

# Self-Driving Networks: Vision, Enablers, Challenges

Stefan Schmid

“We cannot direct the wind,  
but we can adjust the sails.”

(Folklore)

# It`s a Great Time to Be a Networking Researcher!



Rhone and Arve Rivers,  
Switzerland

Credits: George Varghese

# Flexibilities

Along 3 Dimensions



Passau, Germany

Inn, Donau, Ilz

# Flexibilities

Along 3 Dimensions



# It's High Time

Increasing Traffic, Stringent Requirements




Datacenters (“hyper-scale”)



+network

Interconnecting networks:  
a **critical infrastructure**  
of our digital society.

Traffic  
Growth



Source: Facebook

# Requirements vs Reality

Today, dependability requirements stand in contrast with reality:

## 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-thumbbed 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.

## 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 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.

Even tech-savvy companies struggle:



Credits:

Laurent Vanbever, Nate Foster



# Requirements vs Reality

Today, dependability requirements stand in contrast with reality:

## 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-thumbbed 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.

## 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 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.

Even tech-savvy companies struggle:



**Mainly:  
human  
errors!**


Credits:

Laurent Vanbever, Nate Foster

# Reason: Complexity

Especially Under Failures (Policy Compliance)

Example: BGP in  
Microsoft  
datacenter






# Reason: Complexity

Especially Under Failures (Policy Compliance)


Example: BGP in  
Microsoft  
datacenter



# Reason: Complexity

Especially Under Failures (Policy Compliance)

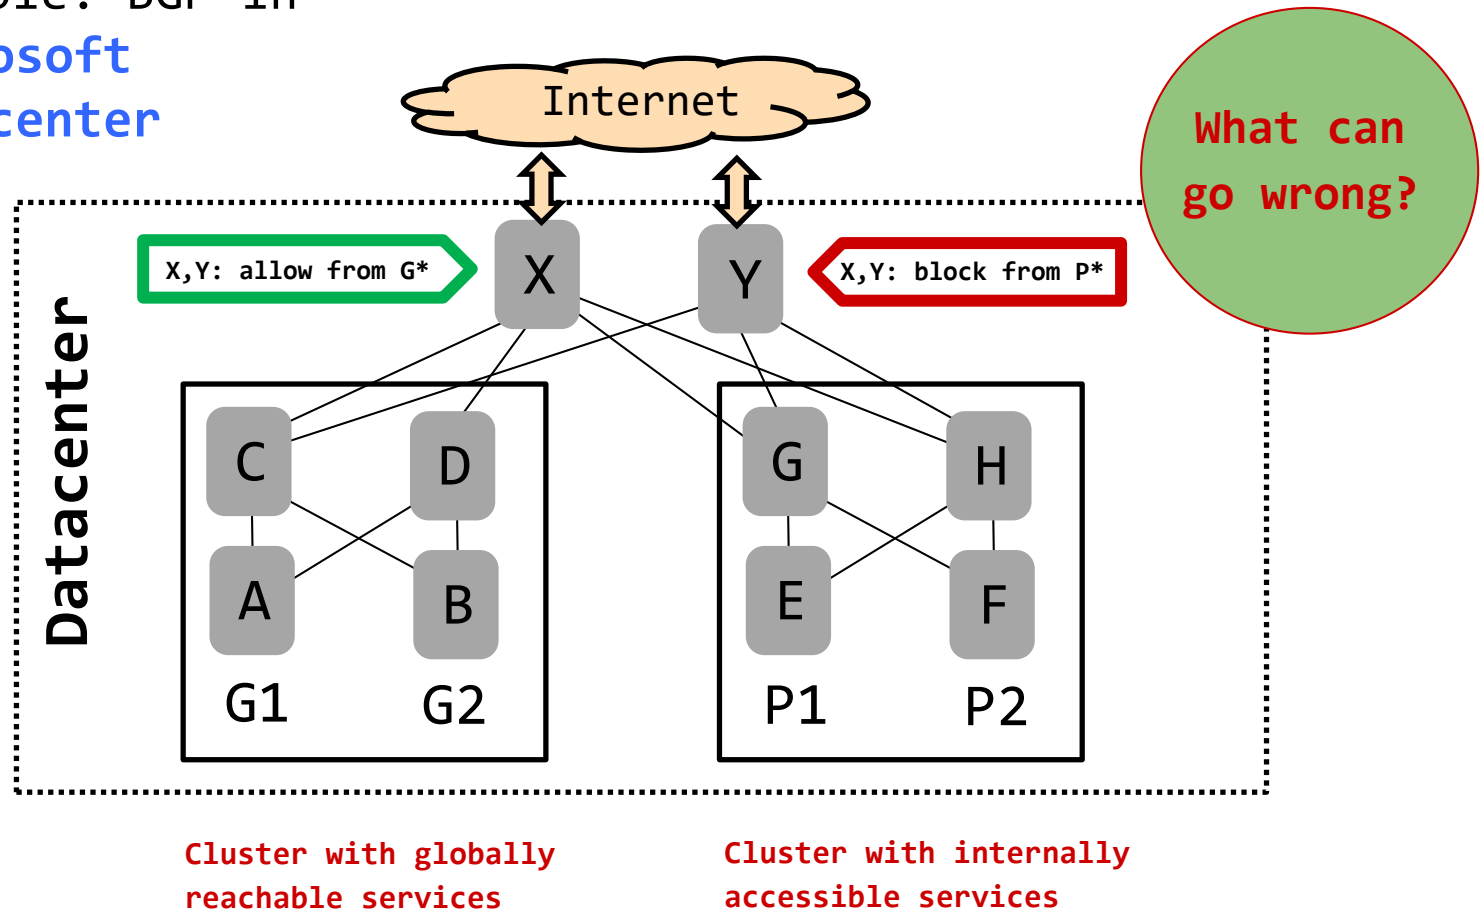
Example: BGP in  
Microsoft  
datacenter



# Reason: Complexity

Especially Under Failures (Policy Compliance)


Example: BGP in  
Microsoft  
datacenter



# Reason: Complexity

Especially Under Failures (Policy Compliance)

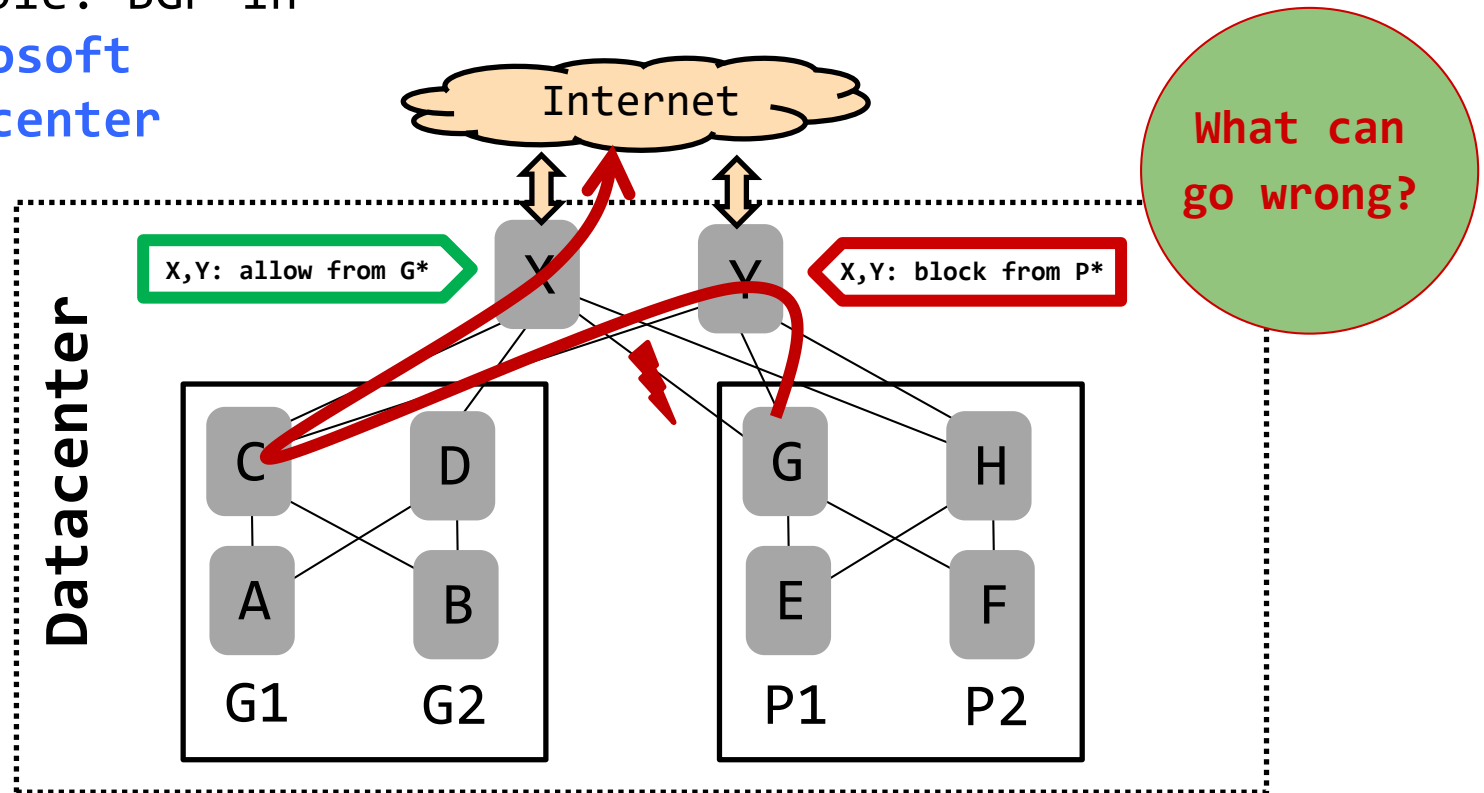
Example: BGP in  
Microsoft  
datacenter



# Reason: Complexity


Especially Under Failures (Policy Compliance)

Example: BGP in  
Microsoft  
datacenter




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.)

# Admin's Responsibilities






# Admin's Responsibilities



→ Reachability: Can traffic from ingress port A reach B?

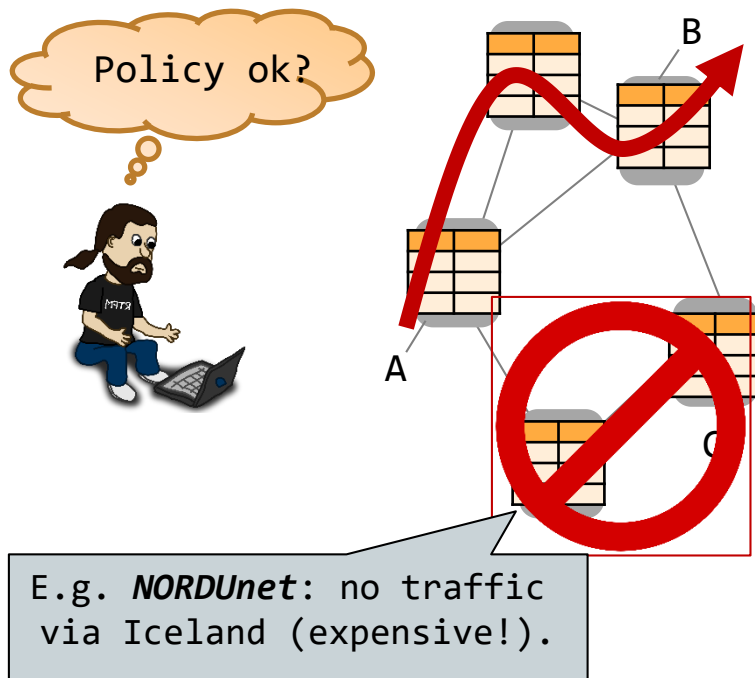
# Admin's Responsibilities



→ Reachability: Can traffic from ingress port A reach B?

→ Loop-freedom: Do forwarding rules imply loop-free routes?

# Admin's Responsibilities




→ Reachability: Can traffic from ingress port A reach B?

→ Loop-freedom: Do forwarding rules imply loop-free routes?

→ Policy: Does traffic from A to B never go via C?

# Admin's Responsibilities




→ Reachability: Can traffic from ingress port A reach B?

→ Loop-freedom: Do forwarding rules imply loop-free routes?

→ Policy: Does traffic from A to B never go via C?

→ Waypoint enforcement: Is traffic from A to B always routed via a node C (e.g., an IDS)?

# Admin's Responsibilities



→ Reachability: Can traffic from ingress port A reach B?

→ Loop-freedom: Do forwarding rules imply loop-free routes?


→ Policy: Does traffic from A to B never go via C?

→ Waypoint enforcement: Is traffic from A to B always routed via a node C (e.g., an IDS)?


*... and everything even under multiple failures?!*

A Modern Approach:

# Automated Whatif Analysis



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



local FFT	Out-I	In-Label	Out-I	op
$\tau_{v_2}$	$(v_2, v_3)$	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
$\tau_{v_2}$	$(v_2, v_3)$	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

$pX \Rightarrow qXX$

$pX \Rightarrow qYX$

$qY \Rightarrow rYY$

$rY \Rightarrow r$

$rX \Rightarrow pX$

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

**Formal language**  
which supports  
**automated analysis**




A Modern Approach:

# Automated Whatif Analysis



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



local FFT	Out-I	In-Label	Out-I	op
$\tau_{v_2}$	$(v_2, v_3)$	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
$\tau_{v_2}$	$(v_2, v_3)$	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

On request or regularly.

Interpretation

$pX \Rightarrow qXX$

$pX \Rightarrow qYX$

$qY \Rightarrow rYY$

$rY \Rightarrow r$


$rX \Rightarrow pX$

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


Formal language  
which supports  
**automated analysis**

A Modern Approach:

# Automated Whatif Analysis



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



local FFT	Out-I	In-Label	Out-I	op
$\tau_{v_2}$	$(v_2, v_3)$	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
$\tau_{v_2}$	$(v_2, v_3)$	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

On request or regularly.

Fix/synthesize

$pX \Rightarrow qXX$

$pX \Rightarrow qYX$

$qY \Rightarrow rYY$

$rY \Rightarrow r$

$rX \Rightarrow pX$

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


Formal language  
which supports  
**automated analysis**

A Modern Approach:

# Automated Whatif Analysis



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



local FFT	Out-I	In-Label	Out-I	op
$\tau_{v_2}$	$(v_2, v_3)$	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
$\tau_{v_2}$	$(v_2, v_3)$	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

On request or regularly.

Fix/synthesize

$pX \Rightarrow qXX$

$pX \Rightarrow qYX$

$qY \Rightarrow rYY$

$rY \Rightarrow r$

$rX \Rightarrow pX$

Router **configuration**  
(Cisco, Juniper,



1 language  
supports  
**red analysis**

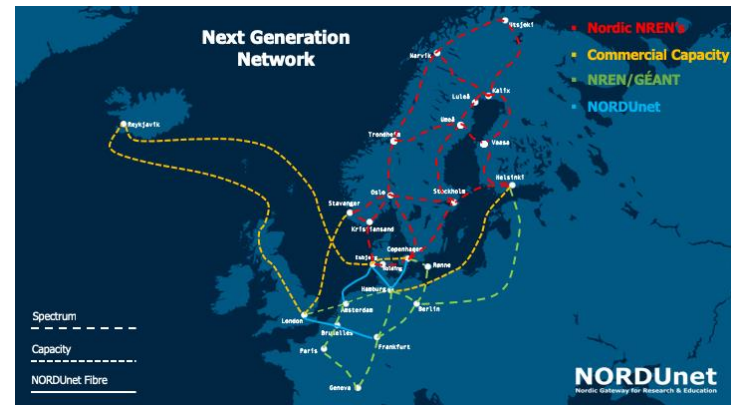
Challenge:

# Hard Even for Computers?

→ NORDUnet: provider for Nordic countries

→ 24 MPLS routers, running Juniper OS

→ More than 30,000 labels!



Case Study:


# Whatif Analysis of MPLS

**FAST!**

What if?!



FT	In-I	In-Label	Out-I	op
$\tau_{v_1}$	$in_1$	$\perp$	$(v_1, v_2)$	$push(10)$
	$in_2$	$\perp$	$(v_1, v_2)$	$push(20)$
$\tau_{v_2}$	$(v_1, v_2)$	10	$(v_2, v_3)$	$swap(11)$
	$(v_1, v_2)$	20	$(v_2, v_3)$	$swap(21)$
$\tau_{v_3}$	$(v_2, v_3)$	11	$(v_3, v_4)$	$swap(12)$
	$(v_2, v_3)$	21	$(v_3, v_4)$	$swap(22)$
$\tau_{v_4}$	$(v_2, v_3)$	11	$(v_3, v_4)$	$swap(12)$
	$(v_2, v_3)$	21	$(v_3, v_4)$	$swap(22)$
$\tau_{v_5}$	$(v_3, v_4)$	12	$out_1$	$pop$
	$(v_3, v_4)$	22	$out_1$	$pop$
$\tau_{v_6}$	$(v_2, v_3)$	40	$(v_5, v_6)$	$pop$
	$(v_2, v_3)$	30	$(v_6, v_7)$	$swap(31)$
$\tau_{v_7}$	$(v_5, v_6)$	30	$(v_6, v_7)$	$swap(31)$
	$(v_5, v_6)$	61	$(v_6, v_7)$	$swap(62)$
$\tau_{v_8}$	$(v_6, v_7)$	71	$(v_7, v_8)$	$swap(72)$
	$(v_6, v_7)$	31	$(v_7, v_8)$	$pop$
$\tau_{v_9}$	$(v_6, v_7)$	62	$(v_7, v_8)$	$swap(11)$
	$(v_6, v_7)$	72	$(v_7, v_8)$	$swap(22)$
$\tau_{v_{10}}$	$(v_3, v_4)$	22	$out_2$	$pop$
	$(v_3, v_4)$	22	$out_2$	$pop$



local FFT	Out-I	In-Label	Out-I	op
$\tau_{v_2}$	$(v_2, v_3)$	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
$\tau_{v_2}$	$(v_2, v_3)$	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

$pX \Rightarrow qXX$

$pX \Rightarrow qYX$

$qY \Rightarrow rYY$

$rY \Rightarrow r$

$rX \Rightarrow pX$

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

MPLS is a **stack-based pushdown system**: can solve with **finite automata-theory**

Case Study:


# Whatif Analysis of MPLS

**FAST!**

What if?!



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



local FFT	Out-I	In-Label	Out-I	op
$\tau_{v_2}$	$(v_2, v_3)$	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
$\tau_{v_2}$	$(v_2, v_3)$	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

$pX \Rightarrow qXX$

$pX \Rightarrow qYX$

$qY \Rightarrow rYY$

$rY \Rightarrow r$

$rX \Rightarrow pX$

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

MPLS is a **stack-based pushdown system**: can solve with **finite automata-theory**

**Counting for congestion? Yes: weighted automata.**



# The 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#]`

Witness

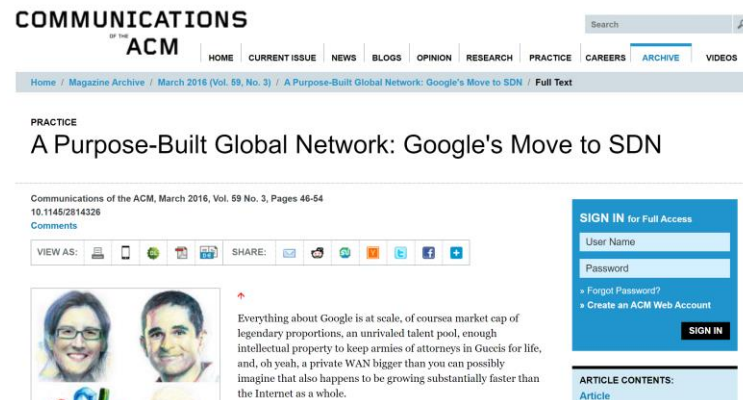
Dozens of networks

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](#)


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

# Another Benefit of Self-Driving Networks: More Adaptive Operation

- Automation and programmability: enables more **adaptable networks**
- 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




## Another Benefit of Self-Driving Networks: More Adaptive Operation



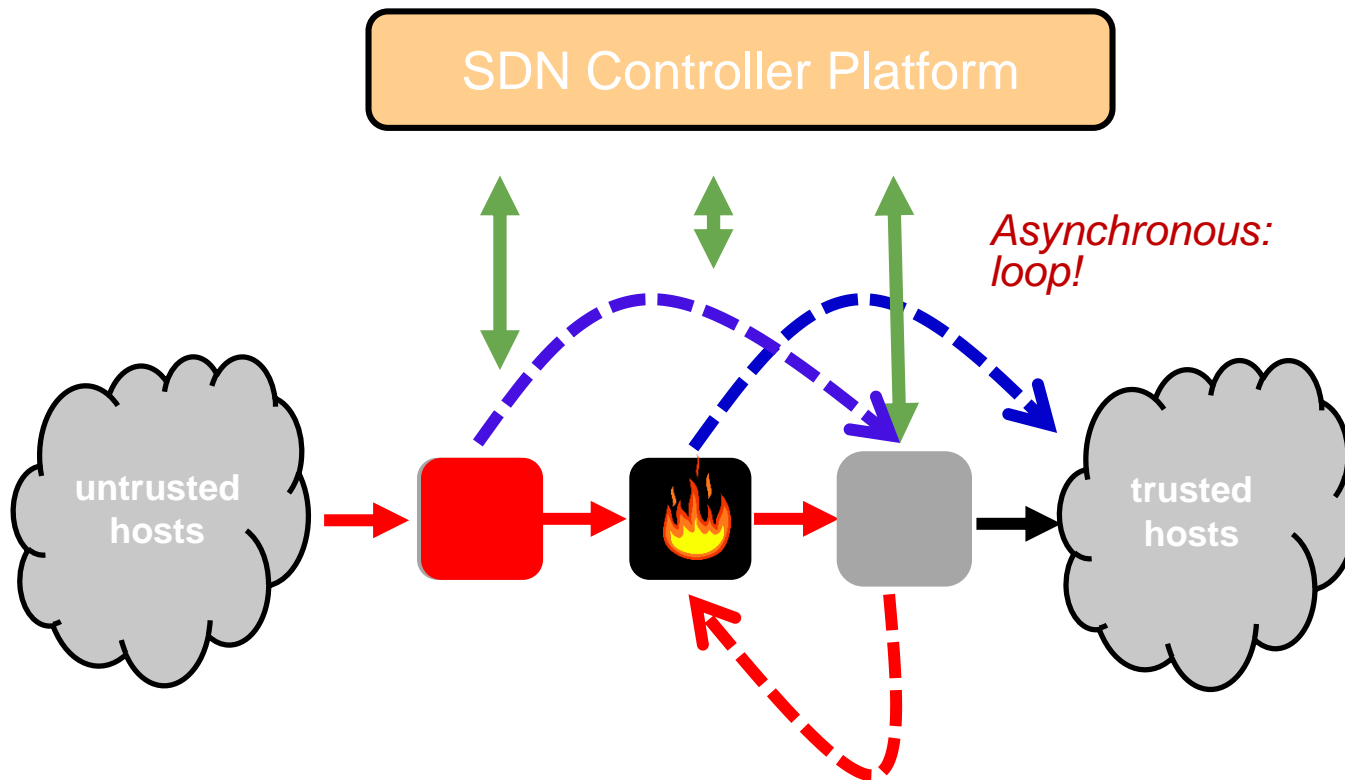
*SDN outsources and consolidates control:  
direct definition of forwarding tables!*

## Another Benefit of Self-Driving Networks: More Adaptive Operation



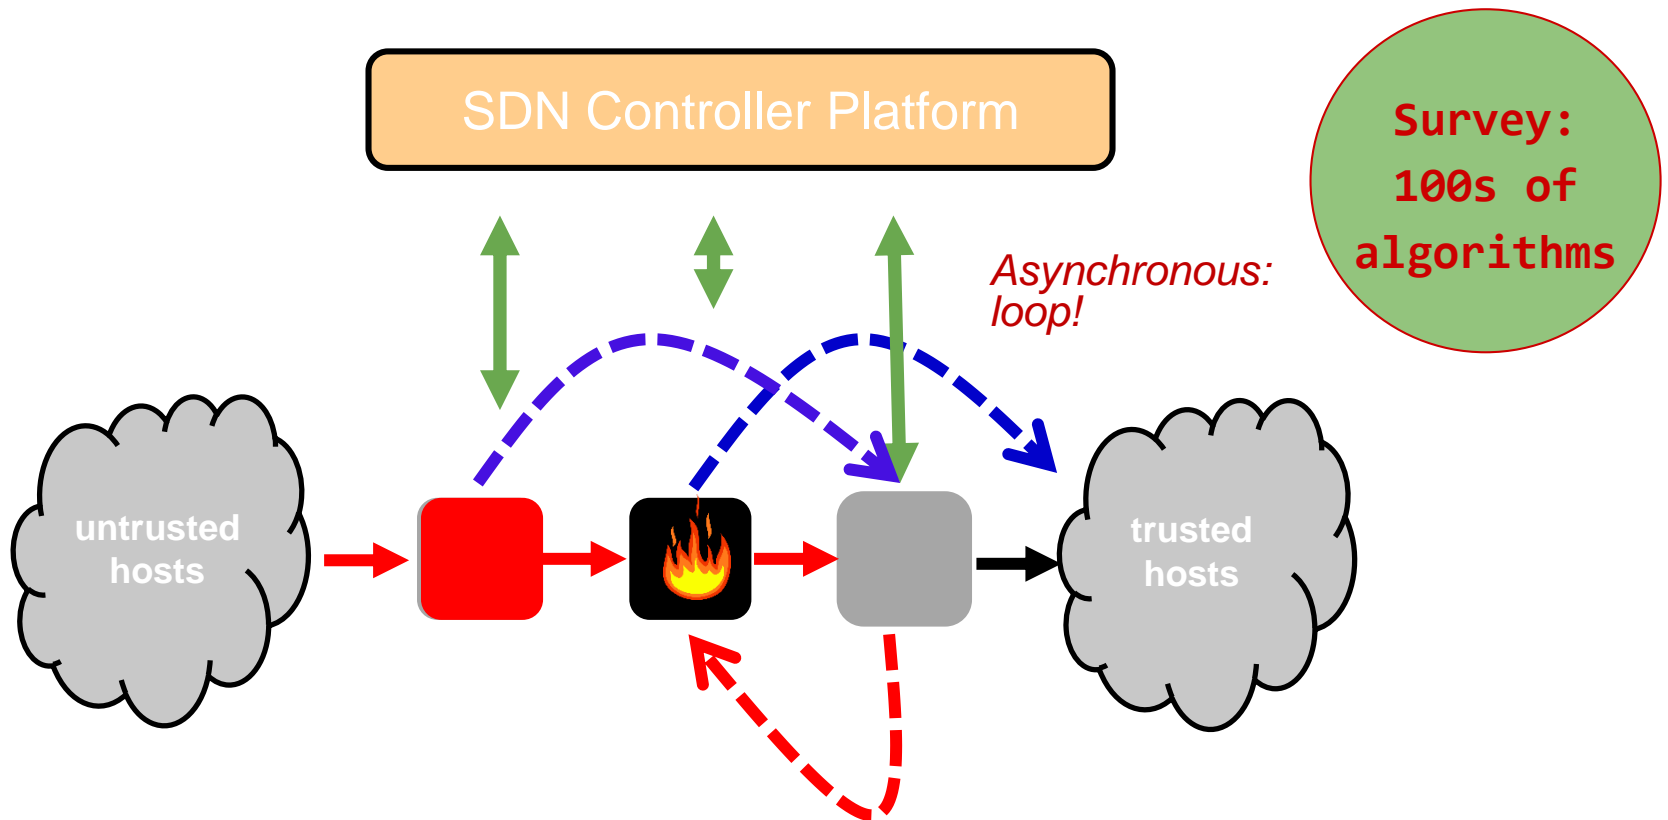
*Can programmatically change route to blue route.*

## Another Benefit of Self-Driving Networks: More Adaptive Operation



*Updates are asynchronous: can be challenging to update the network configuration consistently!*

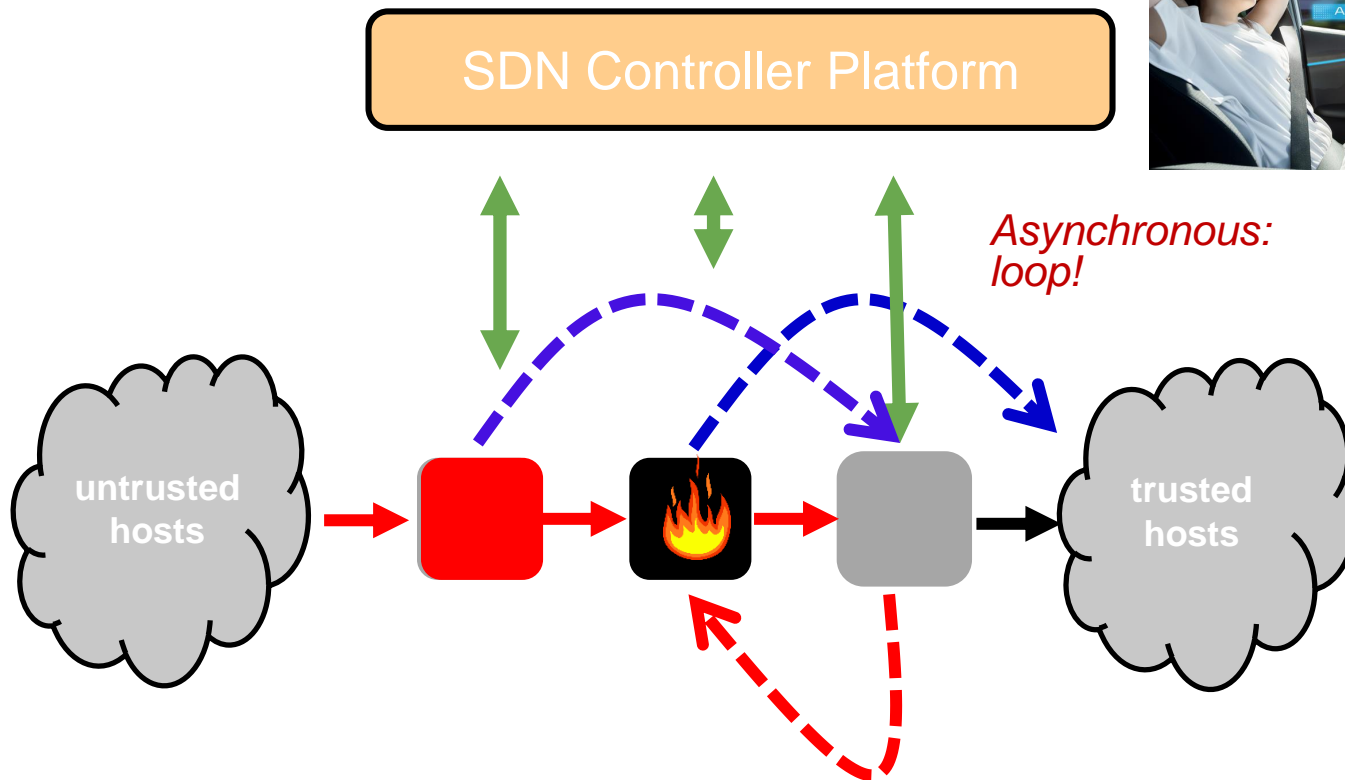
## Another Benefit of Self-Driving Networks: More Adaptive Operation



*Updates are asynchronous: can be challenging to update the network configuration consistently!*



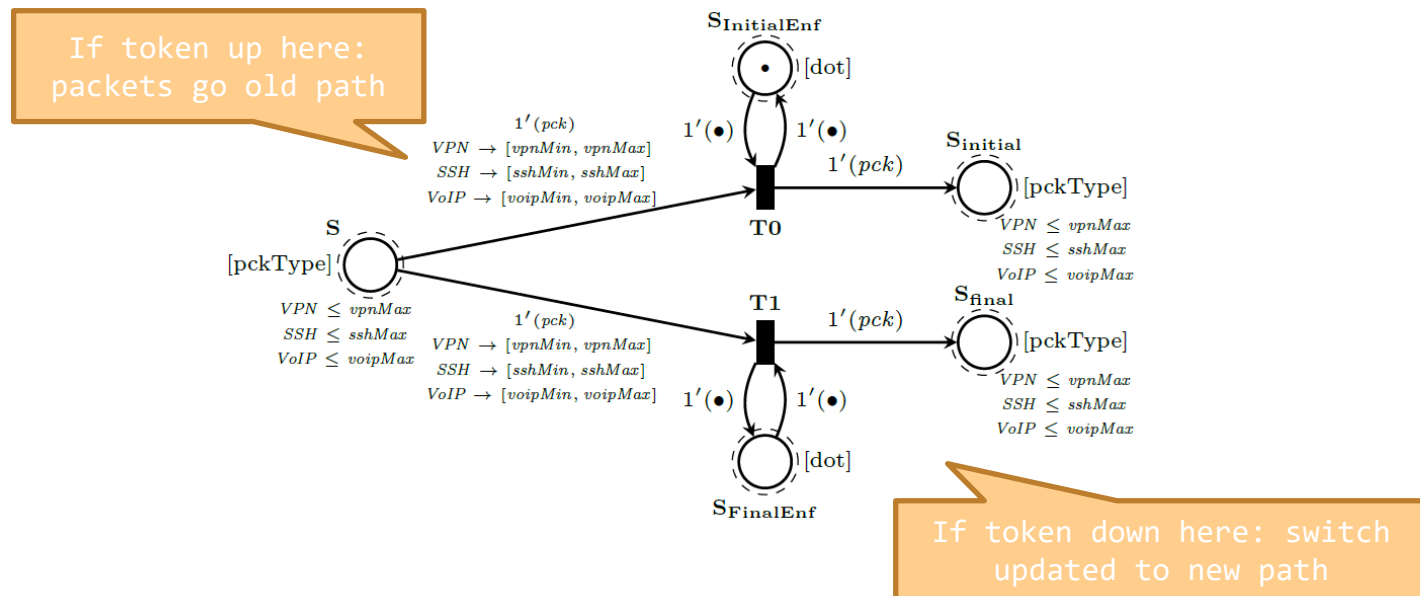
# Again: Formal Approaches Enable More Automated Approaches



*Emerging tools allow to verify and synthesize update schedules for a wide range of networks and objectives.*

# Example: The Latte Tool

- SDN can be modelled as distributed system: **Petri net**
- Latte allows to verify and synthesize update schedules
- Example: Gadget to model switches:




Adaptions also on Lower Layers:  
Emerging Topology Programming

Let's go back to datacentre use case:

# Moore's Law of Datacenters

- Recall: **explosive growth** of demand
- Problem: network equipment reaching capacity limits
  - Transistor density rates stalling
  - “End of **Moore's Law** in networking”
- Hence: more equipment, larger networks
- Resource intensive and: **inefficient**




[1] Source: Microsoft, 2019

Annoying for companies,  
**opportunity** for researchers

# Root Cause


Fixed and Demand-Oblivious Topology

How to interconnect?



# Root Cause

## Fixed and Demand-Oblivious Topology




# Root Cause

## Fixed and Demand-Oblivious Topology




Highway which ignores  
actual traffic:  
**frustrating!**



Many flavors,  
but in common:  
fixed and  
**oblivious** to  
actual demand.

# A Vision


Flexible and Demand-Aware Topologies





# A Vision


## Flexible and Demand-Aware Topologies



## Flexible and Demand-Aware Topologies

demand  
matrix:

	1	2	3	4	5	6	7	8
1					Red			
2						Red		
3							Red	
4								Red
5	Red							
6		Red						
7			Red					
8				Red				




# A Vision

## Flexible and Demand-Aware Topologies

Matches demand


demand  
matrix:

	1	2	3	4	5	6	7	8
1					■			
2						■		
3							■	
4								■
5	■							
6		■						
7			■					
8				■				



# A Vision

## Flexible and Demand-Aware Topologies




# A Vision

## Flexible and Demand-Aware Topologies

Matches demand

new  
demand:

	1	2	3	4	5	6	7	8
1		■						
2	■							
3				■				
4			■					
5						■		
6					■			
7							■	
8								■



e.g.,  
mirrors

new flexible  
interconnect

# A Vision


## Flexible and Demand-Aware Topologies



### Self-Adjusting Networks

new  
demand:

	1	2	3	4	5	6	7	8
1		■						
2	■							
3				■				
4			■					
5						■		
6					■			
7							■	
8								■



e.g.,  
mirrors



new flexible  
interconnect

# The Motivation


Much Structure in the Demand

Empirical studies:

traffic matrices **sparse** and **skewed**




traffic **bursty** over time



Our **hypothesis**: can  
be exploited.

# Sounds Crazy? Emerging Enabling Technology.



Photonics

H2020:

**“Photonics one of only five  
key enabling technologies  
for future prosperity.”**

US National Research Council:

**“Photons are the new  
Electrons.”**





# Enabler

## Novel Reconfigurable Optical Switches

→ **Spectrum** of prototypes

- Different sizes, different reconfiguration times
- From our last years' ACM **SIGCOMM** workshop OptSys



Prototype 1



Prototype 2




Prototype 3

# Example

## Optical Circuit Switch


- Optical Circuit Switch rapid adaption of physical layer
  - Based on rotating mirrors



## Optical Circuit Switch

By Nathan Farrington, SIGCOMM 2010

# The Big Picture



# Unique Position

Demand-Aware, Self-Adjusting Systems


Everywhere, but mainly  
in software



Algorithmic trading




Recommender systems



Neural networks

VS

Our focus:  
in hardware

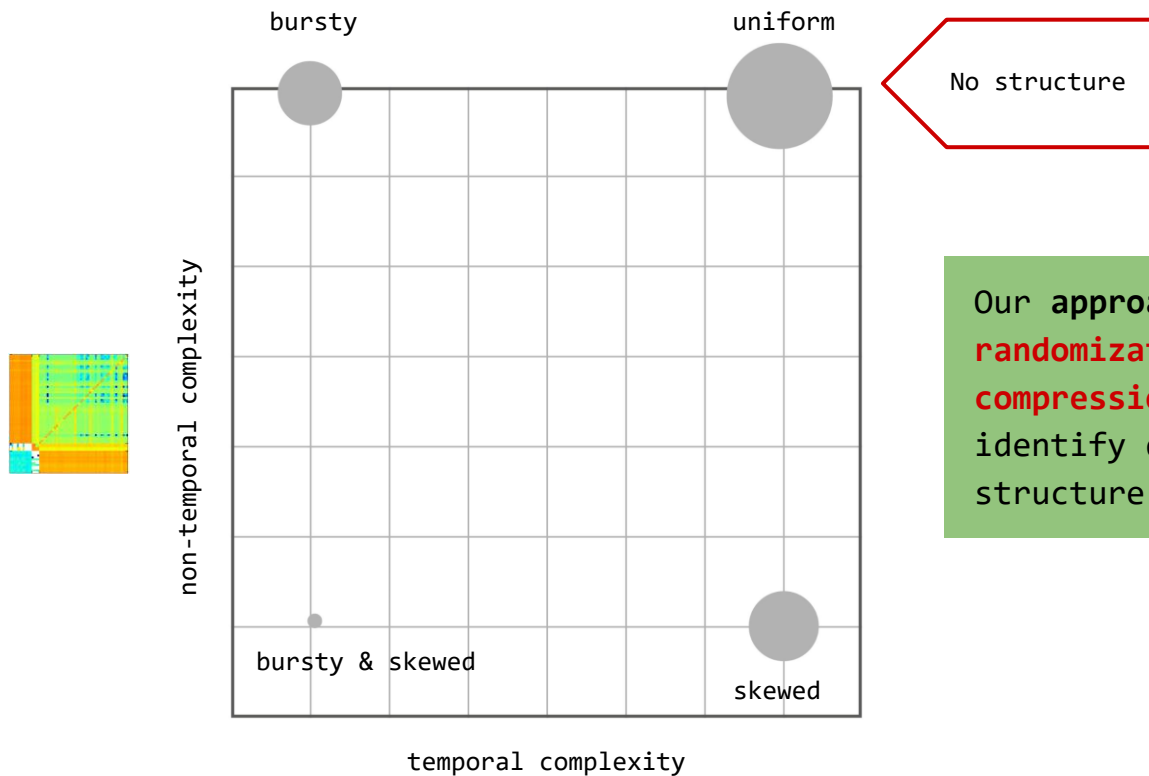


Question 1:


How to Quantify  
such “Structure”  
in the Demand?

## An Information-Theoretic Approach

# Complexity Map




Our **approach**: iterative **randomization and compression** of trace to identify dimensions of structure.




## An Information-Theoretic Approach

# Complexity Map




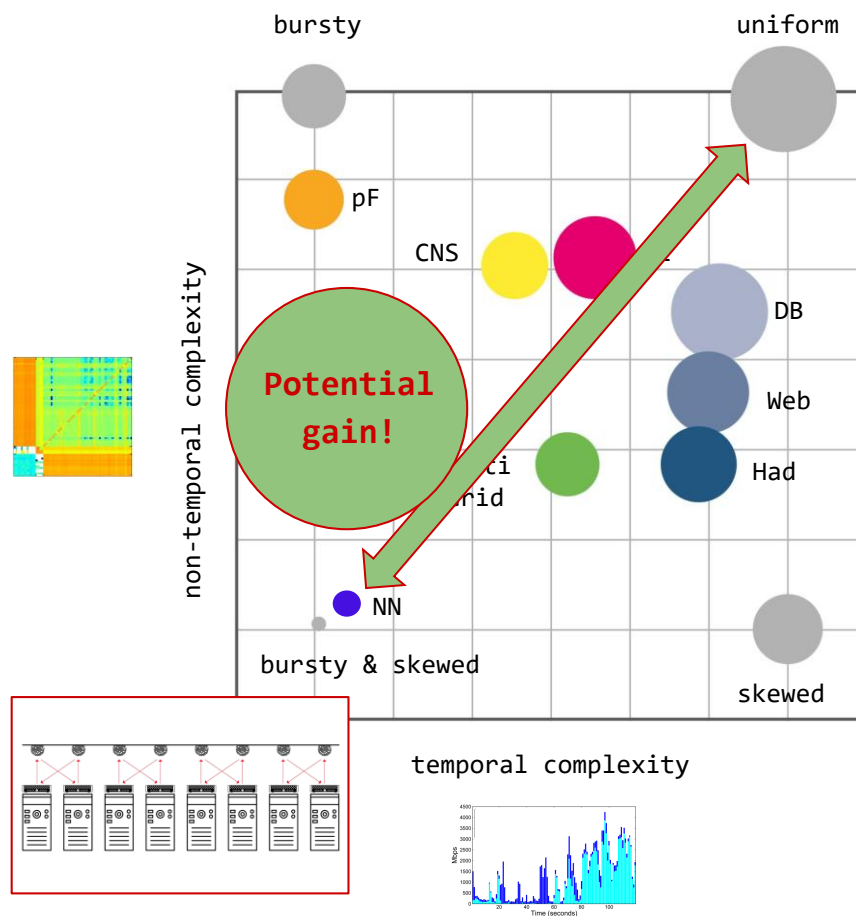
Our approach: iterative  
**randomization and  
compression** of trace to  
identify dimensions of  
structure.

**Different  
structures!**



## An Information-Theoretic Approach

# Complexity Map



Our approach: iterative **randomization** and **compression** of trace to identify dimensions of structure.




Question 2:

Given This Structure,  
What Can Be Achieved?  
Metrics and Algorithms?

A first insight: entropy of the demand.

# Connection to Datastructures

## Self-adjusting BST




26


Our Approach:

# Connection to Datastructures & Coding


Traditional BST  
(Worst-case coding)



Demand-aware BST  
(Huffman coding)




Self-adjusting BST  
(Dynamic Huffman coding)



More structure: improved **access cost** / shorter **codes**

# Connection to Datastructures & Coding

## Self-adjusting BST (Dynamic Huffman coding)




26


Our Approach:

# Connection to Datastructures & Coding


Traditional BST  
(Worst-case coding)



Demand-aware BST  
(Huffman coding)






Self-adjusting BST  
(Dynamic Huffman coding)



More than  
an analogy!

More structure: improved **access cost** / shorter **codes**




Similar **benefits**?


Our Approach:

# Connection to Datastructures & Coding


Traditional BST  
(Worst-case coding)






Demand-aware BST  
(Huffman coding)



Self-adjusting BST  
(Dynamic Huffman coding)



More than  
an analogy!



Generalize methodology:  
... and transfer  
entropy bounds and  
algorithms of data-  
structures to networks.


Results, e.g.:  
Demand-aware networks  
of asymptotically  
optimal route lengths.

Reduced expected **route lengths!**

# Challenges of Software-Defined and Self-Driving Networks


# Example: Security

- Trend: SDN deployed “*in software*”
- E.g., virtual switches in datacenters





# Virtual Switches



Virtual switches reside in the **server's virtualization layer** (e.g., Xen's Dom0).  
Goal: provide connectivity and isolation.

# Complexity: Parsing

Ethernet

LLC

VLAN

MPLS

IPv4

ICMPv4

TCP

UDP

ARP

SCTP

IPv6

ICMPv6

IPv6 ND

GRE

LISP

VXLAN

PBB

IPv6 EXT HDR

TUNNEL-ID

IPv6 ND

IPv6 EXT HDR

IPv6HOPPTS

IPv6ROUTING

IPv6Fragment


IPv6DESTOPT

IPv6ESP

IPv6 AH


RARP

IGMP

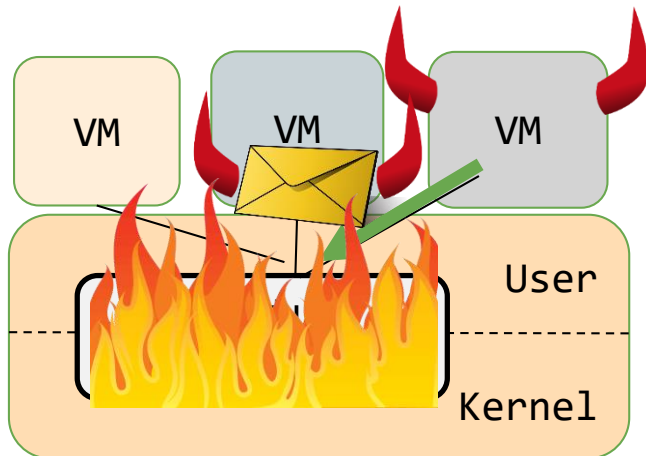


Parser directly faces attacker and vSwitch runs with high security privileges.


# Bears Risks!



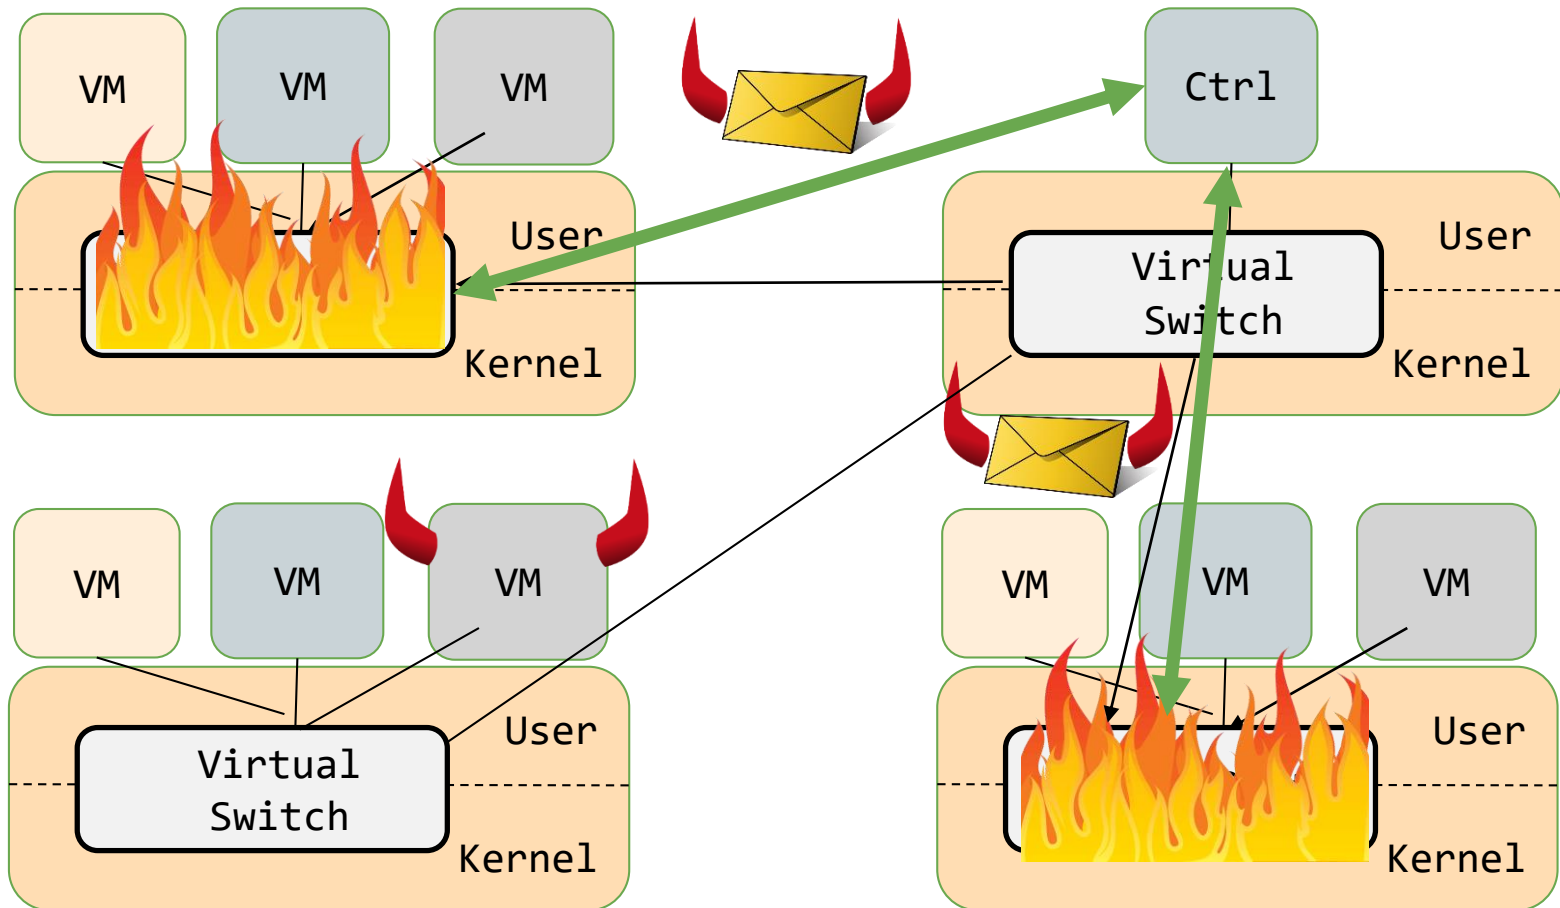
# Bears Risks!



# Bears Risks!



# Bears Risks!



# Limits of Automation?

- What should not or cannot be automated?
- Can networks detect themselves, when the need “help from the operator”?

# Conclusion

- A vision: self-driving networks
- Example 1: policy-compliant networks
  - self-verifying
  - self-repairing
- Example 2: demand-aware topologies
- On both fronts: tip of the iceberg!
- E.g., self-adjusting networks further supported by telemetry (data) and AI (e.g., prediction)



Thank you!



# References

## [AalWiNes: A Fast and Quantitative What-If Analysis Tool for MPLS Networks](#)

Peter Gjøøl 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.

## [On the Complexity of Traffic Traces and Implications](#)

Chen Avin, Manya Ghobadi, Chen Griner, and Stefan Schmid.

ACM **SIGMETRICS** and ACM Performance Evaluation Review (**PER**), Boston, Massachusetts, USA, June 2020.

## [Toward Demand-Aware Networking: A Theory for Self-Adjusting Networks](#) (Editorial)

Chen Avin and Stefan Schmid.

ACM SIGCOMM Computer Communication Review (**CCR**), October 2018.

## [Taking Control of SDN-based Cloud Systems via the Data Plane](#) (Best Paper Award)


Kashyap Thimmaraju, Bhargava Shastry, Tobias Fiebig, Felicitas Hetzelt, Jean-Pierre Seifert, Anja Feldmann, and Stefan Schmid.

ACM Symposium on SDN Research (**SOSR**), Los Angeles, California, USA, March 2018.

# Backup Slides

Case Study:


# MPLS Networks




Original  
Routing

## Case Study:

# MPLS Networks



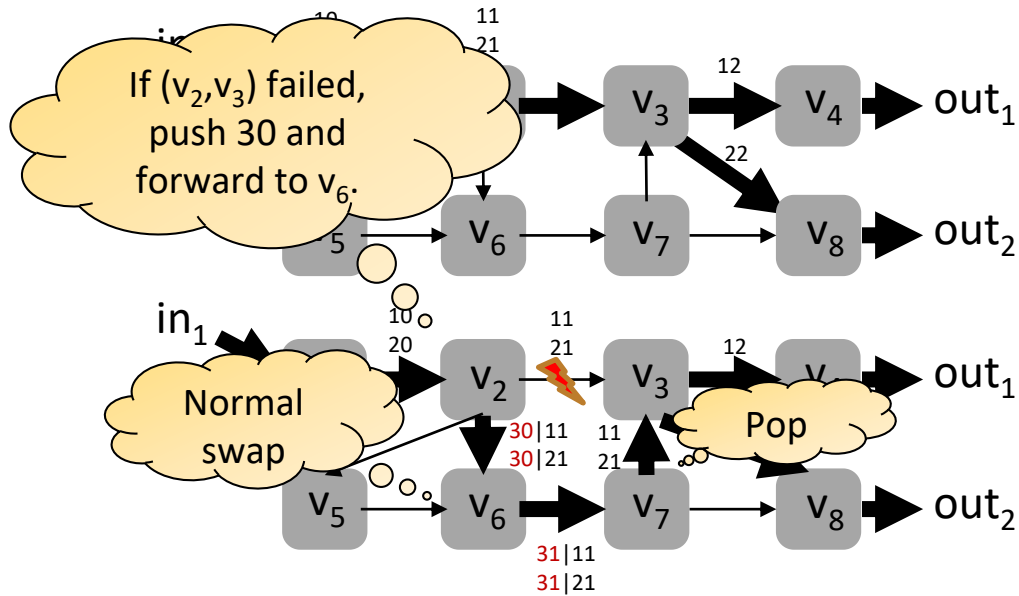
**Original**  
Routing



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

## Case Study:

# MPLS Networks




**Original**  
Routing

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

## Case Study:

# MPLS Networks



**Original**  
Routing

**One failure:**  
push 30: route  
around (v<sub>2</sub>, v<sub>3</sub>)


**Two failures:**  
first push 30:  
route around (v<sub>2</sub>, v<sub>3</sub>)  
**Push recursively**  
40: route around  
(v<sub>2</sub>, v<sub>6</sub>)

# Intuition


Which demand has more structure?

→ Traffic matrices of two different distributed ML applications

→ GPU-to-GPU



VS




Color = communication pair

# Intuition

Which demand has more structure?


→ Traffic matrices of two different distributed ML applications

→ GPU-to-GPU



More uniform

VS



More structure




# Intuition

## Spatial vs temporal structure

→ Two different ways to generate same traffic matrix:

→ Same non-temporal structure

→ Which one has more structure?




# Intuition

## Spatial vs temporal structure

→ Two different ways to generate same traffic matrix:

→ Same non-temporal structure

→ Which one has more structure?




Systematically?

# Trace Complexity

Information-Theoretic Approach


“Shuffle&Compress”



# Trace Complexity

Information-Theoretic Approach

“Shuffle&Compress”




Increasing complexity (systematically randomized)

More structure (compresses better)

# Trace Complexity

Information-Theoretic Approach


“Shuffle&Compress”



# Trace Complexity

Information-Theoretic Approach

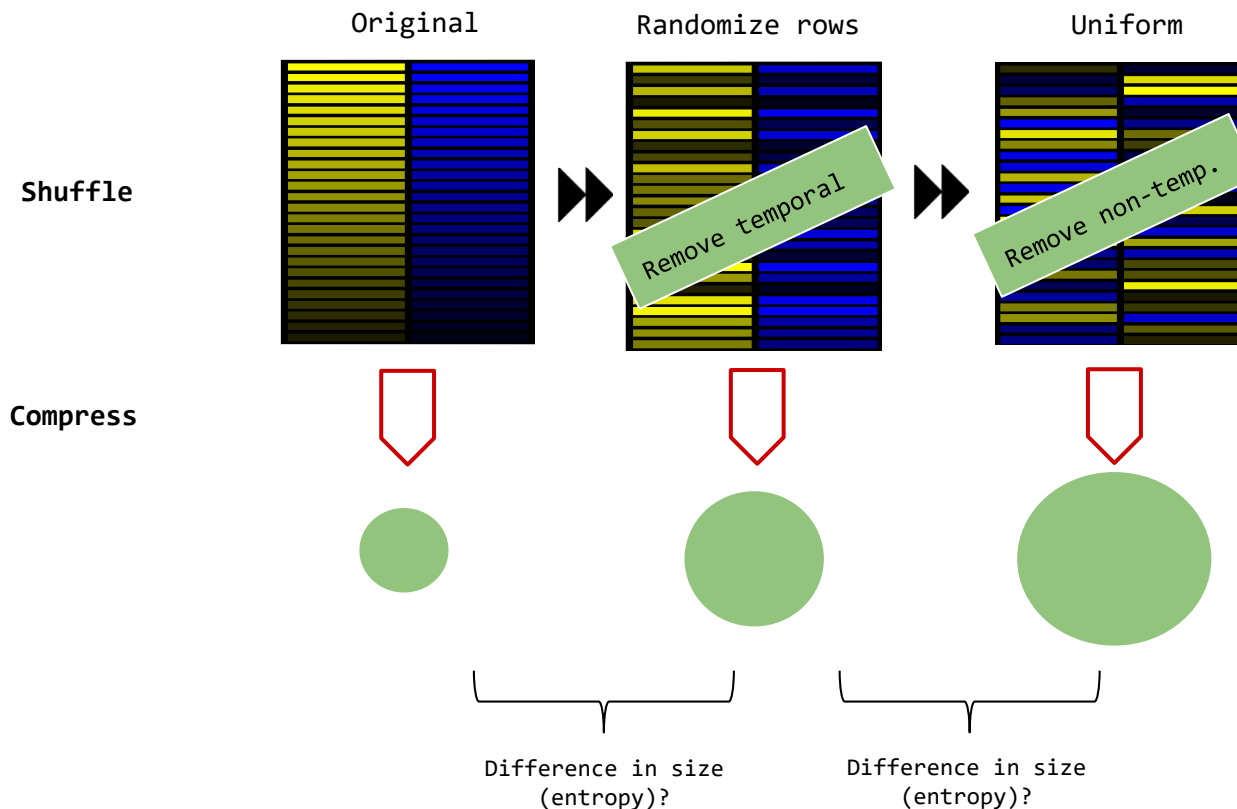
“Shuffle&Compress”



# Trace Complexity

Information-Theoretic Approach

“Shuffle&Compress”



Can be used to define  
2-dimensional  
**complexity map!**


# Bonus Material



Hogwarts Stair



# Bonus Material



Golden Gate Zipper

# Bonus Material

07 May 2021 | 16:55 GMT

## Reconfigurable Optical Networks Will Move Supercomputer Data 100X Faster

Newly designed HPC network cards and software that reshapes topologies on-the-fly will be key to success

By Michelle Hampson



Data illustration: Shutterstock

In HPC