



Robust LFA Protection for Software-Defined Networks (RoLPS)

Daniel Merling, **Steffen Lindner**, Michael Menth

<http://kn.inf.uni-tuebingen.de>



- ▶ Motivation
- ▶ LFAs: State of the Art
- ▶ Robust LFA Protection for Software-Defined Networks (RoLPS)
- ▶ Evaluation
- ▶ Hardware Prototype



- ▶ Packet forwarding in networks is disrupted when a next-hop becomes unreachable
 - Link failure
 - Node failure


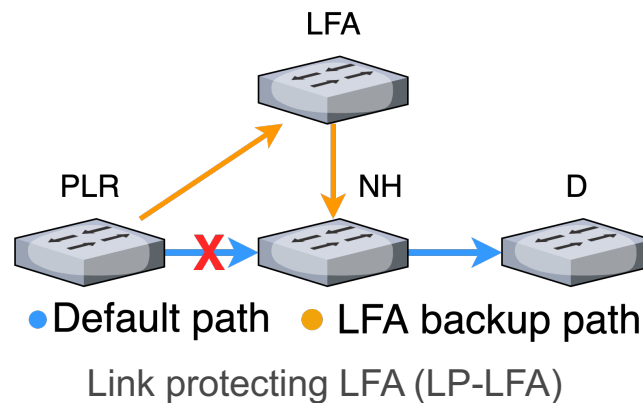
- ▶ Upon failure detection, IGP or controller recomputes forwarding rules
 - Failure detection & computation takes time

- ▶ Fast Reroute (FRR) mechanisms are used in IP networks to quickly reroute packets
 - Pre-computed backup paths are used while forwarding entries are recomputed

- ▶ Desirable: FRR in SDN without controller interaction
 - High coverage
 - Limited forwarding table sizes in SDN forwarding devices

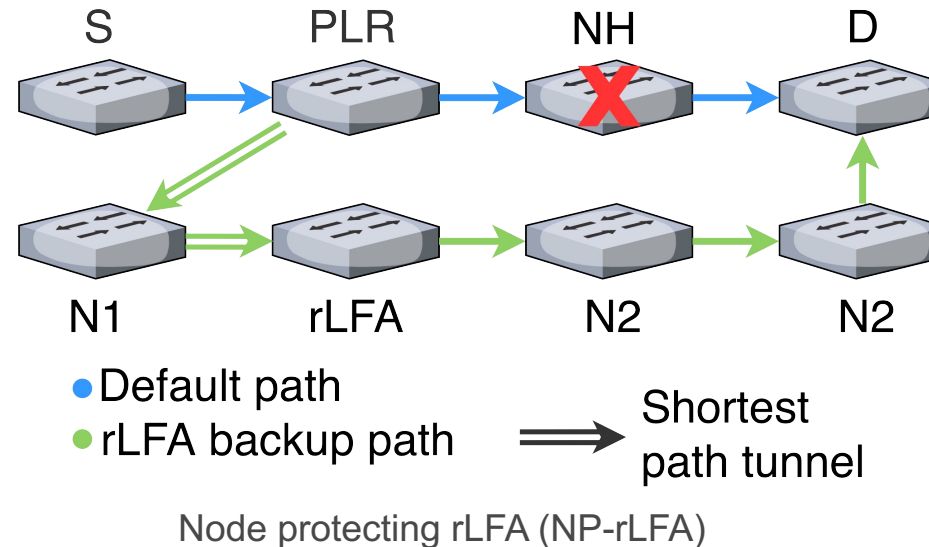
- ▶ Loop-free alternates (LFAs) are a well-known FRR method for IP networks
 - Traffic is sent to alternative next-hops without creating routing loops

- ▶ Two different protection levels
 - Link protection
 - Node protection





- ▶ Sometimes, no (LP/NP-)LFA is available for a given destination
- ▶ Remote LFAs (rLFAs) protect more destinations than LFAs
 - Based on shortest path tunnels to remote nodes





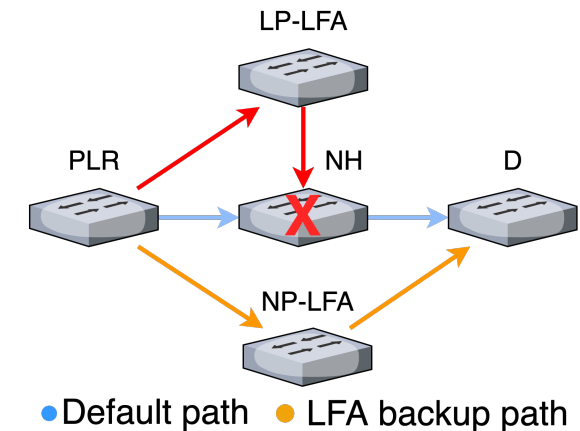
- ▶ Sometimes, even no (LP/NP-)rLFA is available for a given destination

- ▶ Topology-independent LFAs (TI-LFAs) leverage segment routing (SR)
 - Backup path is encoded as stack of forwarding actions in packet
 - Based on IPv6 (Srv6) or MPLS (SR-MPLS)



► (r)LFAs may create routing loops in failure scenarios

- LP-LFAs in node failure cases
 - PLR might be a LP-LFA for its own LP-LFA
- Double link failures
- Single link and single node failure






► Robust LFA Protection for Software-Defined Networks (RoLPS)

- Explicit LFAs
- Advanced Loop Detection
- LFA Ranking

► Explicit LFAs

- Based on explicit tunnels, e.g., unique IP addresses
- Uses rLFAs if available
- Multipoint-to-point tunnels for less forwarding entries



Link protecting eLFA (LP-eLFA)



► Advanced Loop Detection

- Packets should be dropped if they are rerouted more than n times
- Requires only a counter in the packet header
- Implementable in Openflow & P4

► LFA Ranking

- Controller classifies nodes for different PLRs into LP/NP-(e/r)LFAs
- LFAs can be ranked according to their
 - protection level
 - NP is better than LP
 - Complexity
 - Simple LFAs do not require tunneling / additional forwarding entries
 - eLFAs are most complex
- RoLPS ranks LFAs first according to their protection level

Rank	LFA Type
0	NP-LFA
1	NP-rLFA
2	NP-eLFA
3	LP-LFA
4	LP-rLFA
5	LP-eLFA

Table 1: Ranking of LFA types according to protection level and complexity. Preference is given to LFAs with lower rank number.



► Protection Variants

- Many different protection possibilities (loop detection, LP/NP, (e/r)LFAs)
- Naming scheme: {nLD, ALD}-{LP, NP}-{LFA, rLFA, eLFA}

Mechanism	C-LFA (nLD-LP-LFA)	C-rLFA (nLD-LP-rLFA)	LD-LFA (ALD-NP-LFA)	ALD-NP-rLFA	ALD-LP-eLFA	ALD-NP-eLFA
Loop detection			•	•	•	•
Protection against all SLF		o		o	•	•
Protection against all SNF						•
Additional forwarding entries					•	•

Table 2: Properties of protection variants.
Legend: o = only for unit link costs; • = independent of link costs.



- ▶ Performance Evaluation of LFA-Based Protection
 - Evaluated on the Internet topology zoo
 - 205 wide area, commercial, research, and academic networks

- ▶ Metrics of interest
 - Protection coverage
 - Additional forwarding entries



► Protected

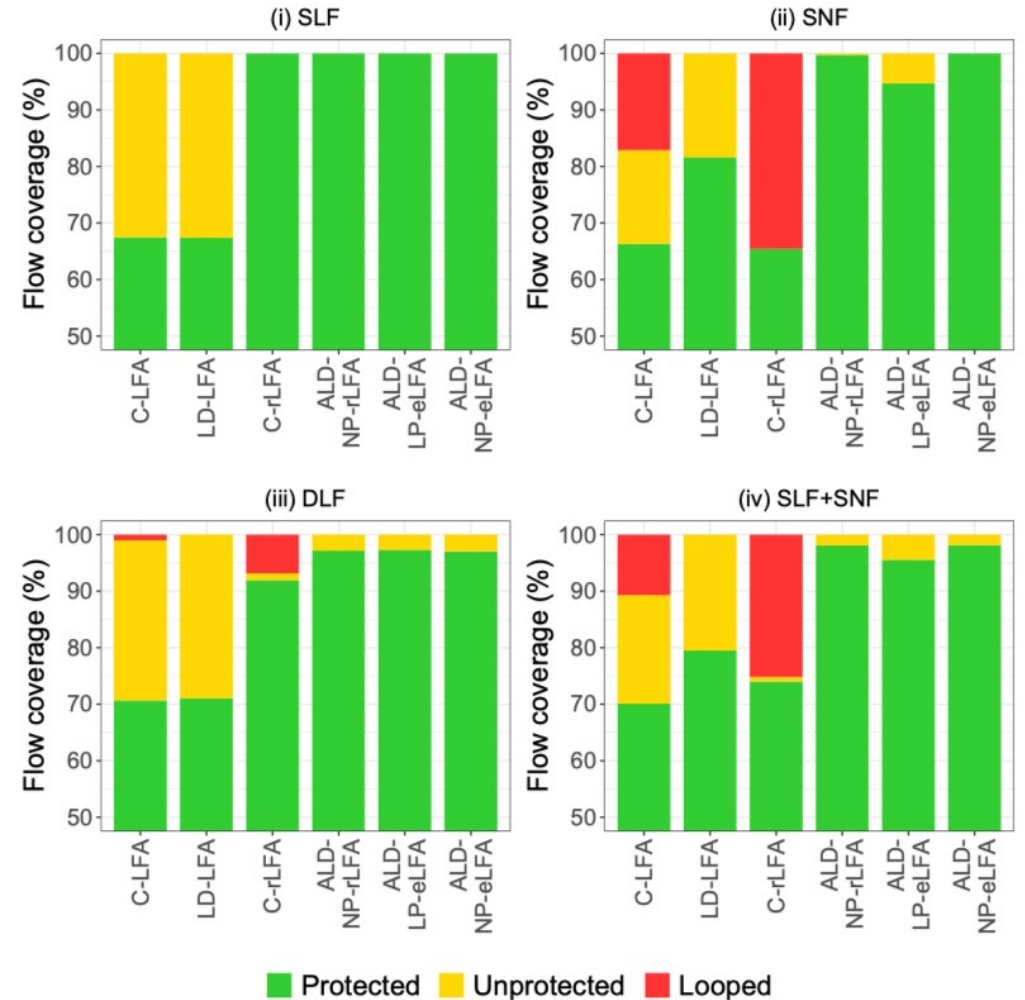
- Packet is successfully delivered
- Packet is dropped to prevent a loop

► Unprotected

- Packet is dropped although the destination is reachable

► Looped

- Microloop was caused by local rerouting





► Single Link Failure (SLF)

- Simple LFAs (C-LFA, LD-LFA) can not cover all single link failures
- (r/e)LFAs cover all single link failures

► Single Node Failure (SNF)


- LP-LFAs (C-LFA, C-rLFA) may result in routing loops
- Only ALD-NP-eLFA covers all destinations

► Dual Link Failure (DLF)

- Variants without loop detection (C-LFA, C-rLFA) may result in microloops
- Simple LFAs (C-LFA, LD-LFA) have lowest coverage

► Single Link + Single Node Failure (SLF + SNF)

- LP-LFAs (C-LFA, C-rLFA) may result in routing loops





► Additional Forwarding Entries

- LFAs and rLFAs are based on shortest paths → no additional forwarding entries
- eLFAs require additional forwarding entries
- Compared to MPLS-facility-backups (MPLS-FB- $\{LP, NP\}$)

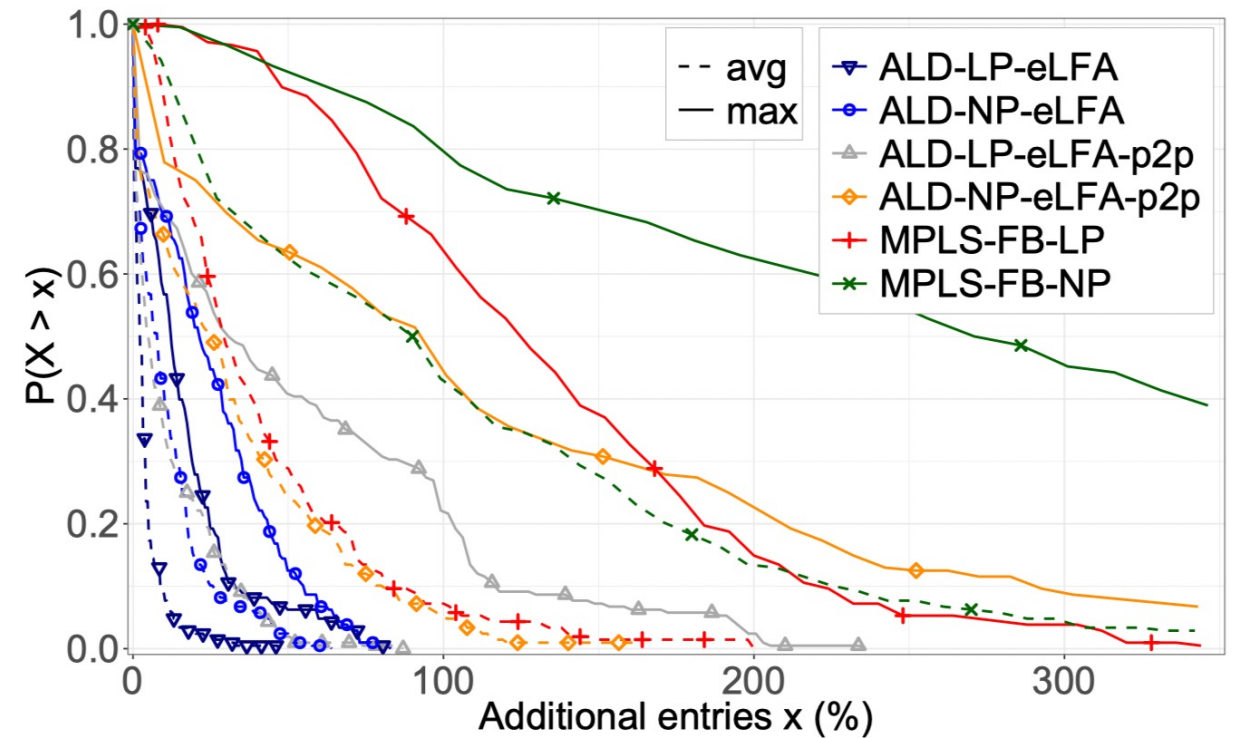


► MPLS-FB-LP

- 55% of networks have at least one node with 120% more additional entries
- 8% of networks have more than 100% additional entries on average


► ALD-LP-eLFA

- No topology with a node that requires more than 80% additional entries
- 95% of networks require less than 15% additional entries on average





- ▶ We implemented ALD-(e/r)LFAs in P4 for the Tofino ASIC with up to 3.2 Tbit/s throughput
 - Tofino generates a special packet when ports are up/down
 - We store this information in registers to apply FRR

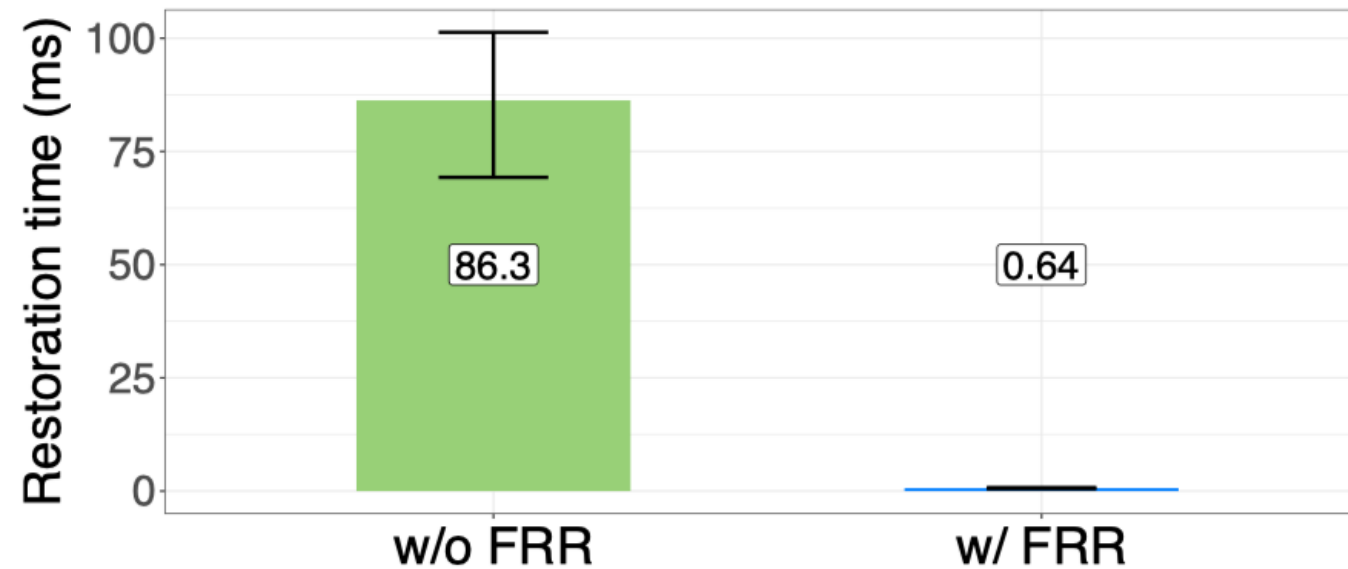


- ▶ RoLPS based SDN controller



► Restoration time

- Time until traffic is received after a failure
- With and without FRR





- ▶ RoLPS leverages (e/r)LFAs with advanced loop detection (ALD)
- ▶ Evaluation shows that existing (r)LFAs do not cover all destinations and may result in routing loops
- ▶ P4-based prototype that features RoLPS-based protection variants and runs at 100 Gbit/s
- ▶ Connectivity is restored in less than 1 ms



<https://github.com/uni-tue-kn/p4-lfa>



<https://ieeexplore.ieee.org/document/9461214>

Published @ TNSM (Special Section on Design and Management of Resilient Networks)

Steffen Lindner

University of Tuebingen

Faculty of Science

Department of Computer Science

Chair of Communication Networks

