Long-term Data Collection and Analysis of Outages at the Edge

John Heidemann
joint work with Lin Quan and Yuri Pradkin

8 February 2013
CAIDA Workshop on Active Internet Measurements
San Diego, California, USA

work supported by DHS S&T, Cyber Security Division

1. How Pings Measure Internet Outages?

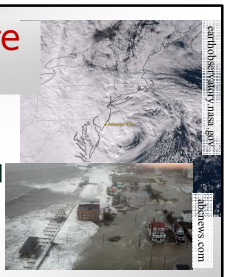

PING 8.8.8.8 (8.8.8.8) 56(84) bytes of data.
64 bytes from 8.8.8.8: icmp_req=1 ttl=251 time=89.6 ms
64 bytes from 8.8.8.8: icmp_req=2 ttl=251 time=83.6 ms
64 bytes from 8.8.8.8: icmp_req=3 ttl=251 time=86.6
^C
--- 8.8.8.8 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2001ms
rtt min/avg/max/mdev = 83.602/86.627/89.641/2.465 ms

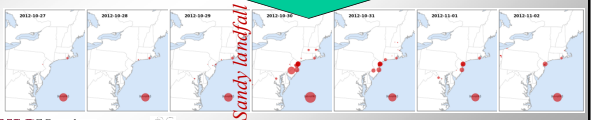

Broader Goal for Outages: Understanding Edge Networks

- quickly know the impact of **natural disasters**
  - Hurricane Sandy, Tōhoku Earthquake 2011, etc.
  - and human ones :-(like Egypt 2011, etc.
- learn about **outage shapes**
  - *wide outages:* many people
  - *long outages:* long time
  - and both
- in ***edge networks*** *(/24* address *blocks,* like 1.2.3.*)
  - most outages are small, *inside* ISPs, *not* from routing
    - e.g.: [Bush et al, IMC 2007]; us: ~70% smaller than routable prefixes
  - want to characterize what people see at home

2. What Role for Long-Term Data Collection?

- real data and new ideas
- collection for serendipity
- "instant" longitudinal study

*=> what properties make **successful** long-term data collection?*

Background: Active Probing with Pings

pings (ICMP echo request)
draw **positive replies** when an IP address is in use

PING 8.8.8.8 (8.8.8.8) 56(84) bytes of data.
64 bytes from 8.8.8.8: icmp_req=1 ttl=251 time=89.6 ms
64 bytes from 8.8.8.8: icmp_req=2 ttl=251 time=83.6 ms
64 bytes from 8.8.8.8: icmp_req=3 ttl=251 time=86.6 ms
^C
--- 8.8.8.8 ping statistics ---
6 packets transmitted, 3 received, 50% packet loss, time 6001ms
rtt min/avg/max/mdev = 83.602/86.627/89.641/2.465 ms

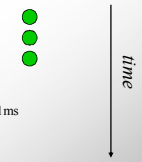

Background: Active Probing with Pings

pings (ICMP echo request)
draw **positive replies** when an IP address is in use

PING 8.8.8.8 (8.8.8.8) 56(84) bytes of data.
64 bytes from 8.8.8.8: icmp_req=1 ttl=251 time=89.6 ms
64 bytes from 8.8.8.8: icmp_req=2 ttl=251 time=83.6 ms
64 bytes from 8.8.8.8: icmp_req=3 ttl=251 time=86.6 ms

or get **negative (non-)replies**
no reply from 8.8.8.8: icmp_req=4
no reply from 8.8.8.8: icmp_req=5
no reply from 8.8.8.8: icmp_req=6
^C
--- 8.8.8.8 ping statistics ---
6 packets transmitted, 3 received, 50% packet loss, time 6001ms
rtt min/avg/max/mdev = 83.602/86.627/89.641/2.465 ms

Pings Tell You Something But Not Everything

positive: block is up

negative: block is down
or
computer crashed
laptop suspended
computer address reassigned
probe or reply lost
firewall enabled

negative replies are ambiguious

So We Probe *Multiple* Addresses

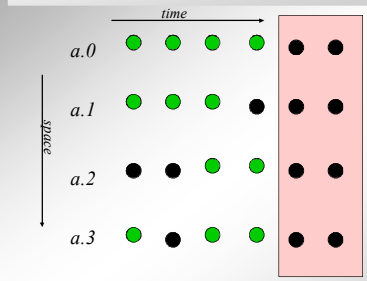

all negative together disambiguates:

*network is **really** down*

Approach: Detect Changes in Ping Response

1. probe multiple addresses in each block frequently

green: positive
black: no response
blue: not probed;
each band is a /24 block

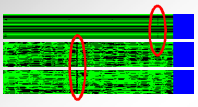

Approach: Detect Changes in Ping Response

1. probe multiple addresses in each block frequently

green: positive
black: no response
blue: not probed;
each band is a /24 block

2. gaps indicate block-level outages

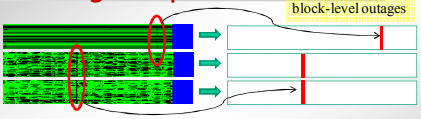

Approach: Detect Changes in Ping Response

1. probe multiple addresses in each block frequently

green: positive
black: no response
blue: not probed;
each band is a /24 block

2. gaps indicate block-level outages

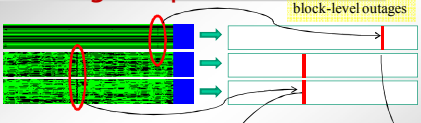

3. show block one per line; reorder to cluster by similarity

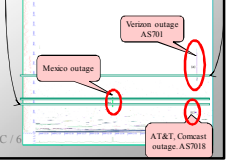

Approach: Detect Changes in Ping Response

1. probe multiple addresses in each block frequently

green: positive
black: no response
blue: not probed;
each band is a /24 block

2. gaps indicate block-level outages

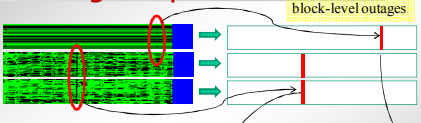

3. show block one per line; reorder to cluster by similarity

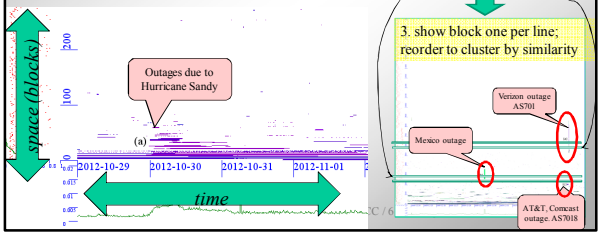

Long-term Data: Internet Surveys

- Internet Surveys
  - sample: 41k blocks (~2% of active address space)
    - half the same from survey to survey
    - half vary, with one-quarter chosen new each time
  - probe for 2 weeks
  - all addresses in each block every 11 minutes
- long-term effort
  - 31 sets to date (it12 to it51), 21 from ≥2 sites, 9 from 3 sites
  - started in 2006, and ongoing through today…
- details and data are available
  - ISI-TR-678b: http://www.isi.edu/~johnh/PAPERS/Quan12a.html
  - data: http://www.isi.edu/ant/traces/

Data About Sandy

long-term data collection happened to cover Sandy

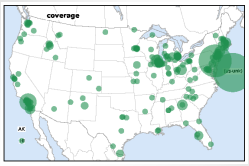

- look at one dataset: internet_address_reprobing_it50j-20121027
- 41,582 /24 blocks
- 11,900 geolocate to US
- 4,117 have enough reponse to analyze
  - 60 of these don't have states

Outages at Sandy Landfall

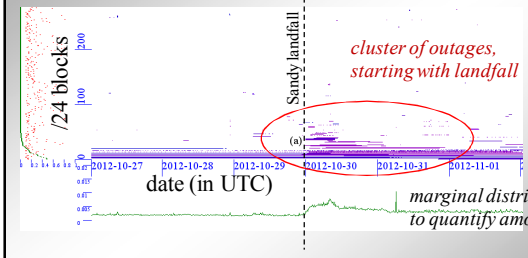

Measuring the Impact

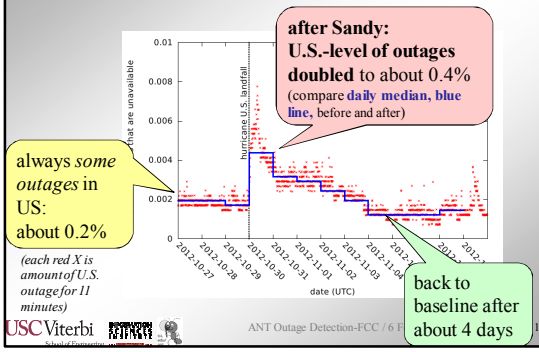

Where Are Outages? NY/NJ

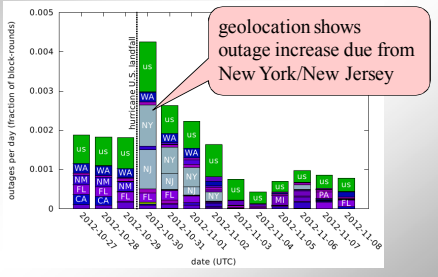

The Northeast, by Day

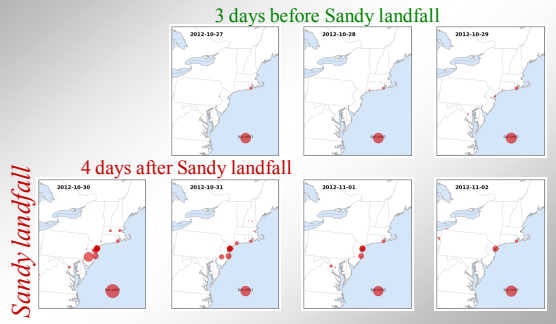

Role of Long-Term Data in Developing Outage Detection

- real data and new ideas
- collection for serendipity
- "instant" longitudinal study

Real Data is Inspirational

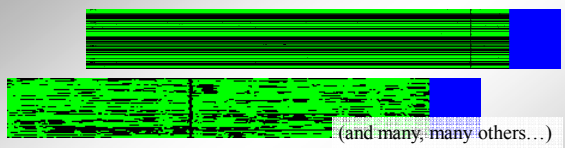

- outage discovery idea came from raw data
  - gee, what's that black vertical line?
  - how can we remove *that* error?
  - hmm…is this something deeper?

Serendipty: Prominent and *Unknown*

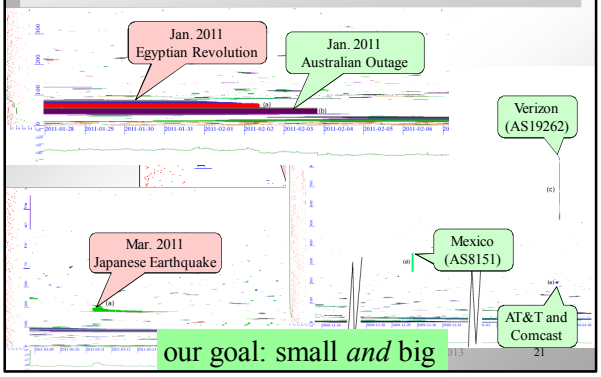

our goal: small *and* big

"Instant" Longitudinal Study

- what's "typical" in the Internet?
- do results vary by site?

⇒ these are "free" if you already have the data

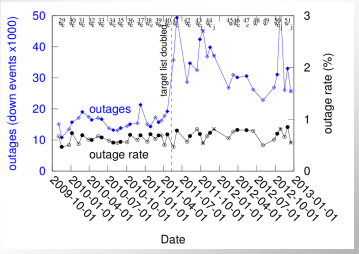

(this data is outages from a single site, so it mixes partial and global outages)

Needs of Longitudinal Data

- regular collection
- consistent, documented methodology
- careful archival
  - checksums (that you check!)
  - backups
- sharing
  - distribution procedures
  - results and what you learn (we have a wiki)

=> a non-trivial amount of work

Cost of Long-term Data Collection

- complaints
- traffic on target
  - survey puts 1 probe / 3 seconds (1400 probes/hour) per block
- optimization can do *much* better
  - we're scaling up outage detection to the whole analyzable Internet: 3.4M blocks
  - optimized probe rate: <20 probes/hour per block
- but specialization incompatible with long-term, *general-purpose* datasets

What Next?

- outage detection from pings works
  - exciting what we see in old data
  - data enabled progress towards Internet-wide detection
- longitudinal data collection important
  - takes care and persistence

- blog: http://ant.isi.edu/blog/
- papers: http://www.isi.edu/ant/pubs
- datasets: http://www.isi.edu/ant/traces
  - can they enable *your* research idea?