Measuring ISP Network Topologies with Rocketfuel

Neil Spring, Ratul Mahajan, and David Wetherall

University of Washington

ACM SIGCOMM 2002

Motivation

We’d like to understand Internet structure and design.

- ISP router-level topologies are designed.

We can’t get the real maps.

- Backbone maps often available in marketing form.
  Severely lacking in router-level detail.
- When we asked for them, they said “no.”

ISP topologies for research

Could extract from a Whole-Internet map:

eg. Skitter, Mercator, Lumeta.

Our philosophy:

- By focusing on an ISP, can get better precision.
- ISPs publish enough information to reconstruct maps.
- End goal is more accurate maps for research.

Rocketfuel overview

Integrate all information exported by an ISP:

- BGP - which prefixes are served
- Traceroute - what the paths are
- DNS - where routers are and what they do

Build detailed maps:

- Backbone
- POPs
- Peering links

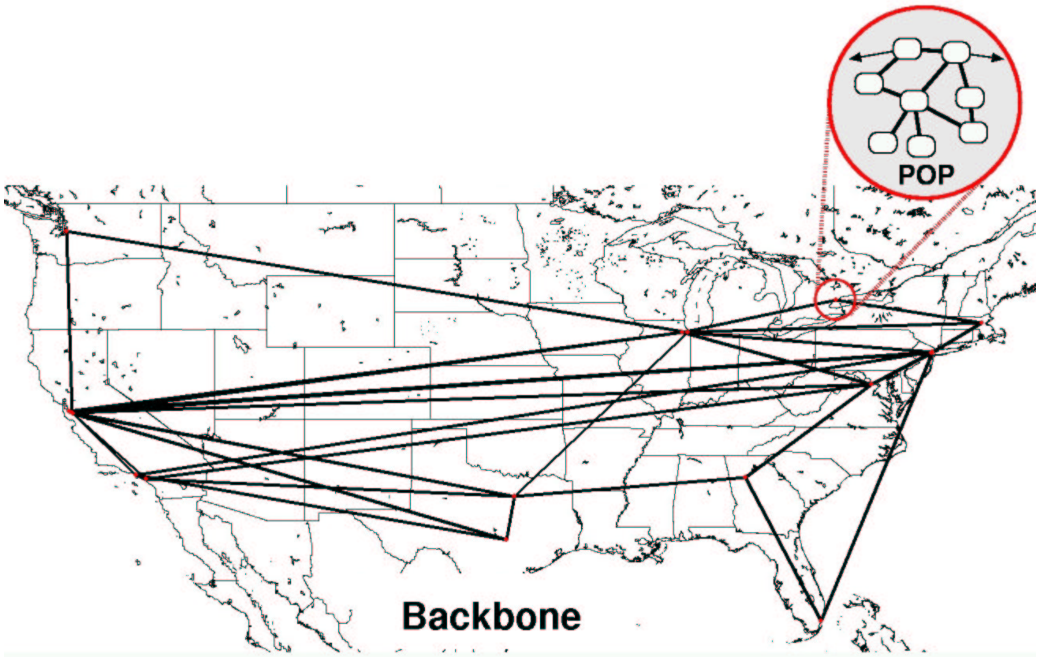

Challenges

Is it possible to collect accurate maps?

- Many public traceroute servers.
- Can’t use them all the time.
- Restricting load, it would take a year.
- Balance accuracy and speed.

Can we make any sense out of the result?

- POP / Backbone structure
- Alias resolution

Choosing traceroutes

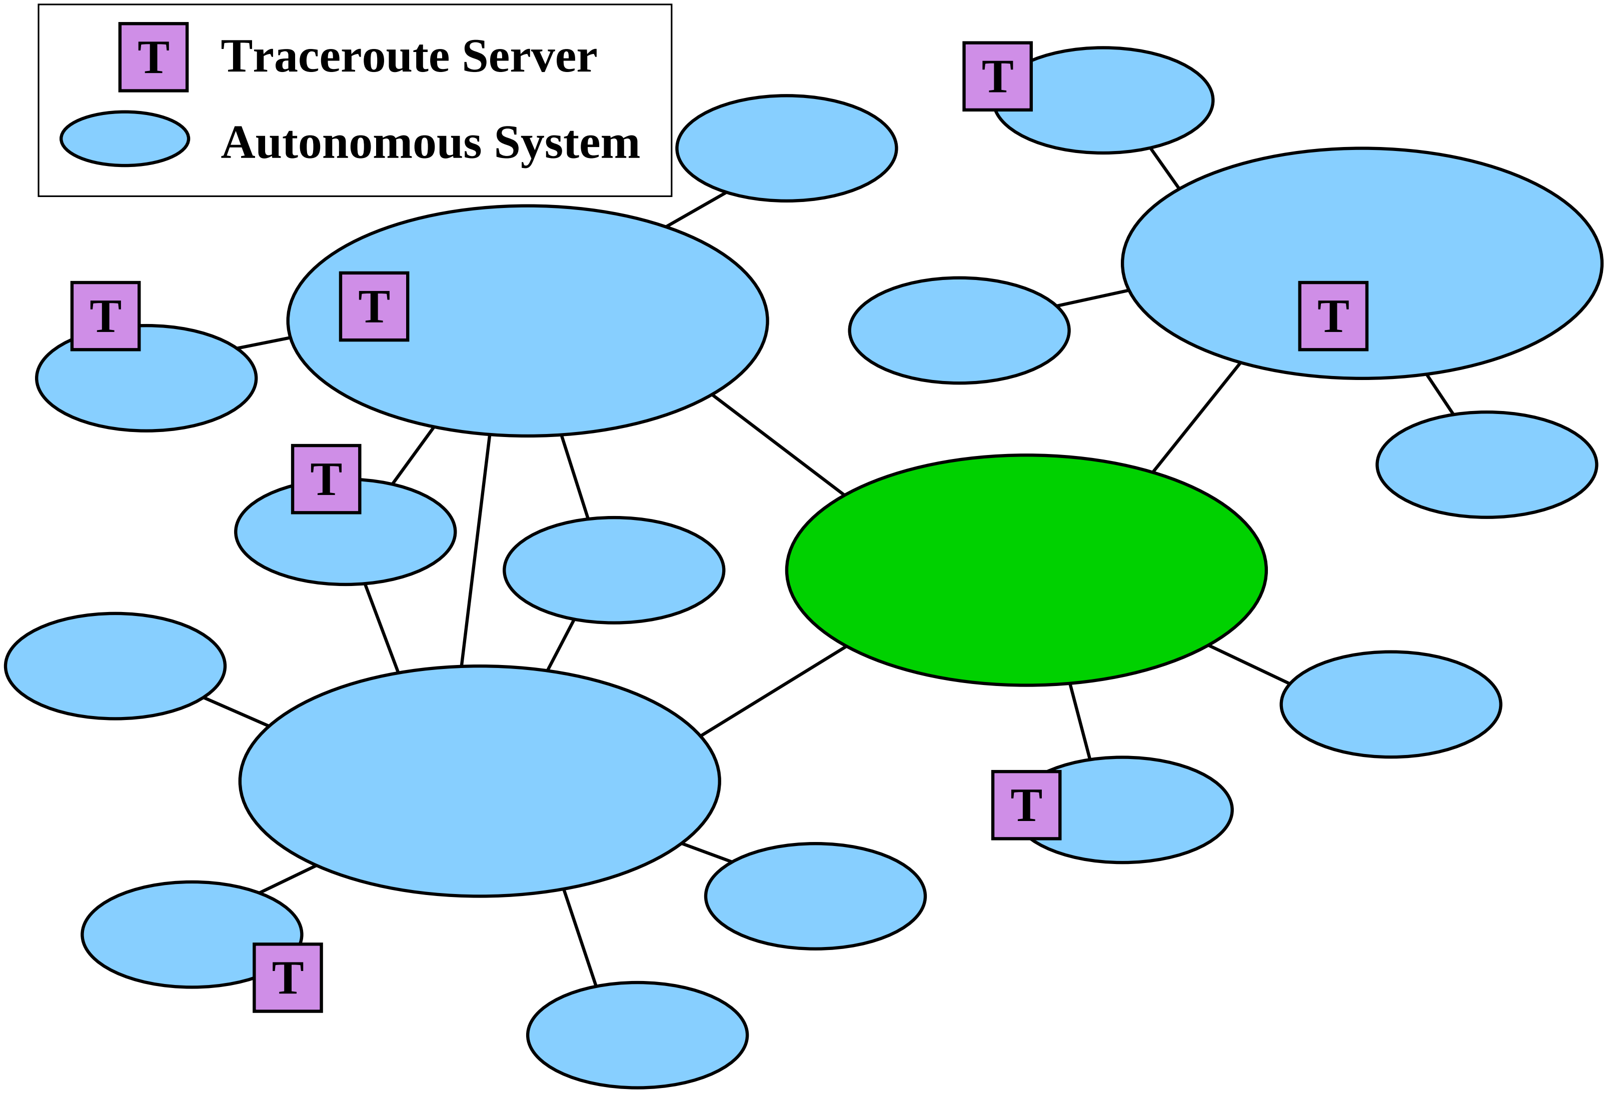

Choosing traceroutes

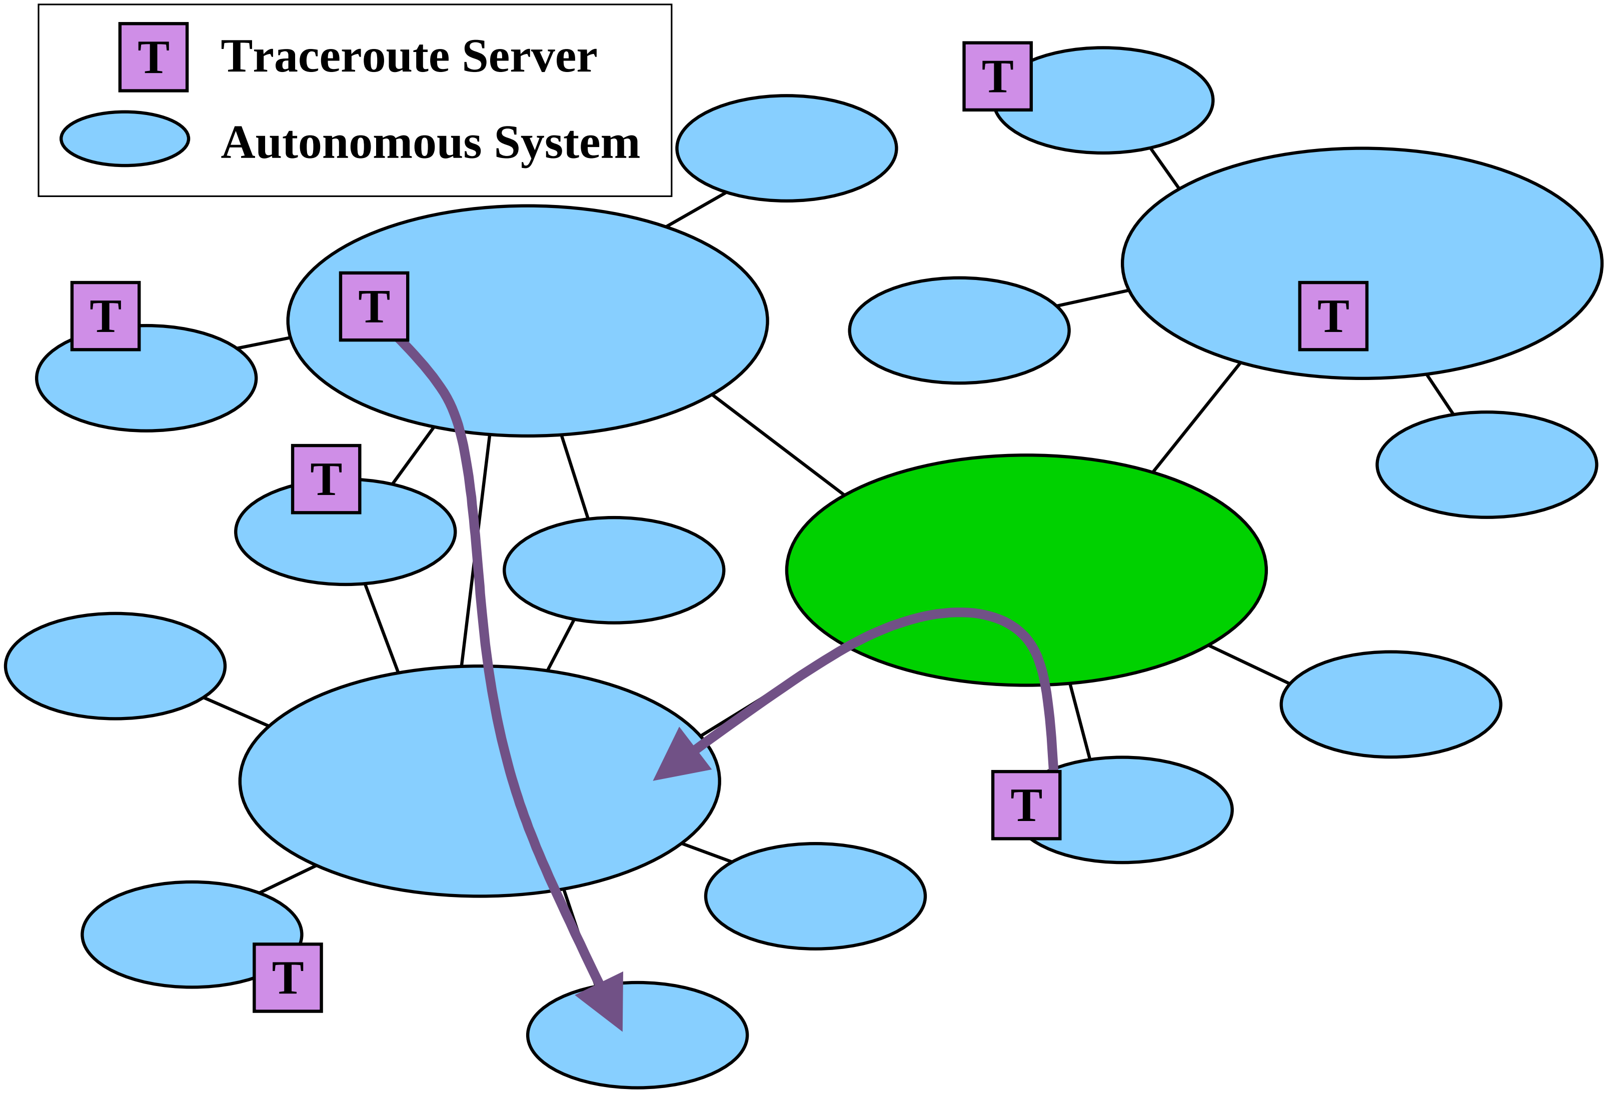

Choosing traceroutes

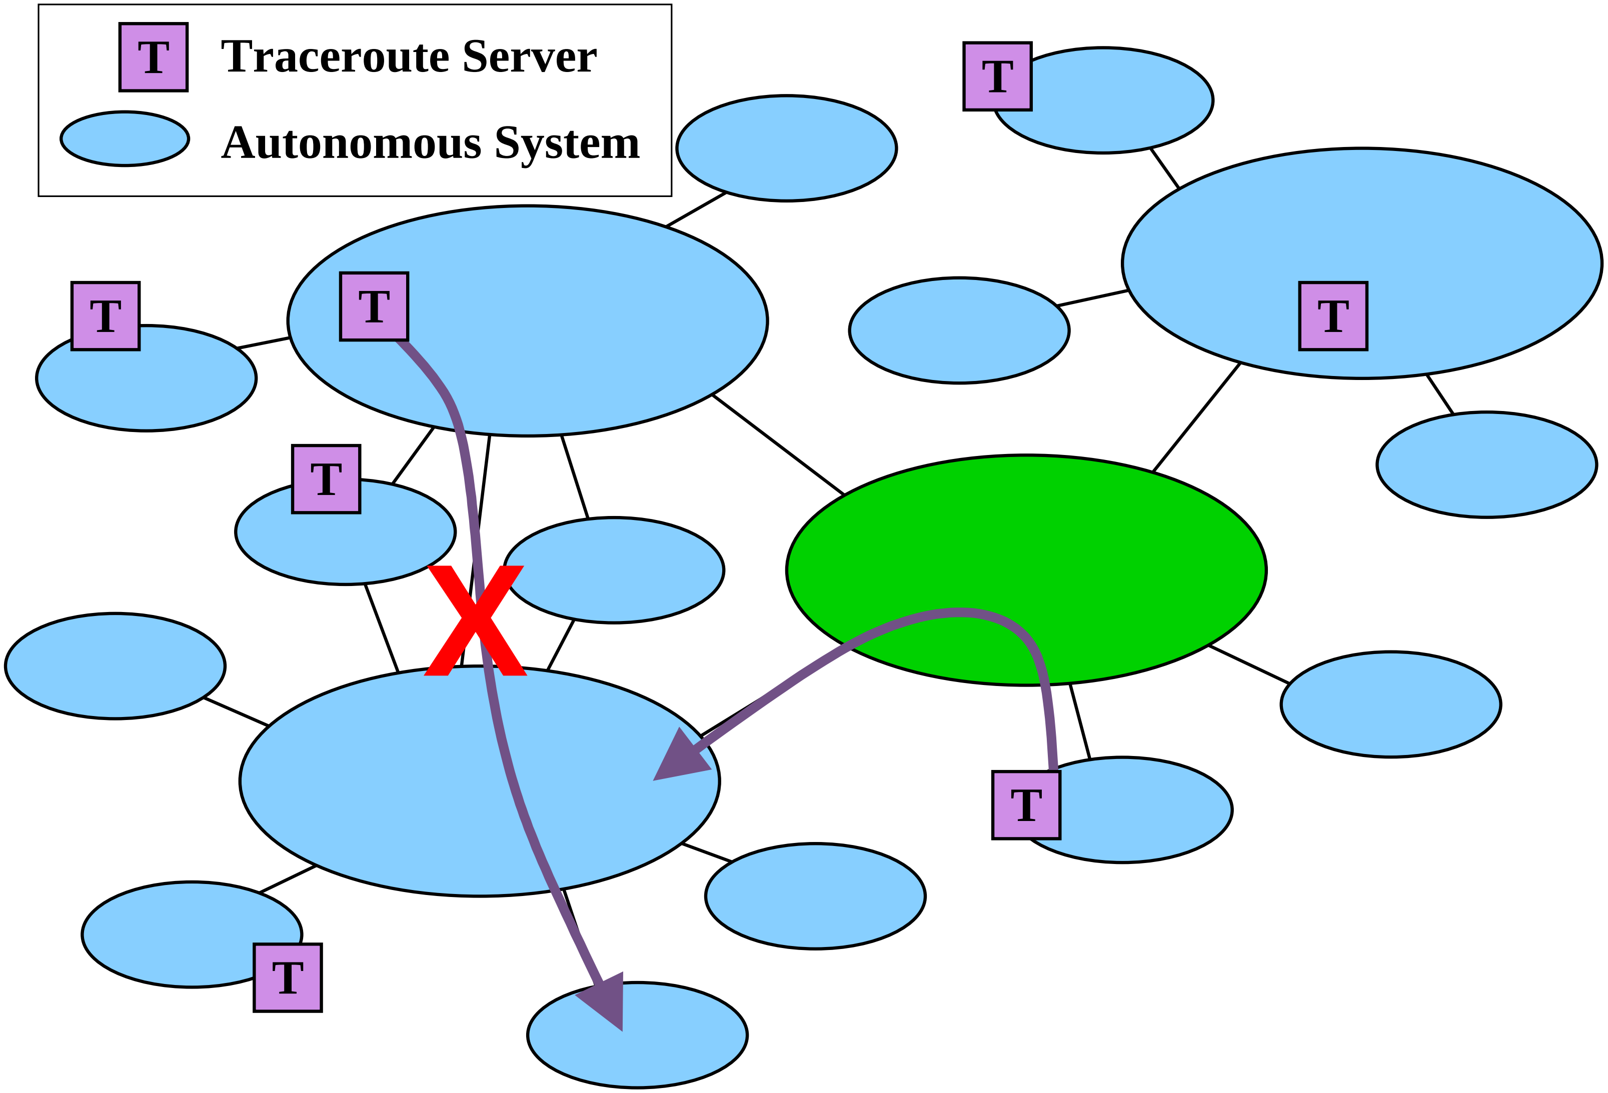

Choosing traceroutes

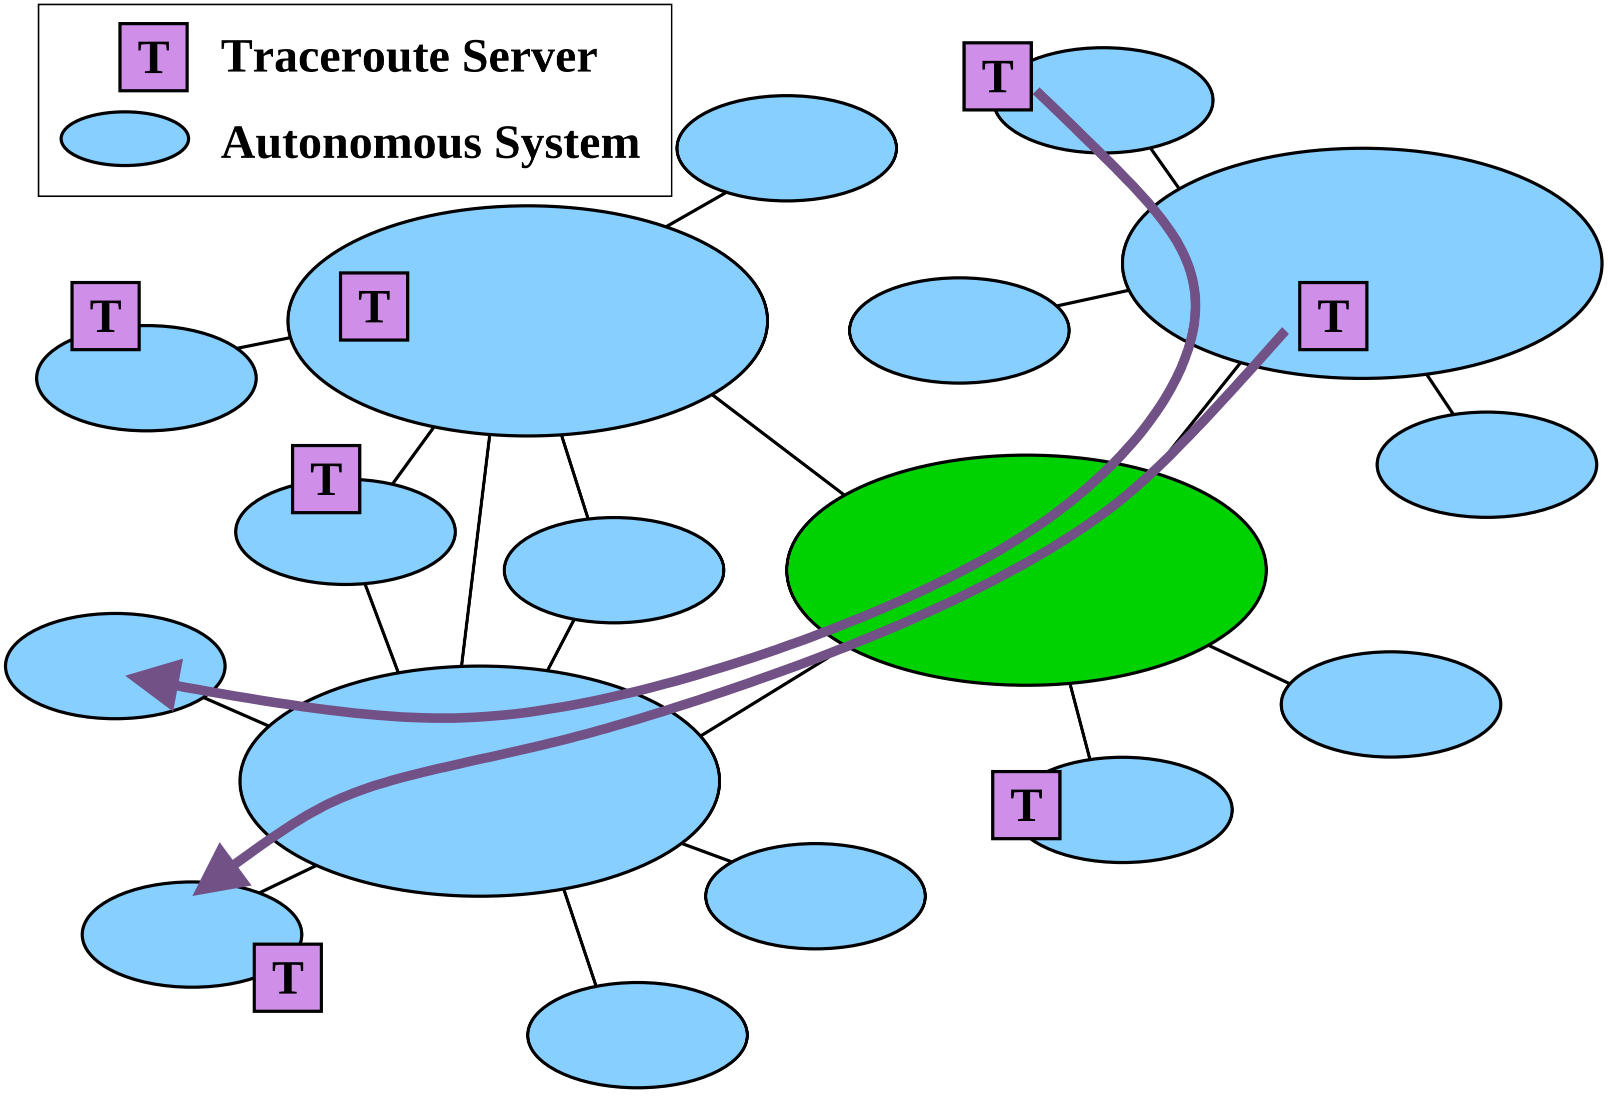

Choosing traceroutes

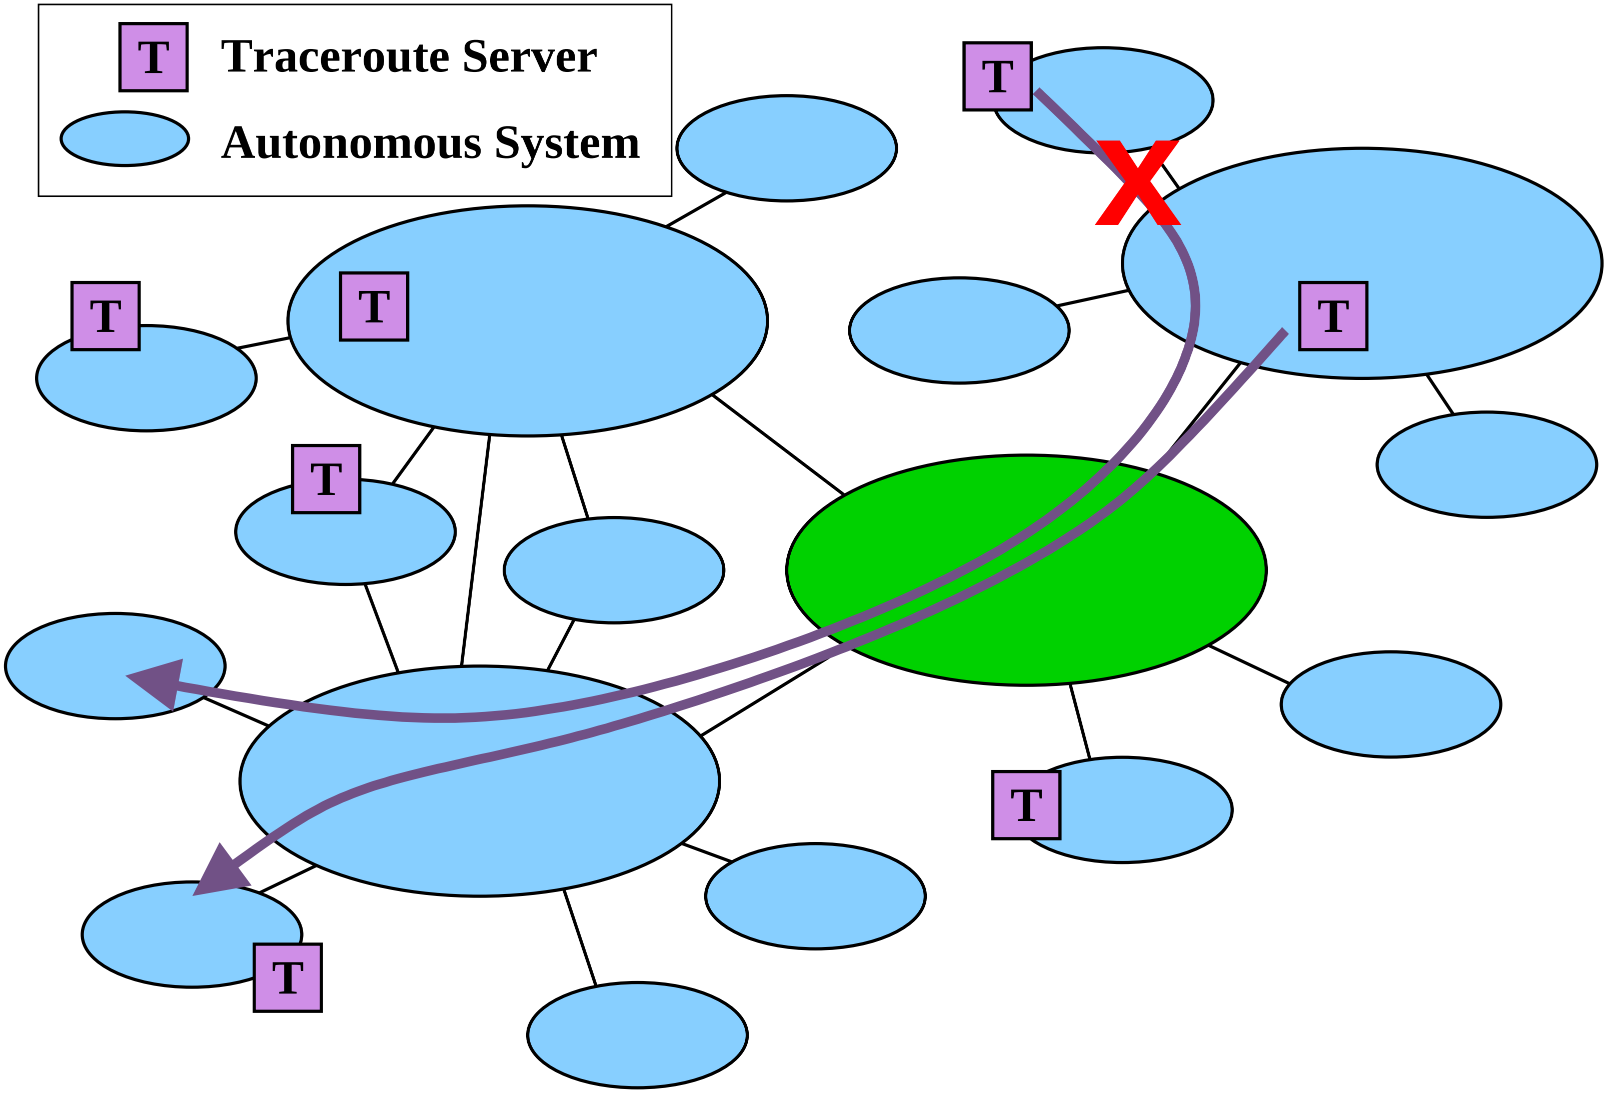

Using BGP to choose traces

RouteViews BGP tables consist of AS-paths to destination prefixes

Destination	AS-Path
	Closer to destination $\longrightarrow$
10.1.0.0/16	**12 20 11 6 7**

Traces likely to traverse AS 6:

- From servers in ASes 12, 20, and 11 to 10.1/16.
- If all known paths to 10.1/16 include AS 6,
  - From anywhere to 10.1/16.
  - From servers in 10.1/16 to anywhere.

Path Reductions

Want to choose unique paths – with new information.

Skip repeated traces of the same path.

Expect the common case:

- Traceroute server has one *ingress point*
- Customer prefix has one *egress point*
- BGP peers have one *early-exit* per ingress.

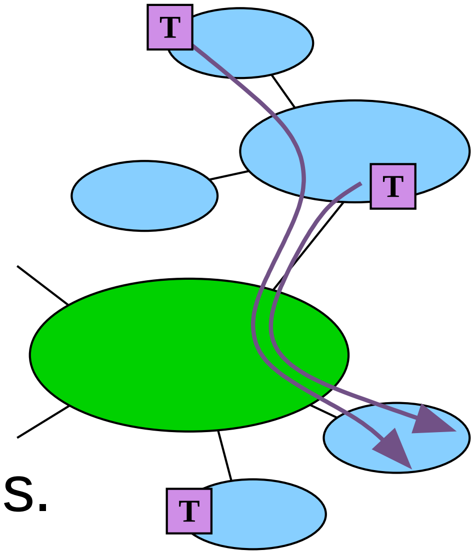

If we're wrong, we might miss some paths.

New servers add paths or share load!

Reduction Effectiveness

- Brute force:
  All servers to all BGP prefixes, disaggregate ISP prefixes.
  90-150 million traceroutes required
- BGP directed probes:
  All traceroutes identifiable from RouteViews.
  0.2-15 million traceroutes required
- Executed after path reduction:
  Traceroutes chosen by Rocketfuel.
  8-300 thousand traceroutes required

Directed probing and path reductions are effective at reducing the number of probes required to map an ISP

Alias resolution problem

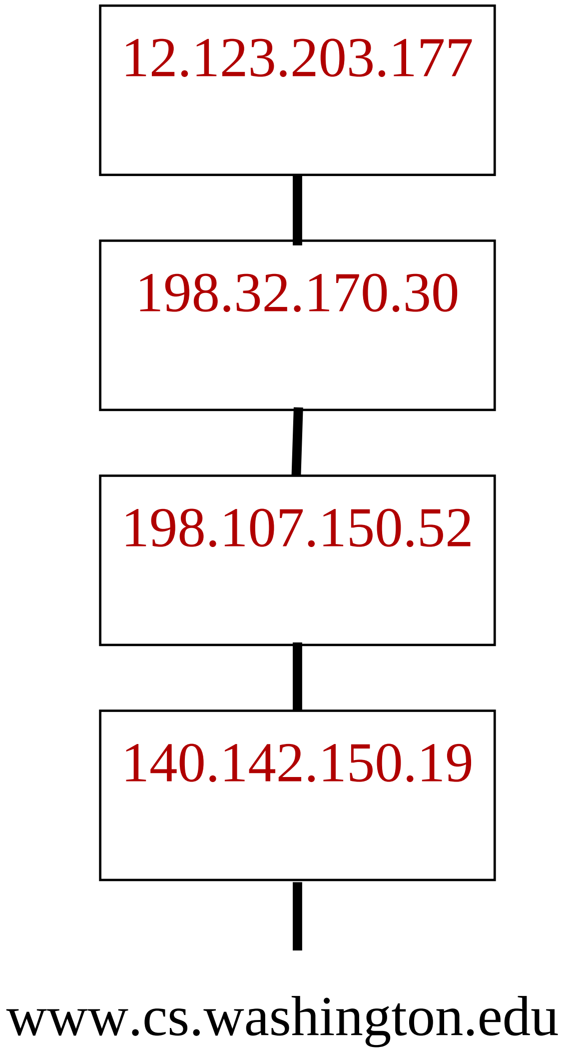

A well-known problem for Internet mapping.

Because traceroute lists IP's, we might think that the path to www.cs and ...

Alias resolution problem

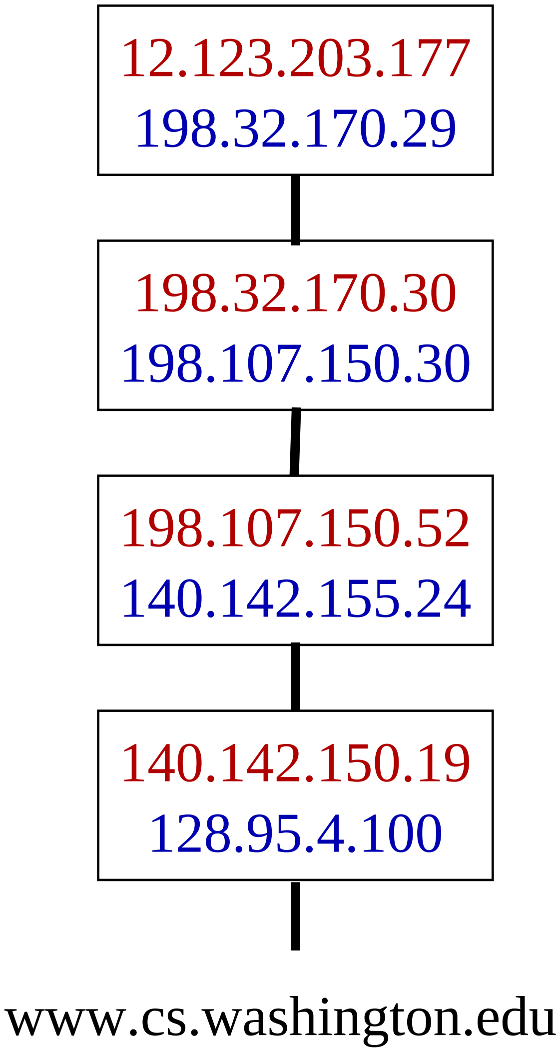

A well-known problem for Internet mapping.

Because traceroute lists IP's, we might think that the path to www.cs and the path from www.cs are different.

Alias resolution finds IP's that belong to the same router.

Our alias resolution solution

Send a packet to each interface to solicit responses.

- Previous work - responses have the same source: Routers often set source address to outgoing interface
- Our approach - responses have nearby IP identifiers: IP ID is commonly set from a counter.

Identify candidate aliases

- Sort by DNS name - find aliases quickly
- Cluster by return TTL - rule out many aliases

Our tool found 2.8 times as many aliases

10% of addresses were unresponsive or unreachable.

Backbone: Telstra

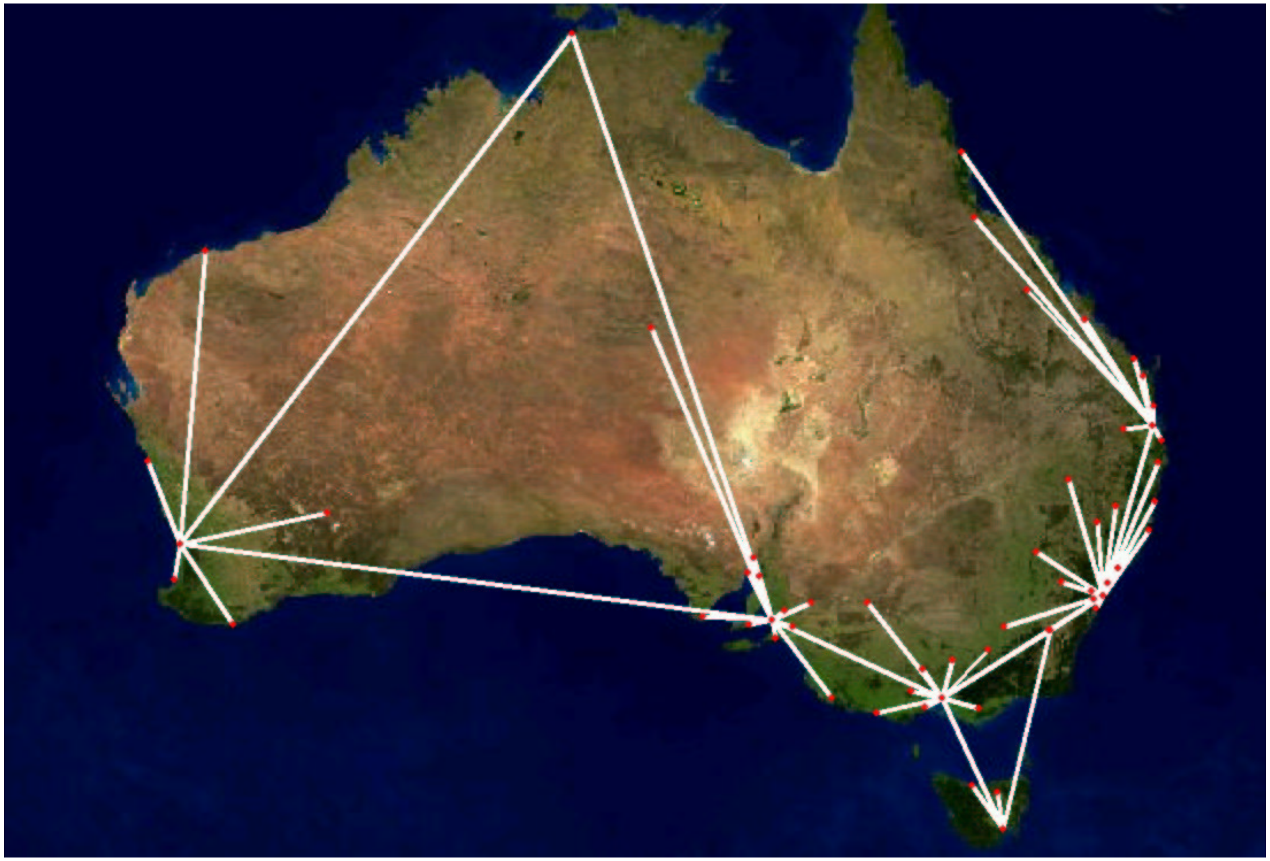

Telstra has hubs in major cities (Sydney, Melbourne, Perth) and spokes elsewhere.

Backbone: EBone

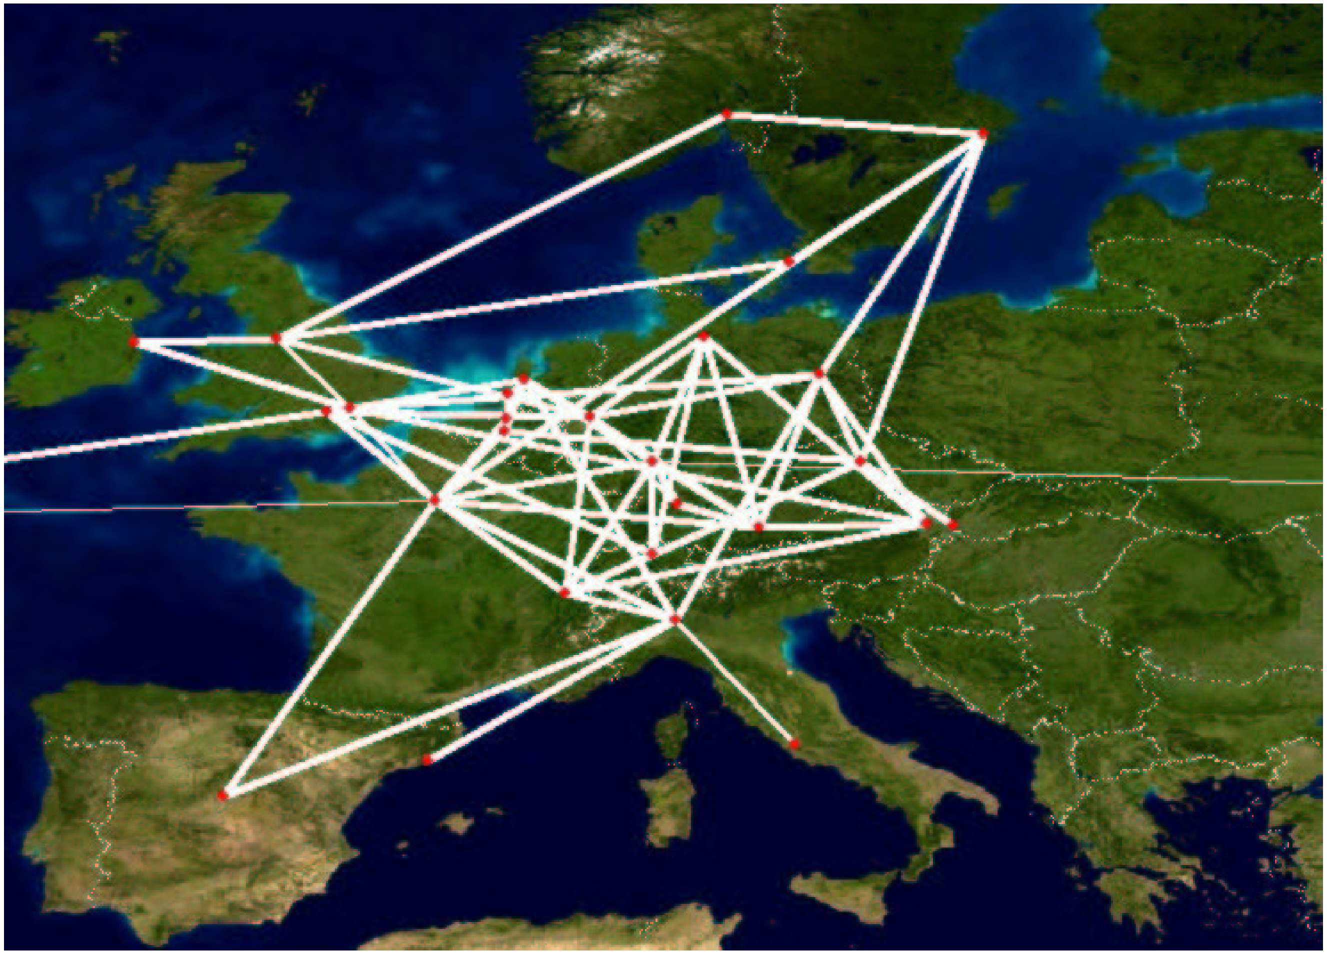

EBone had a highly meshed topology.

Backbone: Level3

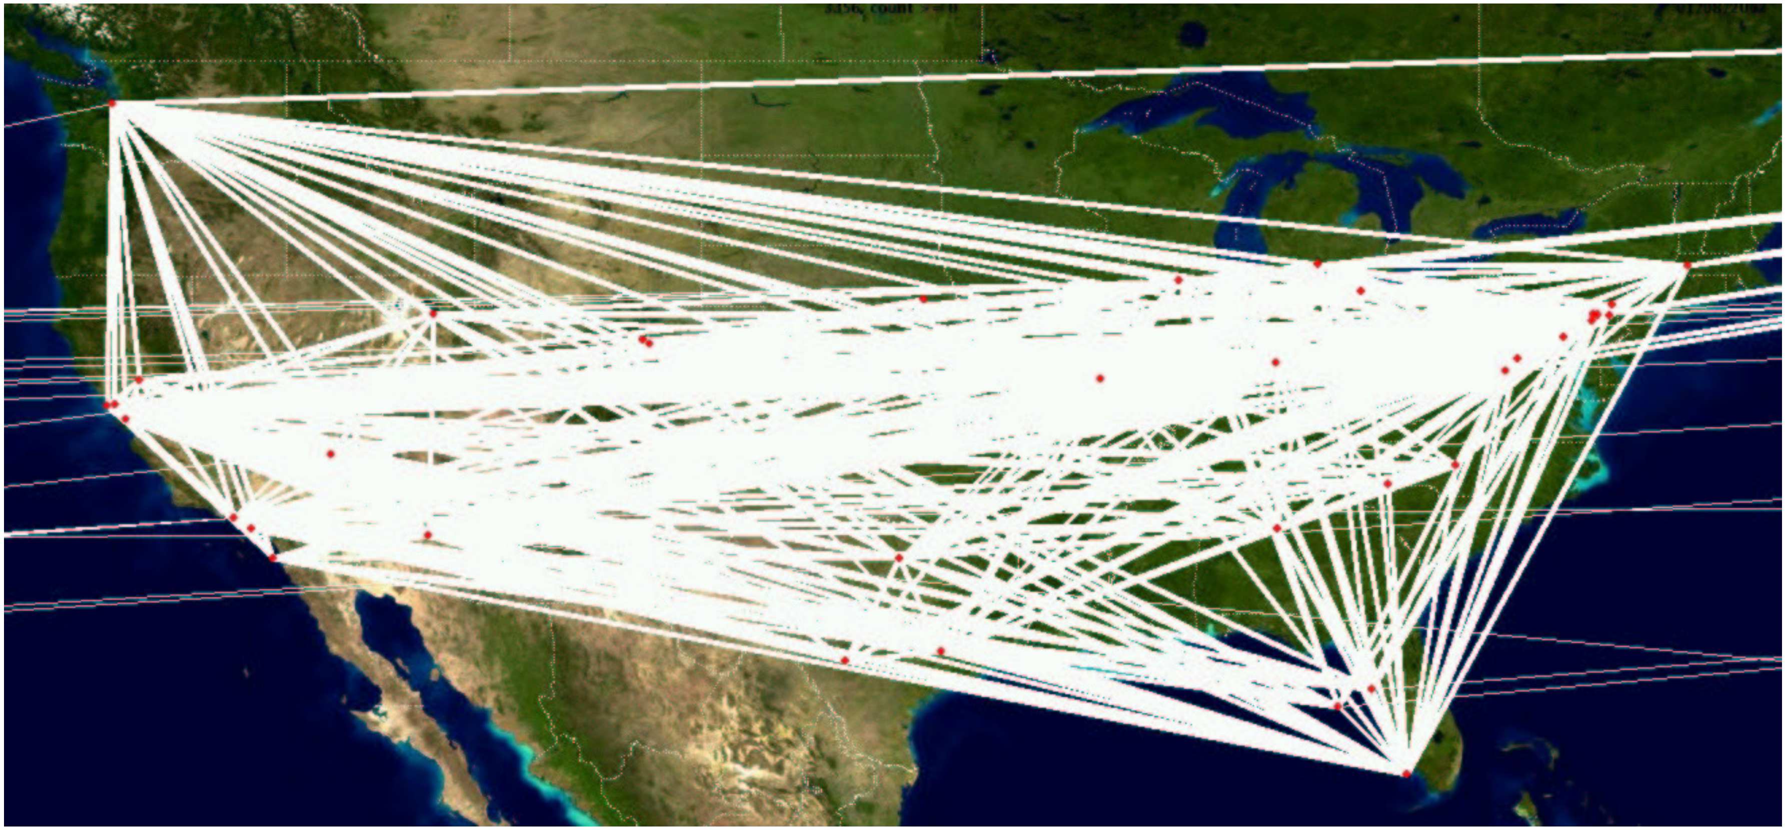

We see the logical topology of Level3's MPLS-based backbone.

Uncovering Structure: POPs

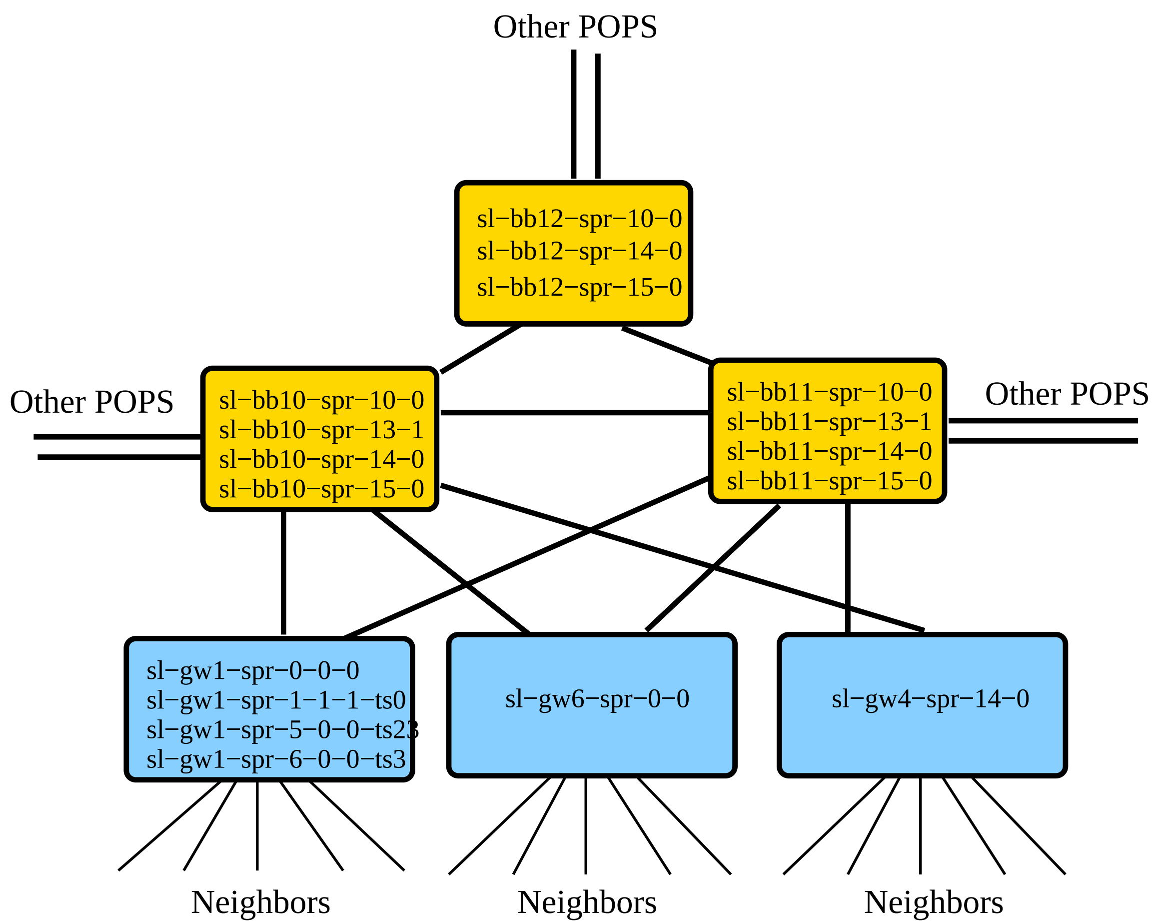

Same structure as “real” Sprint POPs, at a scale we can render.

Map Quality

- Validate with ISPs

  "Good" to "excellent"

- Scan address space to estimate how much better maps might be

  No more than twice as many links.

- Compare to RouteViews

  Found peerings not represented in RouteViews

- Compare to a Skitter dataset

  Skitter maps the Internet, using tens of servers.

  5-10 times as many IPs, links, routers, for a given ISP.

POP Size Distribution

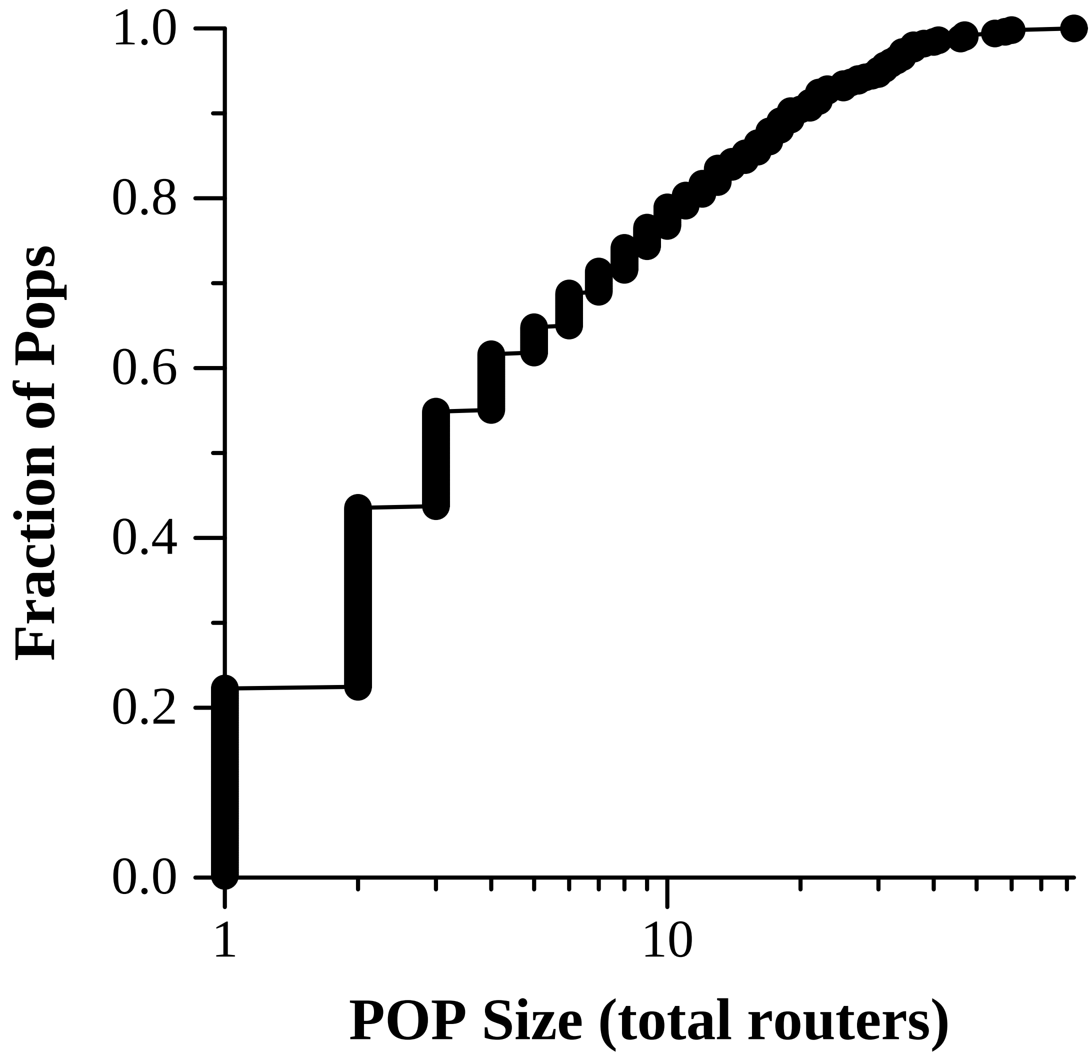

POP Out-Degree Distribution

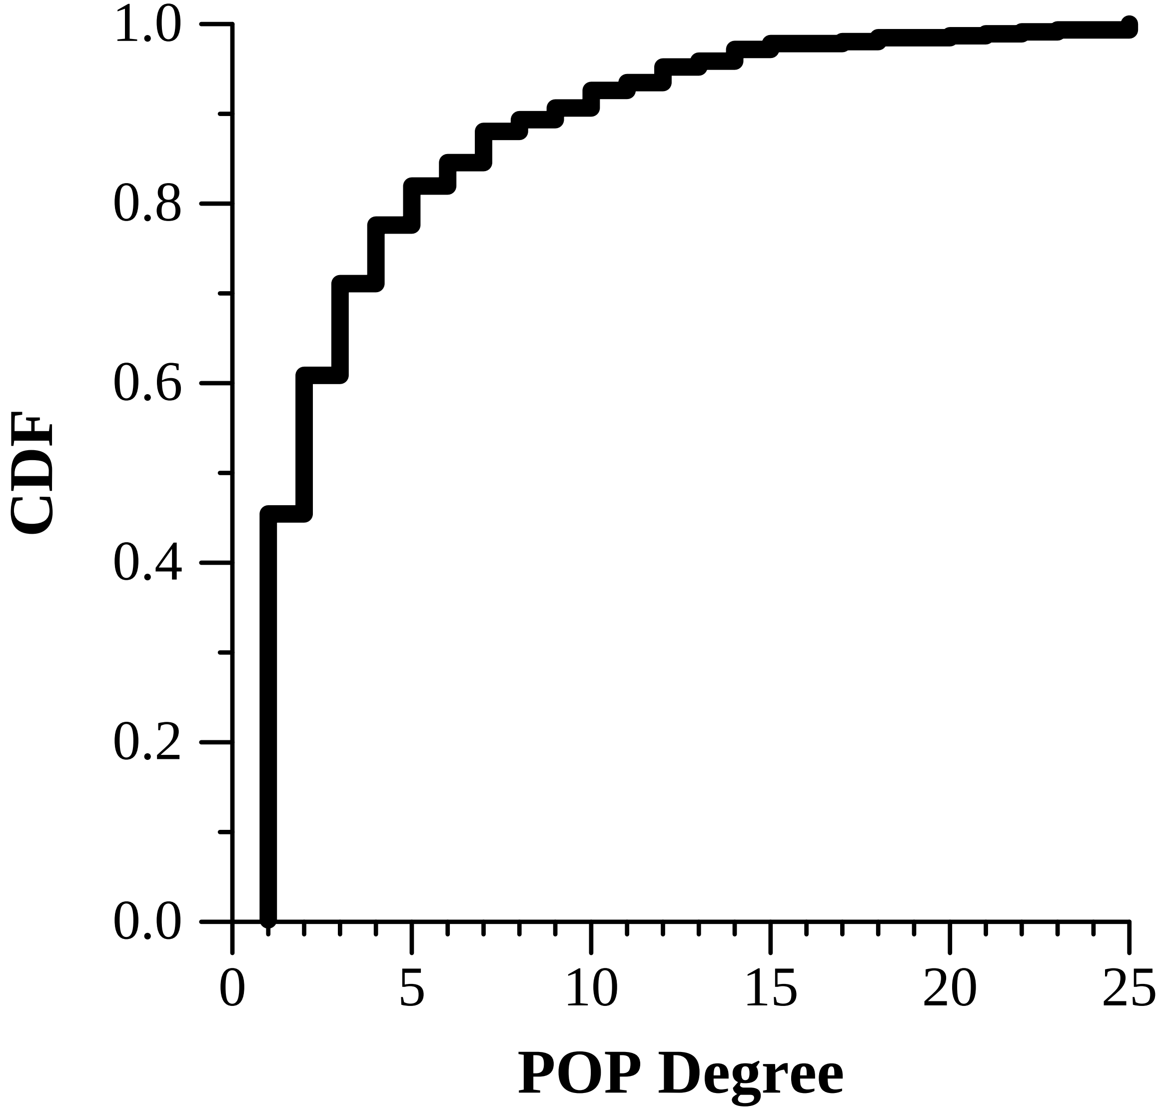

Conclusion

We wanted to understand Internet structure,

We developed techniques to measure router-level maps.

- Trading accuracy for speed where needed
- New alias resolution technique
- Inferred locations and roles from DNS

We learned some things about ISPs, like

- Different backbone styles
- New data about POP size and peerings

See and download the maps at:

http://www.cs.washington.edu/research/networking/rocketfuel/

Out-Degree distributions