



Thomas Sorkin
University of San Francisco

ANALYSIS OF THE SLINGSHOT HIGH-SPEED NETWORK

SIParCS 2024

Mentors: Will Shanks and Storm Knight

July 30, 2024





Purpose of the Project

We have a
supercomputer network,
but we don't know
exactly what it's doing

Table of Contents

BACKGROUND

- What is a High-speed Network?
- What is Slingshot?
- What is Dragonfly?
- What is Derecho? (And Gust)

COUNTER COLLECTION

- dump_counters and lldp
- Putting counters into TimescaleDB
- slingshot-topology-tool

VISUALIZATIONS

- Gust: Graphviz, pyvis, networkx, Gephi, and Grafana
- Derecho: Counter visualizations
- Future Work


BACKGROUND

High-Speed Networks, Slingshot, and Dragonfly




High-Speed Networks (in HPC)

- Need to move a lot of data around very quickly
 - Copper cables/Fiber-optic Cables
 - Switches
 - Blades
- How do you connect this complex system?
 - Ethernet
 - Infiniband
 - Slingshot





Slingshot

- Rosetta Switches
 - 64 ports, 200Gbps
- HPE Ethernet
- Network features:
 - Adaptive Packet Routing
 - Congestion Control
 - Quality of Service




<https://www.nextplatform.com/2019/08/16/how-cray-makes-ethernet-suited-for-hpc-and-ai-with-slingshot/>

Fat Tree




Hypercube

Torus



Dragonfly Topology

- 3-hop Maximum!
 - Highly connected structure
- Based around **groups**
- Is the shortest path the best path?



Derecho and Gust

- Latest NCAR Supercomputer
 - Gust is the testbed system
- Derecho: 2570 nodes, 200 switches
 - 2488 CPU nodes, 82 GPU nodes
- Gust: 18 nodes, 10 switches
 - 16 CPU nodes, 2 GPU nodes



Counter Collection

Counter discovery, Telegraf, and TimescaleDB



Switch Counters

- Main utility: dump_counters
 - Switch-by-switch basis
 - General use counters AND specialized counters
- Link-Layer Discovery Protocol
- Various counters hidden in the file system

```
00, eeg_tx_ok_00_to_127
00, eeg_tx_ok_128_to_255
00, eeg_tx_ok_256_to_511
00, eeg_tx_ok_512_to_1023
00, eeg_tx_ok_1024_to_2047
00, itf_rx_bad_opt
00, itf_rx_ok_undersize
00, itf_rx_ok_64
00, itf_rx_ok_65_to_127
00, itf_rx_ok_128_to_255
00, itf_rx_ok_256_to_511
00, itf_rx_ok_512_to_1023
00, itf_rx_ok_1024_to_2047
00, eeg_sts_tx_bad
00, eeg_itf_sts_rx_bad
00, IfInOctets
00, IfInUcastPkts
00, IfInNUcastPkts
00, IfInErrors
00, IfOutOctets
00, IfOutUcastPkts
```

Consolidating and Aggregating Data

- Poll 200 switches regularly
 - Avoid stressing network devices
 - Fabric managers and admin nodes
- **Telegraf**
 - Tool to collect/send metrics
 - Plug-in based
 - Plug-ins for network telemetry
- **Lightweight Distributed Metric Service (LDMS)**



Storing Counters

- TimescaleDB
 - Postgresql database
- Tables configured by Telegraf plug-in
- Need specificity
 - Account for every switch
 - Account for every port
 - Account for every counter



Discovering the Network Topology

- Need to know:
 - Where each switch connects
 - Over what port(s)
 - Switch group #
- slingshot-topology-tool

xname	type	dst	hostname	port_num
x9000c1r3j11p1	fabric	x9000c3r3j11	unknown	14
x9000c1r3j11p0	fabric	x9000c3r3j11	unknown	15
x9000c1r3j13p1	fabric	x9000c3r1j11	unknown	62
x9000c1r3j13p0	fabric	x9000c3r1j11	unknown	63
x9000c1r3j18p1	fabric	x9000c1r7j24	unknown	43
x9000c1r3j18p0	fabric	x9000c1r7j24	unknown	42
x9000c1r3j20p1	fabric	x9000c1r7j22	unknown	56
x9000c1r3j20p0	fabric	x9000c1r7j22	unknown	57
x9000c1r3j22p1	fabric	x9000c1r7j20	unknown	55
x9000c1r3j22p0	fabric	x9000c1r7j20	unknown	54
x9000c1r3j24p1	fabric	x9000c1r7j18	unknown	36
x9000c1r3j24p0	fabric	x9000c1r7j18	unknown	37
x9000c1r3j8p1	fabric	x9000c1r7j2	unknown	27
x9000c1r3j8p0	fabric	x9000c1r7j2	unknown	26
x9000c1r3j6p1	fabric	x9000c1r7j4	unknown	8
x9000c1r3j6p0	fabric	x9000c1r7j4	unknown	9
x9000c1r3j4p1	fabric	x9000c1r7j6	unknown	7
x9000c1r3j4p0	fabric	x9000c1r7j6	unknown	6
x9000c1r3j2p1	fabric	x9000c1r7j8	unknown	20
x9000c1r3j2p0	fabric	x9000c1r7j8	unknown	21
x9000c1r3j1p1	fabric	x3001c0r36j2	unknown	4
x9000c1r3j1p0	fabric	x3001c0r36j2	unknown	5
x9000c1r3i3p1	fabric	x3001c0r35j2	unknown	23

Visualizations

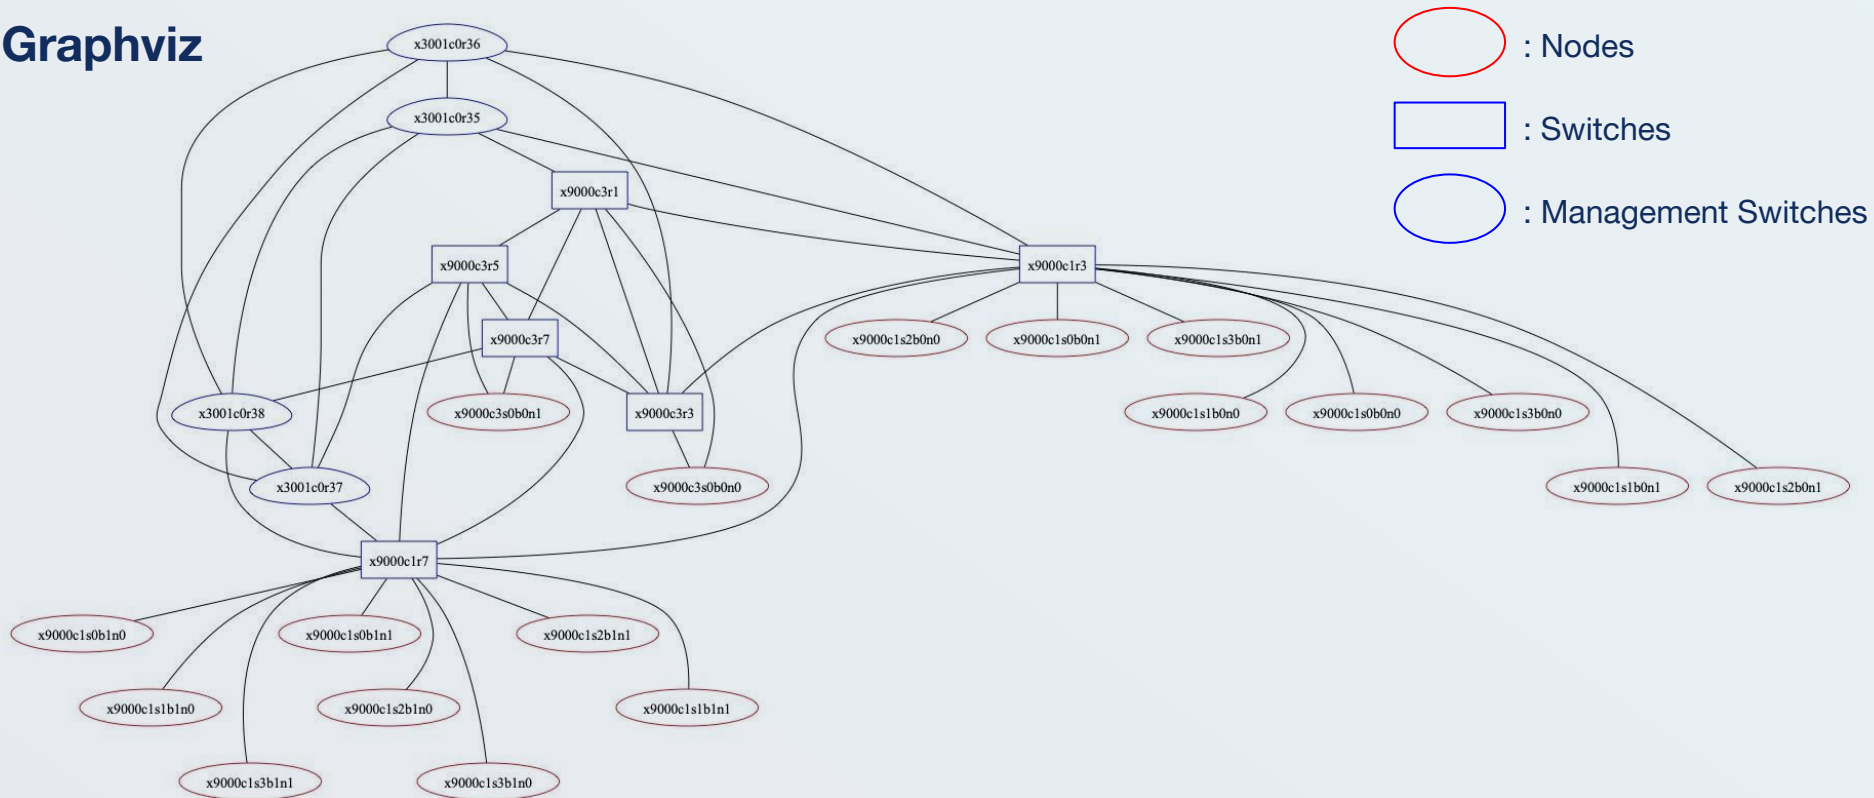
**Node graphs on Gust, Node graphs on Derecho, and Grafana
Heatmaps**



Different Visualization Tools (Gust)

- Graphviz, pyvis, Gephi, and networkx
- Switches grouped, nodes gone, counters added to gust

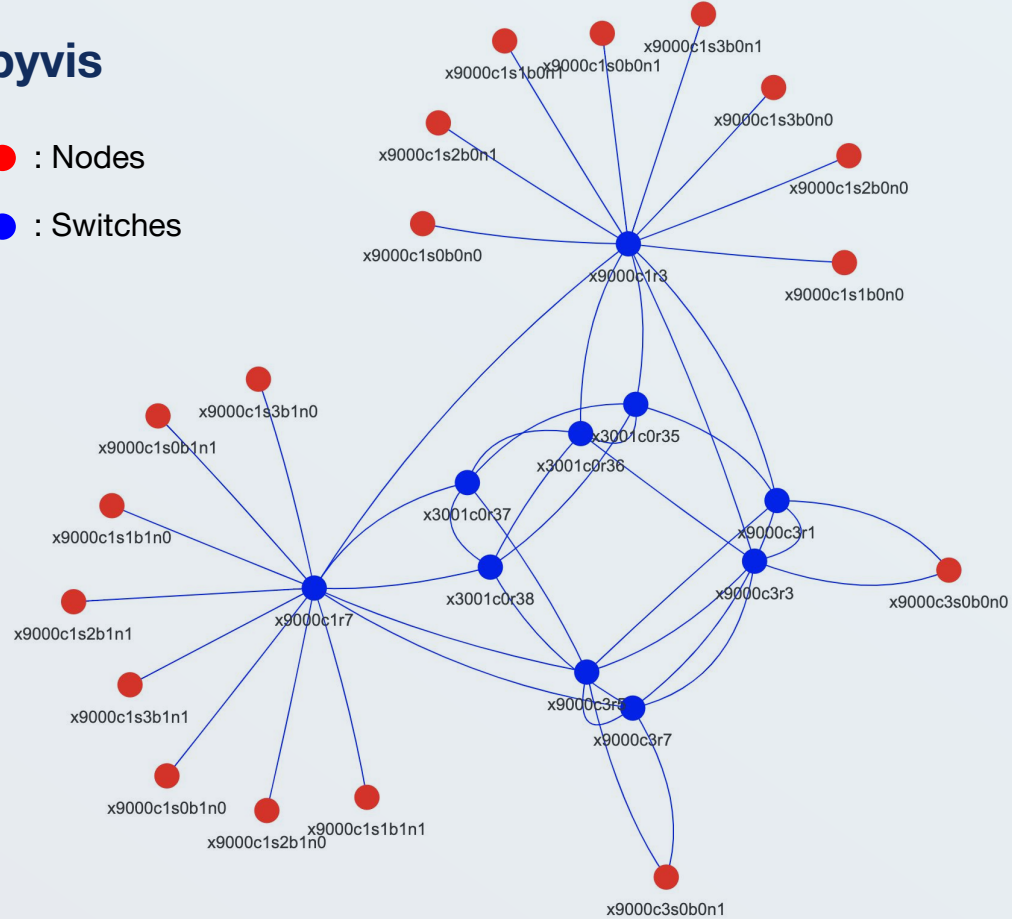
Graphviz



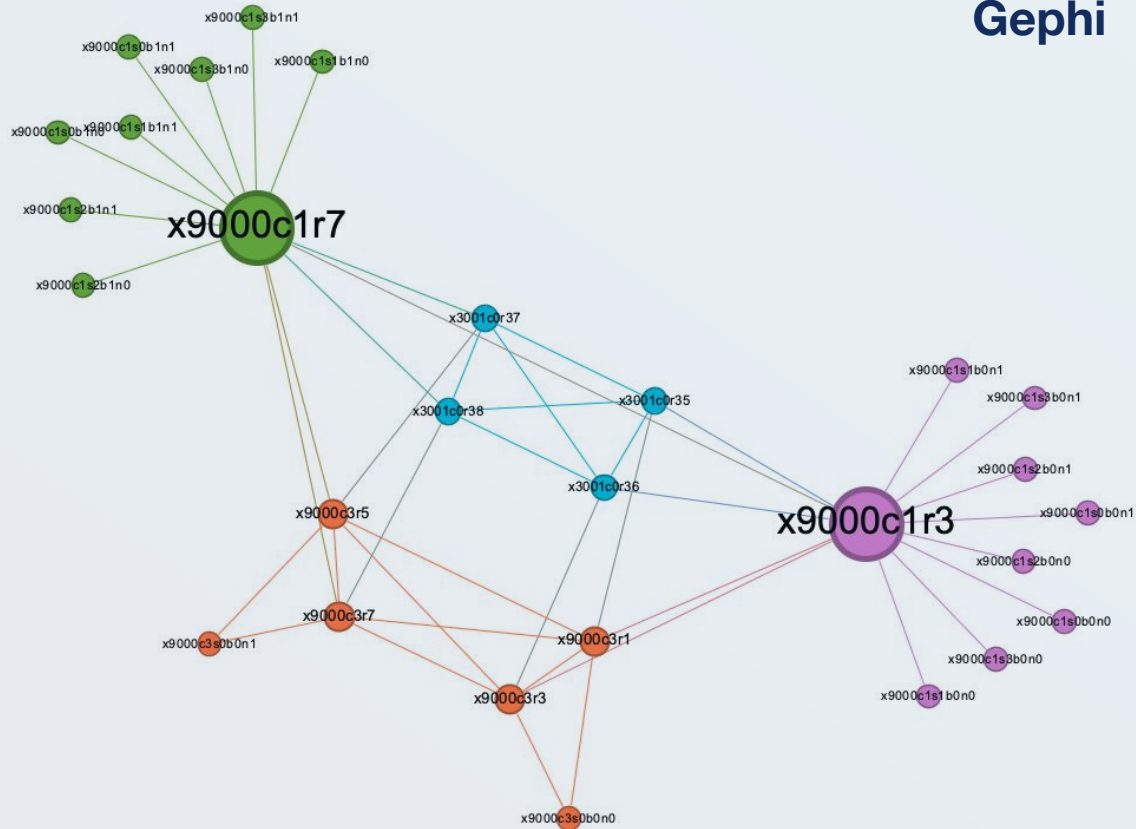
pyvis

● : Nodes

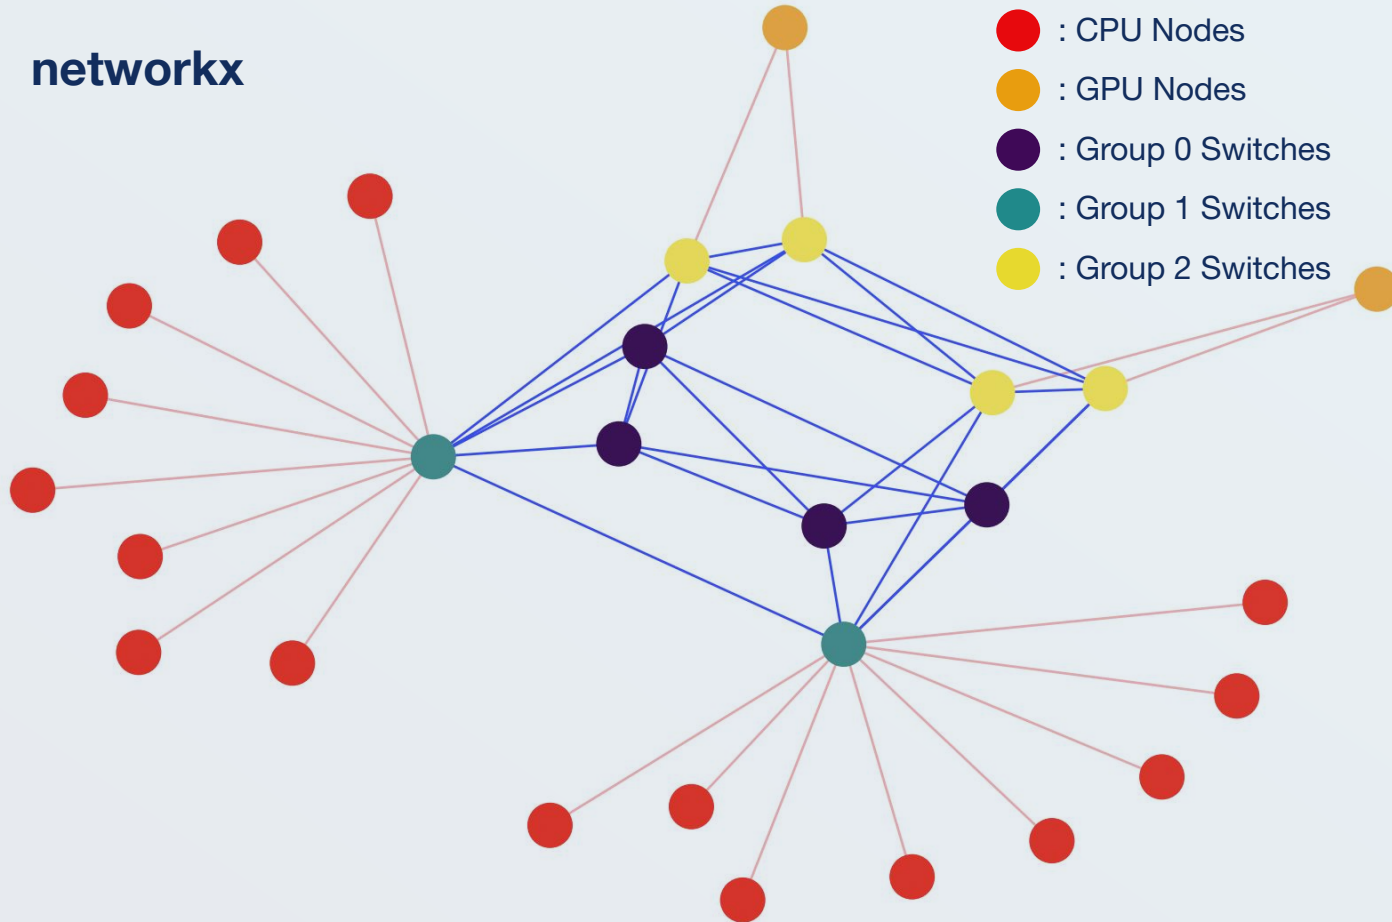
● : Switches




Gephi



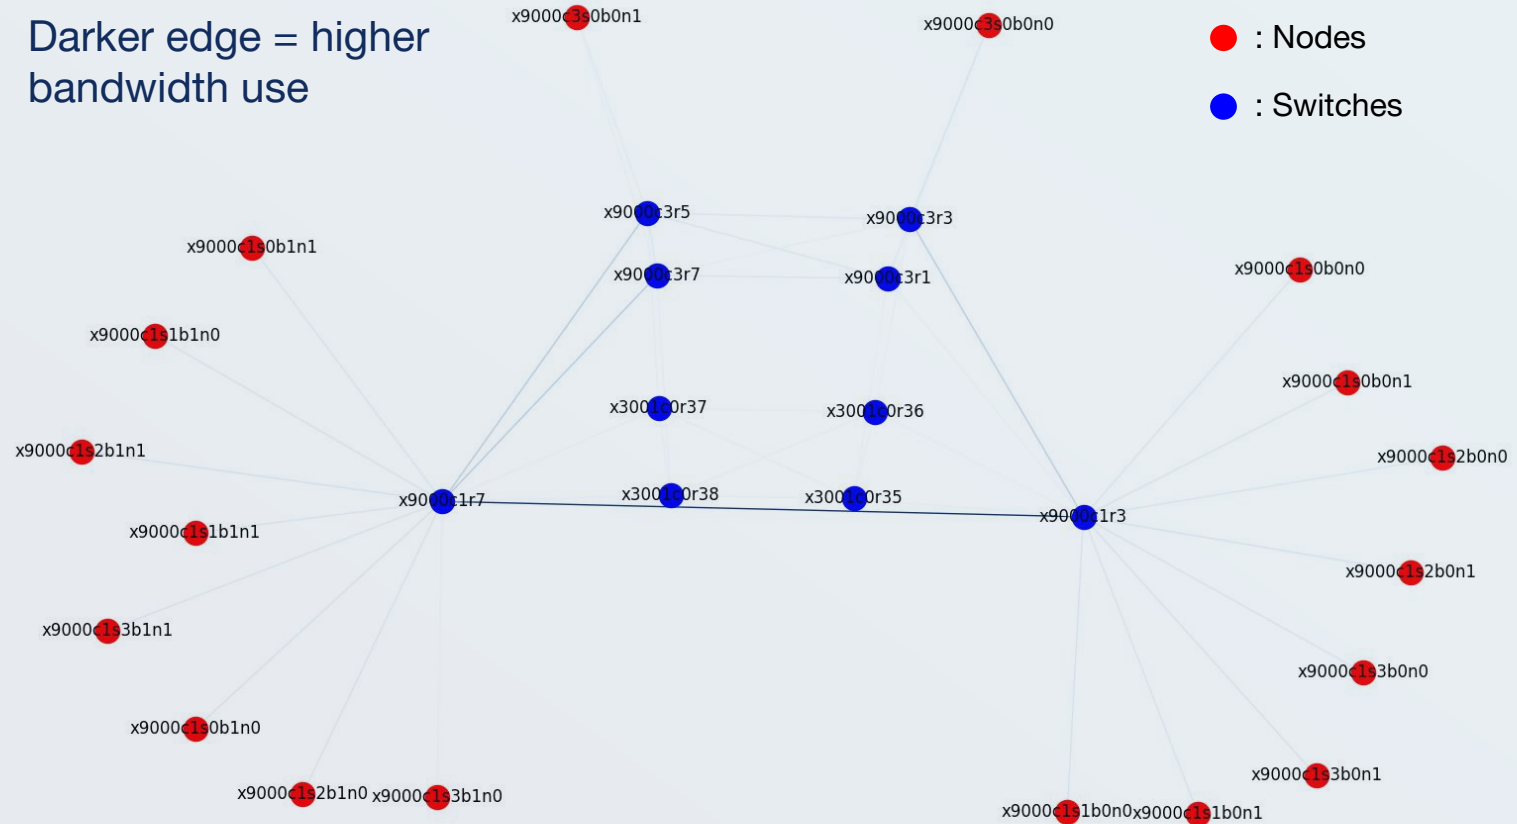
networkx



Switches
Grouped
Together



Darker edge = higher bandwidth use



Visualizing Derecho

- **Much** larger system
 - How do we make it intuitive?
- Node graphs and heatmaps
- Timelapses

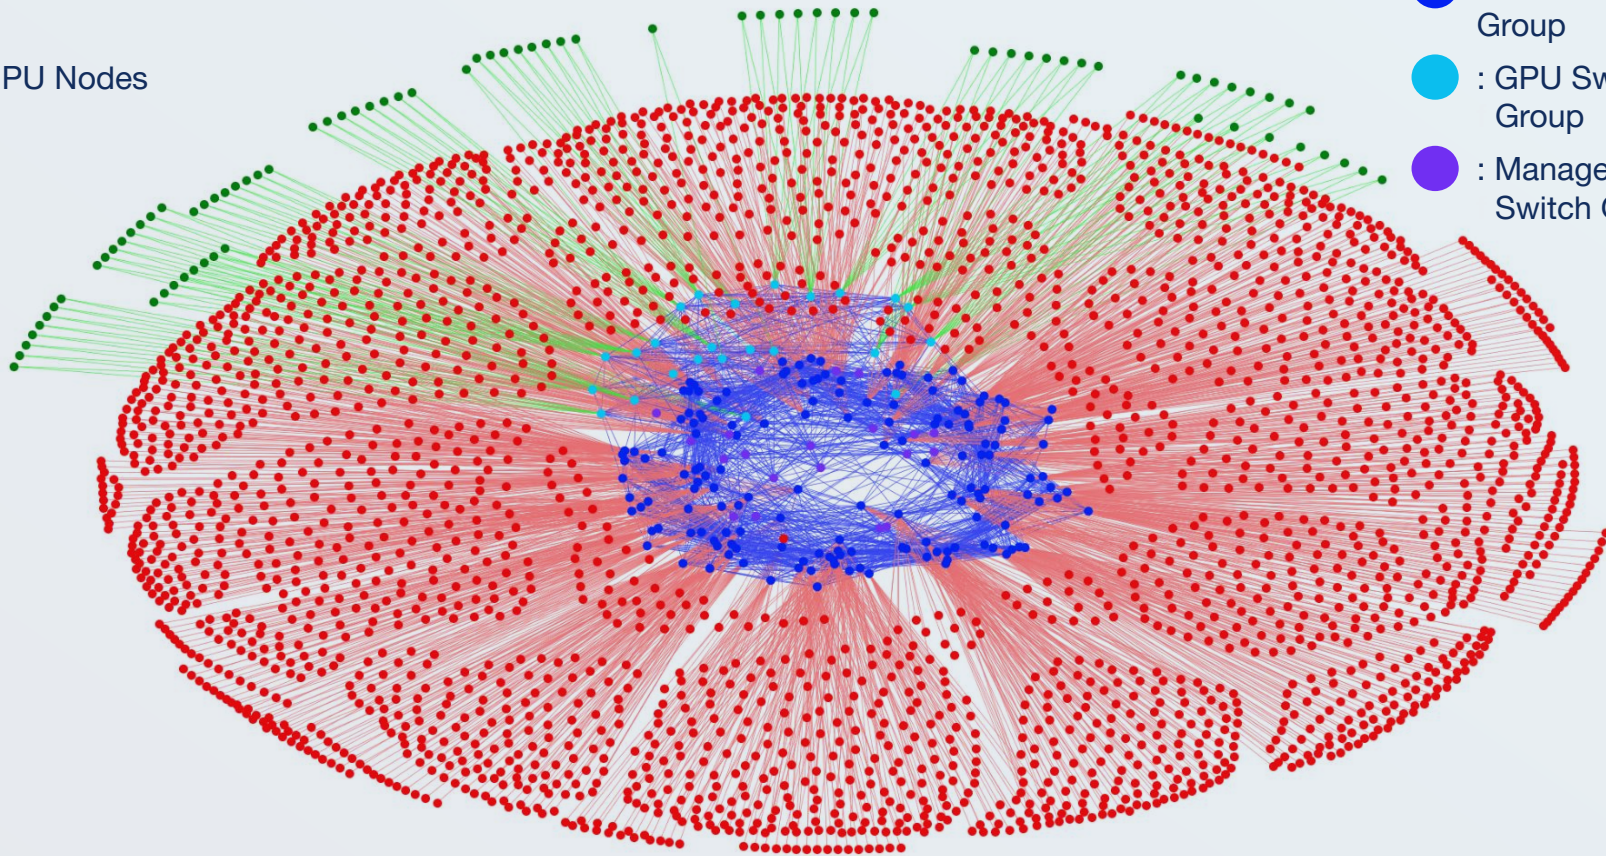
● : CPU Nodes


● : GPU Nodes

● : CPU Switch Group

● : GPU Switch Group

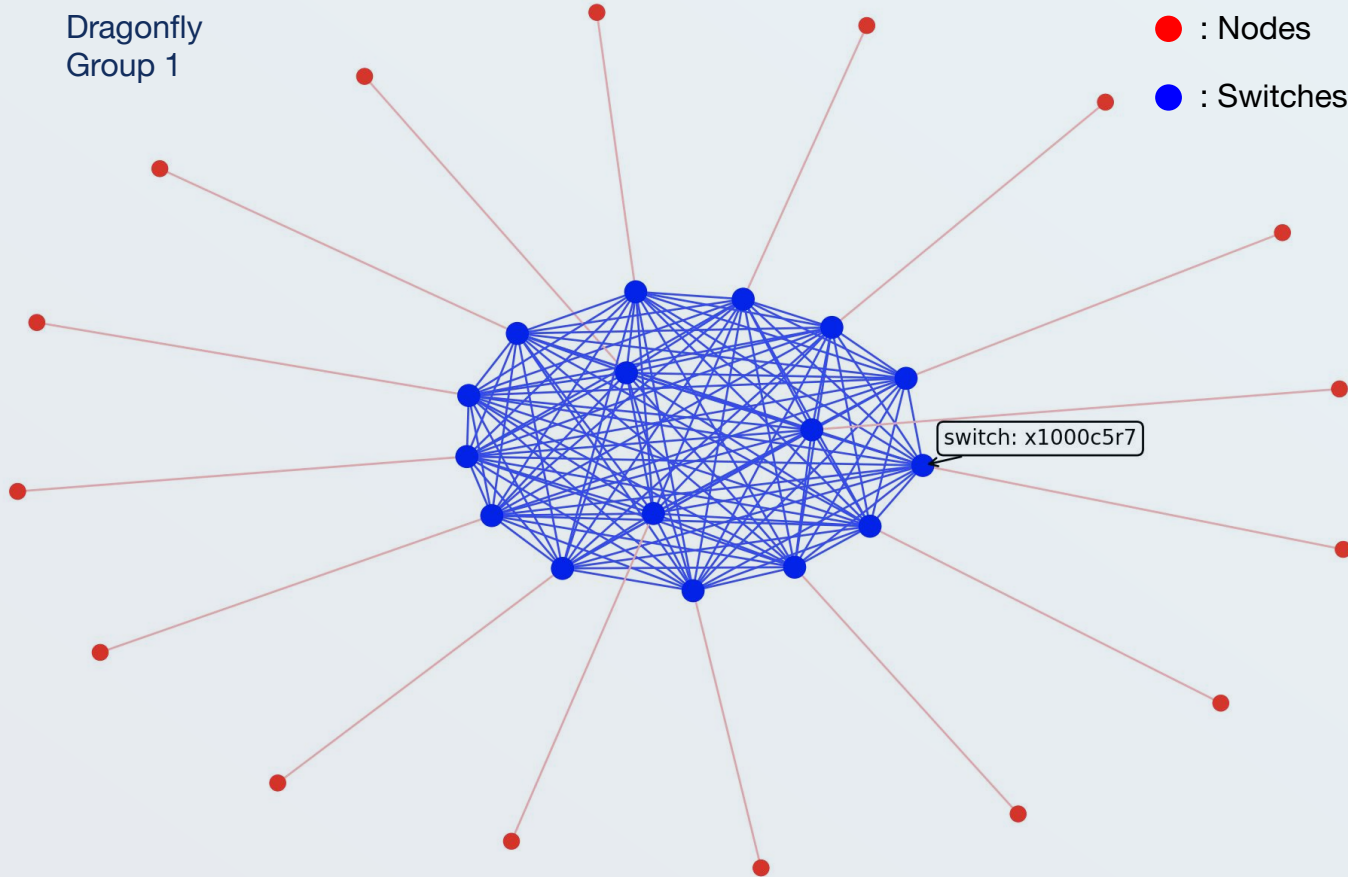
● : Management Switch Group

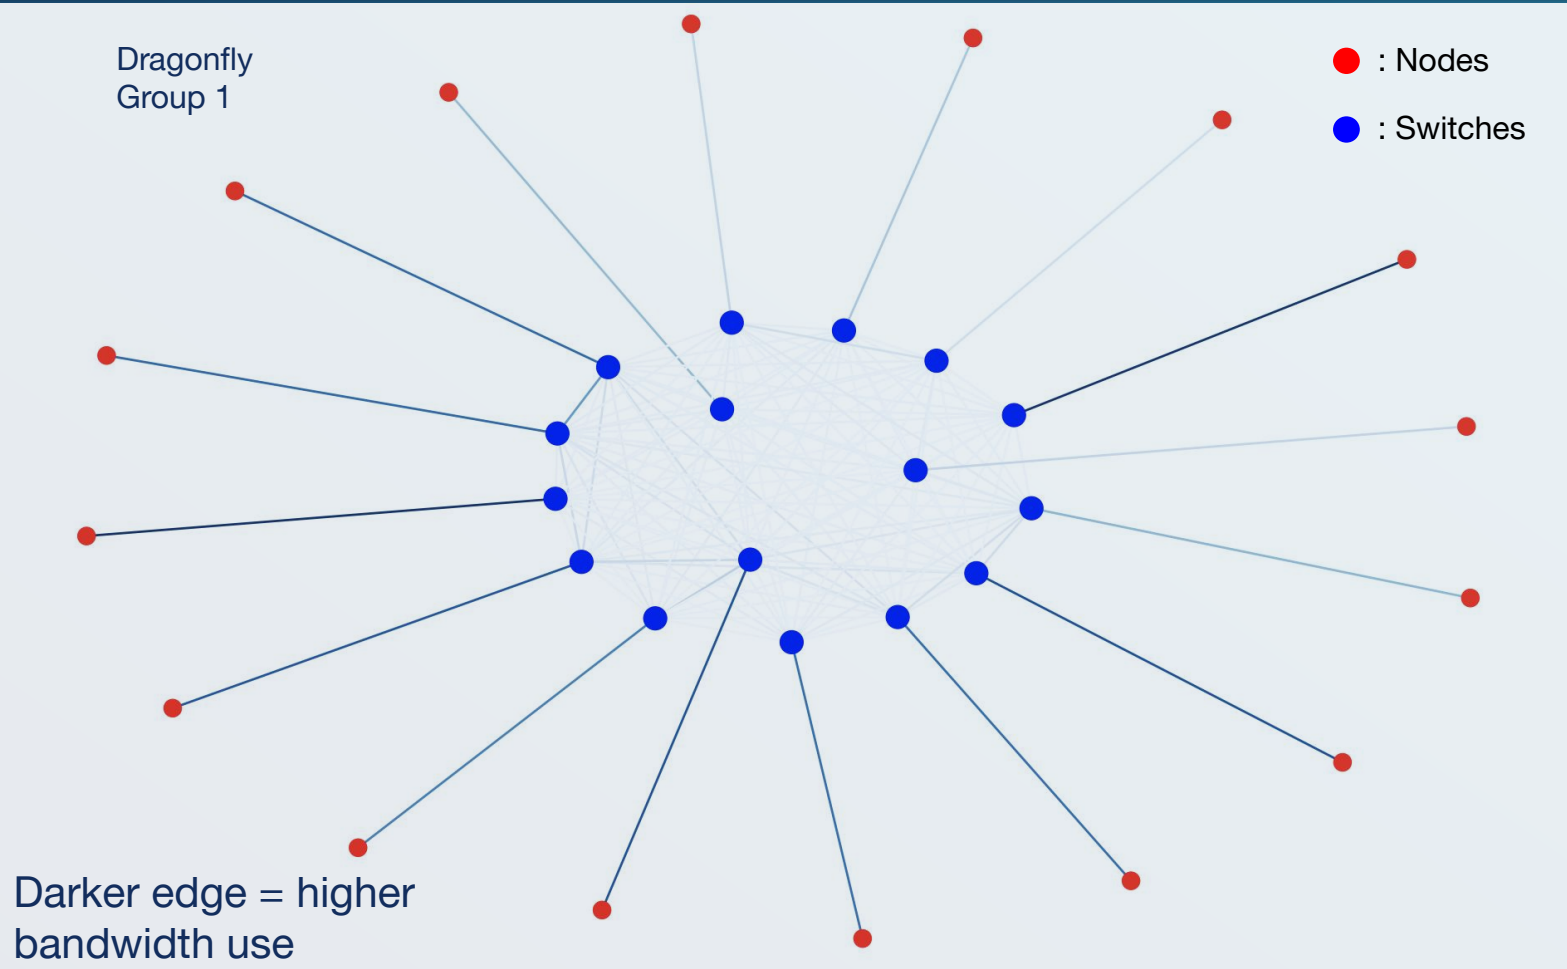




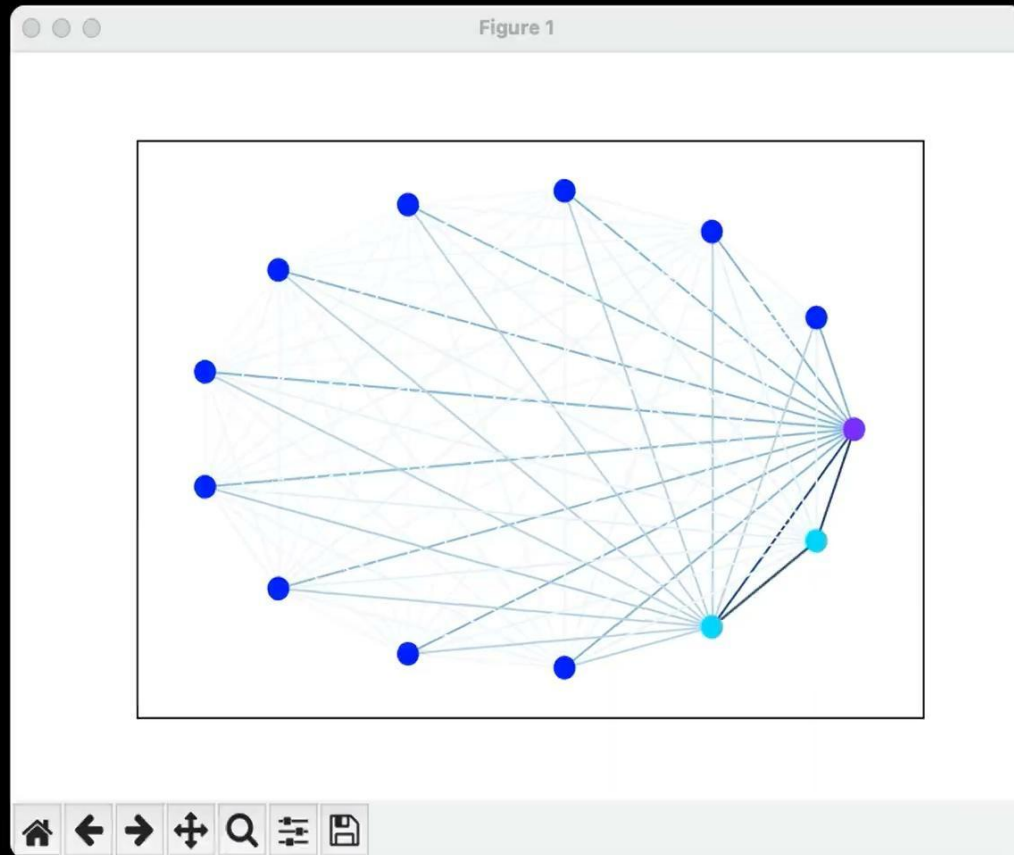
Dragonfly
Group 1

● : Nodes
● : Switches







Full Graph
Visualization,
Switches
Grouped



Overview ⓘ



Flow Interruption Counters ⓘ



Future Work

- More data analysis
 - There are a lot more counters to look at
- Time of day? Day of the week? Job-by-job basis?
- More control/communication over what is running on Derecho
- Timelapses over a specific time interval

Acknowledgements



Thank you!

My mentors **William Shanks** and **Storm Knight**

The entire HPCRD team

The SIParCS 2024 Coordinators:
Virginia Do, Jerry Cyccone, Eva Sosoo, and Jessica Wang

Everyone in my SIParCS cohort for such an enjoyable summer!

