EdgeScope: Exposing the View of the Edge of the Network

David Choffnes

EECS, Northwestern U.

http://aqualab.cs.northwestern.edu/projects/EdgeScope.html

Internet wide studies need…

- Internet scale data
  - BGP is useful, but has serious limitations
  - PlanetLab offers more control, but skewed view
  - End system monitors need widespread adoption
- EdgeScope
  - Make our collection of edge traces available
  - What do we have?
  - When can you have it?

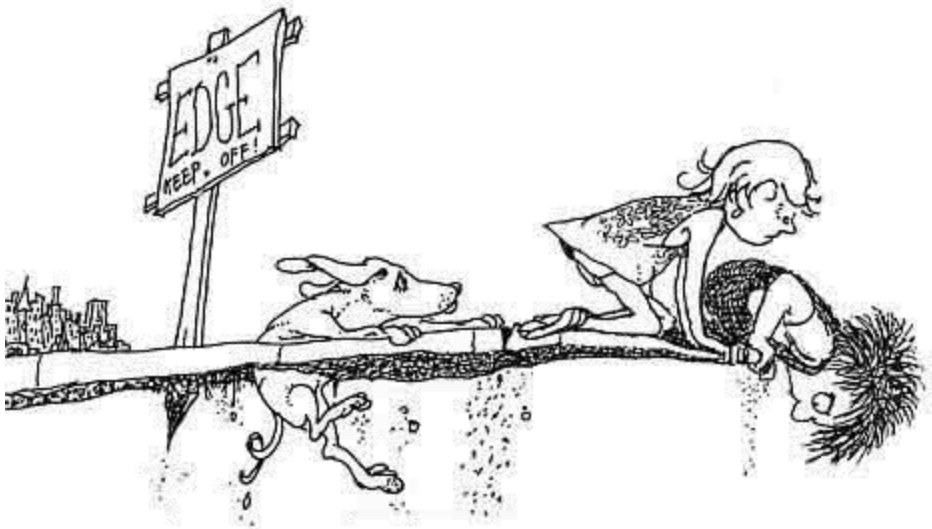

Ono users (IPs)

Where our data comes from

- Ono
  - Uses CDN redirections to inform peer biasing for BitTorrent
  - Installed more than 800,000 times
- NEWS
  - Uses passively gathered BT data to detect, confirm and isolate network events
  - More than 40,000 users
- Coverage
  - ~8k ASNs
  - 54k routable prefixes
  - 200 countries

What we have

- Data types (see website for details)
  - Per-download stats
    - Transfer rates, file-size estimates, state
  - Per-connection stats
    - Transfer rates, cumulative data transferred, seed/leech
  - Global stats
    - Overall transfer rates, session times
  - … other interesting/necessary stuff
    - IP changes
- All of this is sampled every 30 seconds
  - Per-session data sampled every hour and at end
- Traceroutes/pings
    - Uses builtin command, we are playing with v6 traceroutes/pings
    - Limited to a maximum number of measurements in parallel

What this is not

- A platform for controlled experiments
  - Why?
    - Security implications!
- A topology measurement tool
  - But we do have loads of traceroute data
- An arbitrarily extensible data collection system
  - Everything we collect relates to Ono/NEWS performance
  - If it fits, we can add it fairly easily
    - Needs to go through a beta process (usually about a week)
    - Once mainlined, near full adoption within about 4 days

How much is there

- Started (proper) collection in December 2007
- Daily stats (approximate)
  - 3 to 4 GB of compressed data
  - About 10 to 20 GB raw data
  - 2.5-3M traceroutes
  - 100-150M connection samples

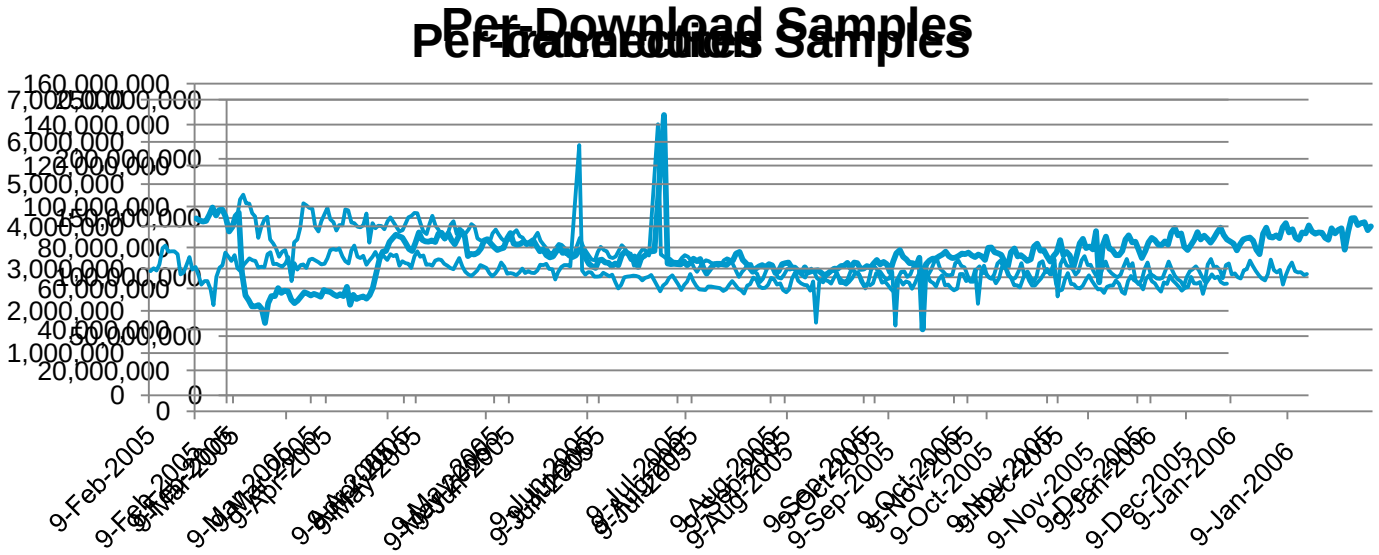

Some cool things we've done with it

- NEWS
- SwarmScreen
- Network positioning (cool? cold? you decide.)
- Topology studies
- Fabian's talk

What can you get?

- Preliminaries
  - CAIDA-style agreement
- Anonymization
  - AS-level detail
  - Prefix-preserving
  - (Maybe) User ids, without location info
- If you need IPs, you have to work with us

When can you get it?

- Ono dataset
  - Now
- AS links (CoNEXT work)
  - Now
- Everything else
  - On-demand anonymization (takes time)
  - Hardware on order
  - Quarantine period (6-12 months)

Some advice

- Before you ask for data
  - Be sure you know what you want
  - Make sure you have space for it
  - Give us time to get it to you
- Working with data at this scale
  - Throwing hardware at the analysis doesn't work
    - Good data structures do work
  - MapReduce isn't always the best fit
    - Especially if you don't have a giant cluster
  - Dynamic languages are a bad idea
    - Seriously, perl is not your friend here
    - Thank me later

Lessons learned

- Privacy is hard
  - No really, this is serious
    - Messing this up will ruin it for everyone
  - We invite new proposals/research in this area
- Scale
  - Cannot do this without good tech support, net ops cooperation, elbow grease
  - MTTFs in mirror(ed array) are closer than they appear
- Other fun stuff
  - Timestamp synchronization
  - UUIDs
  - Users can come and go at any time
    - Sessions and install/uninstall

Link to site

http://aqualab.cs.northwestern.edu/projects/EdgeScope.html