

UMassAmherst



***Informed Detour Selection  
Helps Reliability***

Boulat A. Bash



# Internet Reliability Problem

- Internet service outage probability 1.5% – 3.3% (Paxson, Dahlin, etc)
  - Far short of 0.001% (99.999%, or five nines of reliability) observed in telephone networks
- How can we improve? Add **redundancy**
  - Server redundancy
    - CDNs (Akamai)
  - **Path redundancy**




# Path Redundancy

intermediary 1

server

client

intermediary 2




# Path Redundancy

intermediary 1

server

client

intermediary 2




# Path Redundancy

intermediary 1

server

client

intermediary 2




# Outline

- Introduction and motivation
- Previous work
- Our proposal
- Experimental validation
- Discussion
- Conclusion




# Choosing intermeds: aggressive probing

- Aggressive monitoring of complete overlay graph
  - Example: RON (Resilient Overlay Network)



# Choosing intermeds: aggressive probing


- Aggressive monitoring of complete overlay graph
  - Example: RON (Resilient Overlay Network)
- Problem: monitoring overhead  $\Rightarrow$  lack of scaling






# Choosing intermeds: random choice

- Choose intermediaries randomly
  - No path monitoring, assume uncorrelated paths
  - Example: SOSR (Scalable One-hop Source Routing)




# Choosing intermeds: random choice

- Choose intermediaries randomly
  - No path monitoring, assume uncorrelated paths
  - Example: SOSR (Scalable One-hop Source Routing)



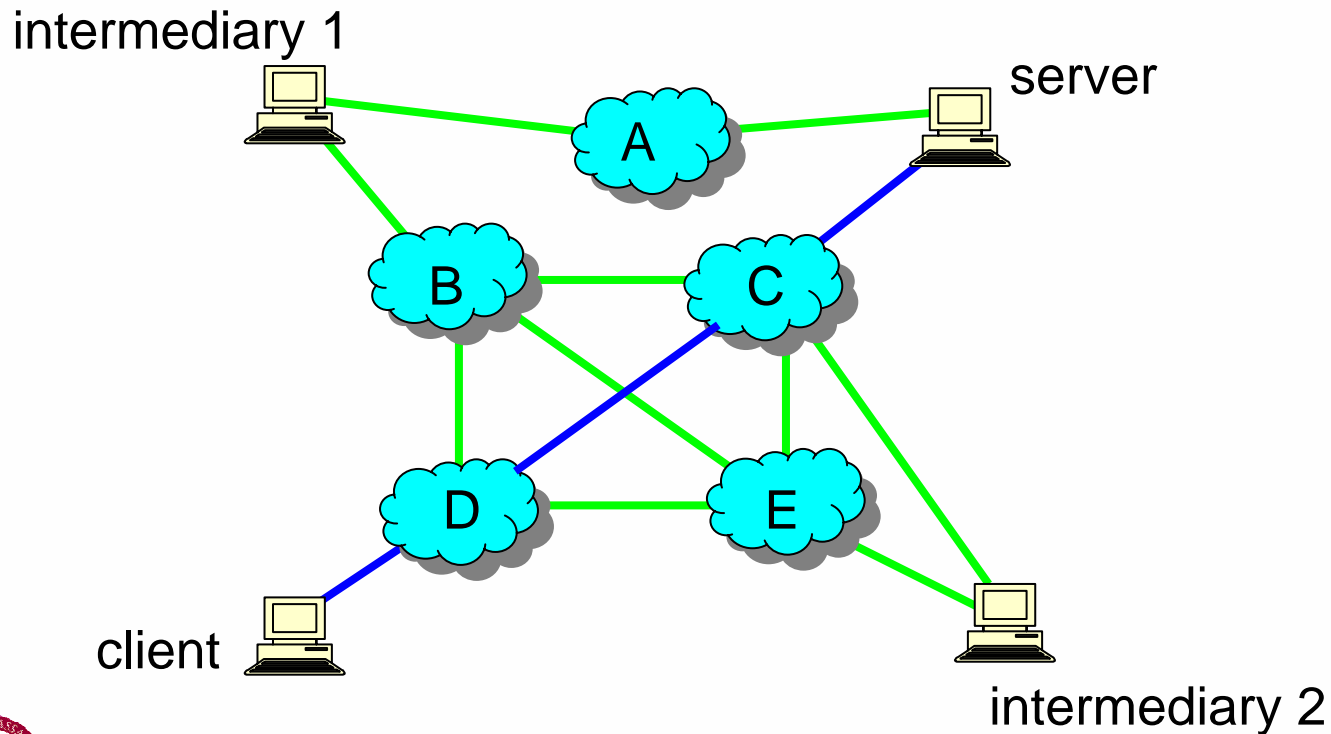
# Choosing intermeds: random choice

- Choose intermediaries randomly
  - No path monitoring, assume non-overlapped paths
  - Example: SOSR (Scalable One-hop Source Routing)
- Problem: overlapped paths  $\Rightarrow$  choose invalid detour



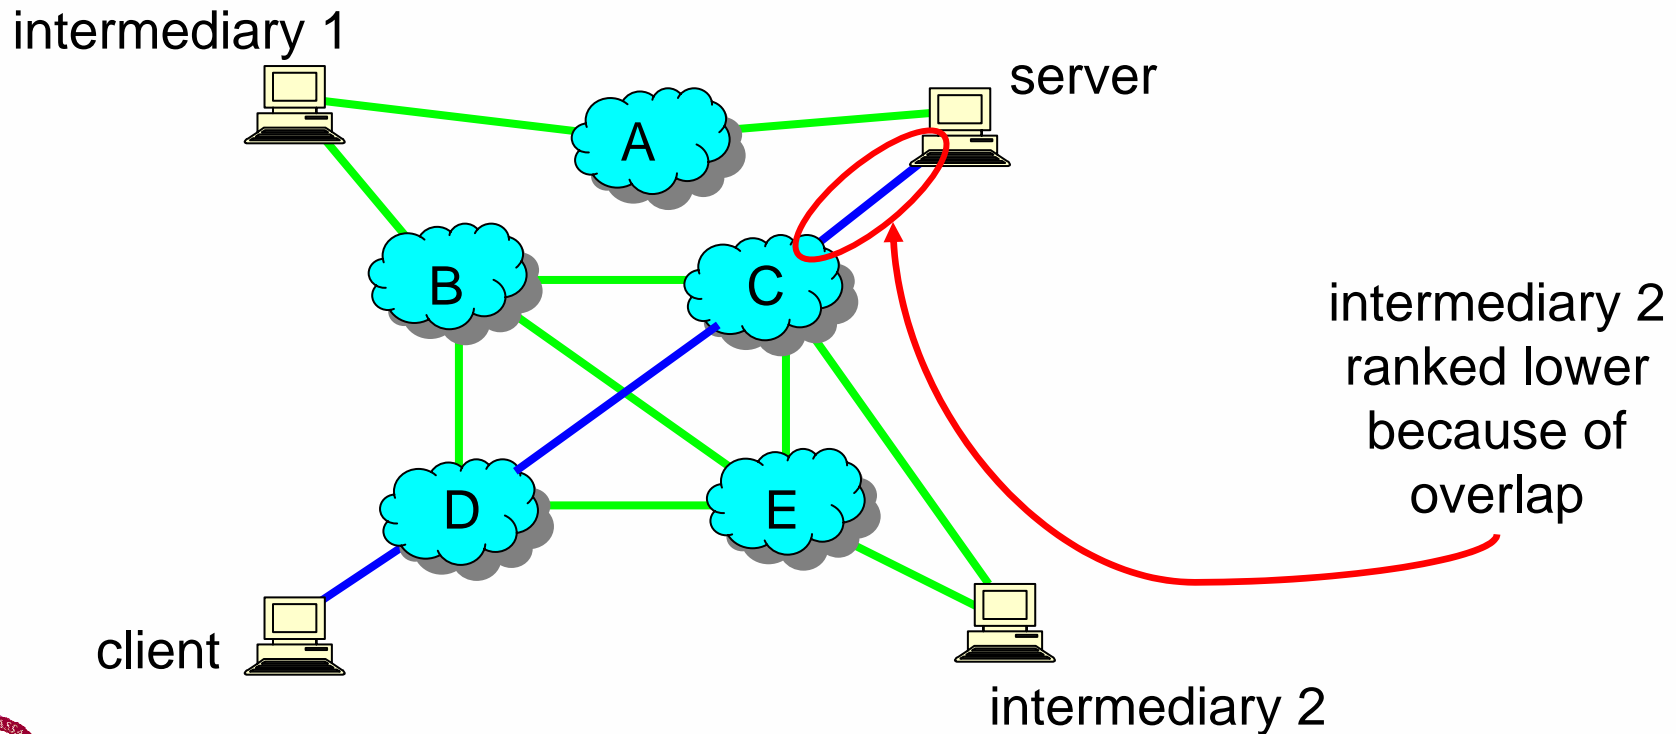
# Our Proposal

- Rank intermediaries by overlap with S-D path
  - Want less overlap
- No need to aggressively monitor: AS/PoP-level paths static over 24 hours



# Our Proposal

- Rank intermediaries by overlap with S-D path
  - Want less overlap
- No need to aggressively monitor: AS/PoP-level paths static over 24 hours




# Experimental Methodology

- Validation study
- PlanetLab-based experiment
  - Monitor 4,269 paths from 46 PL nodes to set of Internet routers for 379 hours
  - Each path monitored by one PL node
    - Ping every 15 seconds, failure if two pings missed
    - Intermediaries ping failed destination every 15 seconds during failure



# Experimental Methodology

- Use collected data and PoP path data (from iPlane) to simulate recovery from detected failures
  - Random and informed choices of intermediaries



# Experiment

- $\approx 55,000$  path outage events
- Mean outage duration:  $\approx 22$  minutes
- Median outage duration: 2 minutes
- Longest outage: over 6 days
- Aggregate path outage probability: 1.48%
- Each client had access to between 83 and 97 intermediaries





# Improving Path Outage

- Impact of intermediary selection methods on path outage

Number of intermediaries	Random (SOSR)	Common PoP count	Common link count
0	1.48%	1.48%	1.48%
1	1.30%	0.76%	0.74%
2	1.15%	0.65%	0.65%
3	1.02%	0.57%	0.57%
4	1.02%	0.52%	0.53%
5	1.02%	0.52%	0.53%



# Improving Path Outage

- Impact of intermediary selection methods on path outage

Number of intermediaries	Random (SOSR)	Common PoP count	Common link count
0	1.48%	1.48%	1.48%
1	1.30%	0.76%	0.74%
2	1.15%	0.65%	0.65%
3	1.02%	0.57%	0.57%
4	1.02%	0.52%	0.53%
5	1.02%	0.52%	0.53%

- Informed methods reduce outage probability by **half!**



# Outline

- Introduction and motivation
- Previous work
- Our proposal
- Experimental validation
- Discussion
- Conclusion



# Discussion

- Why did informed selection outperform random?
  - Paths not independent
- Reasons for path correlations
  - PlanetLab design: universities on Internet2
  - Geography
- No knowledge of failure locations in study
- But can examine how close the sets of responding intermediaries are to what is *expected*
- If paths from intermediary set  $A$  are independent:

$$P(X \in A) \approx P(X \in A \mid X \text{ has valid path})$$



# PlanetLab Design

- On average across failures, for intermediary X:
  - $P(X \text{ at university}) = 92.0\%$
  - $P(X \text{ at university} \mid X \text{ has valid path}) = 91.6\%$
- Connectivity failures not correlated between universities



# Geography

- On average across failures, for intermediary X:
  - Paths originating in the US
    - $P(X \text{ in US}) = 71.6\%$
    - $P(X \text{ in US} \mid X \text{ has valid path}) = 59.8\%$
  - Paths originating in Asia
    - $P(X \text{ in Asia}) = 7.2\%$
    - $P(X \text{ in Asia} \mid X \text{ has valid path}) = 3.7\%$
  - Path originating in Europe
    - $P(X \text{ in Europe}) = 20.8\%$
    - $P(X \text{ in Europe} \mid X \text{ has valid path}) = 26.7\%$
- Thus, evidence of geographic correlation



# Path Outage revisited

- Effect of restricting random selection to intermediaries outside of vantage point geographical area

Num intermeds	Random (SOSR)	Common PoP count	Common link count	Geog. random
0	1.48%	1.48%	1.48%	1.48%
1	1.30%	0.76%	0.74%	1.26%
2	1.15%	0.65%	0.65%	1.09%
3	1.02%	0.57%	0.57%	0.95%
4	1.02%	0.52%	0.53%	0.84%
5	1.02%	0.52%	0.53%	0.75%



# Conclusions

- Demonstrated potential improvement using informed detour selection in reliability over random selection in SOSR
  - Biasing SOSR random selection outside own geographical area captures some benefit of informed method
- Future work
  - Implement informed detour selection mechanism
  - Evaluate path estimation services (Rocketfuel, iPlane)
  - Examine whether fresher path information helps





# Thanks!

## Questions/Comments?

