Some Examples of Network Measurements

- Example 1
  - Data: Traceroute measurements
  - Objective: Inferring Internet topology at the router-level
- Example 2
  - Data: Traceroute measurements
  - Objective: Inferring Internet topology at the level of Autonomous Systems (ASes)
- Example 3
  - Data: BGP measurements
  - Objective: Inferring Internet topology at the level of Autonomous Systems (ASes)

Measurement tool: traceroute

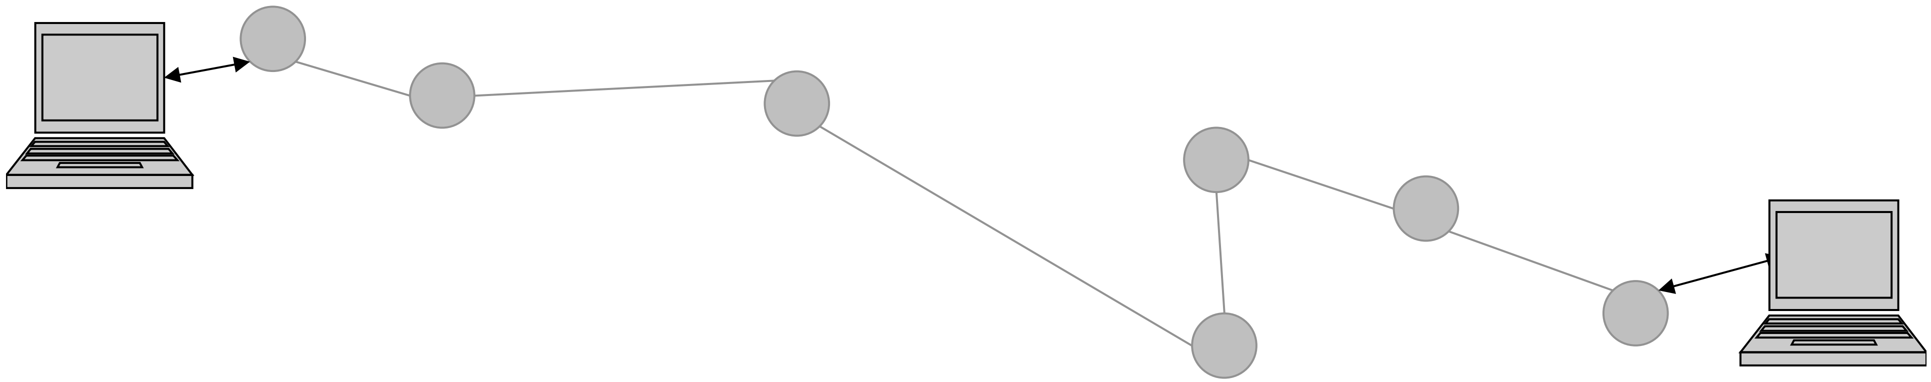

- traceroute www.duke.edu
  - traceroute to www.duke.edu (152.3.189.3), 30 hops max, 60 byte packets
  - 1 fp-core.research.att.com (135.207.16.1) 2 ms 1 ms 1 ms
  - 2 ngx19.research.att.com (135.207.1.19) 1 ms 0 ms 0 ms
  - 3 12.106.32.1 1 ms 1 ms 1 ms
  - 4 12.119.12.73 2 ms 2 ms 2 ms
  - 5 tbr1.n54ny.ip.att.net (12.123.219.129) 4 ms 5 ms 3 ms
  - 6 ggr7.n54ny.ip.att.net (12.122.88.21) 3 ms 3 ms 3 ms
  - 7 192.205.35.98 4 ms 4 ms 8 ms
  - 8 jfk-core-02.inet.qwest.net (205.171.30.5) 3 ms 3 ms 4 ms
  - 9 dca-core-01.inet.qwest.net (67.14.6.201) 11 ms 11 ms 11 ms
  - 10 dca-edge-04.inet.qwest.net (205.171.9.98) 11 ms 15 ms 11 ms
  - 11 gw-dc-mcnc.ncren.net (63.148.128.122) 18 ms 18 ms 18 ms
  - 12 rlgh7600-gw-to-rlgh1-gw.ncren.net (128.109.70.38) 18 ms 18 ms 18 ms
  - 13 roti-gw-to-rlgh7600-gw.ncren.net (128.109.70.18) 20 ms 20 ms 20 ms
  - 14 art1sp-tel1sp.netcom.duke.edu (152.3.219.118) 23 ms 20 ms 20 ms
  - 15 webhost-lb-01.oit.duke.edu (152.3.189.3) 21 ms 38 ms 20 ms
- 1 traceroute measurement: about 1KB

Large-scale traceroute experiments

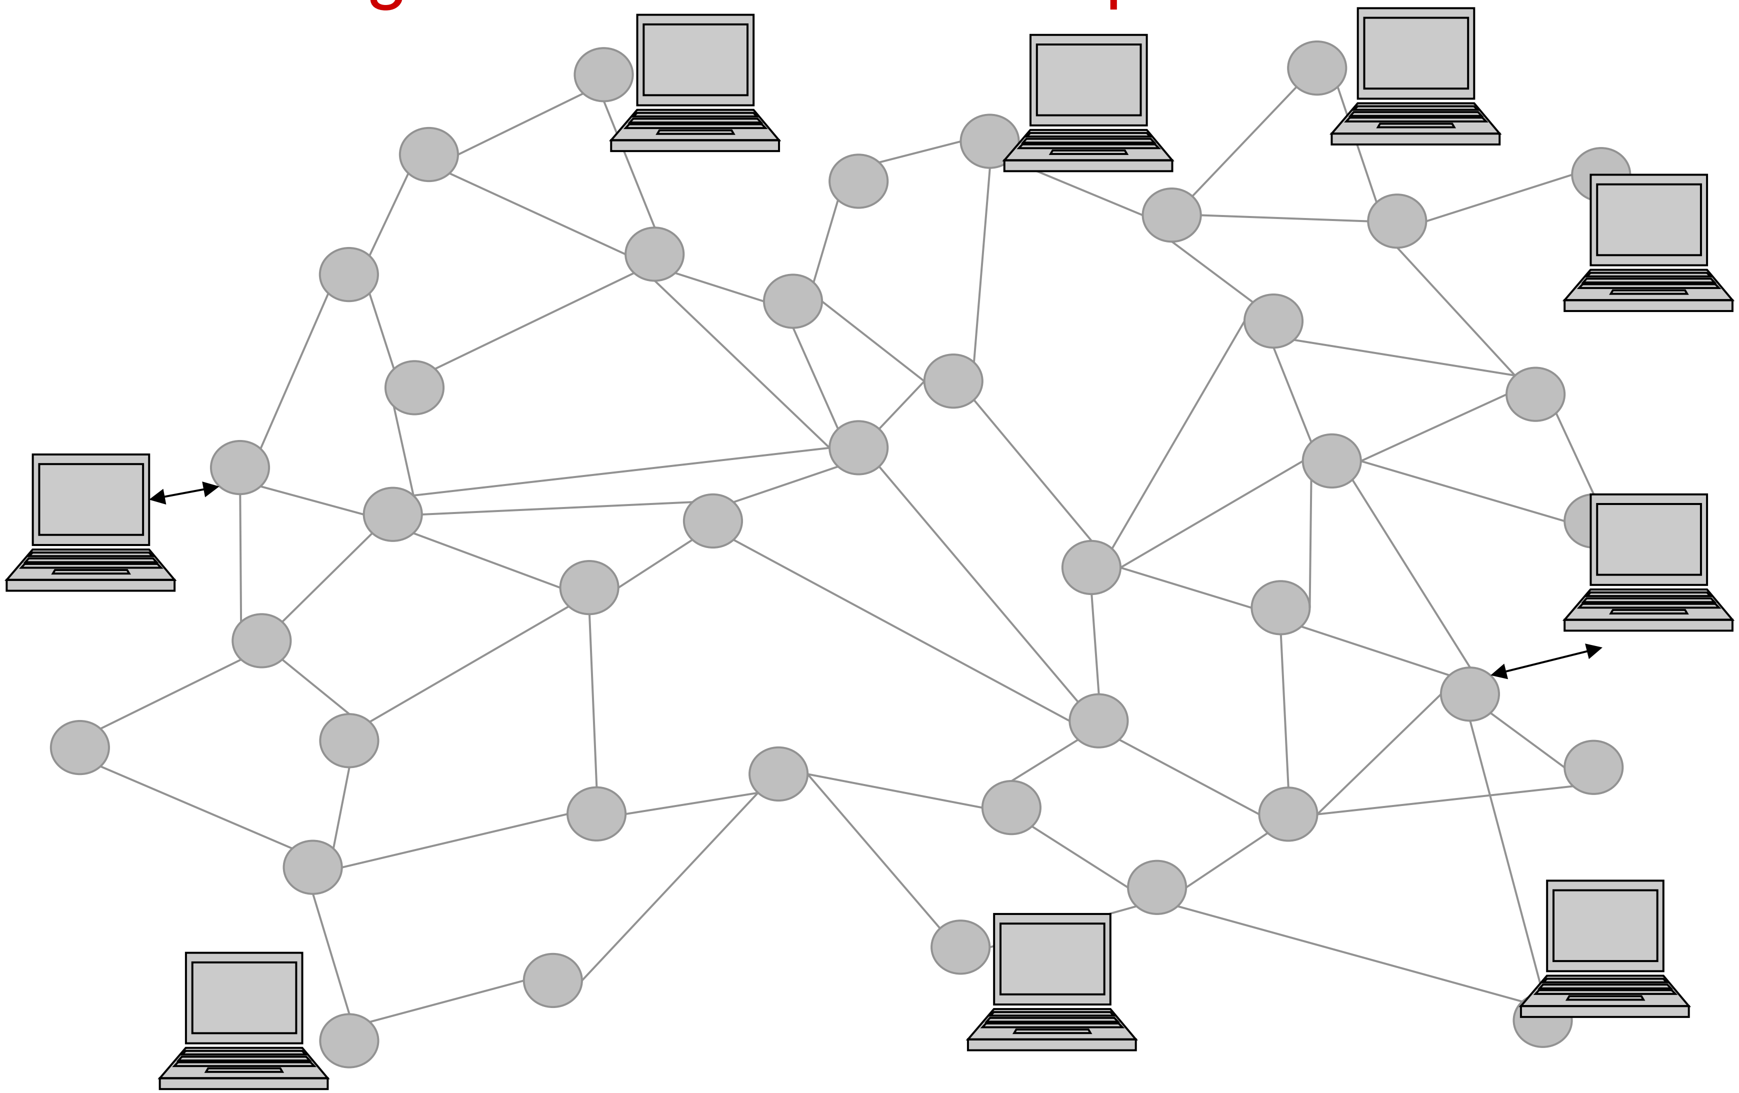

1 million x 1 million traceroutes: 1PB

Two Examples of inferred ISP topology

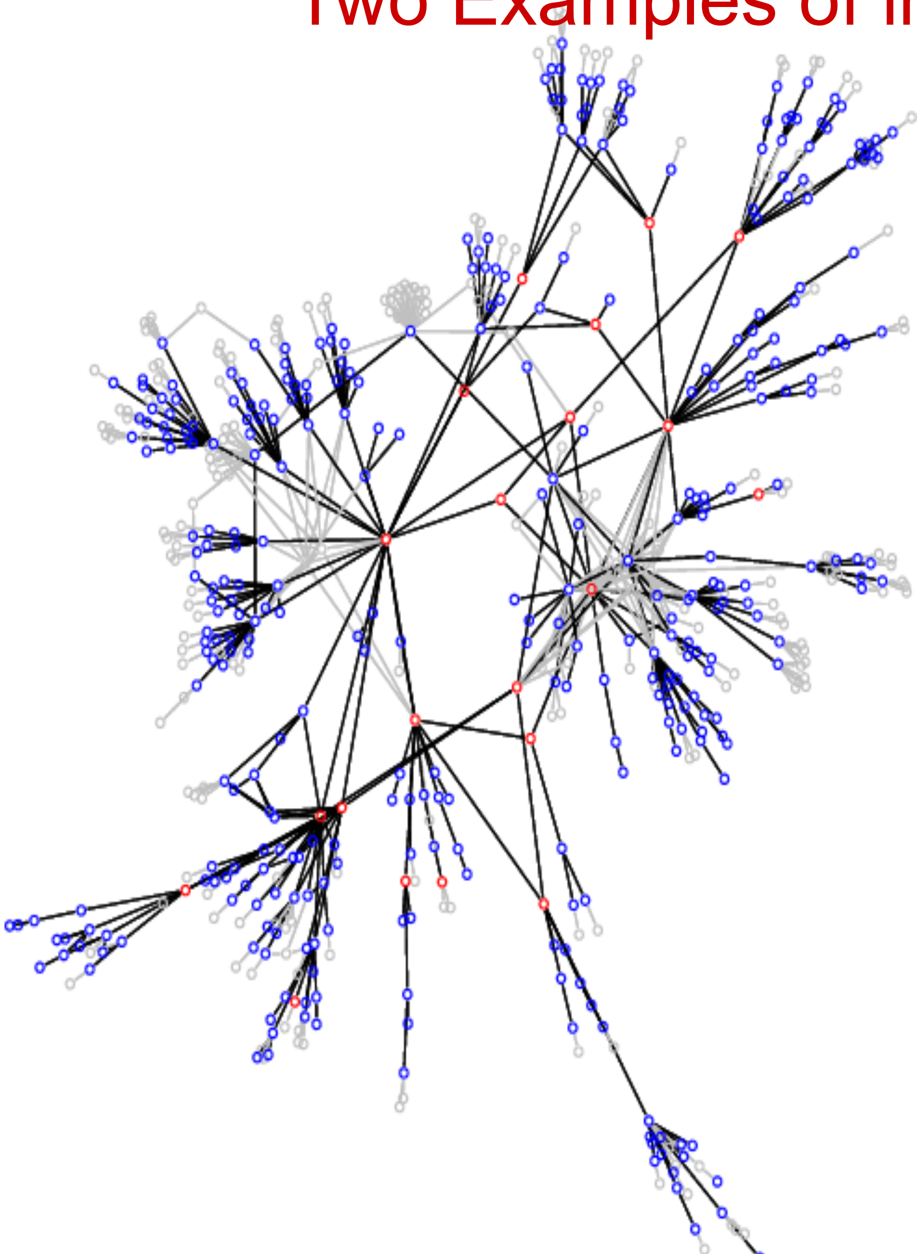

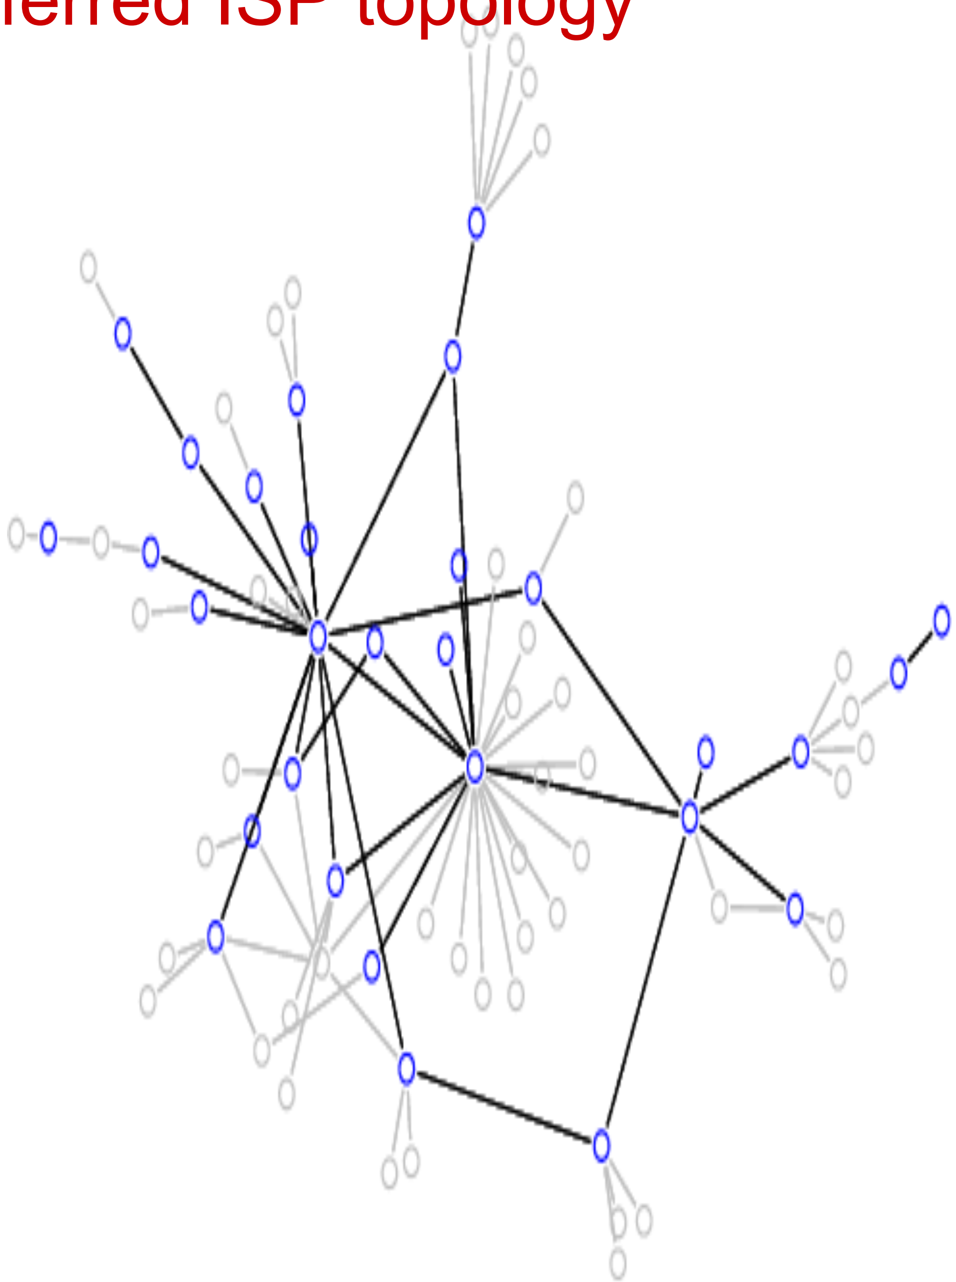

http://www.isi.edu/scan/mercator/mercator.html

About the Traceroute tool (1)

- traceroute is strictly about IP-level connectivity
  - Originally developed by Van Jacobson (1988)
  - Designed to trace out the route to a host
- Using traceroute to map the router-level topology
  - Engineering hack
  - Example of what we can measure, not what we want to measure!
- Basic problem #1: IP alias resolution problem
  - How to map interface IP addresses to IP routers
  - Largely ignored or badly dealt with in the past
  - New efforts in 2008 for better heuristics …

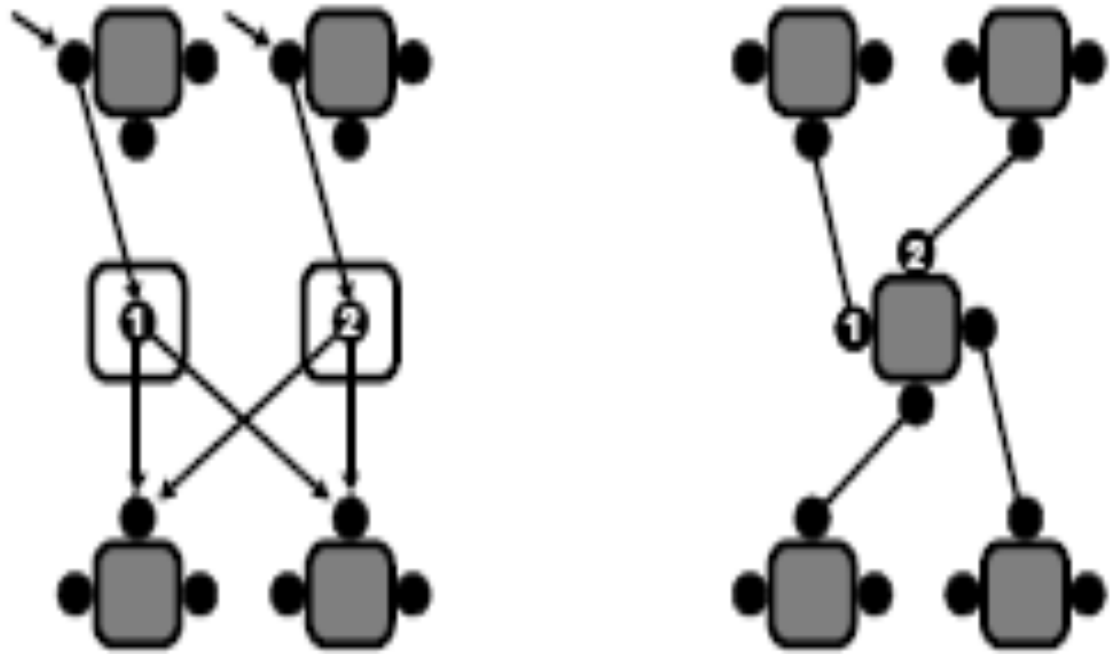

Interfaces 1 and 2 belong to the same router

IP Alias Resolution Problem for Abilene (thanks to Adam Bender)

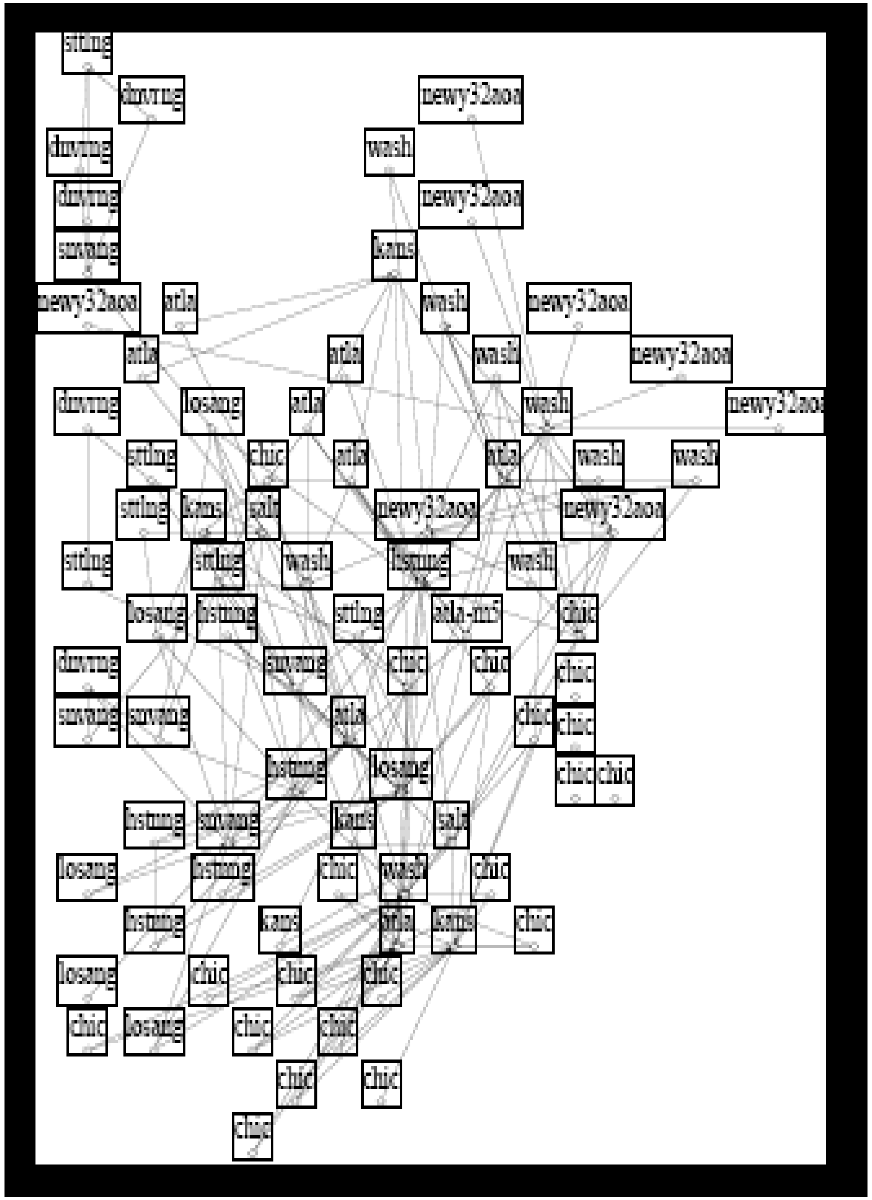

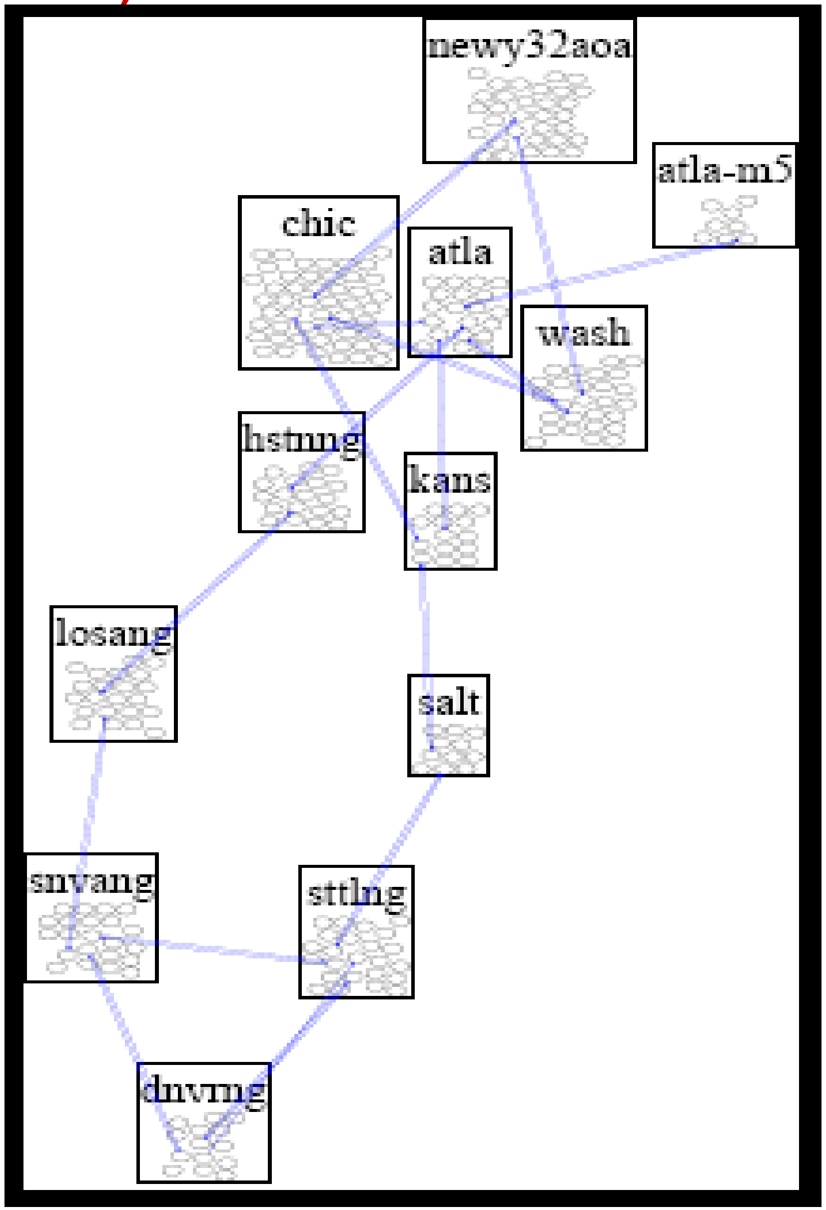

About the Traceroute tool (2)

- traceroute is strictly about IP-level connectivity
- Basic problem #2: Layer-2 technologies (e.g., MPLS, ATM)
  - MPLS is an example of a circuit technology that hides the network's physical infrastructure from IP
  - Sending traceroutes through an opaque Layer-2 cloud results in the "discovery" of high-degree nodes, which are simply an artifact of an imperfect measurement technique.
  - This problem has been largely ignored in all large-scale traceroute experiments to date.

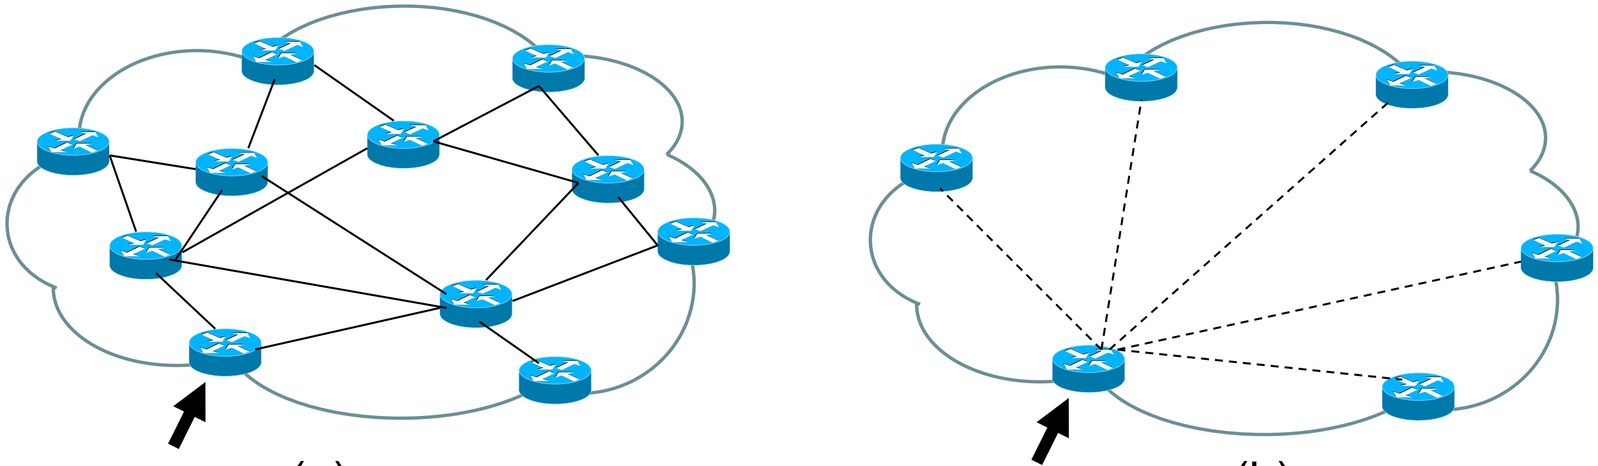

(a) (b)

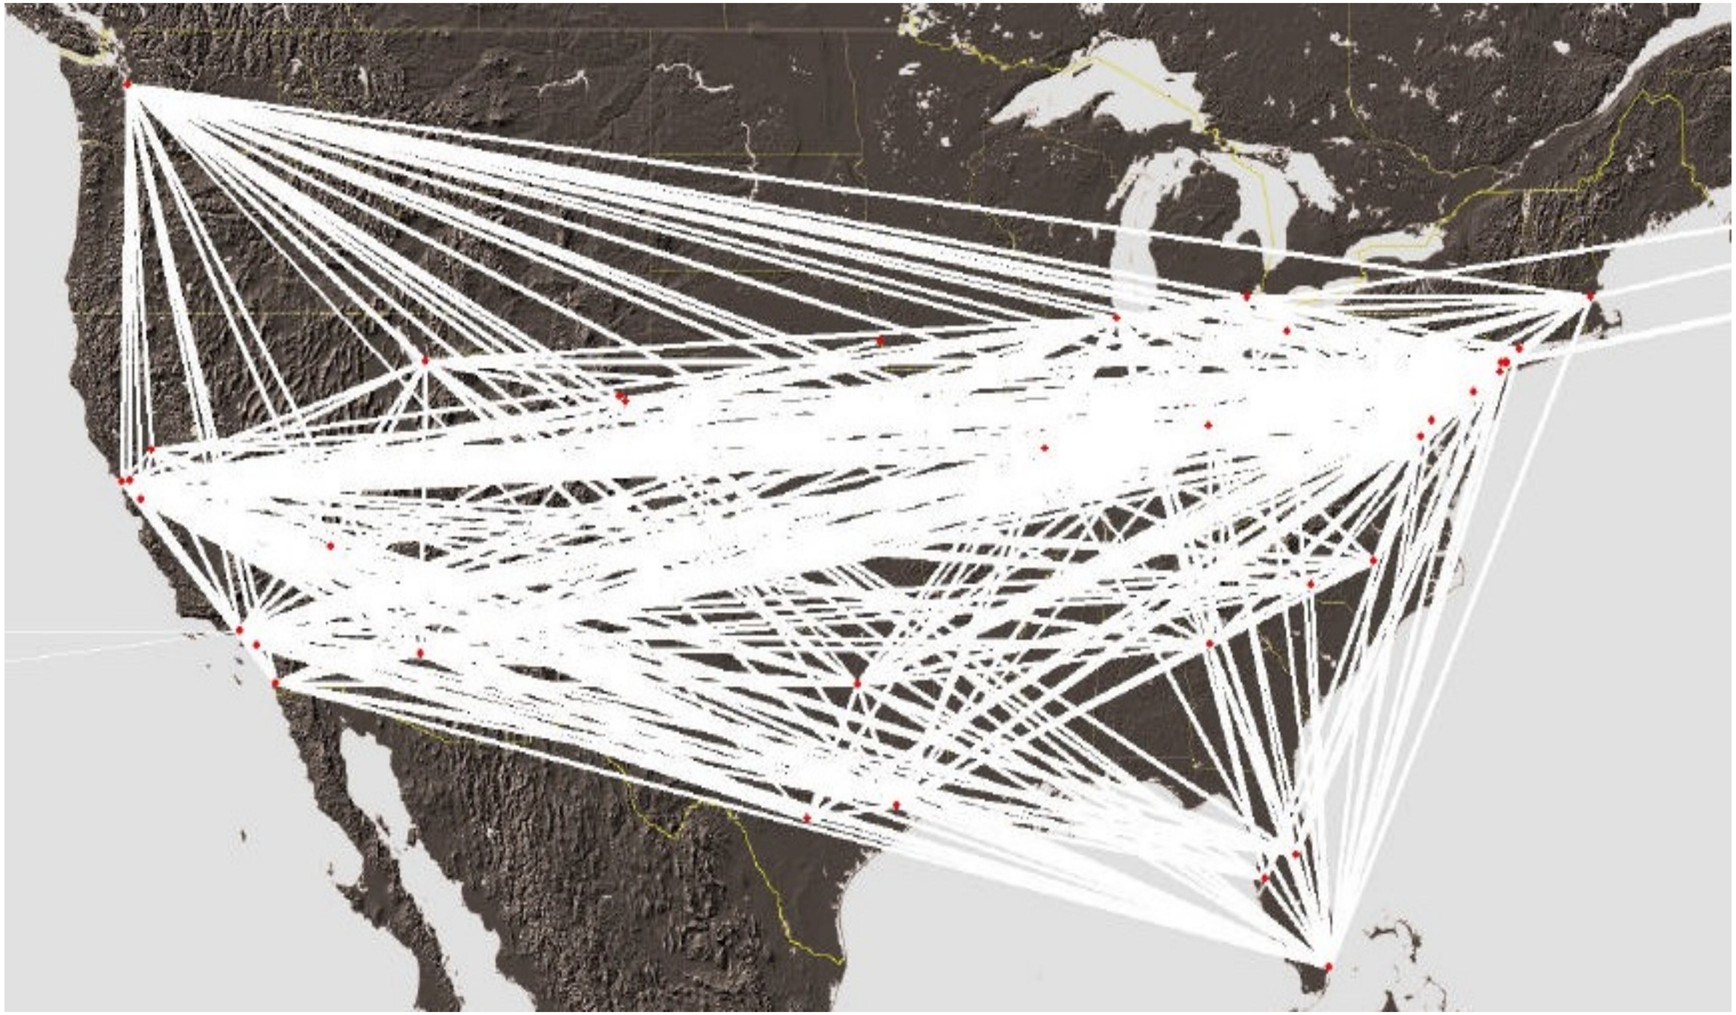

About the Traceroute tool (3)

- The irony of traceroute measurements
  - The high-degree nodes in the middle of the network that traceroute reveals are not for real …
  - If there are high-degree nodes in the network, they can only exist at the edge of the network where they will never be revealed by generic traceroute-based experiments …
- Additional irony
  - Bias in (mathematical abstraction of) traceroute
  - Has been a major focus within CS/Networking literature
  - Non-issue in the presence of above-mentioned problems

Example 1: Lessons learned

- Know your measurement technique!
  - Question: Can you trust the data obtained by your tool?
- Know your data!
  - Critical role of Data Hygiene in the Petabyte Age
  - Corollary: Petabytes of garbage = garbage
  - Data hygiene is often viewed as “dirty/unglamorous” work
  - Question: Can the data be used for the purpose at hand?
- Regarding Example 1:
  - (Current) traceroute measurements are of (very) limited use for inferring router-level connectivity
  - It is unlikely that future traceroute measurements will be more useful for the purpose of router-level inference

A textbook example for what can go wrong …

- *J.-J. Pansiot and D. Grad, "On routes and multicast trees in the Internet," ACM Computer Communication Review 28(1), 1998.*
  - *Original traceroute data -- purpose for using the data is explicitly stated*
  - *Most of the issues with traceroute are listed!*
- M. Faloutsos, P. Faloutsos, and C. Faloutsos, *"On the power-law relationships of the Internet topology",* Proc. ACM SIGCOMM'99, 1999.
  - Rely on the Pansiot-Grad data, but use it for a very different purpose
  - Take the available data at face value, even though Pansiot/Grad list most of the problems
  - There is no scientific basis for the reported power-law findings!
- R. Albert, H. Jeong, and A.-L. Barabasi, "*Error and attack tolerance of complex networks*", Nature, 2000.
  - Do not even cite original data source (i.e., Pansiot/Grad)
  - Take the results of FFF'99 at face value
  - The reported results are all wrong!

Applying lessons to Example 2

- Example 2: Use of traceroute measurements to infer Internet topology at the level of Autonomous Systems (ASes)
- Know your measurement technique!
  - traceroute (see Example 1)
- Know your data!
  - Main source of errors: IP address sharing between BGP neighbors makes mapping traceroute paths to AS paths very difficult
  - Up to 50% of traceroute-derived AS adjacencies appear to be bogus

Applying lessons to Example 2 (cont.)

- Regarding Example 2
  - (Current) traceroute measurements are of (very) limited use for inferring AS-level connectivity
  - Obtaining the "ground truth" is very challenging
  - It is possible that in the future, more targeted traceroute measurements in conjunction with BGP data will be more useful for the purpose of inferring AS-level connectivity

Applying lessons to Example 3

- Example 3: Use of BGP data to infer Internet topology at the level of Autonomous Systems (ASes)
- Know your measurement technique!
  - BGP -- de facto inter-domain routing protocol
  - BGP -- designed to propagate reachability information among ASes, not connectivity information
  - Engineering hack – not designed to obtain connectivity information
  - Example of what we can measure, not what we want to measure!
  - Collect BGP routing information base (RIB) information from as many routers as possible

Applying lessons to Example 3 (cont.)

- Know your data!
  - Examining the hygiene of BGP measurements requires significant commitment and domain knowledge
  - Parts of the available data seem accurate and solid (i.e., customer-provider links, nodes)
  - Parts of the available data are highly problematic and incomplete (i.e., peer-to-peer links)
  - "Ground truth" is hard to come by
- Regarding Example 3
  - (Current) BGP-based measurements are of questionable quality for inferring AS-level connectivity
  - Obtaining the "ground truth" is very challenging
  - It is possible that in the future, more targeted traceroute measurements in conjunction with BGP data will be more useful for the purpose of inferring AS-level connectivity

A Reminder

- Data-driven network analysis in the presence of high-quality data that can be taken at face value
  - *"All models are wrong … but some are useful"* (G.E.P. Box)

- Data-driven network analysis in the presence of highly ambiguous data that should not be taken at face value
  - *"When exactitude is elusive, it is better to be approximately right than certifiably wrong." (B.B. Mandelbrot)*

SOME RELATED REFERENCES

- L. Li, D. Alderson, W. Willinger, and J. Doyle, *A first-principles approach to understanding the Internet's router-level topology, Proc. ACM SIGCOMM 2004*.
- J.C. Doyle, D. Alderson, L. Li, S. Low, M. Roughan, S. Shalunov, R. Tanaka, and W. Willinger. The "robust yet fragile" nature of the Internet. *PNAS 102 (41)*, 2005.
- D. Alderson, L. Li, W. Willinger, J.C. Doyle. Understanding Internet Topology: Principles, Models, and Validation. *ACM/IEEE Trans. on Networking* 13(6), 2005.
- L. Li, D. Alderson, J.C. Doyle, W. Willinger. *Toward a Theory of Scale-Free Networks: Definition, Properties, and Implications. Internet Mathematics 2 (4),* 2006.
- R. Oliveira, D. Pei, W. Willinger, B. Zhang, L. Zhang. *In Search of the elusive Ground Truth: The Internet's AS-level Connectivity Structure. Proc. ACM SIGMETRICS 2008.*
- B. Krishnamurthy and W. Willinger. *What are our standards for validation of measurement-based networking research? Proc. ACM HotMetrics Workshop 2008.*
- W. Willinger, D. Alderson, and J.C. Doyle. *Mathematics and the Internet: A Source of Enormous Confusion and Great Potential. Notices of the AMS, Vol. 56, No. 2, 2009.*