DHS S&T CYBER SECURITY DIVISION
2014 Cyber Security Division R&D Showcase

Towards Understanding Internet Reliability

John Heidemann
University of Southern California / Information Sciences Institute

December 16, 2014

Homeland Security
Science and Technology

Who We Are

USC Viterbi School of Engineering Information Sciences Institute
Colorado State University

outage detection

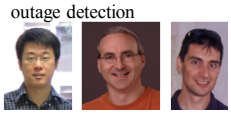

Lin Quan, USC/ISI (now: Bank of China)
John Heidemann, USC/ISI, PI
Yuri Pradkin, USC/ISI

with hosting from
USC/ISI
CSU
Keio University, Japan
Athens U. of Economics and Business

and ongoing collaboration with FCC to evaluate technology

part of the LACREND project:
www.isi.edu/ant/lacrend/

Christos Papadopoulos, CSU co-PI
Xun Fan, USC/ISI
Kaustubh Gadkari, CSU
Zi Hu, USC/ISI
Liang Zhu, USC/ISI

PREDICT www.predict.org

LACREND is part of the DHS PREDICT program

Why Study Internet Outages?

learning about **the Internet**…

…learning about **the World**

Time Warner Cable Says Outages Largely Resolved (NY Times Aug. 27, 2014)

Time Warner Cable said…a problem during routine maintenance caused **a nationwide outage of its Internet service for hours**.
…**There are major outages of at least one telecom provider every year.**…"AT&T had a major outage back in April, Comcast had one last October, Verizon Wireless had several national outages in its 4G

network reliability

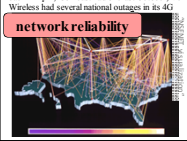

Active Measurement of the Internet

we ping (ICMP echo request) all IPv4

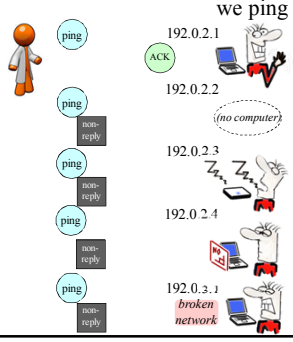

we find computers

and unused space

temporarily unused: (sleeping)

intentionally silent (firewalled)

and *network* ***problems***

Our Goal: Track *All Outages, for All Internet, All The Time*

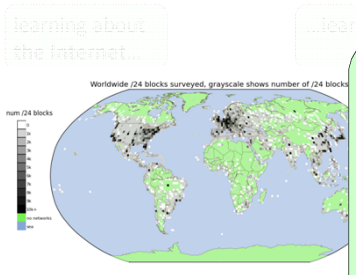

probe the **whole responsive Internet**
4M blocks
each 256 adjacent IPv4 addresses

every 11 minutes

since Dec. 2013

Our Insight: Observing the Internet Informs

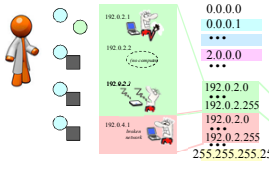

we can probe the **whole Internet**
(the public, unicast, IPv4 Internet)

"only" 4 billion possible addresses
we probe **blocks**: adjacent 256 addrs
(like 192.0.2.0 to 192.0.2.255)
we probe **4 million blocks**

the challenge is: **interpreting the results**
scientifically meaningful results (accuracy!)

and **probing politely**
minimal traffic for sustainable, 24x7 coverage without harming the net (or annoying users!)

Internet Censuses

- since 2003 we've taken *Internet Censuses*
- probe all 4 billion IPv4 addresses over ~2 months
  - pings (ICMP echo request)
  - 0.0.0.0 to 255.255.255.255 (except not private or multicast or reserved)
  - cannot see through firewalls or NAT
- one census is at right:
  - each pixel: average from 65k addrs
  - brightness (red & green): how many respond
  - plotted as Hilbert Curve

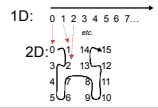

Heidemann et al., "Census and Survey of the Visible Internet", ACM IMC, Oct. 2008 doi.acm.org/10.1145/1452520.1452542

From Census to Outages

can pings detect network outages?
yes…with proper interpretation

responding and stopping …says something

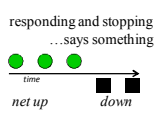

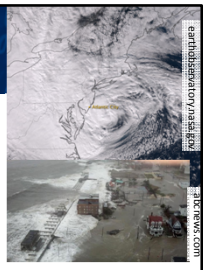

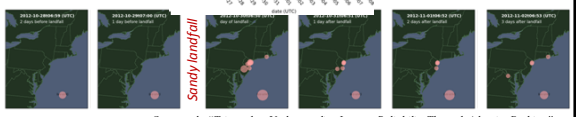

Quan et al., "Trinocular: Understanding Internet Reliability Through Adaptive Probing", ACM SIGCOMM, Aug. 2013 doi.acm.org/10.1145/2486001.248017

Internet Census Over 10 Years

animating IPv4 growth over 10 years…

2011-02-20 (it39w): ICANN fully allocates IPv4 (no more blue!)

2014-08-29 (it61w): utilization rise (more green, less dark)

2006-06-16 (it13w): blue areas are unallocated

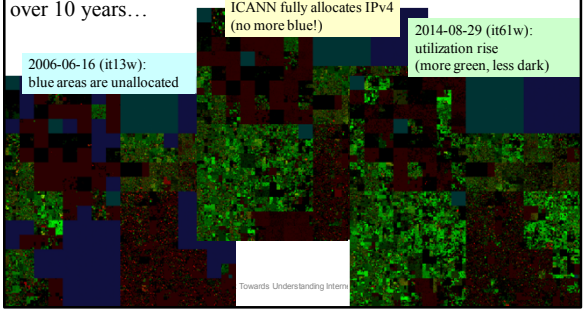

Hurricane Sandy

animating Hurricane Sandy's effects on the Internet

before landfall: few outages

3x outages nationwide on day of landfall

4 days to recover

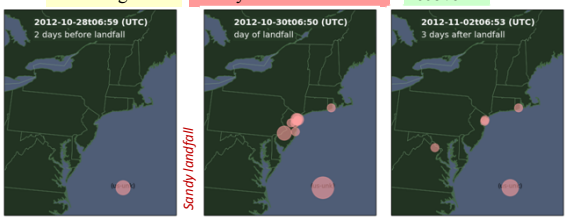

Showcase: Towards Understanding Internet Reliability 10

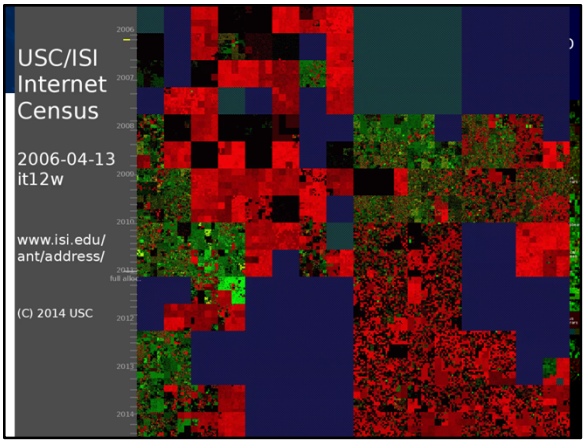

Hurricane Sandy

animating Hurricane Sandy's landfall

- before landfall: typical
- **3x outages nationwide on landfall**
- four days to recover

Our Insight: Observing the Internet Informs

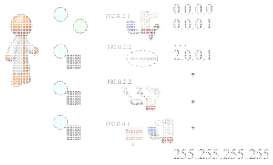

we can probe the **whole Internet**

active probing
=> known precision

the challenge is: **interpreting the results**

observe blocks
=> disambiguate

and **probing politely**

send just enough
=>minimize cost

Case Study: Hurricane Sandy

data pre-dates our 24x7 outage detection
…from our basic science

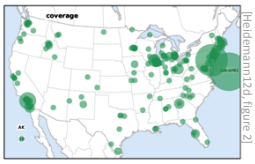

- look at one dataset: internet_address_reprobing_it50j-20121027
- sample of 41,582 blocks
- 11,900 geolocate to US
- 4,117 have enough responses to analyze
  - 60 in US but unknown state: plotted in Atlantic

Observing Blocks to Disambiguate Replies

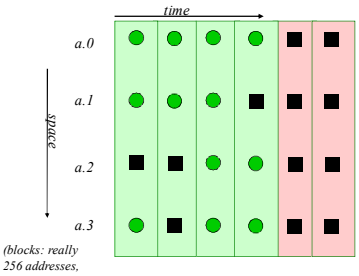

(blocks: really 256 addresses, we show 4 here)

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

multiple probes address ambiguity

***all** negative:* *block is down* !!!

Hurricane Sandy: overall outage rate

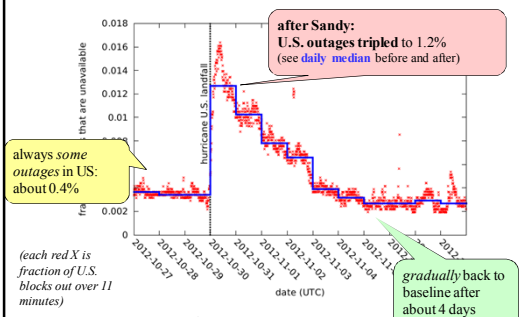

after Sandy:
U.S. outages tripled to 1.2%
(see daily median before and after)

always *some outages* in US: about 0.4%

(each red X is fraction of U.S. blocks out over 11 minutes)

gradually back to baseline after about 4 days

Probing Politely: *Just Enough*

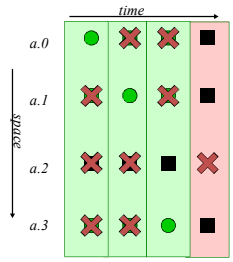

polite: minimal traffic to your net

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

probing politely => observing without harm

Hurricane Sandy: Outages in NY/NJ

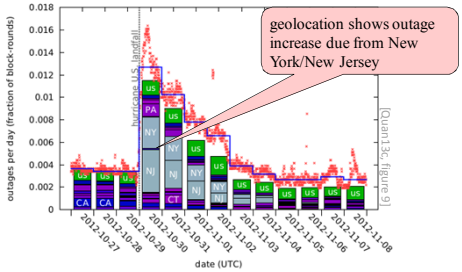

geolocation shows outage increase due from New York/New Jersey

The Northeast, by Day

2 days before Sandy landfall

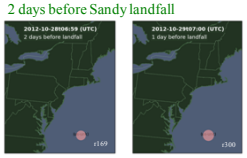

Sandy landfall

observing the Internet informs our understanding of our world

4 days after Sandy landfall

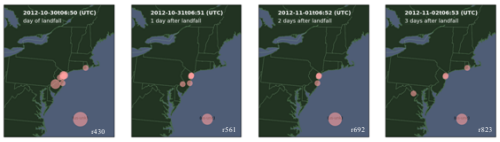

One Week of Global Outages

(taken from 9 months of data)

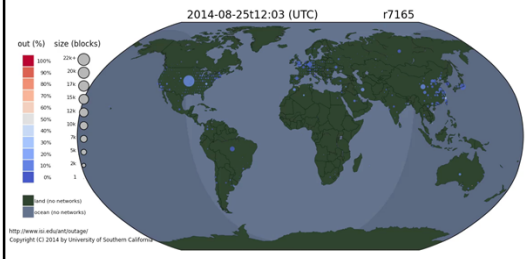

Towards Global Outage Observation

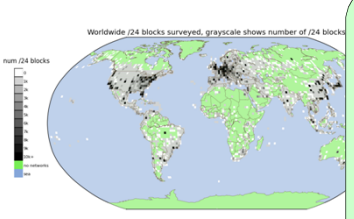

probe the **whole responsive Internet**
4M blocks
each 256 adjacent IPv4 addresses

every 11 minutes

since Dec. 2013

Ongoing Work

- science
  - peer-reviewed publications—7 since 2008
  - data to researchers—as of Oct. 2014: 659 datasets (25TB) to 52 researchers
- operations
  - running 24x7 since Dec. 2013
- technology transfer
  - working with FCC to evaluate methods
    - Rasoul Safavian and John Healy, Bureau of Public Safety and Homeland Security
  - possible application do disaster assessment and improving reliability in general
- future opportunities?
  - can our work help assess other network threats? cybersecurity?

ACM SIGCOMM & Internet Measurements Conferences

PREDICT

One Week of Global Outages

animating outages for the whole Internet

- this dataset:
  - 4M blocks
  - all of 2014q3
- events:
  - Time Warner outage on 2014-08-27 starting 9:20Z
  - ~11 million customers

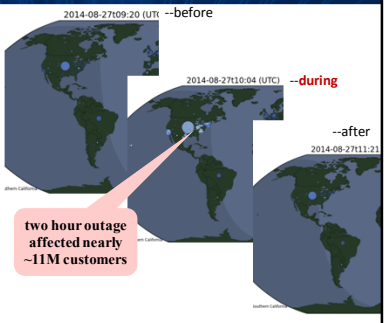

Conclusions

- we can detect Internet outages
  - evidence from Hurricane Sandy and Time Warner
  - in operation
  - technology transition to FCC underway
- DHS support from basic science to technology transfer
- towards a more reliable Internet
- your take?
  - papers: http://www.isi.edu/ant/pubs/
  - our data is available at no cost: https://www.predict.org (more info at http://www.isi.edu/ant/traces/)
  - does our approach apply to **your** mission?

PREDICT www.predict.org