Detecting routing anomalies using RIPE Atlas

Todor Yakimov

Graduate School of Informatics
University of Amsterdam

Wednesday, February 5, 2014

What are routing anomalies?

Introduction

- What are routing anomalies?
  - Incapability of packet delivery to legitimate destinations
  - Delivery of packets to a wrong destination

- Why do they occur?
  - Out of innocent mis-configurations or bugs
  - Government spying or Internet censorship
  - Malicious attackers seeking blackholing, impersonation, interception

What is used to detect such anomalies?

Introduction

- Interior gateway protocol (IGP) environments:
  - All data is under the same administrative control
  - Core tools: ping, traceroute, dig
  - Other tools: Icinga, Nagios
- Exterior Gateway protocol (EGP) environments:
  - Datasets part of different administrative domains
    - Regional Internet Registries (RIR)
    - Remote Route Collectors(RRC), formerly RouteViews
    - RIR Internet numbering assignments datasets
    - Internet Routing Registry (IRR) - RIPE NCC, NTT, Level3, Merit network
  - Tools: Cyclops, PHAS, ARGUS

Research questions

Introduction

Main

"Is it possible to detect filtering, MitM(Man-in-the-Middle) routing attacks, eavesdropping or simply routing policy changes by using RIPE Atlas's historical archives or by using newly-defined active measurements?"

"What other datasets are needed to complement data obtained from RIPE Atlas in the process of accurately detecting the aforementioned Internet routing anomalies?"

RIPE Atlas system specification

- The largest Internet measurement network
  - Public access, everyone can use every probe
  - More than 4800 probes
  - Latest probes are TP-LINK TL-MR3020
  - 1901 IPv4 ASNs covered (4.125%)
  - 139 countries covered (68.137%)
- Centralized reservation, scheduling and storage of measurements
- IPv4/6 Measurement tools
  - ping
  - traceroute
  - dig
  - openssl
  - curl(upcoming)

RIPE Atlas system specification

- GUI for easy usage
- REST API for robust probe selection and measurement specification
- Automatic alerts for ongoing measurement (upcoming)
- Usage limitations - measurements cost credits (hosting probes generate)
  - No more than 175K credits per day
  - Max. 500 probes per measurement
  - No more than 10 ongoing UDMs towards the same target
  - Delay in reserving probes and starting measurements
  - Slight offset in a measurements' interval

Experiment 1 partial results

- Internet censorship
  - DNS blocking
  - Traffic blackholing
- Experiment specification
  - Determine blocking of torrent and news websites - ThePiratebay, TorrentFreak, LiveJournal
  - Approximately 1800 EU probes from unique prefixes used
  - No results with local resolvers considered
  - Traceroute(ICMP echo) to URL
- Experiment detection mechanisms
  - DNS IN A record does not match ip/prefix of website
  - Probe IP, DNS server IP and last-hop IP are from the same ASN

Experiment 1

ThePirateBay.org filtering

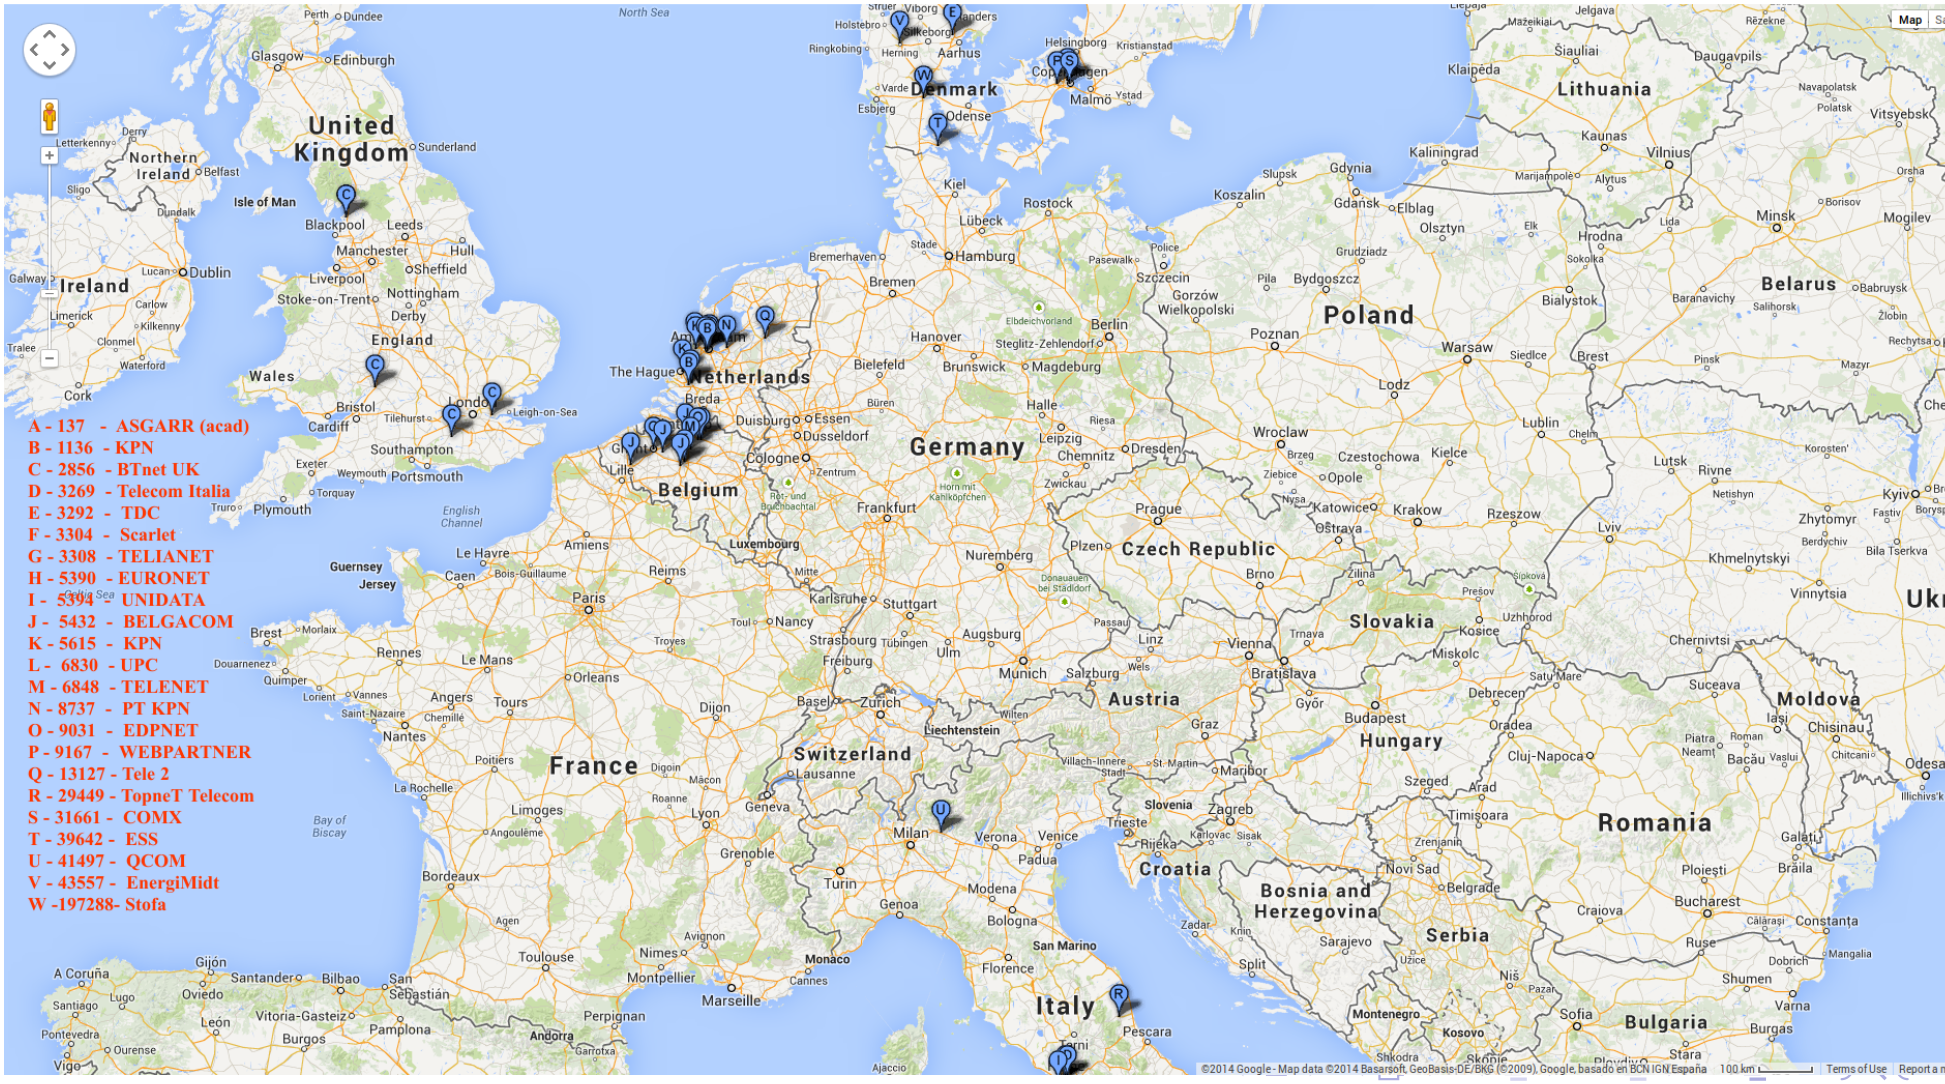

Experiment 2: De-bogonising address space ranges

- De-bogonised IP ranges - previously reserved IPv4 ranges get released and distributed by IANA to RIRs for further assignment
- RIRs first launch debogon projects
  - Control-plane implication analysis: BGP beacons
  - Data-plane implication analysis: background radiation monitoring
- Latest(and last) distributed /8 IPv4 ranges in 2011:
  - APNIC - [36, 39, 42, 49, 101, 103, 106]/8
  - RIPE NCC - 185/8

Experiment 2: De-bogonising address space ranges

- Experiment setup
  - Approximately 3100 world-wide probes used from unique prefixes
  - Ping as measurement
  - Each probes pings both the de-bogonised prefix and another prefix from the same ASN and geo location
- 12 hours scanning of single, currently announced subprefix
  - Subprefixes still advertised by RIRs ASNs with provided pingable targets
  - Pings 20 minute apart from each probe to target prefix
  - Pings 60 minute apart from each probe to reference point
- Successful reachability test:
  - At least one ping reply from host in de-bogonised range
  - At lest one ping reply from reference host

Experiment 2a partial results

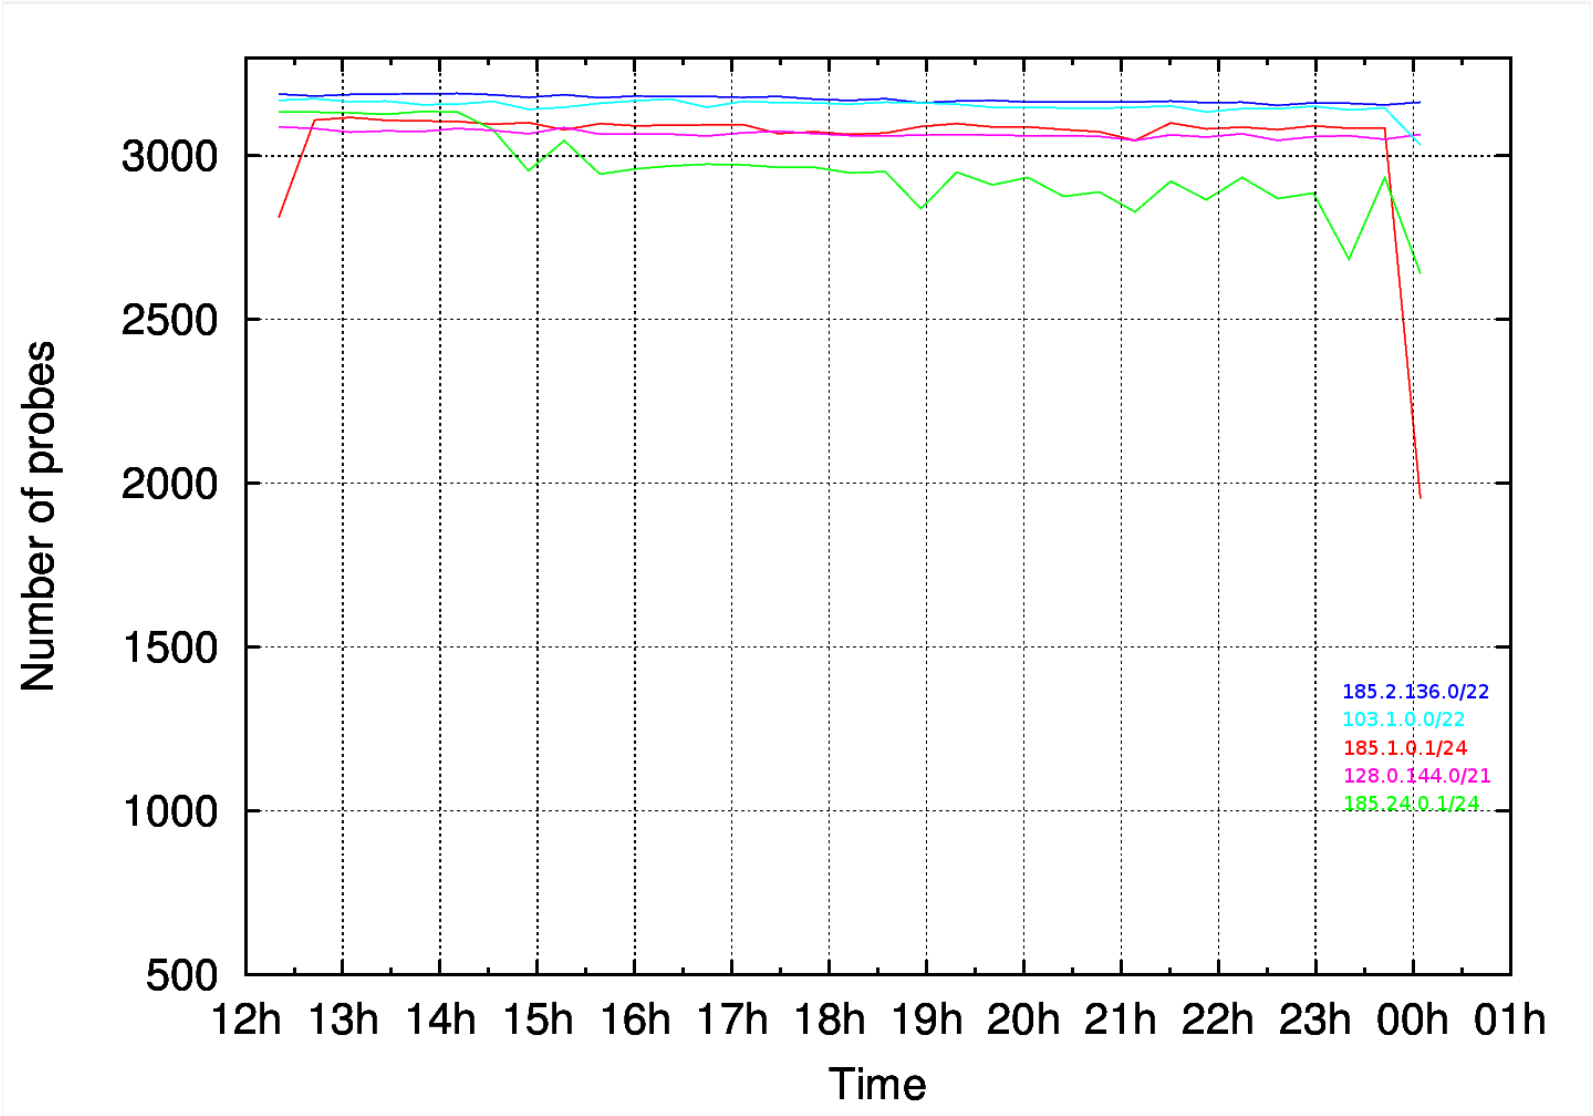

Prefix hijacking detection

- Discarded measurements
  - Each test had 28-115 probes incapable of reaching either host
  - Type 3 replies filtered (if in one set, removed from both)
  - Type 0, 11 considered
- Results - probes incapable of reaching de-bogonised prefix
  - 103.1.0.0/22: Probe 12007 (AS45050 HI-MEDIA France)
  - 128.0.0.0/16: None!
  - 185.1.0.1/24: Probe 12007
  - 185.2.136.0/22: Probes 156 (AS51127 LNET-AS GER), 12007
  - 185.24.0.1/24: Probes 156, 3892 (AS50473 ECO-AS RU), 4532 (ASN2818 BBC UK), 12007

Prefix hijacking

- Falsifying BGP advertisements with the purpose of establishing blackholing, imposture or interception for a given prefix.
  - BGP MOAS or subMOAS conflicts for AS_PATH advertisements with invalid origin
  - No MOAS or subMOAS as invalid transit
  - Keeping a valid route to original prefix destination forms a MitM attack!

Prefix hijacking

- Data-plane detection
  - Monitoring network location
  - Measuring path disagreement with traceroute to target prefix and a reference point
  - Best detection systems use a hybrid approach by correlating control- and data-plane monitoring
  - With data-usage limitations, one really needs to know what to look for

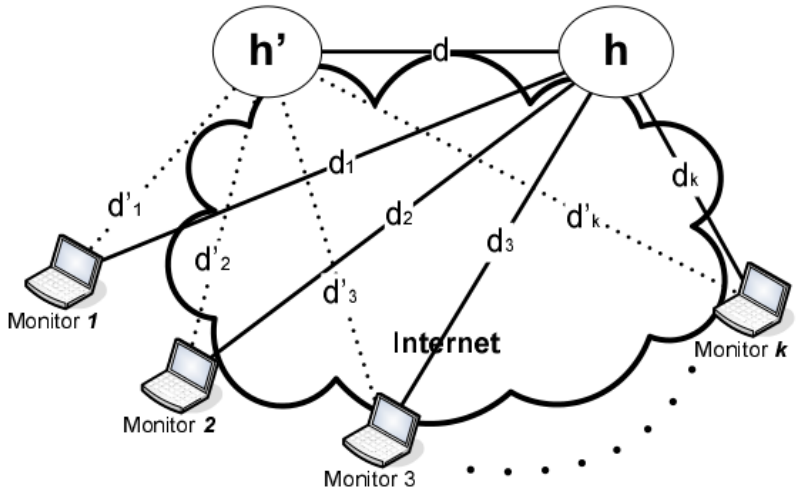

Experimet 3: Prefix hijacking

- Experiment setup
  - Monitored prefix: OS3
  - Traceroute measurement
  - Approximately 1200 world-wide probes from unique ASNs used
    - Unique ASNs
    - ASNn not part of SURFnet's immediate peers
    - Reference point OS3 BSR SURFnet uplink neighbor
    - Hijack simulation: own home probe with more-specific static routes
- Experiment results: data sufficient to detect both network location change and path disagreement

Conclusion

- RIPE Atlas is a robust tool for measuring network anomalies
- Combined with other RIPEStat data, sophisticated vantage point selection is possible
- Large-scale measurement ease of use
- Large-scale measurement scheduling does not suffer too big offsets/delays
- The credit limitations of the system simply makes impossible certain tasks

Questions