




Research Faculty Summit 2018

Systems | Fueling future disruptions



Machine Learning in Azure Networking


(a few sample problems)

David A. Maltz

Distinguished Engineer

Azure Physical Networking Team

dmaltz@microsoft.com



Large Scale Creates Large Problems

- 100,000s of links in each datacenter
- 10,000s of links in each MAN
- 1,000s of links in the WAN

→ **High availability is job number #1 for the network**

At scale, the law of large numbers is not your friend

- Instead of “Occam’s Razor” – the simplest explanation is most likely
- “Murphy’s Law” applies – whatever can go wrong, will

Find the cause of perceived network problems is hard

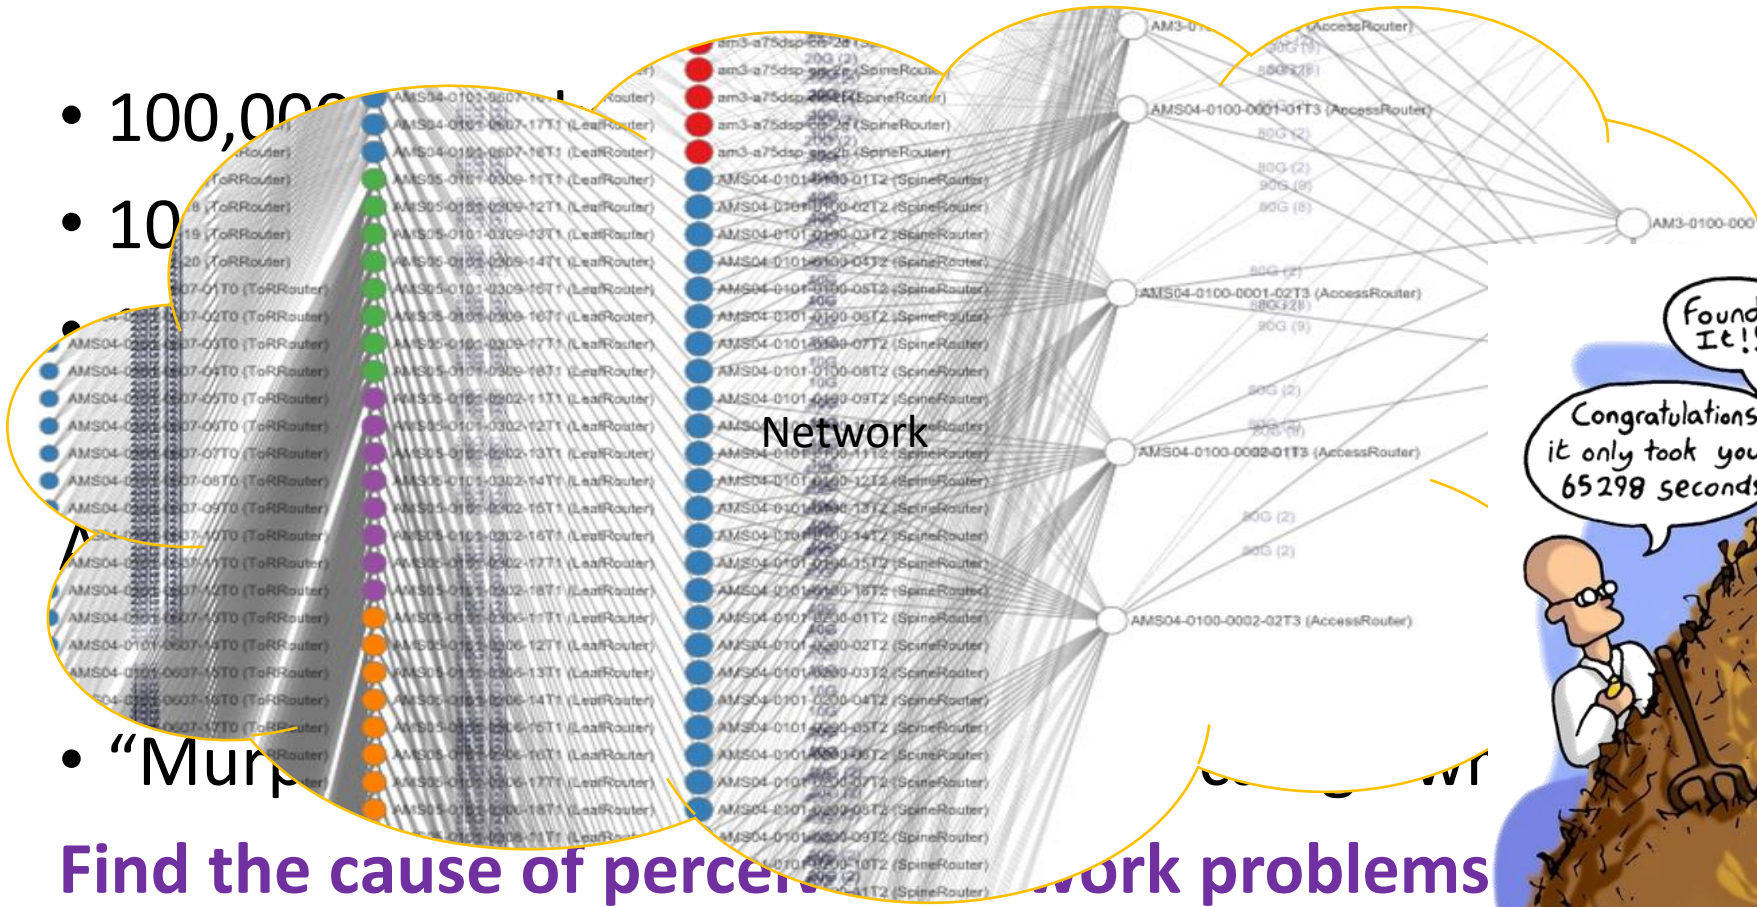
Large Scale Creates Large Problems

- 100,000

- 10

- "Murphy's Law"

Find the cause of percent network problems



Machine Learning in Azure Network

A Few Sample Problems

ML is not the best approach for all problems, regardless of how big and diverse the data set is.

Use ML when...

"I don't understand the underlying physics that causes this; however, I see outcomes, I know good vs bad, and I want to try and understand the outcome"

Use Rules Based when..

"I have a good physical model and understanding of causes"

examine these three further

Problem	Machine Learning	Rules Based System
Optimal Layer-1 topology	✓	
Layer-1 Network Routing Optimization	✓	
Optical Performance Optimization		✓
Demand Estimation	✓	
Network availability	✓	
Fiber segment availability		✓
Probable cause for path degradation		✓
Network Alert Correlation	✓	

Topology: Which cables would you choose?

Problem statement:

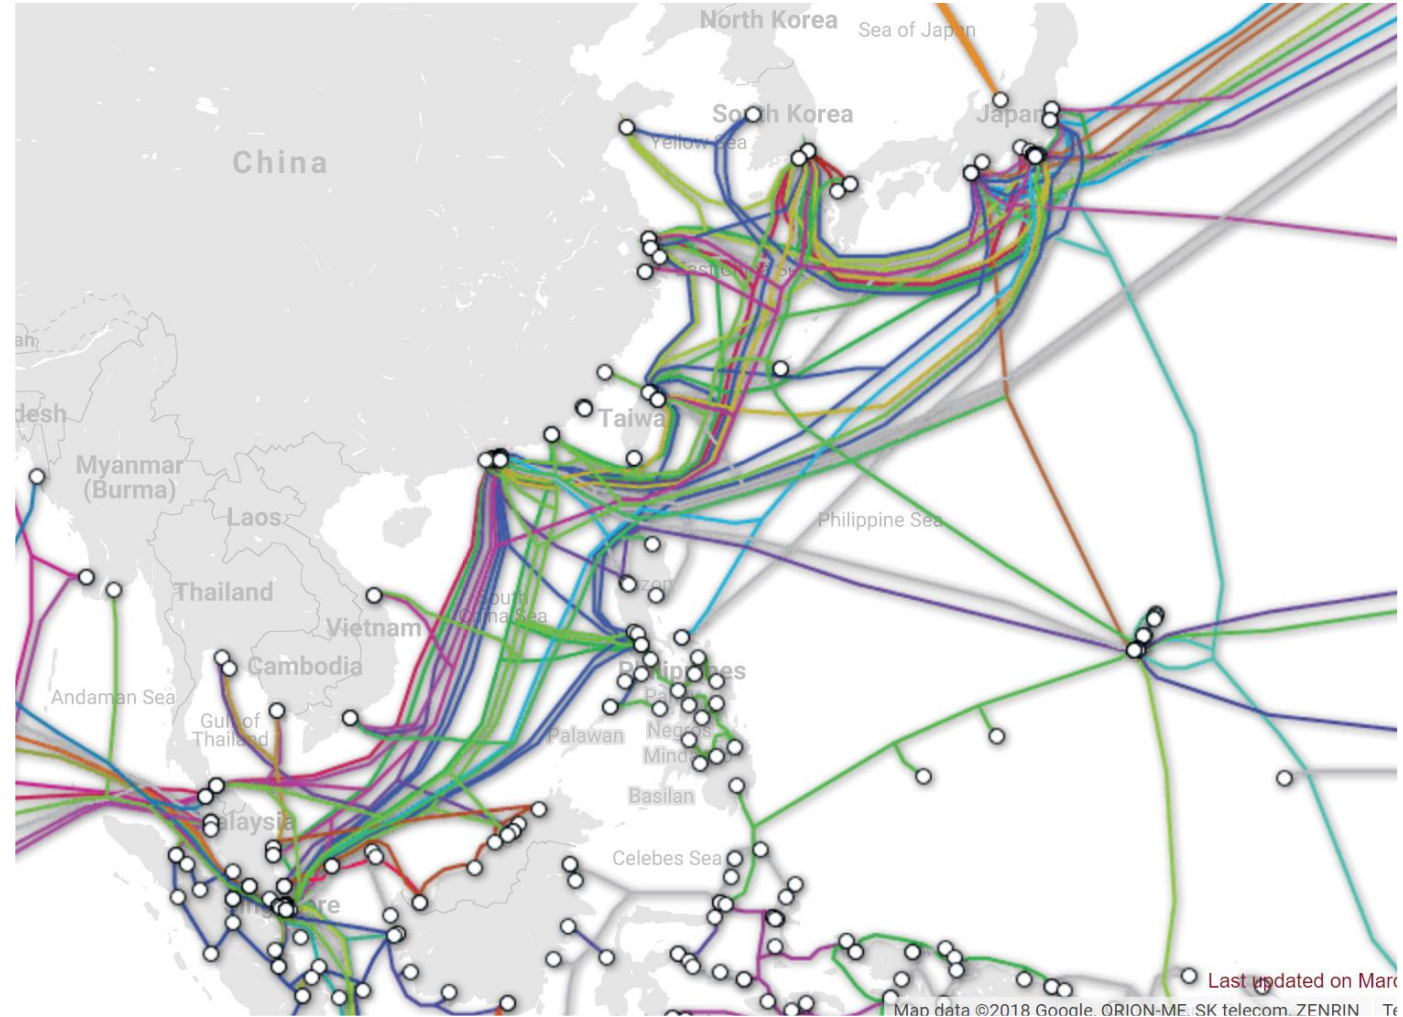
What's Time Between Failure and Time to Repair distribution for a new fiber path?

Approach:

Estimate using historical experience with "similar" paths

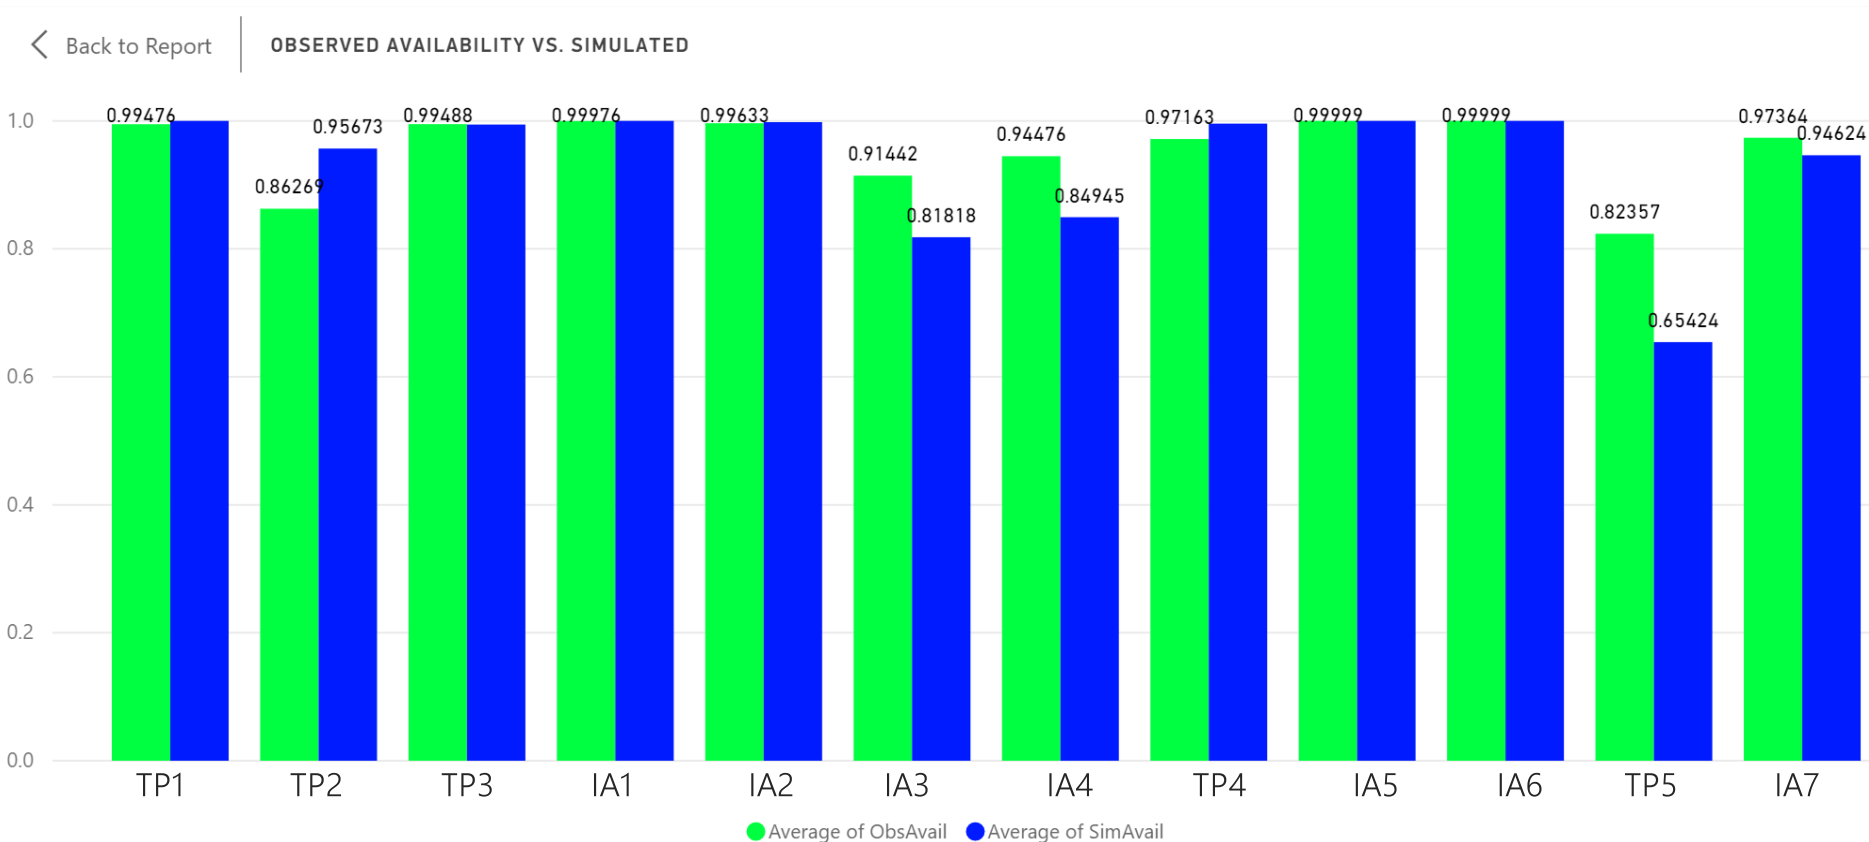
Many, many dimensions...

Topology decisions too important to be influenced by human emotion and relationships



Region and Path Availability

Using only BGP Link States, we measure and predict Layer-1 path Availability



Data Collection

- Streaming BGP telemetry + Azure Blob

Data Clensing

- Classify
- Mine the up/down diff
- Filter maintenance, flapping, etc...
- Subset to geo-locations
- Map to SRLG
- Extract flows and latency
- Bin data and serve into Monte Carlo simulation w/o assuming distribution
- Cluster and prepare candidate add links that are the worst offenders
- Goal seek to discover golden topology.

Machine Learning

- Take that topology, estimate failure trends, and feed back it into the simulation
- Find ranges where topology falls matches SLAs

Machine Learning in Azure Network

Layer-1 Sample Problems

ML is not the best approach for all problems, regardless of how big and diverse the data set is.

Use ML when...

"I don't understand the underlying physics that causes this; however, I see outcomes, I know good vs bad, and I want to try and understand the outcome"

Use Rules Based when..

"I have a good physical model and understanding of causes"

Problem	Machine Learning	Rules Based System
Optimal Layer-1 topology	✓	
Layer-1 Network Routing Optimization	✓	
Optical Performance Optimization		✓
Demand Estimation	✓	
Network availability	✓	
Fiber segment availability		✓
Probable cause for path degradation		✓
Network Alert Correlation	✓	

Wavelength & Performance Optimization

Open Line System (OLS) enables us to tune parameters for each link to maximize bits per second the link can carry

→ Is this a good Machine Learning problem?

Availability is our highest priority therefore we take a cautious approach to optimization

- Monitor proactive metrics (e.g. power levels, BERs)
- Do nothing until you see something bad
- Trigger deterministic behavior based on metric change
- Use physical layer model as reference to guide changes
 - [Gaussian Noise model](#) code implemented as gnpy on [github](#)
 - See contributions to TIP from Mark Filer and MSFT

Physical layer is well understood. No clear role for machine learning in link optimization

Machine Learning in Azure Network

Layer-1 Sample Problems

ML is not the best approach for all problems, regardless of how big and diverse the data set is.

Use ML when...

"I don't understand the underlying physics that causes this; however, I see outcomes, I know good vs bad, and I want to try and understand the outcome"

Use Rules Based when..

"I have a good physical model and understanding of causes"

Problem	Machine Learning	Rules Based System
Optimal Layer-1 topology	✓	
Layer-1 Network Routing Optimization	✓	
Optical Performance Optimization		✓
Demand Estimation	✓	
Network availability	✓	
Fiber segment availability		✓
Probable cause for path degradation		✓
Network Alert Correlation	✓	

Network Availability

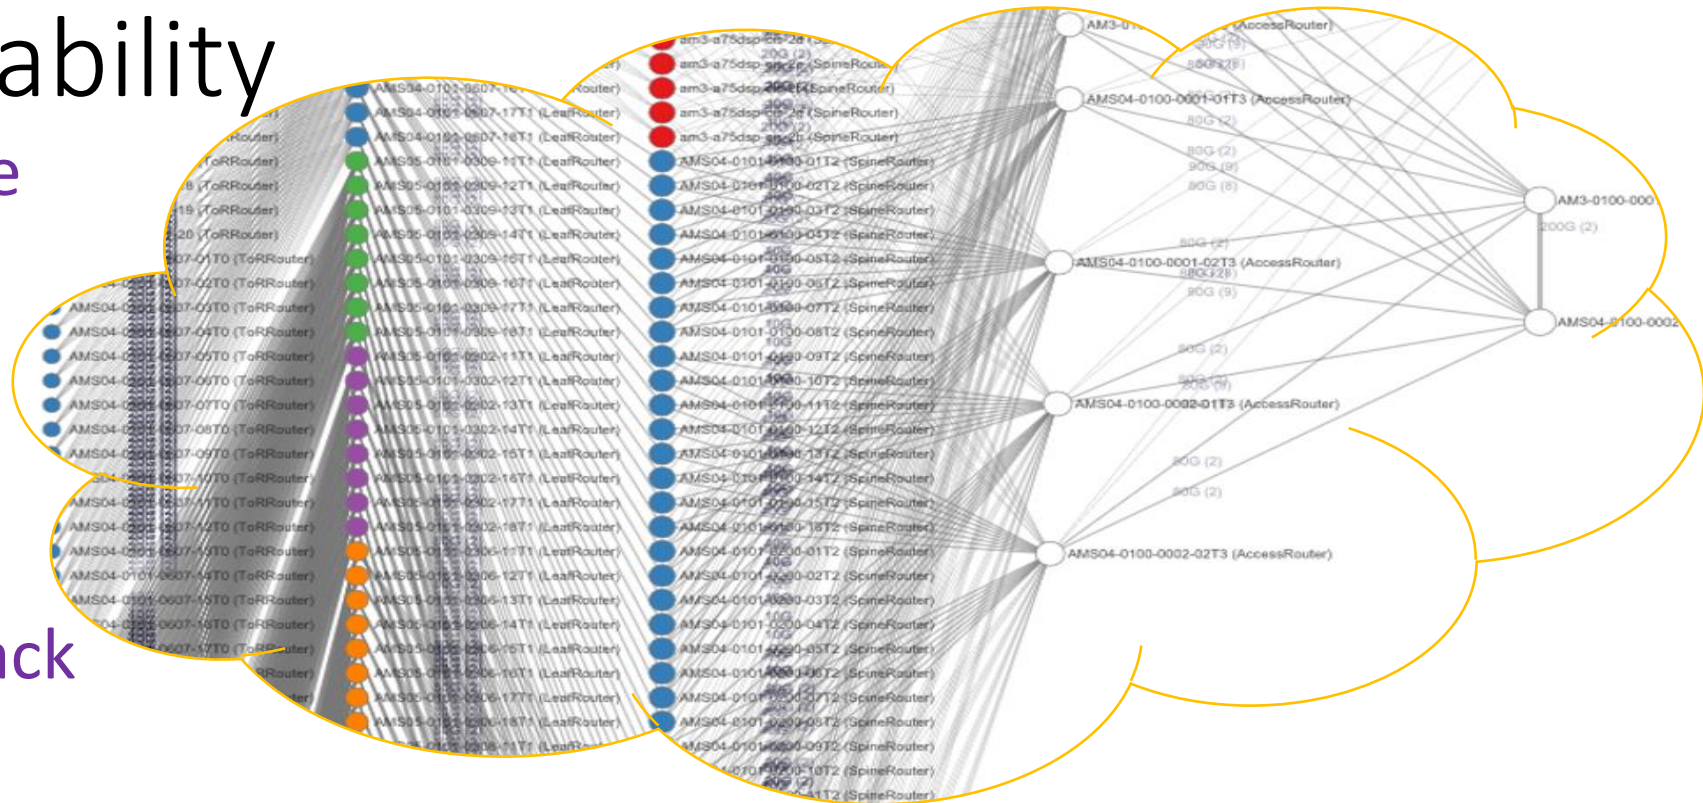
“Gray” switch failures are the worst

- Switch stays in service
- Drops some fraction of packets

Find the needle in haystack

- Pingmesh
- Targeted probe packets
- Error messages sent from switches
- Service-level health metrics

Combine them all to localize problem to most likely switch




Machine Learning in the Azure Network Team

Practitioners Guide

Be prepared to *Cross the Knowledge Gap*

- ML is much more than data science.
- Often, the biggest challenges involve setting up the data pipeline, scrubbing data, monitoring data integrity, classification and scaling.
 - This requires subject matter expertise in the area you are trying to solve. You won't have it. You need to sit with SMEs, extract context, contribute to data pipeline and integrity.
- This is very hard. You need to be more passionate about the problem than your specific subject matter expertise.
- The more overlap you can achieve, the better the outcome.



Thank you!

