Fast Automated What-if Analysis and Updates for Policy-Compliant Networks Even Under Failures Stefan Schmid (TU Berlin)

Fast Automated What-if Analysis and Updates for Policy-Compliant Networks Even Under Failures Stefan Schmid (TU Berlin)

In collaboration with Jiri Srba's team at Aalborg University, Denmark



Networks Are Complex

- Many outages are due to network configuration errors = *human errors*
- Examples (see Ratul Mahajan's NetVerify.Fun blog):
 - The December 2018 CenturyLink outage
 - The June 2020 T-Mobile outage
 - The July 2020 Cloudflare outage
 - The August 2020 CenturyLink outage

Particularly Challenging for Humans: Reasoning about Policy-Compliance under Failures



Particularly Challenging for Humans: Reasoning about Policy-Compliance under Failures







Particularly Challenging for Humans: Reasoning about Policy-Compliance under Failures



The Hope: Automation

• Can we automate the *verification* of the policy-compliance of configurations? Even under *failures*? Or even *synthetize* them?



The Hope: Automation

• Can we automate the *verification* of the policy-compliance of configurations? Even under *failures*? Or even *synthetize* them?





• A main challenge: should be *fast* as network configurations are not only complex for humans but also *computers* (many problems PSPACE-hard).

Roadmap

- A Static Problem: Policy Compliance
 Under Failures
 - AalWiNes: Fast Automated What-if Analysis for MPLS Networks (INFOCOM 2018, ACM CONEXT 2018, ACM CONEXT 2019, TACAS 2021)

- A Dynamic Problem: Scheduling Consistent Network Updates
 - Latte and quantitative extensions (PODC 2015, ICALP 2018, PERFORMANCE 2021)



Roadmap

- A Static Problem: Policy Compliance
 Under Failures
 - AalWiNes: Fast Automated What-if Analysis for MPLS Networks (INFOCOM 2018, ACM CONEXT 2018, ACM CONEXT 2019, TACAS 2021) And SR...
- A Dynamic Problem: Scheduling Consistent Network Updates
 - Latte and quantitative extensions (PODC 2015, ICALP 2018, PERFORMANCE 2021)



How (MPLS) Networks Work

• Forwarding based on top label of label stack



Default routing of two flows

How (MPLS) Networks Work

• Forwarding based on top label of label stack



How (MPLS) Networks Work



Default routing of two flows

Fast Reroute Around 1 Failure

• Forwarding based on top label of label stack (in packet header)



Default routing of two flows

• For failover: push and pop label



One failure: push 30: route around (v_2, v_3)

Fast Reroute Around 1 Failure

• Forwarding based on top label of label stack (in packet header)



Fast Reroute Around 1 Failure

• Forwarding based on top label of label stack (in packet header)



2 Failures: Push *Recursively*



Original Routing

One failure: push 30: route around (v_2, v_3)

Two failures: first push 30: route around (v₂,v₃) *Push recursively* 40: route around (v₂,v₆)

2 Failures: Push *Recursively*







Routers and switches store list of forwarding rules, and conditional failover rules.





Sysadmin responsible for:

• **Reachability:** Can traffic from ingress port A reach egress port B?



Sysadmin responsible for:

- **Reachability:** Can traffic from ingress port A reach egress port B?
- **Loop-freedom:** Are the routes implied by the forwarding rules loop-free?



Sysadmin responsible for:

- **Reachability:** Can traffic from ingress port A reach egress port B?
- **Loop-freedom:** Are the routes implied by the forwarding rules loop-free?
- **Policy:** Is it ensured that traffic from A to B never goes via C?



Sysadmin responsible for:

- **Reachability:** Can traffic from ingress port A reach egress port B?
- **Loop-freedom:** Are the routes implied by the forwarding rules loop-free?
- **Policy:** Is it ensured that traffic from A to B never goes via C?
- Waypoint enforcement: Is it ensured that traffic from A to B is always routed via a node C (e.g., intrusion detection system or a firewall)?



k failures = ssibilities А E.g. IDS

Sysadmin responsible for:

- **Reachability:** Can traffic from ingress port A reach egress port B?
- **Loop-freedom:** Are the routes implied by the forwarding rules loop-free?
- **Policy:** Is it ensured that traffic from A to B never goes via C?
- Waypoint enforcement: Is it ensured that traffic from A to B is always routed via a node C (e.g., intrusion detection system or a firewall)?

... and everything even under multiple failures?!



k failures = ssibilities А E.g. IDS

Sysadmin responsible for:

- **Reachability:** Can traffic from ingress port A reach egress port B?
- **Loop-freedom:** Are the routes implied by the forwarding rules loop-free?
- **Policy:** Is it ensured that traffic from A to B never goes via C?
 - Waypoint enforcement: Is it ensured that traffic from A to B is always routed via a node C (e.g., intrusion detection system or a firewall)?

... and everything even under multiple failures?!

Generalization: service chaining!

Approach: Automation and Formal Methods



Router **configurations** (Cisco, Juniper, etc.) Pushdown Automaton and Prefix Rewriting Systems



Router configurations (Cisco, Juniper, etc.) Pushdown Automaton and Prefix Rewriting Systems

AalWiNes



Online demo: <u>https://demo.aalwines.cs.aau.dk/</u> Source code: <u>https://github.com/DEIS-Tools/AalWiNes</u>

10

Example

Can traffic starting with [] go through s5, under up to k=2 failures?



Why AalWiNes is Fast (Polytime): Automata Theory

- For fast verification, we can use the result by Büchi: the set of all reachable configurations of a pushdown automaton a is regular set
- We hence simply use Nondeterministic Finite Automata (NFAs) when reasoning about the pushdown automata



Julius Richard Büchi 1924-1984 Swiss logician

• The resulting **regular operations** are all **polynomial time**

Case Study: NORDUnet

- Regional service provider
- 24 MPLS routers geographically distributed across several countries
- Running Juniper operating system
- More than 30,000 labels
- Ca. 1 million forwarding rules in our model
- For most queries of operators: answer *within seconds*



Generalizes to Quantitative Properties

- AalWiNes can also be used to test quantitative properties
- If query is satisfied, find trace that minimizes:
 - Hops
 - Latency (based on a latency value per link)
 - Tunnels



- Approach: weighted pushdown automata
 - Fast *poly-time algorithms* exist also for weighted pushdown automata (area of dataflow analysis)
 - Indeed, experiments show: acceptable overhead of weighted (quantitative) analysis

Roadmap

- A Static Problem: Policy Compliance
 Under Failures
 - AalWiNes: Fast Automated What-if Analysis for MPLS Networks (INFOCOM 2018, ACM CONEXT 2018, ACM CONEXT 2019, TACAS 2021)

- A Dynamic Problem: Scheduling Consistent Network Updates
 - Latte and quantitative extensions (PODC 2015, ICALP 2018, PERFORMANCE 2021)



More Adaptable Networks

- Automation and programmability also enables networks to be more adaptable
- Attractive for:

...

- Fine-grained traffic engineering (e.g., at Google)
- Accounting for changes in the demand (*spatio-temporal structure*)
- Security policy changes
- Service relocation
- Maintenance work
- Link/node failures



Introduces a New Challenge: Scheduling Updates



Invariant: Traffic from untrusted hosts to trusted hosts via firewall!

Introduces a New Challenge: Scheduling Updates



Invariant: Traffic from untrusted hosts to trusted hosts via firewall!

Introduces a New Challenge: Scheduling Updates



Invariant: Traffic from untrusted hosts to trusted hosts via firewall!

Latte: Synthesis of Shortest Consistent Update Schedules

- Much work on the design of efficient algorithms for consistent network updates
- Our goal: automated synthesis of fast updates accounting for temporal properties
 - E.g., different packet types have different requirements and *processing times*
 - Builds upon *NetSynth* (gives fixed update order)
- A classic tool to reason about *asynchronous distributed systems*: petri nets
 - Configurations: tokens located at places
- Our extension: Timed-Arc Colored Petri Nets (TACPN)
 - Tokens also contain: *color* information (e.g., different packet *types*) and time information (e.g., modeling *age*)
 - Places and input arcs have time constraints for each color

Latte: Synthesis of Shortest Consistent Update Schedules

- Much work on the design of efficient algorithms for consistent network updates
- Our goal: automated synthesis of fast updates accounting for temporal properties
 - E.g., different packet types have different requirements and *processing times*
 - Builds upon *NetSynth* (gives fixed update order)
- A classic tool to reason about *asynchronous distributed systems*: petri nets
 - Configurations: tokens located at places
- Our extension: Timed-Arc Colored Petri Nets (TACPN)
 - Tokens also contain: *color* information (e.g., different packet *types*) and time information (e.g., modeling *age*)
 - Places and input arcs have time constraints for each color









3 Gadget to model switch update:

How to change between initial and final switch configuration



4 Connecting the pieces: initialization of update sequence for all *n* switches



Analysis

The constructed nets can be analyzed efficiently via their *unfolding* into existing *timed-arc Petri nets*.



Improved Latency of Update Schedules

Compared to conservative delays as produced by NetSynth: over 90% improvement.

Network	Route length	Verification time[s]	Default update time [s]	Optimized update time [s]	Improvement [%]
TLex	4	0.74	3.58	0.25	92.30%
HiberniaIreland	5	1.02	6.05	0.28	95.50%
Harnet	6	1.42	9.08	0.28	96.97%
UniC	7	1.49	12.65	0.28	97.83%
Oxford	8	2.02	16.78	0.28	98.36%
Xeex	10	5.86	26.68	0.28	98.97%
Sunet	11	10.23	32.45	0.28	99.15%
SwitchL3	12	18.88	38.78	0.28	99.29%
Psinet	14	89.67	53.01	0.28	99.48%
Uunet	15	211.86	61.05	0.28	99.55%
Renater2010	16	480.52	69.58	0.28	99.60%
Missouri	25	timeout	171.05	67.10	60.77%
Syringa	35	timeout	336.05	295.35	12.11%
VtlWavenet2011	35	timeout	336.06	295.35	12.11%

- Network topologies from the Topology Zoo
- Experiments run on a 64-bit Ubuntu 18.04 laptop

Improved Latency of Update Schedules

Compared to conservative delays as produced by NetSynth: over 90% improvement.

	Network	Route length	Verification time[s]	Default update time [s]	Optimized update time [s]	Improvement [%]
	TLex	4	0.74	3.58	0.25	92.30%
	HiberniaIreland	5	1.02	6.05	0.28	95.50%
	Harnet	6	1.42	9.08	0.28	96.97%
	UniC	7	1.49	12.65	0.28	97.83%
	Oxford	8	2.02	16.78	0.28	98.36%
	Xeex	10	5.86	26.68	0.28	98.97%
	Sunet	11	10.23	32.45	0.28	99.15%
	a		18.88	38.78	0.28	99.29%
Jp to route length 16, optimal update			date 89.67	53.01	0.28	99.48%
	time can be o	computed.	211.86	61.05	0.28	99.55%
			480.52	69.58	0.28	99.60%
	Missouri	25	timeout	171.05	67.10	60.77%
	Syringa	35	timeout	336.05	295.35	12.11%
	VtlWavenet2011	35	timeout	226.06	205.25	12.11%

Too many updates concurrently: could be tackled with static analysis (future work).

- Network topologies from the Top
- Experiments run on a 64-bit Ubuntu 18.04 laptop

Conclusion

- Finally: networks are moving from manual to more automated operations
- Supported by emerging programmable networks and their solid theoretical foundations and languages
- Automata-theoretical approaches can be used to perform fast what-if analysis of the policy compliance (e.g., P-Rex, *AalWiNes*, etc.)
 - E.g., MPLS networks, but also Segment Routing networks
- More adaptive network operations further require tools for consistent network update scheduling (e.g., *Latte*)
- Current research focus on:
 - Accounting for *quantitative aspects*
 - Improving performance further *with AI*, without losing formal guarantees (e.g., configuration of CEGAR)

Further Reading

The AalWines project https://aalwines.cs.aau.dk/



TAPAAL.net



The TGHAL but of dime a graphical date for dwaring TMP models, simulator for apprivately on the date of without on entropy of the second seco

Netverify.fun

RESEARCH, NETWORK, VERIFICATION

Toward Polynomial-Time Verification of Networks with Infinite State Spaces: An Automata-Theoretic Approach





Jul 20. 2020 · 6 mins read Jiri Srba (View ul 20, 2020 · 6 mins read



ith the increasing scale of communication networks, failures (e.g. link failures) are becoming the norm rather than the exception. Given the critical role such networks play for our digital society it is important to

References

Resilient Capacity-Aware Routing

Stefan Schmid, Nicolas Schnepf and Jiri Srba.

27th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (**TACAS**), Virtual Conference, March 2021.

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

Peter Gjøl Jensen, Morten Konggaard, Dan Kristiansen, Stefan Schmid, Bernhard Clemens Schrenk, and Jiri Srba.

16th ACM International Conference on emerging Networking EXperiments and Technologies (CoNEXT), Barcelona, Spain, December 2020.

Latte: Improving the Latency of Transiently Consistent Network Update Schedules

Mark Glavind, Niels Christensen, Jiri Srba, and Stefan Schmid.

38th International Symposium on Computer Performance, Modeling, Measurements and Evaluation (**PERFORMANCE**) and ACM Performance Evaluation Review (**PER**), Milan, Italy, November 2020.

P-Rex: Fast Verification of MPLS Networks with Multiple Link Failures

Jesper Stenbjerg Jensen, Troels Beck Krogh, Jonas Sand Madsen, Stefan Schmid, Jiri Srba, and Marc Tom Thorgersen. 14th ACM International Conference on emerging Networking EXperiments and Technologies (**CoNEXT**), Heraklion/Crete, Greece, December 2018.

Congestion-Free Rerouting of Flows on DAGs

Saeed Akhoondian Amiri, Szymon Dudycz, Stefan Schmid, and Sebastian Wiederrecht.

45th International Colloquium on Automata, Languages, and Programming (ICALP), Prague, Czech Republic, July 2018.

Polynomial-Time What-If Analysis for Prefix-Manipulating MPLS Networks

Stefan Schmid and Jiri Srba.

37th IEEE Conference on Computer Communications (INFOCOM), Honolulu, Hawaii, USA, April 2018.

DeepMPLS: Fast Analysis of MPLS Configurations Using Deep Learning

Fabien Geyer and Stefan Schmid.

IFIP Networking, Warsaw, Poland, May 2019.



Questions?