WHATIF: Fast and Quantitative What-if Analysis for Dependable Communication Networks (ICT19-045, 2020-2024)

Stefan Schmid

Roadmap

- WHATIF: Motivation and Context
- Current State and Achievements
- Plans

Communication Networks

Critical infrastructure of digital society

- Popularity of **datacentric applications**: health, business, entertainment, social networking, AI/ML, etc.
- IoT and innovative new **smart city** concepts (e.g., Aspern)

Increasingly stringent dependability requirements!

Requirements vs Reality

Entire countries disconnected…

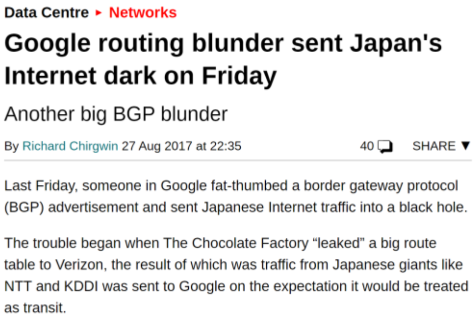

Data Centre ▸ Networks

Google routing blunder sent Japan's Internet dark on Friday

Another big BGP blunder

By Richard Chirgwin 27 Aug 2017 at 22:35 40 SHARE ▼

Last Friday, someone in Google fat-thumbed a border gateway protocol (BGP) advertisement and sent Japanese Internet traffic into a black hole.

The trouble began when The Chocolate Factory "leaked" a big route table to Verizon, the result of which was traffic from Japanese giants like NTT and KDDI was sent to Google on the expectation it would be treated as transit.

… 1000s passengers stranded…

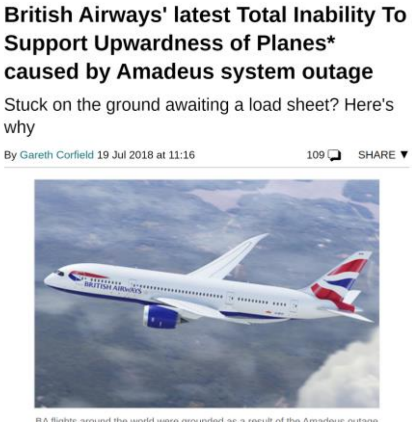

British Airways' latest Total Inability To Support Upwardness of Planes* caused by Amadeus system outage

Stuck on the ground awaiting a load sheet? Here's why

By Gareth Corfield 19 Jul 2018 at 11:16 109 SHARE ▼

BA flights around the world were grounded as a result of the Amadeus outage

… even 911 services affected!

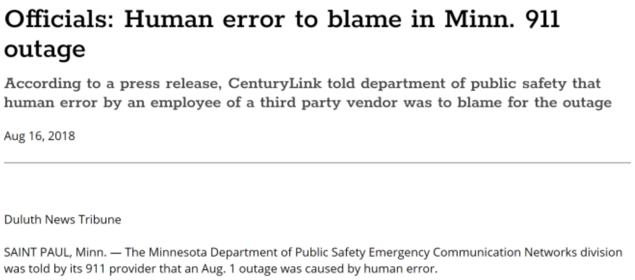

Officials: Human error to blame in Minn. 911 outage

According to a press release, CenturyLink told department of public safety that human error by an employee of a third party vendor was to blame for the outage

Aug 16, 2018

Duluth News Tribune

SAINT PAUL, Minn. — The Minnesota Department of Public Safety Emergency Communication Networks division was told by its 911 provider that an Aug. 1 outage was caused by human error.

Outages simply due to human error! (No attacks...)

Even Tech-Savvy Companies Struggle

github SOCIAL CODING

We discovered a misconfiguration on this pair of switches that caused what's called a "bridge loop" in the network.

A network change was […] executed incorrectly […] more "stuck" volumes and added more requests to the re-mirroring storm.

amazon web services

Go Daddy.COM

Service outage was due to a series of internal network events that corrupted router data tables.

Experienced a network connectivity issue […] interrupted the airline's flight departures, airport processing and reservations systems

United Airlines

Also here: due to human errors.

No Surprise: Networks Are Complex

Manual, device-centric network configurations
(CLI, LANmanager)

Un-evolved best practices
(tcpdump, traceroute - from the 1990s)

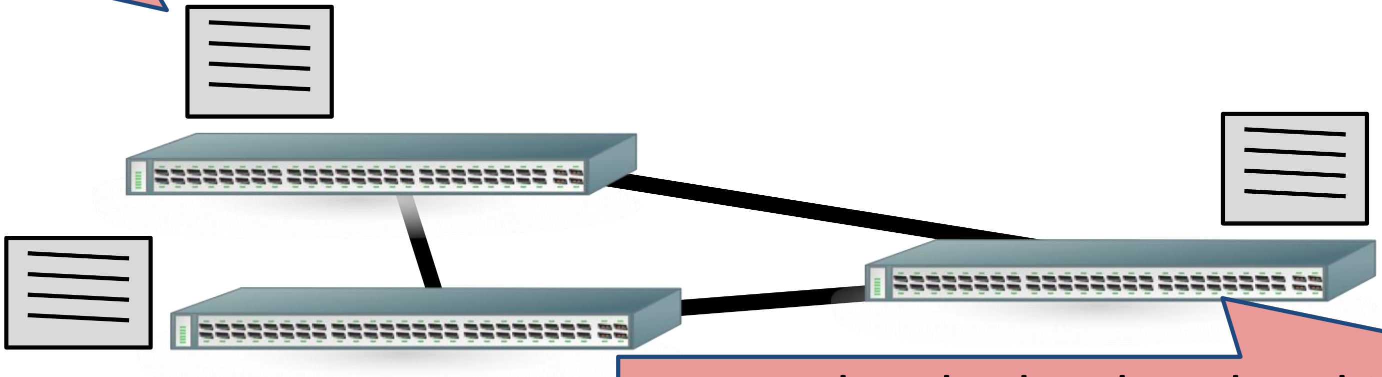

Complex, leaky, low-level interfaces
(VLANs, Spanning Tree, Routing)

500-router network: typically
>1 million lines of configuration

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

Example: BGP in

Datacenter

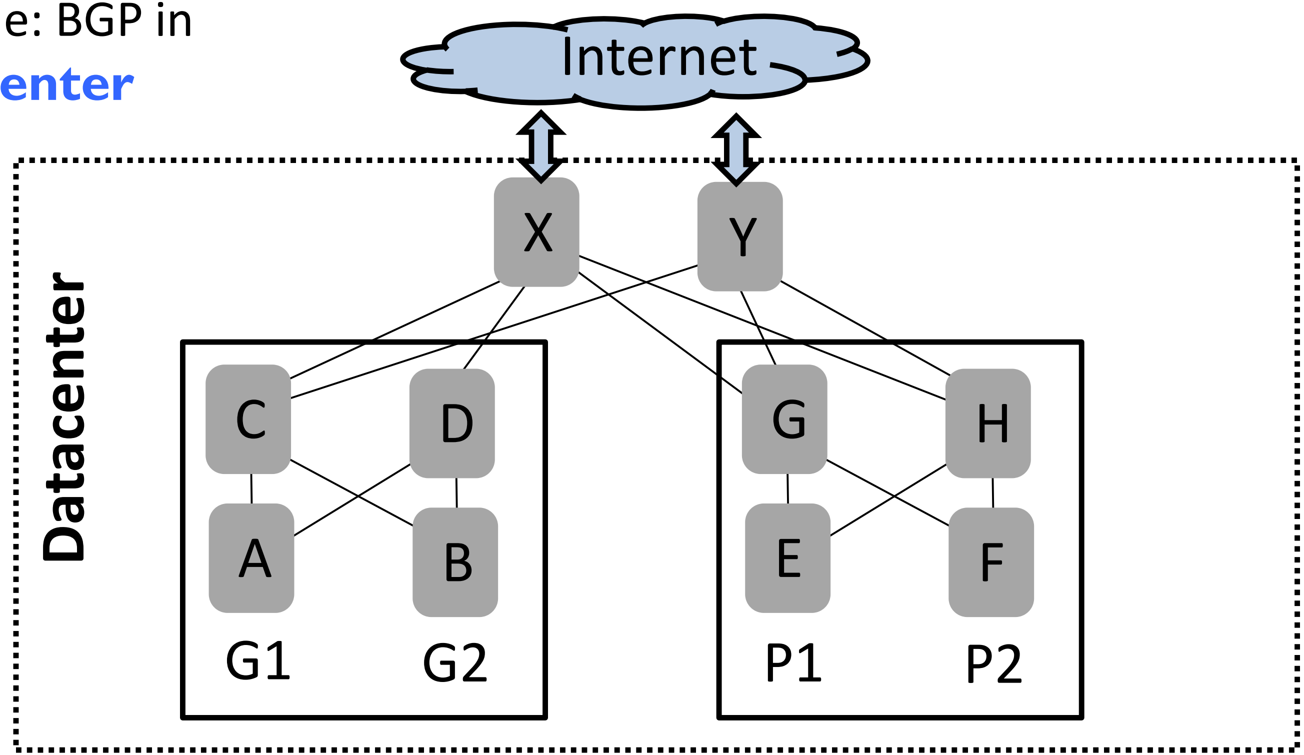

Credits: Beckett et al. (SIGCOMM 2016): Bridging Network-wide Objectives and Device-level Configurations.

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

Example: BGP in **Datacenter**

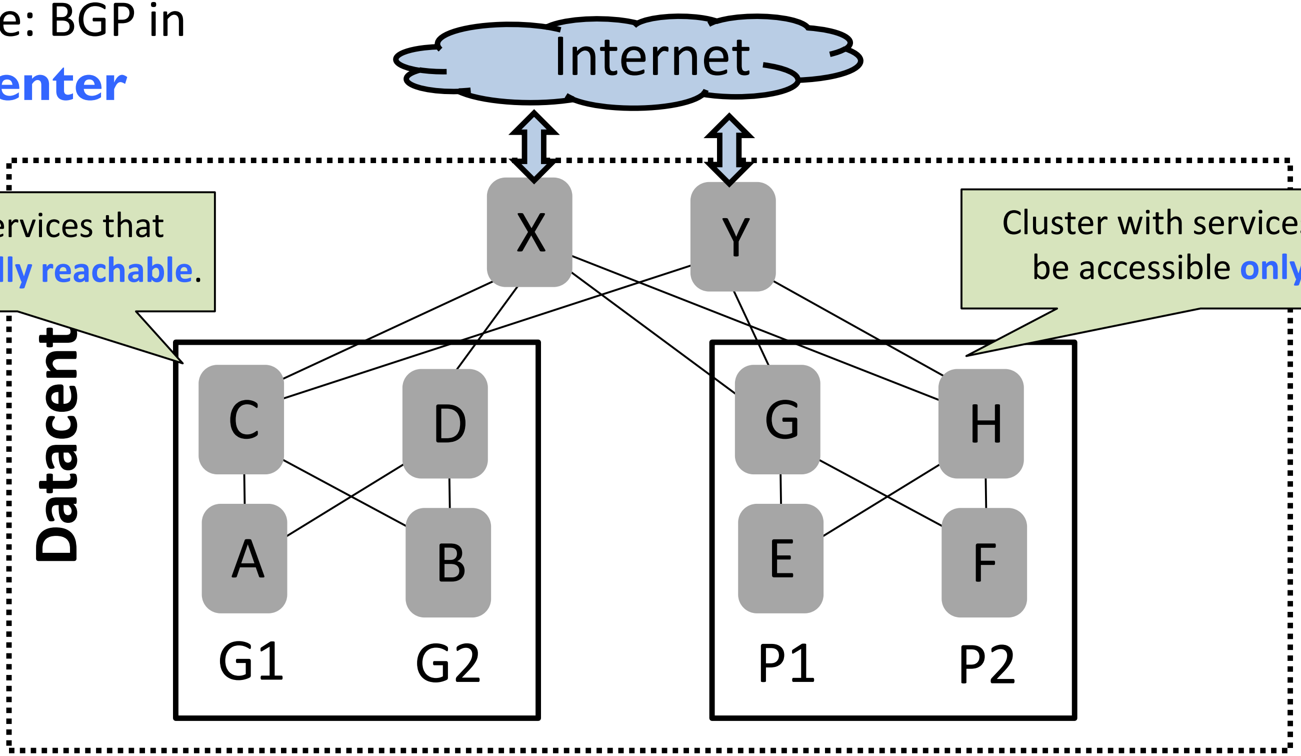

Cluster with services that should be **globally reachable**.

Cluster with services that should be accessible **only internally**.

Credits: Beckett et al. (SIGCOMM 2016): Bridging Network-wide Objectives and Device-level Configurations.

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

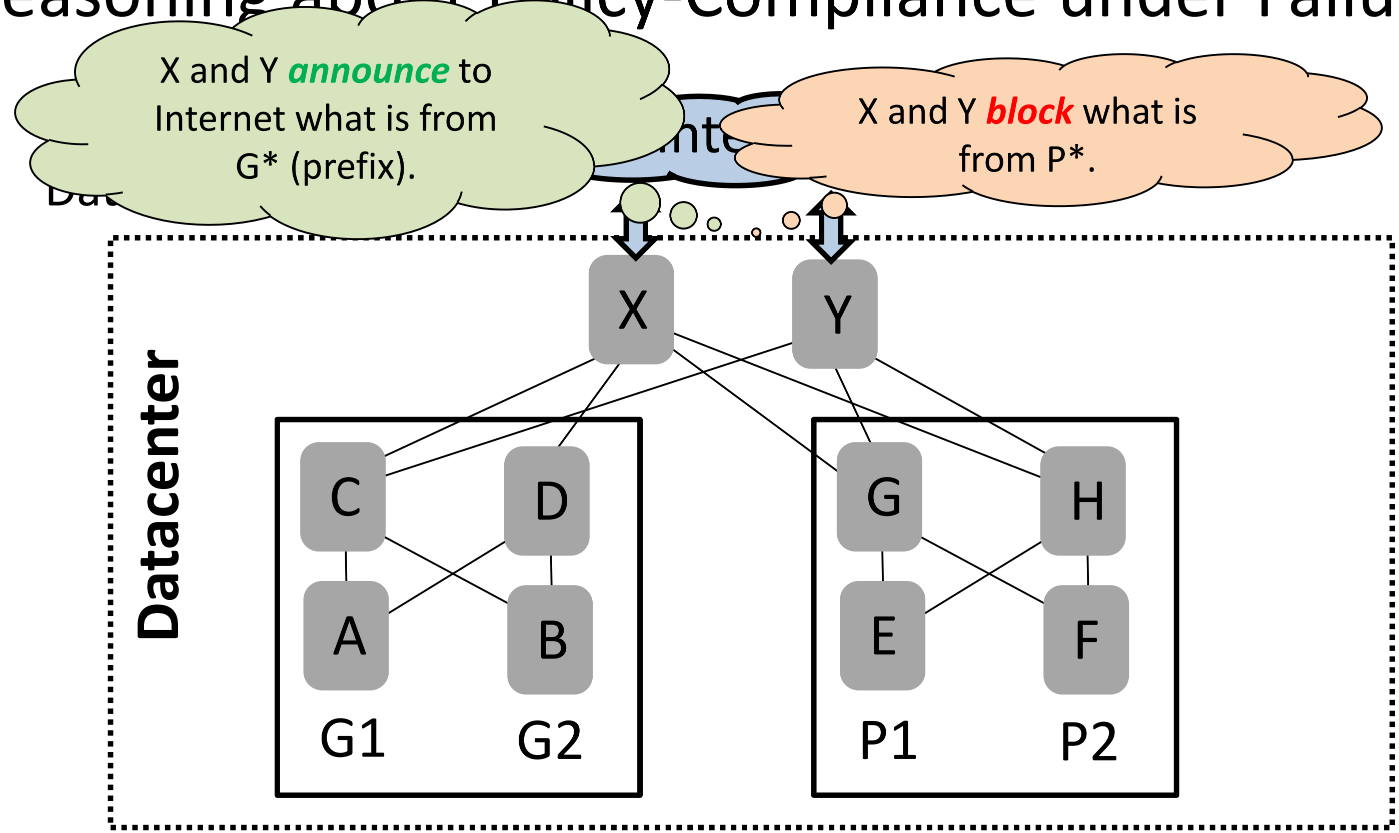

X and Y ***announce*** to Internet what is from G* (prefix).

X and Y ***block*** what is from P*.

Credits: Beckett et al. (SIGCOMM 2016): Bridging Network-wide Objectives and Device-level Configurations.

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

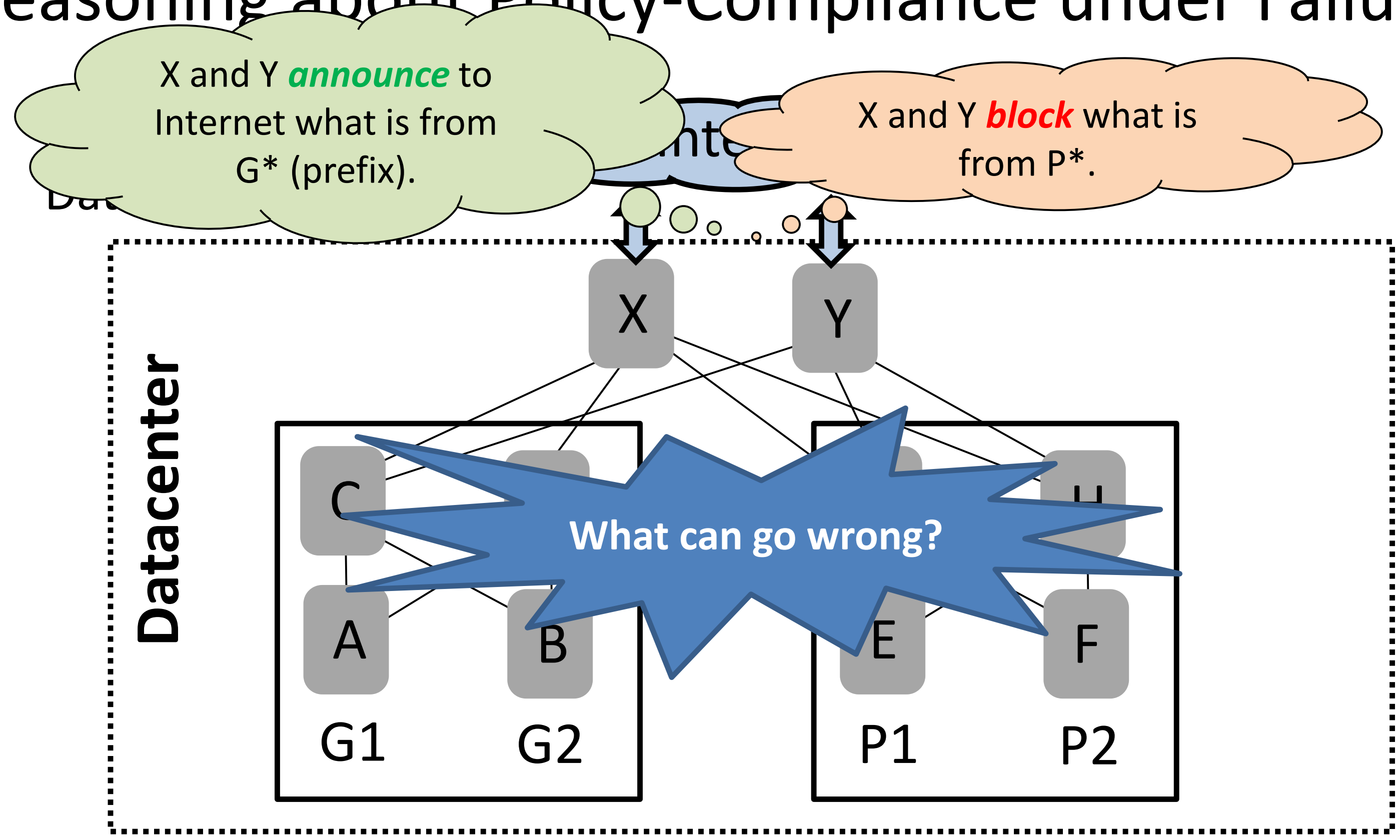
X and Y ***announce*** to Internet what is from G* (prefix).
X and Y ***block*** what is from P*.
Datacenter
X
Y
C
A
B
E
H
F
G1
G2
P1
P2
What can go wrong?

Credits: Beckett et al. (SIGCOMM 2016): Bridging Network-wide Objectives and Device-level Configurations.

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

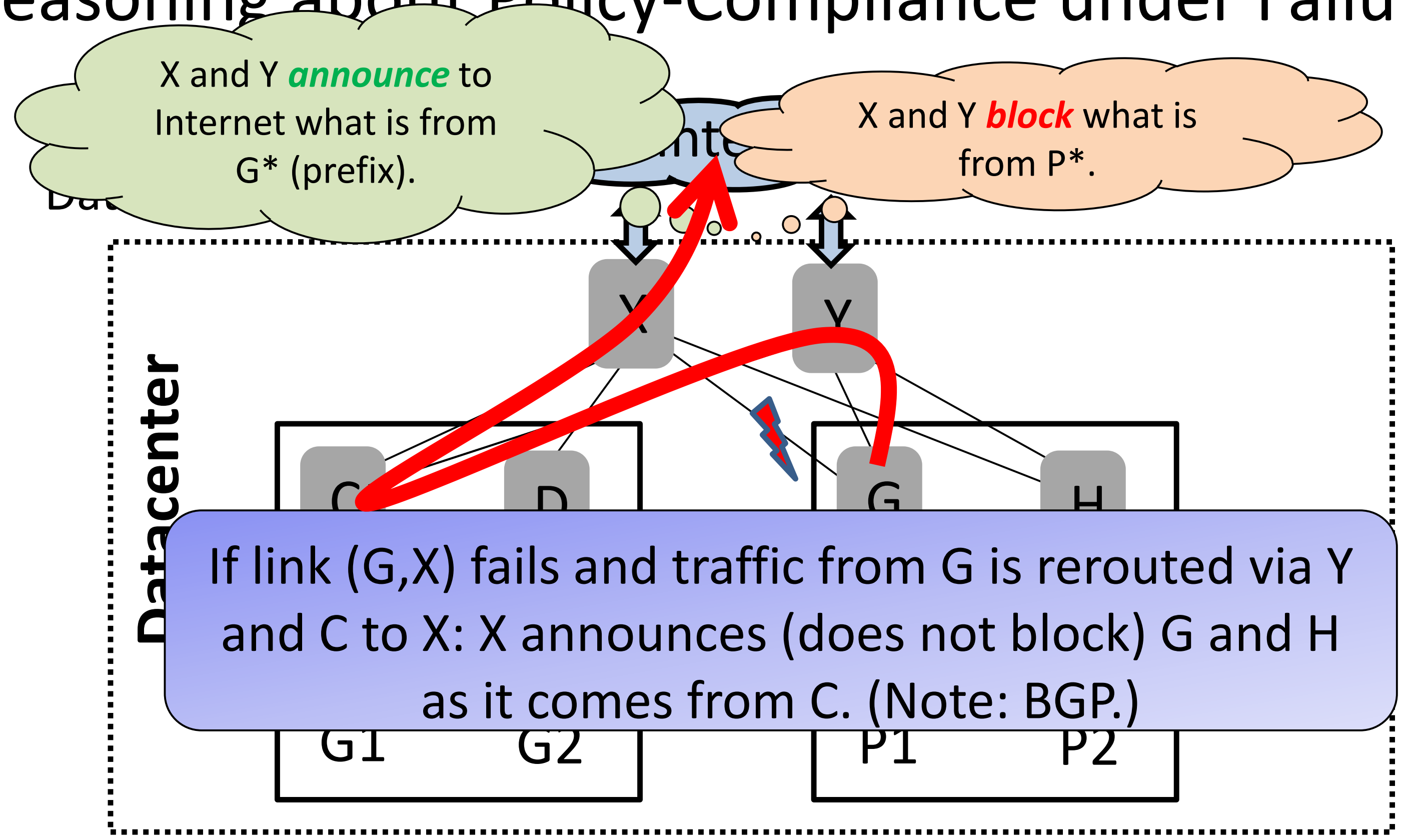

X and Y *announce* to Internet what is from G* (prefix).
X and Y *block* what is from P*.
Inte
Datacenter
X
Y
C
D
G
H
G1
G2
P1
P2

If link (G,X) fails and traffic from G is rerouted via Y and C to X: X announces (does not block) G and H as it comes from C. (Note: BGP.)

Credits: Beckett et al. (SIGCOMM 2016): Bridging Network-wide Objectives and Device-level Configurations.

We’re Falling Behind the Curve: Increasing Complexity, Software from the 90s

- Anecdote **Wall Street bank**: outage of a datacenter
  - Lost revenue measured in **1 mio$/min**
- Quickly, an emergency team was assembled with experts in compute, storage and networking:
  - **The compute team:** ***reams of logs***, written experiments to reproduce and ***isolate the error***
  - **The storage team:** ***system logs*** were affected, ***workaround programs***.
  - “All the **networking team** had were ***two tools invented over twenty years ago*** to merely test end-to-end connectivity. Neither tool could reveal ***problems with the switches***, the ***congestion*** experienced.”

Source: «The world’s fastest and most programmable networks»
White Paper Barefoot Networks

Responsibilities of a Sysadmin

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

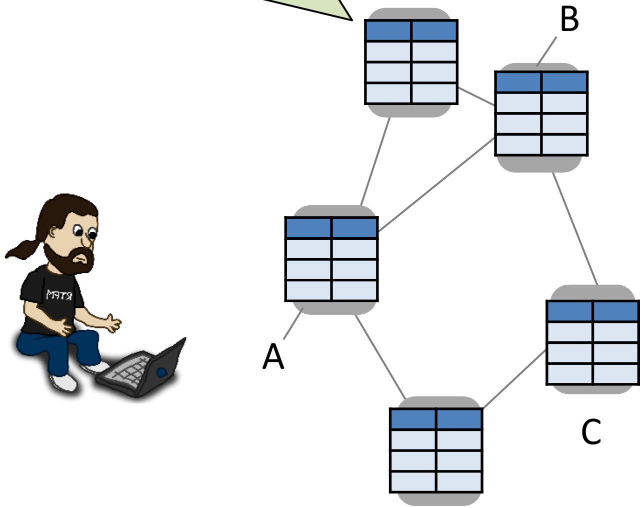

Responsibilities of a Sysadmin

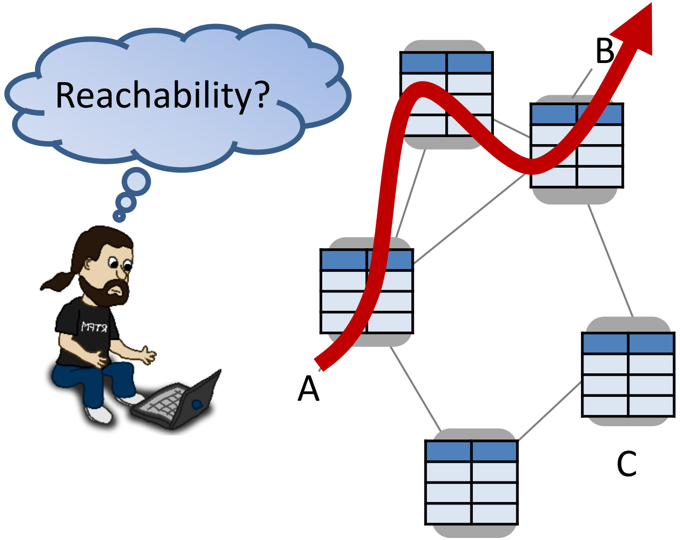

Sysadmin responsible for:

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

Responsibilities of a Sysadmin

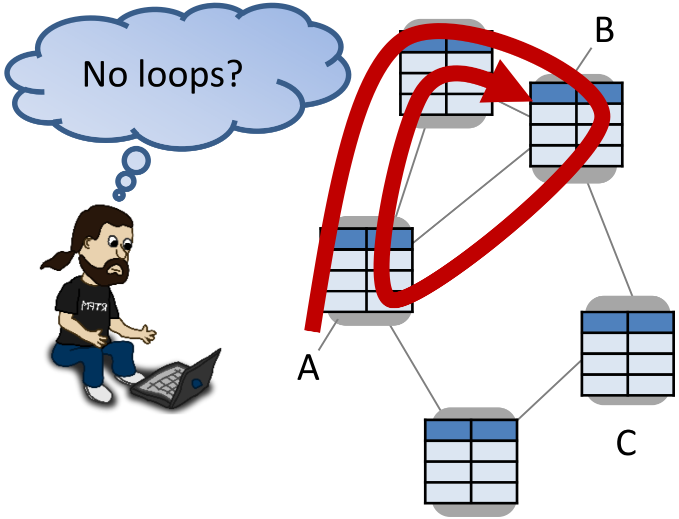

Sysadmin responsible for:

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

Responsibilities of a Sysadmin

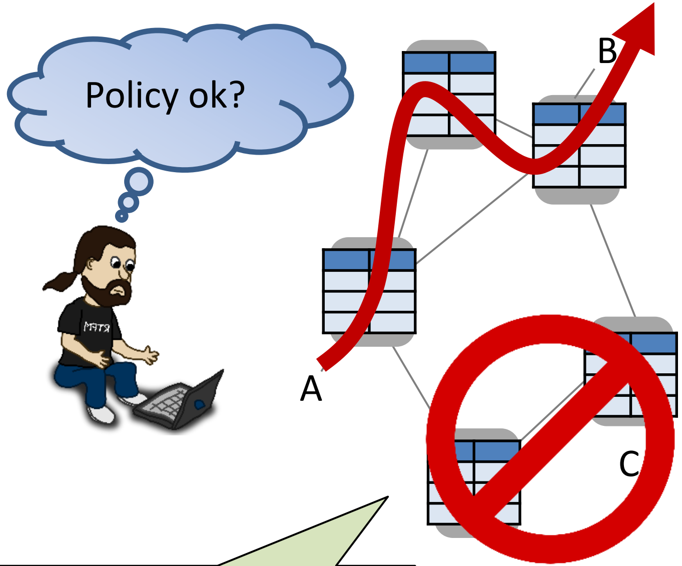

E.g. ***NORDUnet***: no traffic via Iceland (expensive!).

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?

Responsibilities of a Sysadmin

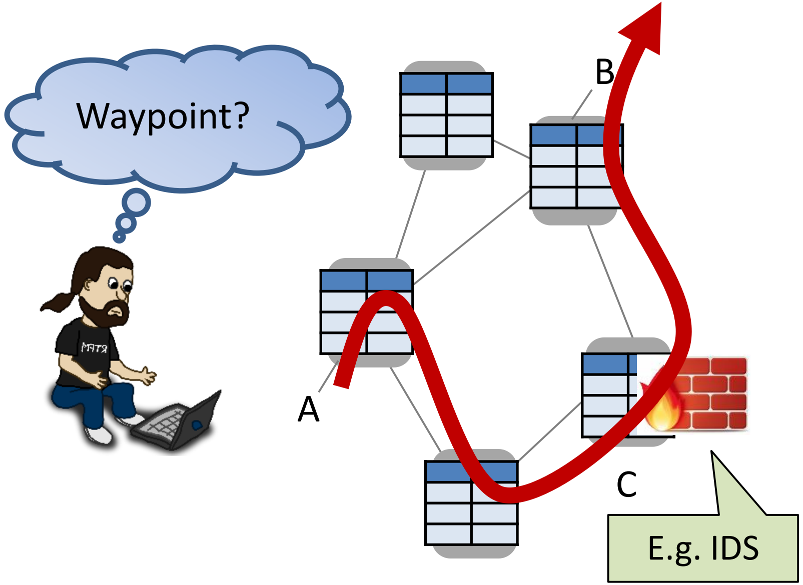

Sysadmin responsible for:

- **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)?

Responsibilities of a Sysadmin

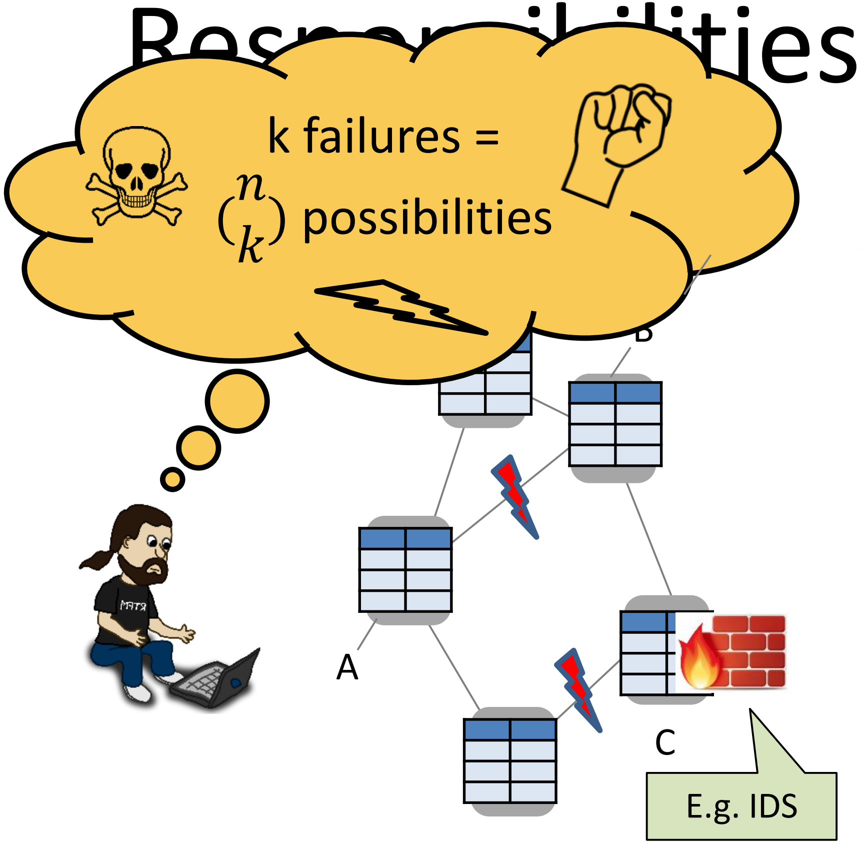

Sysadmin responsible for:

- **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?!

Responsibilities of a Sysadmin

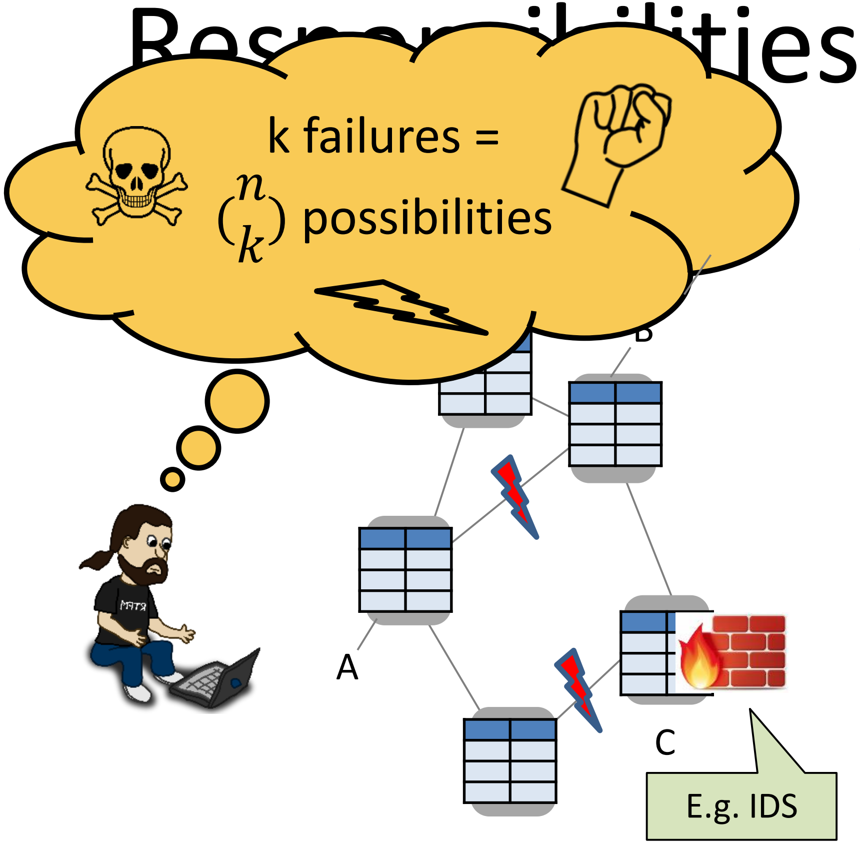
k failures = $\binom{n}{k}$ possibilities

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!

WHATIF: Automation and Formal Methods

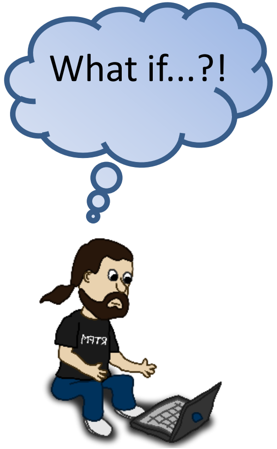

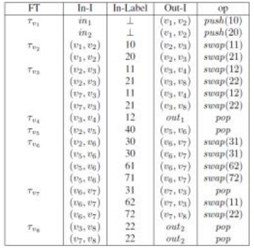

FT	In-I	In-Label	Out-I	op
τ_{v_1}	in_1	$\perp$	(v_1, v_2)	$push(10)$
	in_2	$\perp$	(v_1, v_2)	$push(20)$
τ_{v_2}	(v_1, v_2)	10	(v_2, v_3)	$swap(11)$
	(v_1, v_2)	20	(v_2, v_3)	$swap(21)$
τ_{v_3}	(v_2, v_3)	11	(v_3, v_4)	$swap(12)$
	(v_2, v_3)	21	(v_3, v_8)	$swap(22)$
	(v_7, v_3)	11	(v_3, v_4)	$swap(12)$
	(v_7, v_3)	21	(v_3, v_8)	$swap(22)$
τ_{v_4}	(v_3, v_4)	12	out_1	pop
τ_{v_5}	(v_2, v_5)	40	(v_5, v_6)	pop
τ_{v_6}	(v_2, v_6)	30	(v_6, v_7)	$swap(31)$
	(v_5, v_6)	30	(v_6, v_7)	$swap(31)$
	(v_5, v_6)	61	(v_6, v_7)	$swap(62)$
	(v_5, v_6)	71	(v_6, v_7)	$swap(72)$
τ_{v_7}	(v_6, v_7)	31	(v_7, v_3)	pop
	(v_6, v_7)	62	(v_7, v_3)	$swap(11)$
	(v_6, v_7)	72	(v_7, v_8)	$swap(22)$
τ_{v_8}	(v_3, v_8)	22	out_2	pop
	(v_7, v_8)	22	out_2	pop

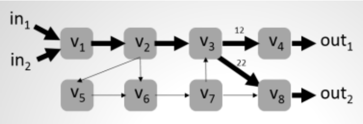

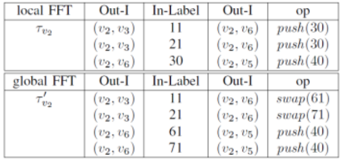

local FFT	Out-I	In-Label	Out-I	op
τ_{v_2}	(v_2, v_3)	11	(v_2, v_6)	$push(30)$
	(v_2, v_3)	21	(v_2, v_6)	$push(30)$
	(v_2, v_6)	30	(v_2, v_5)	$push(40)$

global FFT	Out-I	In-Label	Out-I	op
τ'_{v_2}	(v_2, v_3)	11	(v_2, v_6)	$swap(61)$
	(v_2, v_3)	21	(v_2, v_6)	$swap(71)$
	(v_2, v_6)	61	(v_2, v_5)	$push(40)$
	(v_2, v_6)	71	(v_2, v_5)	$push(40)$

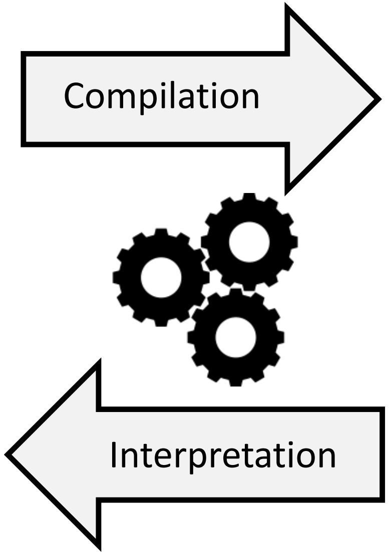

Compilation

Interpretation

$$\mathrm{pX} \Rightarrow \mathrm{qXX}$$
$$pX \Rightarrow qYX$$
$$qY \Rightarrow rYY$$
$$rY \Rightarrow r$$
$$rX \Rightarrow pX$$

Router **configurations**
(Cisco, Juniper, etc.)

Pushdown Automaton and
Prefix Rewriting Systems

WHATIF: Automat… …ods

Use cases: Sysadmin ***issues queries*** to test certain properties, or do it on a ***regular basis*** automatically!

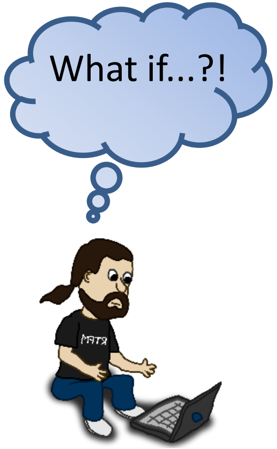

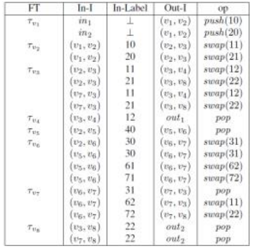

FT	In-I	In-Label	Out-I	op
τ_{v_1}	in_1	$\bot$	(v_1, v_2)	$push(10)$
	in_2	$\bot$	(v_1, v_2)	$push(20)$
τ_{v_2}	(v_1, v_2)	10	(v_2, v_3)	$swap(11)$
	(v_1, v_2)	20	(v_2, v_3)	$swap(21)$
τ_{v_3}	(v_2, v_3)	11	(v_3, v_4)	$swap(12)$
	(v_2, v_3)	21	(v_3, v_8)	$swap(22)$
	(v_7, v_3)	11	(v_3, v_4)	$swap(12)$
	(v_7, v_3)	21	(v_3, v_8)	$swap(22)$
τ_{v_4}	(v_3, v_4)	12	out_1	pop
τ_{v_5}	(v_2, v_5)	40	(v_5, v_6)	pop
τ_{v_6}	(v_2, v_6)	30	(v_6, v_7)	$swap(31)$
	(v_5, v_6)	30	(v_6, v_7)	$swap(31)$
	(v_5, v_6)	61	(v_6, v_7)	$swap(62)$
	(v_5, v_6)	71	(v_6, v_7)	$swap(72)$
τ_{v_7}	(v_6, v_7)	31	(v_7, v_3)	pop
	(v_6, v_7)	62	(v_7, v_3)	$swap(11)$
	(v_6, v_7)	72	(v_7, v_8)	$swap(22)$
τ_{v_8}	(v_3, v_8)	22	out_2	pop
	(v_7, v_8)	22	out_2	pop

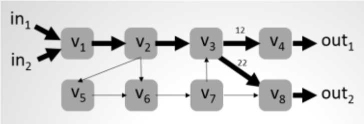

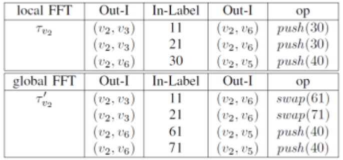

local FFT	Out-I	In-Label	Out-I	op
τ_{v_2}	(v_2, v_3)	11	(v_2, v_6)	$push(30)$
	(v_2, v_3)	21	(v_2, v_6)	$push(30)$
	(v_2, v_6)	30	(v_2, v_5)	$push(40)$

global FFT	Out-I	In-Label	Out-I	op
τ'_{v_2}	(v_2, v_3)	11	(v_2, v_6)	$swap(61)$
	(v_2, v_3)	21	(v_2, v_6)	$swap(71)$
	(v_2, v_6)	61	(v_2, v_5)	$push(40)$
	(v_2, v_6)	71	(v_2, v_5)	$push(40)$

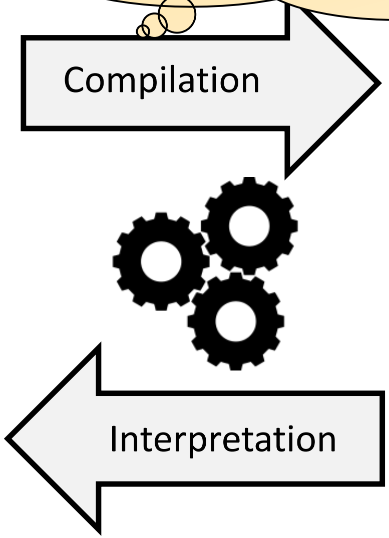

Compilation

Interpretation

$$\mathrm{pX} \Rightarrow \mathrm{qXX}$$

$$pX \Rightarrow qYX$$

$$qY \Rightarrow rYY$$

$$rY \Rightarrow r$$

$$rX \Rightarrow pX$$

Router **configurations**
(Cisco, Juniper, etc.)

Pushdown Automaton and
Prefix Rewriting Systems

Roadmap

- WHATIF: Motivation and Context
- Current State and Achievements
- Plans

Work Packages

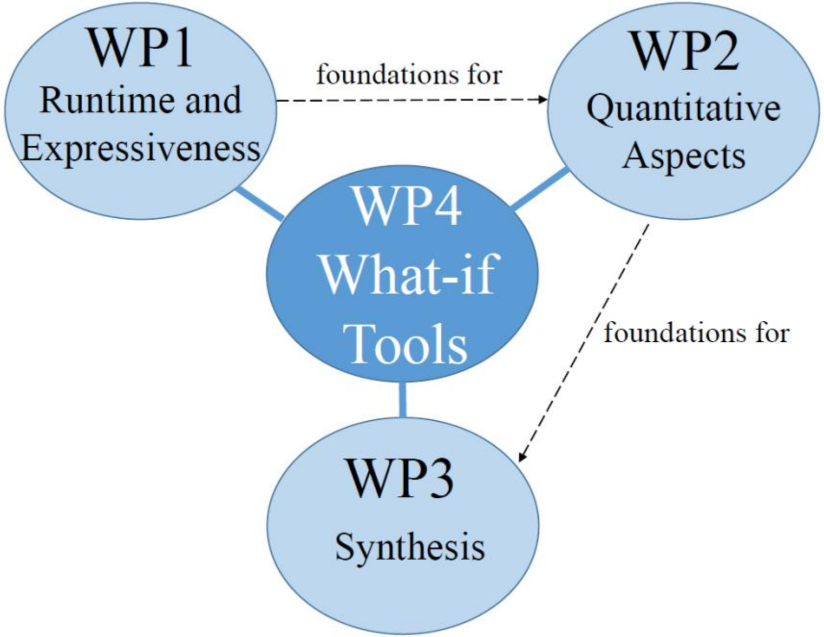

Some Publications (1)

WP2

The Hazard Value: A Quantitative Network Connectivity Measure Accounting for Failures
Pieter Cuijpers, Stefan Schmid, Nicolas Schnepf, and Jiri Srba.
52nd IEEE/IFIP International Conference on Dependable Systems and Networks (**DSN**), Baltimore, Maryland, USA, June 2022.

WP3

On the Price of Locality in Static Fast Rerouting
Klaus-Tycho Foerster, Juho Hirvonen, Yvonne-Anne Pignolet, Stefan Schmid, and Gilles Tredan.
52nd IEEE/IFIP International Conference on Dependable Systems and Networks (**DSN**), Baltimore, Maryland, USA, June 2022.

WP3

NetStack: A Game Approach to Synthesizing Consistent Network Updates
Stefan Schmid, Bernhard Schrenk, and Alvaro Torralba.
IFIP Networking, Catania, Italy, June 2022.

WP1

Automata Theoretic Approach to Verification of MPLS Networks under Link Failures
Peter Gjøl Jensen, Jesper Stenbjerg Jensen, Troels Beck Krogh, Jonas Sand Madsen, Stefan Schmid, Jiri Srba, Marc Tom Thorgersen, and Ingo van Duijn.
IEEE/ACM Transactions on Networking (**TON**), 2021.

WP1

A Survey of Fast-Recovery Mechanisms in Packet-Switched Networks
Marco Chiesa, Andrzej Kamisinski, Jacek Rak, Gabor Retvari, and Stefan Schmid.
IEEE Communications Surveys and Tutorials (**COMST**), 2021.

Some Publications (2)

WP1

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.

WP2

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.

WP3

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.

WP1

Faster Pushdown Reachability Analysis with Applications in Network Verification
Peter Gjøl Jensen, Stefan Schmid, Morten Konggaard Schou, Jirí Srba, Juan Vanerio, and Ingo van Duijn.
19th International Symposium on Automated Technology for Verification and Analysis (**ATVA**), Gold Coast, Australia, October 2021.

WP4

AalWiNes Tool

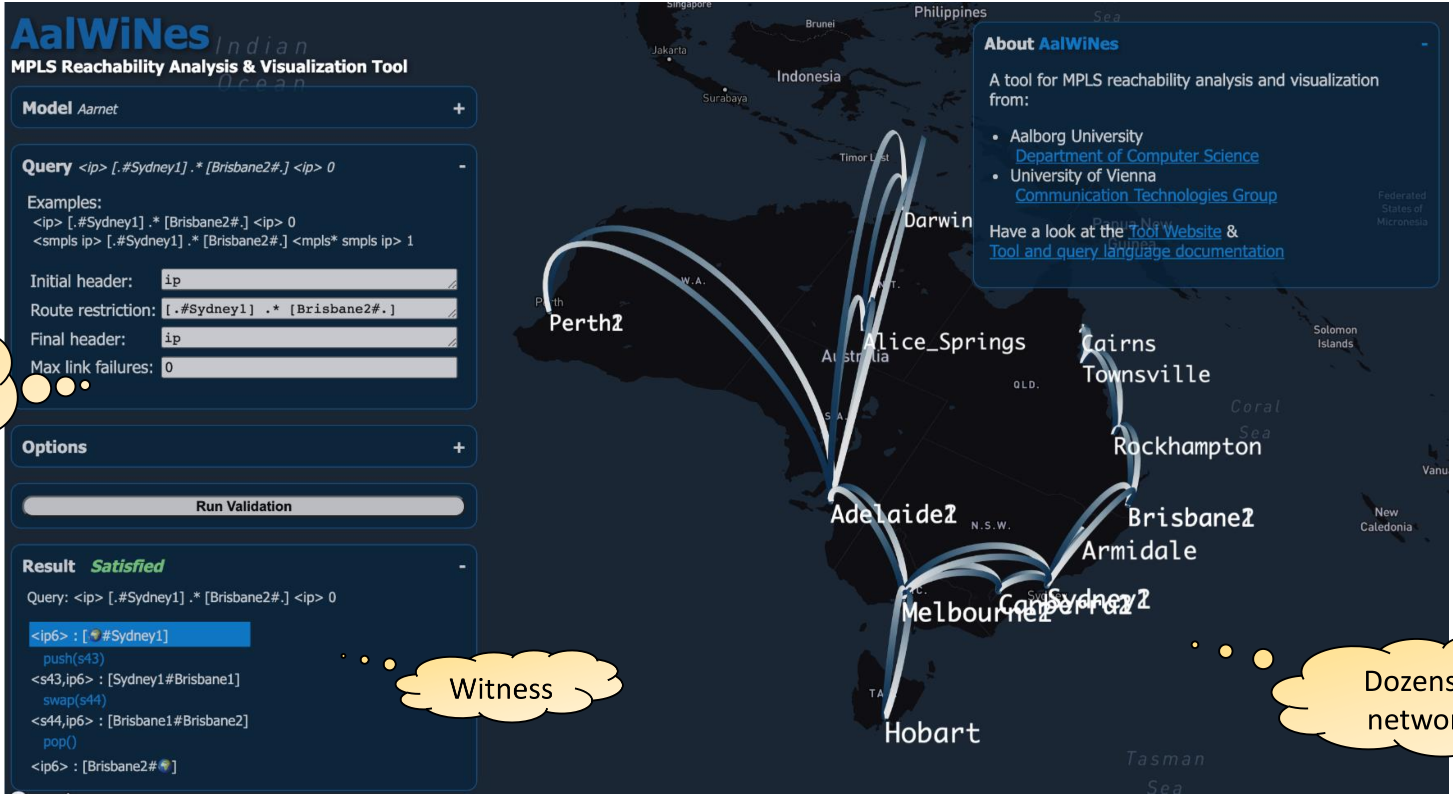

AalWiNes

MPLS Reachability Analysis & Visualization Tool

Model *Aarnet* +

Query *<ip> [.#Sydney1] .* [Brisbane2#.] <ip> 0* -

Examples:
<ip> [.#Sydney1] .* [Brisbane2#.] <ip> 0
<smpls ip> [.#Sydney1] .* [Brisbane2#.] <mpls* smpls ip> 1

Initial header: ip
Route restriction: [.#Sydney1] .* [Brisbane2#.]
Final header: ip
Max link failures: 0

Options +

Run Validation

Result ***Satisfied*** -

Query: <ip> [.#Sydney1] .* [Brisbane2#.] <ip> 0

```
<ip6> : [🌍#Sydney1]
    push(s43)
<s43,ip6> : [Sydney1#Brisbane1]
    swap(s44)
<s44,ip6> : [Brisbane1#Brisbane2]
    pop()
<ip6> : [Brisbane2#🌍]
```

About AalWiNes -

A tool for MPLS reachability analysis and visualization from:

- Aalborg University
  Department of Computer Science
- University of Vienna
  Communication Technologies Group

Have a look at the Tool Website & Tool and query language documentation

Query: regular expression

Witness

Dozens of networks

AalWiNes

Part 1: Parses query and constructs Push-Down System (PDS)

- In Python 3

Part 2: Reachability analysis of constructed PDS

- Using ***Moped*** tool

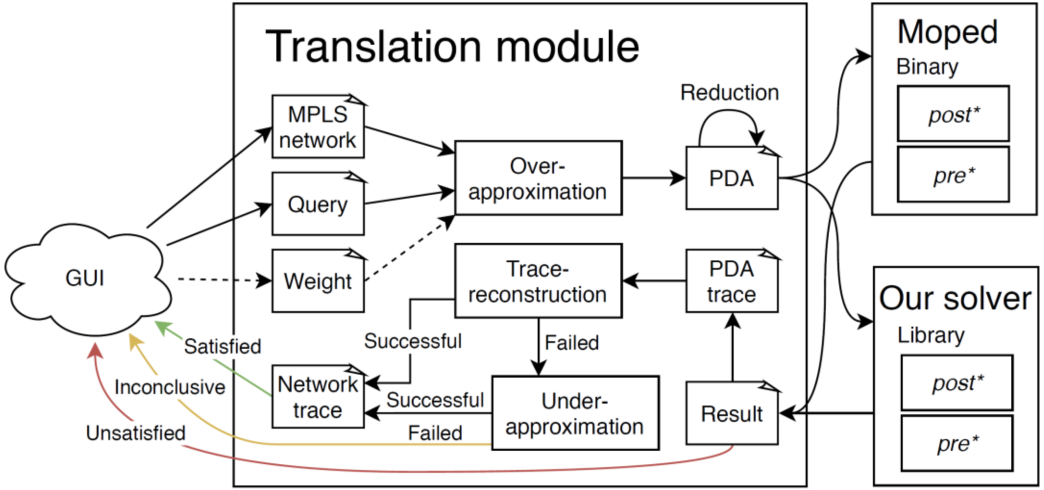

Resp. our new weighted extension and much faster implementation in C++.

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*

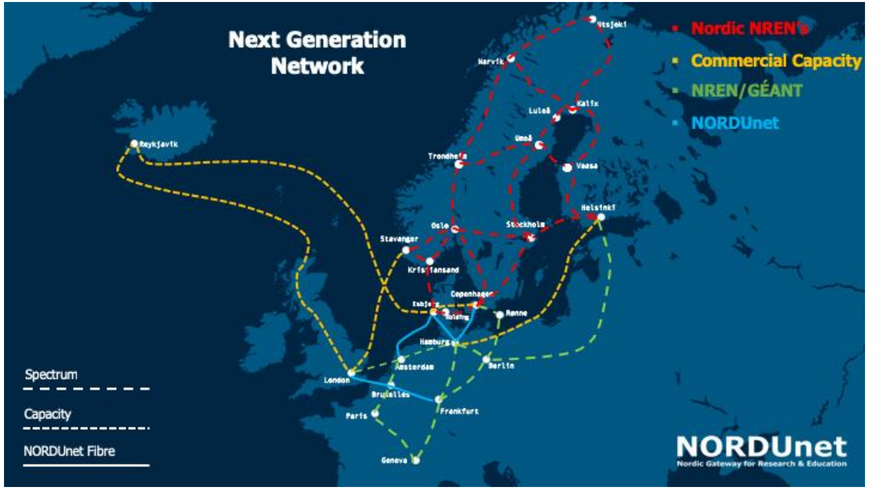

Patent filed for R-MPLS

Some news coverage

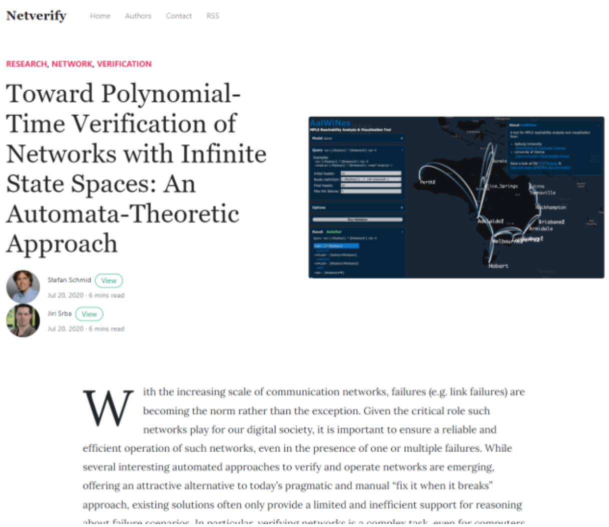

Netverify Home Authors Contact RSS

RESEARCH, NETWORK, VERIFICATION

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

Stefan Schmid View
Jul 20, 2020 · 6 mins read

Jiri Srba View
Jul 20, 2020 · 6 mins read

With 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 ensure a reliable and efficient operation of such networks, even in the presence of one or multiple failures. While several interesting automated approaches to verify and operate networks are emerging, offering an attractive alternative to today's pragmatic and manual "fix it when it breaks" approach, existing solutions often only provide a limited and inefficient support for reasoning about failure scenarios. In particular, verifying networks is a complex task, even for computers

Zero Outage
INDUSTRY STANDARD

NEWS

Home / News / Interviews / Interview with Stefan Schmid, Prof. at University of Vienna

INTERVIEW WITH STEFAN SCHMID, PROF. AT UNIVERSITY OF VIENNA

Roadmap

- WHATIF: Motivation and Context
- Current State and Achievements
- Plans

Next Steps

- Impact:
  - Slowly ready for demos/presentations (e.g., industry,)
  - E.g., at IMAGINE
- Research:
  - Support modelling congestion
  - Tools for quantitative update synthesis

TWIN TRANS-FORMATION
IMAGINE 22
Die Konferenz der IKT-Community Österreichs.
Bundesministerium Klimaschutz, Umwelt, Energie, Mobilität, Innovation und Technologie
23. JUNI 2022
GEMEINSAM MIT DEM AUSTRIAN DATA DAY21

Thank you!

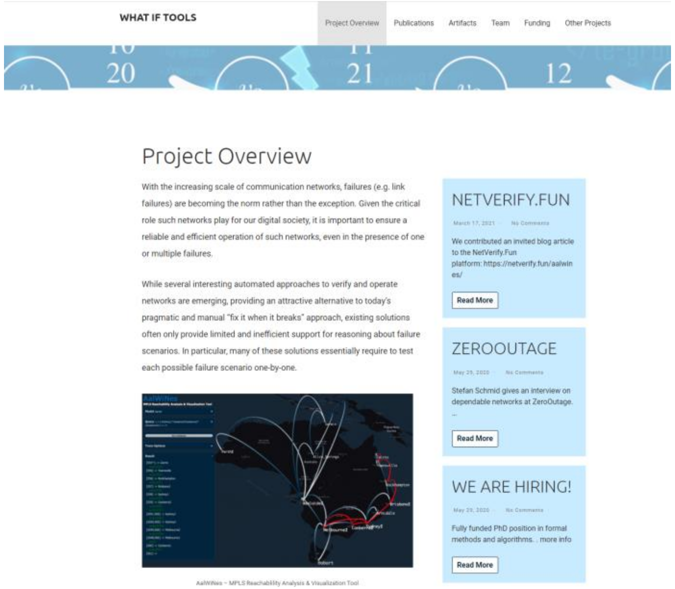

Project Overview

With 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 ensure a reliable and efficient operation of such networks, even in the presence of one or multiple failures.

While several interesting automated approaches to verify and operate networks are emerging, providing an attractive alternative to today's pragmatic and manual "fix it when it breaks" approach, existing solutions often only provide limited and inefficient support for reasoning about failure scenarios. In particular, many of these solutions essentially require to test each possible failure scenario one-by-one.

AalWiNes – MPLS Reachability Analysis & Visualization Tool

NETVERIFY.FUN

March 17, 2021 · No Comments

We contributed an invited blog article to the NetVerify.Fun platform: https://netverify.fun/aalwines/

Read More

ZEROOUTAGE

May 29, 2020 · No Comments

Stefan Schmid gives an interview on dependable networks at ZeroOutage. ...

Read More

WE ARE HIRING!

May 29, 2020 · No Comments

Fully funded PhD position in formal methods and algorithms. . more info

Read More

Questions?