Leveraging Diversity for Resiliency

Roch Guerin
University of Pennsylvania

Acknowledgment

- This is based on joint work with
  - Sangeetha Abdu Jyothi
  - Behnaz Arzani
  - Alejandro Ribeiro

and as usual, mistakes are all mine

- <u>References</u>

[1] B. Arzani, R. Guerin, and A. Ribeiro, "A Distributed Routing Protocol for Predictable Rates in Wireless Mesh Networks." Proc. IEEE ICNP 2012

[2] S. Abdu Jyothi and R. Guerin, "Characterizing Internet Path Diversity." Under preparation.

This work was supported by NSF Grant CNS-1116039

The Premise

- Most communications require an uninterrupted path from source to destination
  - Individual paths are sensitive to disruptions affecting any one of their components
  - The availability of multiple paths can mitigate the impact of disruptions
- Routing is the network mechanisms that is primarily responsible for the discovery and usage of multiple paths
  - Either reactive or proactive
- **Question**: Can we design a communication infrastructure and routing protocols capable of systematically leveraging multiple path to improve resiliency (to failures & attacks)?
- **Focus**: Two extremes – the broader Internet and wireless mesh networks (the *"middle"* is much easier to handle)
  - Scale and complexity of Internet policies
  - Unpredictable nature of the wireless channel

Internet-scale Multipath

- First question: Is it there?
  - What topological diversity in the Internet?
  - How well can Internet routing protocols and policies exploit that diversity?
- Assessing Internet path diversity
  - Internet topology map (various source, *e.g.,* CAIDA, Cyclops)
  - Characterizing path diversity on a ~30k+ nodes and ~100k+ links graph

Characterizing the Internet Map

- A thick waisted Internet

Tier	# ASes	%
1	14	0.04%
2	12,397	31.7%
3	17,895	45.9%
4	4,944	12.7%
5	529	1.4%
6	102	0.26%

- A long-tailed connectivity

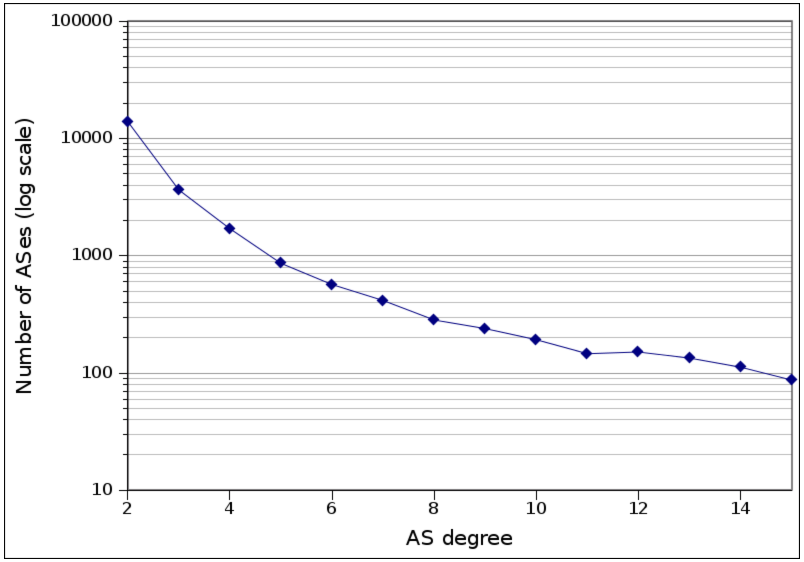

	Tier 1	Tier 2	Tier 3	Tier 4	Tier 5	Tier 6
degree	1394.6	9.4	3.2	2.3	2.2	1.5

More on AS Reachability/Connectivity

Mean number of ASes reachable in 5 hops or less

Tier	Customers	Peers	Providers	Total
1	30444	35508	0	35812
2	353	5325	36021	36025
3	56	2789	35423	35445
4	36	1793	30782	30895
5	2	1468	13213	13887
6	1	779	1541	2296

Internet Path Diversity

- Depends primarily on number of *providers*

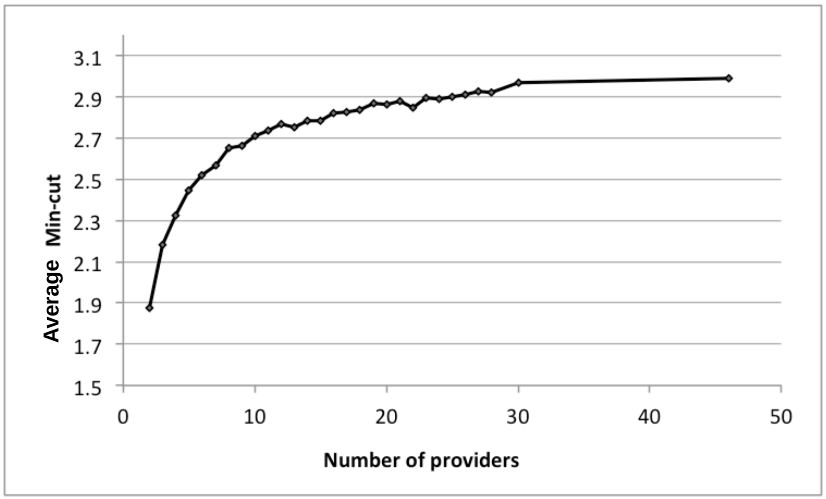

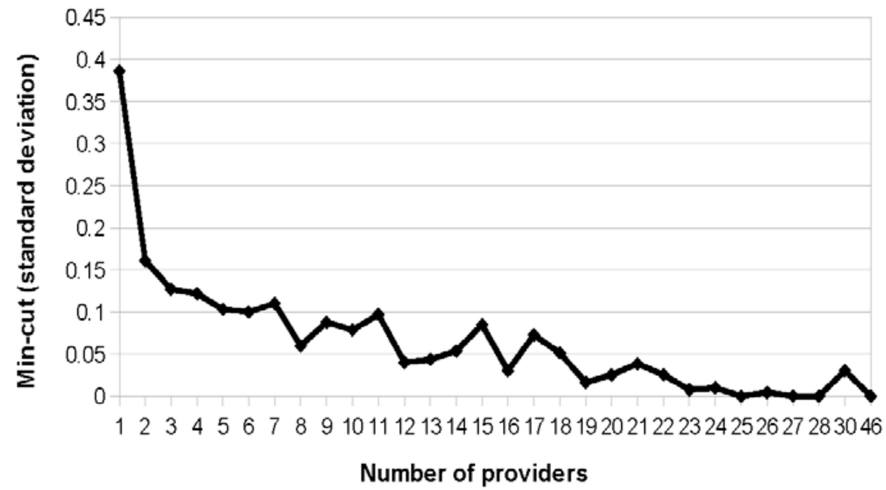

and is not far from the best feasible given connectivity constraints

- Adding a provider, *pretty much any provider*, improves path diversity, though with a law of diminishing returns
  - Adding a 3rd provider improves min-cut by about 0.3 on average, but adding a 4th only yields an average improvement of 0.13

Conclusion and Challenges

- There is already quite a bit of path diversity in the current Internet, and this in spite of the many constraints that BGP policies impose
- However, this does not mean that it can be efficiently exploited to ensure resiliency
  - BGP is a single path protocol
  - BGP path exploration can take a *long* time to converge and “switch” to an alternate path
- Ideally, we should augment BGP to maintain multiple *active* paths, but
  - Need to preserve scalability
  - Maintain consistency with routing policies
  - Ensure backward compatibility and incremental deployment

Wireless Multipath

- The nature of wireless links means that multiple paths are pretty much a given in wireless mesh networks
- The challenge is that wireless resources often behave erratically
  - √ Small time-scale (few μsecs to few msecs): fading & intereference
  - – Meso time-scale (10msecs to few secs): shadowing and channel gain variations
  - √ Large time-scale (secs to mins): hard failures and configuration changes
- Effective mechanisms exist to overcome small and large time-scale variations, but meso time-scales are harder to handle
  - We can deploy a first responder network, but it may not be able to provide predictable communications (reasonably stable end-to-end transmission rates)
- Multipath solutions can help minimize transmission rate variations
  - Computing paths with target average rate guarantees while minimizing rate variance
    - Multipaths are used ***jointly***
  - A distributed optimization based only on local link information

Rate Maximization vs. Rate Stabilization

(Best Path vs. Most Stable Multipath)

Instantaneous rates realization

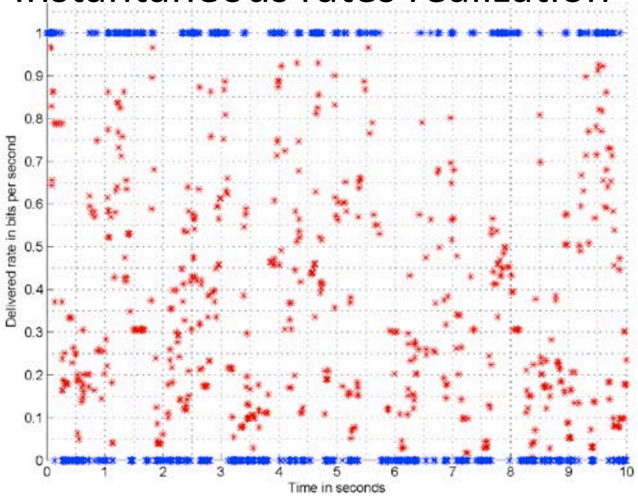

Instantaneous rates distributions

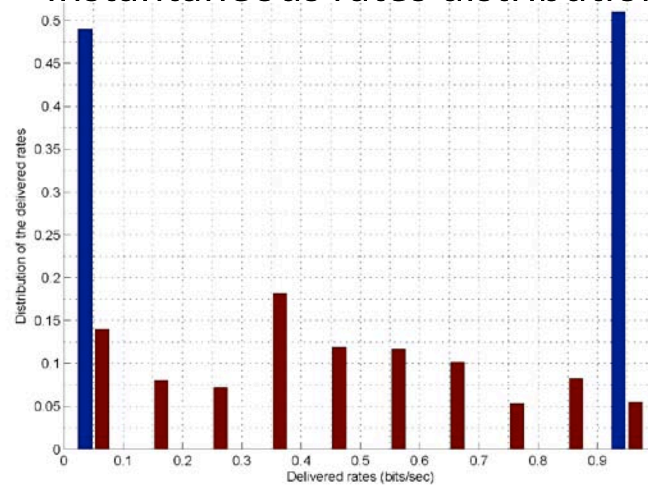

Mean rate penalty

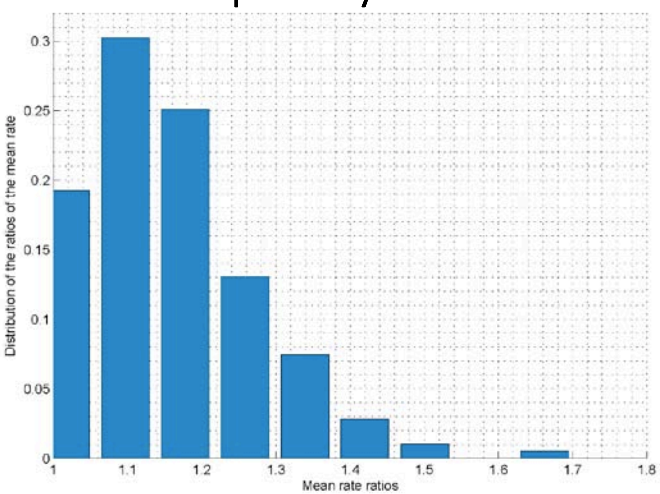

Rate variance benefit

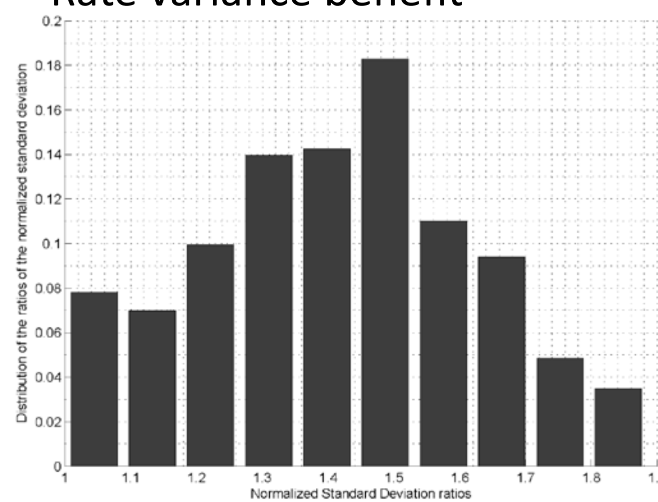

Dead Channel Time Distribution

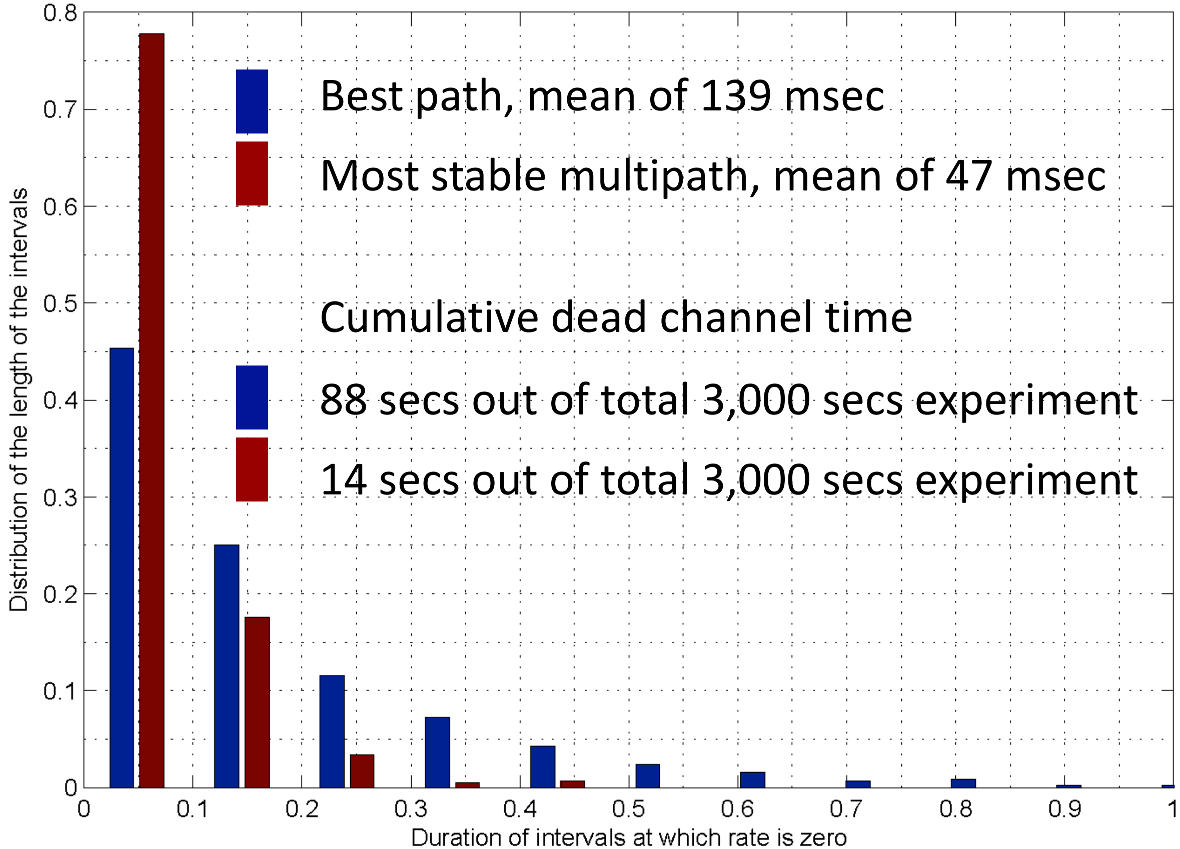

Conclusion and Challenges

- Multipath solutions that account for the inherent variability of wireless channels can be computed in a distributed fashion using only local information
  - Basic link statistics and a distributed computation that can easily piggyback standard shortest path computations
- Benefits include significant improvements in rate stability, including reduction in dead channel time, at the cost of only a slight reduction in channel throughput
- Practical deployments will, however, require
  - Tight integration with channel statistics estimation procedures
  - Effective recomputation triggers (when to ride-out variations vs. acknowledging significant changes) that are closely coupled to routing loop prevention mechanisms

More Generally

- Whether at the scale of the Internet or in a local mesh network, multipath solutions can improve resiliency if we are able to *simultaneously* use multiple paths
  - Long reaction time to changes at an Internet-scale
  - Relatively short-lived fluctuations in a mesh network
- This raises multiple challenges that can realistically only be solved by *jointly* involving the network and end-systems
  - Most applications, and in particular TCP, don't do well when packets are spread "blindly" over multiple paths

“Blind” Multipath TCP

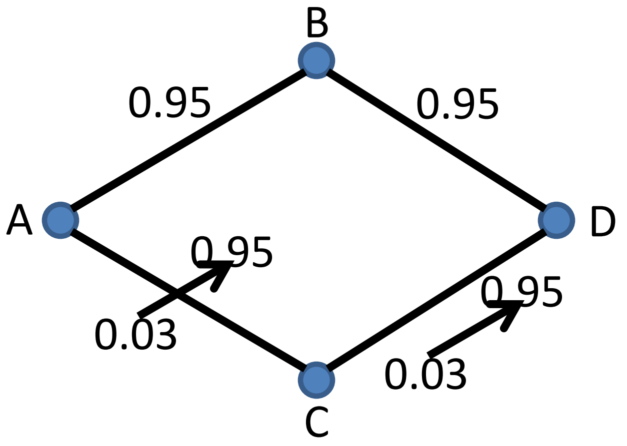

- A-D traffic split based on ratio that achieve minimum rate variance

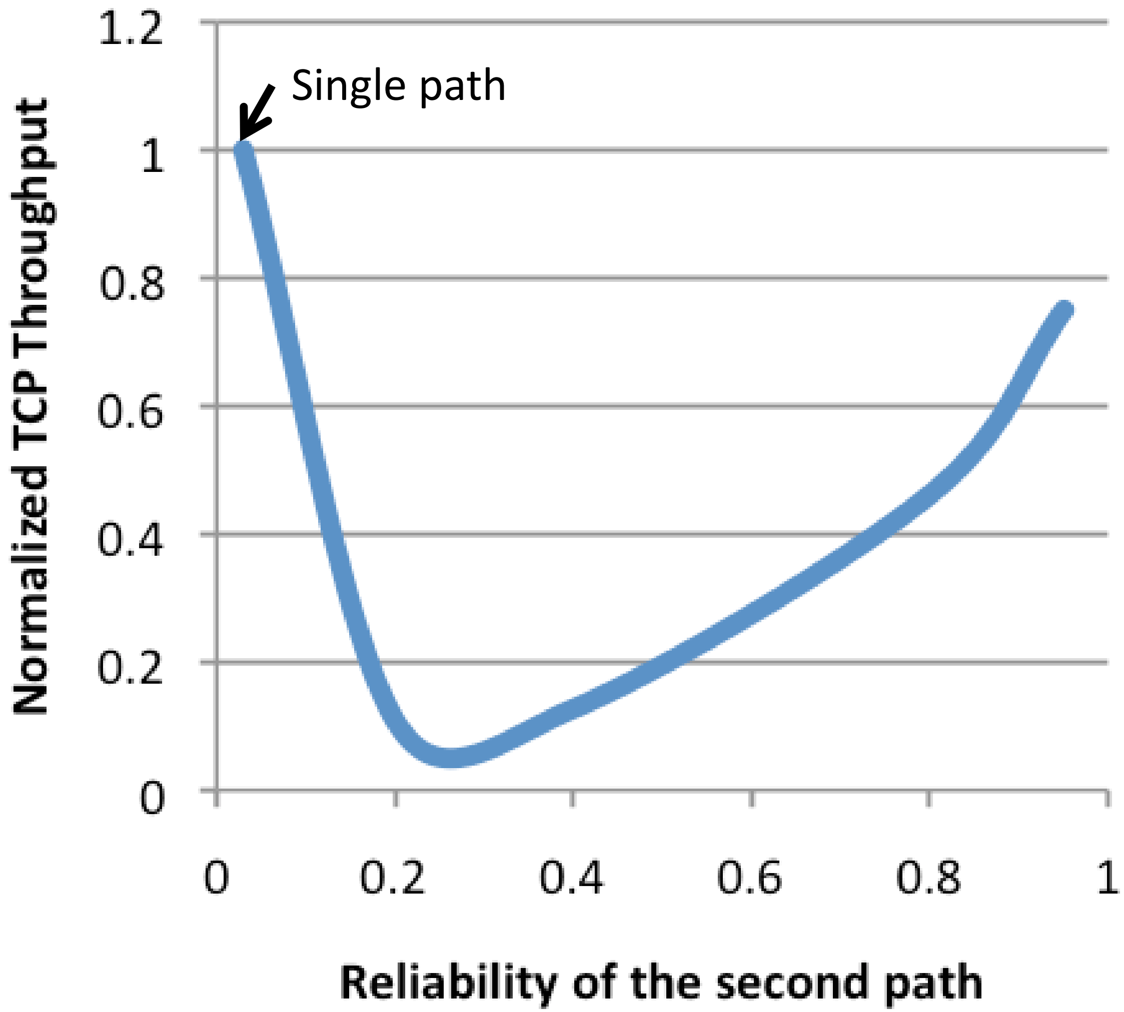

More Generally

- Whether at the scale of the Internet or in a local mesh network, multipath solutions can improve resiliency if we are able to *simultaneously* use multiple paths
  - Long reaction time to changes at an Internet-scale
  - Relatively short-lived fluctuations in a mesh network
- This raises multiple challenges that can realistically only be solved by *jointly* involving the network and end-systems
  - Most applications, and in particular TCP, don't do well when packets are spread "blindly" over multiple paths
  - Some form of coding, and in particular *in-network* coding, can help alleviate those issues
  - But even coding wont help if we keep using bad paths for too long, *i.e.,* we need the network to intelligently distribute packets across paths, or at the very least to stop using bad paths reasonably quickly