Active Measurements: traceroute

Tools: Traceroute

- Exploit TTL (Time to Live) feature of IP
  - When a router receives a packet with TTL=1, packet is discarded and ICMP_time_exceeded returned to sender
- Operational uses:
  - Can use traceroute towards own domain to check reachability
    - list of traceroute servers: http://www.traceroute.org
  - Debug internal topology databases
  - Detect routing loops, partitions, and other anomalies
  - Research, e.g. Internet mapping projects

Traceroute

- In IP, no explicit way to determine route from source to destination
- traceroute: expose intermediate routers

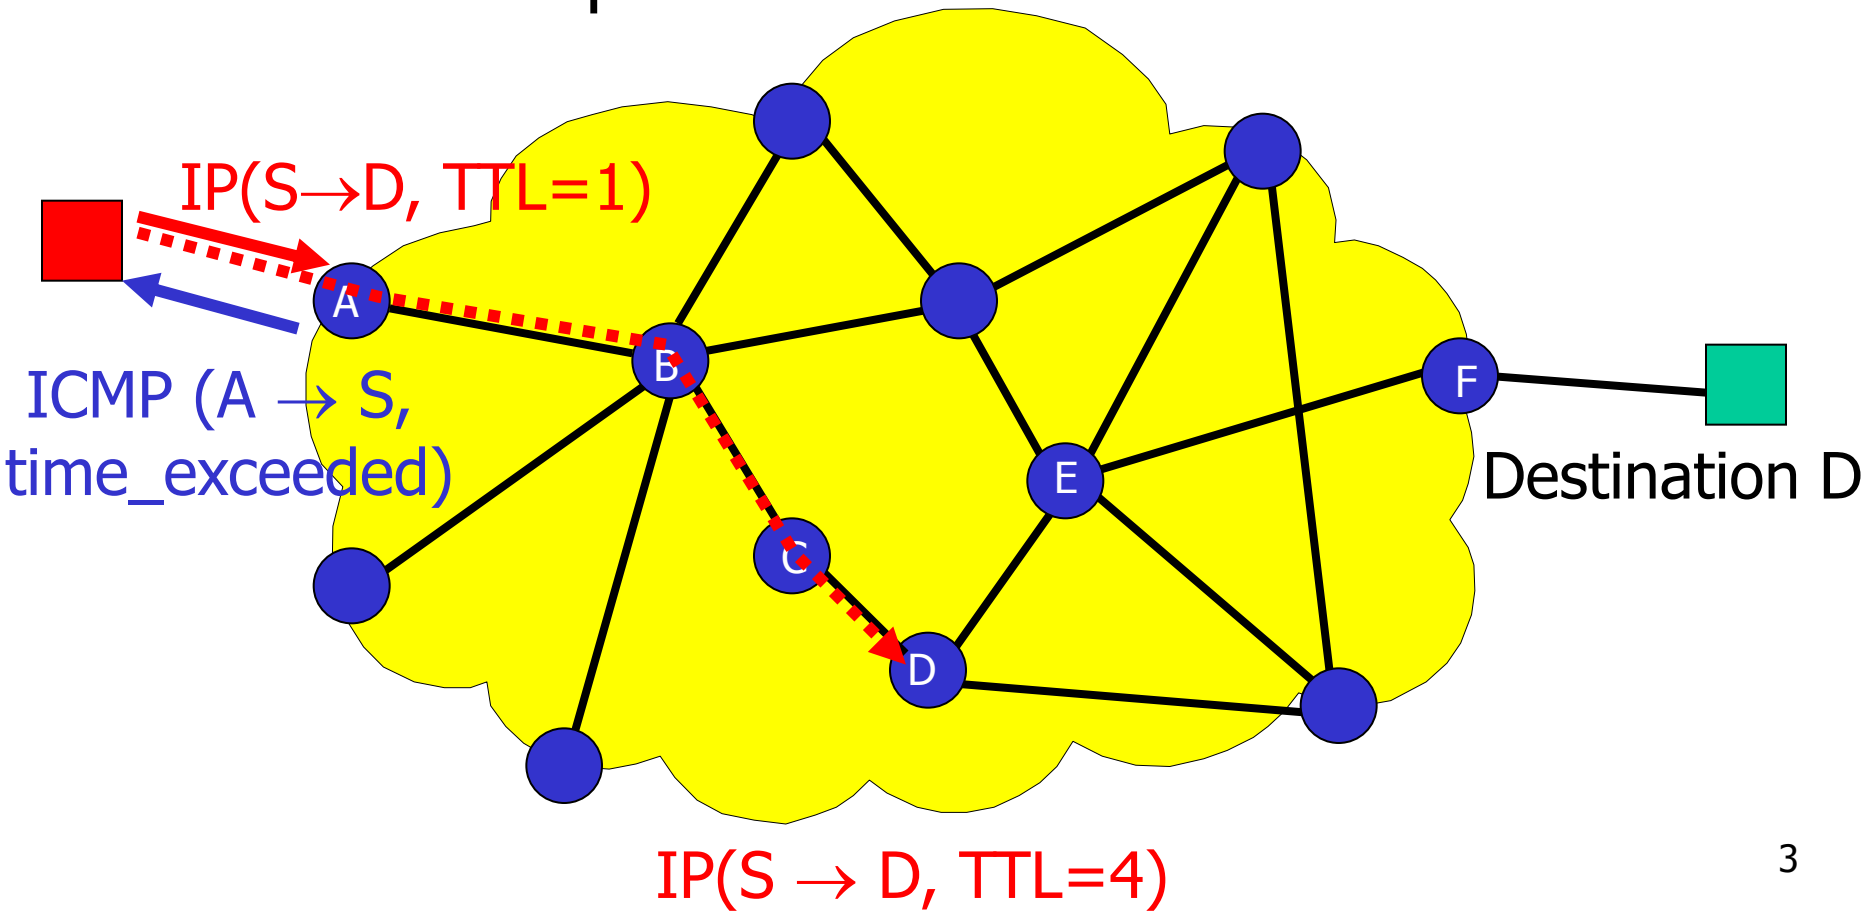

Traceroute: Sample output

```
<chips [ ~ ]>traceroute degas.eecs.berkeley.edu
traceroute to robotics.eecs.berkeley.edu (128.32.239.38), 30 hops max, 40 byte packets
 1  oden (135.207.31.1)  1 ms  1 ms  1 ms
 2  * * *
 3  argus (192.20.225.225)  4 ms  3 ms  4 ms
 4  Serial1-4.GW4.EWR1.ALTER.NET (157.130.0.177)  3 ms  4 ms  4 ms
 5  117.ATM5-0.XR1.EWR1.ALTER.NET (152.63.25.194)  4 ms  4 ms  5 ms
 6  193.at-2-0-0.XR1.NYC9.ALTER.NET (152.63.17.226)  4 ms (ttl=249!)  6 ms (ttl=249!)  4 ms (ttl=249!)
 7  0.so-2-1-0.XL1.NYC9.ALTER.NET (152.63.23.137)  4 ms  4 ms  4 ms
 8  POS6-0.BR3.NYC9.ALTER.NET (152.63.24.97)  6 ms  6 ms  4 ms
 9  acr2-atm3-0-0-0.NewYorknyr.cw.net (206.24.193.245)  4 ms (ttl=246!)  7 ms (ttl=246!)  5 ms (ttl=246!)
10  acr1-loopback.SanFranciscosfd.cw.net (206.24.210.61)  77 ms (ttl=245!)  74 ms (ttl=245!)  96 ms (ttl=245!)
11  cenic.SanFranciscosfd.cw.net (206.24.211.134)  75 ms (ttl=244!)  74 ms (ttl=244!)  75 ms (ttl=244!)
12  BERK-7507--BERK.POS.calren2.net (198.32.249.69)  72 ms (ttl=238!)  72 ms (ttl=238!)  72 ms (ttl=238!)
13  pos1-0.inr-000-eva.Berkeley.EDU (128.32.0.89)  73 ms (ttl=237!)  72 ms (ttl=237!)  72 ms (ttl=237!)
14  vlan199.inr-202-doecev.Berkeley.EDU (128.32.0.203)  72 ms (ttl=236!)  73 ms (ttl=236!)  72 ms (ttl=236!)
15  * 128.32.255.126 (128.32.255.126)  72 ms (ttl=235!)  74 ms (ttl=235!)
16  GE.cory-gw.EECS.Berkeley.EDU (169.229.1.46)  73 ms (ttl=9!)  74 ms (ttl=9!)  72 ms (ttl=9!)
17  robotics.EECS.Berkeley.EDU (128.32.239.38)  73 ms (ttl=233!)  73 ms (ttl=233!)  73 ms (ttl=233!)
```

ICMP disabled

TTL=249 is unexpected (should be initial_ICMP_TTL-(hop#-1)= 255-(6-1)=250)

RTT of three probes per hop

Traceroute: Limitations

- No guarantee that every packet will follow same path
  - Inferred path might be "mix" of paths followed by probe packets
- No guarantee that paths are symmetric
  - Unidirectional link weights, hot-potato routing
  - No way to answer question: on what route would a packet reach me?
- Reports interfaces, not routers
  - May not be able to identify two different interfaces on the same router
- Topological sampling
  - Limitation in sources and destinations
  - Not all links can be seen, e.g. backup

Load balancing

- Per-packet load balancers may mislead traceroute
- Hash-based balancing will show non-existing paths to traceroute

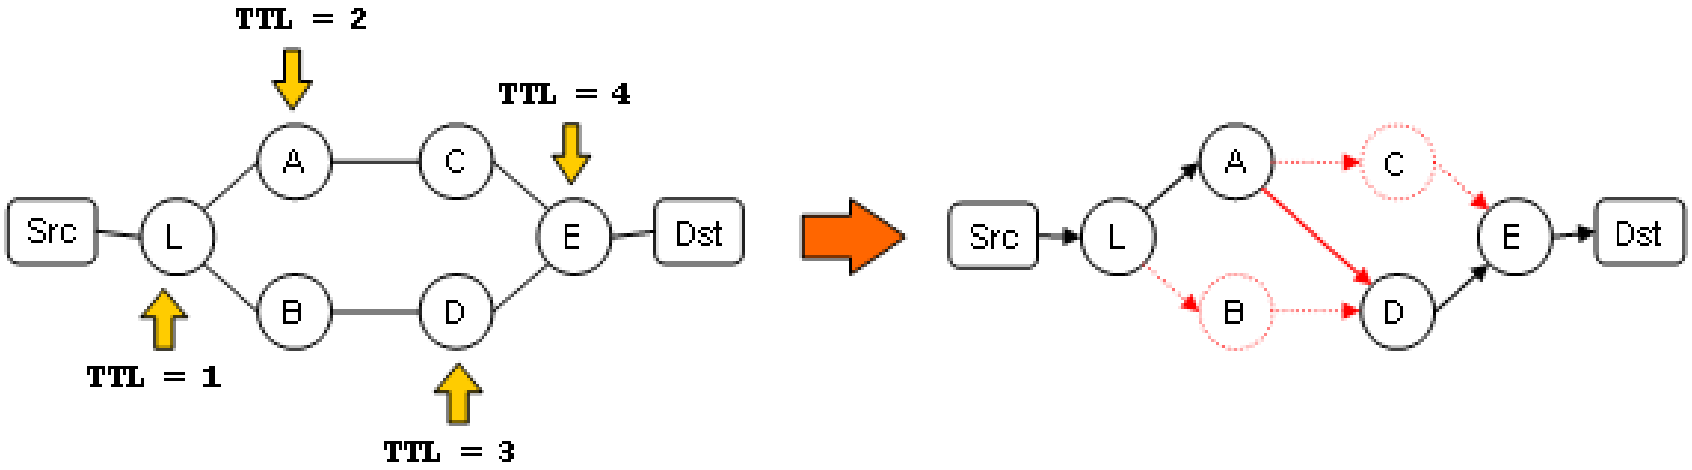

Load balancing (2)

- Flow-level load balancers rely on grey fields to identify a flow
- Packet-level load balancers
- Paris traceroute tries to identify balanced paths by changing header values not used for balancing

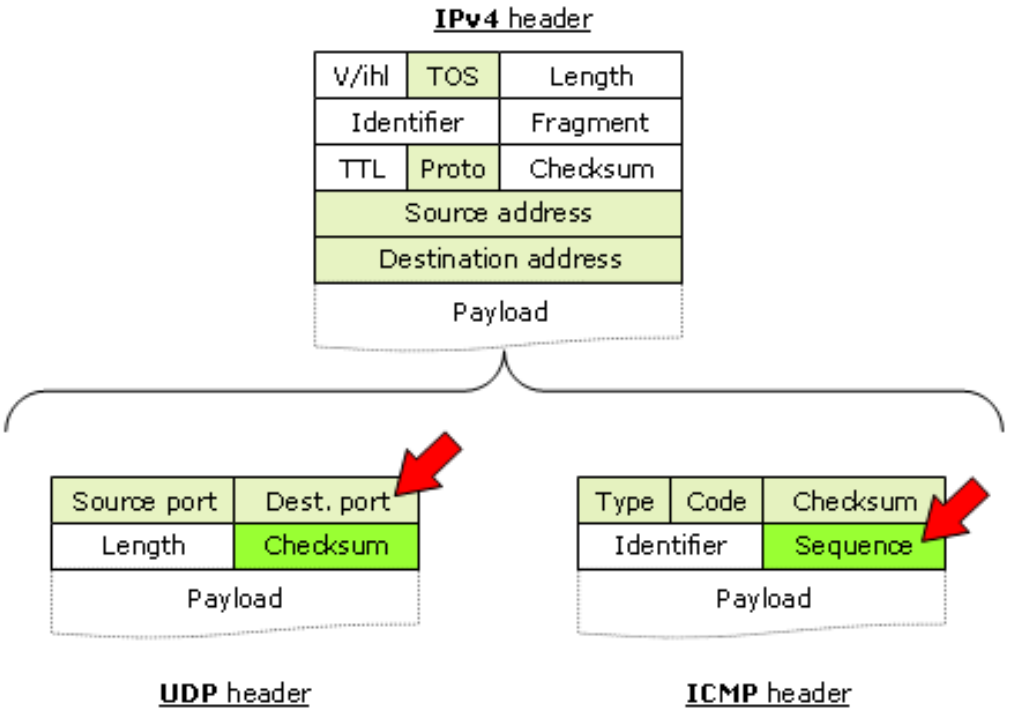

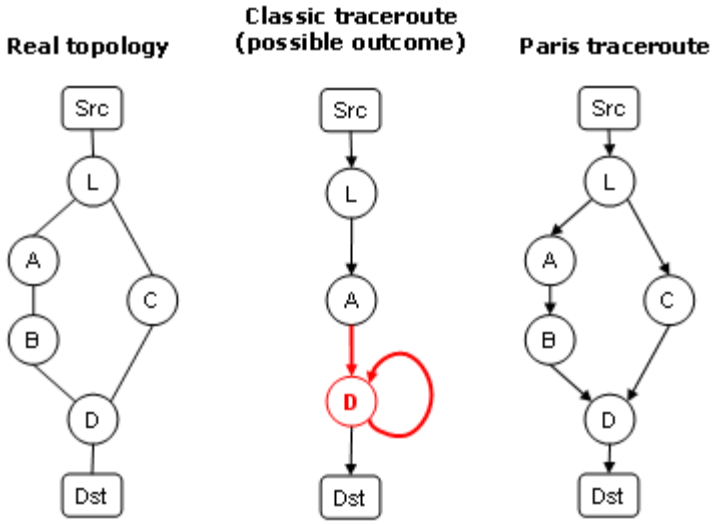

Routing path asymmetry

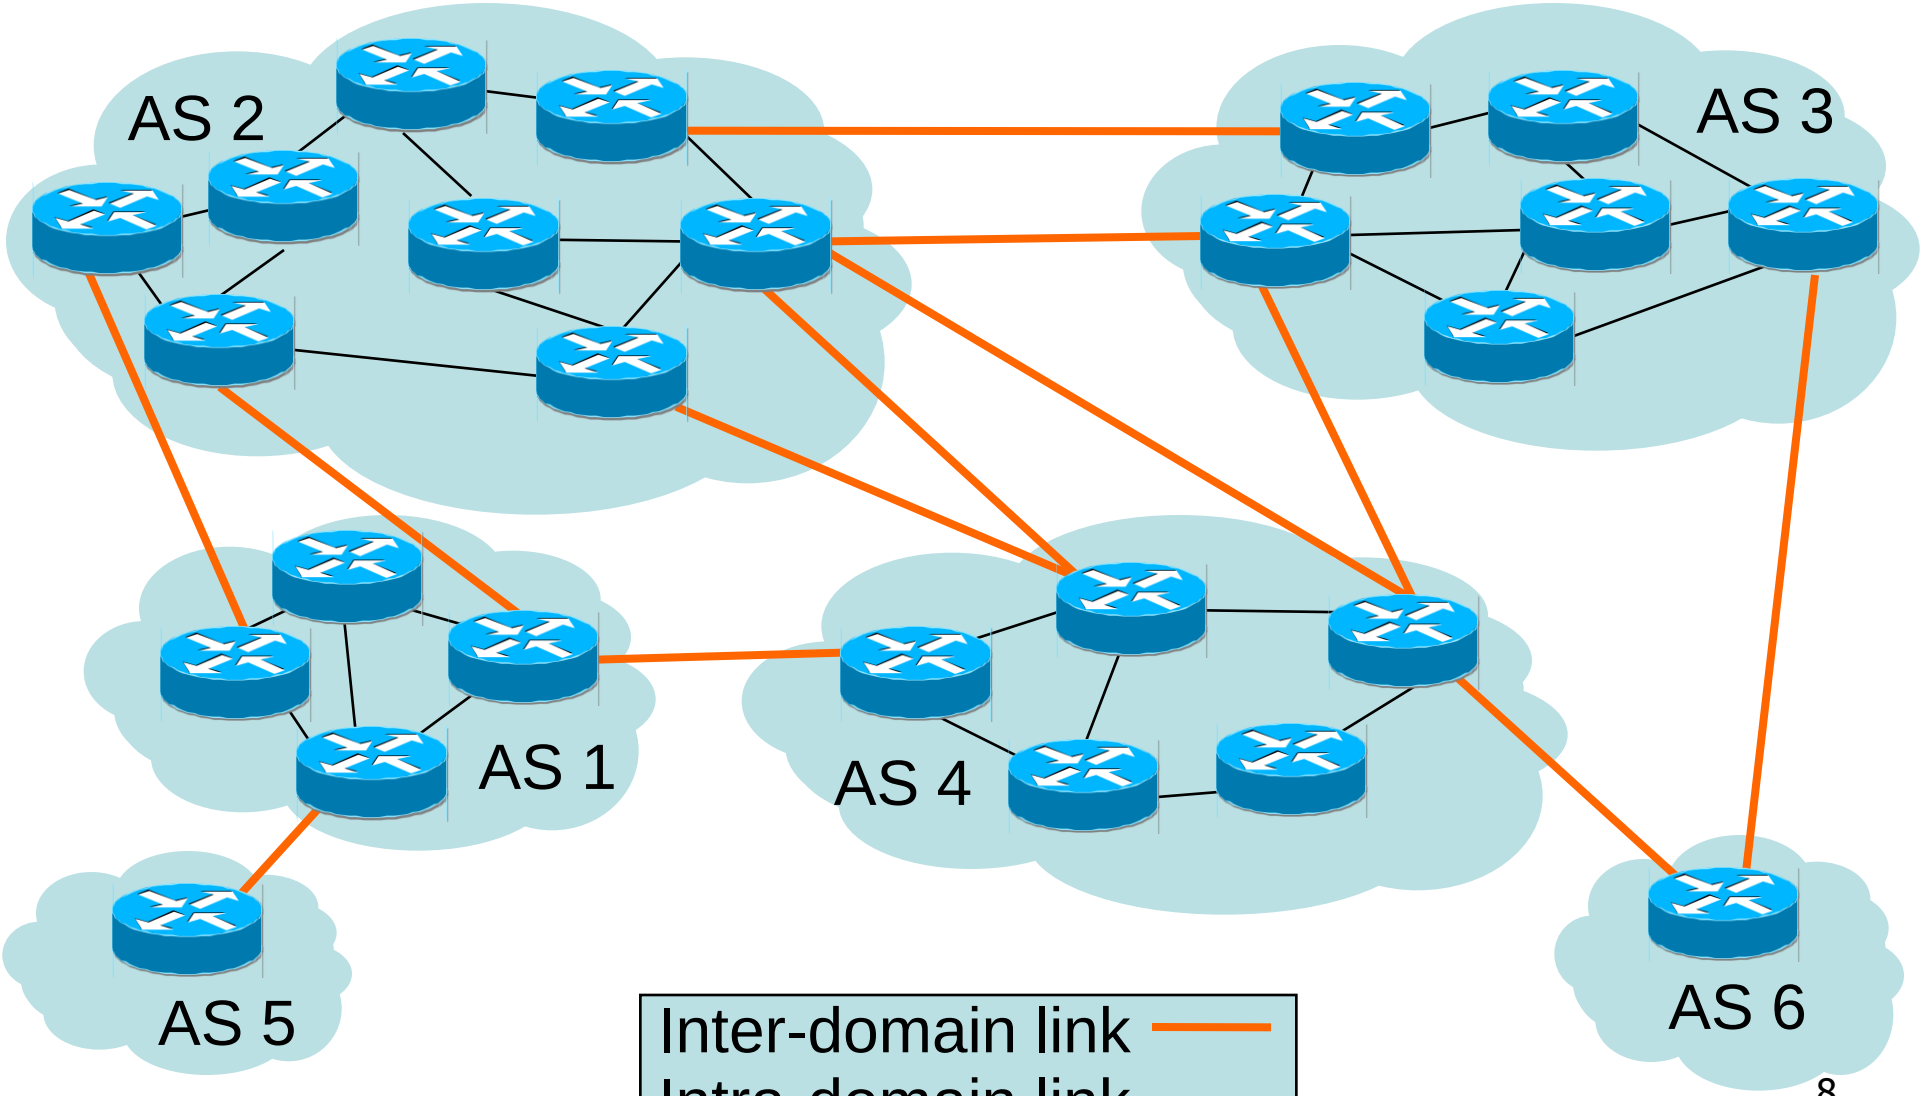

Routing path asymmetry (2)

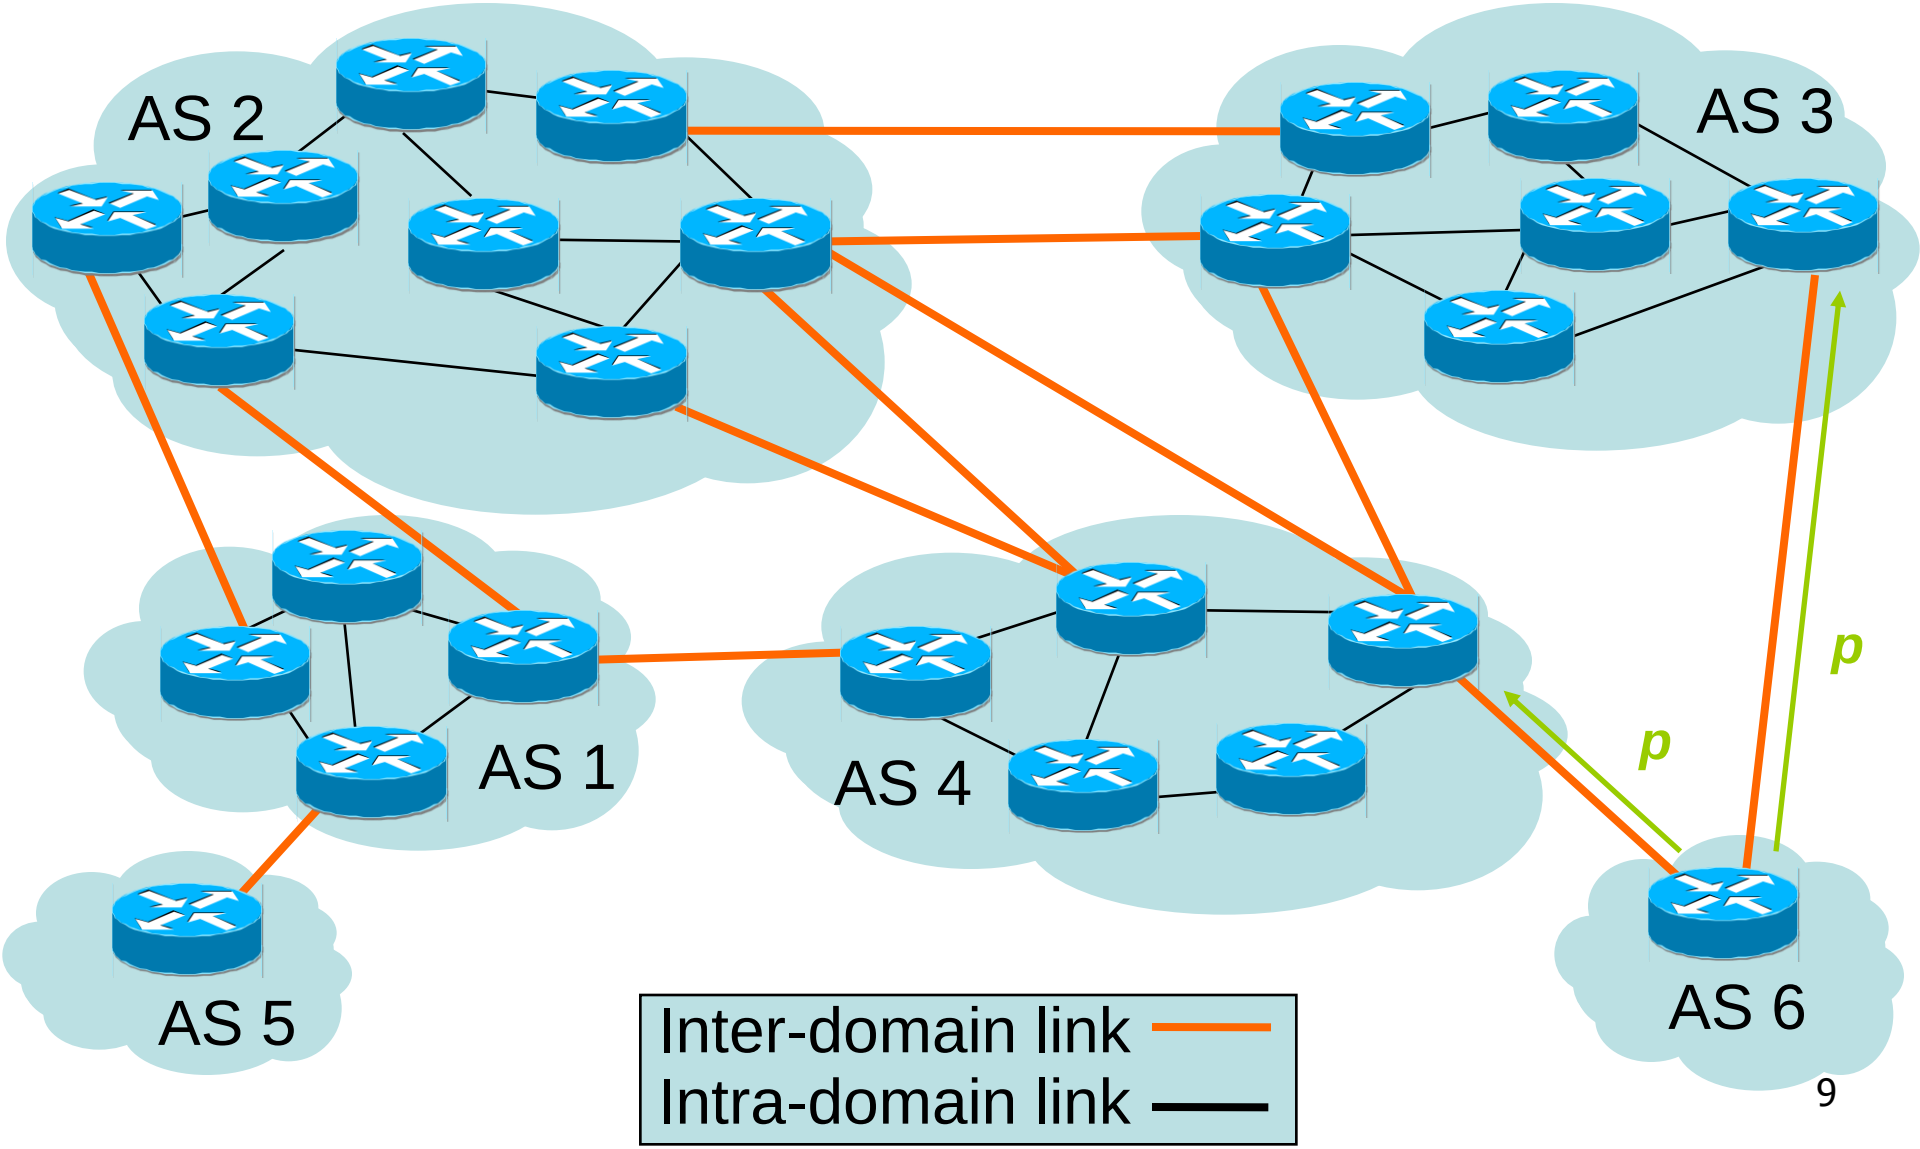

Routing path asymmetry (3)

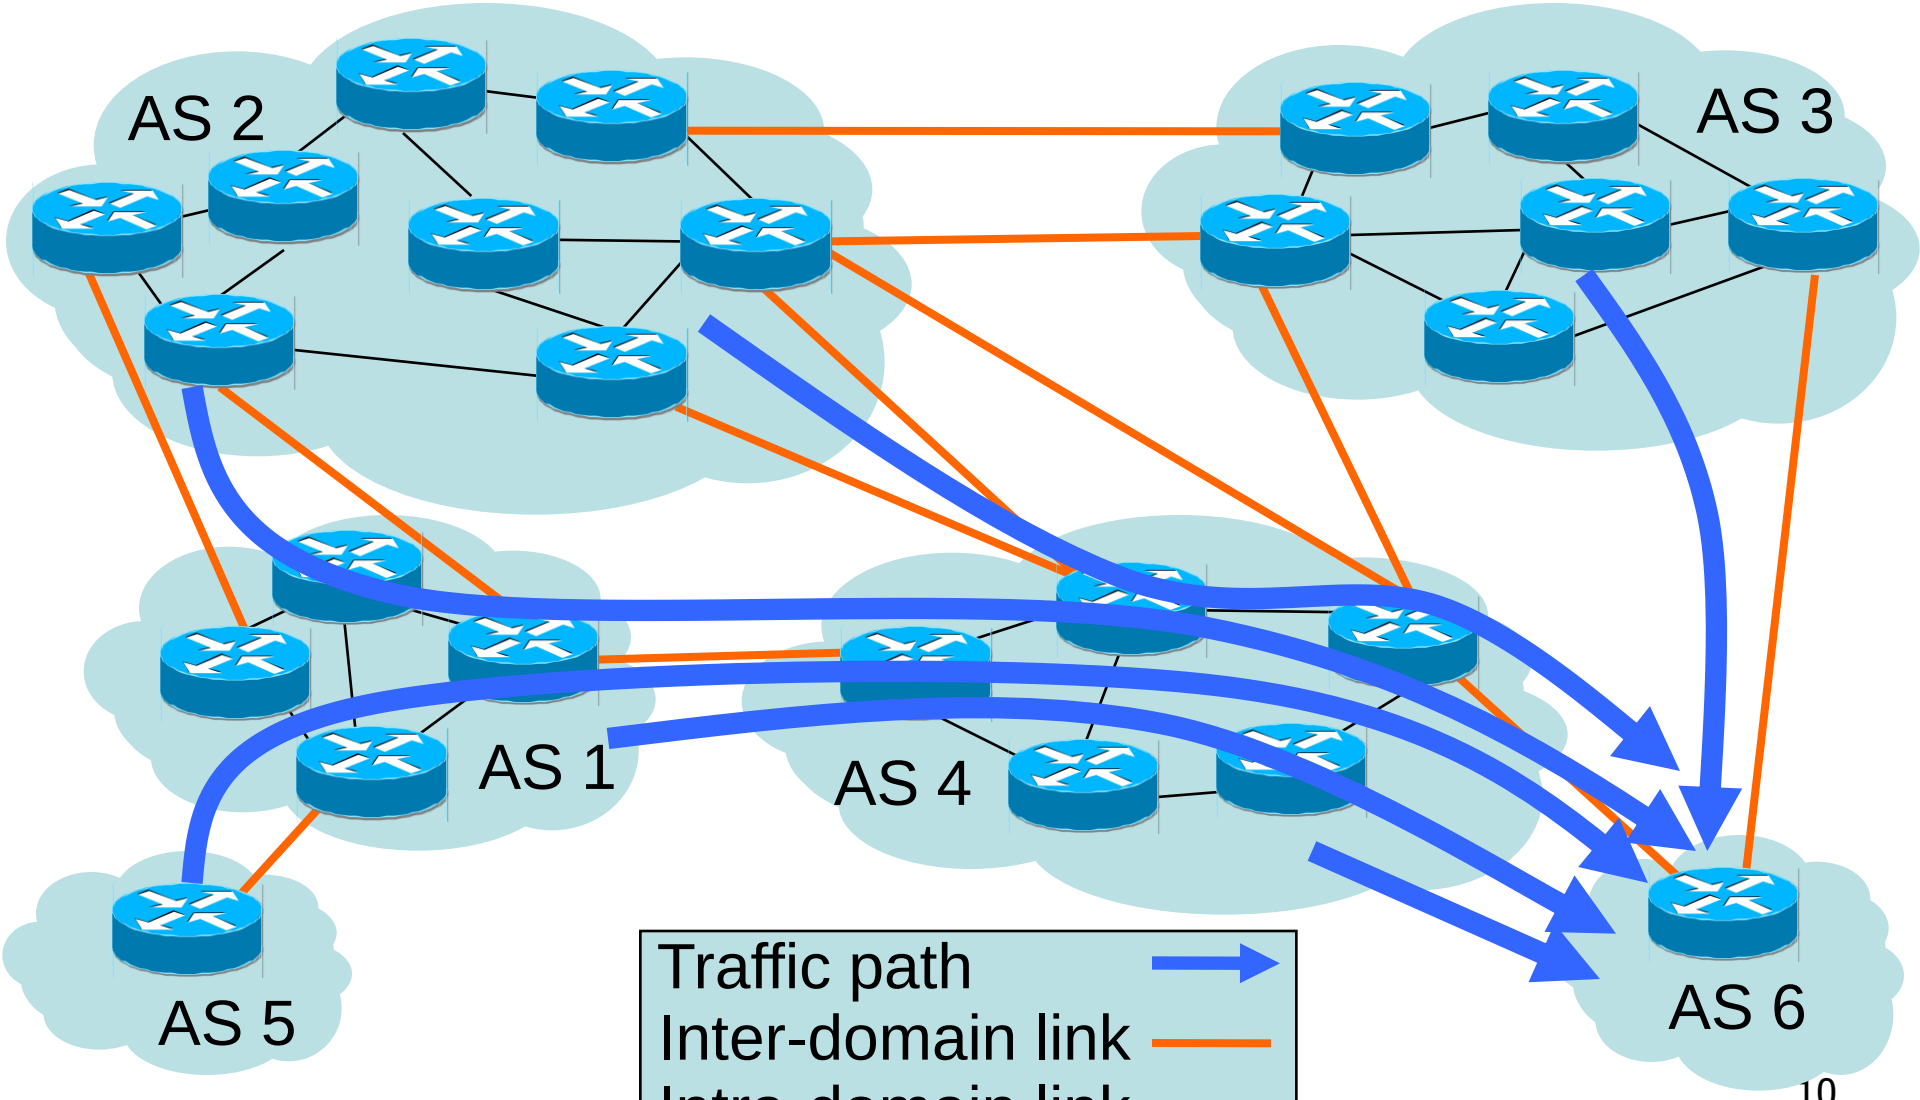

Routing path asymmetry (4)

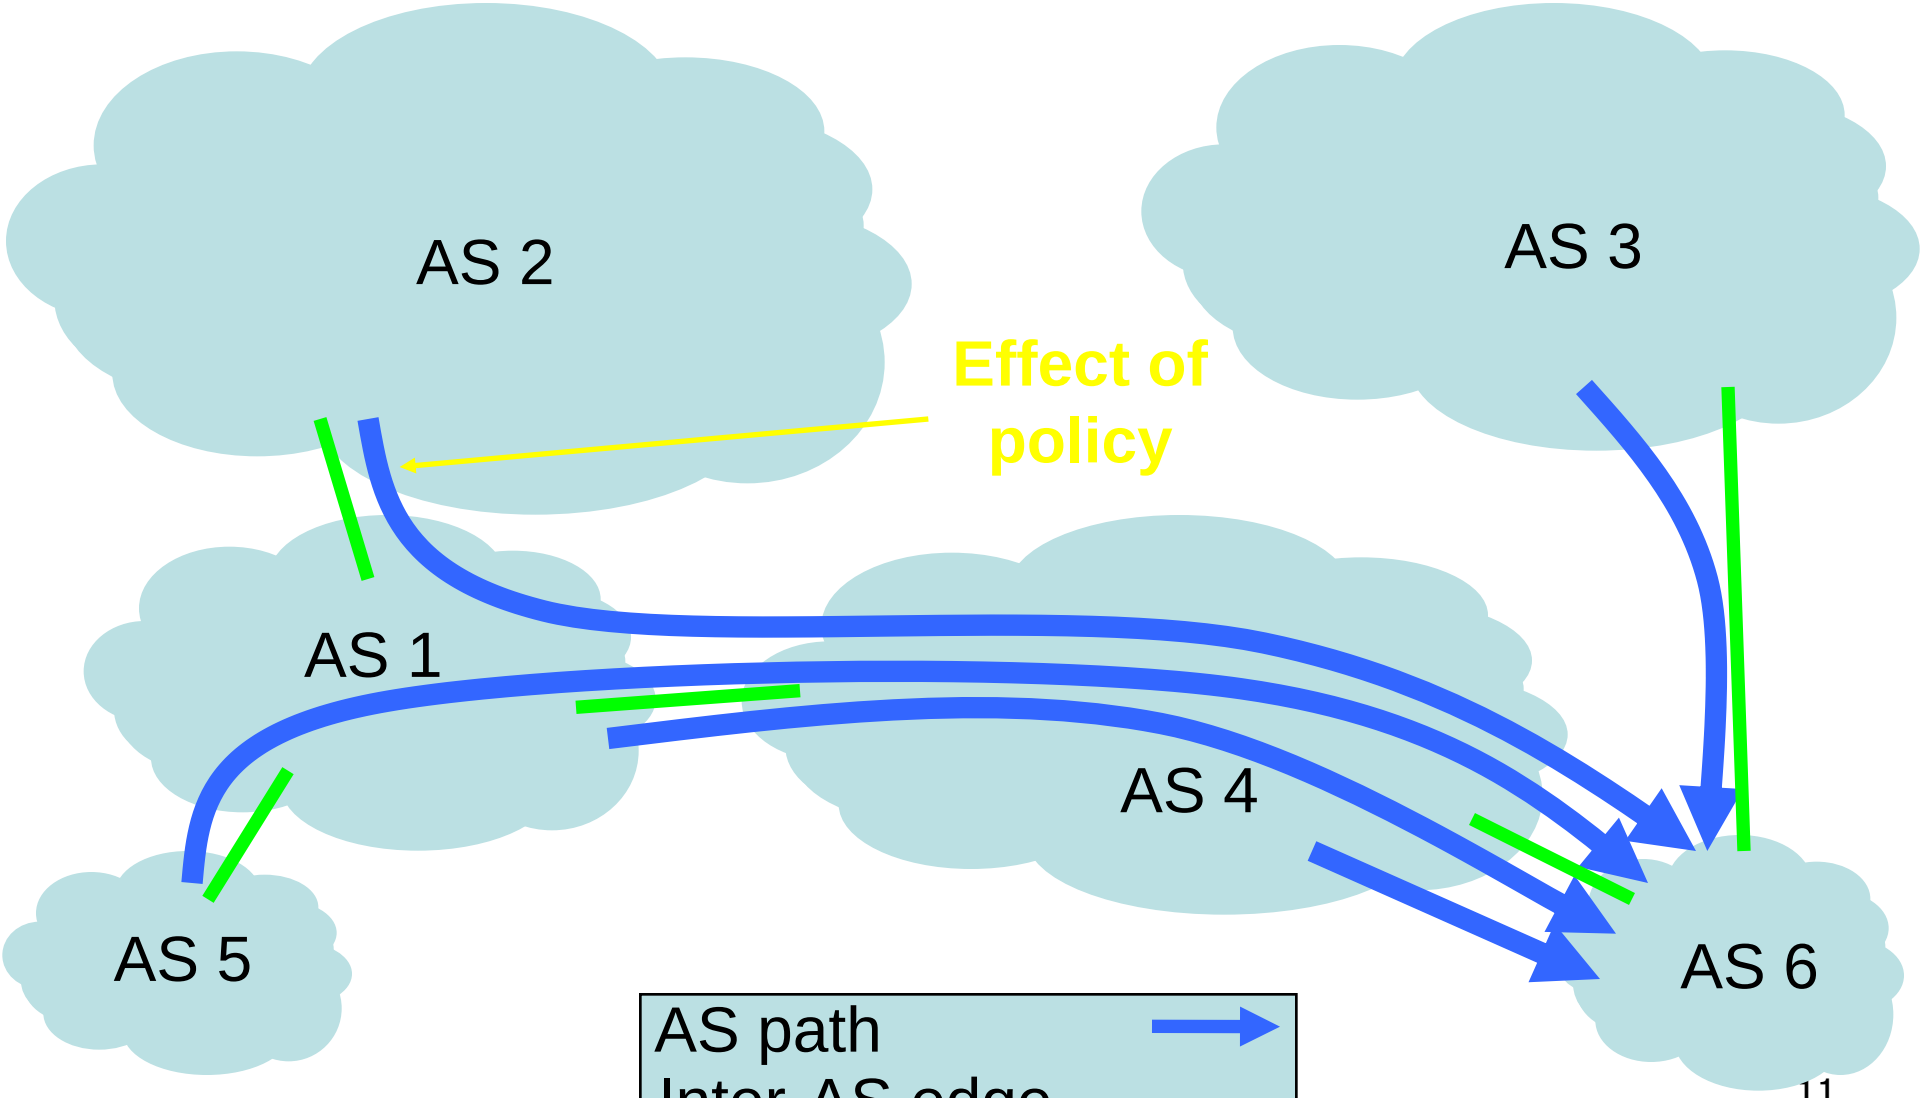

Path asymmetry

- Reachability does not require symmetric paths
- Many factors explain asymmetry
  - Routing policies
    - Routing defined on a per destination prefix
    - Policies may rely on different granularities
  - Static/default routes
    - Many ISPs rely on static routes to compensate for routing failures
  - Intradomain routing (hot-potato)
    - Different entry point may mean different exit point in the network
  - Multiple links/load balancing
    - For different flows that share the same destination prefix, multiple paths may be used

IP aliasing

- Routers have more than one network interface
- How to distinguish multiple interfaces from the same router?
  - Fingerprinting: send probe packets and infer responses come from the same router
    - iffinder: send TCP/UDP towards unreachable port and compare source IP with ICMP reply
    - IPID counter: increase rate in IPID as a host signature
    - DNS
  - Analytics: inference from topology graph
    - Path alignment
    - Record-route option

Topological sampling

- How much probing to sample a given network?
- Bias in traceroute measurements:
  - Sources:
    - Most paths are shared close to sources
  - Destinations:
    - End of paths suffer from NAT/firewalls
    - Finding host that answer to probes is hard
  - Graph properties:
    - Power-law node degree
    - Local connectivity

Topological sampling (2)

- Modeling Internet topology:
  - Random graph or power-law?

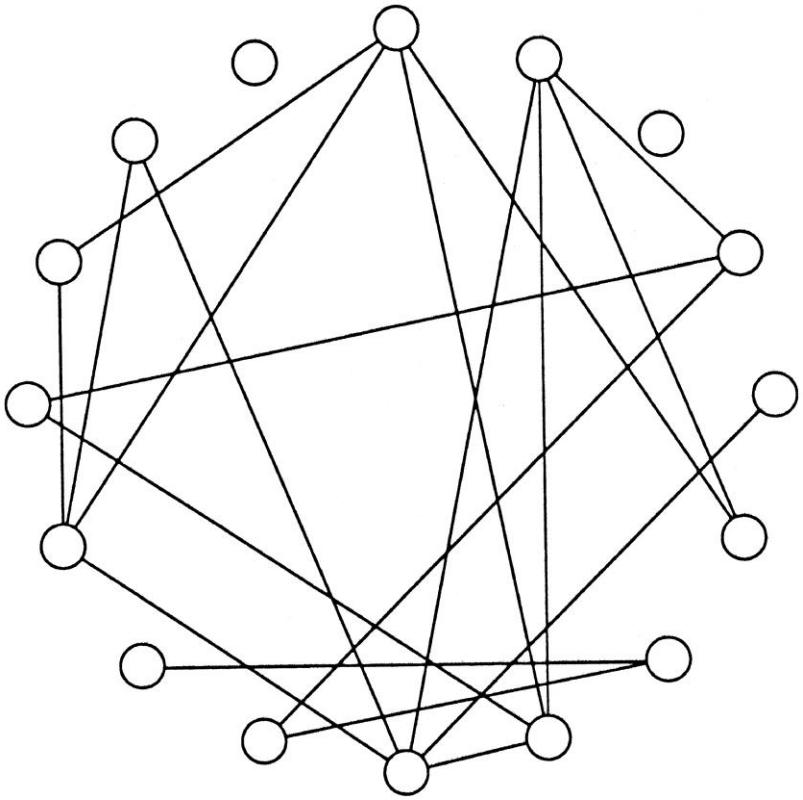

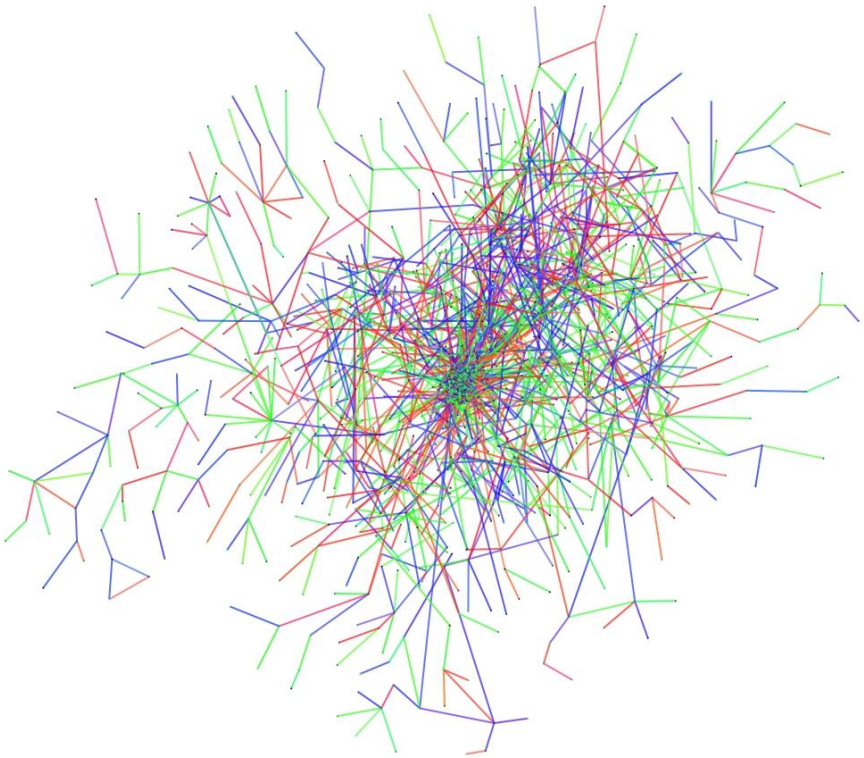

Topological sampling (3)

- Local connectivity properties
  - Clustering coefficient:
    - Number of triangles: directly connected triples/ connected triples
    - How close is the local connectivity from a clique?
  - Assortativity
    - Degree correlation between nodes of a given degree:
      - Nodes of large degrees tend to connect to similar degree nodes, e.g. social networks (assortative network)
      - Small degree nodes and large degree nodes tend to connect to each other, e.g. Internet (disassortative network)

Traceroute: Discussion

- Impact of routing dynamics
  - Traceroute paths need to be correlated with routing data
  - Routing is about dynamics: Internet is constantly changing
- Interface-level map, not router-level
  - Inferring topological properties from traceroute data requires care
  - Ground truth, e.g. MRINFO
- Topological sampling
  - Bias in core and edge sampling
- Reachability
  - Data plane and control plane do see different things
  - Packets more happy than what routing shows