tulip: A tool for locating performance problems along Internet paths

Ratul Mahajan

Neil Spring

David Wetherall

Tom Anderson

University of Washington

Where is the problem along this path?

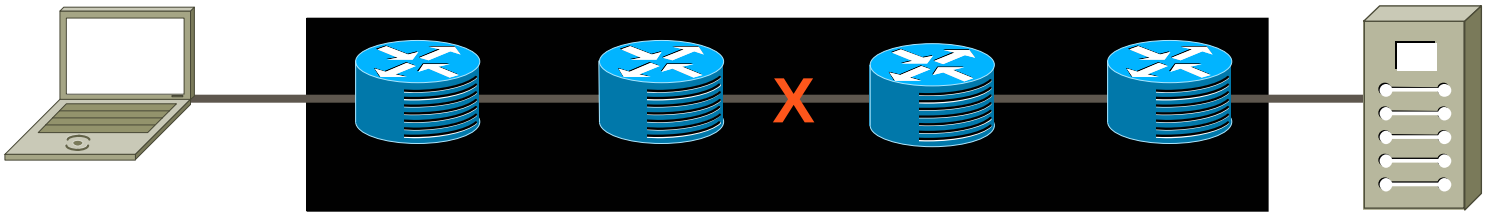

The Internet is a black box

- Problems such as
  - packet loss, queuing (congestion) and significant reordering

Limitations of existing tools

- SNMP stats are limited in scope
  - can be used only inside your domain
  - granularity issues
- *ping and traceroute*-like tools don't deal with path asymmetry

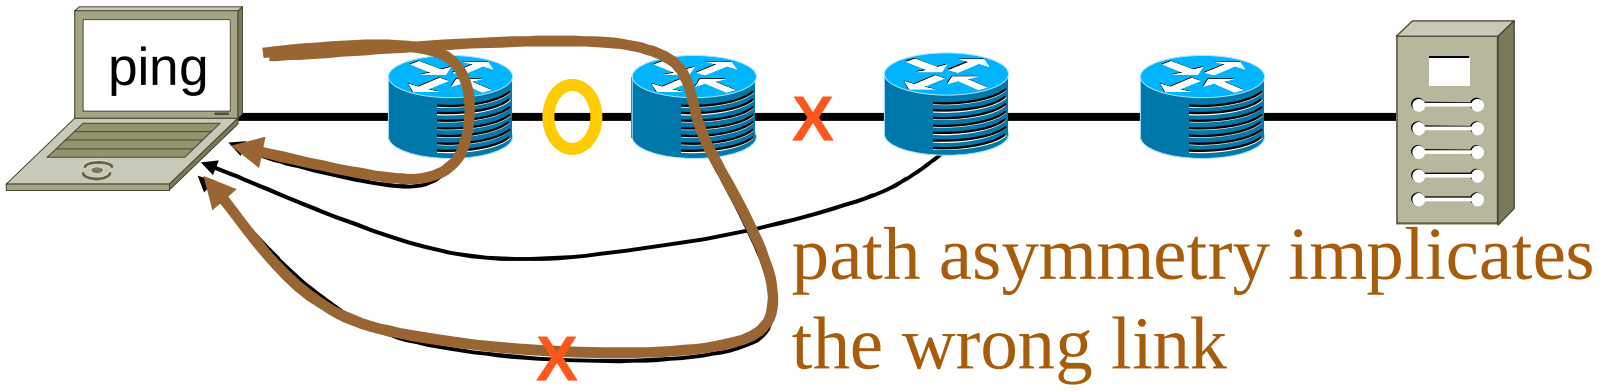

Overview of tulip

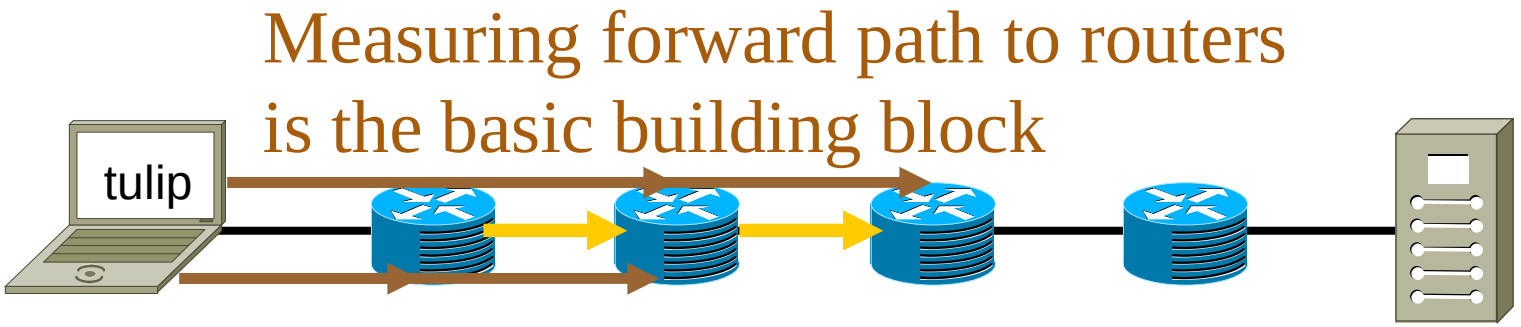

- Localizes reordering, queuing and loss
  - works from a host to an arbitrary IP address
- Infers link properties by subtracting path properties
  - path to router must be a prefix of the path to dst.
- Uses ICMP timestamps and IP identifiers

Queuing on the forward path [cing]

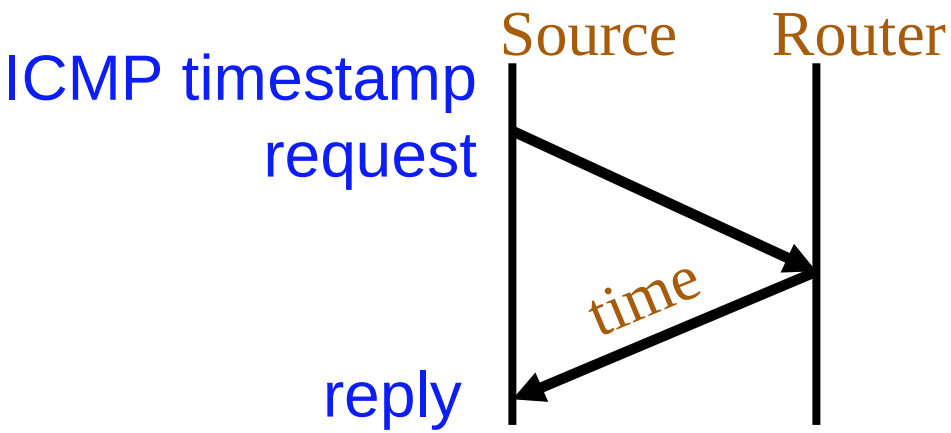

- ICMP timestamps to access router's clock
  - 1 ms resolution; supported by over 90% routers
- Queuing inferred from delay variation
  - engineering: clock calibration, response generation time

Loss on the forward path

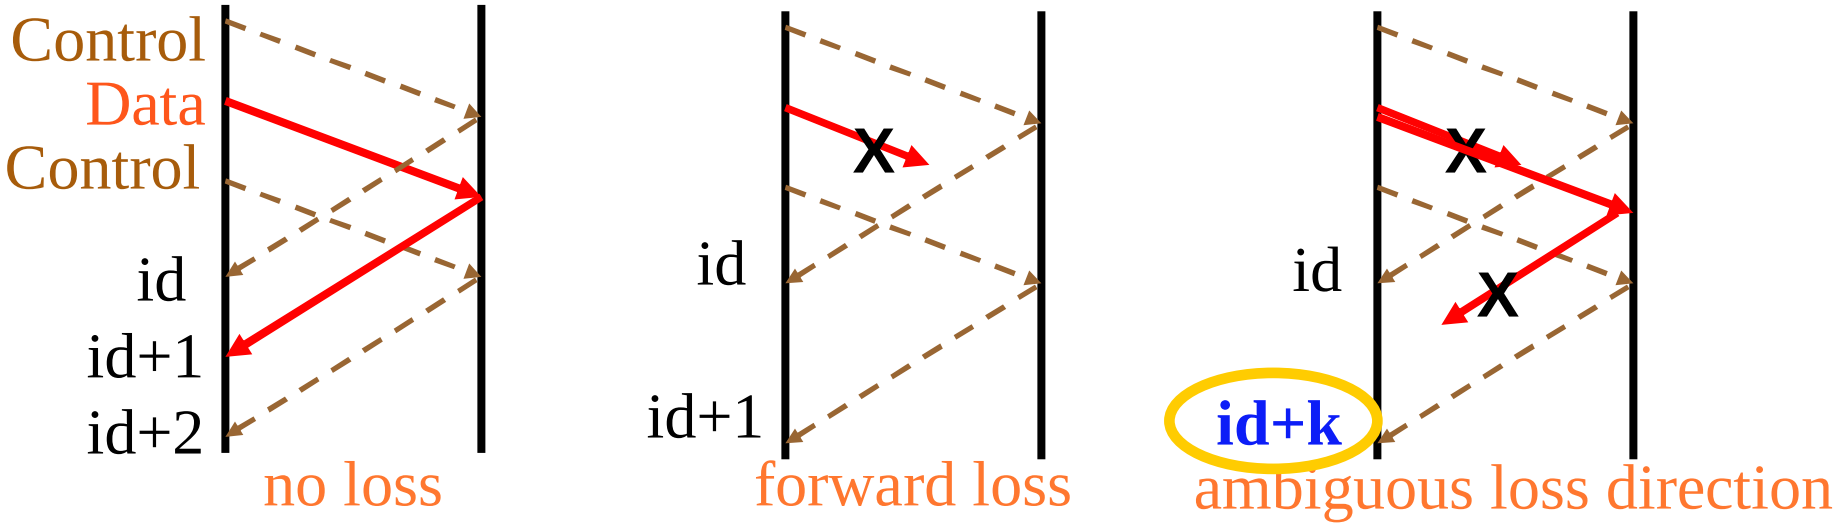

- Loss measurements use the IP identifier field
  - 70% of routers implement IP-ID as a counter
- Unambiguous forward data packet loss when control responses get consecutive IP-IDs
- Robust to response rate-limiting at the routers

Implementation

- Implemented in ruby, on top of Scriptroute
  - runs on Linux, FreeBSD and OS X
- Runs in three phases
  1. discover routers along the path
  2. test which routers support fwd path diagnosis
  3. measure paths to routers
- Reports properties of one-way and round trip paths to intermediate routers

% tulip loss ratul.dhcp.nanog31.cnet.com

1 eureka-GE1-5.cac.washington.edu (128.208.4.100) **rt=0.000 (0/1000)**
fw=0.000 (0/1000) co=0.966 (966/1000) ro=0.000 (0/1000)

2 uwbr1-GE0-1.cac.washington.edu (140.142.150.23) rt=0.000 (0/1000)

3 prs1-wes-ge-1-0-0-0.pnw-gigapop.net (198.107.151.30) rt=0.000 (0/1000)

4 att-pwave-1.peer.pnw-gigapop.net (198.32.170.29) rt=0.001 (1/1000)

5 tbr1-p012502.st6wa.ip.att.net (12.123.203.170) rt=0.000 (0/1000)

6 tbr2-cl1.sffca.ip.att.net (12.122.12.113) rt=0.000 (0/1000)

7 gar4-p390.sffca.ip.att.net (12.123.13.178) rt=0.000 (0/1000)
fw=0.000 (0/1000) co=0.000 (0/1000) ro=0.000 (0/1000)

8 idf22-gsr12-1-pos-7-0.rwc1.attens.net (12.122.255.218) rt=0.000 (0/1000)

9 mdf3-bi4k-1-eth-1-1.rwc1.attens.net (216.148.209.62) rt=0.000 (0/1000)
fw=0.000 (0/1000) co=0.953 (953/1000) ro=0.000 (0/1000)

10 mdf3-bi4k-2-ve-57.rwc1.attens.net (63.241.72.146) rt=0.000 (0/1000)
fw=0.000 (0/1000) co=0.909 (905/996) ro=0.004 (4/1000)

11 ge4-7.cn-sfo1-6506-1.cnet.com (216.239.127.5) rt=0.000 (0/1000)
fw=0.000 (0/1000) co=0.988 (987/999) ro=0.001 (1/1000)

12 ge1-1.365-sfo1-6506-1.cnet.com (216.239.127.2) rt=0.000 (0/1000)
fw=0.000 (0/1000) co=0.965 (964/999) ro=0.001 (1/1000)

13 ge1-3-0.border.nanog31.cnet.com (216.239.118.2) **rt=0.000 (0/1000)**

14 gateway.nanog31.cnet.com (216.239.118.130) **rt=0.001 (1/676)**
fw=0.001 (1/676) co=0.950 (640/674) ro=0.001 (1/675)

15 ratul.dhcp.nanog31.merit.net (192.35.166.145) **rt=0.015 (14/904)**
fw=0.000 (0/904) co=0.989 (877/887) ro=0.003 (3/890)

Tulip in action

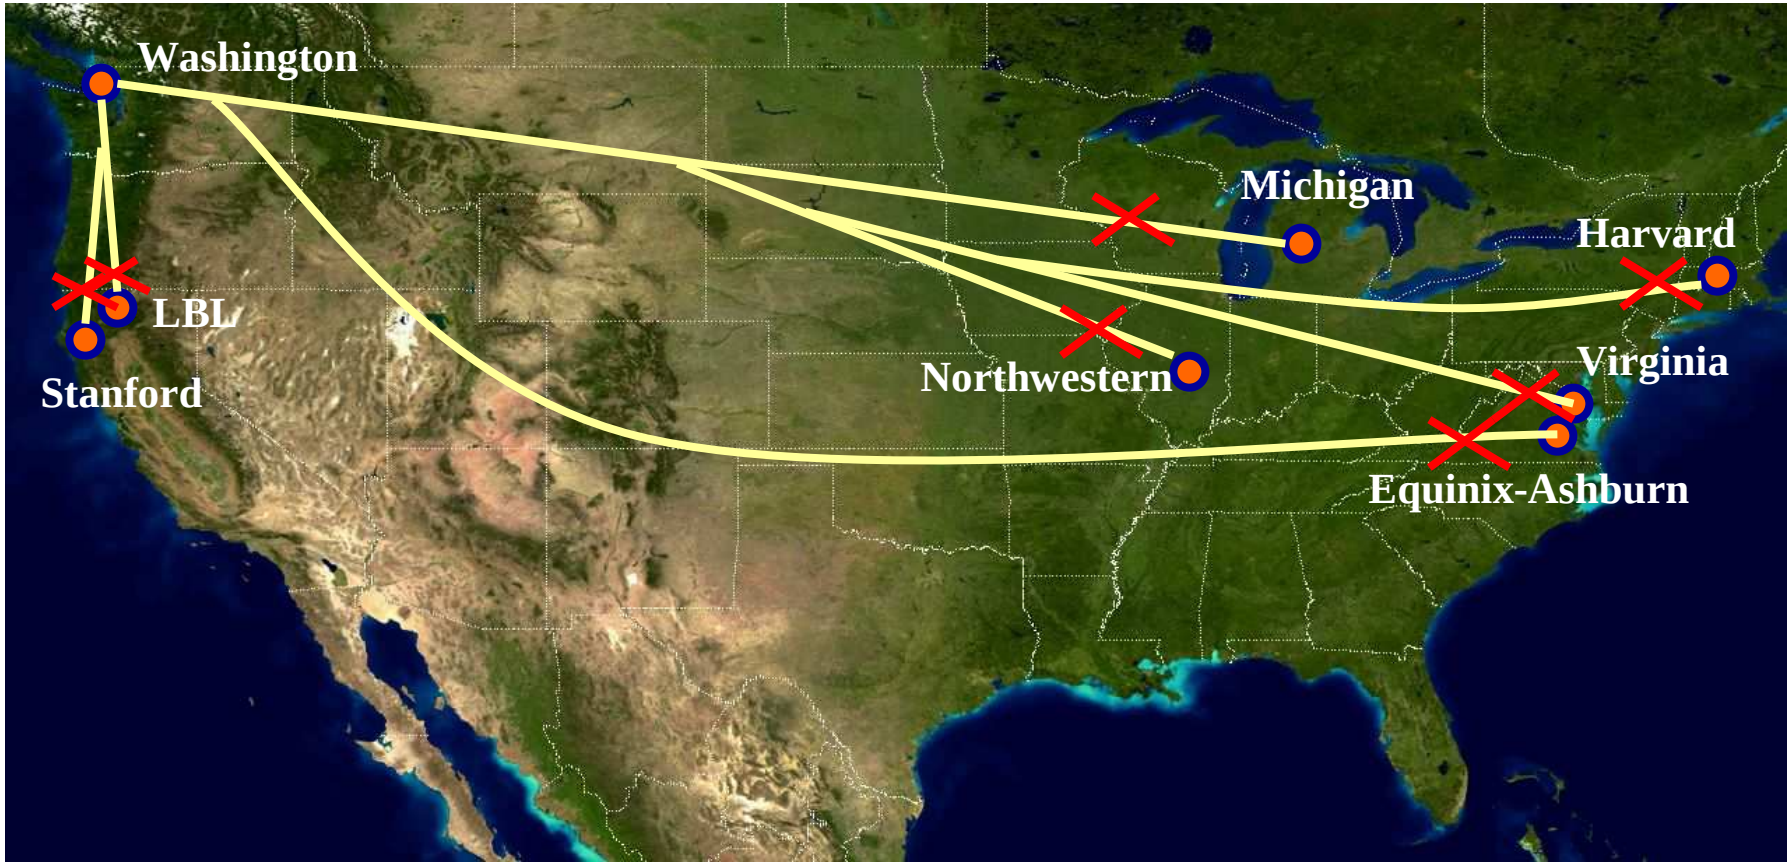

Consistency along the path (queuing)

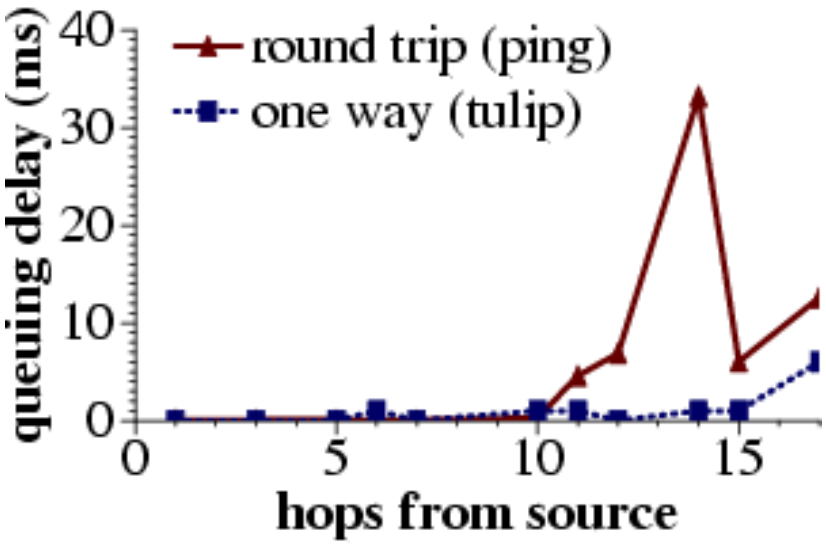

median queuing delay to intermediate routers in an example path

- One-way measurements are consistent
- Round trip measurements can be polluted

Evaluation

- Is it accurate?
  - end-to-end correctness
  - internal consistency

  (see the paper)

- What is the fault localization granularity?

Localization granularity

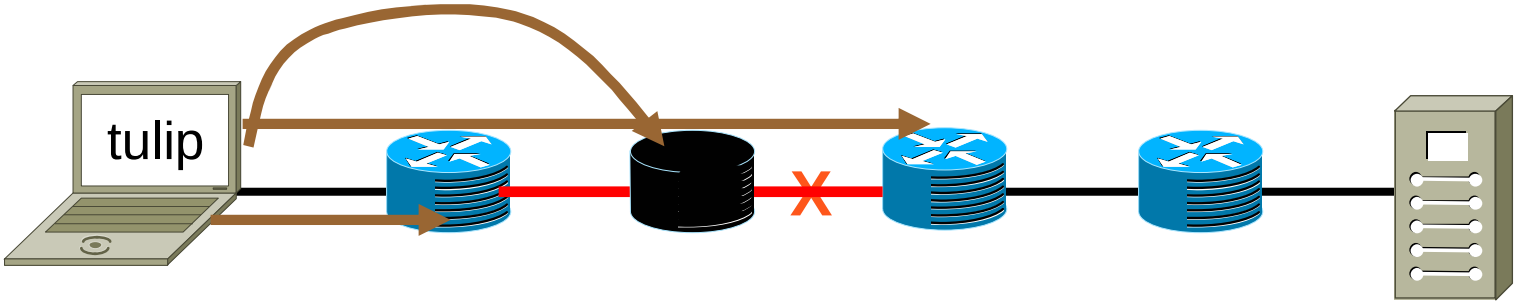

- Location of the fault is uncertain when
  - a router does not support a required feature
  - probes take a non-prefix path to a router

Localization granularity (2)

- Median is 2 hops for loss and 4 hops for queuing

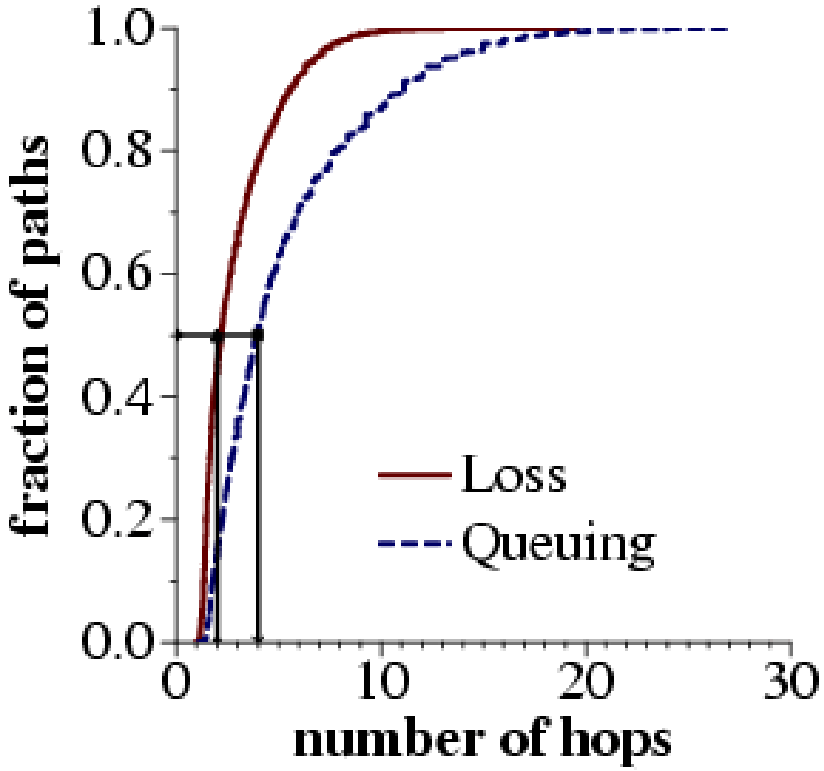

- Round trip probing can further improve granularity

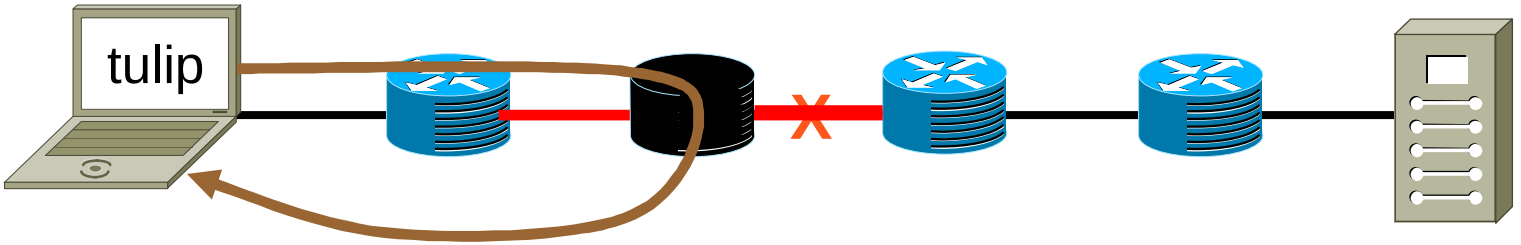

Improving effectiveness

- Turn on the features used by tulip
  - let yourself and other operators diagnose paths through your network
- Simple router changes can further improve the diagnostic ability of tulip
  - granularity of queuing localization
  - supporting multiple simultaneous loss rate measurements

Better timing information

- Problem: ICMP timestamp requests address the router directly
  - reduces diagnosis granularity when the prefix path property does not hold
- Fix: insert timestamps in TTL-expired messages
  - backwards compatible, incrementally deployable
  - 32 unused bits in the TTL-expired messages

Better counter support

- Problem: IP-ID is a shared counter
  - what if all of you start using tulip?
- Fix: maintain N (constant) counters
  - hash source address and IP-ID of the probe to select the counter
  - backwards compatible, incrementally deployable (today, N=1)

Summary

- Tulip is a new performance diagnosis tool
  - finds where packets are lost, reordered or queued to within 2-4 hops on average
  - compatible with asymmetric routing

- Download it at http://www.cs.washington.edu/research/networking/tulip