





AalWiNes: A Fast and Quantitative What-If Analysis Tool for MPLS Networks

Peter G. Jensen¹, Dan Kristiansen¹, Stefan Schmid², Morten K. Schou¹, Bernhard C. Schrenk², Jiří Srba¹

> ¹Aalborg University Denmark

²University of Vienna Austria

AalWiNes = Aalborg Wien Network verification suite

Motivation

- Operation of MPLS networks is an error prone task
- Hard to reason about behaviour in case of failures (what-if)
- Past work focus on logical properties P-Rex [CoNEXT '18]
- Quantitative aspects are also important



Our Contribution

- Tool for dataplane verification of MPLS networks
- Behaviour of **fast reroute** in case of multiple **failures**
- Over-approximation with **polynomial-time** algorithm
- Analyse quantitative aspects like latency and hop-count
- Visualize results in a cross-platform GUI

MPLS Network Model



Router	<i>e</i> _{in}	Label	Priority	<i>e</i> _{out}	Operation		
v_0	<i>e</i> ₀	ip ₁	1	<i>e</i> ₁	push(s20)		
	e_0	ip ₁	1	e_2	push(s10)		
	e_0	s 40	1	e_1	$\mathtt{swap}(s41)$		
v_1	e_2	s 10	1	<i>e</i> ₃	$\mathtt{swap}(\mathtt{s}11)$		
v_2	<i>e</i> ₁	s20	1	e_4	$\mathtt{swap}(s21)$		
	<i>e</i> ₁	s41	1	e_5	$\mathtt{swap}(s42)$		
	e_1	s20	2	e_5	$\texttt{swap}(s21) \circ \texttt{push}(30)$		
v_3	e ₃	s11	1	<i>e</i> ₇	pop		
	e_4	s21	1	<i>e</i> 7	pop		
	<i>e</i> ₆	s43	1	<i>e</i> ₇	swap(s44)		
	e_6	s21	1	<i>e</i> ₇	pop		
v_4	e_5	30	1	e_6	pop		
	e_5	s42	1	<i>e</i> ₆	$\mathtt{swap}(s43)$		

 $\sigma_0 = (e_0, ip_1) \ (e_1, s20 \circ ip_1) \ (e_4, s21 \circ ip_1) \ (e_7, ip_1)$

MPLS Network Model – Link Failure



Router	<i>e_{in}</i>	Label	Priority	<i>e</i> _{out}	Operation	
v_0	e_0	ip ₁	1	<i>e</i> ₁	push(s20)	
	e_0	ip ₁	1	e_2	push(s10)	
	e_0	s 40	1	<i>e</i> ₁	$\mathtt{swap}(s41)$	
v_1	e_2	s 10	1	<i>e</i> ₃	$\mathtt{swap}(s11)$	
v_2	e_1	s20	1	e_4	$\mathtt{swap}(s21)$	
	e_1	s 41	1	<i>e</i> 5	$\mathtt{swap}(s42)$	
	e_1	s20	2	e_5	$\texttt{swap}(s21) \circ \texttt{push}(30)$	
v_3	e ₃	s11	1	<i>e</i> ₇	pop	
	e_4	s21	1	<i>e</i> ₇	pop	
	e_6	s43	1	<i>e</i> ₇	$\mathtt{swap}(s44)$	
	e_6	s21	1	<i>e</i> ₇	рор	
v_4	e_5	30	1	<i>e</i> ₆	pop	
	e_5	s42	1	e_6	swap(s43)	

$$\begin{split} \sigma_0 &= (e_0, ip_1) \; (e_1, s20 \circ ip_1) \; (e_4, s21 \circ ip_1) \; (e_7, ip_1) \\ \sigma_1 &= (e_0, ip_1) \; (e_1, s20 \circ ip_1) \; (e_5, 30 \circ s21 \circ ip_1) (e_6, s21 \circ ip_1) \; (e_7, ip_1) \\ \sigma_2 &= (e_0, ip_1) \; (e_2, s10 \circ ip_1) \; (e_3, s11 \circ ip_1) \; (e_7, ip_1) \\ \sigma_3 &= (e_0, s40 \circ ip_1) \; (e_1, s41 \circ ip_1) \; (e_5, s42 \circ ip_1) \; (e_6, s43 \circ ip_1) \; (e_7, s44 \circ ip_1) \end{split}$$

Query Language

- Query language: $\langle a \rangle b \langle c \rangle k$
 - Regular expressions
 - *a*: Initial header
 - **b**: Path of links
 - C: Final header
- Does there exist a set of failed links F s.t. $|F| \le k$, where a trace satisfies $\langle a \rangle b \langle c \rangle$?

Query Language

• Examples of properties

Witness traces

- Reachability: $\langle ip \rangle$ [.# v_0].* [v_3 #.] $\langle ip \rangle$ 0 σ_0, σ_2
- Waypointing: $\langle ip \rangle$ [.# v_0].* [.# v_4].* [v_3 #.] $\langle ip \rangle$ 1 σ_1
- Avoid edge: $\langle ip \rangle$ [.# v_0] [^ v_2 # v_3]* [v_3 #.] $\langle ip \rangle$ 2 σ_1, σ_2
- Transparency: $\langle s40 \text{ ip} \rangle$ [.# v_0].* [v_3 #.] $\langle \text{mpls}^+ \text{ smpls ip} \rangle 0$ no



$$\begin{split} \sigma_0 &= (e_0, ip_1) \; (e_1, s20 \circ ip_1) \; (e_4, s21 \circ ip_1) \; (e_7, ip_1) \\ \sigma_1 &= (e_0, ip_1) \; (e_1, s20 \circ ip_1) \; (e_5, 30 \circ s21 \circ ip_1) (e_6, s21 \circ ip_1) \; (e_7, ip_1) \\ \sigma_2 &= (e_0, ip_1) \; (e_2, s10 \circ ip_1) \; (e_3, s11 \circ ip_1) \; (e_7, ip_1) \\ \sigma_3 &= (e_0, s40 \circ ip_1) \; (e_1, s41 \circ ip_1) \; (e_5, s42 \circ ip_1) \; (e_6, s43 \circ ip_1) \; (e_7, s44 \circ ip_1) \end{split}$$

Quantitative Extension

- If query is satisfied, find trace that minimizes:
 - Hops
 - Latency (based on a latency value per link)
 - Failures
 - Tunnels
- Combinations: $expr ::= p | a * expr | expr_1 + expr_2$ $(expr_1, expr_2, \dots, expr_n)$
- Example: (*Hops*, *Failures* + 3 · *Tunnels*)

Pushdown Automata Reachability



Handling Link Failures

- Over-approximation of behavior under failures
- Allow *k* failures at each router independently
- Check soundness of witness trace



• Inconclusive in < 1% of experiments



Implementation + Demo



Online demo: <u>https://demo.aalwines.cs.aau.dk/</u>

Source code: <u>https://github.com/DEIS-Tools/AalWiNes</u>



Performance Evaluation

- Real-world network with 31 routers and > 250 000 forwarding rules
- Specific queries requested by network operator

Query	Moped	Dual	Failures
$\langle \texttt{smpls ip} \rangle \; [\cdot \# R6] \; \cdot^* \; [\cdot \# R4] \; \langle \texttt{smpls ip} \rangle \; 1$	9.57	0.82	41.23
$\langle \texttt{smpls ip} \rangle [\cdot \# R2] \cdot^* [\cdot \# R18] \langle (\texttt{mpls}^* \texttt{smpls})? \texttt{ip} \rangle 1$	9.29	0.86	31.76
$\langle \texttt{ip} \rangle \left[\cdot \# R 0 \right] \cdot^* \left[\cdot \# R 4 \right] \langle \texttt{ip} \rangle 0$	0.88	0.01	0.02
$\langle [\$449550] ip \rangle [\cdot \#R0] \cdot^* [\cdot \#R5] \cdot^* [\cdot \#R1] \langle ip \rangle 0$	1.66	0.02	0.03
$\left< [\$449550] ip \right> [\cdot \#R0] \cdot^* [\cdot \#R5] \cdot^* [\cdot \#R1] \left< ip \right> 1$	6.08	0.05	0.06
$\langle \texttt{smpls?ip} \rangle \cdot^* \langle \cdot \texttt{smplsip} \rangle 0$	89.37	14.73	432.66

- Our solver ('Dual') outperforms state-of-the-art modelchecker Moped
- Acceptable overhead of quantitative analysis ('Failures')

Performance Evaluation

- Synthetic networks (based on Topology Zoo)
- More than 5600 experiments

• Quantitative analysis ('Failures') has fewer inconclusive cases

