



Scalable Resilience for Software-Defined Networking Using Loop-Free Alternates with Loop Detection


Wolfgang Braun and Michael Menth

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



Loop-Free Alternates (LFAs)

- Assumption: destination-based forwarding
- ▶ Locally reroute traffic in case of failures
 - Use fast-reroute function in OpenFlow
- ▶ Loop-free alternates (LFA)
 - Avoid loops for single link failures
 - But may cause loops for node failures and multiple failures
- ▶ Classification: LFAs that avoid loops for
 - Link failures (loop-free condition, LFC)
 - Node failures NPC (node-protecting condition, NPC)
 - Multiple failures (downstream condition, DSC)



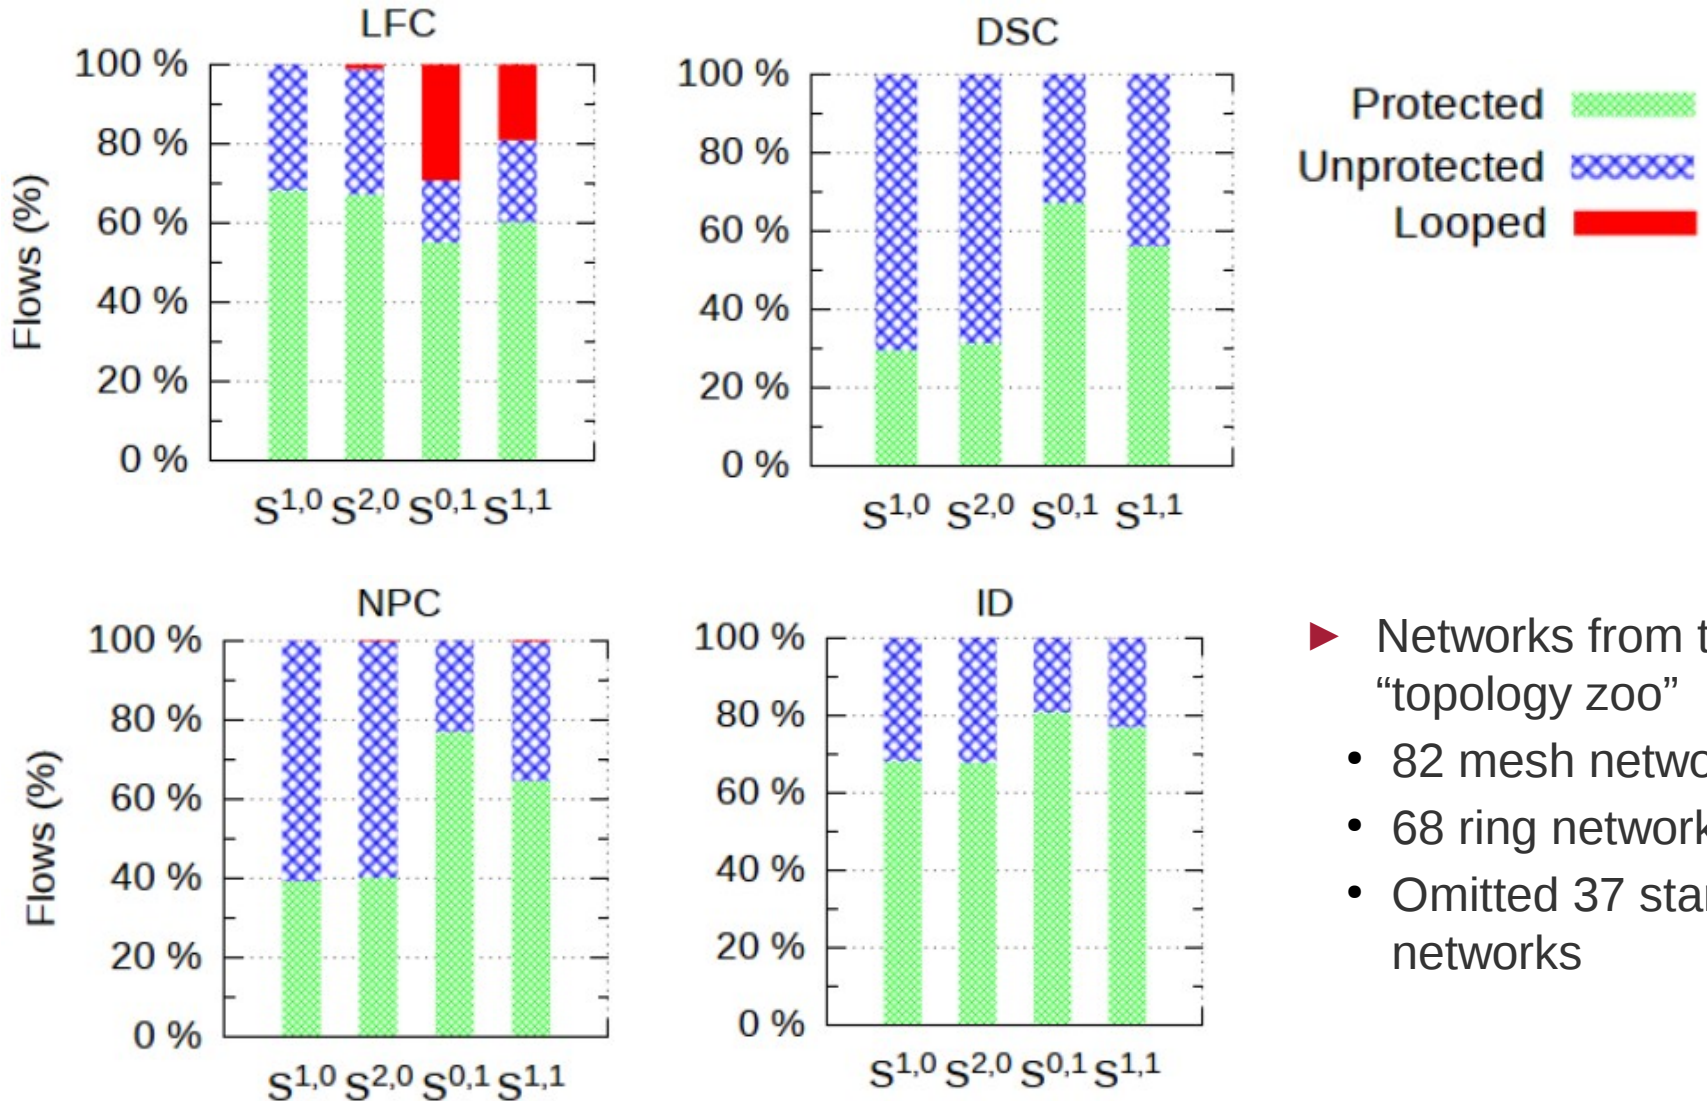


- ▶ Loop detection principle
 - If packet is rerouted, add node ID to packet header
 - Nodes drop arriving packets with their own ID in the header

- ▶ Implementation in OpenFlow (LFA-ID)
 - Assign IDs to nodes
 - Add ID label to packet header
 - Each bit in the ID label stands for a node ID
 - If pkt is rerouted, set node's bit in ID label
 - Add one flow entry to switch that
 - Drops pkts with own node ID in pkt hdr



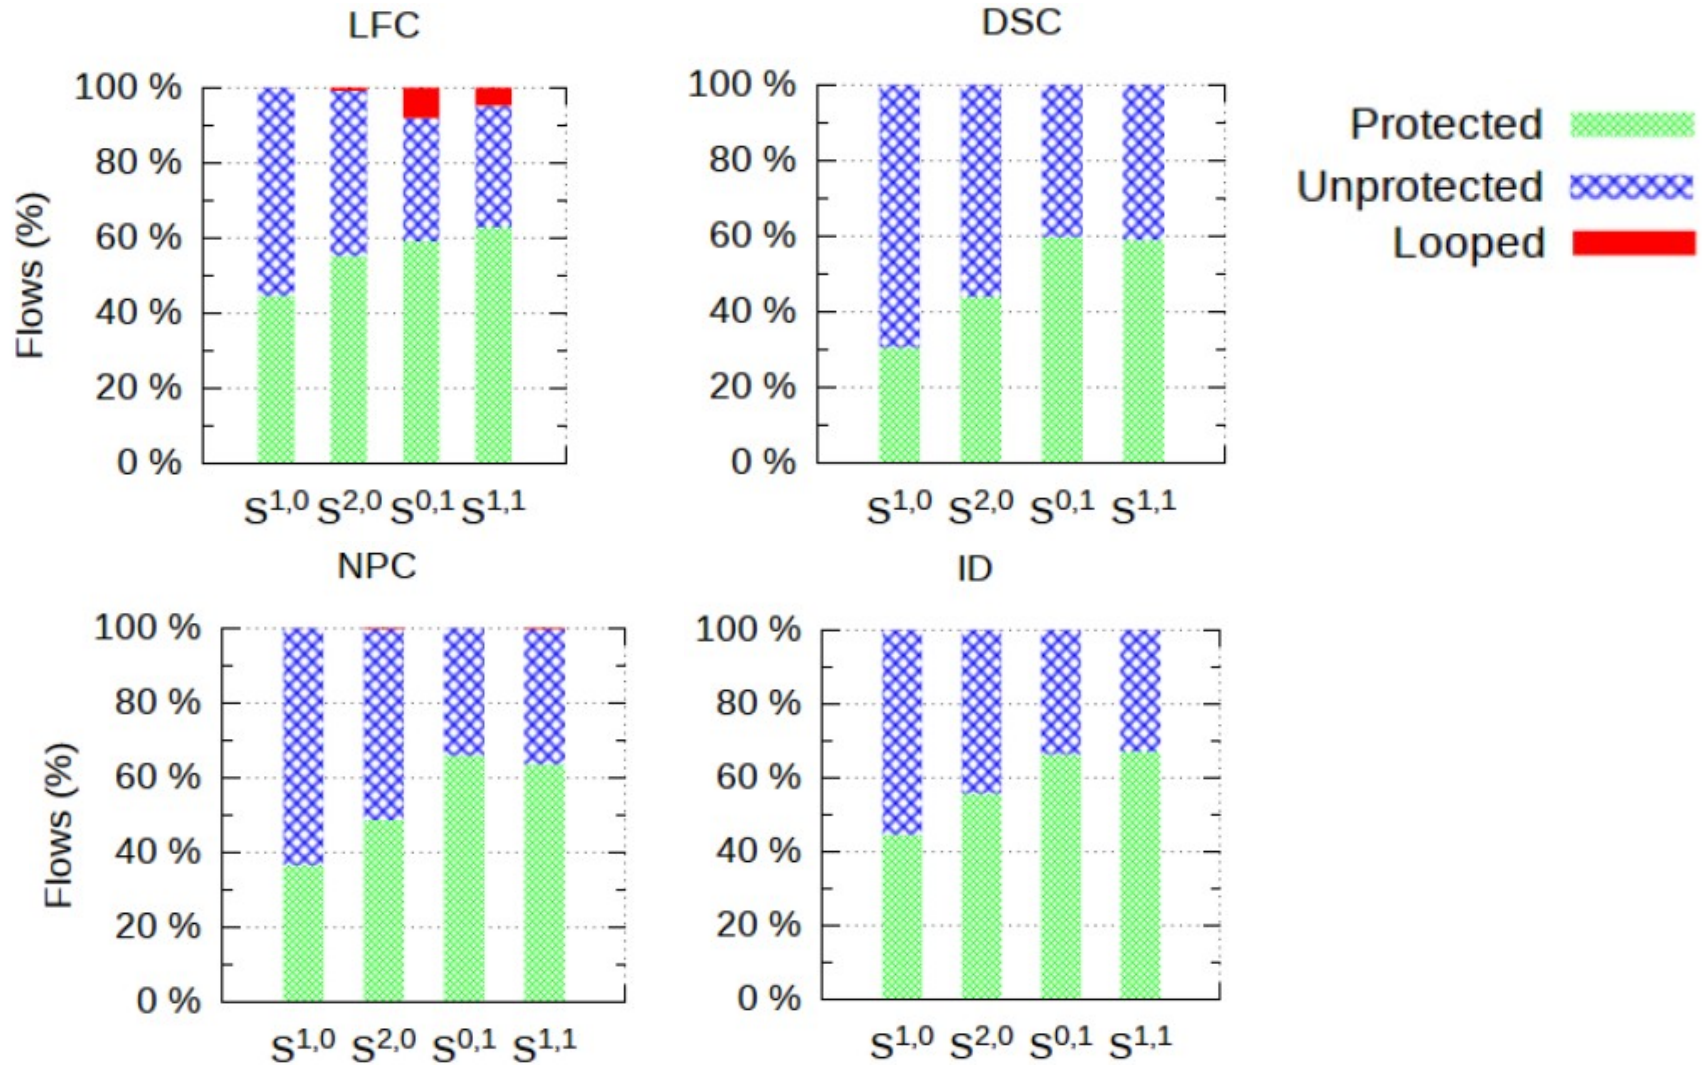
LFAs in Mesh Topologies



- ▶ Networks from the “topology zoo”
 - 82 mesh networks
 - 68 ring networks
 - Omitted 37 star networks




LFAs in Ring Topologies






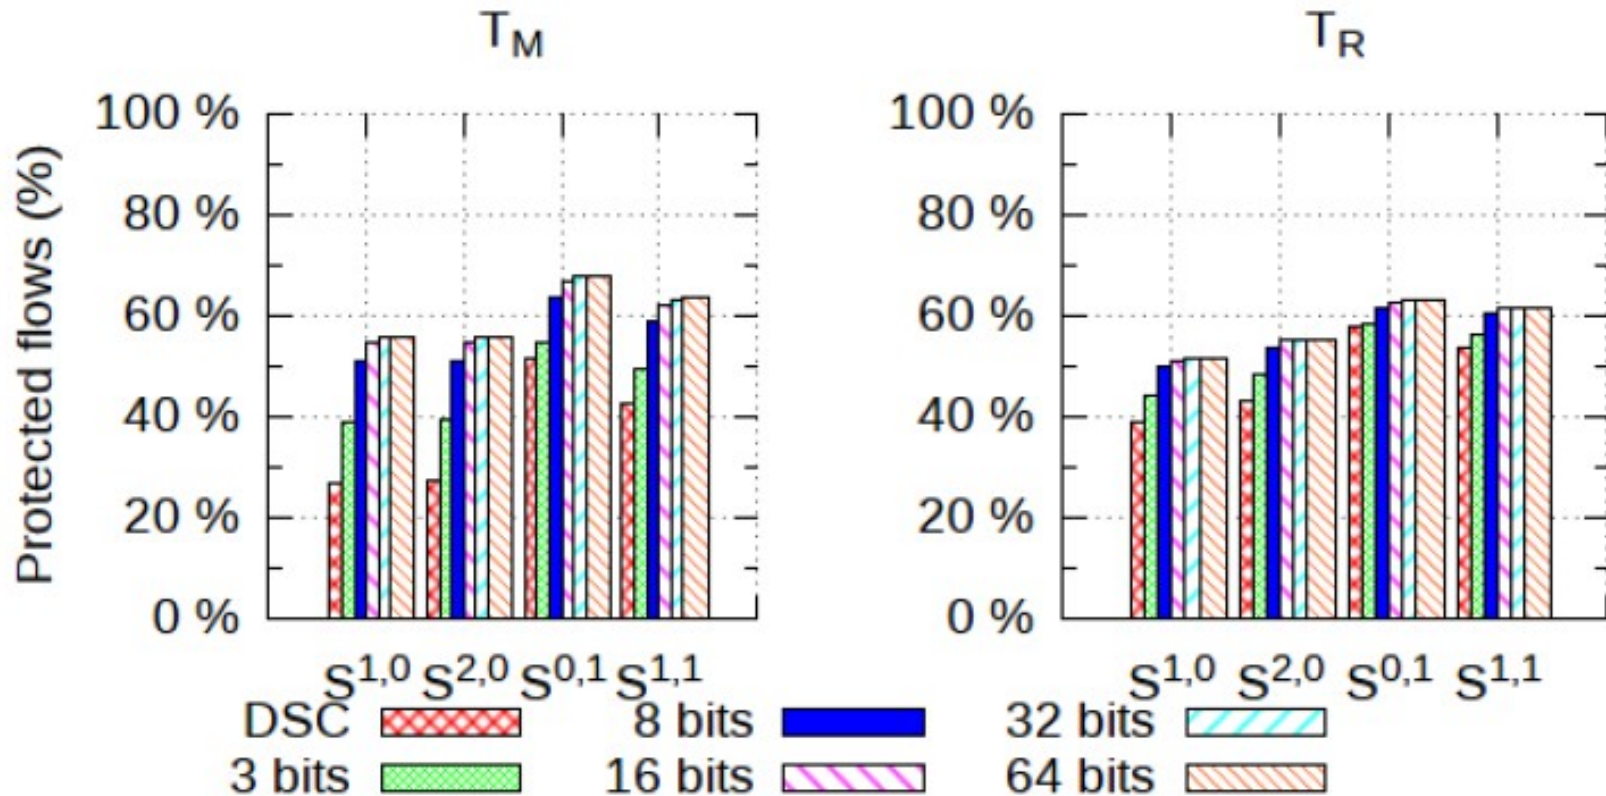
How to deal with large networks?

- ▶ Use large ID labels
- ⇒ Large overhead



- ▶ Use small ID labels and map one bit to multiple nodes
- ⇒ Packets may be dropped unnecessarily






- ▶ Results for 14 networks with 50 or more nodes
- ▶ Label size of 8 to 16 bits can be sufficient for large networks
 - MPLS labels may be used



	Fat-tree	BCube	DCell
LFC	70.8% - 81.9%	18.8% - 19.1%	21% - 50%
LFC-NPC	70.8% - 81.9%	18.8% - 19.1%	21% - 50%
LFC-DSC	70.8% - 81.9%	18.8% - 19.1%	0% - 28%
LFA-ID	70.8% - 81.9%	18.8% - 19.1%	21% - 50%



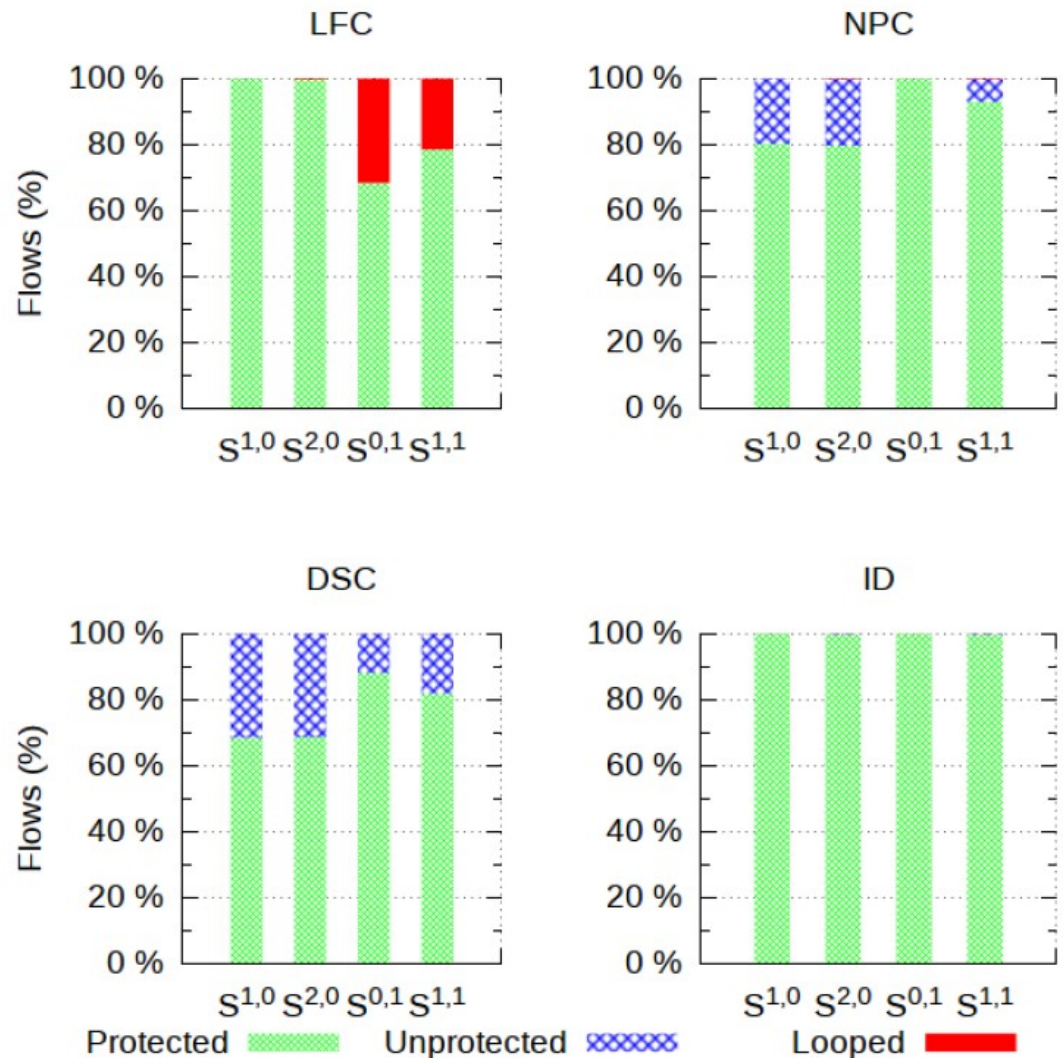
Link-Augmented Fat-Tree





LFAs in Link-Augmented Fat-Tree

- ▶ Full protection for single link failures with LFC
- ▶ But many loops in case of node failures and multiple failures
- ▶ LFA-ID almost protects 100% even for multiple failures (< 0.3% of flows cannot be protected)





- ▶ LFAs for OpenFlow-based SDN with loop detection
 - ID label
 - 1 flow entry per switch

- ▶ Coverage depends on network topology
 - Almost full failure coverage in link-augmented fat-tree topologies

- ▶ Future work: add remote LFAs
 - Complete coverage
 - Keep overhead low