


Multipath Transport

D. Towsley, Y.-s. Lim, C. Liu
Umass-Amherst


Traditional single path approach

- ❑ requestor uses one path from content source
- ❑ congestion control (TCP) provides "some" robustness to traffic changes
- ❑ does not handle faults, mobility, attacks



Multi-source data transport

- provide requestor "set" of sources (NAs)
- network provides paths
- requestor balances load across sources/paths
- exhibits desirable load balancing properties
- robust to failures, mobility, attacks
- responsive to prices



Approaches

- independent path/source control
(independent TCP connections)
 - BitTorrent model
- coordinated/path source control
 - better load balancing
 - fairness to single path sessions
 - controller designs exist based on fluid models
 - capacity increases as #paths/sources increases

Path reselection

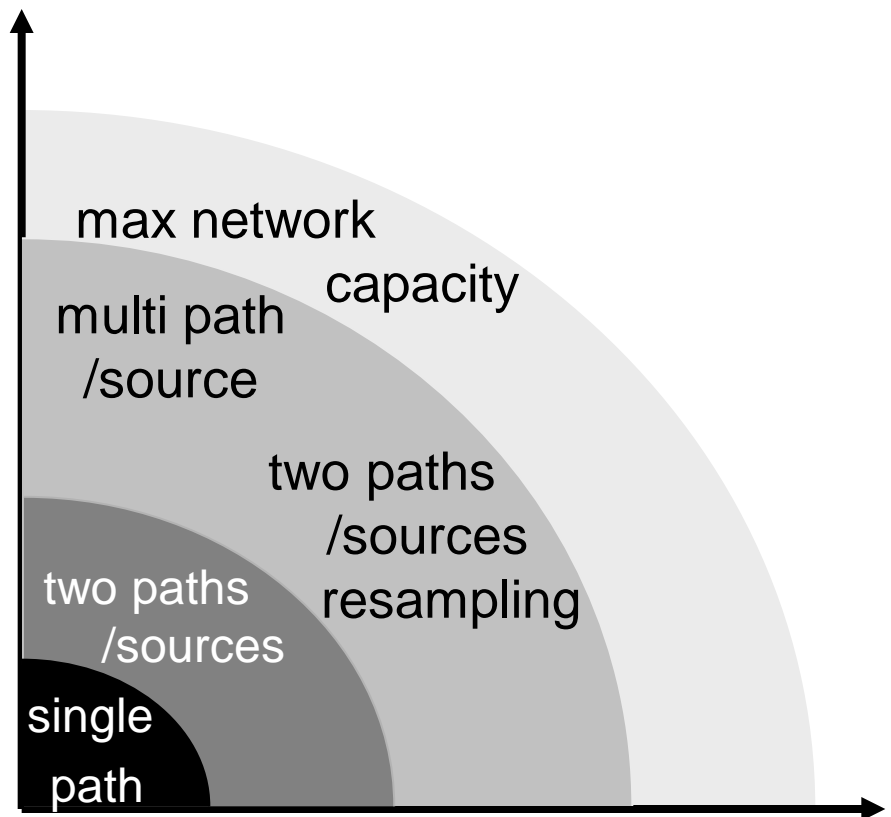
- the more paths/sources the better but ...
do we need to use them all?
 - 2 per session
 - periodically add new path/source at random
 - drop worst one

Theorem: Under random path/source reselection, coordinated multi-transport achieves same capacity as when using "all" paths/sources

Performance under multi-transport

- capacity increases with number of paths/sources per session
- two + random resampling achieves same capacity as using all paths/sources

(Key, etal. 2007,2011)



Recent work

- controller design: issues
 - small # flows => low utilization
 - path/source flappiness
- proper choice of increase/decrease rules

- receiver-based control
 - needed for multiple sources

Future directions

- true multi-source transport
 - coordinated receiver control
- short flows vs. long flows
- integration with GNRS
- increase robustness to path breaks
- experimentation on multi-homed systems