How Dynamic is the ISPs Address Space? Towards Internet-Wide DHCP Churn Estimation

Giovane C. M. Moura⋆, Carlos Gañán, Qasim Lone, Payam Poursaied, Hadi Asghari, and Michel van Eeten

giovane.moura@sidn.nl

Delft University of Technology

⋆SIDN Labs

IRTF & ISOC Workshop on
Research and Applications of Internet Measurements (RAIM)
Saturday, October 31, 2015
Yokohama, Japan

Paper in a nutshell

- **Problem:** bot counts based on IPs counts are flawed
  - Why? DHCP churn
  - Torpig paper showed that DE bots 4x more IPs than US
- **Fix:** need to compensate for DHCP churn
- **OK, so how to measure churn?**
  - It's been done (passively), small scale (not ISP wide)
  - Need to scale-up ; ISP-independent

Internet-wide DHCP churn measurement

Our Method

1. **Probe:** continuous ICMP probes on entire ASes, every 10min
   - Based on the Internet Census paper
2. **DHCP session estimation:** Interpolate consecutively ack'ed packets
   - Missing ack: session expired
   - More complex: see paper
3. **Validation:** mid-size ISP (1 M IP addresses)
   - Radius Logs vs measured DHCP sessions
   - 2 weeks period

Ground truth: what we try to measure

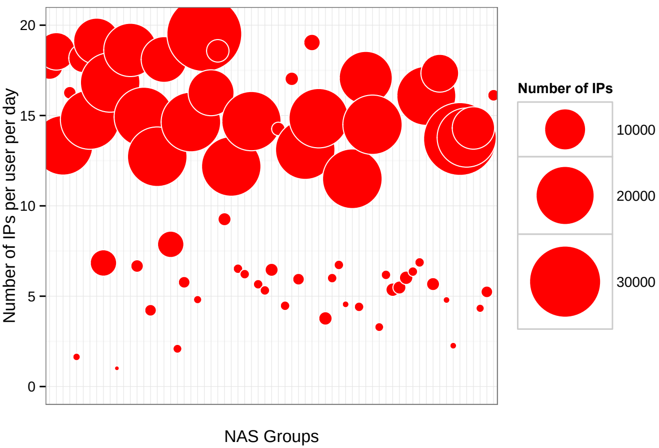

Figure: Pools of addresses (NAS) and average daily User/IP

Validation

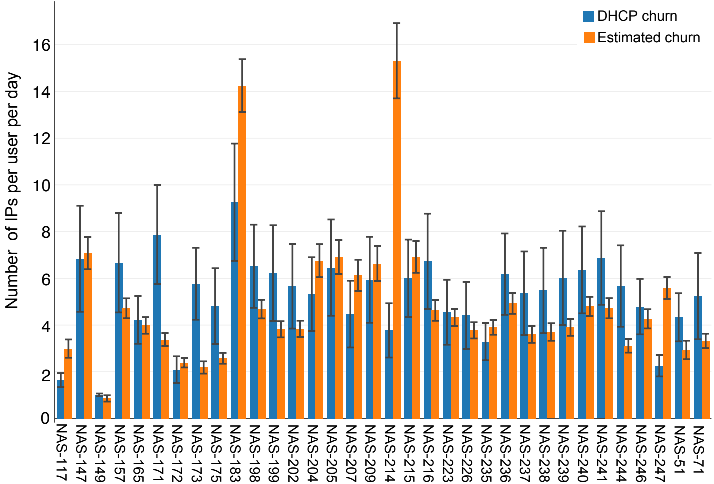

- 72.3% average precision in our model
- Simple method that works on a highly dynamic network

Now, measure other ASes

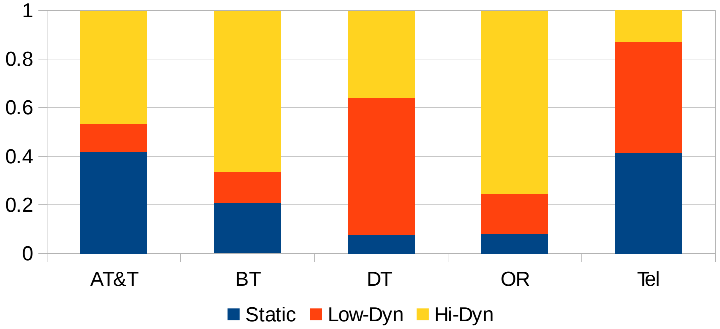

- Employed k-means to 5 ASes of large ISPs
  - Fastrack Elsevier ComCom Paper (under review)
- 2nd validation: RIPE Atlas (works better)
- Next: normalize bot counts