


Reliable routing

Brighten Godfrey
CS 538 October 4 2012





The Internet is messy in practice

- Transient loops
- Persistent loops
- Asymmetry
- Instability

How to look inside a black box

Looking inside a black box



End-to-end measurement from vantage points combined with careful statistics

A standard for the field

- End-to-End Effects of Internet Path Selection [Savage '99]
- RON [Anderson '01]
- Related area: **network tomography**

Many resources now available

- PlanetLab, Seattle P2P testbed, RouteViews, DIMES, CAIDA, ...

Name	Description
adv	Advanced Network & Services, Armonk, NY
austr	University of Melbourne, Australia
austr2	University of Newcastle, Australia
batman	National Center for Atmospheric Research, Boulder, CO
bnl	Brookhaven National Lab, NY
bsdi	Berkeley Software Design, Colorado Springs, CO
connix	Caravela Software, Middlefield, CT
harv	Harvard University, Cambridge, MA
inria	INRIA, Sophia, France
korea	Pohang Institute of Science and Technology, South Korea
lbl	Lawrence Berkeley Lab, CA
lbli	LBL computer connected via ISDN, CA
mid	MIDnet, Lincoln, NE
mit	Massachusetts Institute of Technology, Cambridge, MA
ncar	National Center for Atmospheric Research, Boulder, CO
near	NEARnet, Cambridge, Massachusetts
nrao	National Radio Astronomy Observatory, Charlottesville, VA
oce	Oce-van der Grinten, Venlo, The Netherlands
panix	Public Access Networks Corporation, New York, NY
pubnix	Pix Technologies Corp., Fairfax, VA
rain	RAINet, Portland, Oregon
sandia	Sandia National Lab, Livermore, CA
sdsc	San Diego Supercomputer Center, CA
sintef1	University of Trondheim, Norway
sintef2	University of Trondheim, Norway
sri	SRI International, Menlo Park, CA
ucl	University College, London, U.K.
ucla	University of California, Los Angeles
ucol	University of Colorado, Boulder
ukc	University of Kent, Canterbury, U.K.
umann	University of Mannheim, Germany
umont	University of Montreal, Canada
unij	University of Nijmegen, The Netherlands
usc	University of Southern California, Los Angeles
ustutt	University of Stuttgart, Germany
wustl	Washington University, St. Louis, MO
xor	XOR Network Engineering, East Boulder, CO

[Paxson's vantage points]

Can we make the Internet reliable?



Faster routing convergence

- Tune timers so BGP/OSPF converge faster
- But still transient disconnection (loops, black holes)

Backup paths in forwarding table

- Don't have to wait for control plane convergence
- E.g., MPLS Fast Reroute, R-BGP
- But can't protect against all failure scenarios

Key question: Can we achieve **perfect reliability** in routing & forwarding?

- (except due to congestion or physical layer corruption)



It is possible to avoid transient routing problems

- BGP, OSPF, RIP, ISIS, ..., all have loops & black holes during convergence
- but that is **not fundamentally necessary**

Don't wait for global convergence

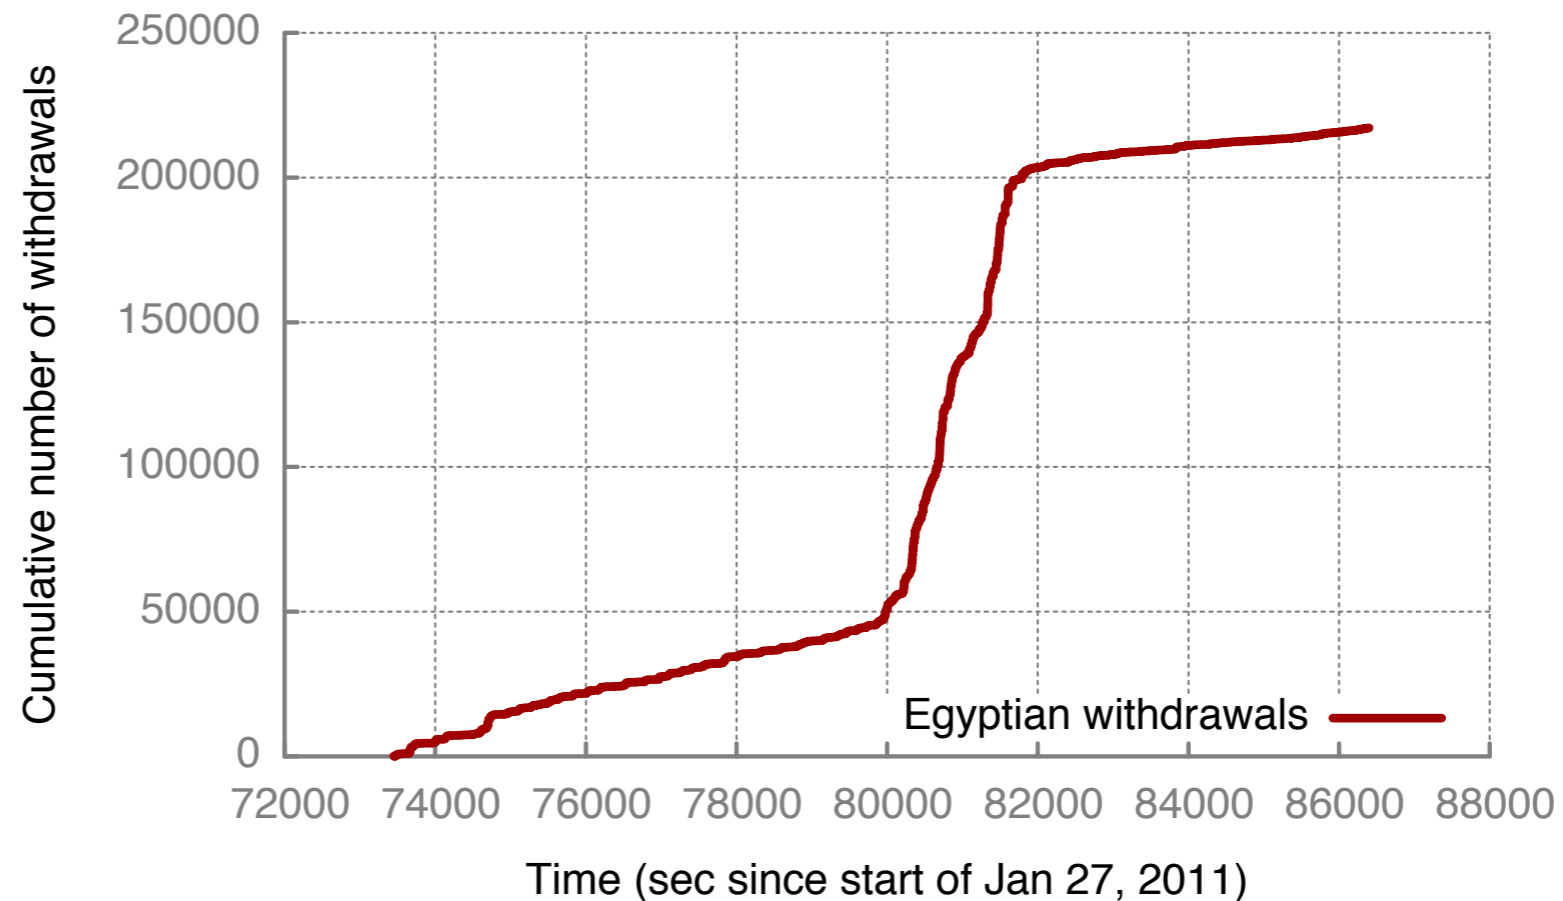
- Separate two functions:
 - long-term topology distribution
 - handling transient changes
- Trick: carry topology updates in packet

Assignment 1



Returned by tomorrow

Part 1 result:



Part 2 discussion:

- Why did adding the proxy increase throughput?



Jon Peterson

- Distinguished Engineer, Neustar Inc.
- Talk title: “Governing an Internet”
- In 2405 again

