


Collecting and Visualizing Outages Over the Long Haul

John Heidemann

joint work with Yuri Pradkin, Lin Quan, Abdulla Alwabel
University of Southern California / Information Sciences Institute
CAIDA / AIMS Workshop / San Diego, 2016-03-01

Copyright © 2017 by John Heidemann
Release terms: CC-BY-NC 4.0 international



USC Viterbi School of Engineering

Our Objective: Edge Network Outages, 24x7

- target:
 - **outages in edge networks**, at end users
 - IPv4 /24 address blocks (like 192.0.2.*)
 - for public, visible Internet (not private networks, not behind firewalls)
 - for the whole world
 - **24x7 for years—for the long haul**
- non-targets:
 - core networks and routing
 - individual addresses

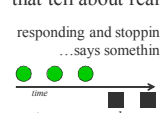
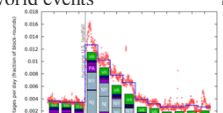
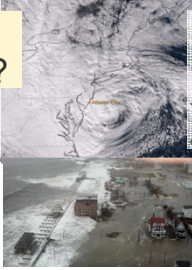
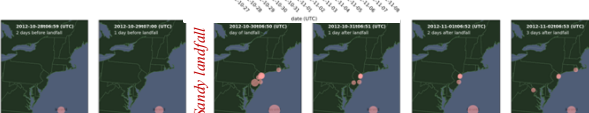
USC Viterbi School of Engineering

Long Haul Outages/ 2016-03-02 4

Can the Internet tell about the Real World?

yes...with interpretation
pinging the Internet tells network outages that tell about real world events

responding and stopping ...says something







Quan et al., "Trinocular: Understanding Internet Reliability Through Adaptive Probing", ACM SIGCOMM, Aug. 2013 doi.acm.org/10.1145/2486001.248017

USC Viterbi School of Engineering

Our Method: Active Probing with Trinocular

- active probing
 - ICMP echo request (pings)
 - sender controls result precision
- adaptive to the target
 - probe *just enough* probes
 - minimize traffic to target (reduce complaints!)
- result: *Trinocular* outage detection
 - running 24x7 since Nov. 2013



USC Viterbi School of Engineering

Long Haul Outages/ 2016-03-02 5

Our Objective: Edge Network Outages

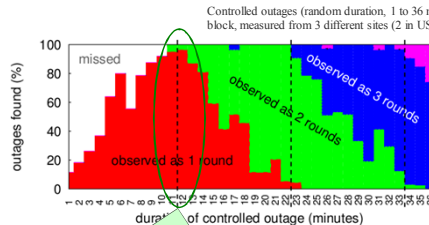
- target:
 - **outages in edge networks**, at end users
 - IPv4 /24 address blocks (like 192.0.2.*)
 - for public, visible Internet (not private networks, not behind firewalls)
 - for the whole world
- non-targets:
 - core networks and routing
 - individual addresses

USC Viterbi School of Engineering

Long Haul Outages/ 2016-03-02 3

Good things About Active: precise guarantees

Controlled outages (random duration, 1 to 36 minutes) in test block, measured from 3 different sites (2 in US, 1 in Japan).



We detect **all** outages longer than 11 minutes (the probing interval)

USC Viterbi School of Engineering

Long Haul Outages/ 2016-03-02 6

Limitations of Active: “only” the public Internet

- we see *the public Internet*
- we cannot see:
 - private networks and behind NAT boxes (but we see the NAT itself)
 - behind firewalls (if they drop pings)
- but... the public Internet is pretty big
 - about 4M /24 blocks

Bad Things About Active: they see you


- because you probe them, they see you
- 24x7
- responses
 - “you’re making my Internet slow”
 - stop DDoS’ing me
 - or “stop [curseword] me, you [curseword]”
- => **implication for the long haul**
 - people *will* drop your traffic
 - our response: look for “gone dark” networks
 - (not real outages!)
 - and an opt-out list with canned responses

Limitations of Active: “only” the public Internet

- we see *the public Internet*
- we cannot see:
 - private networks and behind NAT boxes (but we see the NAT itself)
 - behind firewalls (if they drop pings)
- but... the public Internet is pretty big
 - about 4M /24 blocks
- => **implication for the long haul**
 - it will change (blocks go in and out of use)
 - our response: quarterly refresh our hitlist; seed from ISI Internet censuses (that probe *everything*)

Unmeasurable Blocks Over Time

- we now detect and filter “gone dark” blocks
- **unmeasurable**: never responding for 7+ days



[Alwabel15a, figure 2d]: unmeasurable blocks from A17 (2014q3)

Bad Things About Active: they see you

- because you probe them, they see you
- 24x7
- responses
 - “you’re making my Internet slow”
 - stop DDoS’ing me
 - or “stop [curseword] me, you [curseword]”


Long-Haul Challenge: Visualizing All That Data

- use linear ordering algorithm to identify patterns
 - Quan, Heidemann, and Pradkin. Visualizing Sparse Internet Events: Network Outages and Route Changes. *Computing*, V. 96 (N. 1), pp. 39-51, January, 2014. <<http://dx.doi.org/10.1007/s00607-013-0283-7>>
 - with in-browser viewer

<https://ant.isi.edu/outage/browse/>


[Quan14d, figure 3: The Tohoku earthquake. Dataset: internet survey ii39c]

- problem: $O(n^2)$ algorithm; cannot run on 4M blocks
 - work-around: downsample data 1:200



Efficient Visualization

- **new linear ordering algorithm** $O(n \log n \log m)$
 - where n is the number of blocks and m the duration
- approach:
 - map clustering to sorting: $O(n \log n)$ in time
 - sort on *multi-timescale bitmap*: $O(\log m)$ in space
- goal: put “similar” blocks near each other



Conclusions

- the **long haul is great**
 - 2 years of data
 - analysis is starting... a long way to go, but *we're learning a lot*
- the **long haul is not easy**
 - the target is changing
 - a lot of data to visualize
- want to know more?
 - data: <https://ant.isi.edu/datasets/outage> or <https://impactcybertrust.edu>
 - visualization: <https://ant.isi.edu/outage/browse/>
 - ISI-TR-701 “Evaluating Externally Visible Outages” [Alwabel15a]; more in progress

Multi-Timescale Mapping

- input: outage timeseries from 5 /24 blocks
 - b1 1111 1110 1111 1111
 - b2 1111 1111 1111 1110
 - b3 1111 1100 1111 1111
 - b4 1111 1100 0111 1111
 - b5 1111 1110 1111 1111
- example multi-timescale mapping for b4
 - b4 1111 1100 0111 1111
 - 1 1 0 0 1 1 1
 - 1 0 1 1
 - 1 1
 - 1 ⇒ 1 - 11 - 1011 - 1110 0111 - 1111 1100 0111 1111
- apply to all blocks and sort to get linear ordering
 - b2 1 - 11 - 1111 - 1111 1110 - 1111 1111 1111 1110
 - b1 1 - 11 - 1111 - 1110 1111 - 1111 1110 1111 1111
 - b5 1 - 11 - 1111 - 1110 1111 - 1111 1110 1111 1111
 - b3 1 - 11 - 1011 - 1110 1111 - 1111 1100 1111 1111
 - b4 1 - 11 - 1011 - 1110 0111 - 1111 1100 0111 1111

Visualization Result and Demo

- new output is in our web-based browser
- browser functional, but improvements planned
 - better interactivity
 - drill-down to find IP addresses, AS, geolocation

Display of *it17all*, US blocks, including the 2014-08 **Time Warner Outage**.

