


Can we trust our computer networks?

Stefan Schmid (Uni Wien)





The Internet 50 Years Ago



- *Connectivity between fixed locations / “super computers”*
- *For researchers : Simple applications like email and file transfer*

The Internet: A Success Story

Today:

- Connectivity between humans, **machines**, datacenters, or even **things**
- **Heterogeneous**: e-commerce, VoD, science, etc.
- Wireless and **mobile** endpoints
- ***It hardly changed! But now: mission-critical infrastructure***



So how secure are our networks?



The Internet at first sight:

- Monumental
- Passed the “Test-of-Time”
- Should not and cannot be changed

So how secure are our networks?



The Internet at first sight:

- Monumental
- Passed the “Test-of-Time”
- Should not and cannot be changed



The Internet at first sight:

- Antique
- Brittle
- More and more successful attacks

On Security Assumptions...

- Internet in 80s: based on **trust**
- Danny Hillis, TED talk, Feb. 2013, “There were two Dannys. *I knew both*. Not everyone knows everyone, but there was an atmosphere of trust.”

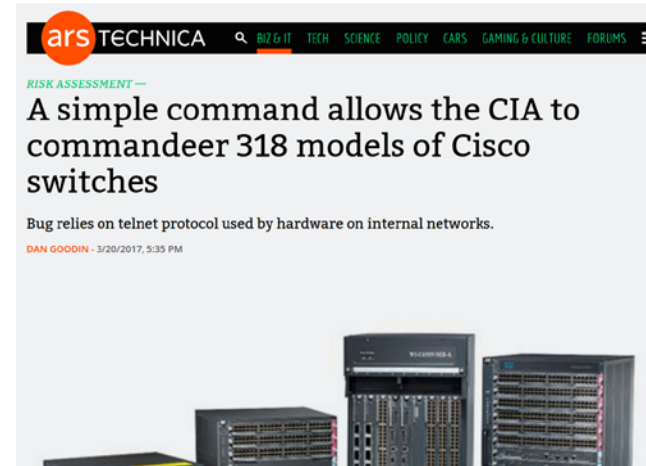


Reality: New Types of Exploits

(TS//SI//NF) Such operations involving **supply-chain interdiction** are some of the most productive operations in TAO, because they pre-position access points into hard target networks around the world.



(TS//SI//NF) Left: Intercepted packages are opened carefully; Right: A “load station” implants a beacon



- **Hardware backdoors** and exploits
- We even need to ask ourselves: how can we *build a secure network if the underlying hardware can be insecure?*

The Big Trend and Challenge: Data

...sensors generate 6GB of data every hour...



AI enabled:

- collision risk prediction
- eight on-board cameras
- six radar emitters
- twelve ultrasonic sensors
- IMU sensor for autonomous driving
- computer power of 22 Macbook Pros

Many Data-Centric Applications



NETFLIX



Many Data-Centric Applications

Datacenters (“hyper-scale”)



NETFLIX



+network

Many Data-Centric Applications




NETFLIX

Datacenters (“hyper-scale”)



+network



Source: Facebook

Many Data-Centric Applications




NETFLIX

Datacenters (“hyper-scale”)



+network

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




Source: Facebook

Key to Success: Resource Sharing


Key to Success: Resource Sharing

That is: **virtualization**


Datacenter Architecture




Datacenter Architecture




Datacenter Architecture




Datacenter Architecture



Datacenter Architecture




Security Requires Isolation on *All Levels*




State-of-the-Art Datacenter Networks

Network Virtualization Today




State-of-the-art: overlays, **tunneling** (e.g., **VxLAN**, VLAN, MPLS, ...)

Virtual Switches: Networking VMs



So far: Network Virtualization Complex and Inflexible


- Configuring tunnels/overlays today is *complex*
- And *inflexible*, e.g., VM migration
- Requires *manual* work



Configuring Today's Networks is Hard:

Case Study Microsoft Datacenter

Example: BGP in
Datacenter (!)




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

Configuring Today's Networks is Hard:

Case Study Microsoft Datacenter

Example: BGP in
Datacenter (!)

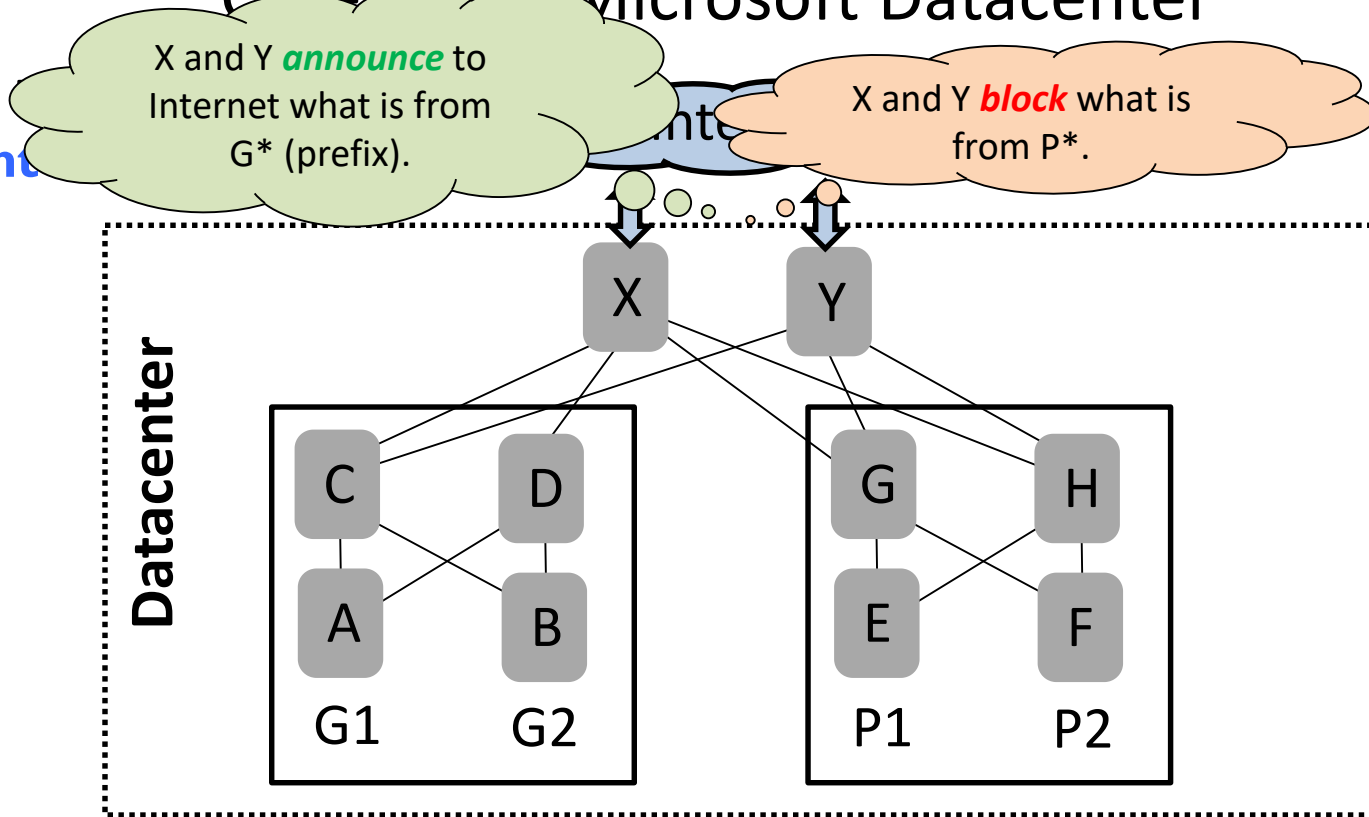


Credits: Beckett et al. (SIGCOMM 2016): Bridging Network-wide Objectives and Device-level Configurations.

Configuring Today's Networks is Hard:

Case Study: Microsoft Datacenter

Example:
Datacenter

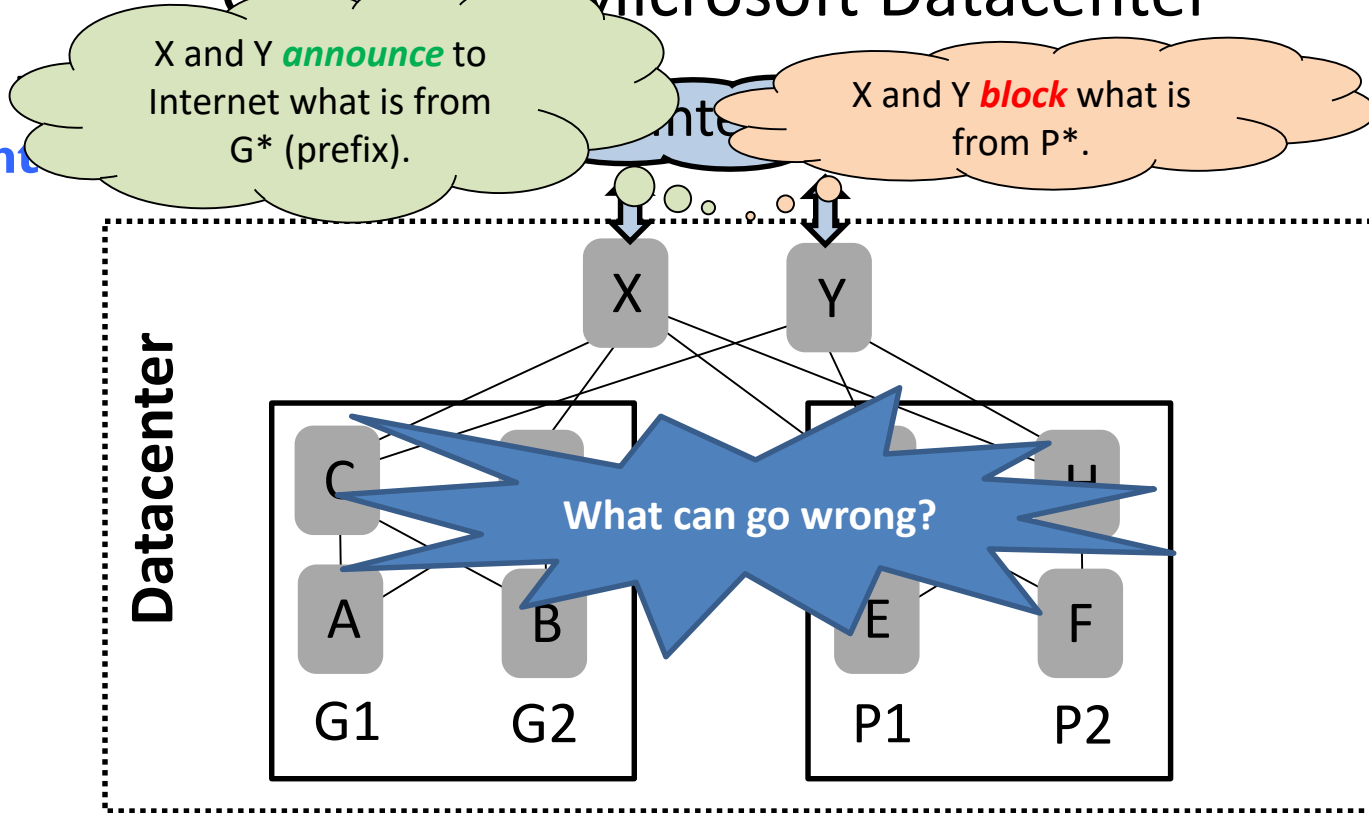


Credits: Beckett et al. (SIGCOMM 2016): Bridging Network-wide Objectives and Device-level Configurations.

Configuring Today's Networks is Hard:

Case Study: Microsoft Datacenter

Example:
Datacenter




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

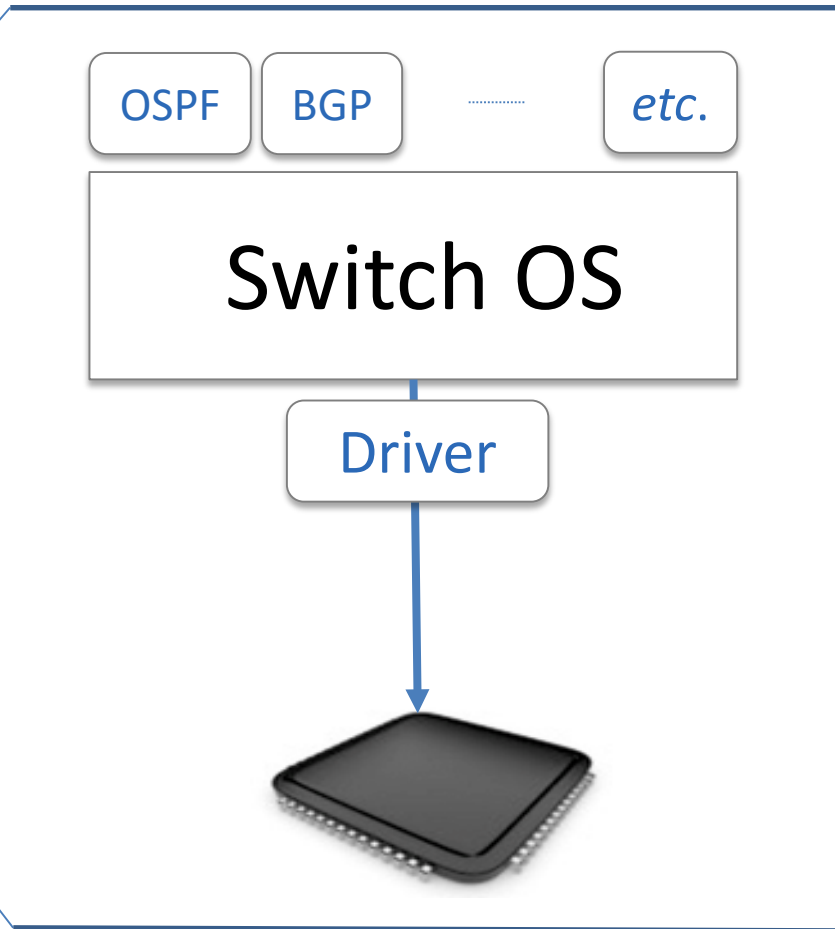
Configuring Today's Networks is Hard:

Case Study: Microsoft Datacenter

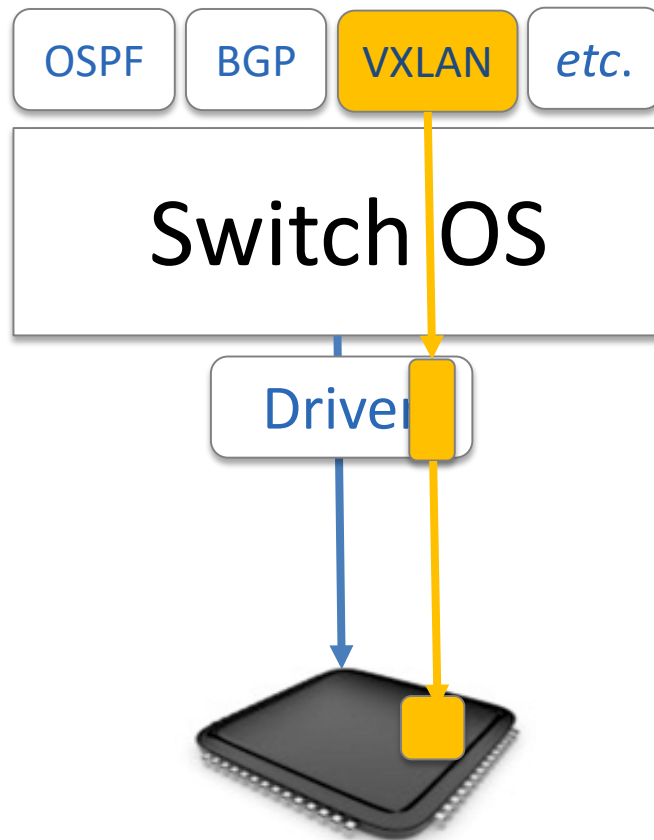

Example:
Datacenter



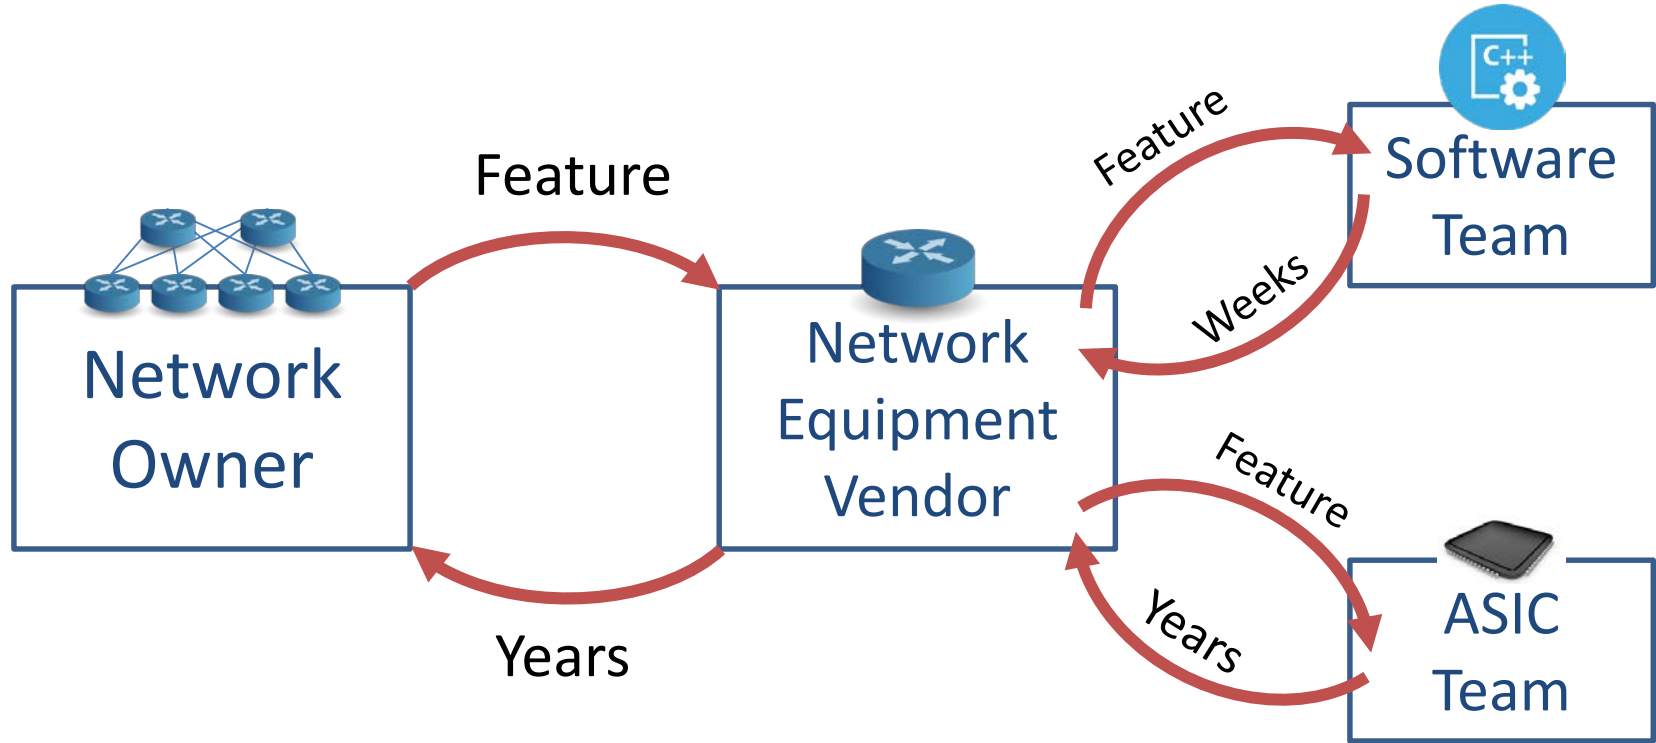
Another problem:
innovation is slow...



Another problem:
innovation is slow...



VxLAN: Took Years...



Slow Innovation...

Operator says:

**I need extended VTP
(VLAN Trunking
Protocol) /
a 3rd spanport
etc. !**


Cisco's answer:

Buy one of these!



Slow Innovation...

Operator says:

A blue speech bubble with a white border and a tail pointing towards the bottom-left.

**I need
something
better than STP
for my data-
center...**

Cisco's answer:


A red speech bubble with a white border and a tail pointing towards the bottom-right.

**We don't
have that!**

Trends in Networking: Opportunities

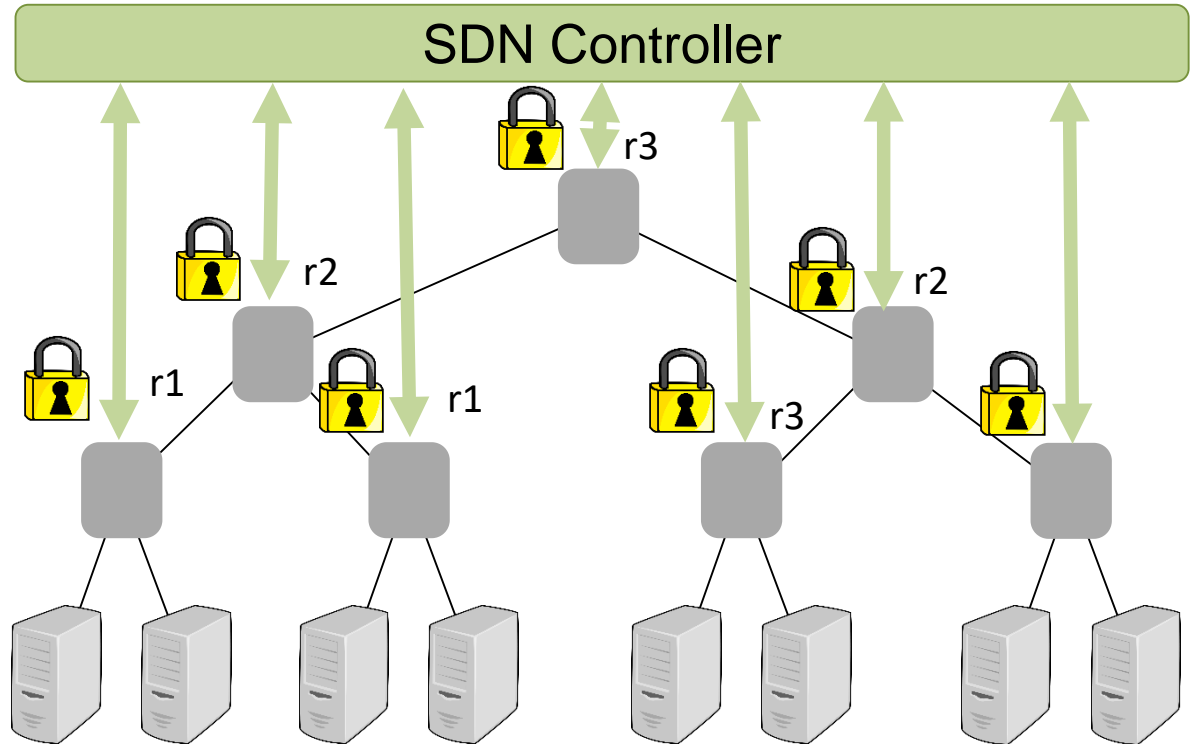
Emerging Software-Defined Networks

- SDN = “The **Linux** of Networking”
 - *Open* interfaces
- **Centralized** and programmatic control
- Fine-grained control, lots of flexibilities
- **Killer application**: network virtualization




Emerging Software-Defined Networks

- SDN = “The **Linux** of Networking”
 - *Open* interfaces
- **Centralized** and programmatic control
- Fine-grained control, lots of flexibilities
- **Killer application**: network virtualization
- Secure communication




Emerging Software-Defined Networks

- SDN = “The **Linux** of Networking”
 - *Open* interfaces
- **Centralized** and programmatic control
- Fine-grained control, lots of flexibilities
- **Killer application**: network virtualization
- Secure communication
- Also **programmable dataplane**
 - Introducing VxLAN easy!




Case Study: Secure Trajectory Sampling

- Traditionally: e.g., **trajectory sampling**
 - Sample packets with $\text{hash}(\text{imm. header}) \in [x, y]$
 - See routes of **some packets**




Case Study: Secure Trajectory Sampling

- Traditionally: e.g., **trajectory sampling**
 - Sample packets with $\text{hash}(\text{imm. header}) \in [x, y]$
 - See routes of **some** packets




Case Study: Secure Trajectory Sampling

- Traditionally: e.g., **trajectory sampling**
 - Sample packets with $\text{hash}(\text{imm. header}) \in [x, y]$
 - See routes of **some** packets
 - Others not!** (Usually later...)




Case Study: Secure Trajectory Sampling

- Traditionally: e.g., **trajectory sampling**
 - Sample packets with $\text{hash}(\text{imm. header}) \in [x,y]$
 - See routes of **some packets**
 - Others not!** (Usually later...)
- What can we do if switches may be **malicious**?
 - Problem: all switches sample the **same space**: known!
 - Can exploit, e.g., **know when unobserved**.



Case Study: Secure Trajectory Sampling


- Solution: **adversarial trajectory sampling with SDN**
- Idea:
 - Use **secure** channels between controller and switches to distribute hash ranges
 - Give **different hash ranges** to different switches, but add some **redundancy**: risk of being caught!
- In general: obtaining live data from the network **becomes easier!**



Trends in Networking: Challenges


Challenge 1: New Threat Vectors

- **Controller** may be attacked or exploited
- E.g., introduces new **covert communication** channels:
 - Communication along existing connections but *bypassing security elements* in the dataplane




Challenge 1: New Threat Vectors

- **Controller** may be attacked or exploited
- E.g., introduces new **covert communication** channels:
 - Communication along existing connections but *bypassing security elements* in the dataplane




Challenge 2: Security of vSwitch



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

The Underlying Problem: Complexity



Number of parsed high-level protocols constantly increases...

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

IPv6HOPOPTS

IPv6ROUTING

IPv6Fragment


IPv6DESTOPT

IPv6ESP


IPv6 AH

RARP


IGMP




Enables Very Low-Cost Attacks




Enables Very Low-Cost Attacks



Enables Very Low-Cost Attacks




Enables Very Low-Cost Attacks



Hopes: MTS, SCION, and Automation?

Hope 1: Better Isolation Mechanisms

- Idea for better *isolation*: put vSwitch in a VM
- But what about *performance*?
- Or container?




Hope 2: Successful Clean Slate Approaches




E.g., the SCION project

Hope 3: Automated What-If Analysis Tools




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_4)	$swap(22)$
	(v_1, v_3)	11	(v_3, v_4)	$swap(12)$
	(v_1, v_3)	21	(v_3, v_4)	$swap(22)$
τ_{v_4}	(v_3, v_4)	12	out_1	pop
τ_{v_5}	(v_2, v_3)	40	(v_5, v_6)	pop
τ_{v_6}	(v_5, v_6)	30	(v_6, v_7)	$swap(31)$
	(v_5, v_6)	30	(v_6, v_7)	$swap(31)$
τ_{v_7}	(v_5, v_6)	64	(v_6, v_7)	$swap(62)$
	(v_5, v_6)	71	(v_6, v_7)	$swap(72)$
τ_{v_8}	(v_6, v_7)	34	(v_7, v_8)	pop
	(v_6, v_7)	62	(v_7, v_8)	$swap(11)$
τ_{v_9}	(v_6, v_7)	72	(v_7, v_8)	$swap(22)$
	(v_3, v_8)	22	out_2	pop
	(v_1, 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_3)	$push(40)$
	(v_2, v_6)	71	(v_2, v_3)	$push(40)$

Network **configurations**



$pX \Rightarrow qXX$
 $pX \Rightarrow qYX$
 $qY \Rightarrow rYY$
 $rY \Rightarrow r$
 $rX \Rightarrow pX$

Automated reasoning (logic)

E.g., P-Rex

But Many Other Fronts Where
Solutions are Needed!

E.g., BitCoin Network




The screenshot shows the coindesk website. The top navigation bar includes links for Blockchain 101, Technology, Markets, Business, Data & Research, Events, and a Search bar. A yellow banner below the navigation bar reads: "Register for our upcoming webinar looking at liquidity, risk & derivatives in crypto assets". The main content area features a large image of fiber optic lights with the headline: "Researchers Uncover Bitcoin 'Attack' That Could Slow or Stop Lightning Payments". To the right of the image is a table titled "coindesk | data" listing the prices and percentage changes for several cryptocurrencies.


Cryptocurrency	Price	% Change
Bitcoin (BTC)	\$7,473.0	-0.42%
Ethereum (ETH)	\$161.2	+0.29%
Litecoin (LTC)	\$50.0	+1.99%
XRP	\$0.275	+4.3326%
Bitcoin Cash (BCH)	\$215.7	+3.56%

*Hijacking Routes in Payment Channel Networks:
A Predictability Tradeoff*

Attracting Transaction Routes



Attracting Transaction Routes



By announcing low fees, can attract significant fraction of transactions on offchain networks!

Conclusion

- Networks are *critical* backbone of our digital society
- Cloud computing introduces new security *challenges*
- Opportunities and threats of network *virtualization* and *programmable* networks
- Security of network also important in other emerging applications, e.g., *bitcoin*

Further Reading

[MTS: Bringing Multi-Tenancy to Virtual Switches](#)

Kashyap Thimmaraju, Saad Hermak, Gabor Retvari, and Stefan Schmid.

USENIX Annual Technical Conference (**ATC**), Renton, Washington, USA, July 2019.

[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.

[Outsmarting Network Security with SDN Teleportation](#)

Kashyap Thimmaraju, Liron Schiff, and Stefan Schmid.

2nd IEEE European Symposium on Security and Privacy (**EuroS&P**), Paris, France, April 2017.

[Preacher: Network Policy Checker for Adversarial Environments](#)

Kashyap Thimmaraju, Liron Schiff, and Stefan Schmid.

38th International Symposium on Reliable Distributed Systems (**SRDS**), Lyon, France, October 2019.

[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 Thorgeresen.

14th International Conference on emerging Networking EXperiments and Technologies (**CoNEXT**), Heraklion, Greece, December 2018.

Hijacking Routes in Payment Channel Networks: A Predictability Tradeoff

And

Saar Tochner and Aviv Zohar
The Hebrew University of Jerusalem
{saar,tochner}@cs.huji.ac.il

Stefan Schmid
Faculty of Computer Science, University of Vienna
stefan_schmid@univie.ac.at

Abstract—Off-chain transaction networks can mitigate the scalability issues of today's trustless electronic cash systems such as Bitcoin. However, these peer-to-peer networks also introduce a new attack surface which is not well-understood today. This paper identifies and analyzes, a novel Denial-of-Service attack which is based on route hijacking, i.e., which exploits the way transactions are routed and executed along the created channels of the network. This attack is conceptually interesting as even a limited attacker that manipulates the topology through the creation of new channels can navigate tradeoffs related to the way

done using bidirectional payment channels that only require direct communications between a handful of nodes, while the blockchain is used only rarely, to establish or terminate channels. As an incentive to participate in others' transactions, the nodes obtain a small fee from every transaction that was routed through their channels. Over the last few years, payment channel networks such as Lightning [24], Ripple [4], and Raiden [23] have been implemented, deployed and have started growing.