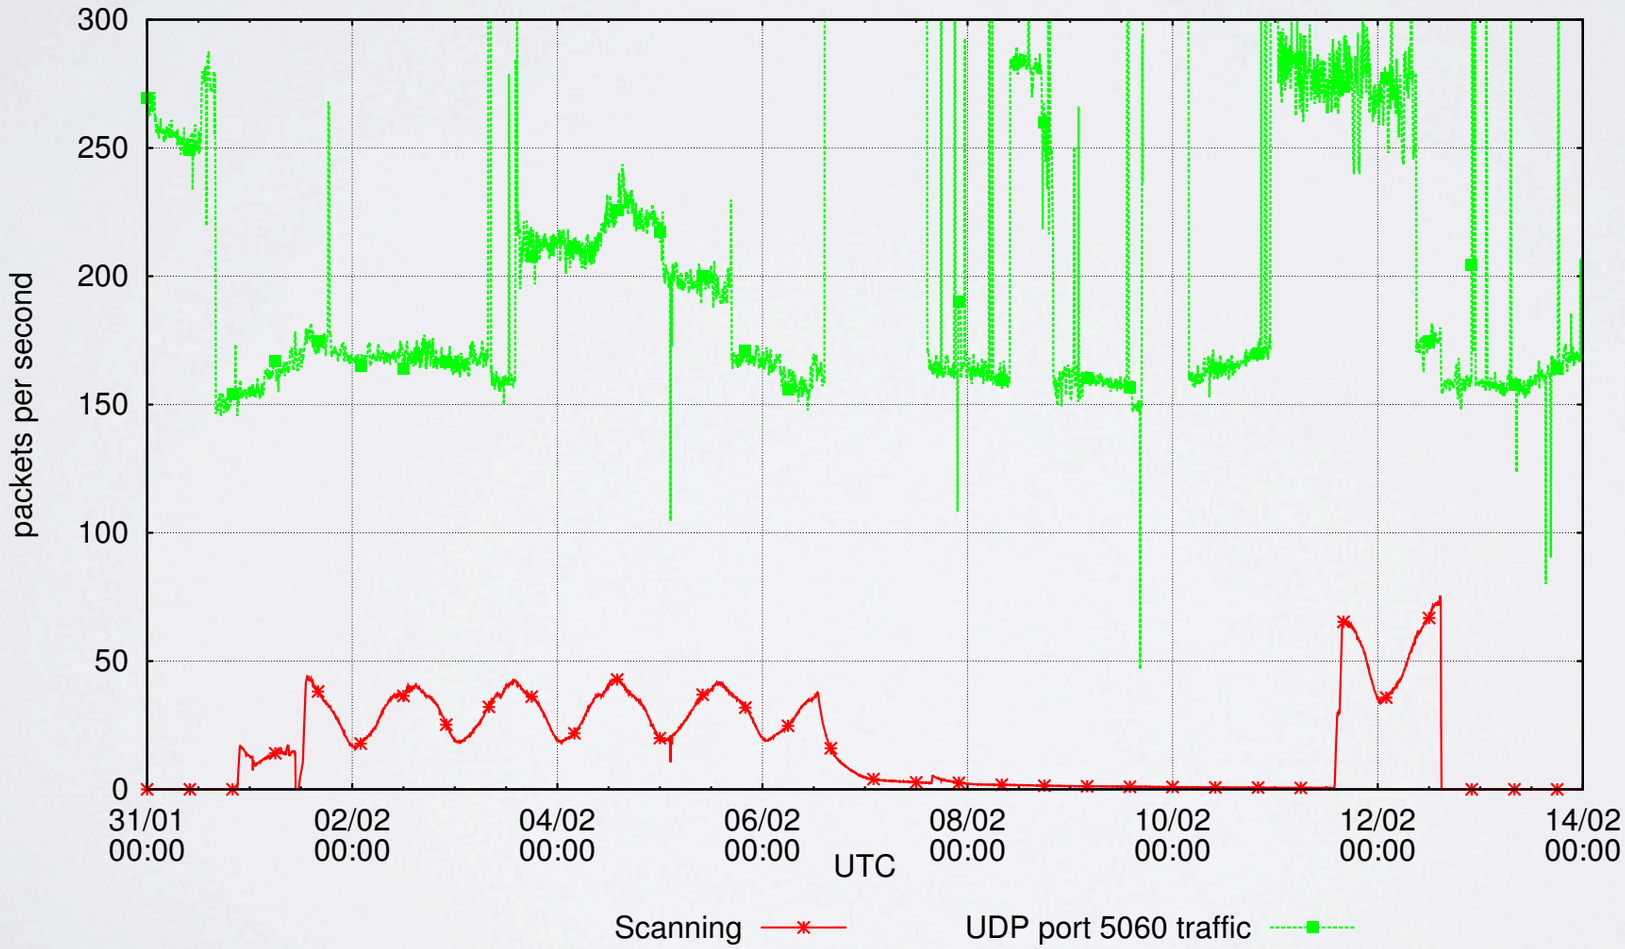

DO OTHER *"SIPSCAN-LIKE"* SCANS EXIST?

WE FOUND NONE (YET)

we still believe they are there...

A VIEW FROM A /8 DARKNET

port UDP 5060

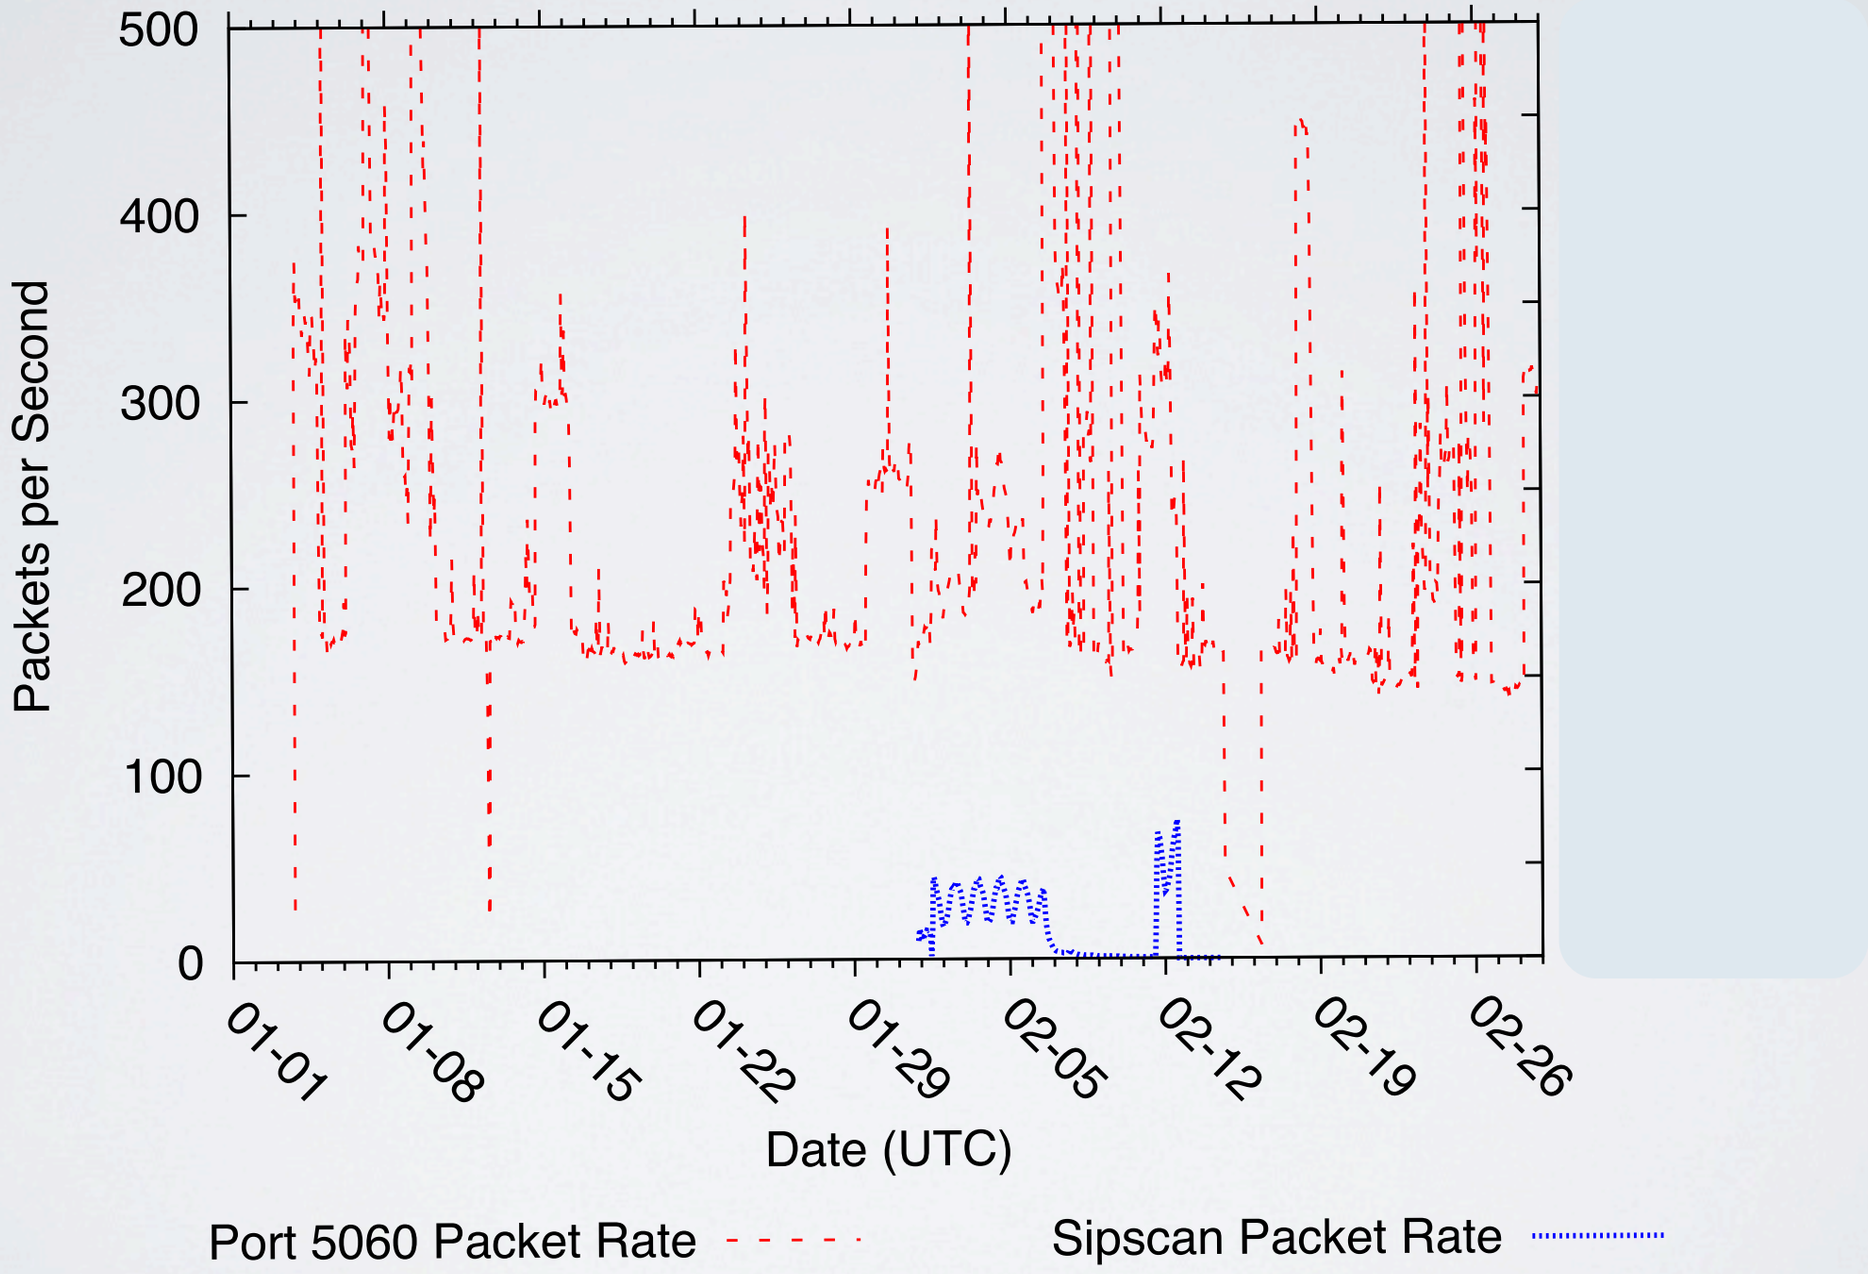

A VIEW FROM A /8 DARKNET

port UDP 5060

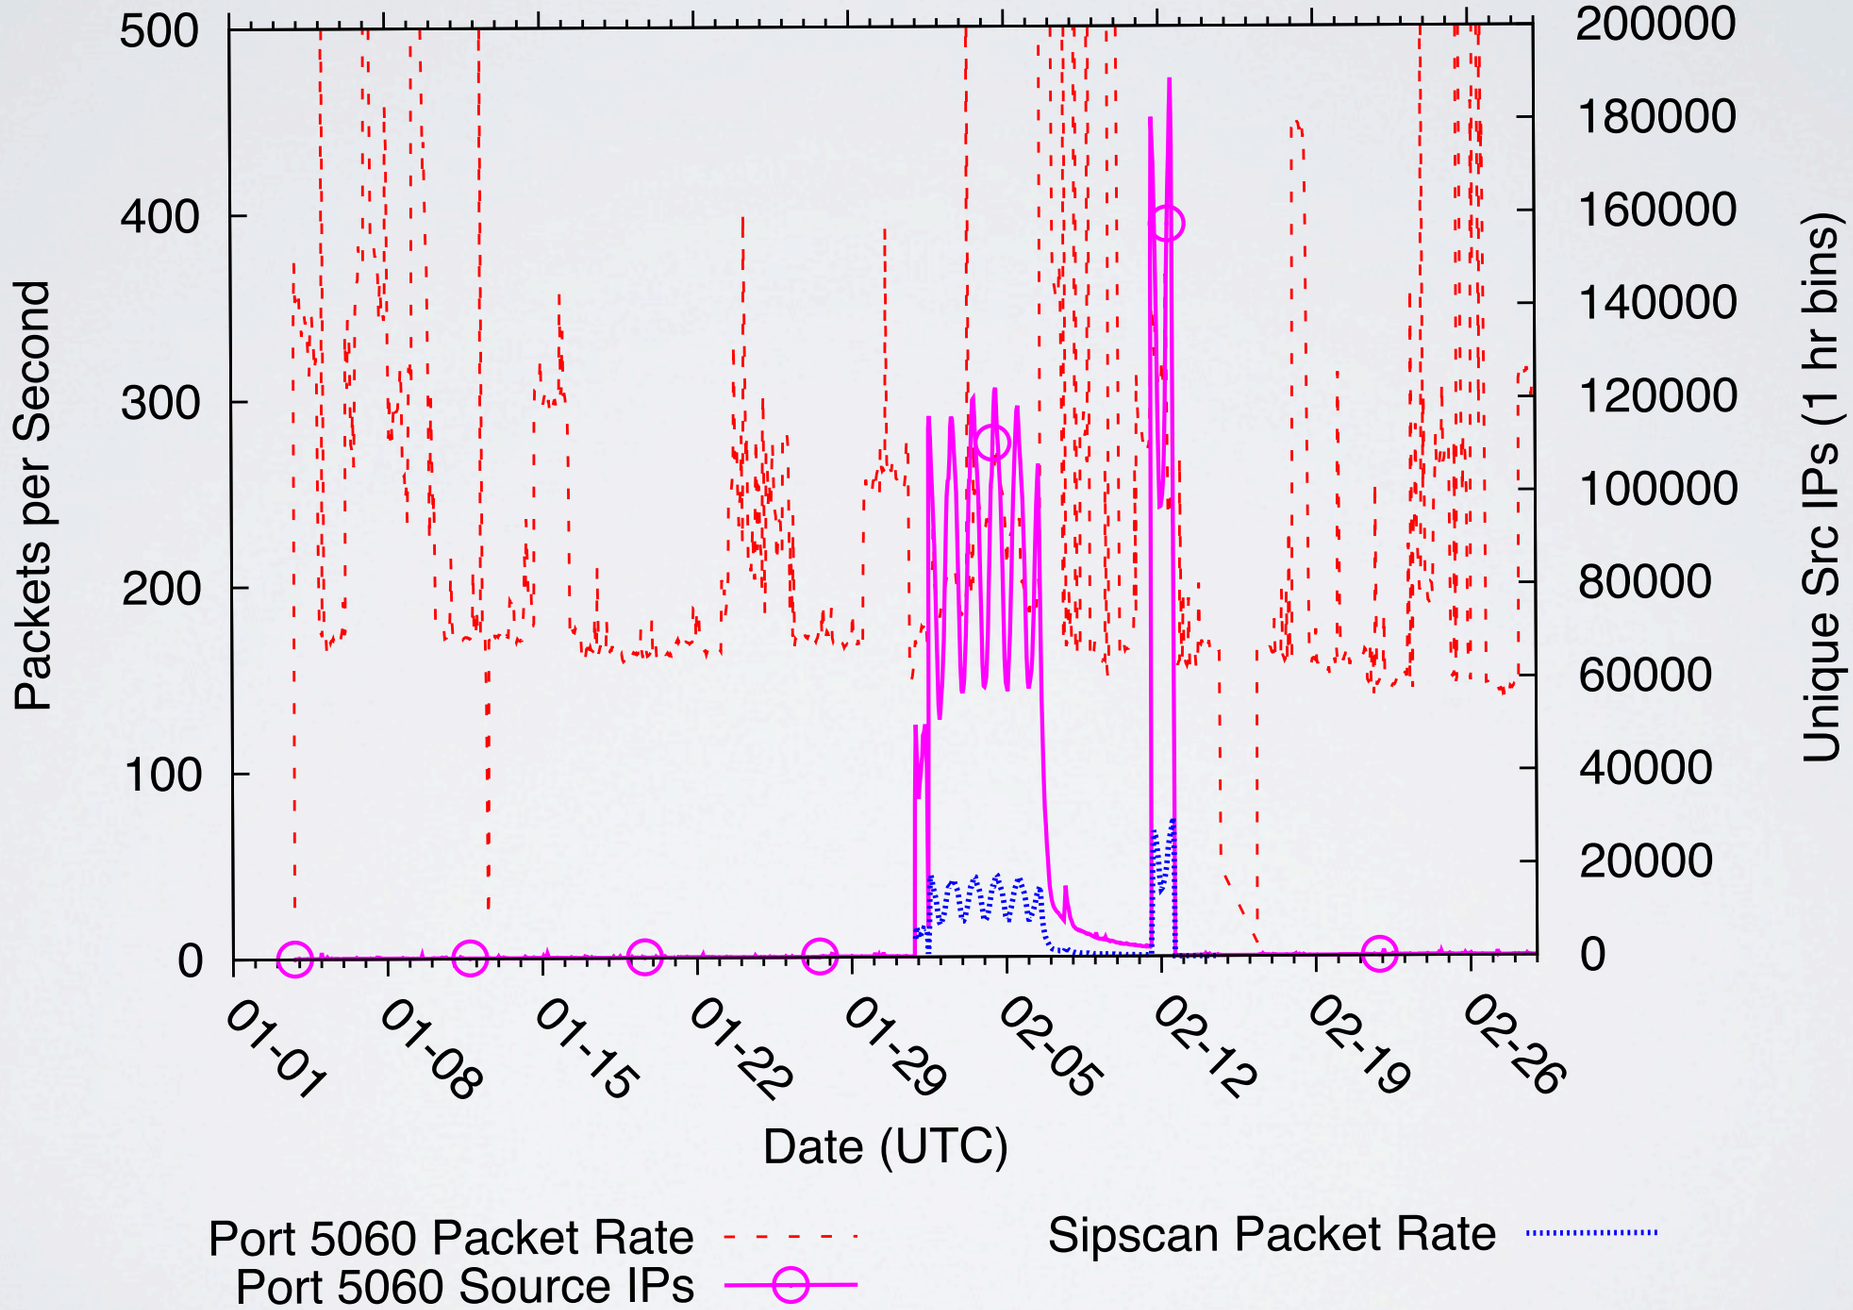

WOULD IT WORK ON A /24 ?

UNLIKELY

- *1 Source IP ~ every hour*
- *IBR not uniformly distributed among /24s*

EVEN ON A /8..

..THERE ARE ISSUES

- *Popular ports*
  - *TCP 80: ~25k distinct source IPs per hour*
  - *TCP 445: ~96k distinct source IPs per hour*
- *Blacklisting*

DATA SHARING

a possible strategy

- A distributed metric
  - Based on observation
    - several Sipscan source IPs hit our /8 only once
    - however, every bot was probing at least 15 other /8 networks
    - recurring bots were approx. probing other 255 networks before hitting our /8 again
  - **Different networks** would observe an unusual amount of **common source IPs** in short time intervals

- **Other ideas?**

DATA SHARING

ISSUES and STRATEGIES

- Large amount of data
- Privacy
  - SEPIA? *http://sepia.ee.ethz.ch*
- One-way unsolicited traffic in live networks
  - More data
  - Immune to blacklisting
  - Useful in identifying hosted bots and botnet
  - Useful in observing reaction of victims

THANKS

alberto@caida.org

Cooperative Association for Internet Data Analysis
University of California San Diego