

SHEAR – A Highly Available and Flexible Network Architecture

Marrying Distributed and Logically Centralized Control Planes

Michael Markovitch

Joint work with Stefan Schmid

Department of Communications Systems Engineering


Ben-Gurion University

TU Berlin Internet Network Architectures

Network Flexibility and Availability

- Current networks are hard to manage and complicate to configure
 - Low flexibility
 - Traffic Engineering not widely deployed
- Availability based on distributed protocols
 - Relatively long time to become stable after failure

SDN deployment?



Must upgrade to full SDN?

Current Situation & Related work

- Upgrade networks incrementally
- Abstraction to make a hybrid network look like a logical SDN
 - A hybrid network as a temporary step on the way to full SDN
- SDN controller which can affect legacy non SDN parts of the network
 - Workaround, not an architecture

[HK, 2014] [LKZXJ, 2015] [LCSSF, 2014] [LAZLRJ, 2013] [VVB, 2014]

Transitional networks: Challenges

- Replacement of all nodes in a network can take a very long time
 - Especially in LANs (and campus networks)
- SDN nodes not cheap
 - Further impeding full upgrade
- More complicated than just managing the old distributed network
 - The worst of two worlds

SHEAR - ARchitecture

- What it is:
 - An architecture for a hybrid network
 - Combination of distributed and centralized control planes
 - Not a workaround
- What it is not:
 - Not an SDN controller application
 - Not a logical abstraction of a full SDN
 - Such an abstraction can be applied


Hybrid control and data planes

- Hybrid data plane:
 - Small partial deployment of SDN nodes
 - Extract control over path diversity
 - Simplifies traffic engineering
- Hybrid control plane
 - Monitor network state with distributed protocols
 - React to events with centralized controller
 - No waiting for stabilization

Extract path diversity


- Make the legacy network simple
 - No multiple paths in legacy connected components
- Multiple paths through SDN nodes only
 - Locate SDN nodes to control paths
 - Minimum feedback vertex set problem (NP-hard)
- Use fundamental cycles to locate SDN nodes
 - Goal: no cycles remaining
 - Very fast

Fundamental Cycles




- Select node which is part of most fundamental cycles

Fundamental Cycles



- Repeat until no cycles remain

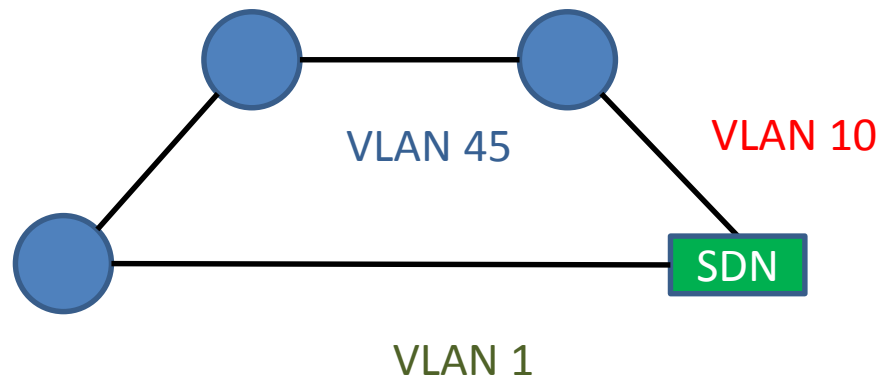
Fundamental Cycles



- Not the minimal number of SDN nodes
 - Advantage: Network divided

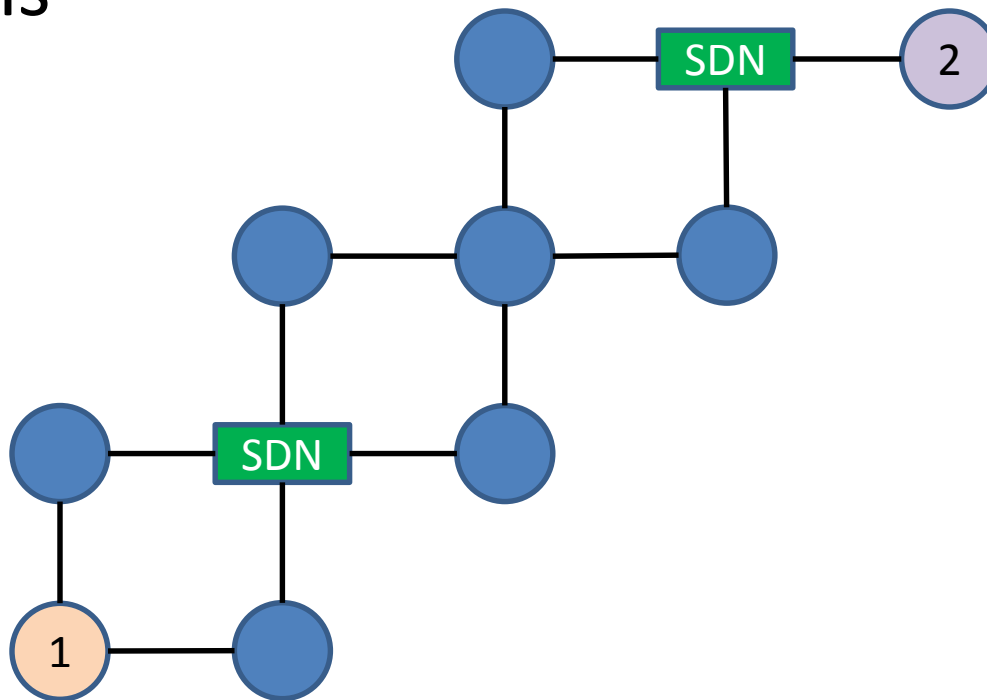
Pathlets

- Pathlet
 - Defined by combination of source port of SDN node and VLAN (defining the other end)
 - No loops – Pathlets are well defined in a connected component
- Every legacy node has at least one pathlet to a SDN node



Pathlet stitching


- Combine Multiple Pathlets to define end to end paths



End to End paths achieved through Combination of pathlets

Spanning trees

- Monitor the connectivity between nodes via Multiple Spanning Trees
 - Each spanning tree used to monitor group of links and nodes between root and a SDN mode
- Roots as far as possible from SDN nodes
 - **Contrary to current best practices**
 - Allows for fault localization



SHEAR recap

- Path diversity controlled by SDN nodes
 - Extracted from the legacy network
- End to end paths go through SDN nodes
 - Pathlet stitching defines paths
- Monitoring and fault localization
 - Via MSTP

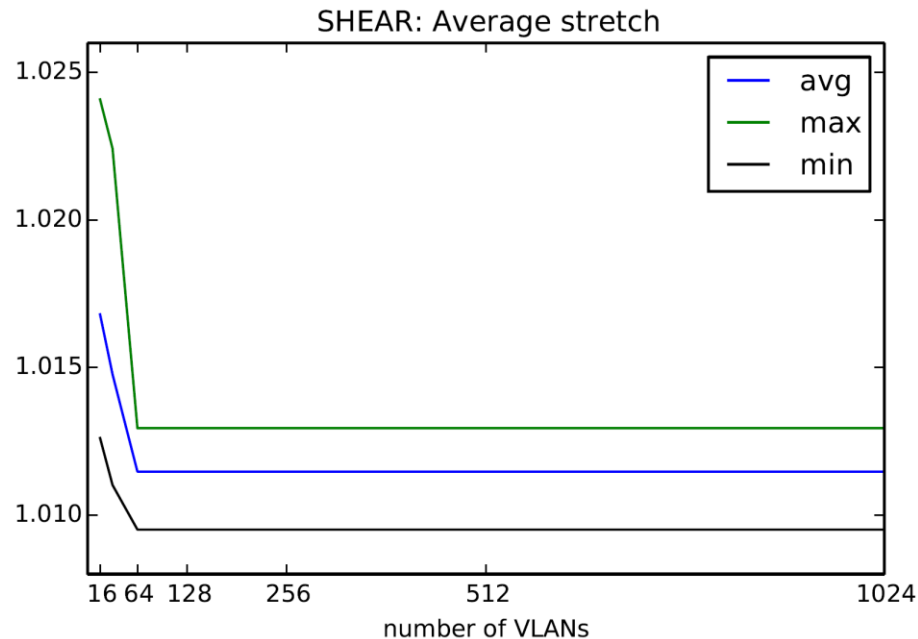
Simulations: Nodes needed

- Cisco topologies follow recommendations for high availability campus networks
- Enterprise is a real large campus network

Topology	Number of nodes	Nodes to upgrade	Ratio
Cisco medium	122	11	0.0901
Cisco large	242	20	0.0826
Enterprise	1577	29	0.0183


Simulations: Stretch

- For the real enterprise network with 29 SDN nodes, measure stretch factor of paths (against optimal paths)



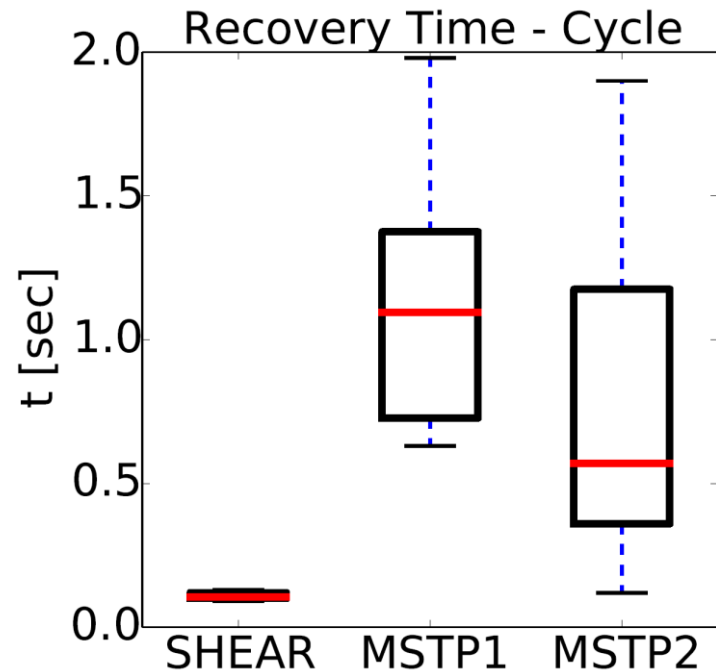
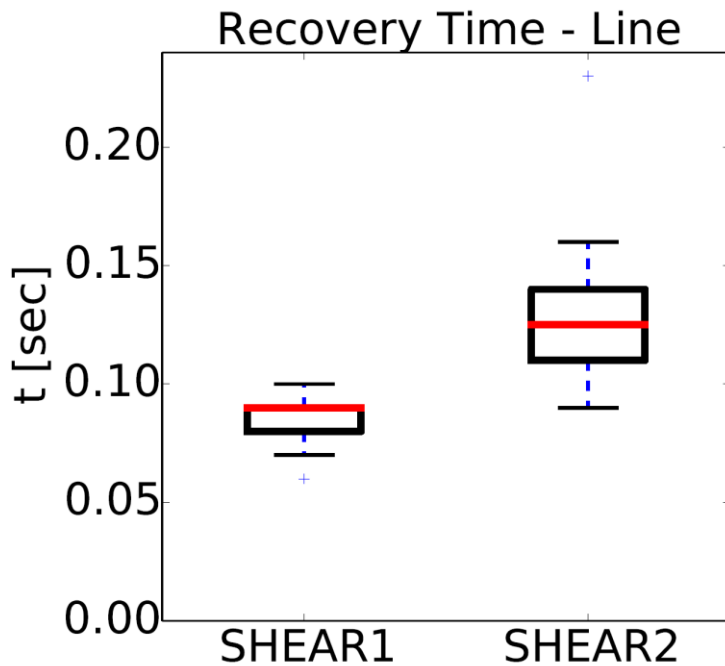
Testbed

- Nodes 3 and 7 spanning tree roots
- Some of the switches with very old firmware (pre standard MSTP)
- Simple controller application over RYU



Testbed: Recovery from failure

- A link was shutdown at distances of 2 and 3 hops from the SDN nodes
 - Recovery times were measured using ping



Conclusions

- A hybrid network can be a long term solution in order to achieve SDN abstraction
- Distributed protocols used to measure distances from known nodes (beacons)
 - Shown with MSTP, *can work with any routing algorithm*
- Failure recovery times competitive with any current distributed mechanisms
 - Distributed protocols only propagate failure data

Questions?

- markomic@post.bgu.ac.il