Network verification: Lessons learned and outlook

Ratul Mahajan

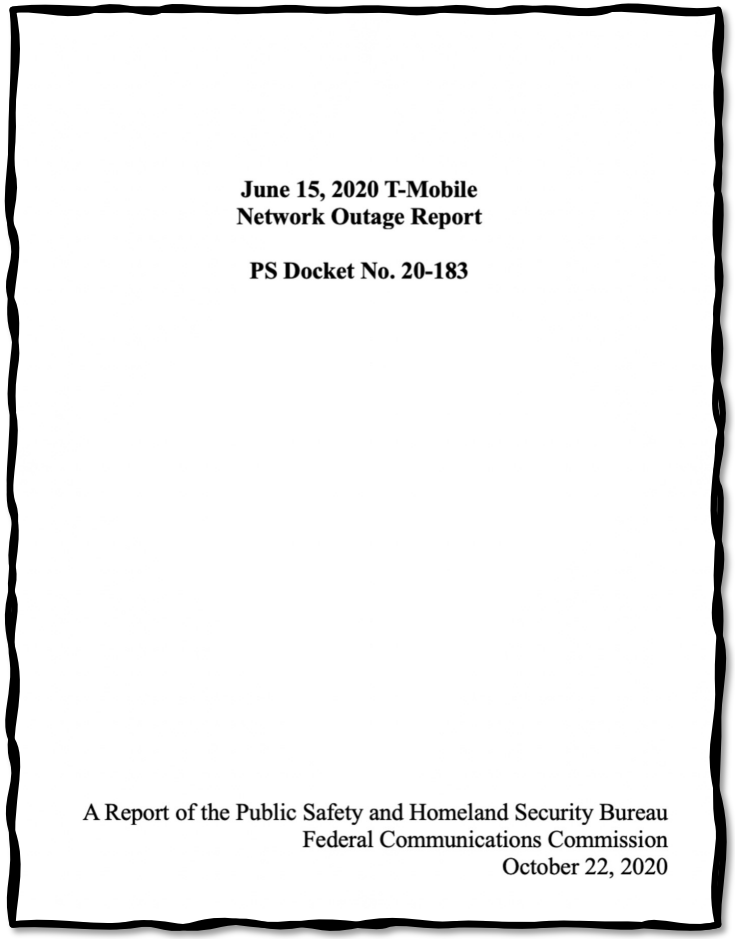

"At least 41% of all calls that attempted to use T-Mobile's network during the outage failed, including at least 23,621 failed calls to 911."

"[An old woman] who has dementia, could not reach [her son] after her car would not start and her roadside-assistance provider could not call her to clarify her location; she was stranded for seven hours"

Anatomy of the outage (illustration)

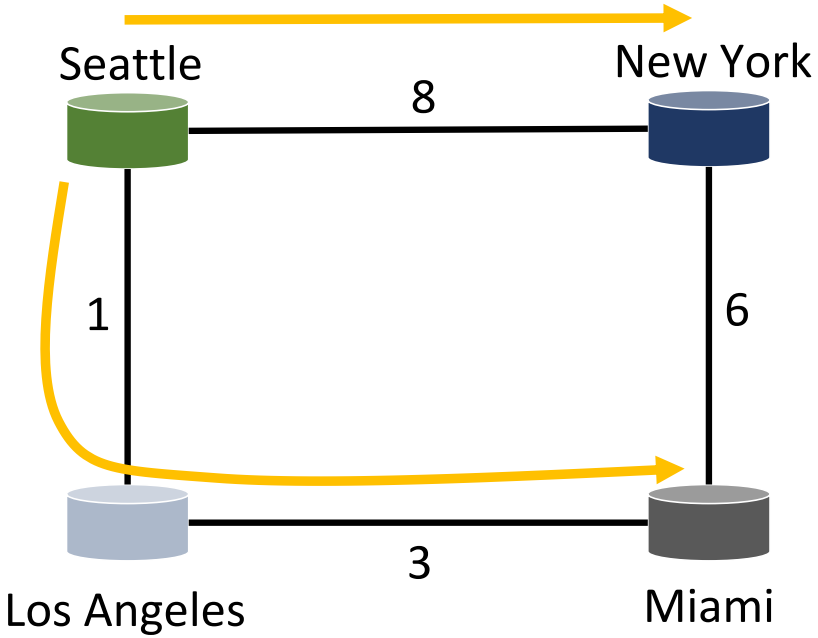

Anatomy of the outage (illustration)

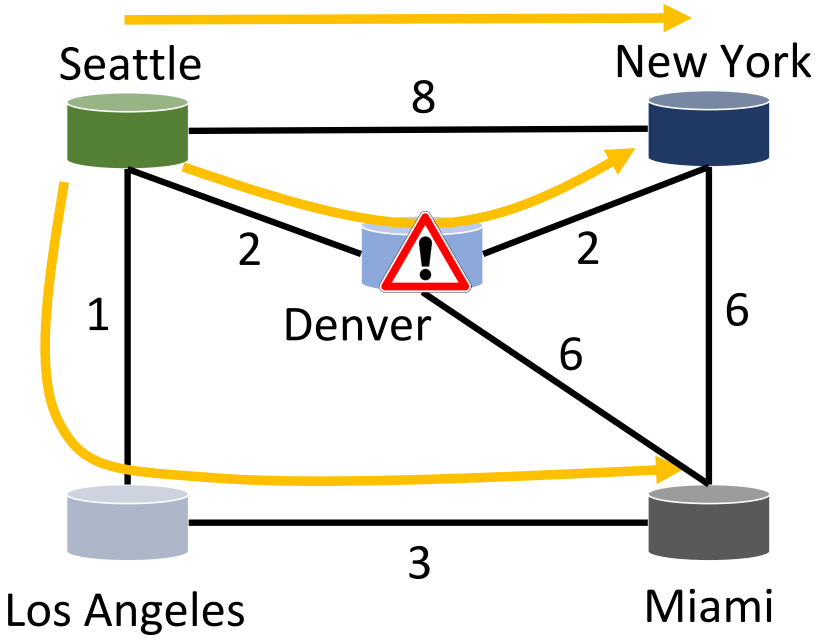

Anatomy of the outage (illustration)

What if T-Mobile could guarantee that no traffic will transit Denver?

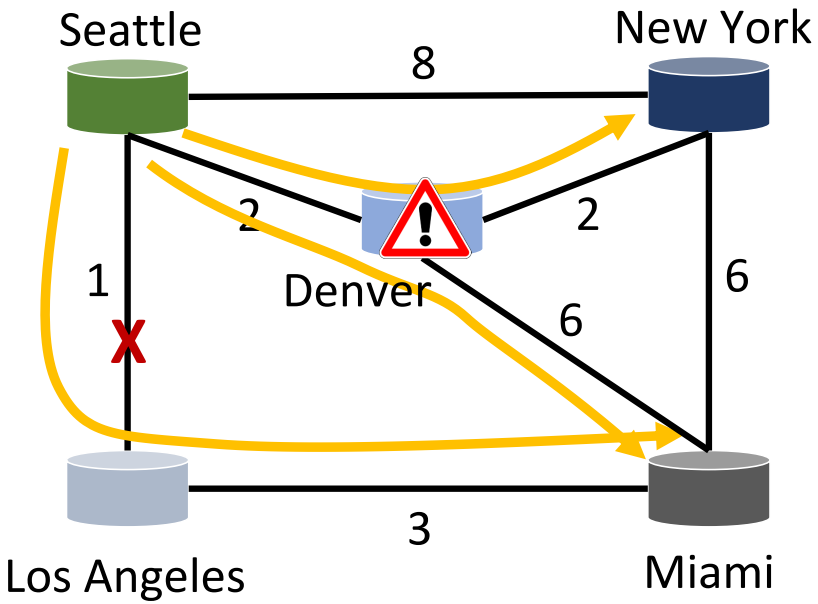

What if T-Mobile could predict the impact of link failure?

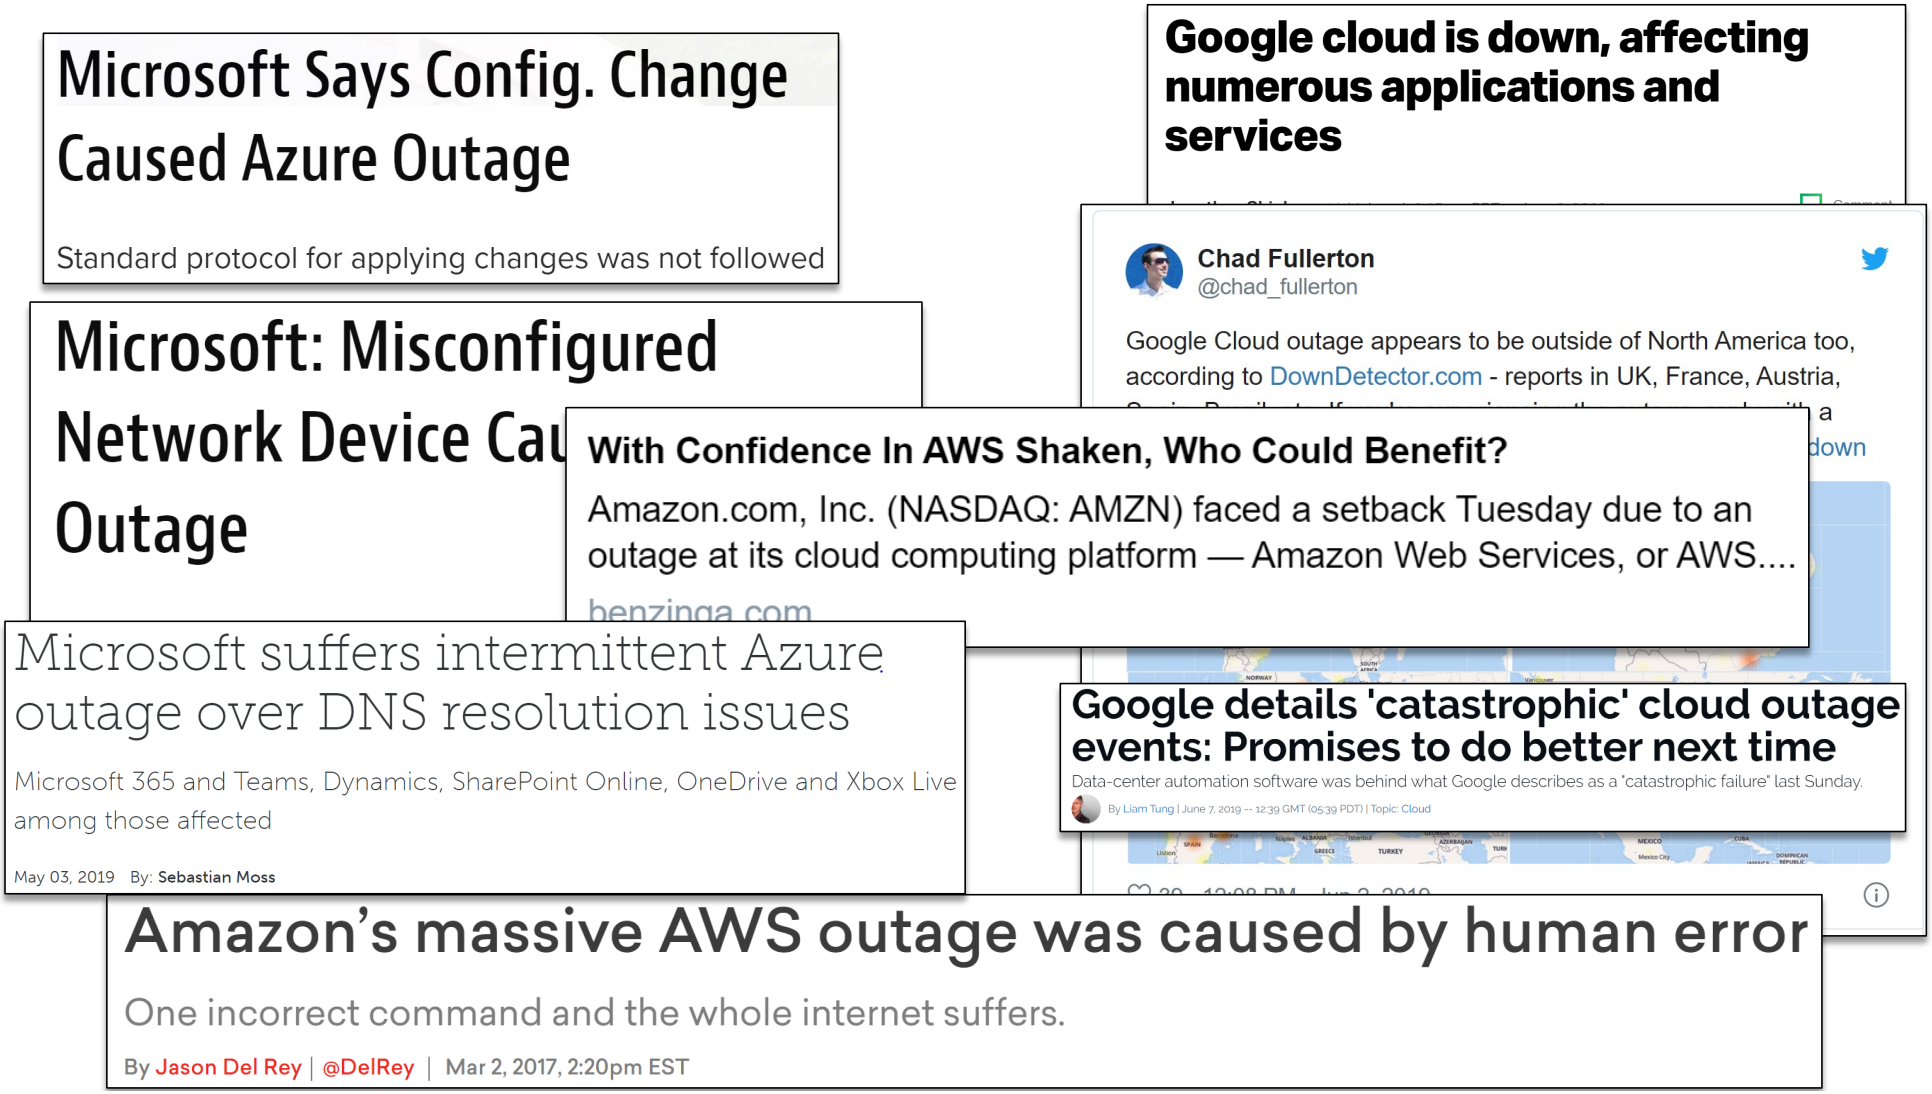
Microsoft Says Config. Change Caused Azure Outage
Standard protocol for applying changes was not followed
Google cloud is down, affecting numerous applications and services
Chad Fullerton
@chad_fullerton
Google Cloud outage appears to be outside of North America too, according to DownDetector.com - reports in UK, France, Austria,
Microsoft: Misconfigured Network Device Cau
Outage
With Confidence In AWS Shaken, Who Could Benefit?
Amazon.com, Inc. (NASDAQ: AMZN) faced a setback Tuesday due to an outage at its cloud computing platform — Amazon Web Services, or AWS....
benzinga.com
Microsoft suffers intermittent Azure outage over DNS resolution issues
Microsoft 365 and Teams, Dynamics, SharePoint Online, OneDrive and Xbox Live among those affected
May 03, 2019 By: Sebastian Moss
Google details 'catastrophic' cloud outage events: Promises to do better next time
Data-center automation software was behind what Google describes as a "catastrophic failure" last Sunday.
Amazon's massive AWS outage was caused by human error
One incorrect command and the whole internet suffers.
By Jason Del Rey | @DelRey | Mar 2, 2017, 2:20pm EST

Network verification to the rescue

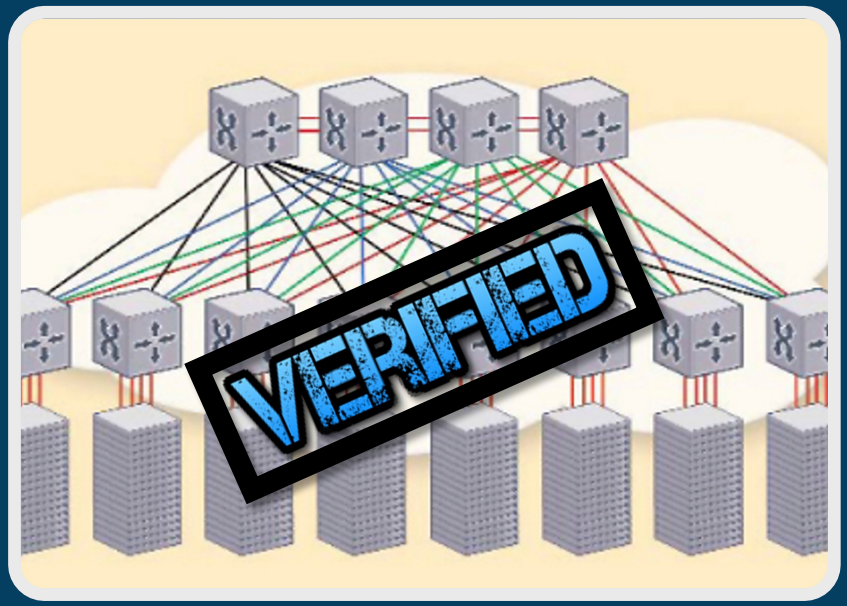

Guarantee network behavior *

* Some behaviors under some assumptions

How network verification slices the problem

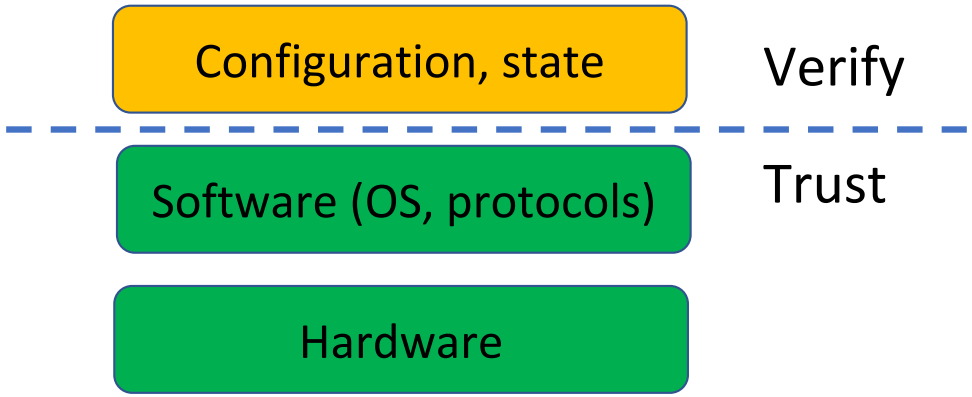

The “haystack” of network behaviors is HUGE

Large scale

$O(10^3)$ devices
$O(10^6)$ routes
$O(10^9)$ packets

Complex interactions

Distributed routing
Protocol redistribution
Rich route filtering

Batfish: A production-grade network verifier

Op
Us
The
Fou

(NSDI 2015)

A General Approach to Network Configuration Analysis

Ari Fogel *Stanley Fung* *Luis Pedrosa* *Meg Walraed-Sullivan*

Ramesh Govindan *Ratul Mahajan* *Todd Millstein*

University of California, Los Angeles University of Southern California Microsoft Research

https://github.com/batfish/batfish

Batfish does proactive network verification

Verify configuration changes ***before*** they affect the network

Batfish's original 4-stage pipeline

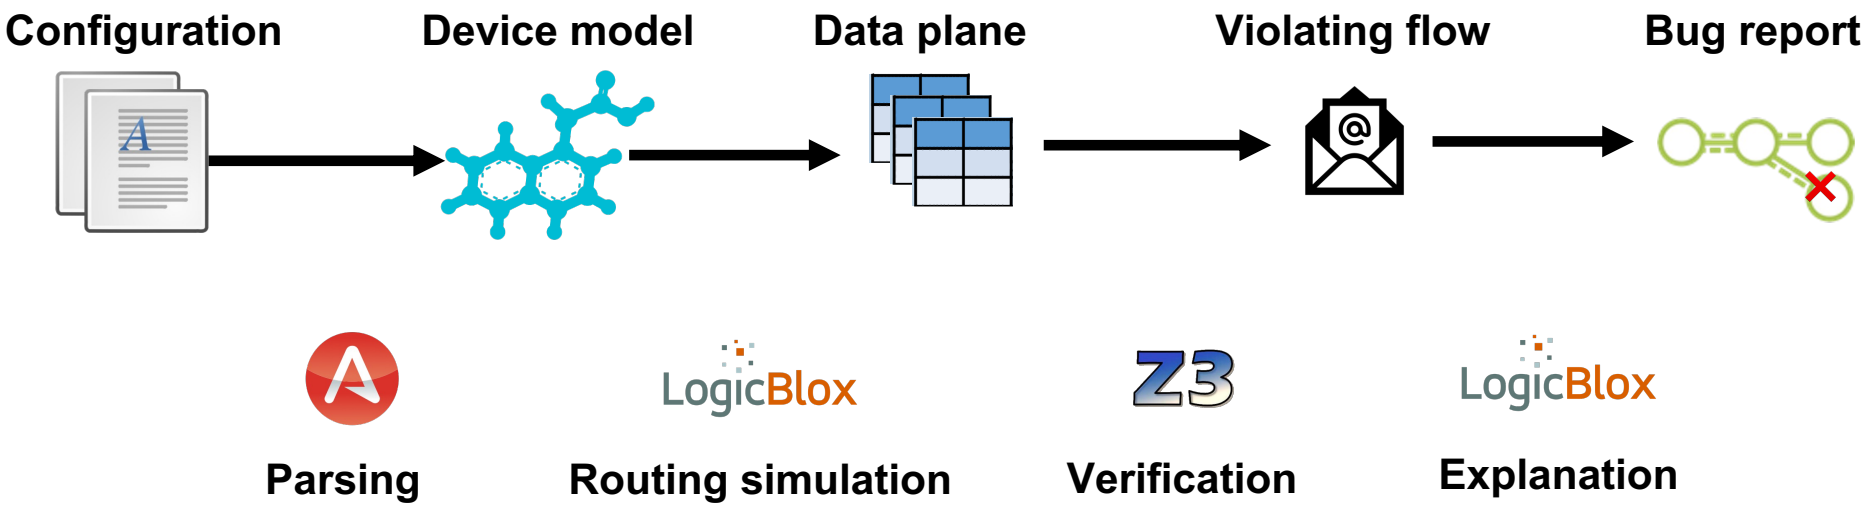

Batfish's original 4-stage pipeline

2023

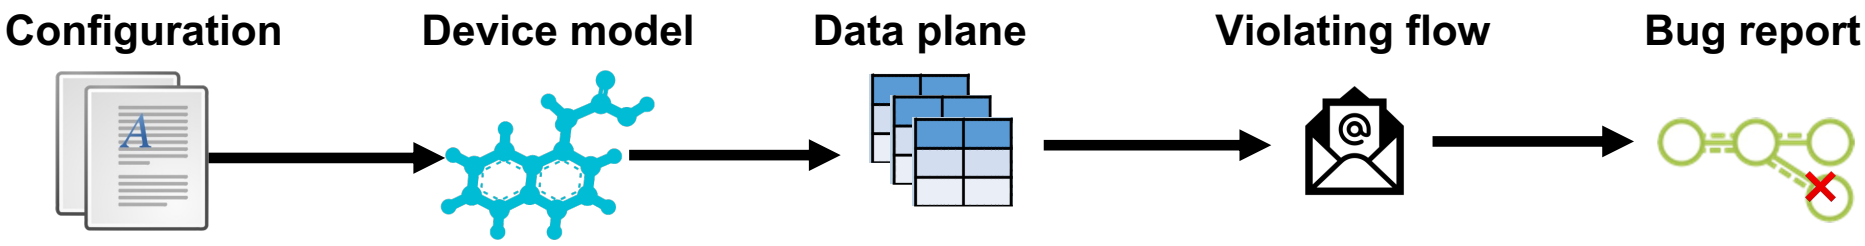

New! Fidelity

New! Parsing

New! Routing simulation

New! Verification

New! Explanation

1500x faster, **400x** larger networks

(SIGCOMM 2023)

Experience Track

Lessons from the evolution of the Batfish configuration analysis tool

Matt Brown
Intentionet

Ari Fogel
Intentionet

Daniel Halperin
Intentionet

Victor Heorhiadi
Intentionet

Ratul Mahajan
Intentionet
University of Washington

Todd Millstein
Intentionet
UCLA

Lesson 1: Datalog was great for prototyping, but not for production use

Three key challenges:

1. **Expressiveness**
2. **Performance**
3. **Deterministic convergence**

Solution: replace Datalog with imperative code

Parsing

Routing simulation

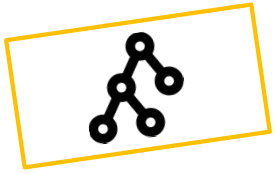

Verification

Explanation

Lesson 2: Model fidelity is hard, but not why you think

Concern*:* “Every software version will have different semantics!”

Reality*:* The real challenge is **undocumented semantics**

Solution: New stage to benchmark Batfish against an emulator

Fidelity

Parsing

Routing simulation

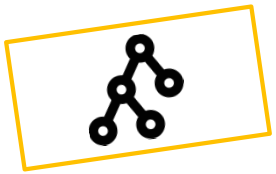

Verification

Explanation

Lesson 3: Usability is hard for reasons you think, and then some

Ambiguity: “Hosts **A** can reach hosts **B**”

- ALL **applications** can reach SOME **DNS server** (e.g., in the same AZ)
- SOME **SNMP collector** can reach **infrastructure elements**
- ALL **service frontends** can reach ALL **backend VIPs**

Solution: custom assertions for each use case.

What’s next for network verification?

Make it an *effective* and *universal* practice

Key hurdles

- Lack of network automation
- Lack of expertise
- Lack of precise specifications

Network verification is only as good as its invariants

Network change

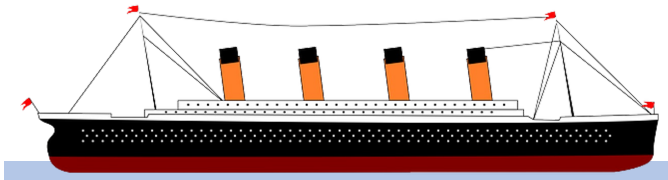

invariants

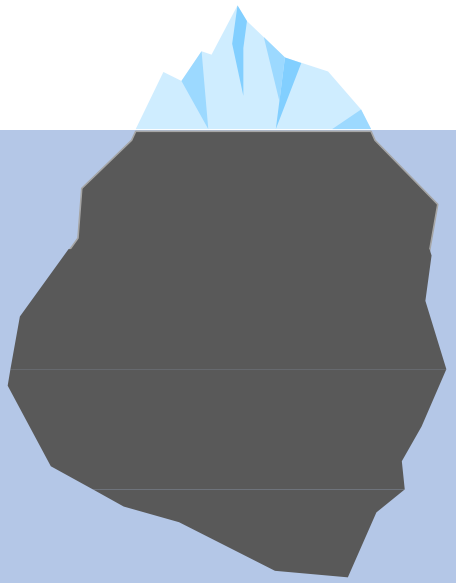

untested
network
behaviors

Network verification is only as good as its invariants

This article was published on: 10/4/21

Home / Featured / Facebook outage triggered by BGP configuration issue as services fail for 6 billion

Featured Read This

Facebook outage triggered by BGP configuration issue as services fail for 6 billion

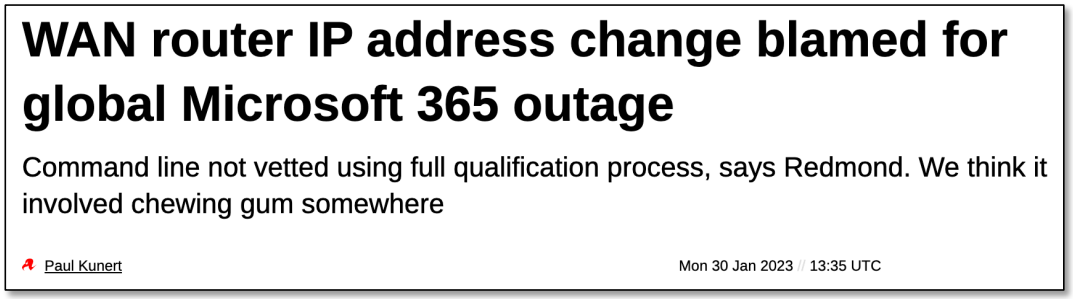

WAN router IP address change blamed for global Microsoft 365 outage

Command line not vetted using full qualification process, says Redmond. We think it involved chewing gum somewhere

Paul Kunert Mon 30 Jan 2023 // 13:35 UTC

Inspiration from code coverage

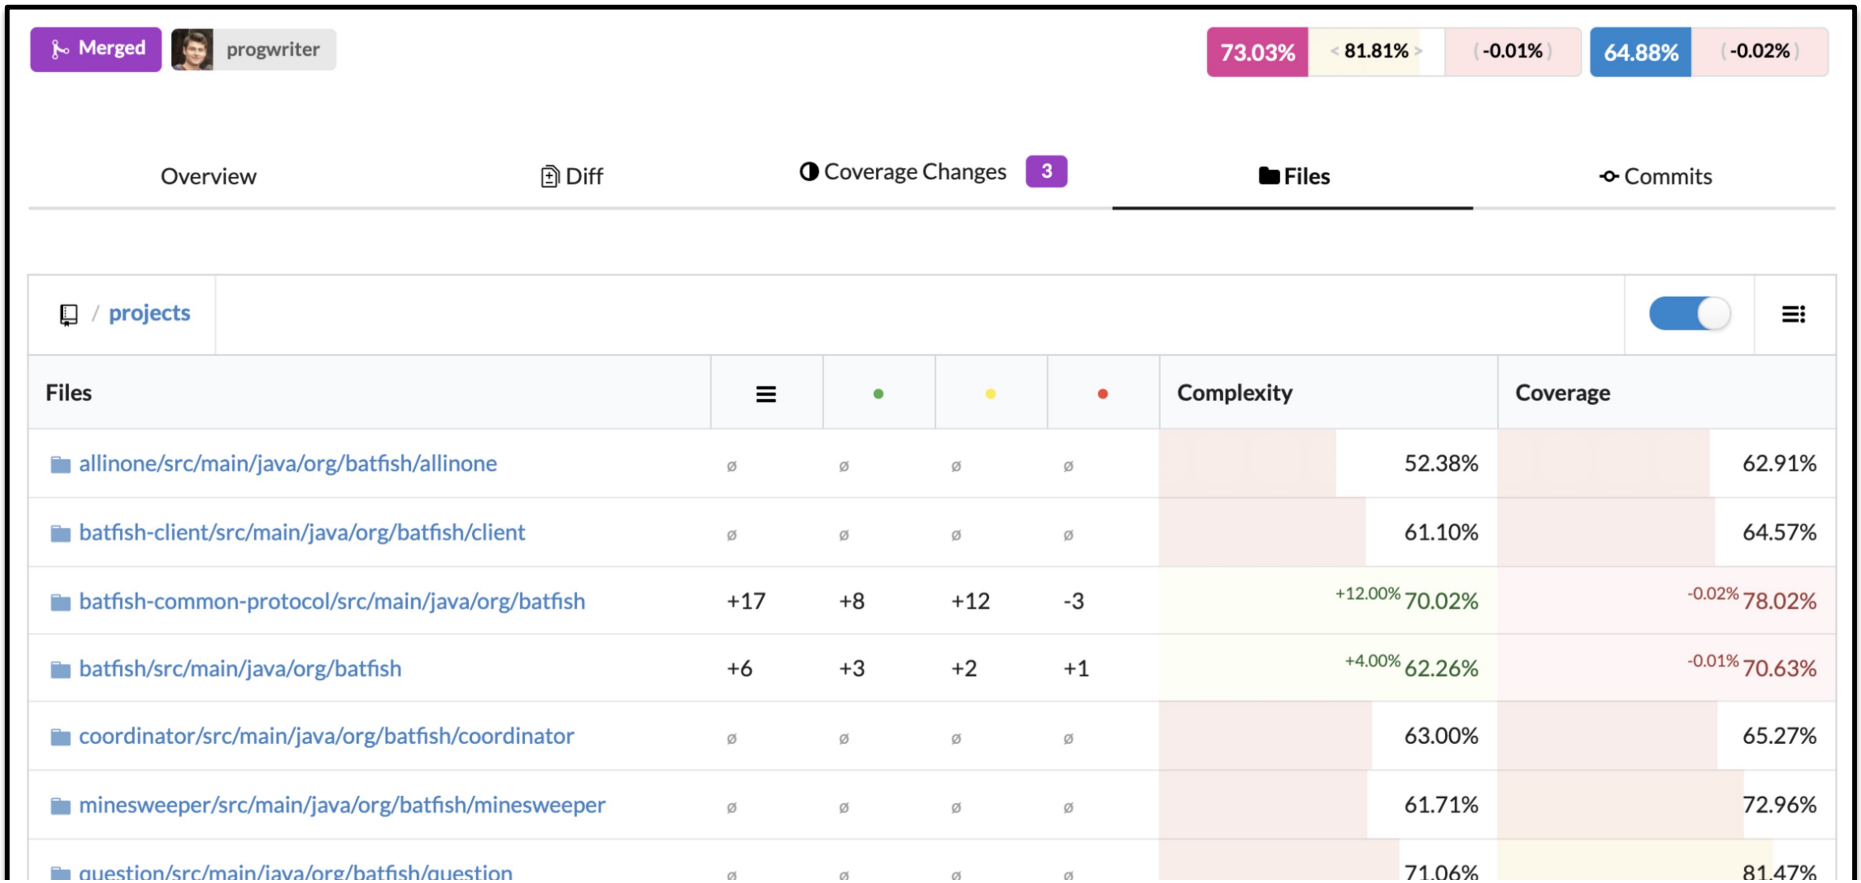

NetCov: Coverage for network configurations

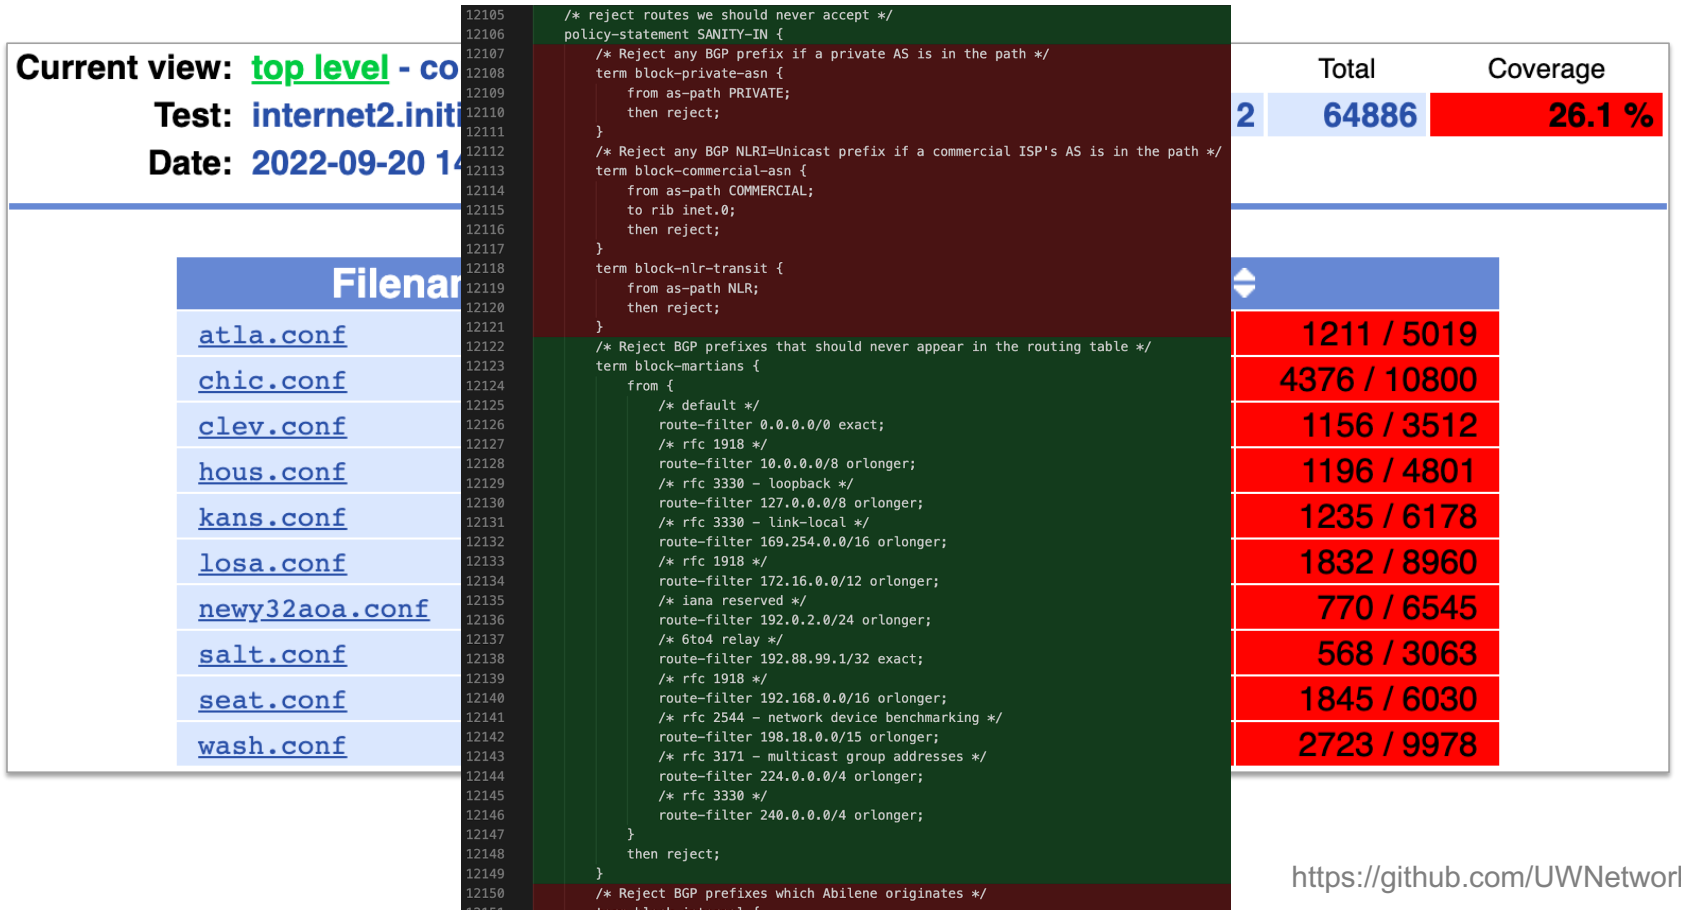

https://github.com/UWNetworksLab/netcov

NetCov maps tested data plane state to covered config lines

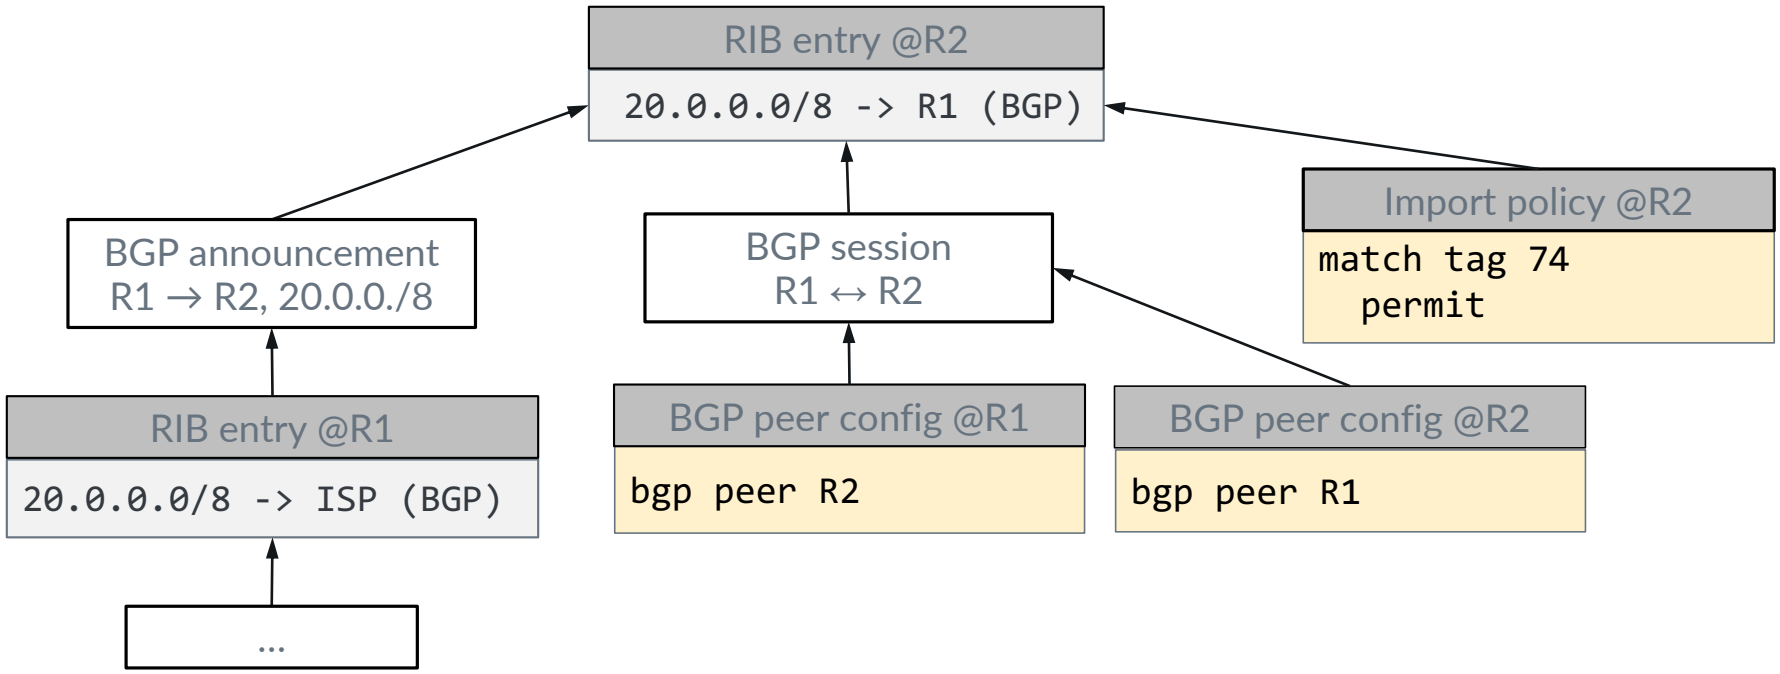

NetCov maps tested data plane state to covered config lines

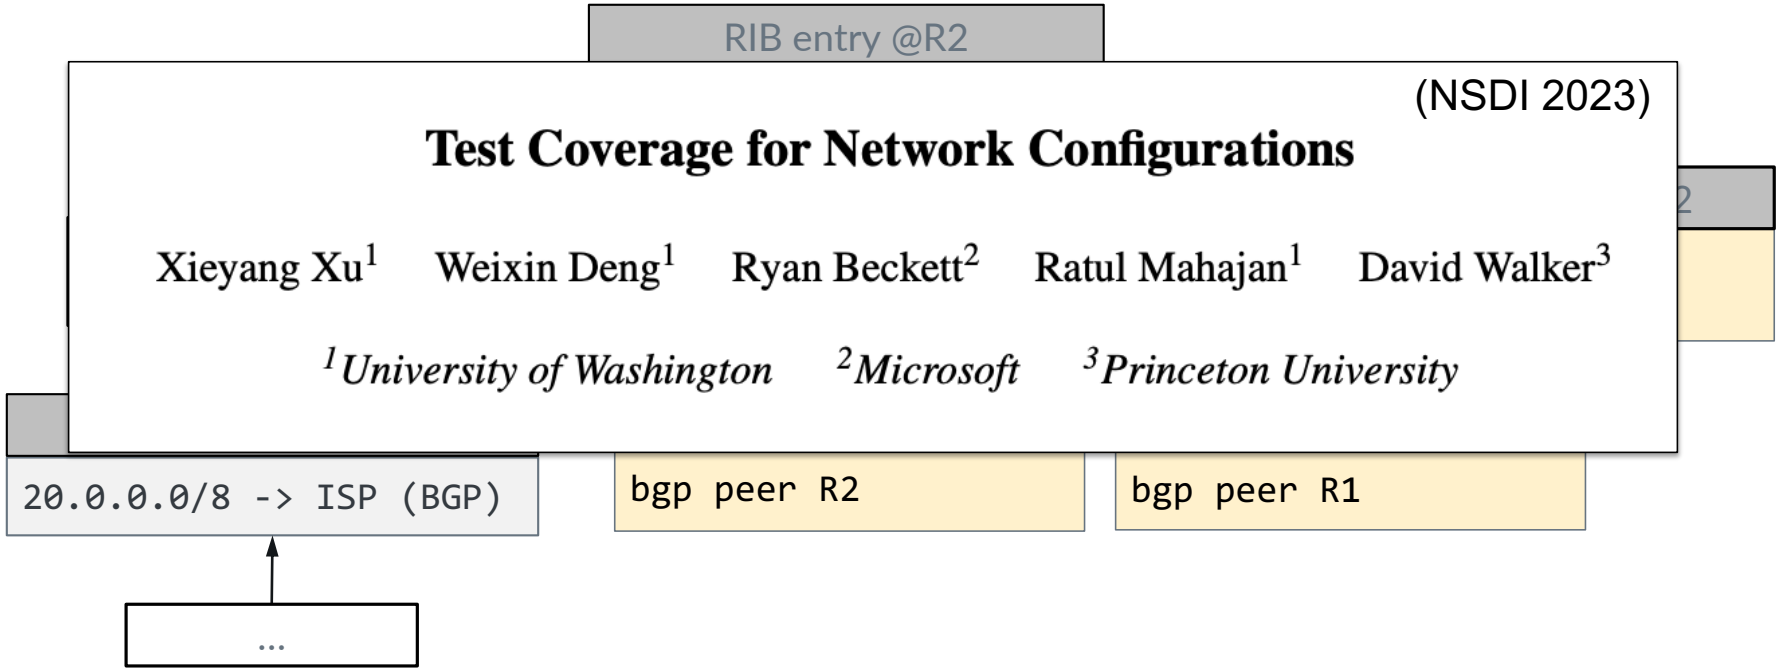

Toward “specification-less” verification

Insight: Network’s spec may not be known but a change’s intent is

- The change should have no impact on reachability
- The change should make the new subnet reachable from here
- The change should make traffic on *path1* take *path2*

Solution: Differential network verification

- A relational language for network changes
- A evaluation procedure based on finite state transducers

w/ Xieyang Xu, Zak Kincaid, Arvind Krishnamurthy, David Walker, and Yifei Yuan

Summary

Network verification is key to high network availability

First generation of tools have taught us a lot about what (does not) work

Must focus now on making network verification an effective and universal practice