

Adaptive Source Routing

Oliver Michel, Ashish Vulimiri, P. Brighten Godfrey
University of Illinois at Urbana-Champaign



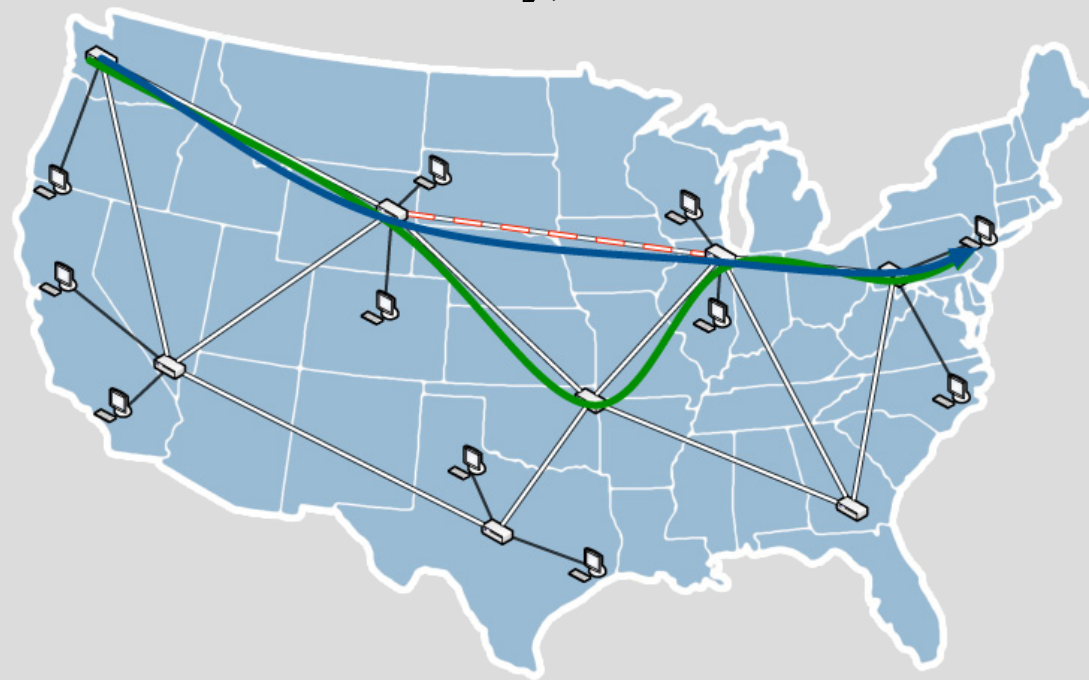
ILLINOIS

Motivation

- Today's routing techniques on the Internet rely completely on decisions taken within the network
- Lacking an end-to-end view, today's strategies often react slowly to dynamics and do not take into account the type of traffic routed
- Let sources define the path along a packet traverses the network!

Strategy

- Make routing decisions based on continuous end-to-end measurements to optimize end-to-end performance
- Sources can optimize different metrics such as latency, loss rate or throughput
- Continuous probing let sources detect outages or decreased performance immediately and choose a different path without waiting for the routing protocol to re-converge



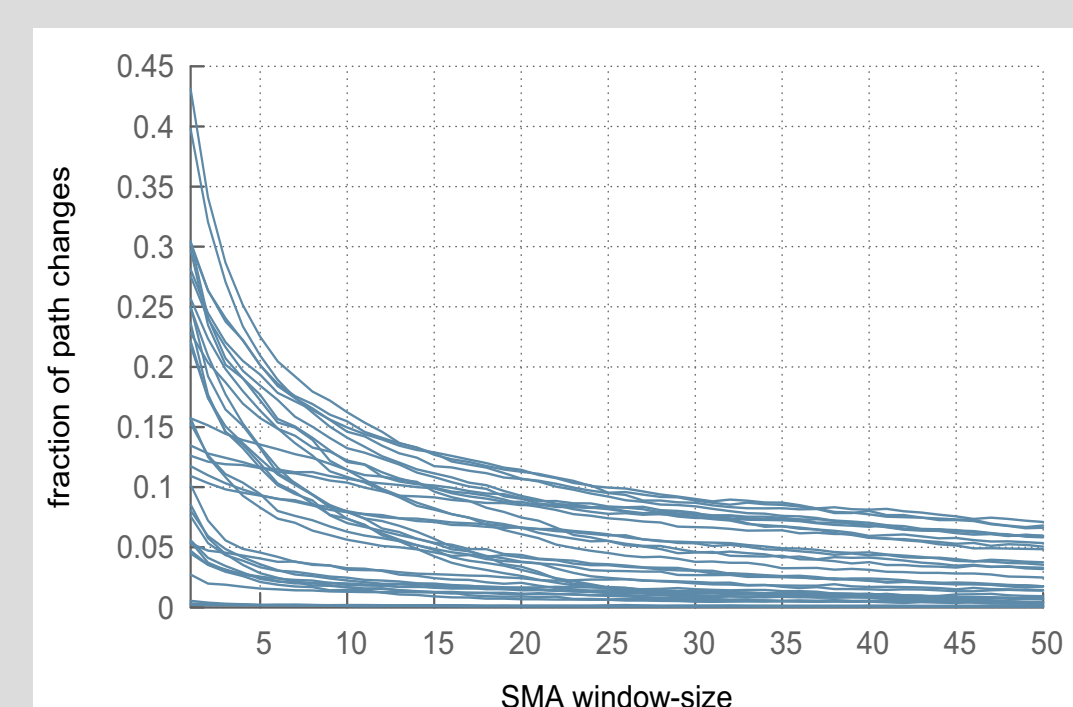
Redundant Multi-Path Optimization

- Duplicate packets over multiple best performing paths simultaneously
- Only consider the packet which arrives first
- Overhead for certain use-cases reasonable (e.G. TCP handshakes, DNS requests, API calls)
- Often corrects previous "incorrect" path-decisions
- Fast adaption yields to virtually no packet loss in case of path failure

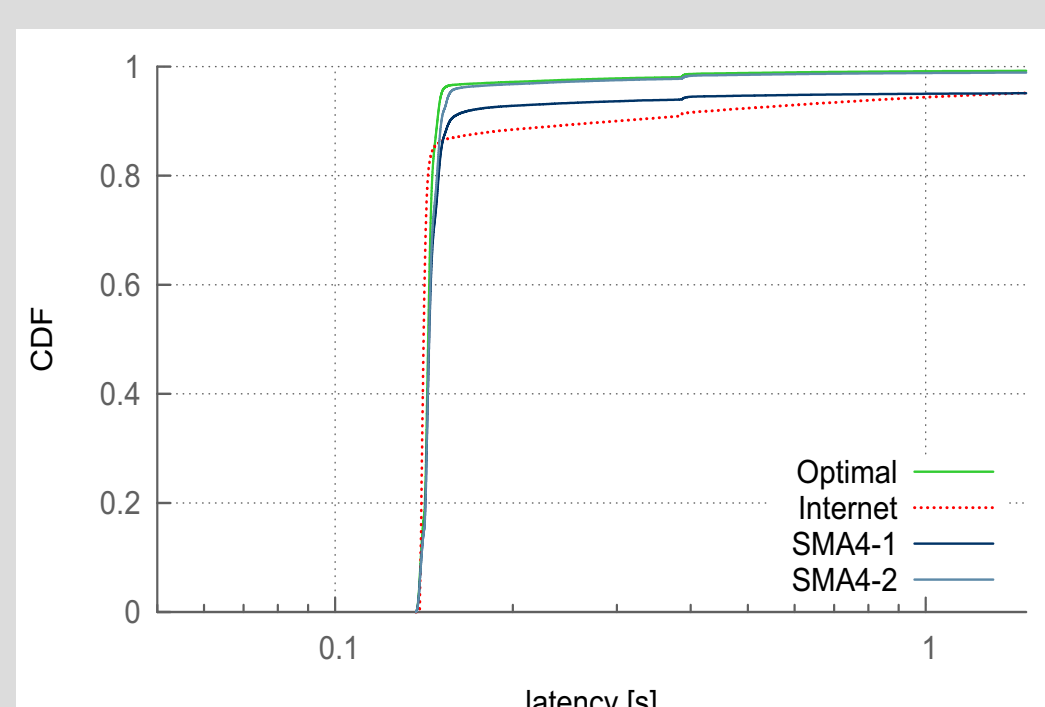
Prototype System

- Implementation of end-host software continuously monitoring path performance and selecting best paths based on simple moving average
- Set of best paths used for actual data-transmission in different traffic scenarios and data-rates (e.G. DNS Requests, VoIP)
- Once a path is selected for actual transmission, data packets are also used as probe-packets yielding higher probing resolutions

Trace Analysis



Fraction of path-changes (y-axis) when using SMA window sizes up to 50



CDF of latencies seen in PlanetLab overlay after applying SMA on trace data with 1 and 2 simultaneous paths

- Collection of extensive trace data on PlanetLab to determine reasonable path-selection strategies on both overlay and layer-2 deployments

Round-based algorithms to determine good paths

- simple average over all seen latencies (AVG)
- simple moving average (SMA)
- exponentially weighted moving average (EWMA)
- Use best performing path from previous round (PREV)

Live Run

- Deployment of Prototype on GENI's OpenFlow core with using of overlay paths at the same time
- Sender/Receiver Software on ProtoGENI nodes at Stanford/GPO
- Transmission of actual payload data
- Live-Run of SMA algorithm with window-size 5 while still collecting trace-data for later offline-analysis

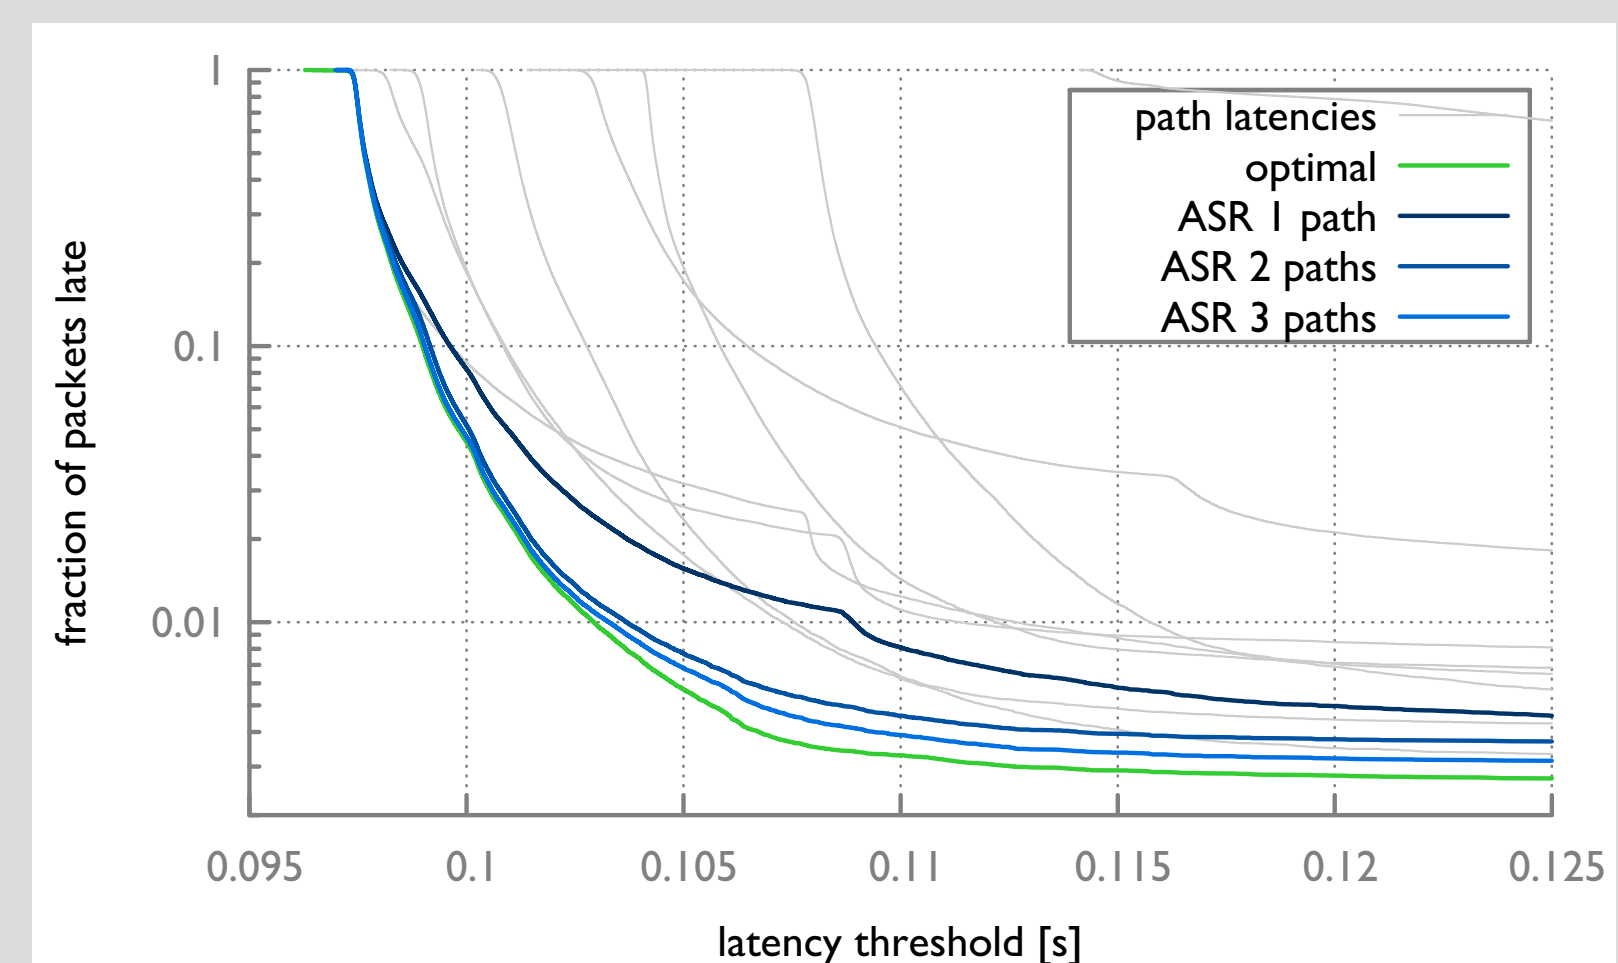
Results

Trace Analysis


- SMA and PREV perform best in the average case
- Incorrect path decisions are made when a path's latency is bad but it's packet loss is minimal and vice versa

Redundant Multi-Path Optimization

- Improves overall-performance dramatically especially in the latency-distribution's tail by achieving close-to optimal results
- Prevents packet-losses up to a high level even when using only two different paths



Latencies seen in live-run on mixed L2/overlay deployment for 48h using SMA (window size = 5) for path selection and different duplication levels (data rate = 32 kbit/s)



Loss rates seen in live-run over 48h compared to simultaneously measured direct-internet path for verification and reference (data rate = 32kbit/s)

Conclusion

- Positive verification of previous results (presented at GEC8/9) from PlanetLab on L2-deployment reducing unpredictability and actual live run instead of trace analysis only
- Especially for slow packets better performance compared to average of all available paths
- Continuous probing allows dynamic adaptability to decreasing path-performance and outages

Future Work

- Analysis of new traces collected on GENI in order to determine better path selection techniques than basic moving average in this prototype
- Optimization of various metrics for different traffic-types (e.G. large downloads may optimize for throughput and not for latency while real-time applications reduce latency)
- Further Investigation on the threshold when it is reasonable to duplicate traffic along paths