Broadening Network Performance: the Edge and Outages

John Heidemann

joint work with Lin Quan and Yuri Pradkin

USC/Information Sciences Institute

20 February 2014

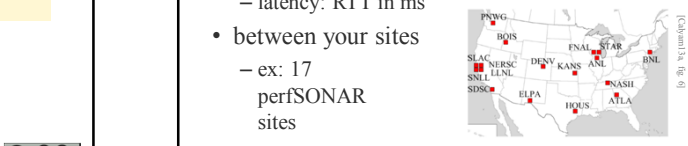

[Calyam'13a, fig. 6]

What Performance Matters?

- traditional network performance:
  - bandwidth: bits/s
  - latency: RTT in ms
- between your sites
  - ex: 17 perfSONAR sites
- find persistent performance problems

What Performance Matters? for **End-Users**?

- ~~traditional~~ **broader** network performance:
  - bandwidth: bits/s
  - latency: RTT in ms
- between your sites
  - ex: 17 perfSONAR sites

- **– reliability**
- **from everywhere**
  - 4 billion IPv4
  - increasing IPv6

What Performance Matters? for **End-Users**? for **Services**?

- ~~traditional~~ **broader** network performance:
  - bandwidth: bits/s
  - latency: RTT in ms
- between your sites
  - ex: 17 perfSONAR sites

- **– reliability**
- **– tail behavior**
- **from everywhere**
- **to everywhere**
  - 4 billion IPv4
  - increasing IPv6

Broadening Network Performance

- network performance can be broader
  - bandwidth and latency… *yes* but also **reliability**
  - in the core…*yes* but also **to the edge**
  - median…*yes* but also **in the tail**

Broadening Network Performance

- network performance can be broader
  - bandwidth and latency… *yes* but also **reliability**
  - in the core…*yes* but also **to the edge**
  - median…*yes* but also **in the tail**
- recent work: Trinocular
  - active probes to all /24 networks
  - to observe network outages

(reliability and the edge: this talk)

(about the tail: not this talk, see recent work from Google)

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

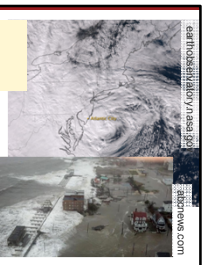

Broader Results

- 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

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

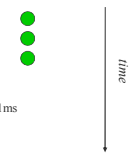

Background: Active Probing with Pings

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

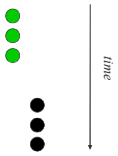

Approach: Active Probing, Adaptive to the Target

- 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

Basic Probing: Multiple Addresses to Disambiguate

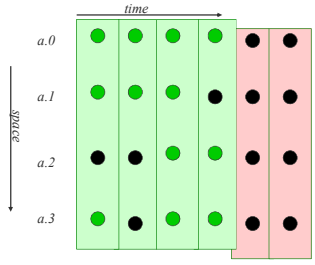

multiple probes address ambiguity

negative: block is down

or
computer crashed
laptop suspended
computer address reassigned
probe or reply lost
firewall enabled

Adaptive Probing: Probe *Just Enough*

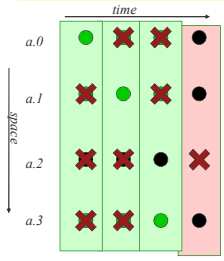

positive responses => block is up
but don't need all 4 to learn

1. instead: probe one by one
2. find **one is up** => **stop early**
3. if **try is down** => **try again**
   => **stop less early**
4. **several fail** => **block down**

adaptive probing uses Bayesian inference informed by model of block response

Based on Model of Blocks and Bayesian Inference

- model each block
  - E(b): addresses that ever respond
  - A(E(b)): Pr[those addresses will respond]
- build model from Internet census
- refine A(E(b)) online
  - concurrent with outage detection
  - exponential weighted moving average
- track *belief* block is up

probe result	prior U'	$P(probe\|U')$	reason
n	U	$1-A(E(b))$	inactive addr.
p	U	$A(E(b))$	active addr.
n	$\bar{U}$	$1-(1-\ell)^{\|b\|}$	non-response to block
p	$\bar{U}$	$(1-\ell)^{\|b\|}$	lone router?

$$B'(\bar{U}) = \frac{P(p|\bar{U})B(\bar{U})}{P(p|\bar{U})B(\bar{U}) + P(p|U)B(U)}$$

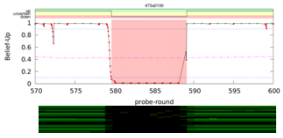

Key Properties of Trinocular

- Trinocular: active probing to detect Internet edge outages
  - **principled**: probe only when needed (informed by Bayesian inference)
  - **precise**: outage duration ±330s (half of probing interval)
  - **parsimonious**: only +0.7% background radiation (at target /24, per Trinocular instance)

(details: "Trinocular: Understanding Internet Reliability Through Adaptive Probing", Quan, Heidemann, Pradkin, SIGCOMM Aug. 2013)

Principled: When to Probe Informed by Bayesian Inference

model: every responding |E(b)|=111, active A(E(b))=0.515
this block is sparse but consistent, so *only a few probes needed*

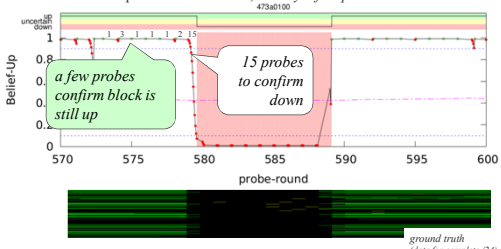

More Intermittent Block

model: every responding |E(b)|=234, active A(E(b))=0.241
this block is dense but intermittent, so *sometimes need more probes*

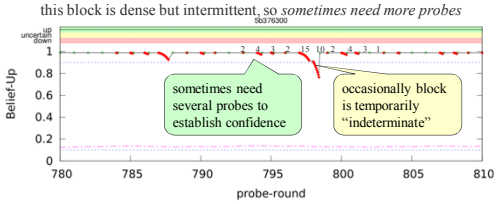

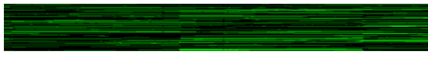

Precise: Detect All Outages?

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

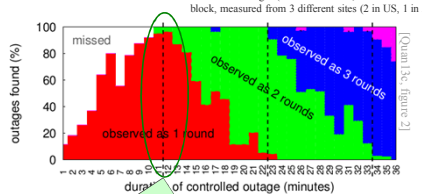

We detect **all** outages longer than 11 minutes (the probing interval)

Parsimonious: Probing Rate

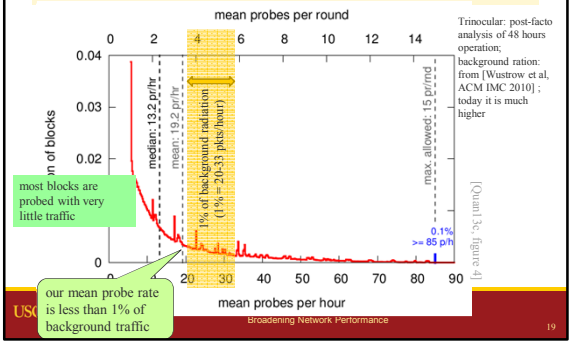

Data About Hurricane Sandy

long-term data collection happened to cover Sandy

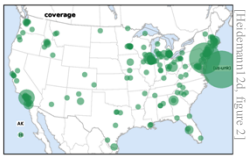

- look at one dataset: internet_address_reprobing_it50j-20121027
- 41,582 /24 blocks
- 11,900 geolocate to US
- 4,117 have enough reponses to analyze
  - 60 of these don't have states

Outages at Sandy Landfall

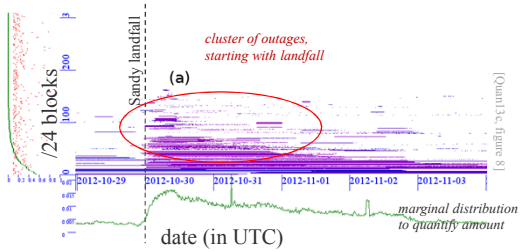

Measuring the Impact

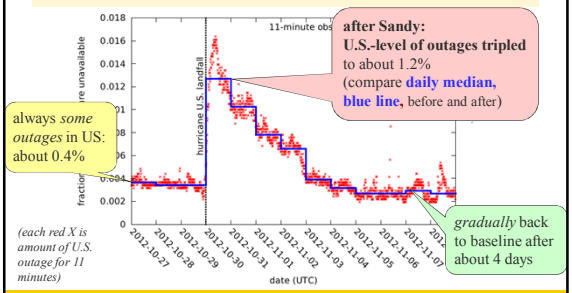

Hurricane Sandy: outages in NY/NJ

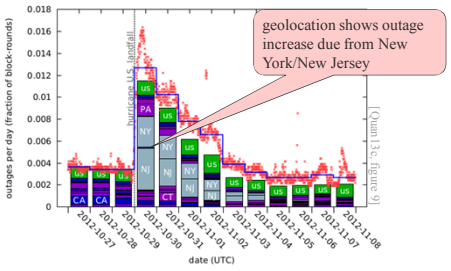

The Northeast, by Day

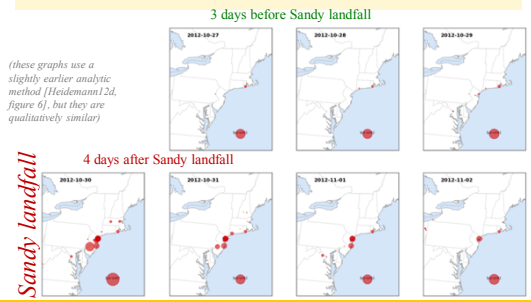

Broader Outages: Prominent and *Unknown*

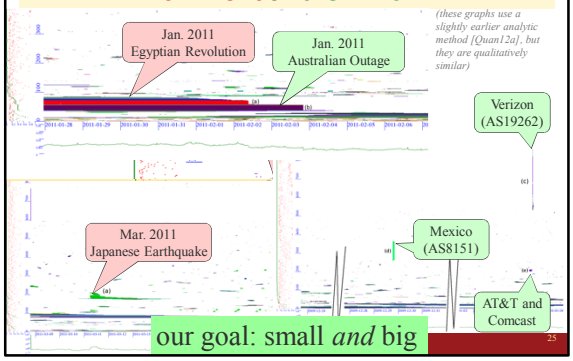

our goal: small *and* big

25

What Next?

- outage detection from pings works
  - Trinocular probing 3.4M /24s (24x7 since Dec. 2013)
- network performance includes **reliability**
- much to learn from **the edge**
- are *reliability* and *the edge* relevant to your nets?

- updates: http://ant.isi.edu/blog/
- papers: http://www.isi.edu/ant/pubs
- datasets: http://www.isi.edu/ant/traces