Diagnosing the Location of Bogon Filters

Randy Bush
Internet Initiative Japan (IIJ)

James Hiebert
National Oceanic and Atmospheric Administration

Olaf Maennel
University of Adelaide

Matthew Roughan
University of Adelaide

Steve Uhlig
Delft University of Technology

Outline

- Advertising a new prefix
- Methodology
- In-probes
- Out-probes
- Relationship in- and out-probes
- Further work

Problem: "Bogon filters"

- ISPs often filter unallocated address space to protect themselves from malicious attacks and unwanted traffic
- Over time unallocated address space may become allocated and legitimately announced address space…
- Problem: Filters need to be updated but seem often not to be

Objectives

- Develop methodology that is capable of detecting and locating filters that are blocking newly allocated address space
- Analyze reachability status of a newly allocated prefix
- For the experiment, ARIN loaned us
  96.0.0.0/16 97.64.0.0/16
  98.128.0.0/16 99.192.0.0/16

Testing reachability of a new prefix

- Terminology:
  - Test-prefix: newly allocated prefix to be tested
  - Anchor-prefix: well-established prefix whose reachability should be fine
  - Probe-site: router that announces both the test-prefix and the anchor-prefix

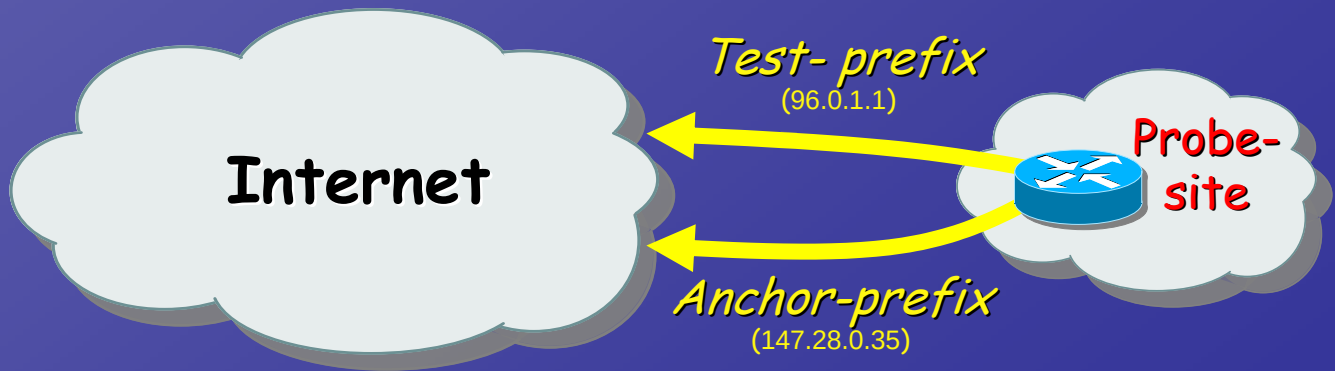

Overview: Approach

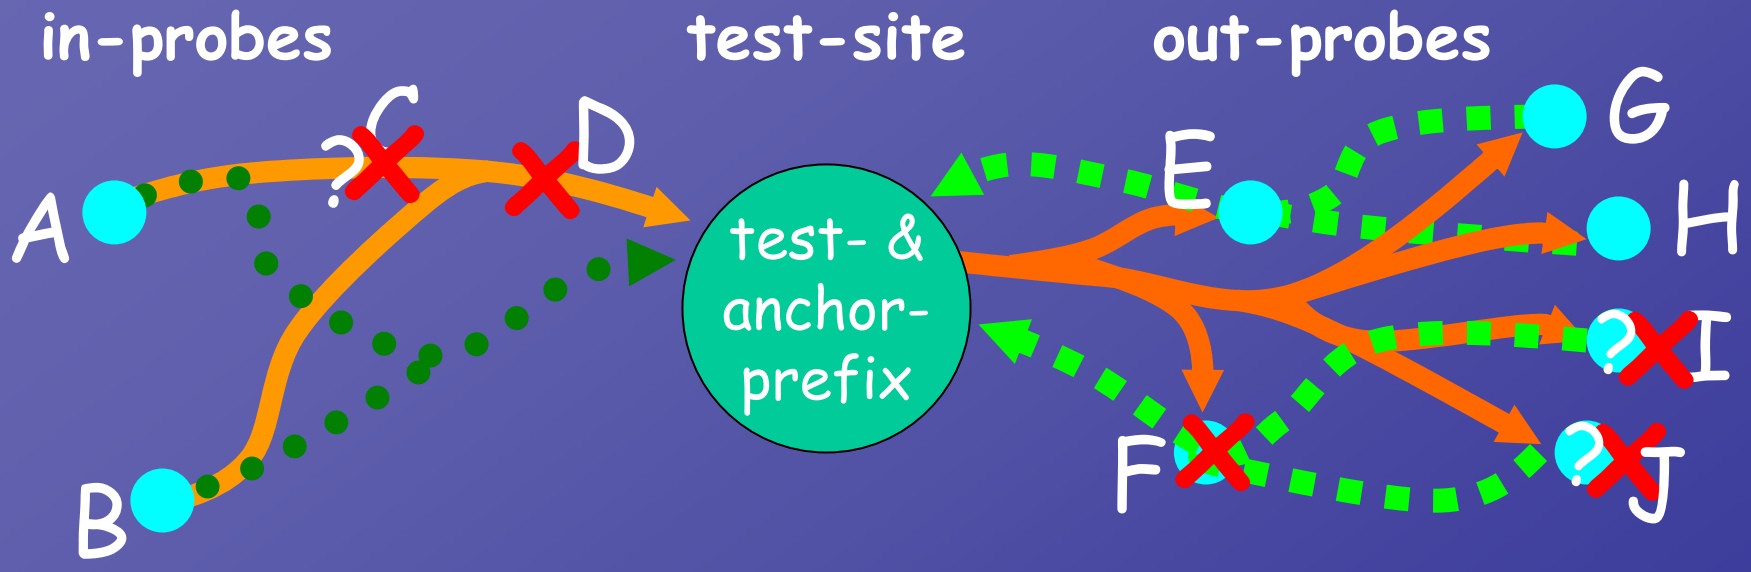

- In-probes : traceroutes from public traceroute servers to test- & anchor-prefix
- Out-probes : traceroutes *from test-site* towards pingable IPs. Source addresses are both test-IP and anchor-IP

Testing reachability of a new prefix: In-Probes

- Two IPs hosted at the same location:
  - anchor IP : well established, hopefully unfiltered
  - test IP : newly allocated address
- Assume that they are propagated in the same way (as they are announced from the same location)
- From each traceroute-server run two traceroutes : to test-IP and to anchor-IP

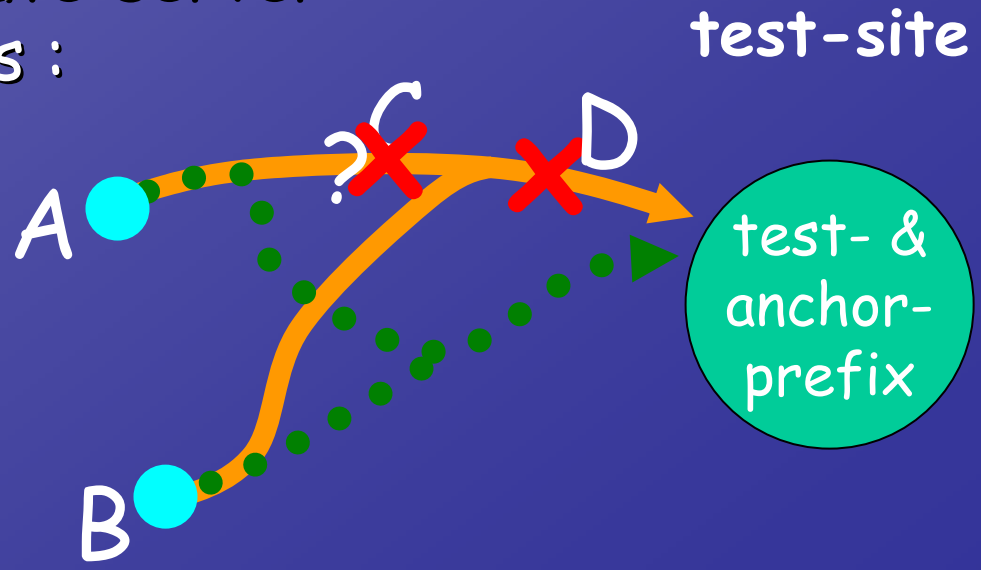

In-Probes: Principles

- In-probes give reachability information towards the test and anchor prefixes
- If traceroute from test-prefix traceroute diverges at some point, we build a list of possible candidates that might filter.

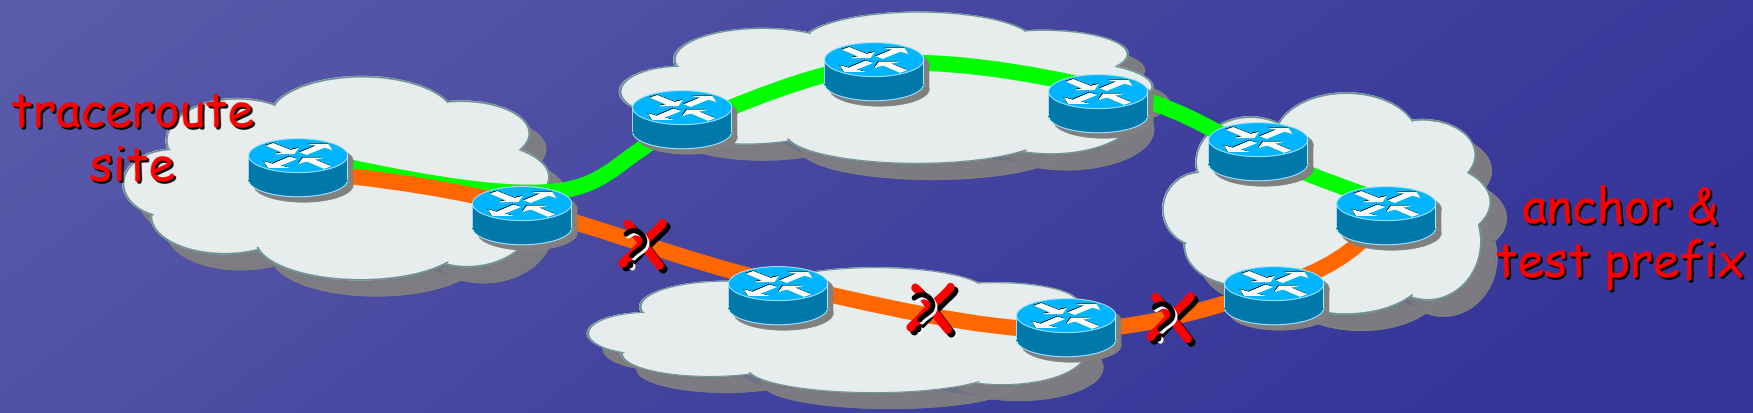

In-Probes: Limitation

- Catch only filters that are between public traceroute-server/looking glass and test-site.

=> can only find limited number of filters,
but identifies intermediate ASs that filter.

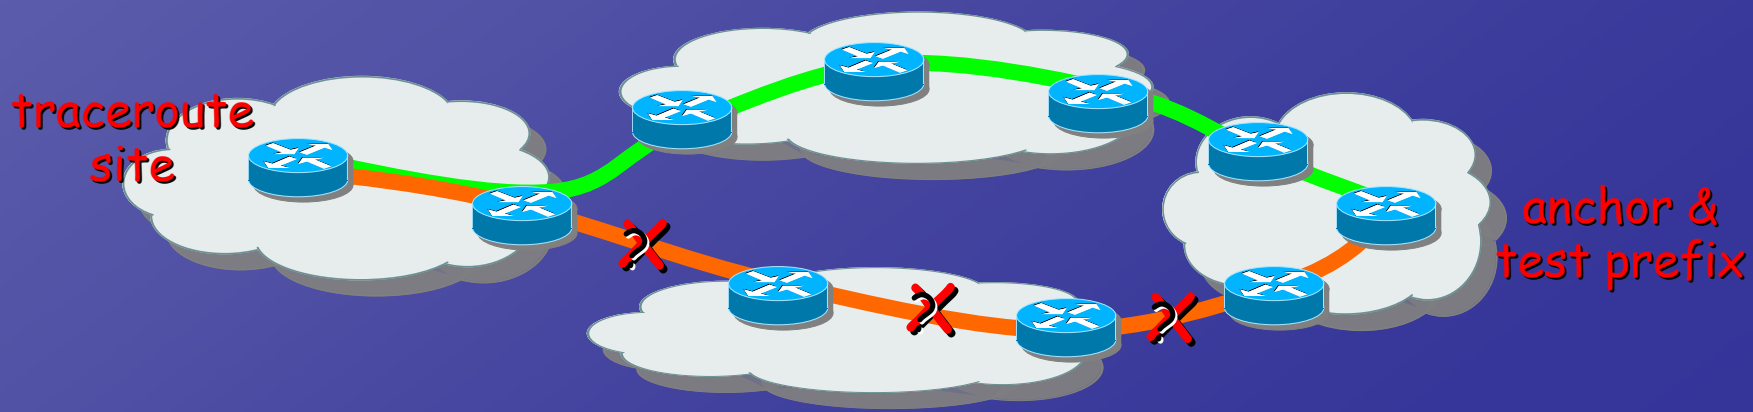

In-Probes: measurements

- Advertise test and anchor prefixes from 4 probe-sites: Seattle (USA), Munich (DE), Wellington (NZ), Tokyo (JPN)
- 480 public traceroute serves and PlanetLab nodes. Mainly US & Europe, but covering 56 countries
- Many volunteers from NANOG posting

In-Probes: results

- Categories:

- "good" (anchor and test take exactly same path)
  - 66.9%
- "diverging inside" (anchor and test take different paths)
  - 20.6%
- Test stops, but anchor ok
  - 8.6%
- Failure (either anchor or anchor and test failed)
  - 3.9%

In-Probes: results

- Derive candidate links, eliminate unlikely candidates.
- Remaining candidate links:
  - ~ 34 ASs that may contain incorrectly configured filters.

http://psg.com/filter-candidates.txt

In-Probes: evaluation

- Advantages:
  - traceroutes go around bogon filters
  - known details about IP-level path
- Disadvantages:
  - traceroute site MUST be "behind" bogon filter
  - Never enough traceroute sites available

- Goal: test as many ASs as possible for reachability
- Solution: "out-probes"

Testing for usable reachability: Out-Probes

- Out-probe : ping and traceroute performed *from* test-IP and anchor-IP towards external IP addresses
- Return-Path is of interest, but unknown
- What we learn is which AS has connectivity
- Why it works:
  - high AS coverage
  - only usable connectivity

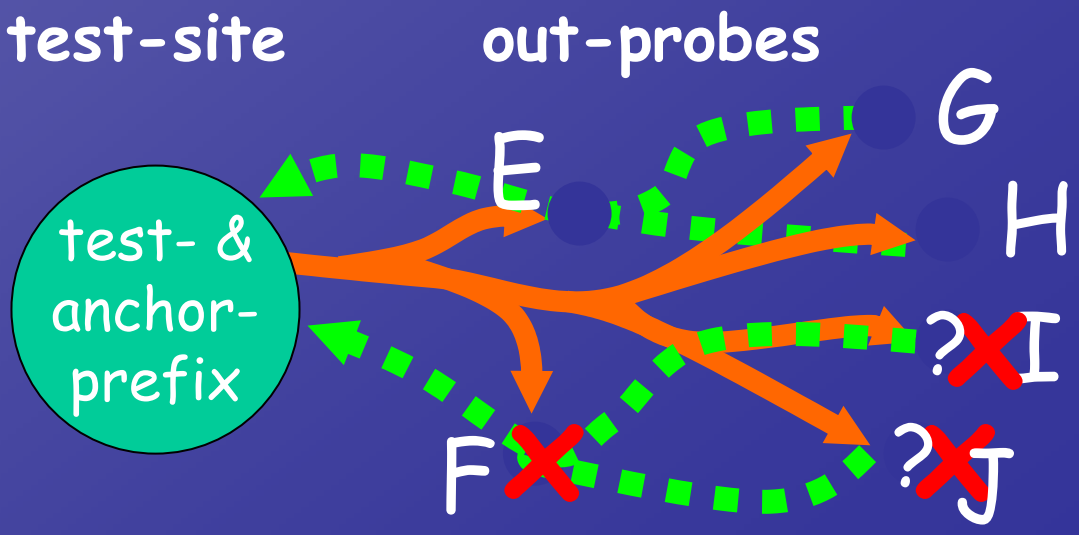

Out-Probes: measurements

- Perform ping from test-sites (test-IP and anchor-IP) towards a large set of pingable-IP addresses (46,569) in 18,574 different ASs
- If ping comes back => usable reachability :)
  (~85% of all probes)
- If ping does not come back => annotate anchor link with "proximity" index.
  (~10% of all probes)
  => roughly 2,500 ASs (!) (depends on probe site)

- (~5% not pingable anymore, e.g., dial-up)

Out-Probes: measurements

- Build filtering likelihood index based on "proximity" and per AS observations

- x-axis index: aggregate all observations, normalize, and weight with "proximity index"

=> provides ordering of ASs that are likely to filter

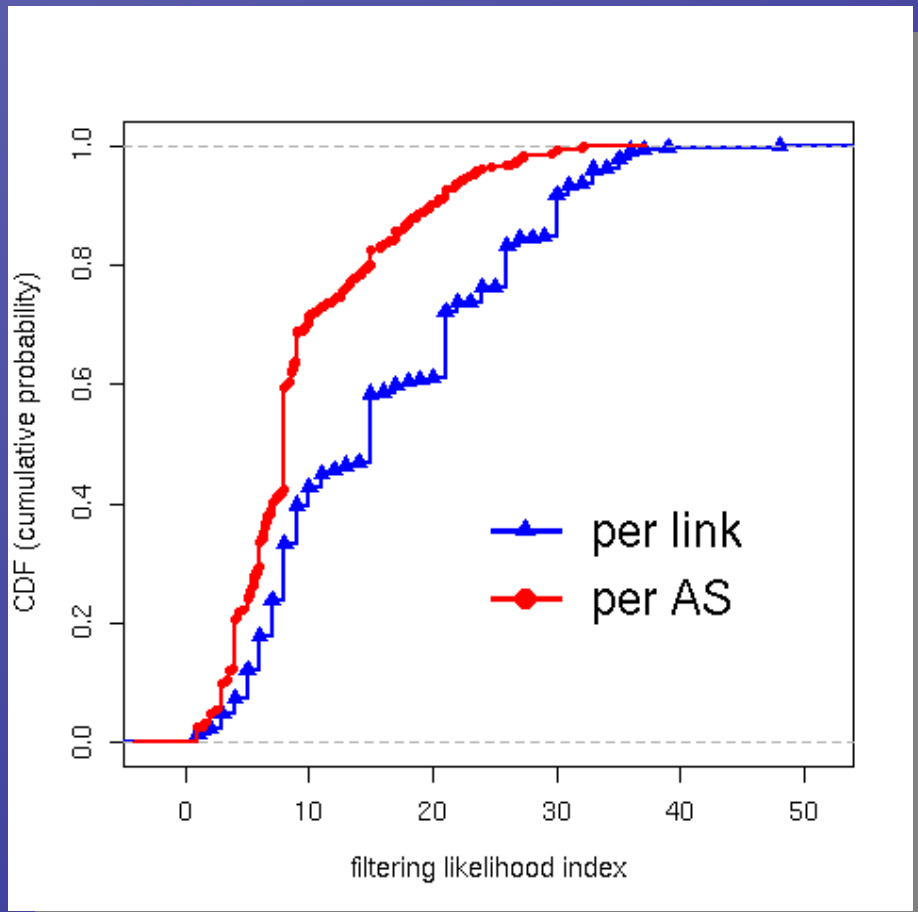

Out-Probes: Initial validation

- We derived 443 candidate ASs that are likely to filter.

- manual search for 15 traceroute servers within those 443 candidate ASs:
  - 7 filter
  - 5 do not filters themselves, but have no usable [up-stream] connectivity.

  => 12 out of 15 (=80%) correctly identified

  3 failed, but validation was taken at different time. Thus, ASs might have changed filter in meantime.

Summary: In- and Out-Probes

- Out-probes tell about "usable reachability":
  - Find areas of non-reachability
  - Larger coverage (currently > 85% of Internet ASs)
  - No information about: return path and thus non-optimal paths
- In-probes tell us about filters on the path:
  - Reachability available - goal: detect intermediate filters
  - Smaller coverage
  - Many traceroute servers are needed at the "edge"

Further Work

- Sent list of candidate suspected bogon filtering links to ISPs, waiting for their feedback to validate our analysis
- Increasing number of in-probes to have more information about location of bogon filters and their number
- How accurate can we be in identifying bogon filters using measurements?
- How would we quantify that accuracy?
- How many out-probes are needed/useful?

Results - Out-Probes

- We can identify unreachable places: Via out-probes we can see if an IP is not well routed.
- Aside from small issues related to ICMP, we know that if the probe doesn't come back that there is NO usable connectivity. That's simple and straight forward.
- It is possible to achieve a reasonable coverage of the Internet (<18k ASs).
- The methodology produces usable results.

Results - In-Probes

- We can go a step further and detect places where there is "non-optimal" connectivity.
- Keep in mind that with the in-probes we mainly look at traceroutes that BOTH reach the destination.
- We would very much like more validation by the operator community

How you can help...

- We plan to establish an ongoing service.
- For that we need:
  - pingable addresses
    Tell us about addresses that we can ping once in a while and we make sure that you have connectivity to newly allocated prefixes.
  - traceroute servers
    Tell us about traceroute servers, so that we can improve the quality of our inference.

Thanks To

- ARIN for IP space and commissioning research
- CityLink – NZ, a test site
- IIJ - JP, a test site
- SpaceNet - DE, a test site
- PSGnet – US, a test site
- Universities of Adelaide & Delft
- NSF award ANI-0221435
- Australian Research Council grant DP0557066