UCL Université catholique de Louvain

UNIVERSITÉ DE STRASBOURG

MERLIN

Measure the Router Level of the INternet

Pascal Mérindol, Benoit Donnet,
Jean-Jacques Pansiot, Matthew Luckie, Young Hyun

Kaiserslautern - June 2011

WAND Network Research Group

Next Generation Internet 2011

caida

Agenda

- Topology Discovery Background
- Limitations using `mrinfo-rec`
- A new probing tool: MERLIN
- Deployment and IGMP filtering
- Towards the MERLIN platform ?

ip networking lab
uci, louvain-la-neuve, belgium

LSiiT

Topology Discovery

- Internet seen as a dynamic graph
  - ✓ of IP interfaces
    - ➡ `traceroute, route_record`
  - ✓ of routers
    - ➡ alias resolution : `ally, iffinder, ...`
  - ✓ of Autonomous Systems
    - ➡ IP to AS mapping (`routeview project`), router to AS mappping ?
- Goals
  - ✓ IP network models & simulations
  - ✓ *ground truth* input for topology generation

mrinfo

- Topology discovery using `mrinfo`
  - Uses IGMP messages
    - ✓ **`ASK_NEIGHBORS`**
    - ✓ **`NEIGHBORS_REPLY`**
  - Output
    - ✓ All multicast interfaces of a given router
    - ✓ All multicast neighbors/links
- `mrinfo` applied recursively
  - **`mrinfo-rec`**
    - ✓ probe all neighbors
    - ✓ daily based

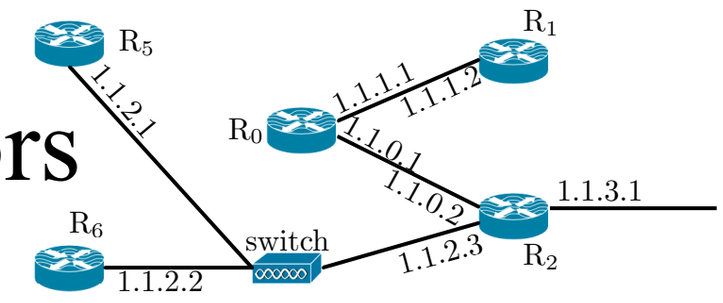

1.1.0.2 [version 12.4]
1.1.0.2 → 1.1.0.1 [1/0/pim/querier]
1.1.2.3 → **1.1.2.1** [1/0/pim/querier]
1.1.2.3 → **1.1.2.2** [1/0/pim/querier]
1.1.3.1 → 0.0.0.0 [1/0/pim/leaf]

mrinfo-rec

- Global Limitations
  - multicast scope
  - IGMP filtering (local and transit)

- Technical Limitations
  - IGMP fragmentation
  - lack of multiplexing (no port number)

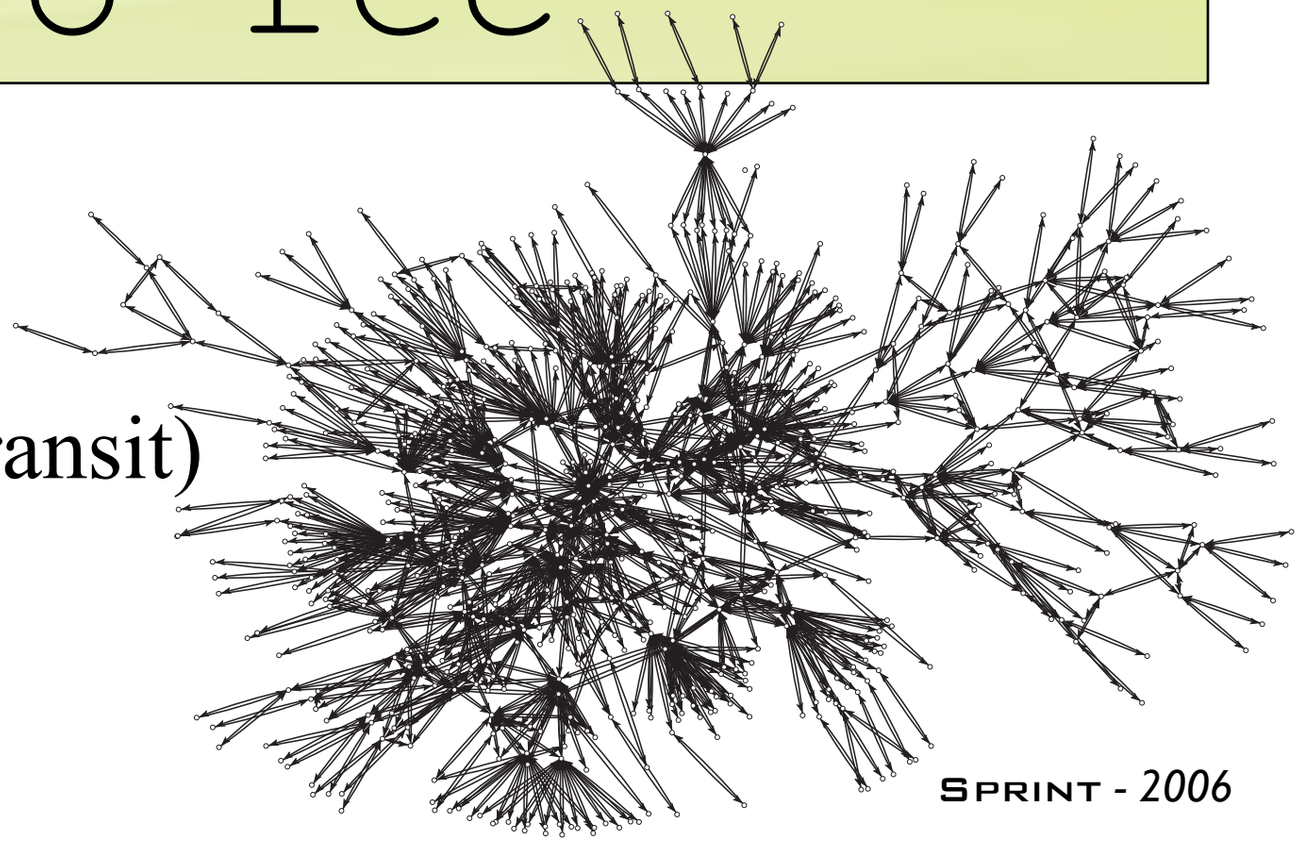

- Advantages
  - network friendly probing: 1 probe injected per router
  - native router level vision: no need for alias resolution
  - forwarding independent: backup links visible *[IMC2009]*
  - layer-2 vision: distinguish the IP layer over MAC *[IMC2010]*

Limitations

- `mrinfo-rec`: ~ 4 years of daily collected data
  - ➡ ~10000 routers ~100000 IP ~300-800 AS

- Only a single vantage point in Strasbourg
  - IGMP transit filtering issue: some (borders) routers do not forward IGMP requests/replies

- IGMP fragmentation: large Cisco routers «IGMP-fragment» their responses (576 bytes at maximum)

- No multiplexing: use multiple IP addresses or ignore replies where target IP ≠ reply IP ?

IGMP fragmentation

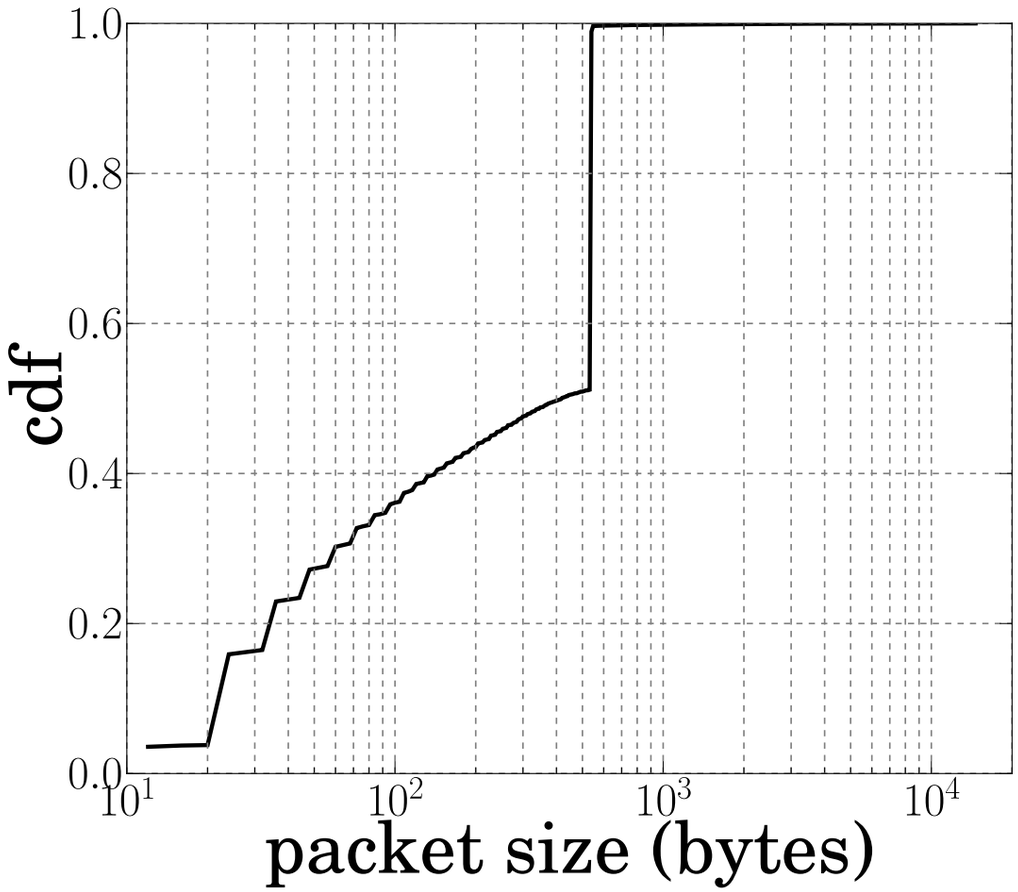

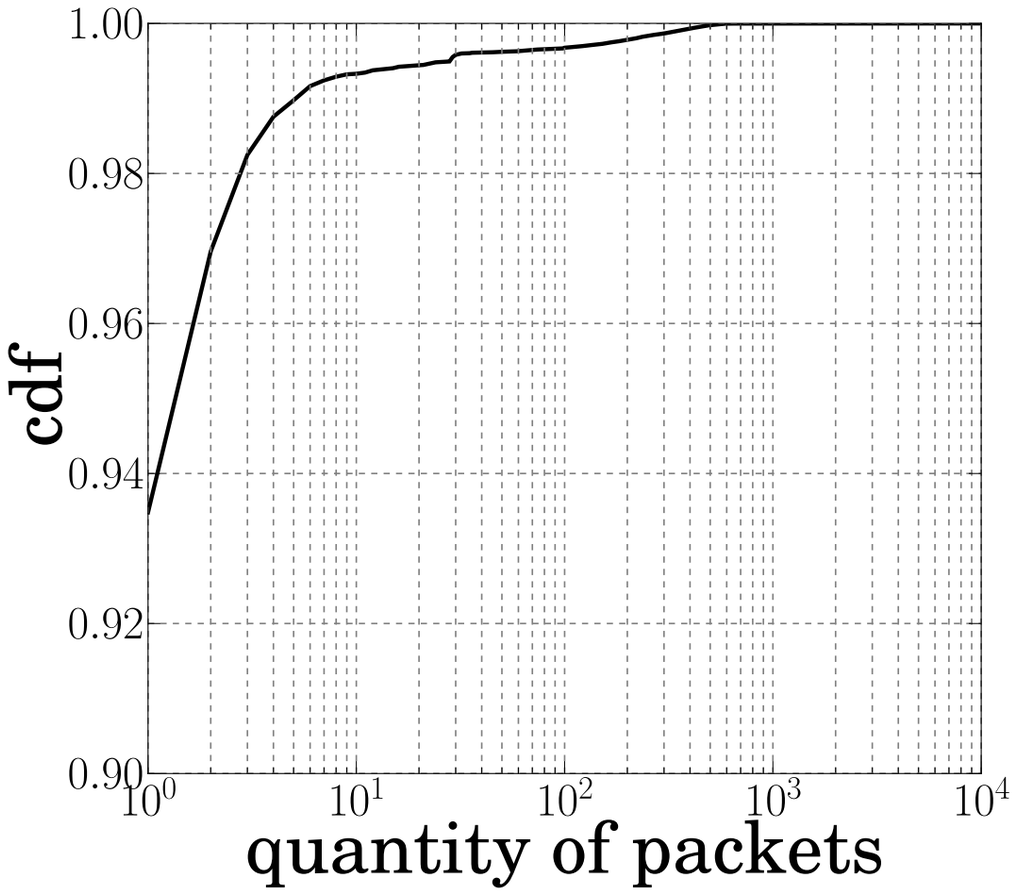

$$|\text{headers}| + \sum_{i=1}^{n} (8 + 4 \times m_i).$$

- A few number of routers generates fragments (~6%)
- ...but they generates almost half of the replying traffic !

Limitations per router brand

- Fingerprints ability

Probable brand	Version[4]	Proportion
Cisco IOS	11.*, 12.*, 15.*	78.25%
Juniper	3.255	7.61%
Not classified	[0-9].*, 21.3, 21.95, 37.90, 60.1, 76.0	13.12%

- Cisco routers «IGMP-fragments»:
  - how to collect subsequent responses ?
- Juniper routers IP-fragments: OK transparent for mrinfo-rec but not correct according to the draft

- Some non Cisco routers (~10%) have an «instead of» behavior: the IP of reply is not the one targeted!
  - how to speed up the probing process ?

MERLIN

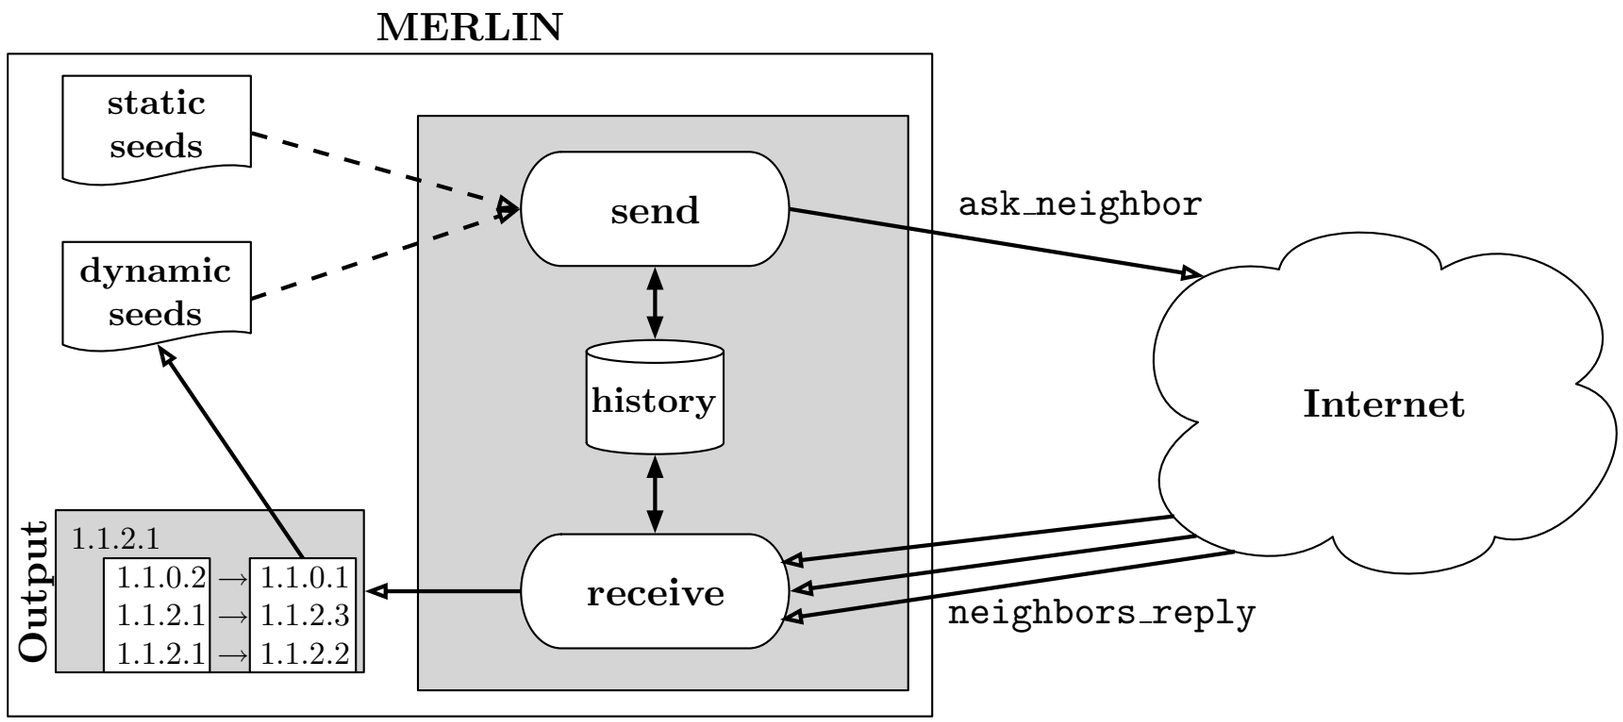

- Two parallel processes: *send & receive*
  - replies are indexed on the *src addr* (multiplexing)
  - fragments having the same *src addr* are merged (fragmentation)
- History process to avoid probing redundancy
  - hash based for performance (for all local IP interface)
- Two seeding lists: *static & recursive*
  - recursive first approach

Reprobing risk and Calibration

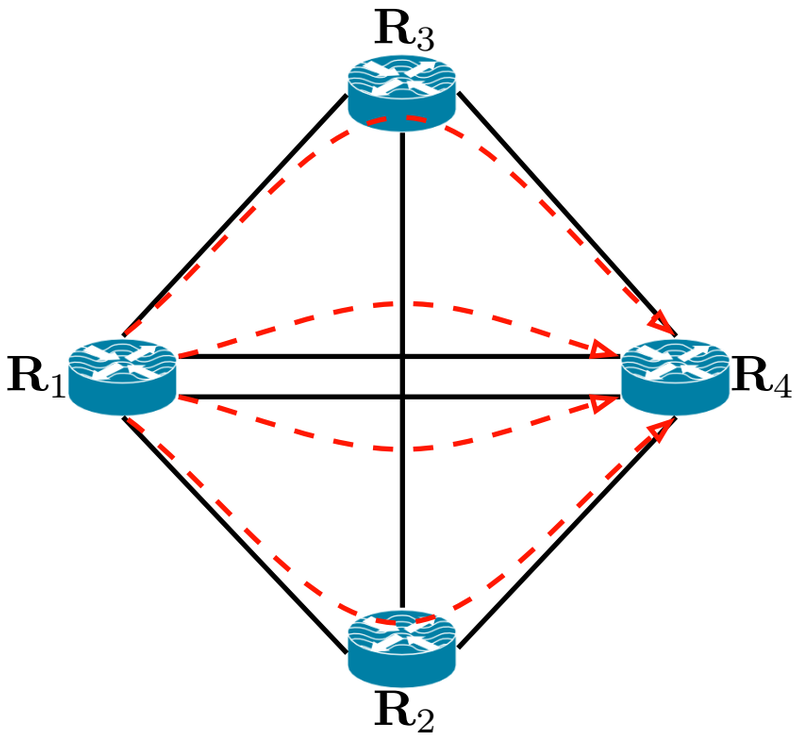

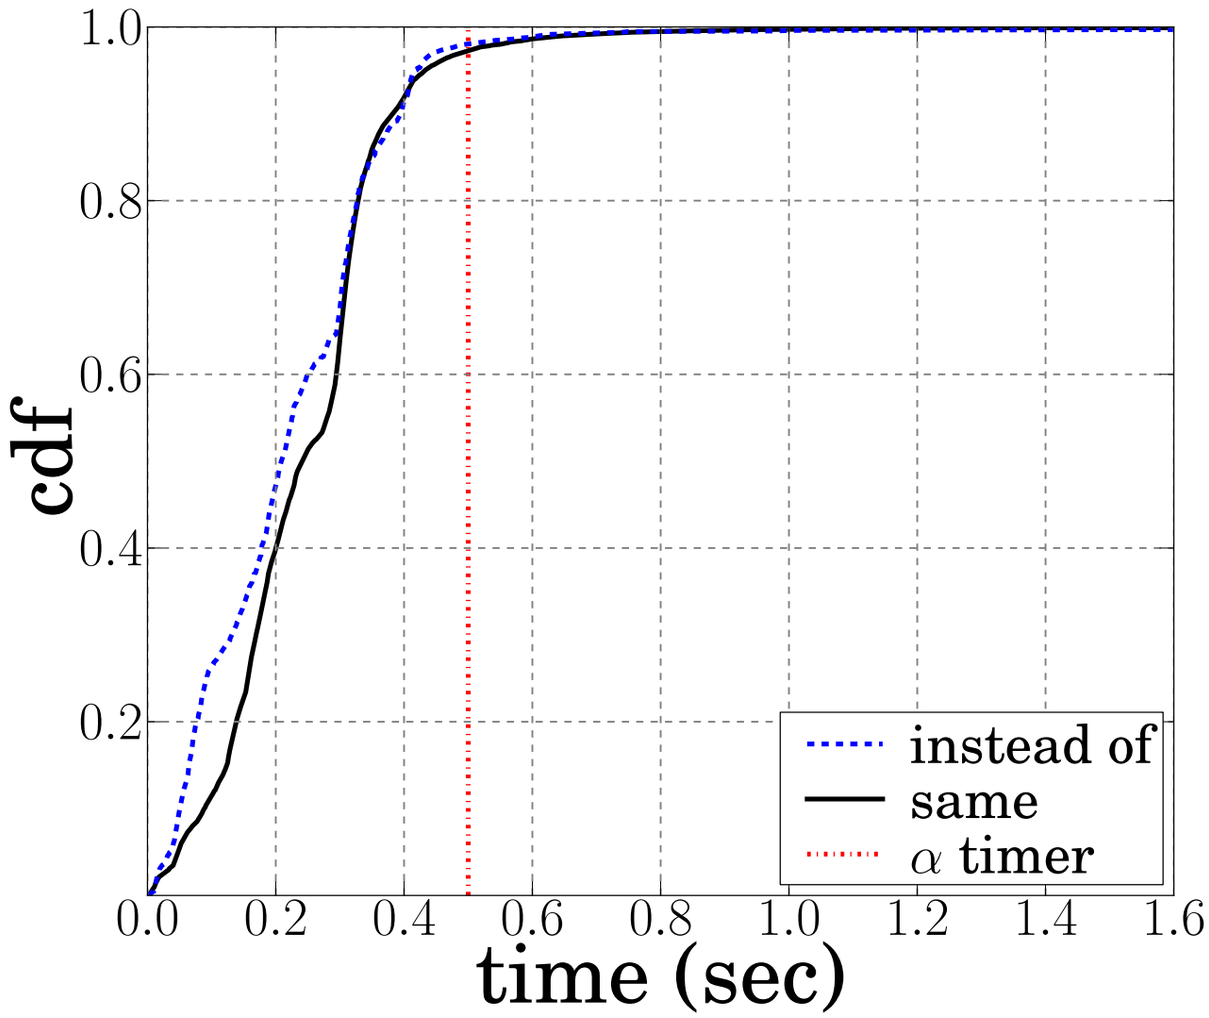

- Two probing modes:
  - recursive, α=*0.5* sec ➥ to elapse probes and reduce reprobing risk
  - static, β=*0.05* sec ➥ to speed up the probing campaign when the reprobing risk is low
- Replies are flushed every 5 sec: fragments reassembling (~0.1 sec)

MERLIN behavior

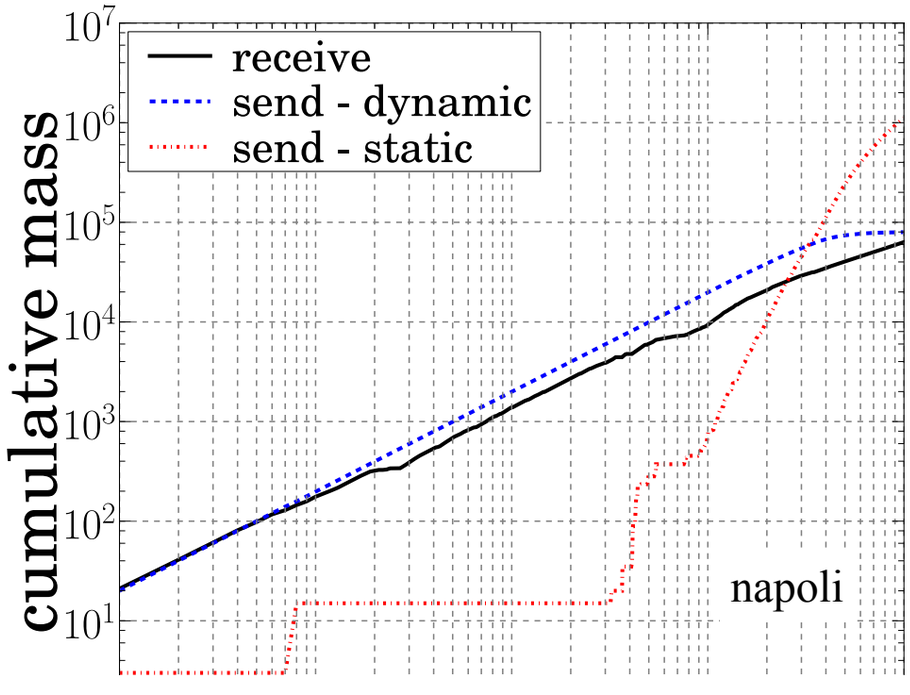

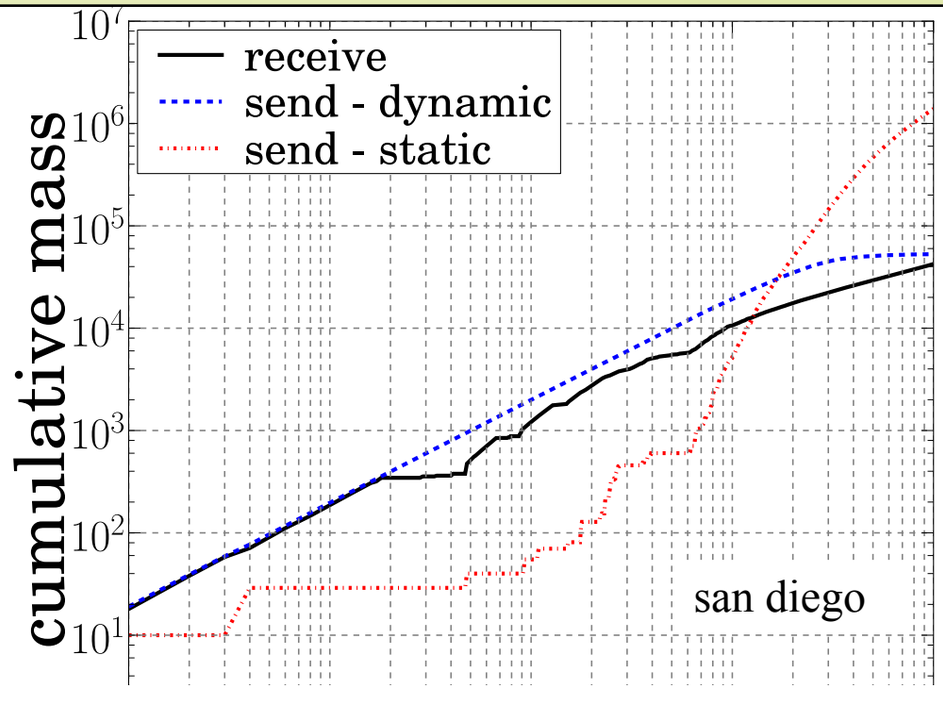

- Recursion does the job first and then static list finishes it...

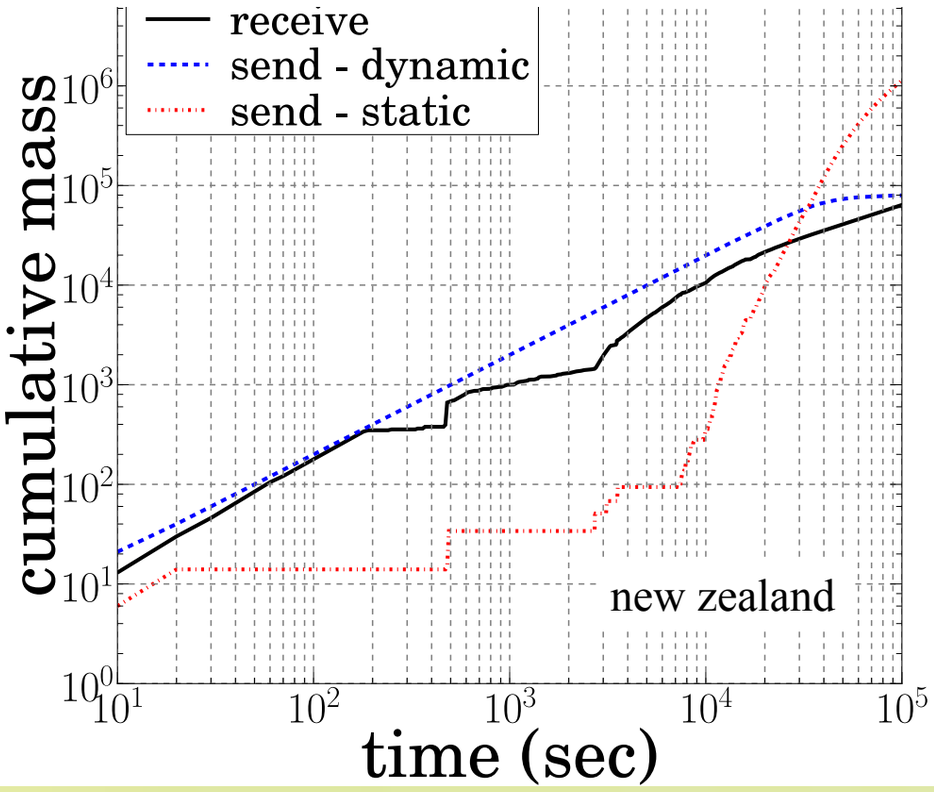

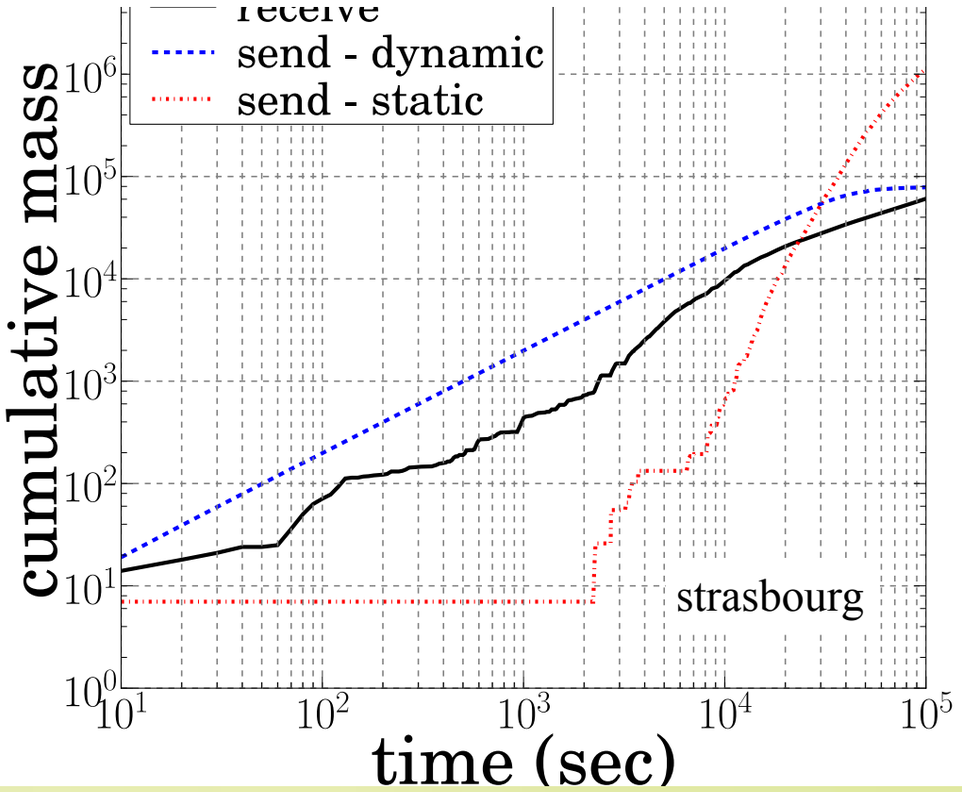

Deployment

- 6 vantage points:
  - (Louvain-la-Neuve - Belgium, Napoli - Italy, Strasbourg - France), two in North America (San Diego - USA, Redwood City - USA), and one in Oceania (Hamilton - New Zealand)

- The probing hitlist is made of:
  - 1.2 M Caida's Archipelago addresses;
  - ``missing middle'' IP (Archipelago);
  - 3,580 addresses from known topologies;
  - 24,429 addresses from a Tier-1 ISP;
  - 155,674 Reverse Traceroute addresses;
  - 224,762 mrinfo-rec addresses replying on the four previous datasets.

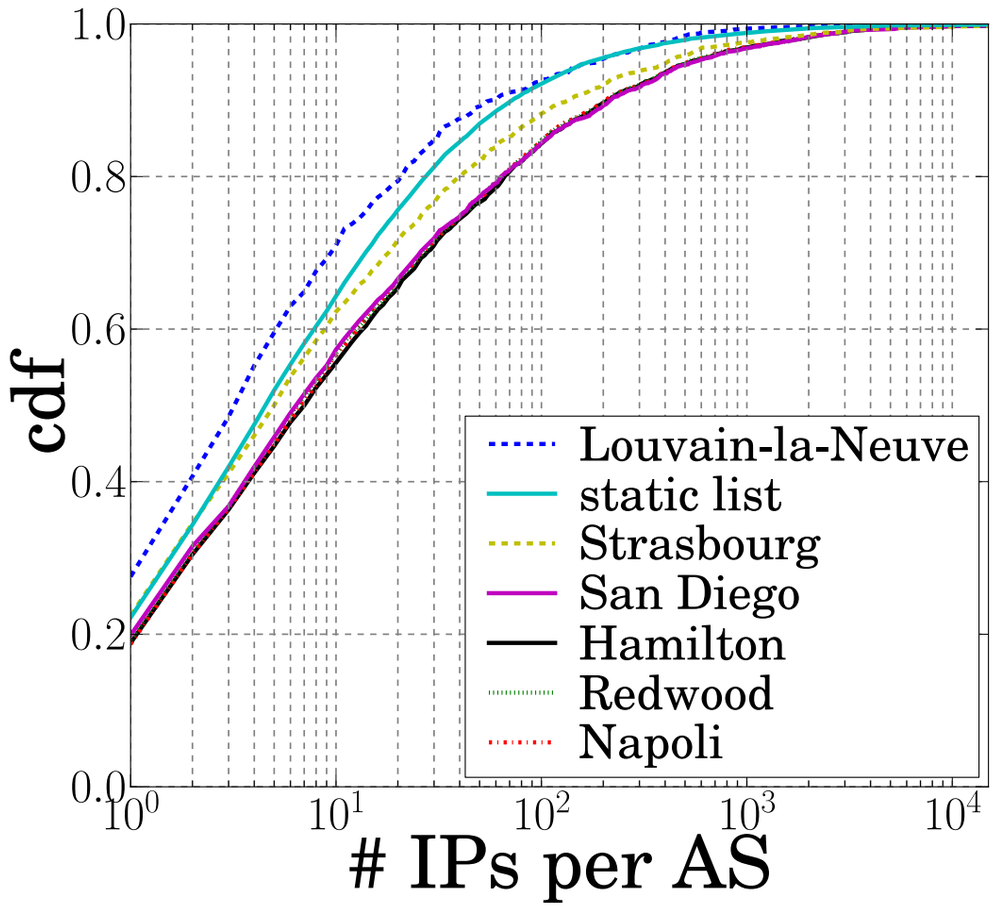

- ~50,000 unique routers in 3000 ASes
- A global and non uniform coverage greater than 5%

Unicast lacks

- The reply's *src addr* may not appear in the list of interfaces

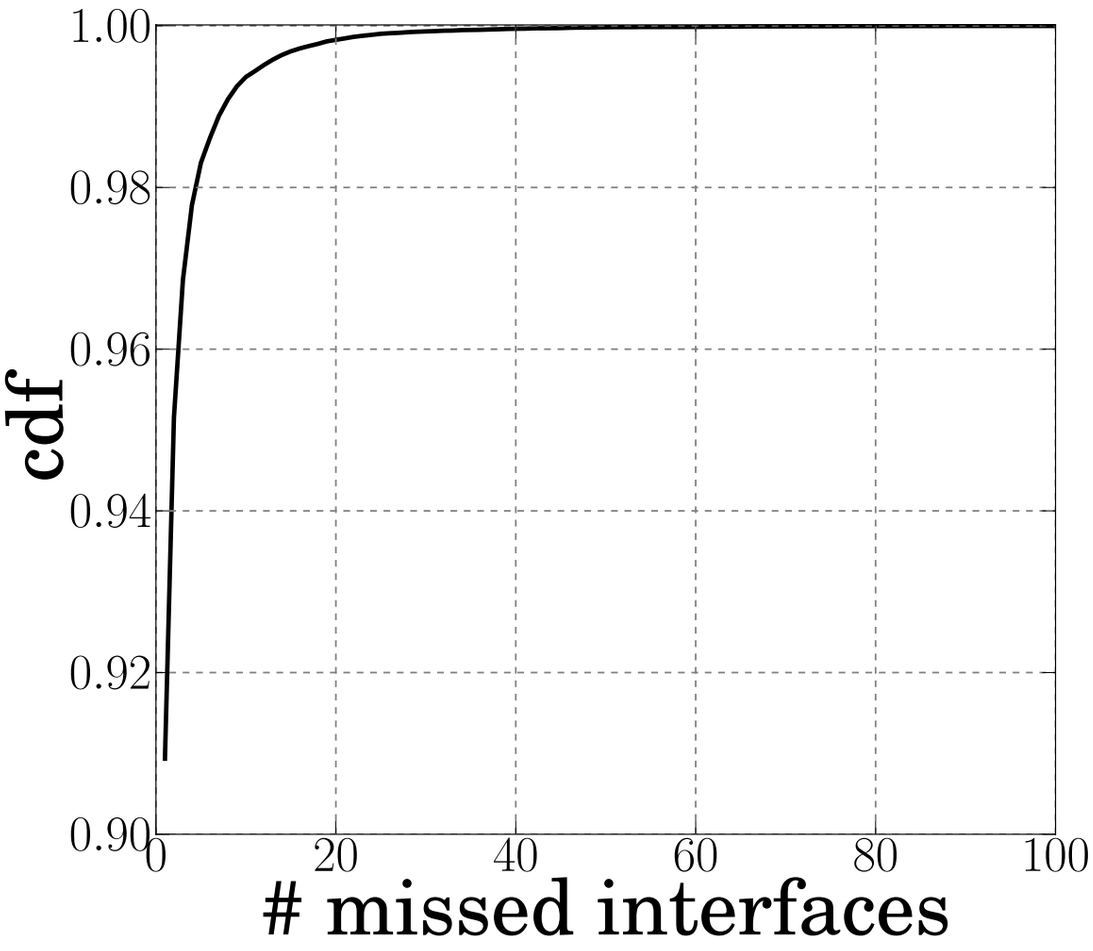

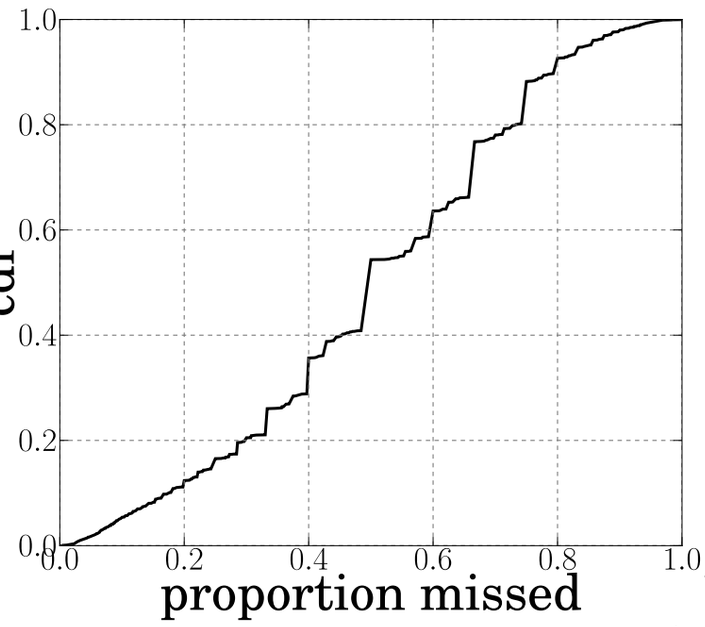

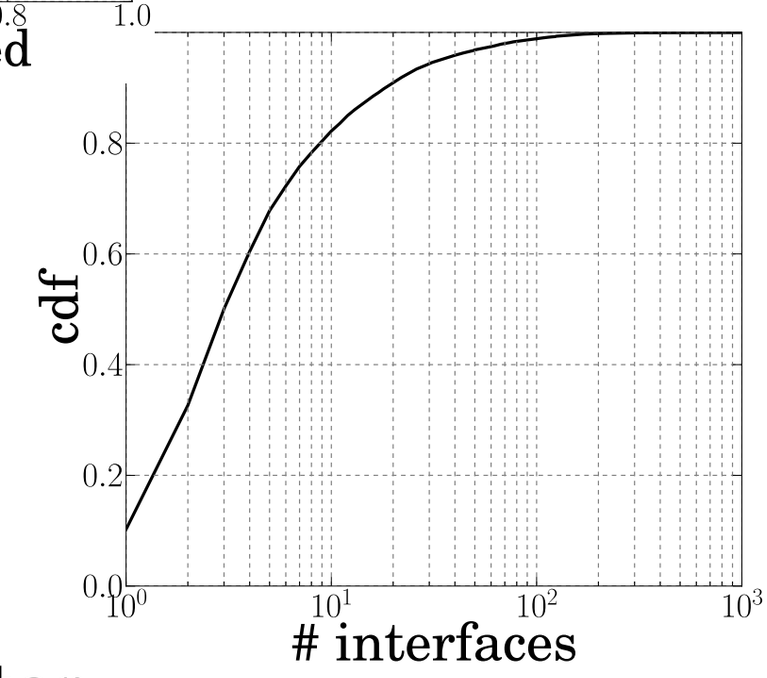

- **IGMP unicast alias resolution**
  - missing unicast IP are added to the router

IGMP filtering: monitor's utility

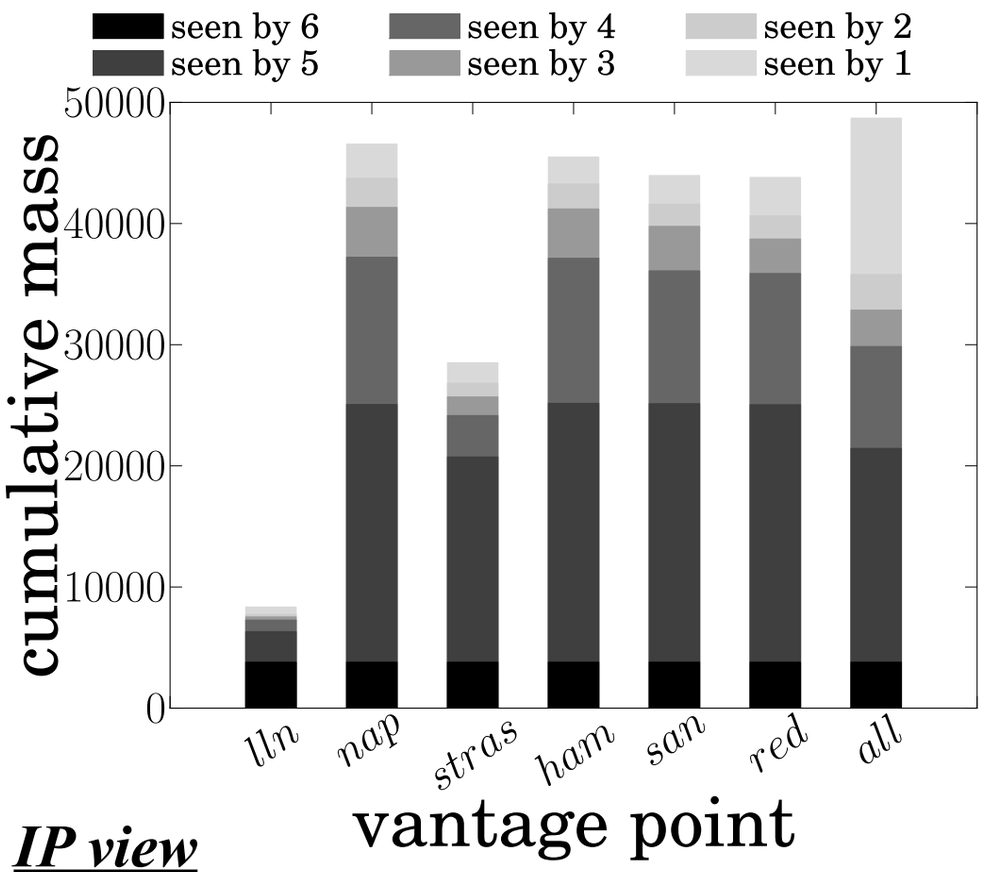

IP view

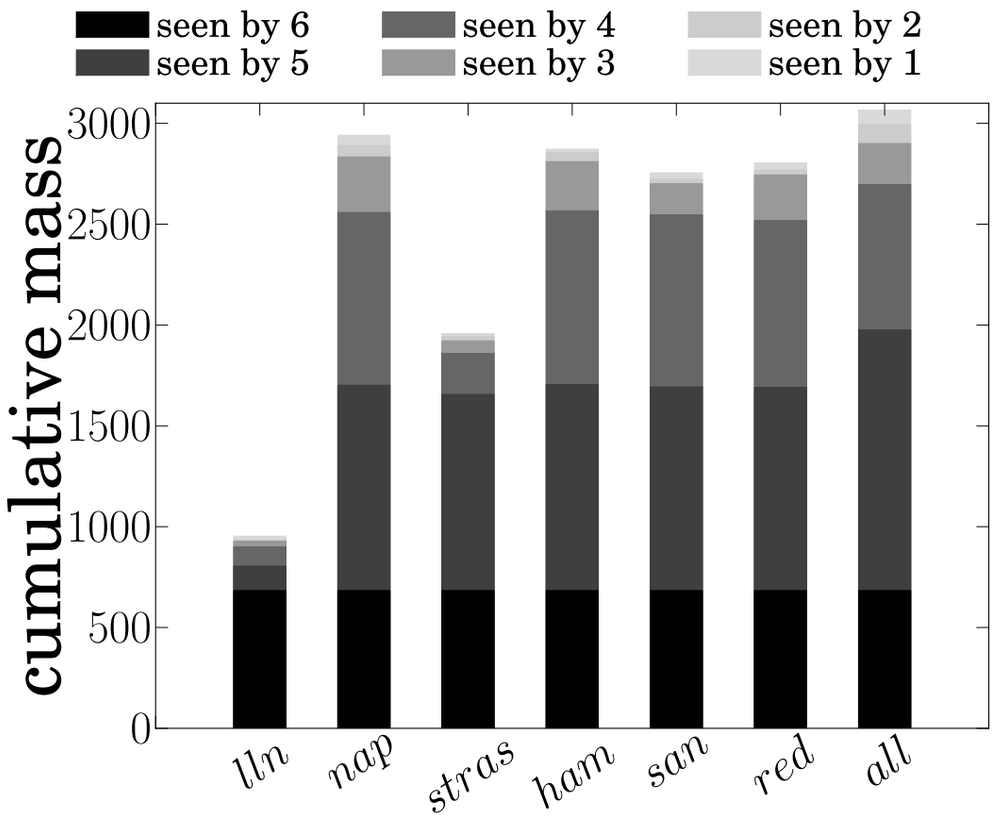

AS view

- The utility of using several vantage point (vp) is high: *seen by 1*
  - Some vps are less subject to IGMP filtering
  - Each vp brings its unique contribution
  - Even inside a given AS, the utility of several vps is high

The MERLIN platform

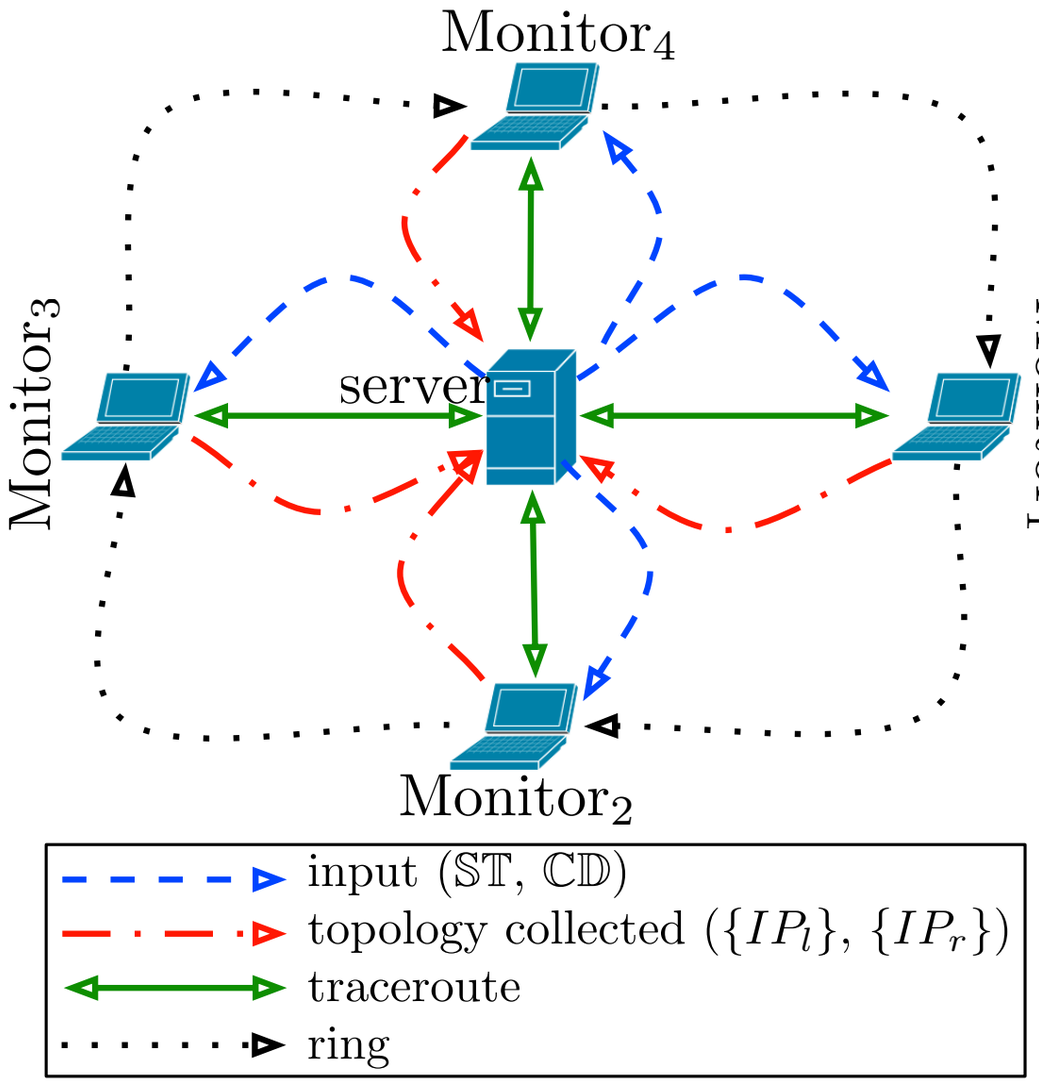

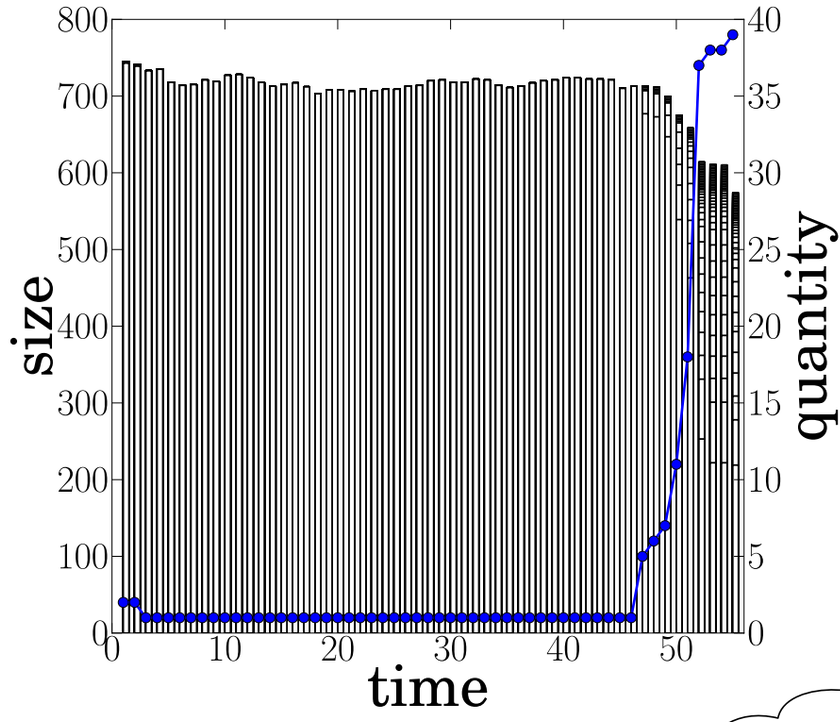

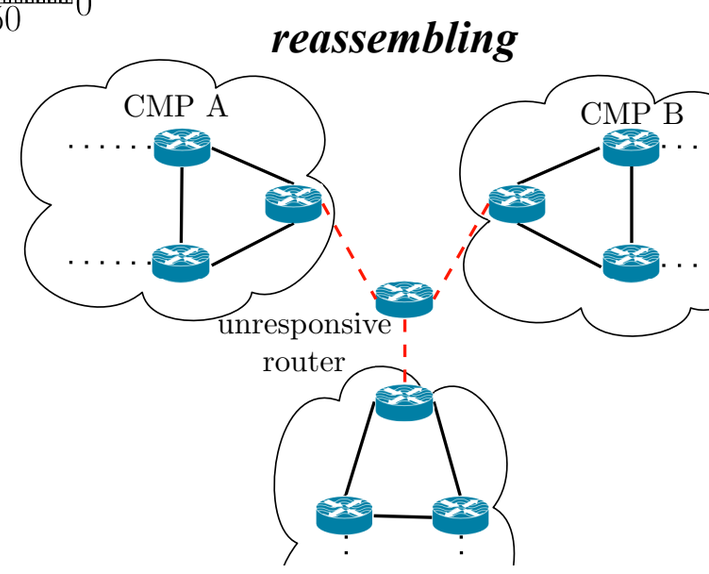

- MERLIN monitors can be coordinated via a central server
  - avoid redundancy and improve efficiency
  - use active and targeted traceroute for seeding and reassembling (+ alias resolution)
- The MERLIN platform targets multicast enabled AS cores

Conclusion

- **IGMP probing** is a useful for several reasons
  - **describe a connected multicast topology at the router level (no need for alias resolution)**
  - **can discover backup links (no forwarding dependence)**
  - **able to natively infer L2 devices (hybrid bipartite graph)**
  - **efficient probing scheme**

- **MERLIN** solves mrinfo and mrinfo-rec lacks
  - **technical issues: fragmentation and multiplexing**
  - **is fed per traceroute and recursive seeds**
  - **IGMP filtering and unicast lacks can be solved**
  - **can be plugged in an client/server platform**

Questions ?

http://svnet.u-strasbg.fr/merlin/

- ✓ Pietro Marchetta, Pascal Mérindol, Benoit Donnet, Antonio Pescapé and Jean-Jacques Pansiot.
  Topology Discovery at the Router Level: A New Hybrid Tool Targeting ISP Networks.
  In IEEE JSAC, Special Issue on Measurement of Internet Topologies, 2011.

- ✓ Pascal Mérindol, Benoit Donnet, Jean-Jacques Pansiot, Olivier Bonaventure.
  On the Impact of Layer-2 on Node Degree Distribution
  In Proc. ACM/USENIX Internet Measurement Conference (IMC), November 2010.

- ✓ Jean-Jacques Pansiot, Pascal Mérindol, Benoit Donnet, and Olivier Bonaventure.
  Extracting Intra-Domain Topology from mrinfo Probing
  In Proc. Passive and Active Measurement Conference (PAM), April 2010.