CSC2209
Computer Networks

Routing around Failures

Stefan Saroiu
Computer Science
University of Toronto

Administrivia

- The first 16 reviews have been graded
  - you should have received the grades already
- Vint Cerf's talk at 3pm
  - Meeting with students
- Mark Corner is visiting on Friday (DTNs)
  - Stratis asked whether students can get organized

Measurement study of path failures

- Conducted a week long measurement study
  - probed 3,153 destinations from 67 Planetlab sites
  - each destination is probed from exactly one node

- Goal is to answer the following:
  - How often do paths fail?
  - Where do failures occur?
  - How long do failures last?

Slide from Krishna Gummadi

Choosing destinations

- Understand how network paths to servers and broadband hosts differ
  - it has implications for different workloads/apps
    - Web transfers between servers and broadband hosts
    - VOIP apps between broadband hosts

- 3153 destinations:
  - 378 popular web servers
  - 1,139 broadband hosts
  - 1,636 randomly selected IPs

Slide from Krishna Gummadi

How often do paths fail?

- Failures do happen, but not frequently
  - on average each path sees 6 failures/week
  - server paths see 4 failures/week
  - broadband paths see 7 failures/week
- Most paths see at least one failure in a week
  - 85% of all paths
  - 78% of server paths
  - 88% of broadband paths
- Is this good news or bad news?

Slide from Krishna Gummadi

Categories of failure locations

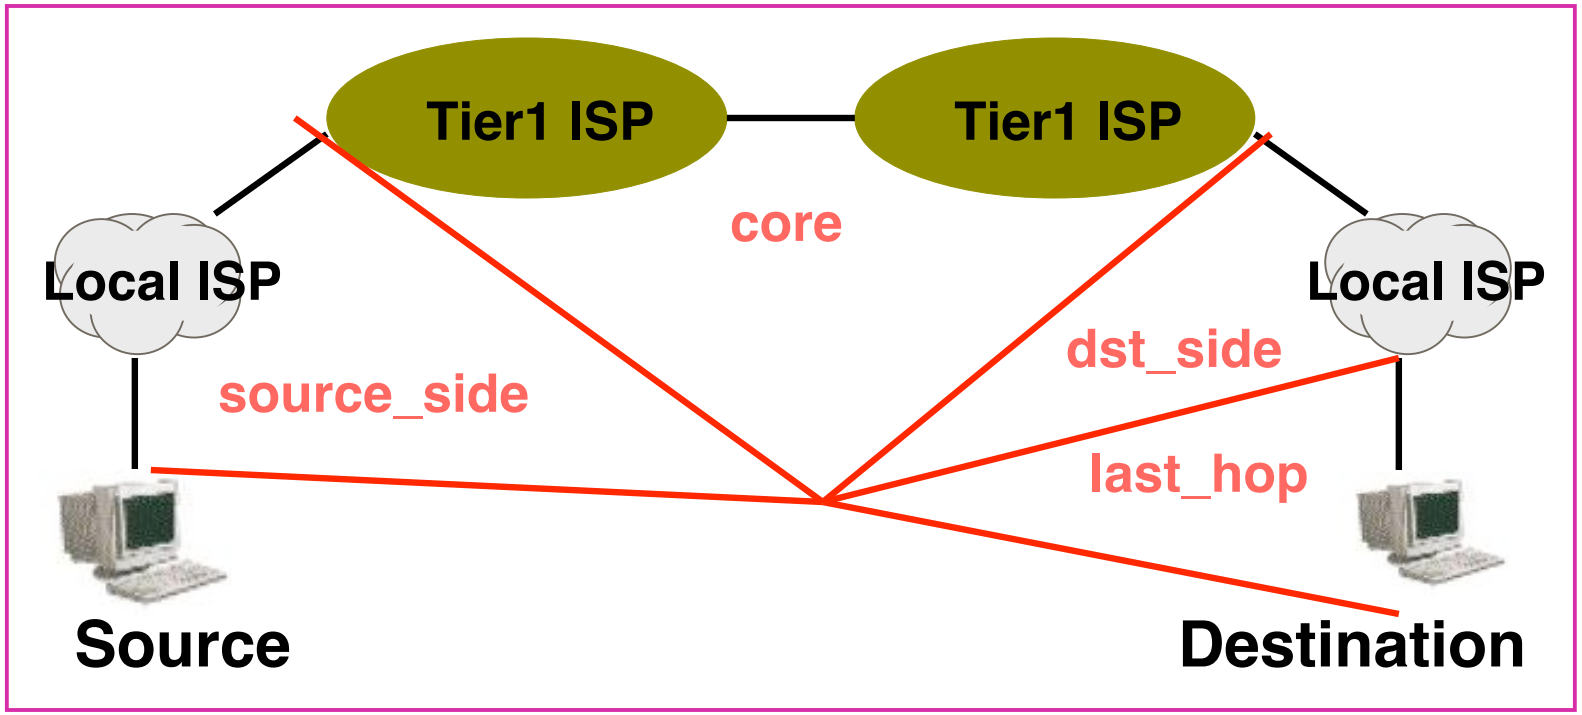

- Categories help distinguish between core and edge failures

Slide from Krishna Gummadi

Where do paths fail?

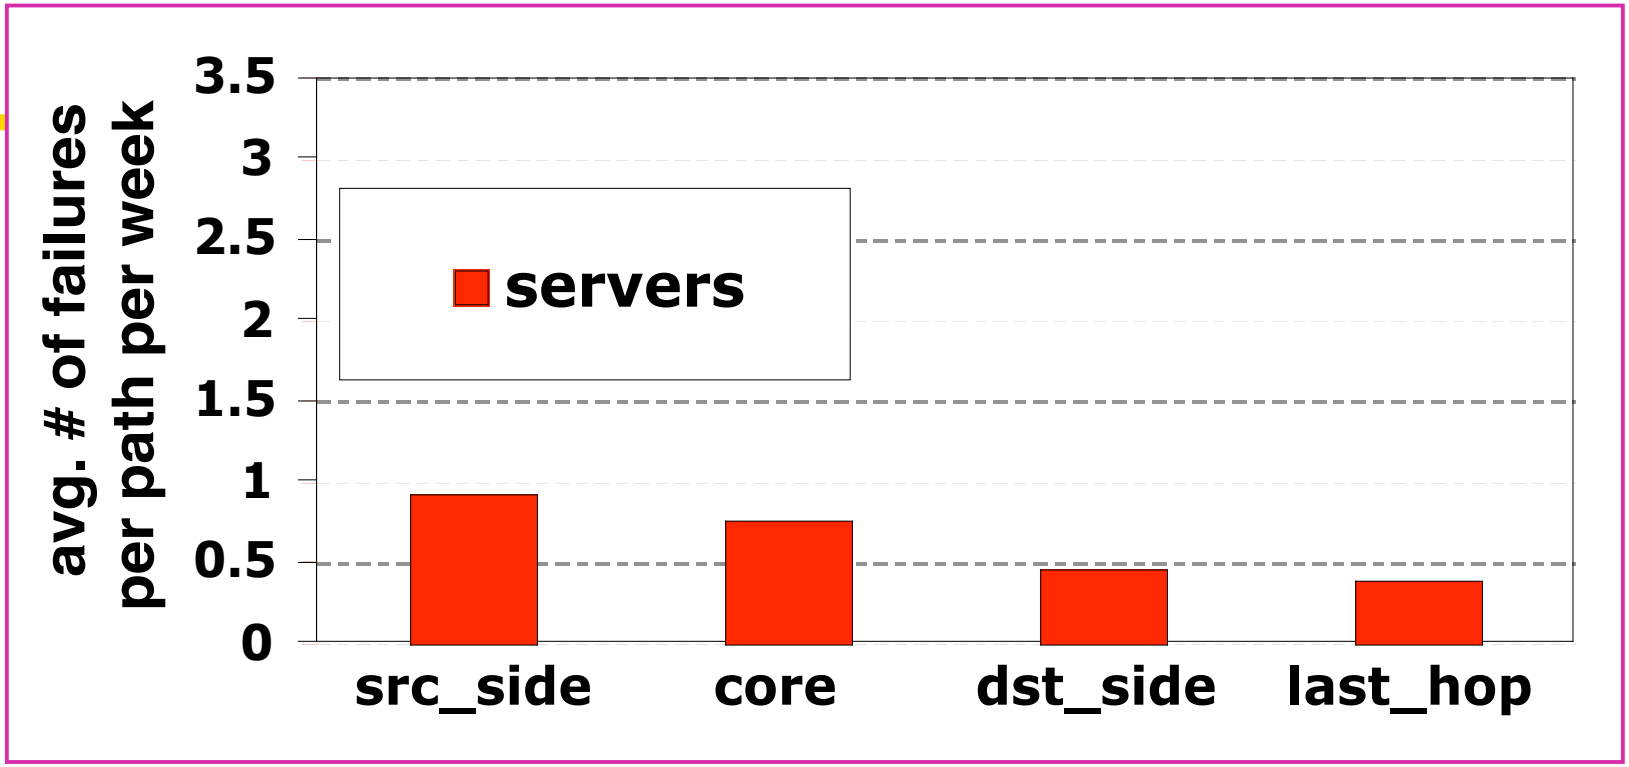

- Server path failures occur throughout the network
  - very few (16%) last_hop failures
  - suggests network is the dominating cause for server unavailability

Slide from Krishna Gummadi

Where do paths fail?

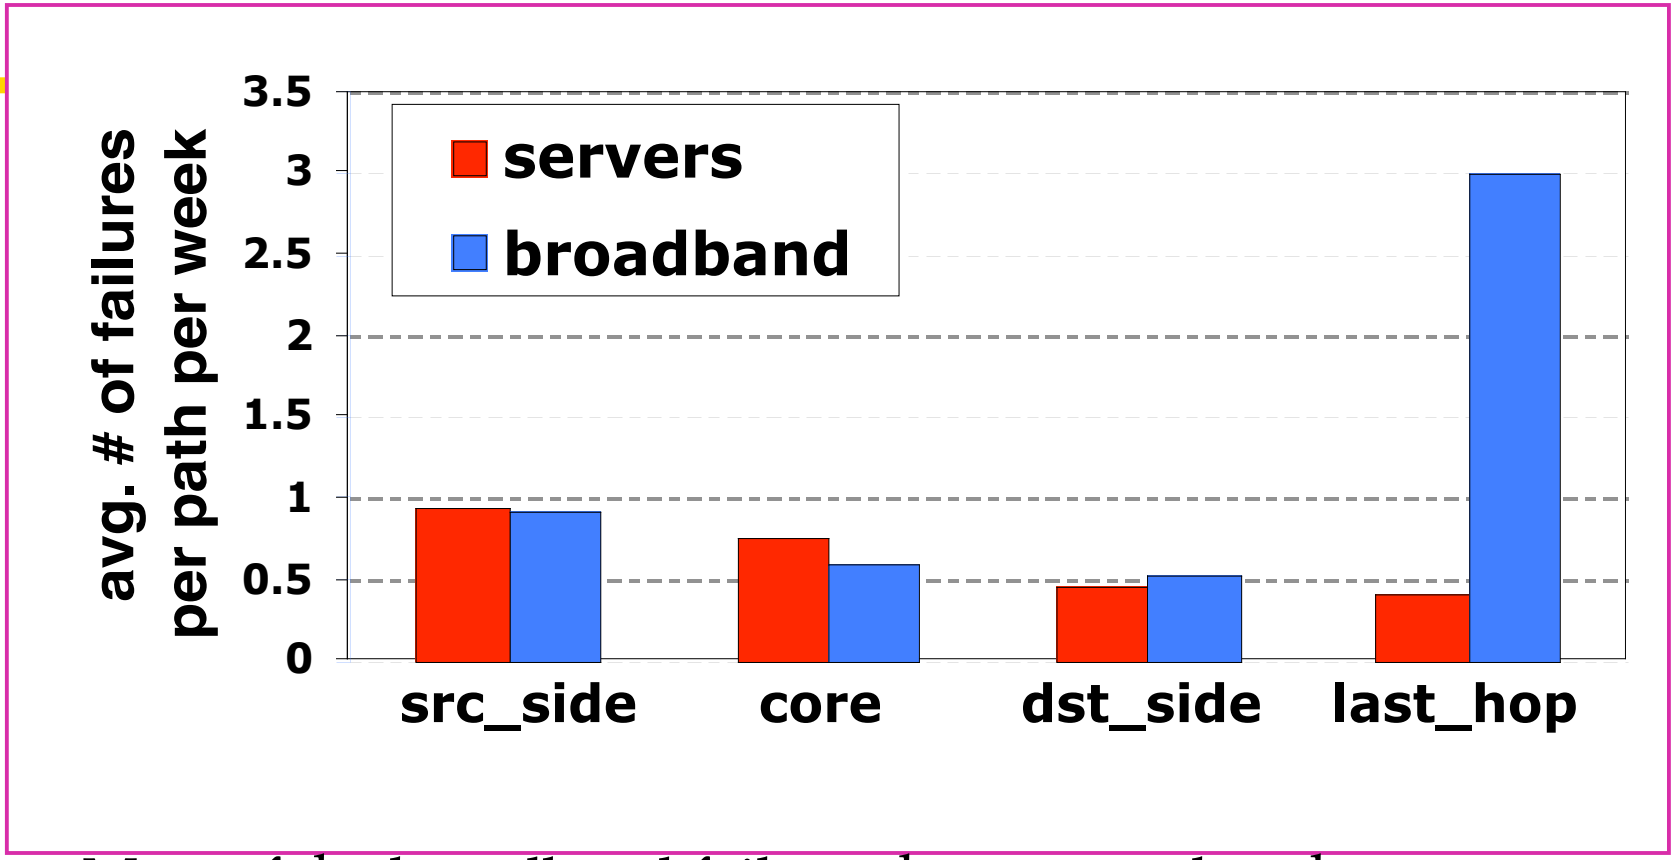

- Most of the broadband failures happen on last_hop
- Excluding last_hop, server and broadband paths see similar number of failures

Slide from Krishna Gummadi

How long do failures last?

- Failure durations are highly skewed
- Majority of failures are short
  - median failure duration: 1-2 min for all paths
  - median path availability: 99.9% for all paths
- A non-negligible fraction of paths see long failures
  - tend to occur on last_hop
  - mean path availabilty: 99.6% for servers and 94.4% for broadband

Slide from Krishna Gummadi

How Robust is Internet Routing?

Paxson 95-97	• 3.3% of all routes had serious problems
Labovitz 97-00	• 10% of routes available < 95% of the time • 65% of routes available < 99.9% of the time • 3-min minimum detection+recovery time; often 15 mins • 40% of outages took 30+ mins to repair
Chandra 01	• 5% of faults last more than 2.75 hours

1. Slow outage detection and recovery
2. Inability to detect badly performing paths
3. Inability to efficiently leverage redundant paths
4. Inability to perform application-specific routing
5. Inability to express sophisticated routing policy

Slide from Hari Balakrishnan' course notes

Implications for indirection routing

- Failures happen often enough that they are worth fixing
- But, they are rare enough that recovery schemes should be inexpensive under normal conditions
- Failures near the end-nodes limit the performance of indirection routing
  - good news: servers see very few failures near them
  - bad news: broadband hosts see many last_hop failures

Slide from Krishna Gummadi

Indirect Routing Service

- What are possible designs of an Indirect Routing Service?

Indirect Routing Service

- What are possible designs of an Indirect Routing Service?
  - AT&T or other Tier-1 networks deploy proxies
  - How about handling last-hop failures for broadband?

Indirect Routing Service

- Seems to make sense for routing around failures
- Does it make sense for performance?

Indirect Routing Service

- Seems to make sense for routing around failures
- Does it make sense for performance?
  - It depends:
    - WWW: not really, most flows are small
    - P2P: not really, most flows are big
    - VoIP: maybe for avoiding lossy or high-jitter paths

Additional Unexplored Issues

- Indirect routing has been shown to be useful whenever analyzed in isolation:
  - What if RON or SOSR become popular tomorrow?

Additional Unexplored Issues

- Indirect routing has been shown to be useful whenever analyzed in isolation:
  - What if RON or SOSR become popular tomorrow?

- How much overlap is there among best indirect paths?

Additional Unexplored Issues

- Indirect routing has been shown to be useful whenever analyzed in isolation:
  - What if RON or SOSR become popular tomorrow?
- How much overlap is there among best indirect paths?
- Are we just measuring Internet2 effects?