ciena
Experience. Outcomes.

Using P4 INT to feed the ML Monster

ONF Connect

David Bainbridge
Karthick Ramanarayanan
Yan Zhuang

December 4 – 6, 2018

https://www.demilked.com/artists-redraw-children-drawings-inspiration-monster-project/

“It’s not magic, it’s math.”

– Many different people

The Plan

Collect P4 in-band telemetry (INT) and understand if the data can be used to predict network performance utilizing machine learning technologies

Compare results of predictions based on P4 INT to predictions based on traditional, polled, performance metrics

Utilize predictions for closed loop policy and correction

Network Under Test

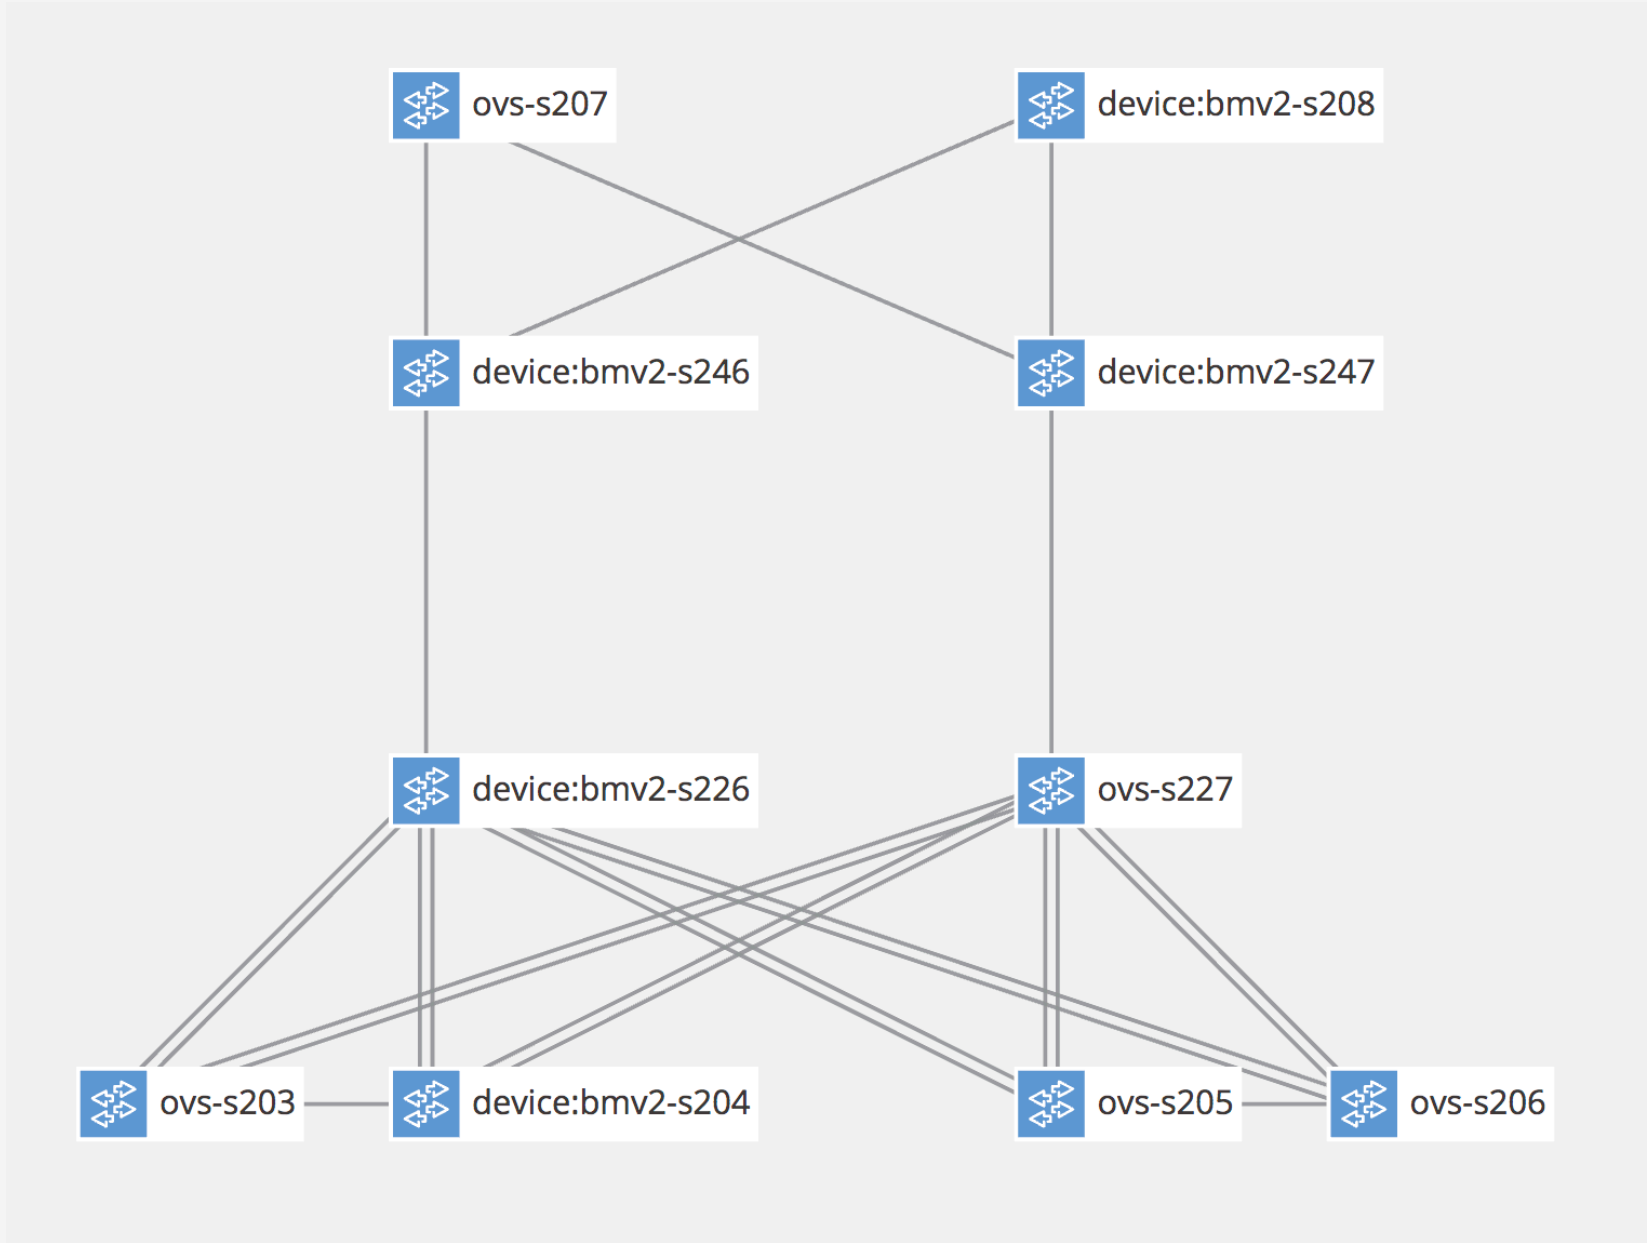

Network Under Test with Data Injection/Collection Points

- Traffic injection/collection point

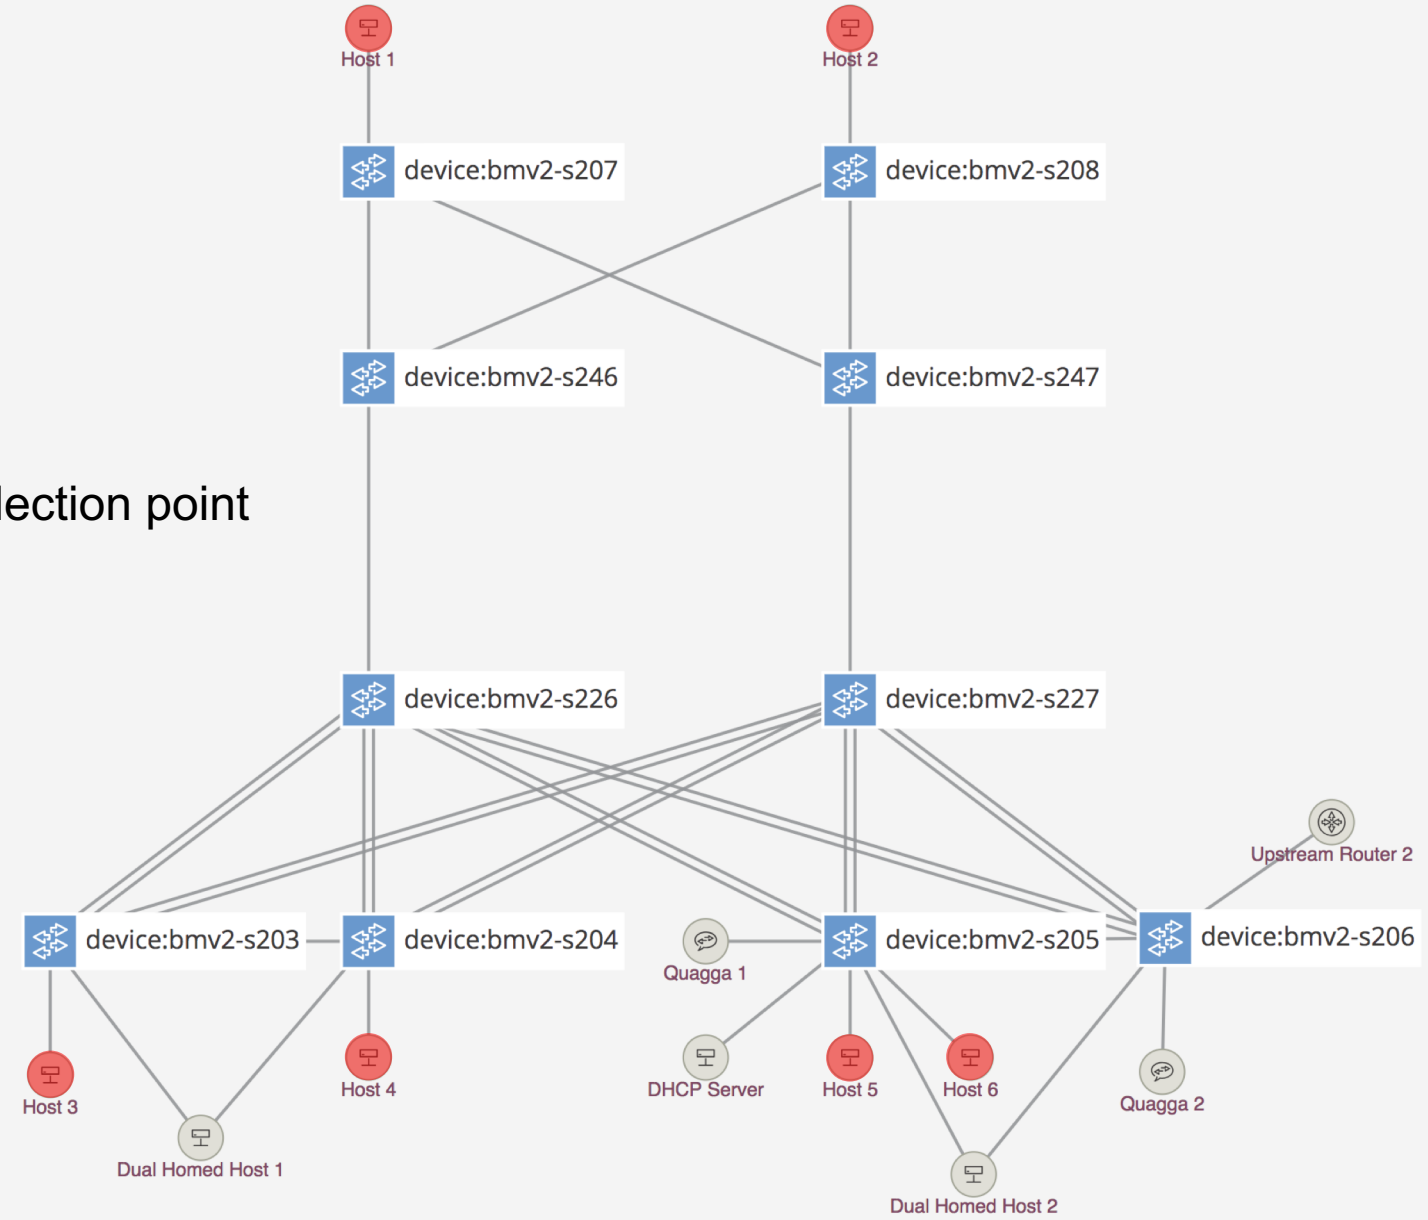

P4 Hop Latency

Originally only able to retrieve ingress/egress port information and hop latency

Collected P4 INT data and transformed it into uniform table

- **Collector ID –** MAC address of host from which the data was collected
- **Sequence Number –** Each INT sample represents a path through the network, this “sequence number” allowed correlation by collection sample
- **Timestamp –** Time at which the sample was collected
- **Switch ID –** Switch for which this row represents
- **Ingress Port –** Switch port ort on which the packet was received
- **Egress Port –** Switch port on which the packet was transmitted
- **Hop Latency –** Time the packet spent from ingress to egress

Using Hop Latency to Predict Hop Latency

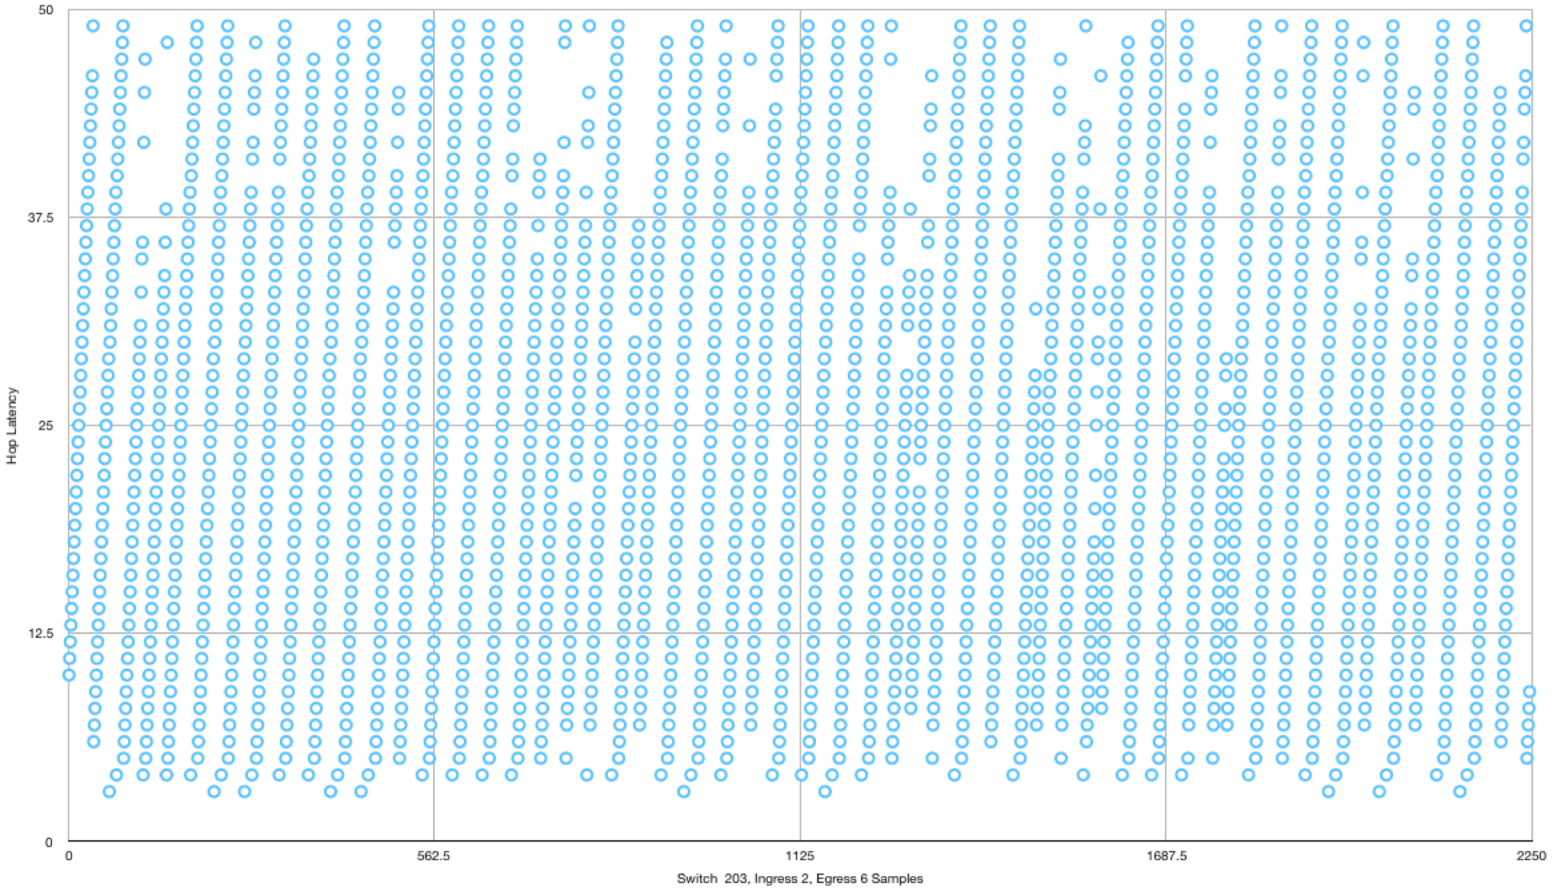

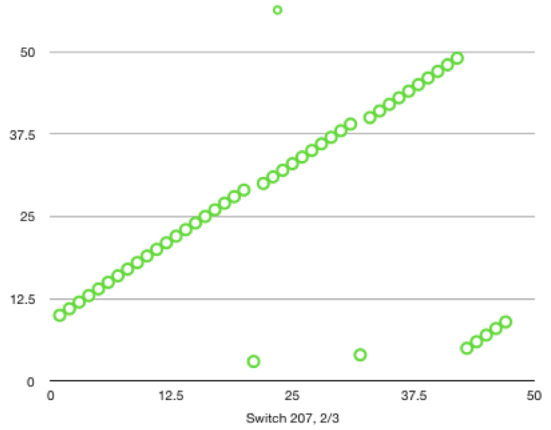

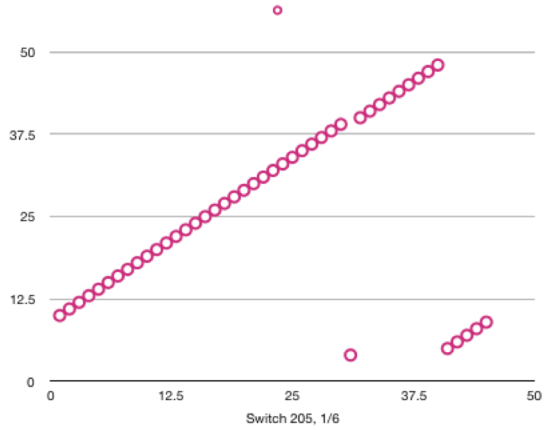

Incorporating Packet Path Into Predictive Data

Reconstruct Packet Path = Collector + Sequence Number + Ingress + Egress

Bucketing by Switch

- 10 Switches == 10 buckets
- Bucket values == hop latency if in sequence, else 0

Encoding Packet Paths

- Enumerate all possible paths through network
  - 4037913 (1! + 2! + … N!)
- Practically, this example only had about 15 unique paths

Hop Latency By Switch

SW203	SW204	SW205	SW206	SW207	SW208	SW226	SW227	SW246	SW247
13	0	0	0	0	11	10	0	11	0
31	0	0	0	0	26	9	0	10	0
0	0	0	0	10	15	0	0	0	13
0	463	0	0	0	42	0	11	0	17
0	0	11	0	0	0	0	12	0	14

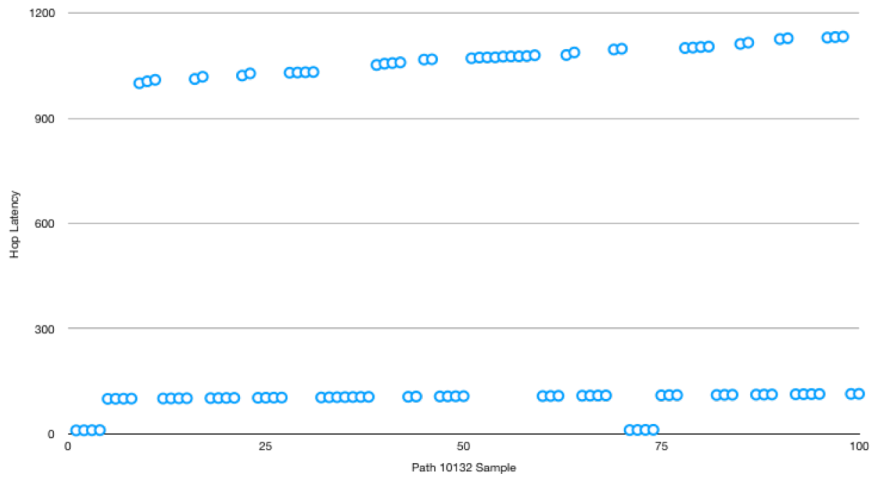

Connected Devices Correlate

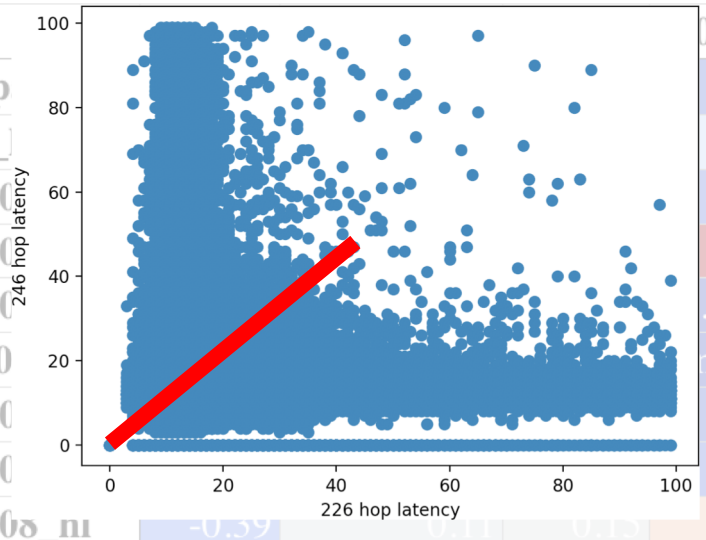

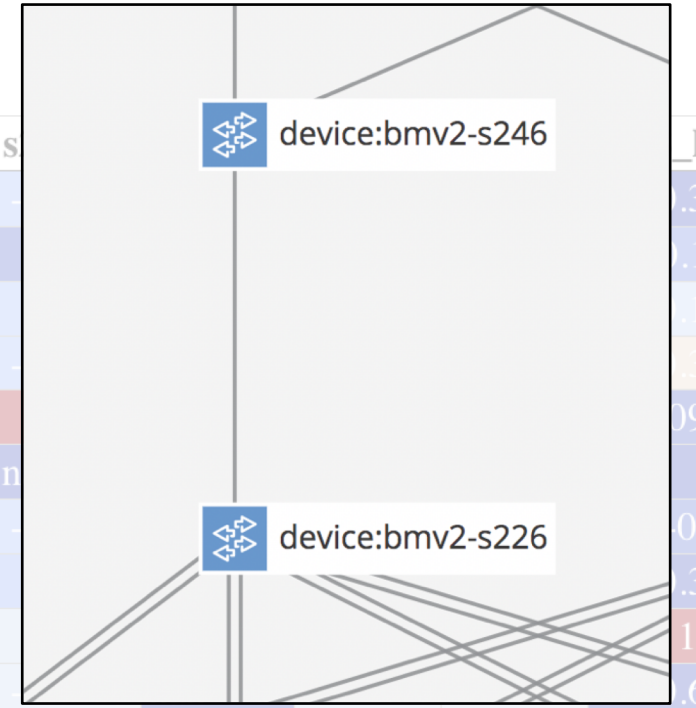

	s226_hl	s227_hl	s246_hl	s247_hl
p	0.45	0.26	0.6	nan
hop_	0.068	0.023	0.011	nan
s2	0.037	-0.36	0.038	nan
s2	-0.25	-0.016	-0.39	nan
s2	-0.0039	-0.011	-0.0024	nan
s20	nan	nan	nan	nan
s2	-0.01	0.45	0.13	nan
s2	0.019	0.44	0.16	nan
s20s_nl	-0.69	0.13	-0.58	nan
s226_hl	1.0	-0.51	0.78	nan
s227_hl	-0.51	1.0	-0.43	nan
s246_hl	0.78	-0.43	1.0	nan
s247_hl	nan	nan	nan	nan

Incorporating Packet Path Into Predictive Data

Reconstruct Packet Path = Collector + Sequence Number + Ingress + Egres

Bucketing by Switch

- 10 Switches == 10 buckets
- Bucket values == hop latency if in se ct

Encoding Packet Path

- Enu ers a po ible aths through network
  - 379 (2! + N!)
- Practically, this example only ha about 2 u ique paths

Hop Laten By

SW2	SW	SW205	6	SW207	SW208	SW226	SW227	SW246	SW247
13		0	0	0	11	10	0	11	0
1	0	0	0	0	26	9	0	10	0
0	0	0	0	10	15	0	0	0	13
0	463	0	0	0	42	0	11	0	17
0	0		0	0	0	0	12	0	14

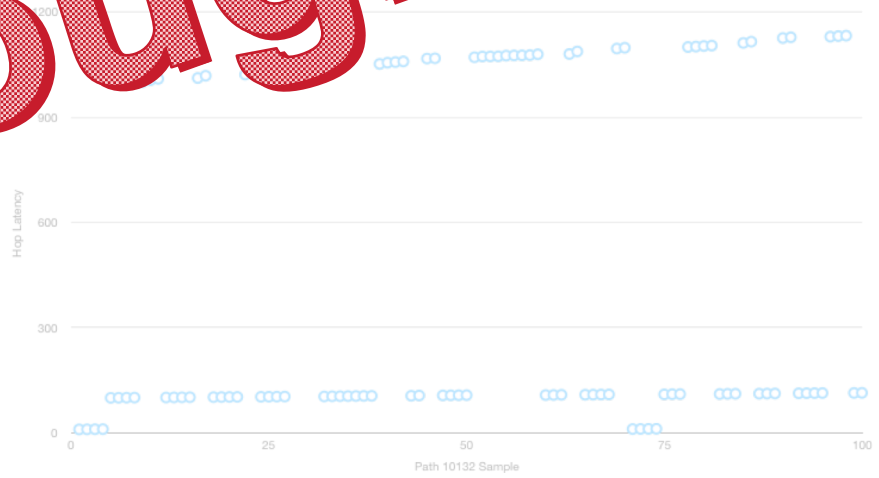

Hop Latency Not Enough

Oh, that is the Queue ID … Correlation

	switchid	ingress_port	egress_port	queue_occup	hop_latency
switchid	1.0	-0.16	-0.14	-0.19	0.079
ingress_port	-0.16	1.0	-0.17	0.025	-0.0037
egress_port	-0.14	-0.17	1.0	0.0065	-0.044
queue_occup	-0.19	0.025	0.0065	1.0	0.3
hop_latency	0.079	-0.0037	-0.044	0.3	1.0

Spearman Correlation

	switchid	ingress_port	egress_port	queue_occup	hop_latency
switchid	1.0	-0.1	-0.11	-0.089	-0.11
ingress_port	-0.1	1.0	-0.092	0.063	-0.0002
egress_port	-0.11	-0.092	1.0	-0.042	0.03
queue_occup	-0.089	0.063	-0.042	1.0	0.78
hop_latency	-0.11	-0.0002	0.03	0.78	1.0

Pearson Correlation

Correlate the World

	queue_occup	hop_latency	s203_hl	s203_qo	s204_hl	s204_qo	s205_hl	s206_hl	s206_qo	s207_hl	s207_qo	s208_hl	s208_qo	s226_hl	s226_qo	s227_hl	s227_qo	s246_hl	s246_qo
queue_occup	1.0	0.75	0.071	0.093	0.75	0.98	-0.00032	0.055	0.069	0.06	0.067	0.029	0.012	-0.014	0.0044	0.01	0.008	-0.0064	0.0093
hop_latency	0.75	1.0	0.096	0.082	0.98	0.74	0.0023	0.041	0.044	0.045	0.046	0.036	0.017	0.00064	0.013	0.0043	0.0086	-0.0052	0.013
s203_hl	0.071	0.096	1.0	0.84	-0.0077	-0.01	0.0018	-0.0025	-0.00092	-0.0032	-0.0017	0.0008	-0.00081	0.0007	-0.00033	-0.007	0.00093	0.00074	0.00058
s203_qo	0.093	0.082	0.84	1.0	-0.004	-0.0045	0.0012	-0.00095	-0.00044	-0.0016	-0.00092	-0.00015	-0.0006	0.00044	-0.00021	-0.0029	0.00042	0.00096	0.0011
s204_hl	0.75	0.98	-0.0077	-0.004	1.0	0.78	-0.0038	-0.0096	-0.0033	-0.0081	-0.003	0.016	-0.0018	-0.018	-0.0024	0.0029	0.00063	-0.024	-0.0036
s204_qo	0.98	0.74	-0.01	-0.0045	0.78	1.0	-0.0026	-0.011	-0.0039	-0.0047	-0.0025	0.022	0.00015	-0.021	-0.0019	0.0089	-0.0012	-0.019	-0.0028
s205_hl	-0.00032	0.0023	0.0018	0.0012	-0.0038	-0.0026	1.0	-0.00069	-0.0003	-0.0011	-0.00042	0.0014	0.00065	-0.00078	-7.7E-05	-0.0025	-0.00022	-0.0004	-0.00011
s206_hl	0.055	0.041	-0.0025	-0.00095	-0.0096	-0.011	-0.00069	1.0	0.92	-9.9E-05	-1.8E-05	-0.0025	-0.00053	-0.0042	-0.00022	0.019	0.0016	-0.002	-0.00044
s206_qo	0.069	0.044	-0.00092	-0.00044	-0.0033	-0.0039	-0.0003	0.92	1.0	-0.00017	-0.00042	-0.0012	-0.00042	-0.0027	-0.0003	0.0082	0.00033	-0.0019	-0.00045
s207_hl	0.06	0.045	-0.0032	-0.0016	-0.0081	-0.0047	-0.0011	-9.9E-05	-0.00017	1.0	0.92	-0.0035	4.9E-05	-0.003	-0.0002	0.017	0.00069	-0.0007	5.3E-05
s207_qo	0.067	0.046	-0.0017	-0.00092	-0.003	-0.0025	-0.00042	-1.8E-05	-0.00042	0.92	1.0	-0.00068	0.00011	-0.0023	-0.0002	0.0077	0.0002	-0.0017	-0.00027
s208_hl	0.029	0.036	0.0008	-0.00015	0.016	0.022	0.0014	-0.0025	-0.0012	-0.0035	-0.00068	1.0	0.85	-0.019	-0.0015	0.011	0.00096	-0.014	-0.0022
s208_qo	0.012	0.017	-0.00081	-0.0006	-0.0018	0.00015	0.00065	-0.00053	-0.00042	4.9E-05	0.00011	0.85	1.0	-0.0022	-0.00017	0.002	-0.00012	-0.0017	-0.00026
s226_hl	-0.014	0.00064	0.0007	0.00044	-0.018	-0.021	-0.00078	-0.0042	-0.0027	-0.003	-0.0023	-0.019	-0.0022	1.0	0.84	-0.024	-0.0015	0.092	0.049
s226_qo	0.0044	0.013	-0.00033	-0.00021	-0.0024	-0.0019	-7.7E-05	-0.00022	-0.0003	-0.0002	-0.0002	-0.0015	-0.00017	0.84	1.0	-0.0019	-0.00012	0.075	0.05
s227_hl	0.01	0.0043	-0.007	-0.0029	0.0029	0.0089	-0.0025	0.019	0.0082	0.017	0.0077	0.011	0.002	-0.024	-0.0019	1.0	0.88	-0.019	-0.0029
s227_qo	0.008	0.0086	0.00093	0.00042	0.00063	-0.0012	-0.00022	0.0016	0.00033	0.00069	0.0002	0.00096	-0.00012	-0.0015	-0.00012	0.88	1.0	-0.0012	-0.00018
s246_hl	-0.0064	-0.0052	0.00074	0.00096	-0.024	-0.019	-0.0004	-0.002	-0.0019	-0.0007	-0.0017	-0.014	-0.0017	0.092	0.075	-0.019	-0.0012	1.0	0.91
s246_qo	0.0093	0.013	0.00058	0.0011	-0.0036	-0.0028	-0.00011	-0.00044	-0.00045	5.3E-05	-0.00027	-0.0022	-0.00026	0.049	0.05	-0.0029	-0.00018	0.91	1.0

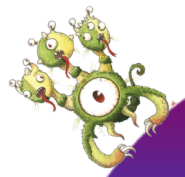

Global Queue Occupancy v. Hop Latency

Queue Occupancy Status Was There All the Time

- Looking for correlation between Queue Occupancy and Hop Latency
- While general correlation seems to exist, many outliers

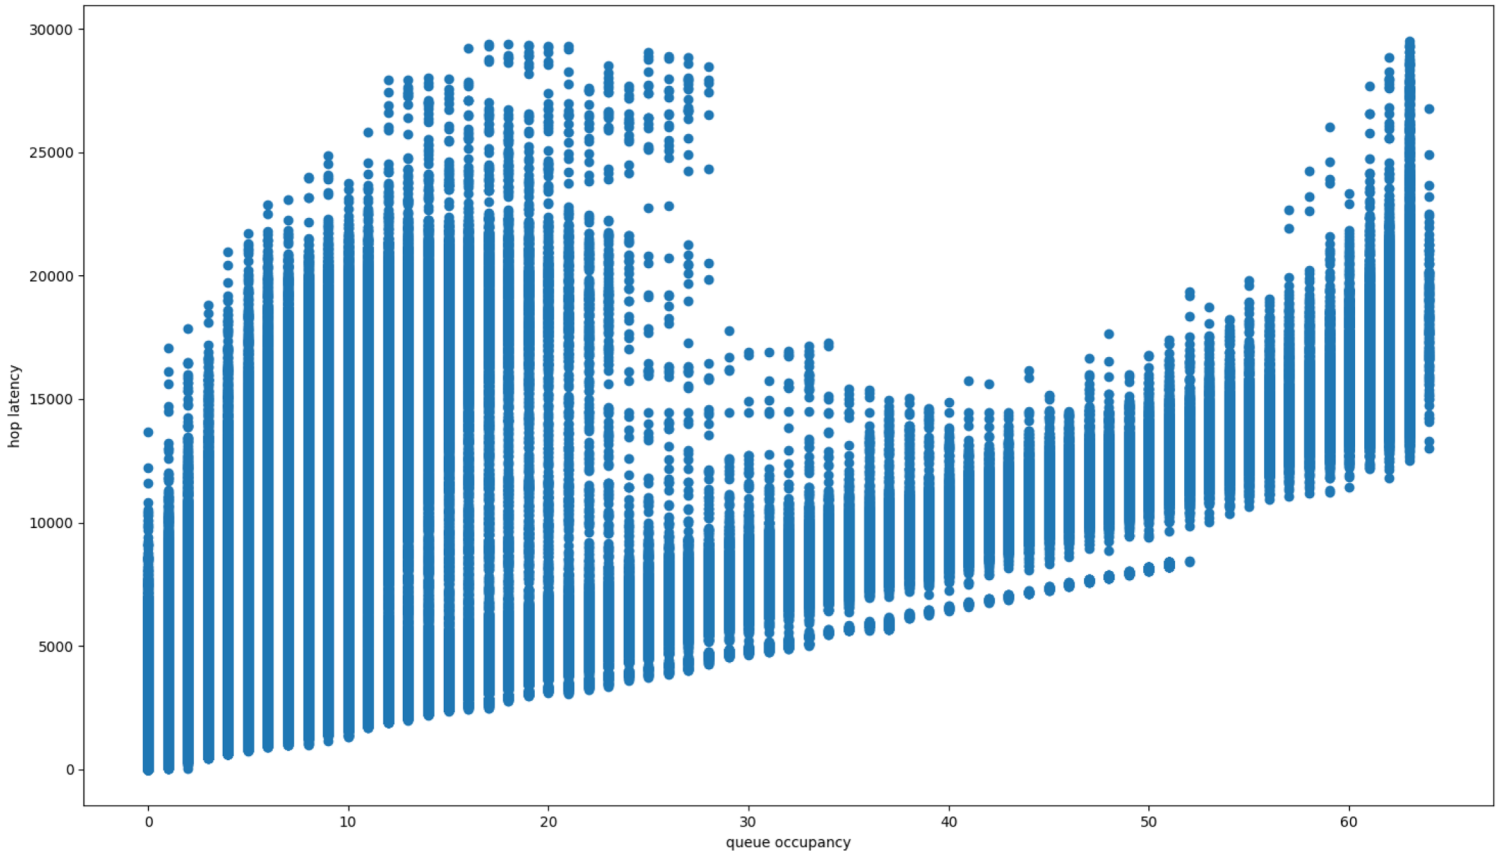

Queue Occupancy v. Hop Latency by Switch

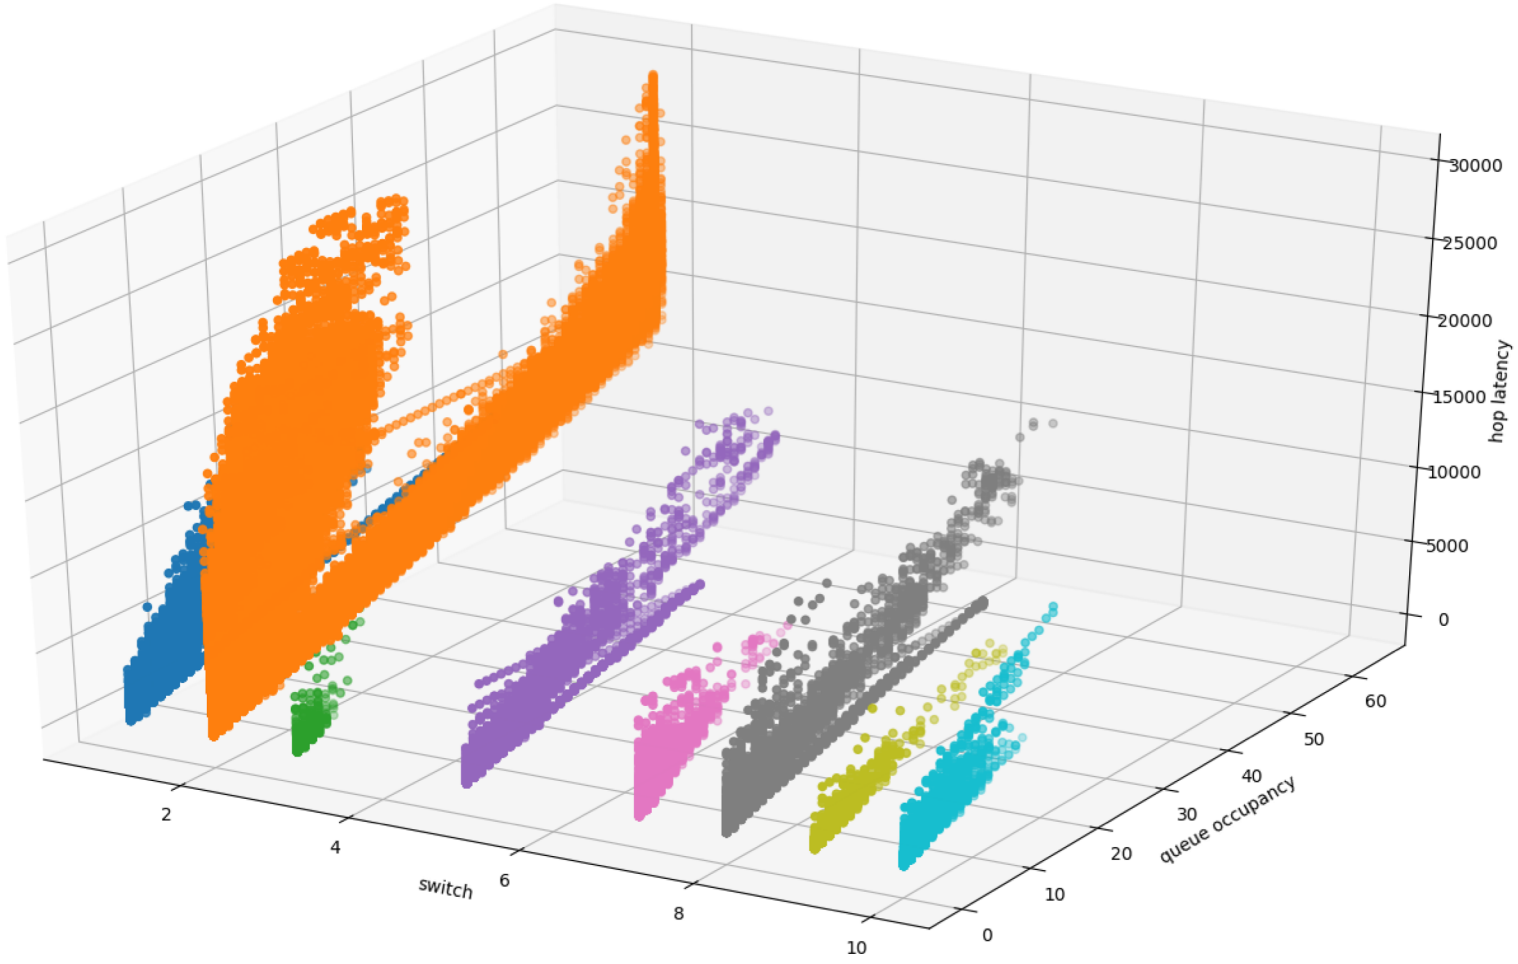

Training

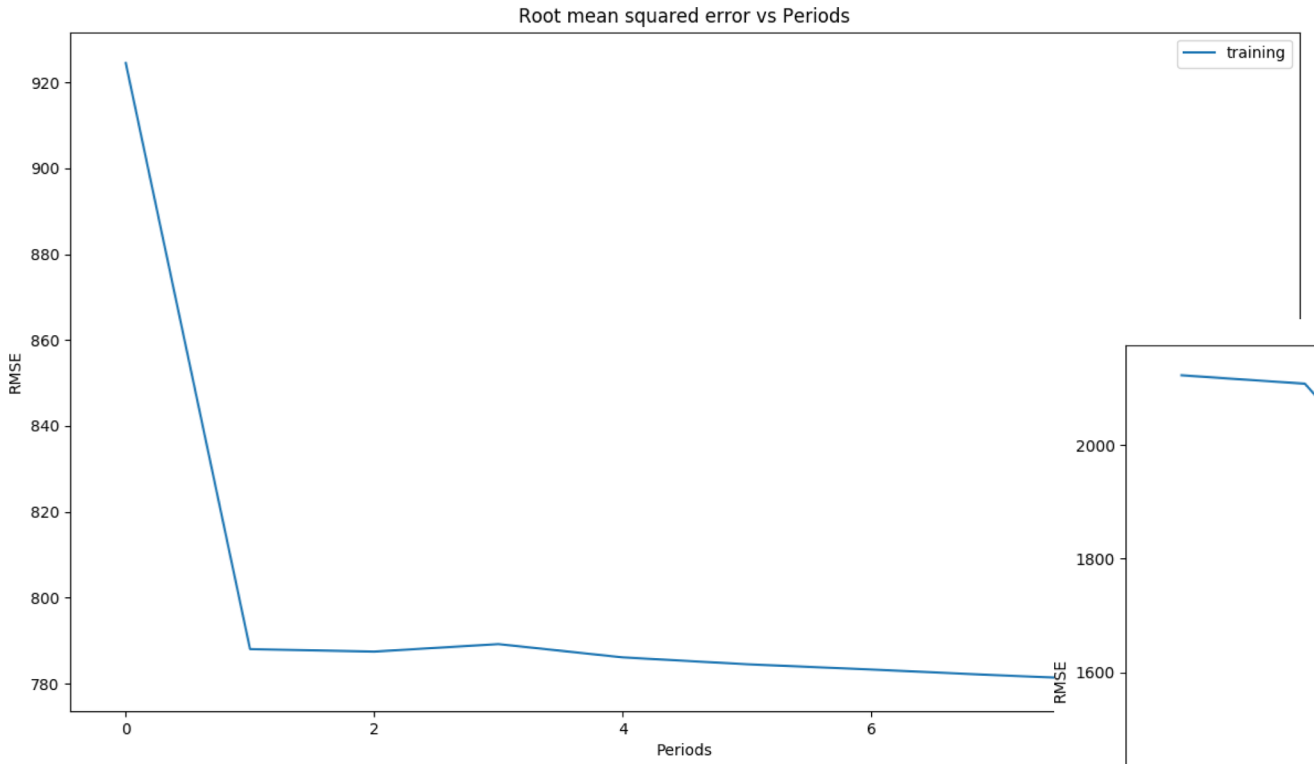

Without Path

With Path

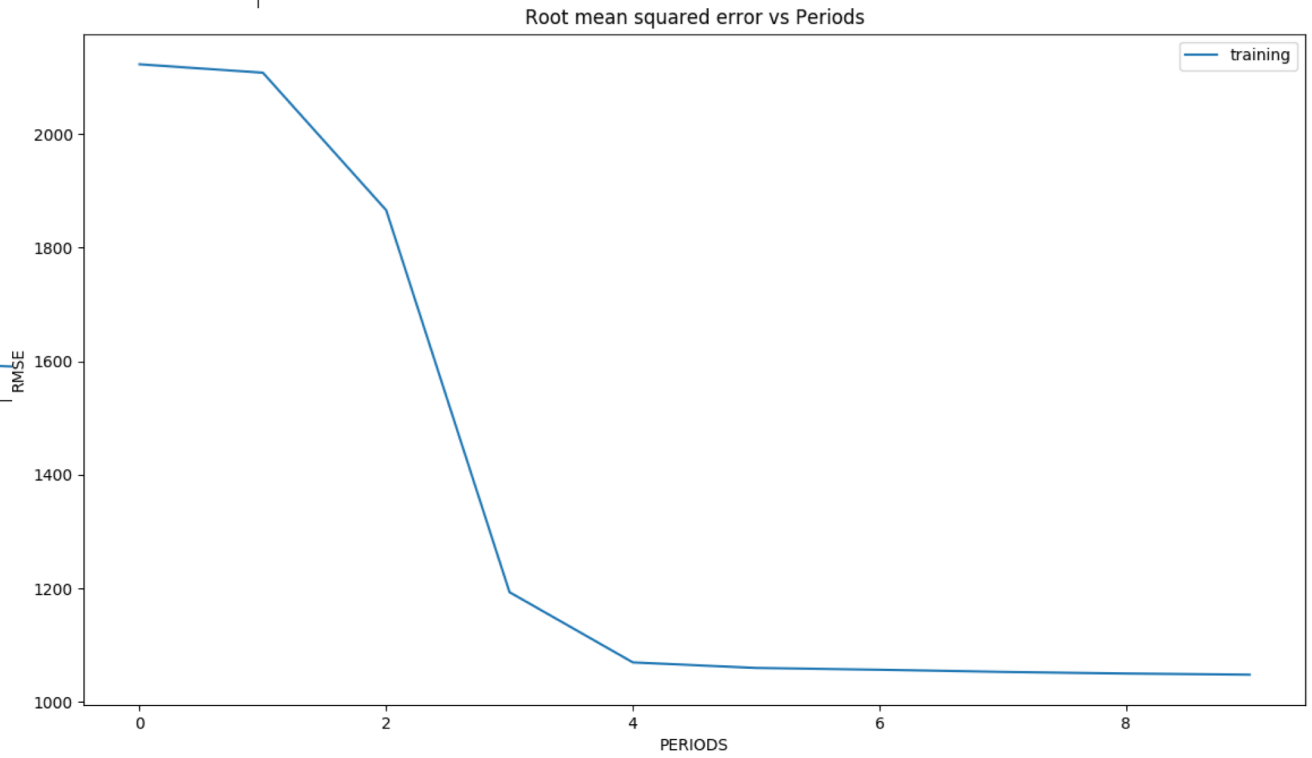

Predicted Hop Latency (no path component)

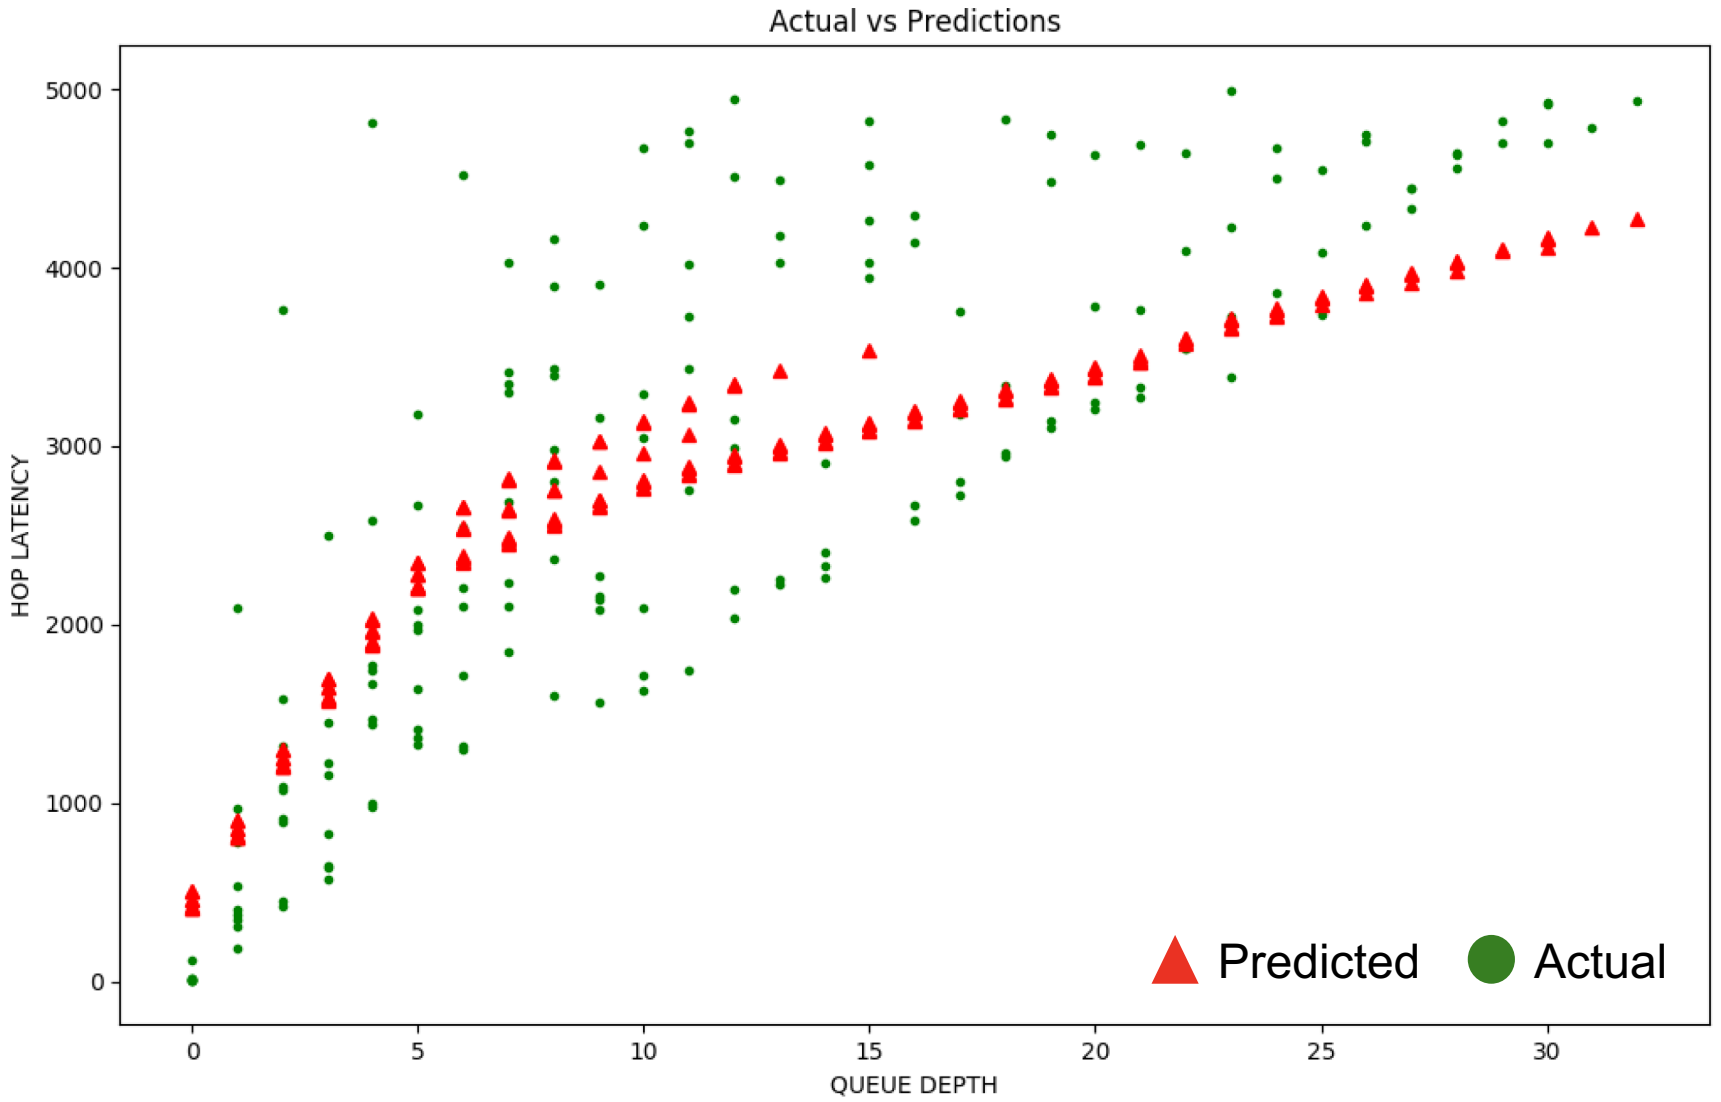

Predicted Hop Latency (with path component)

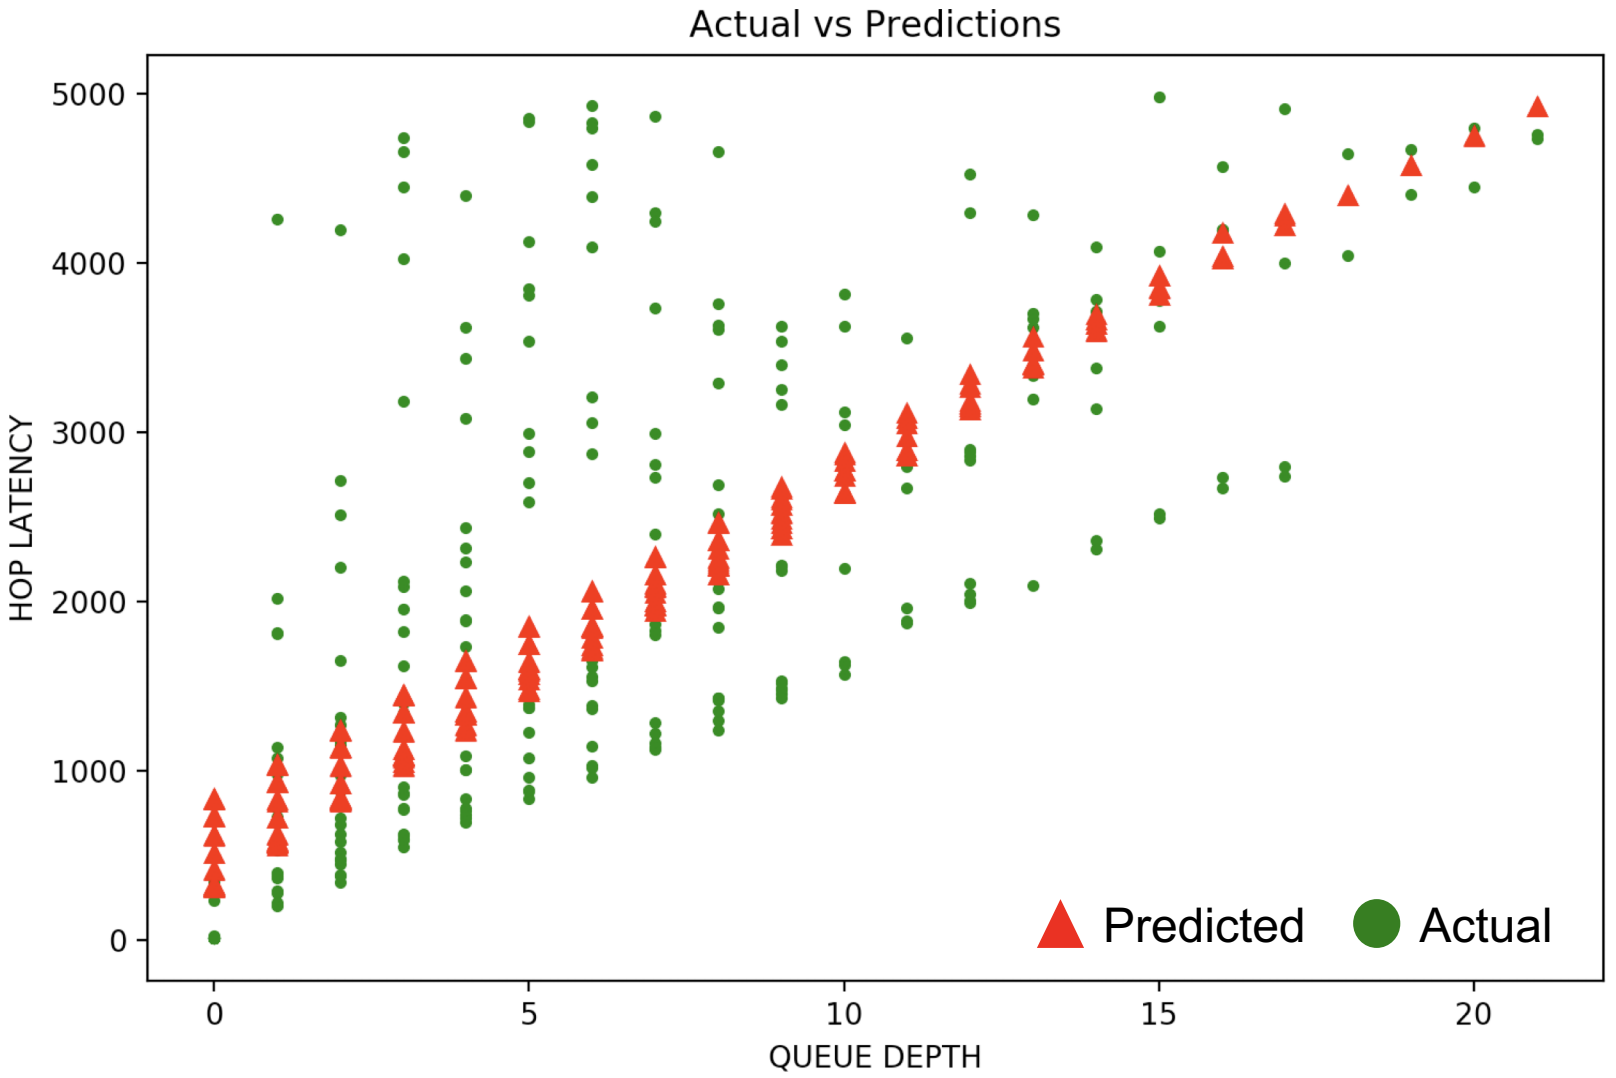

Incorporating ONOS Port Statistics

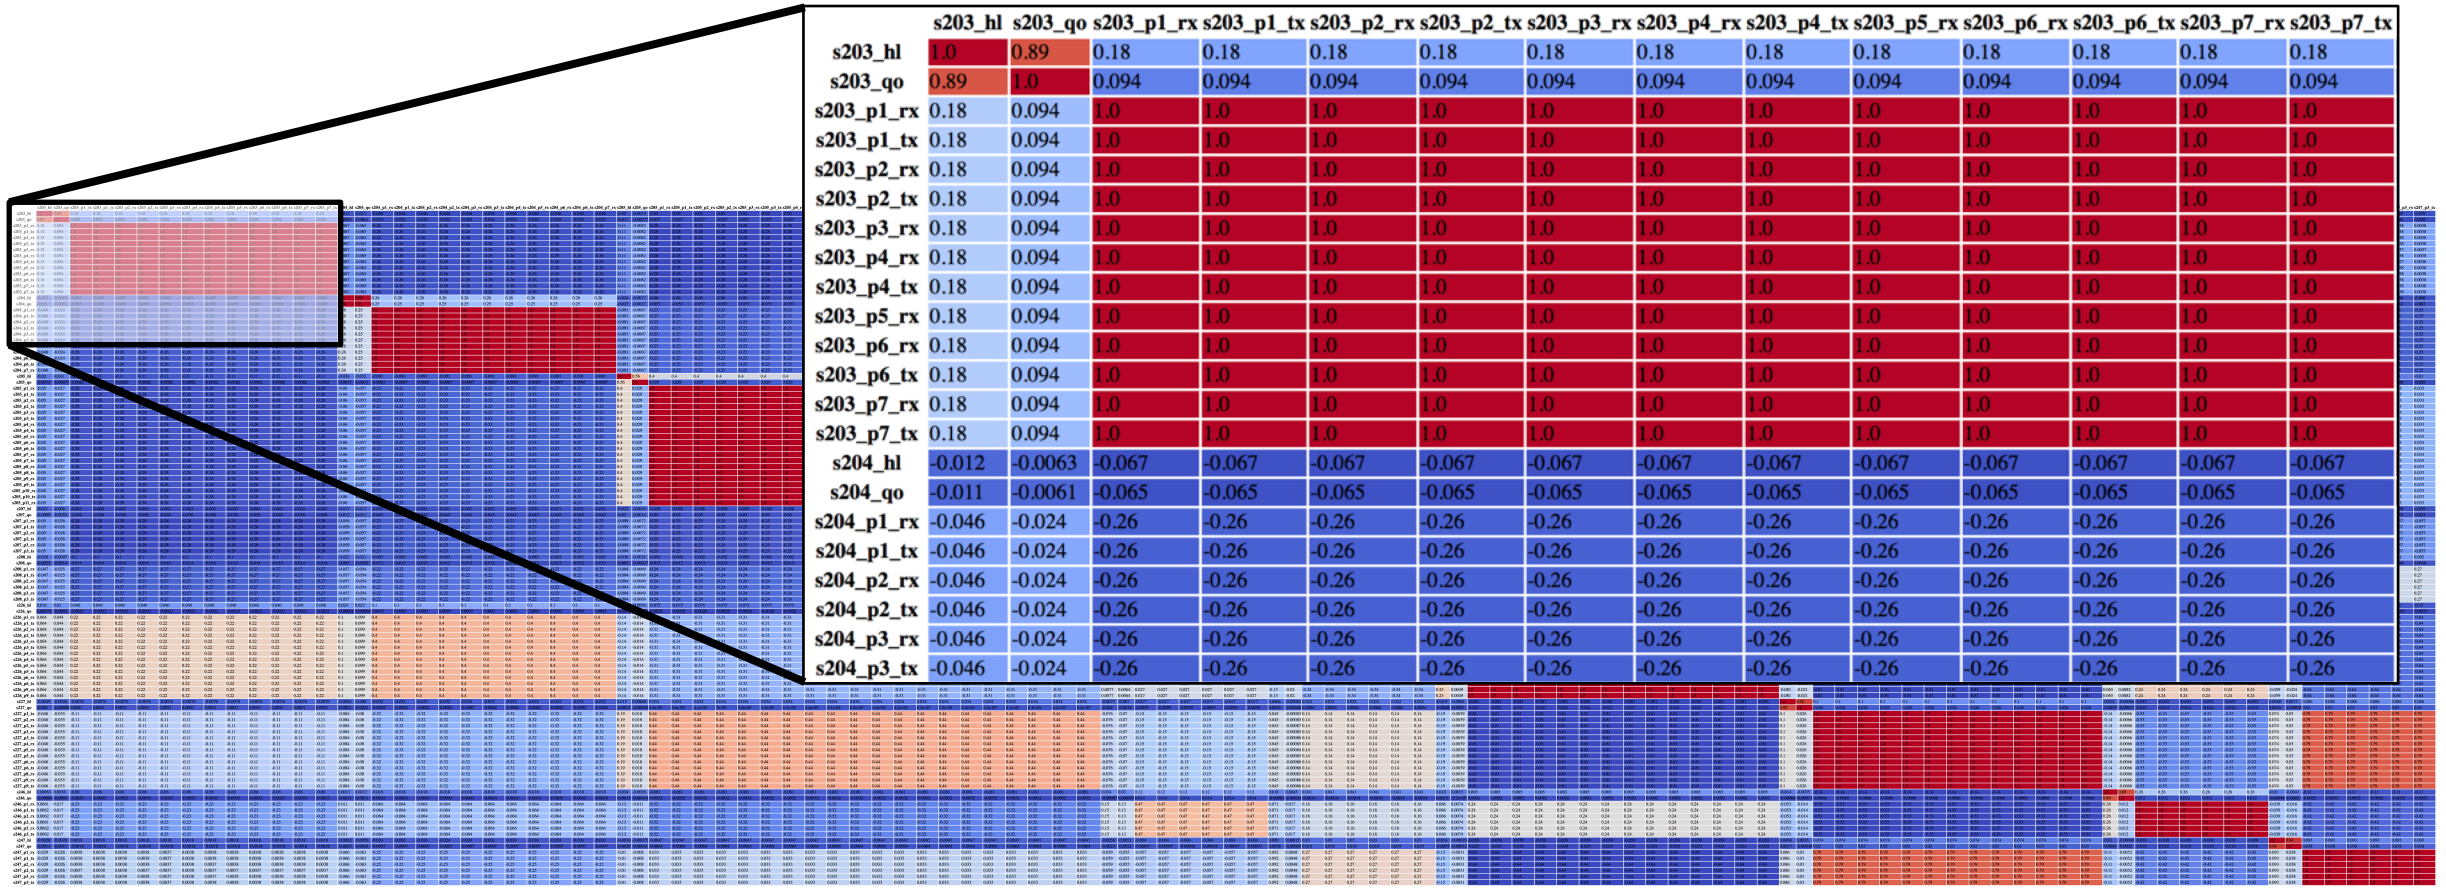

	s203_hl	s203_qo	s203_p1_rx	s203_p1_tx	s203_p2_rx	s203_p2_tx	s203_p3_rx	s203_p4_rx	s203_p4_tx	s203_p5_rx	s203_p6_rx	s203_p6_tx	s203_p7_rx	s203_p7_tx
s203_hl	1.0	0.89	0.18	0.18	0.18	0.18	0.18	0.18	0.18	0.18	0.18	0.18	0.18	0.18
s203_qo	0.89	1.0	0.094	0.094	0.094	0.094	0.094	0.094	0.094	0.094	0.094	0.094	0.094	0.094
s203_p1_rx	0.18	0.094	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
s203_p1_tx	0.18	0.094	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
s203_p2_rx	0.18	0.094	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
s203_p2_tx	0.18	0.094	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
s203_p3_rx	0.18	0.094	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
s203_p4_rx	0.18	0.094	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
s203_p4_tx	0.18	0.094	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
s203_p5_rx	0.18	0.094	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
s203_p6_rx	0.18	0.094	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
s203_p6_tx	0.18	0.094	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
s203_p7_rx	0.18	0.094	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
s203_p7_tx	0.18	0.094	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
s204_hl	-0.012	-0.0063	-0.067	-0.067	-0.067	-0.067	-0.067	-0.067	-0.067	-0.067	-0.067	-0.067	-0.067	-0.067
s204_qo	-0.011	-0.0061	-0.065	-0.065	-0.065	-0.065	-0.065	-0.065	-0.065	-0.065	-0.065	-0.065	-0.065	-0.065
s204_p1_rx	-0.046	-0.024	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26
s204_p1_tx	-0.046	-0.024	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26
s204_p2_rx	-0.046	-0.024	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26
s204_p2_tx	-0.046	-0.024	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26
s204_p3_rx	-0.046	-0.024	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26
s204_p3_tx	-0.046	-0.024	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26	-0.26

* Some [useless] columns/rows removed

Post Mortem

Data Volume

- Around 6 million records in around 10 mins on small 10 node test network
- Need to work on sampling and understand when to increase or decrease data collected as well as sample rate

Knowledge Acquisition

- Data trained and weights calculated on test network not directly applicable to other networks
- Would need to train on each network
- Possible network paths is an important characteristic for prediction
- Far too much up front, deployment specific data analysis required for practical deployments

Still Unclear INT is providing better data for prediction

- More data, certainly
- Actual packet paths, real observations

Direction – Adaptive Model Generation (simplified)

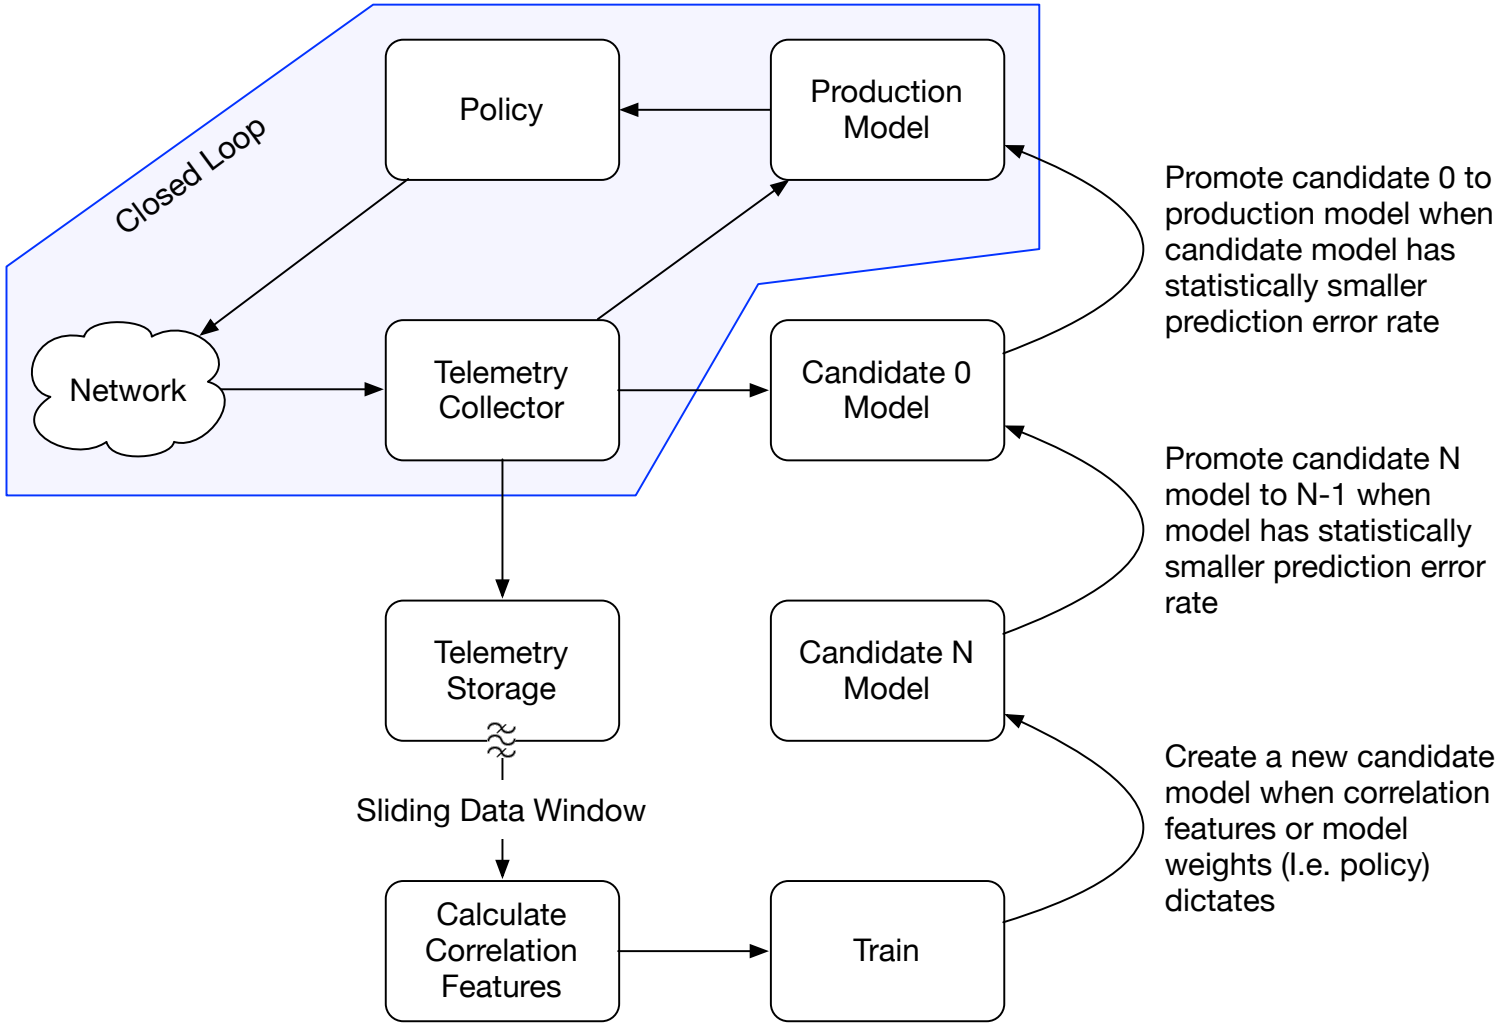

Direction – Adaptive Closed Loop

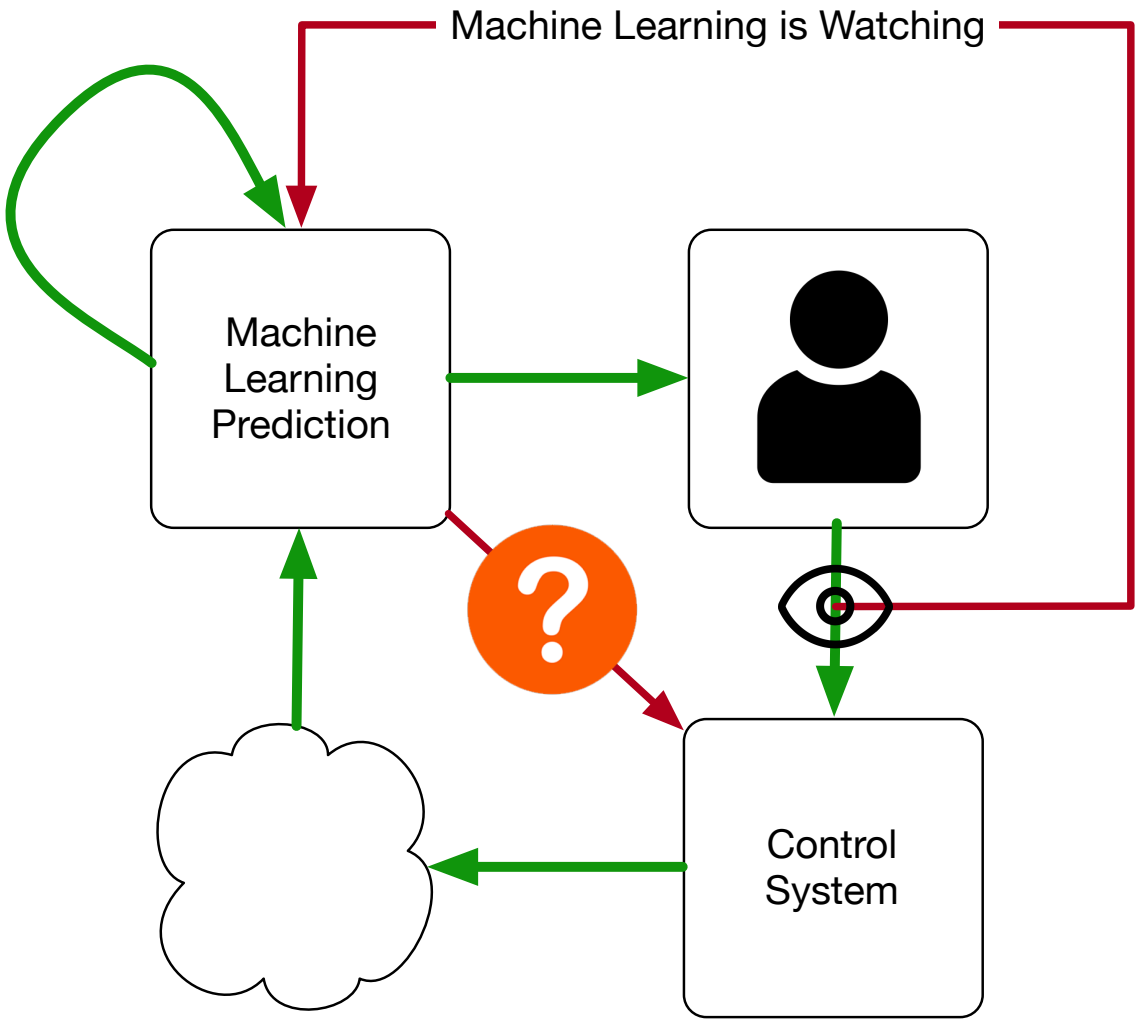

Direction - Adaptive Machine Learning

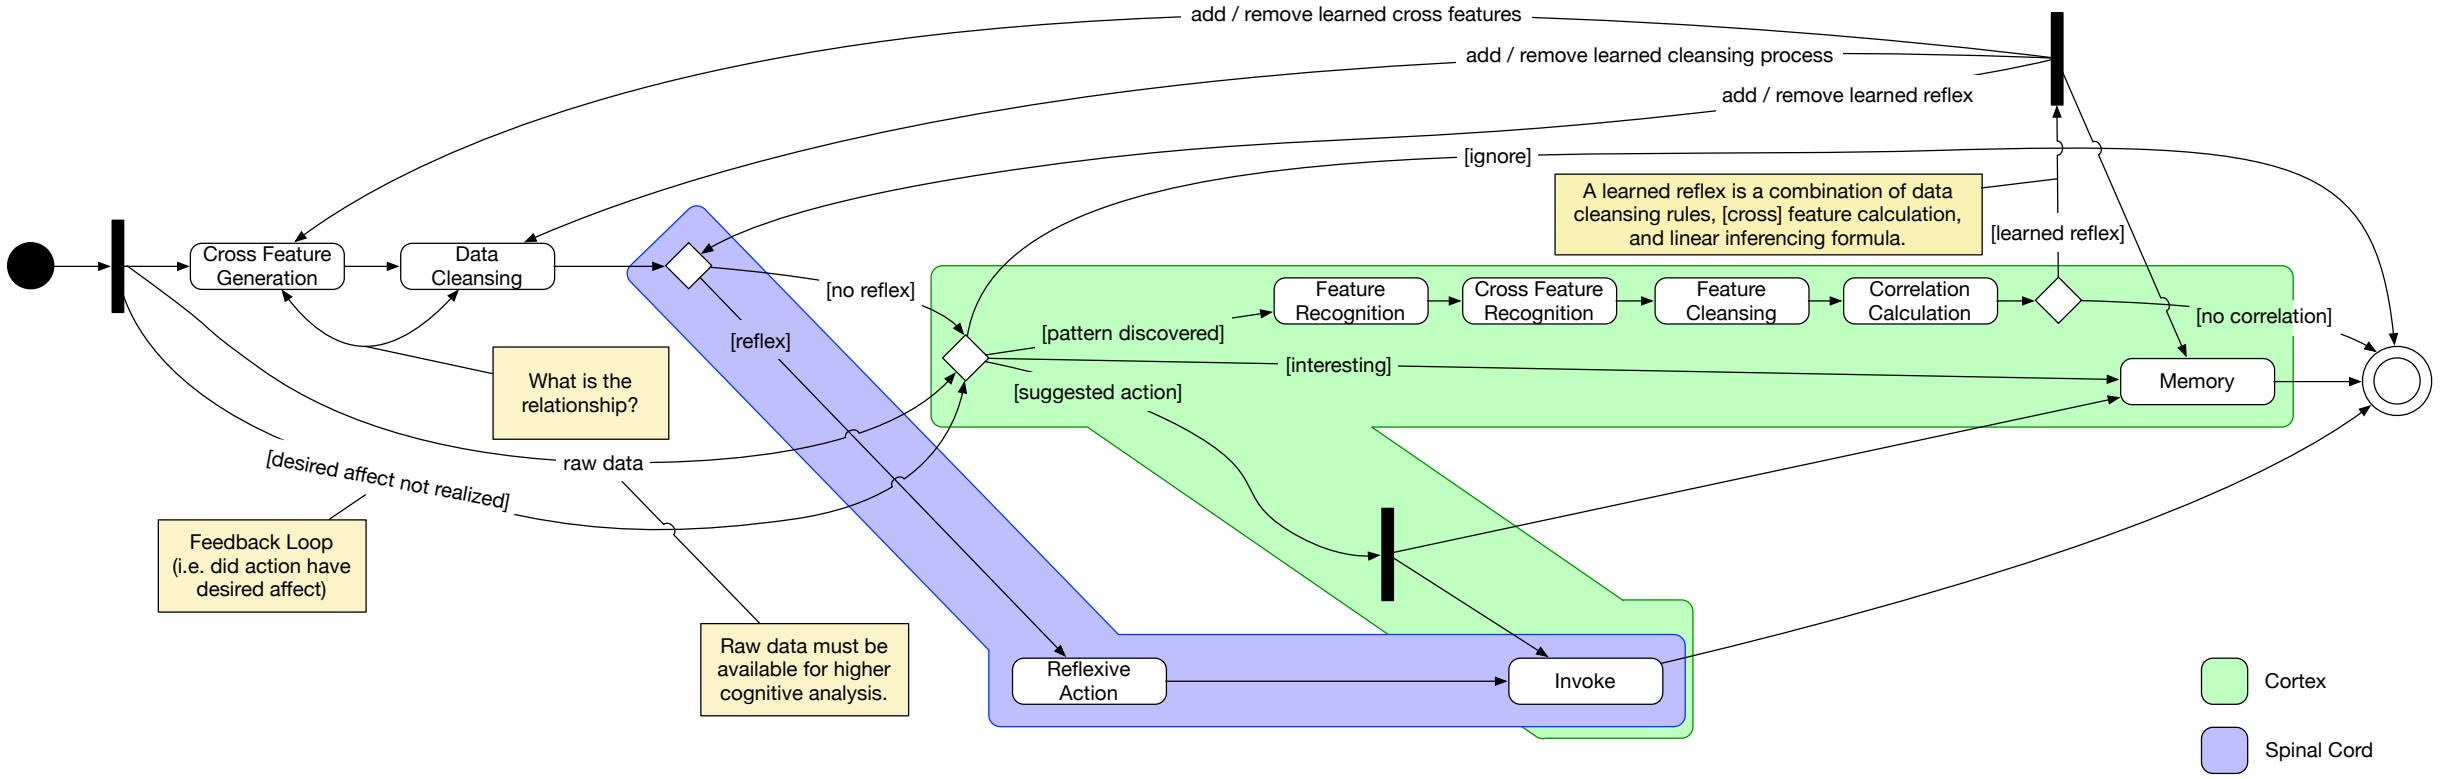

"As data scientists, our job is to extract signal from noise."

Daniel Tunkelang, Consultant / Advisor

"If Your Data Is Bad, Your Machine Learning Tools Are Useless"

Thomas C. Redman - https://hbr.org/2018/04/if-your-data-is-bad-your-machine-learning-tools-are-useless

Coming soon to open source …

Plan to open source the code that was used for this experiment

- Need to clean things up a bit to make it easier to consume
- Will post to ONF mailing list when it is available

Mèsi