Improving Data Center Visibility With Machine Learning

SIParCS 2024

Analiese Gonzalez
Mentors: Ben Matthews & Jenett Tillotson

July 30, 2024

This material is based upon work supported by the NSF National Center for Atmospheric Research, which is a major facility sponsored by the U.S. National Science Foundation under Cooperative Agreement No. 1852977.

Overview

1. Main Objective
2. Background
3. Data
4. Machine Learning
5. Future Work
6. Acknowledgements

Main Objective

Is there a way to effectively predict node failures using machine learning?

What insight can machine learning provide about our supercomputers?

What is a supercomputer?

- Very powerful
- Processes large amounts of data
- Casper - 121 nodes
- Derecho - 2570 nodes
    - Includes 3 CDUs (Cooling Distribution Unit)

What is a node?

- Individual computers
  - Each have its own “brain” (CPU or GPU)
  - Own RAM
  - Own Storage Space

- Casper has 121
- Derecho has 2570

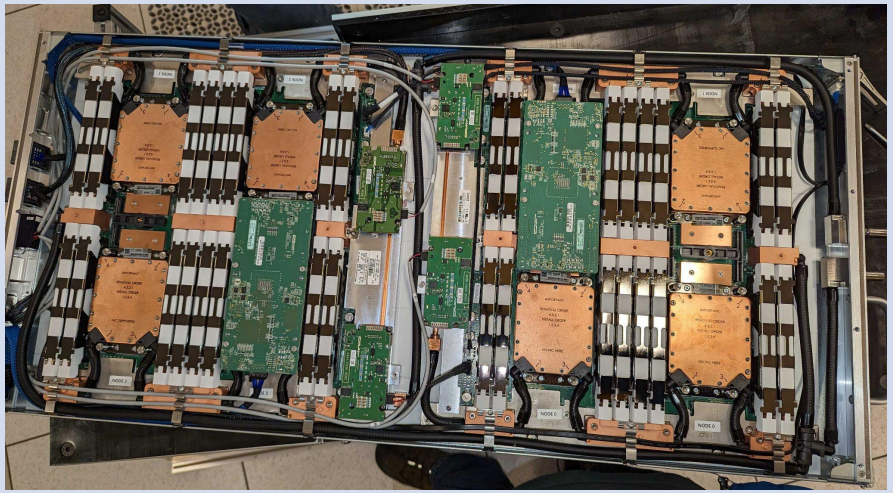

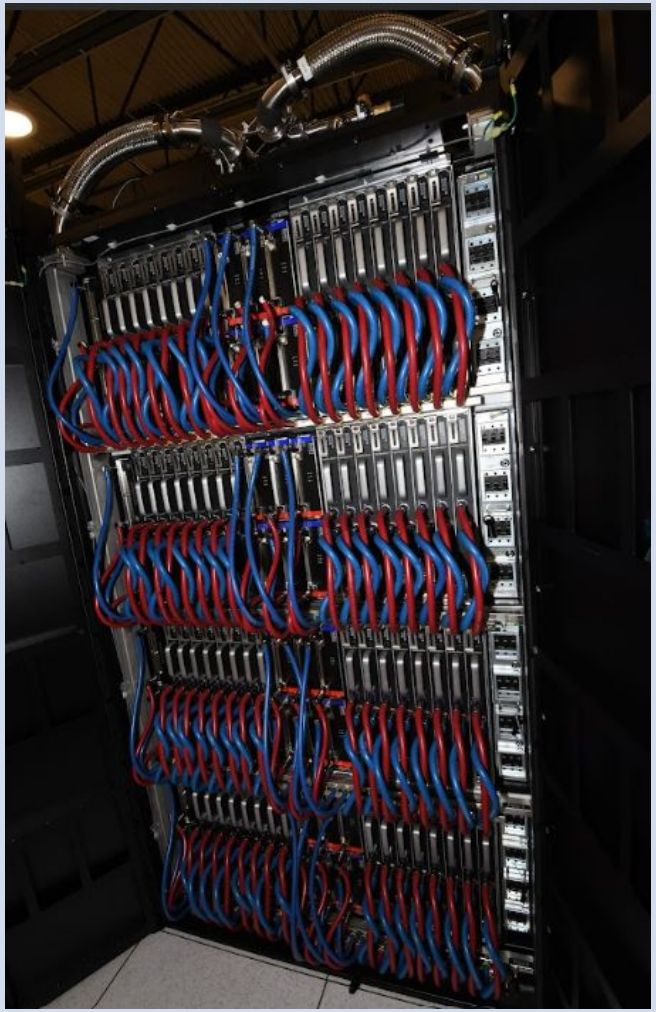

Glimpse of the Data

- Timescale Database
  - Time Series Data
  - 20 tables
  - Cheyenne, Casper, Gust, and Derecho
  - CDU data, node data

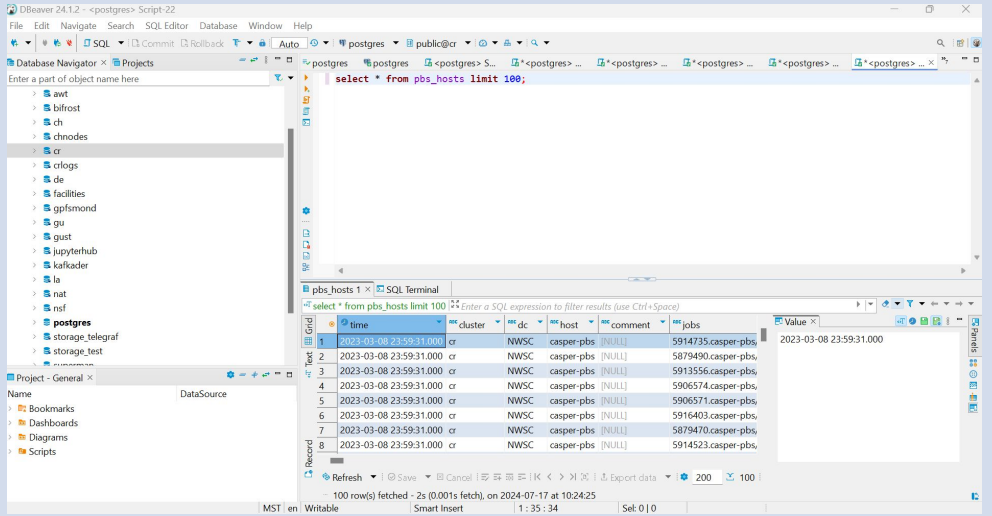

- Used SQL in DBeaver to query and export data
- Derecho CDU data & node data
- Casper per node data

Casper node data example

Grid | Text | Record

	time	cluster	dc	host	comment	jobs	last_state_change_time	last_used_time	license
1	2023-03-08 23:59:31.000	cr	NWSC	casper-pbs	[NULL]	5914735.casper-pbs/0, 5913255.casper-pbs/1, 58	1,678,298,704	1,678,316,193	I
2	2023-03-08 23:59:31.000	cr	NWSC	casper-pbs	[NULL]	5879490.casper-pbs/0, 5879648.casper-pbs/1, 59	1,678,239,928	1,678,319,571	I
3	2023-03-08 23:59:31.000	cr	NWSC	casper-pbs	[NULL]	5913556.casper-pbs/0, 5911637.casper-pbs/1, 59	1,678,151,801	1,678,306,083	I
4	2023-03-08 23:59:31.000	cr	NWSC	casper-pbs	[NULL]	5906574.casper-pbs/0, 5906575.casper-pbs/1, 59	1,678,165,977	1,678,319,571	I
5	2023-03-08 23:59:31.000	cr	NWSC	casper-pbs	[NULL]	5906571.casper-pbs/0, 5914499.casper-pbs/1, 59	1,678,314,797	1,678,314,796	I
6	2023-03-08 23:59:31.000	cr	NWSC	casper-pbs	[NULL]	5916403.casper-pbs/0, 5879489.casper-pbs/1, 59	1,678,201,671	1,678,317,481	I
7	2023-03-08 23:59:31.000	cr	NWSC	casper-pbs	[NULL]	5879470.casper-pbs/0, 5879640.casper-pbs/1, 58	1,678,184,803	1,678,318,748	I
8	2023-03-08 23:59:31.000	cr	NWSC	casper-pbs	[NULL]	5914523.casper-pbs/0, 5904807.casper-pbs/1, 58	1,678,227,255	1,678,319,872	I
9	2023-03-08 23:59:31.000	cr	NWSC	casper-pbs	[NULL]	5906544.casper-pbs/0, 5904977.casper-pbs/1, 59	1,678,215,685	1,678,314,534	I
10	2023-03-08 23:59:31.000	cr	NWSC	casper-pbs	[NULL]	5916464.casper-pbs/0, 5906524.casper-pbs/1, 59	1,678,286,243	1,678,319,569	I
11	2023-03-08 23:59:31.000	cr	NWSC	casper-pbs	[NULL]	5906386.casper-pbs/0, 5912216.casper-pbs/1, 59	1,678,179,896	1,678,315,607	I
12	2023-03-08 23:59:31.000	cr	NWSC	casper-pbs	[NULL]	5879496.casper-pbs/0, 5904974.casper-pbs/1, 59	1,678,174,142	1,678,319,571	I
13	2023-03-08 23:59:31.000	cr	NWSC	casper-pbs	[NULL]	5916488.casper-pbs/0, 5914716.casper-pbs/1, 59	1,678,226,021	1,678,318,346	I
14	2023-03-08 23:59:31.000	cr	NWSC	casper-pbs	[NULL]	5904284.casper-pbs/0, 5912011.casper-pbs/1, 59	1,678,302,095	1,678,316,196	I
15	2023-03-08 23:59:31.000	cr	NWSC	casper-pbs	[NULL]	5916431.casper-pbs/0, 5879494.casper-pbs/1, 59	1,677,816,920	1,678,318,274	I
16	2023-03-08 23:59:31.000	cr	NWSC	casper-pbs	[NULL]	5894303.casper-pbs/0, 5914053.casper-pbs/1, 58	1,678,313,511	1,678,313,508	I
17	2023-03-08 23:59:31.000	cr	NWSC	casper-pbs	[NULL]	5914061.casper-pbs/0, 5894242.casper-pbs/1, 58	1,677,516,397	1,678,315,529	I
18	2023-03-08 23:59:31.000	cr	NWSC	casper-pbs	[NULL]	5879578.casper-pbs/0, 5916407.casper-pbs/1, 59	1,678,298,237	1,678,318,307	I
19	2023-03-08 23:59:31.000	cr	NWSC	casper-pbs	[NULL]	5914265.casper-pbs/0, 5906356.casper-pbs/1, 59	1,678,311,770	1,678,317,307	I
20	2023-03-08 23:59:31.000	cr	NWSC	casper-pbs	[NULL]	5906383.casper-pbs/0, 5904952.casper-pbs/1, 59	1,678,226,393	1,678,318,339	I
21	2023-03-08 23:59:31.000	cr	NWSC	casper-pbs	[NULL]	[NULL]	1,677,881,426	1,678,280,347	I
22	2023-03-08 23:59:31.000	cr	NWSC	casper-pbs	[NULL]	5904953.casper-pbs/0, 5914215.casper-pbs/1, 59	1,677,861,356	1,678,315,339	I
23	2023-03-08 23:59:31.000	cr	NWSC	casper-pbs	[NULL]	5911675.casper-pbs/0, 5914228.casper-pbs/1, 59	1,678,226,192	1,678,319,951	I

Derecho CDU data example

	time	chassis	host	rack	CoolingSubsystem0ACInputFrequency	CoolingSubsystem0ACInputLine1Current	CoolingSubsystem0ACInputLi
1	2023-08-28 23:58:04.000	c1	x1100c1	x1100	60	10.1	
2	2023-08-28 23:58:04.000	c1	x1000c1	x1000	60	16	
3	2023-08-28 23:56:02.000	c1	x1007c1	x1007	60	16.5	
4	2023-08-28 23:56:02.000	c1	x1100c1	x1100	60	10.1	
5	2023-08-28 23:56:02.000	c1	x1000c1	x1000	60	16.1	
6	2023-08-28 23:54:06.000	c1	x1100c1	x1100	60	10	
7	2023-08-28 23:52:03.000	c1	x1007c1	x1007	60	16.6	
8	2023-08-28 23:52:03.000	c1	x1000c1	x1000	60	16	
9	2023-08-28 23:52:02.000	c1	x1100c1	x1100	60	10.1	
10	2023-08-28 23:50:04.000	c1	x1000c1	x1000	60	15.9	
11	2023-08-28 23:50:03.000	c1	x1100c1	x1100	60	10	
12	2023-08-28 23:50:03.000	c1	x1007c1	x1007	60	16.5	
13	2023-08-28 23:48:04.000	c1	x1100c1	x1100	60	10.1	
14	2023-08-28 23:48:04.000	c1	x1000c1	x1000	60	16	
15	2023-08-28 23:48:04.000	c1	x1007c1	x1007	60	16.5	
16	2023-08-28 23:46:05.000	c1	x1100c1	x1100	60	10	
17	2023-08-28 23:46:04.000	c1	x1007c1	x1007	60	16.5	
18	2023-08-28 23:46:04.000	c1	x1000c1	x1000	60	16	
19	2023-08-28 23:44:07.000	c1	x1007c1	x1007	60	16.6	
20	2023-08-28 23:44:07.000	c1	x1100c1	x1100	60	9.9	
21	2023-08-28 23:44:06.000	c1	x1000c1	x1000	60	16	
22	2023-08-28 23:42:06.000	c1	x1007c1	x1007	60	16.8	
23	2023-08-28 23:42:06.000	c1	x1000c1	x1000	60	16.1	

K-Means Clustering Model Algorithm

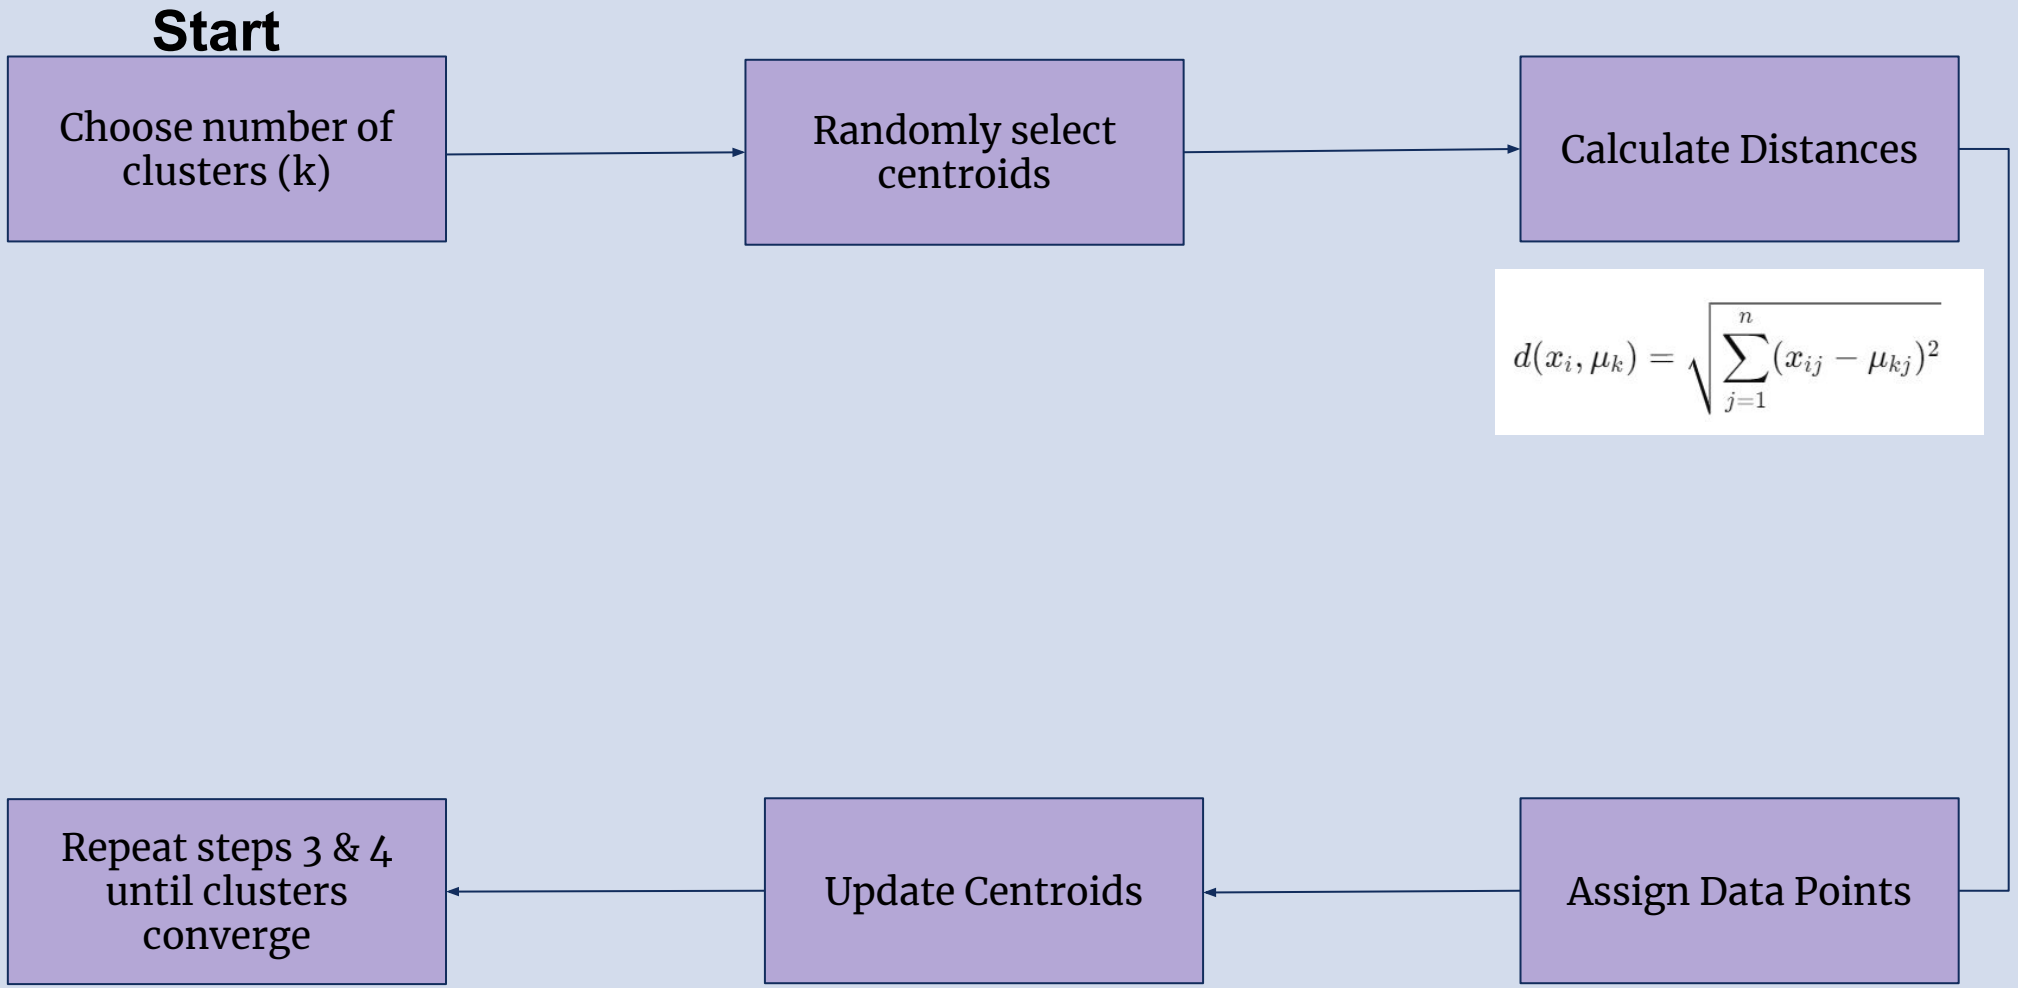

K-Means Graphs-Derecho

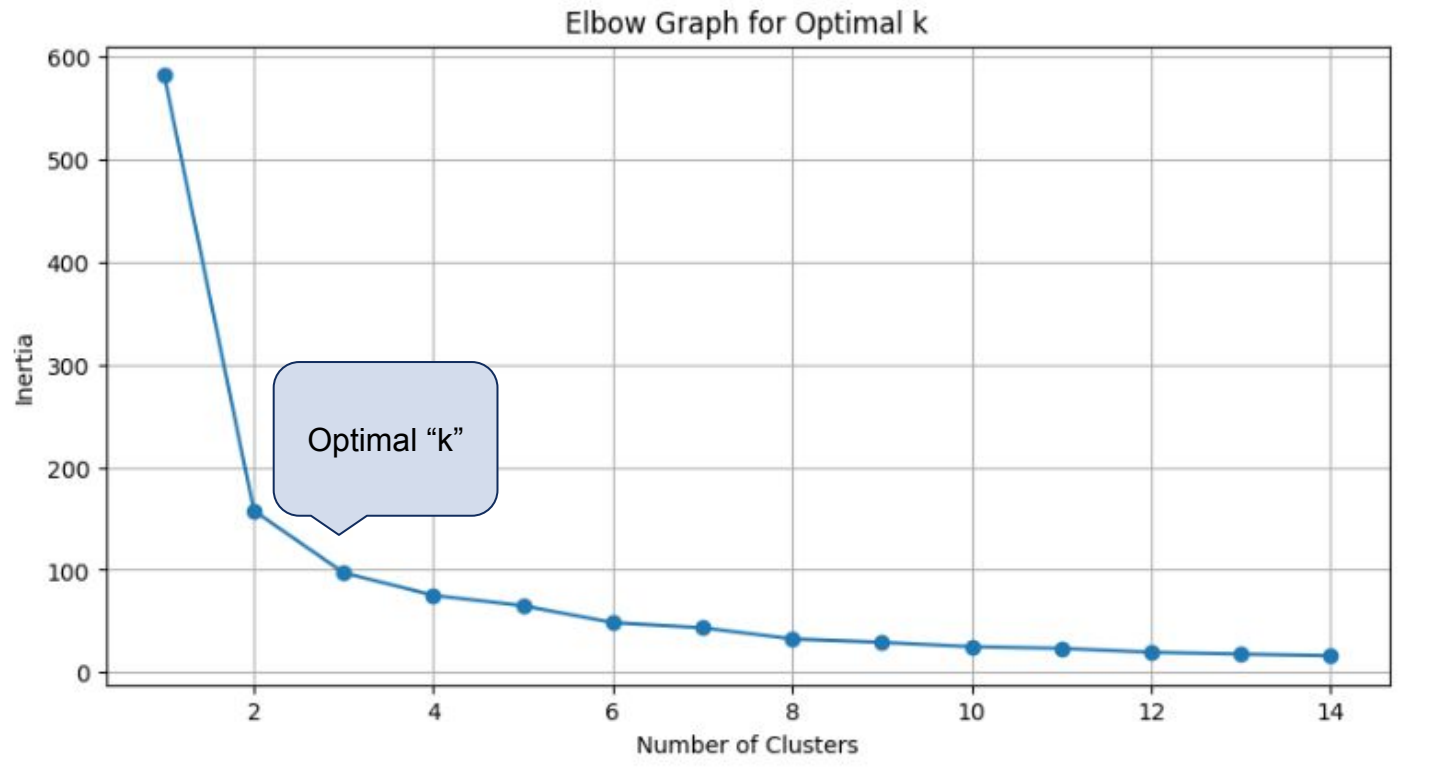

- Used to find optimal “k” or clusters
- Elbow “bends” at around 3

K-Means Graphs-Derecho

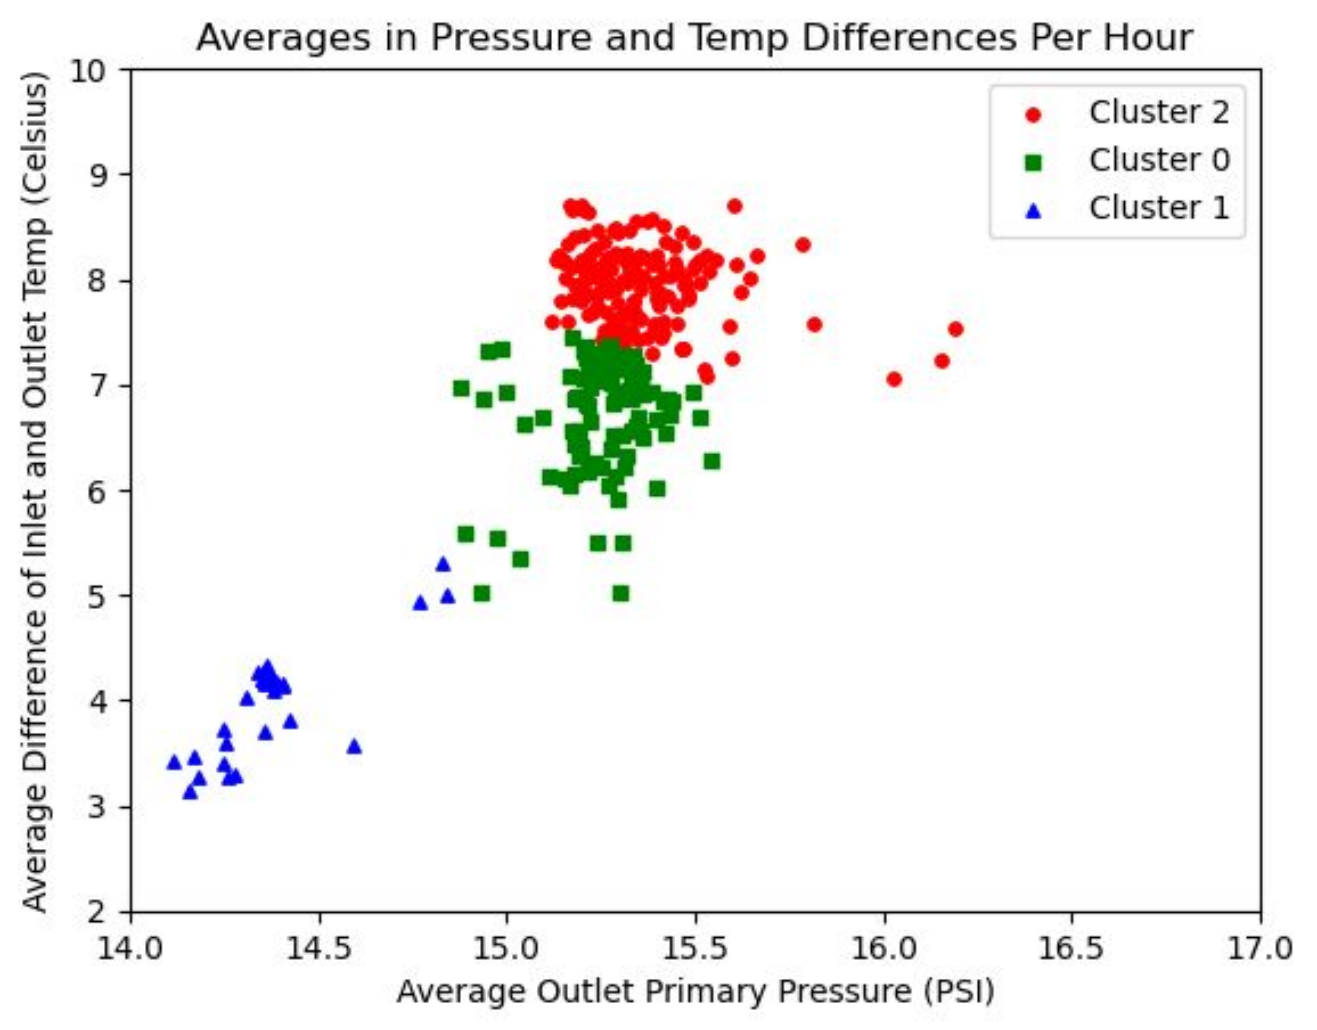

- x1100 reporting 0 for data
- Discovered broken sensor
- 291 points
- Outlet Primary Pressure - pressure of the fluid when it exits the cdu and enters the primary cooling loop
- Inlet and Outlet Temp - Temp of the liquid as it enters and exits system

K-Means Graphs-Derecho

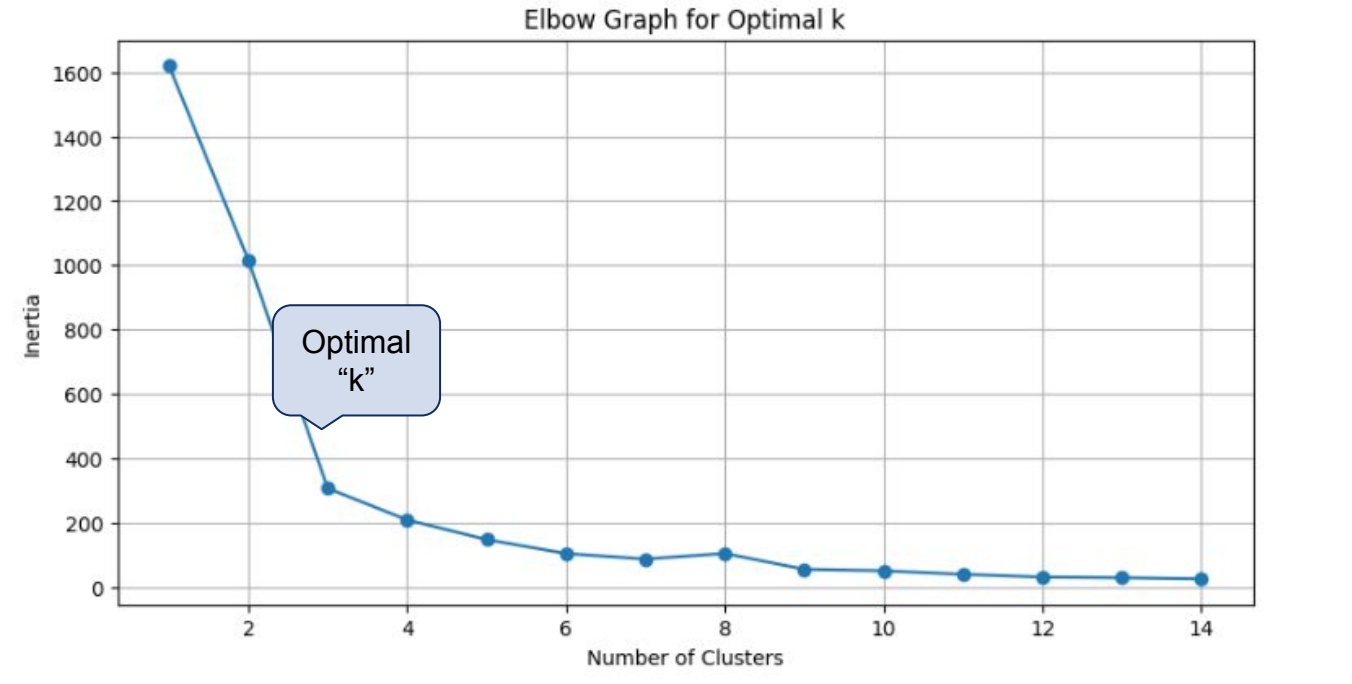

- Chose 3 clusters
- Elbow “bends” at around 3

K-Means Graphs-Derecho

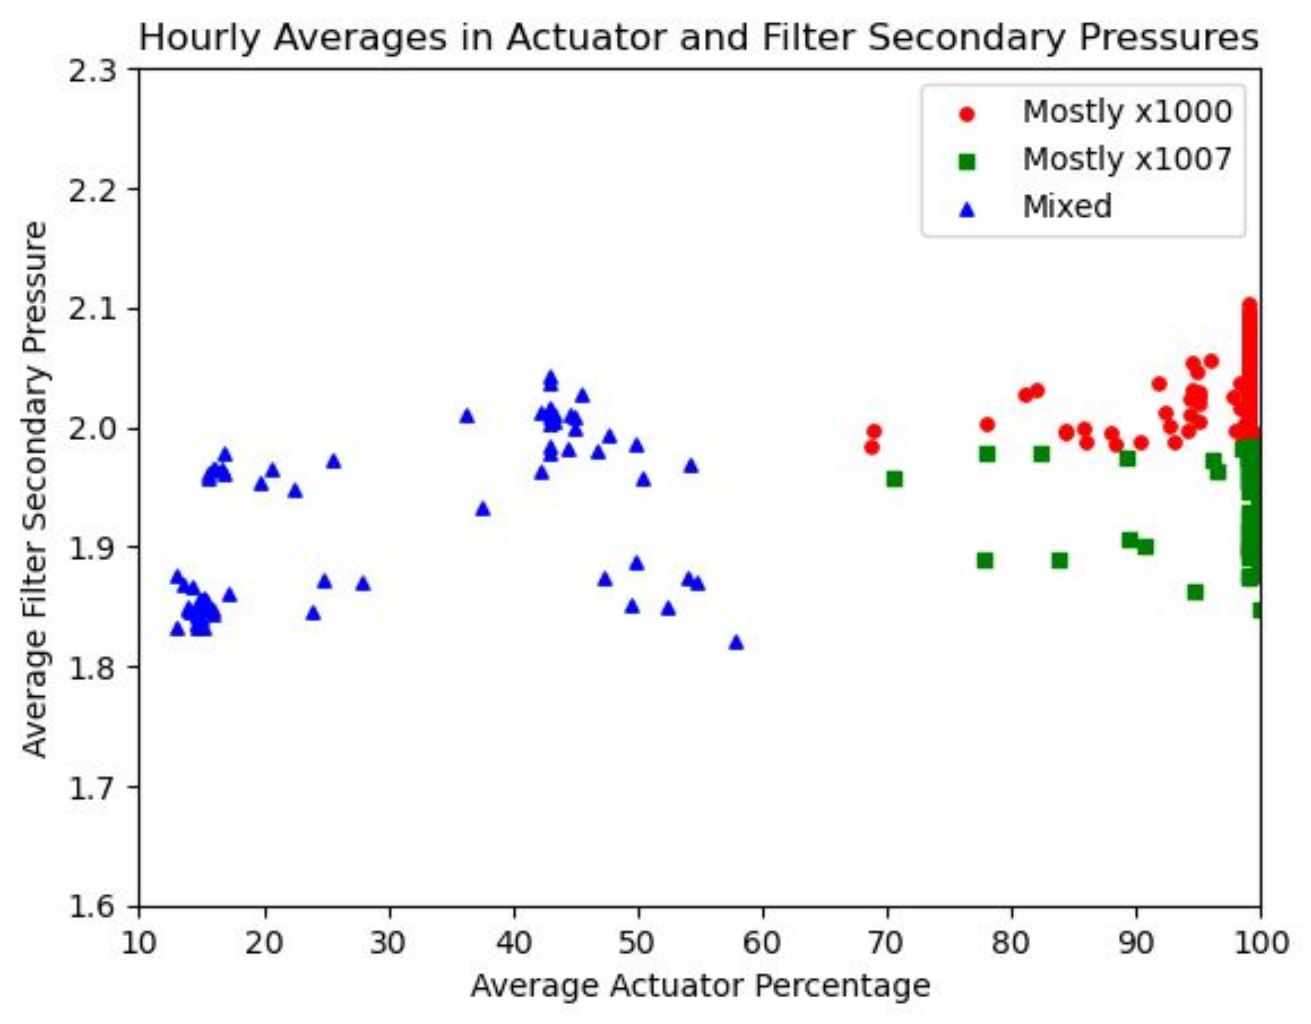

- Excluded x1100
- Actuator - allows water to flow in
- Filter Secondary Pressure - Pressure of cooling fluid
- One CDU rack was misconfigured

K-Means Graphs-Casper

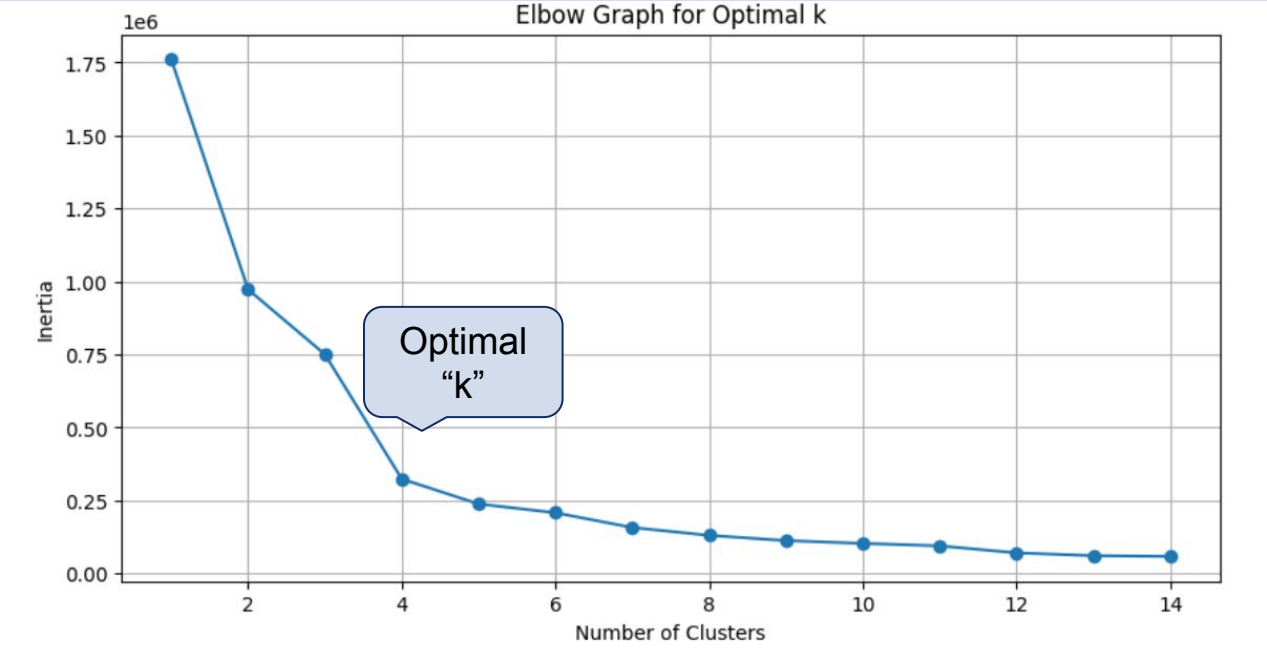

- Chose 4 as optimal “k”
- Could possibly have chosen 2

K-Means Graphs-Casper

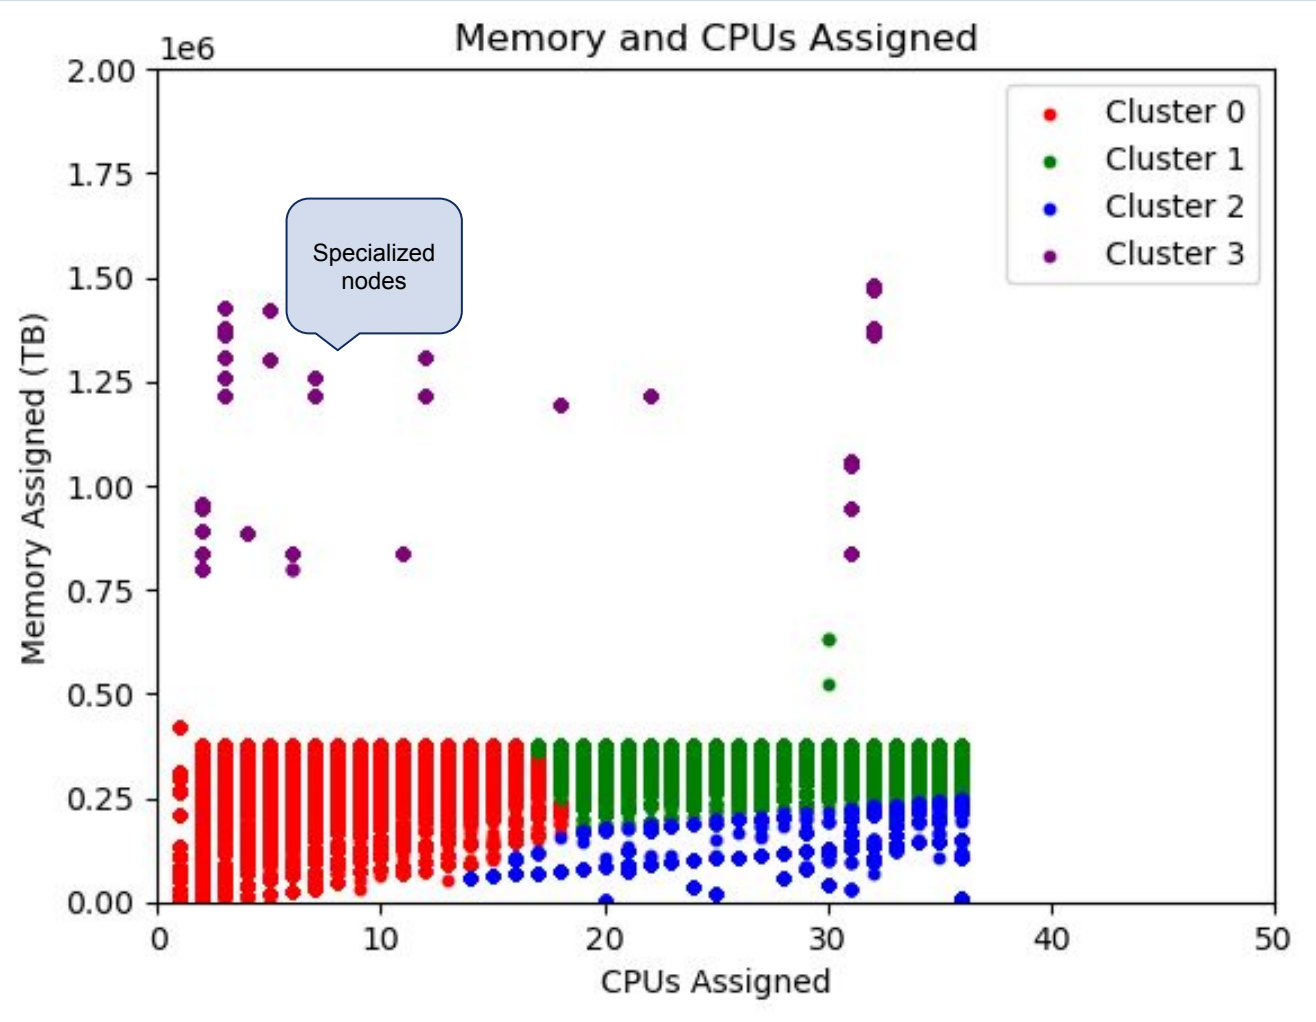

- Sept 6-15 2023
- Most ask for max amount of memory
- Most nodes are HPC nodes
- Purple (diamond) cluster is specialized nodes

Neural Networks

- Inspired by the human brain
- Neurons receive inputs
- Learns by changing the weights between neurons
- Efficient with anomaly detection, complex data, can be trained on historical data

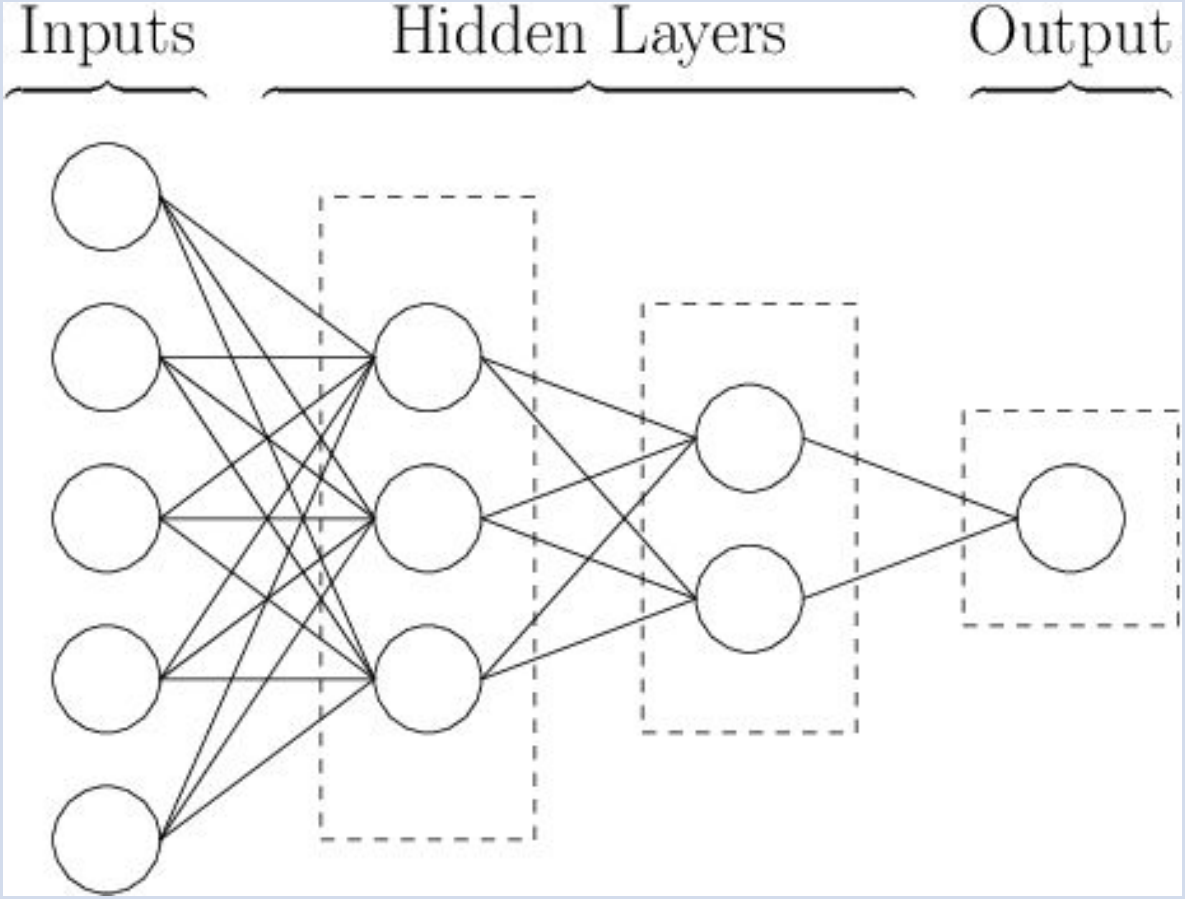

https://www.researchgate.net/figure/Simple-neural-network-diagram_fig1_332158639

Neural Networks

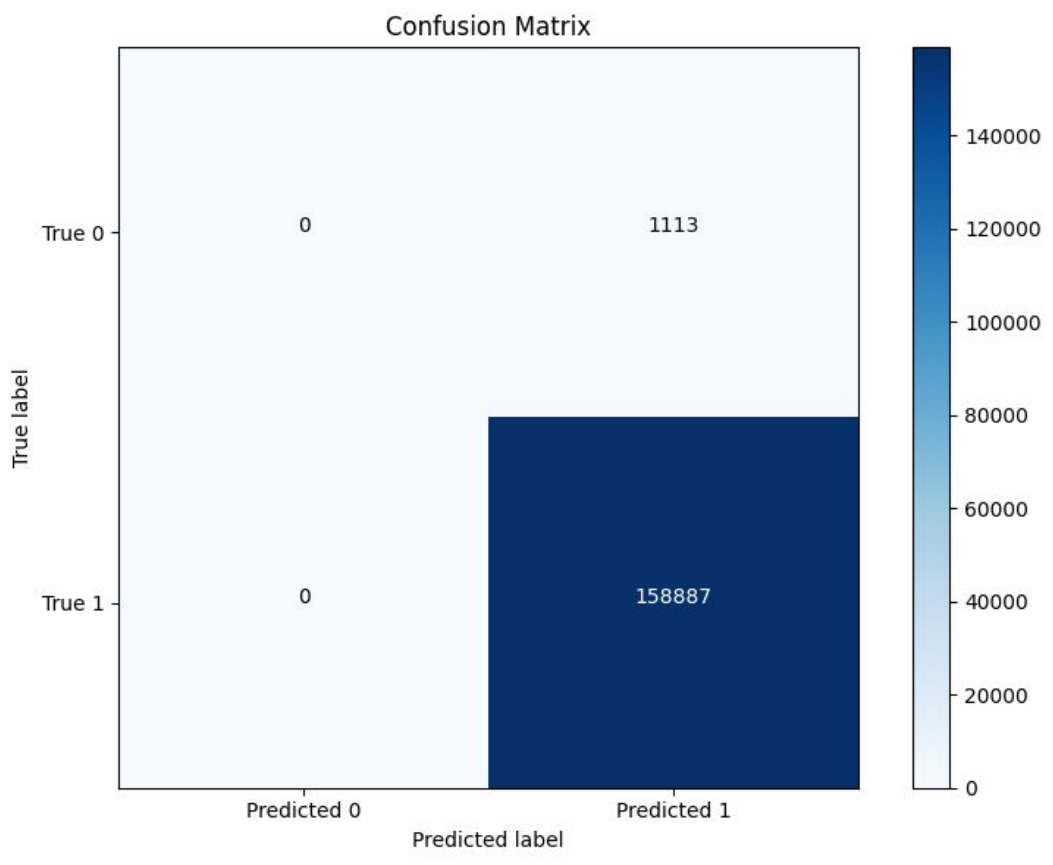

- First model
- Derecho CDU data
- Model never predicted 0

Metric	Value
Mean Absolute Error	0.012361268613487481
Mean Squared Error	0.006910000116743188
Root Mean Squared Error	0.08312641046468437
R-squared	-0.000309586524963378 9

Neural Networks

Casper Model:

Parameter	Value
Window size	50.0
Batch size	16.0
Epochs	50.0
Mean Absolute Error	0.008267107579070884
Mean Squared Error	0.008267107579070884
Root Mean Squared Error	0.09092363597586099
R-squared	0.523836128720933

Derecho Model:

Parameter	Value
Window size	50.0
Batch size	16.0
Epochs	50.0
Mean Absolute Error	0.003593501030048069
Mean Squared Error	0.003593501030048069
Root Mean Squared Error	0.05994581745249679
R-squared	0.8053383024007601

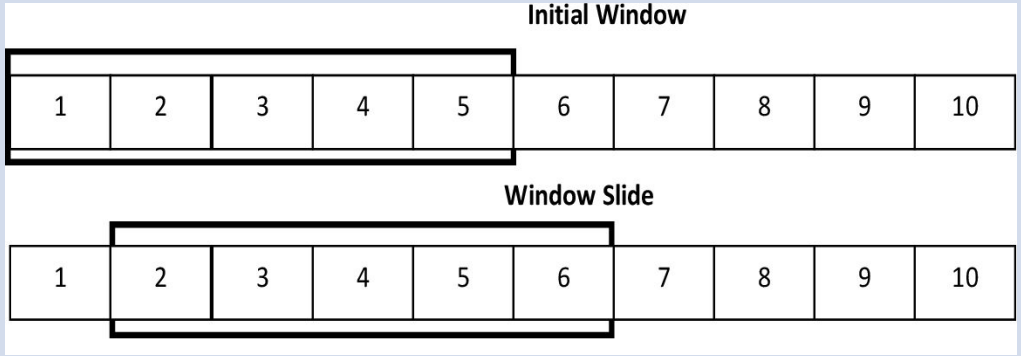

https://www.researchgate.net/figure/Simple-neural-network-diagram_fig1_332158639

Heat Maps

- last_state_change_time first for both machines
  - May be resolving nodes too quickly
- resources_assigned towards bottom for Casper but top for Derecho
  - Possibly because Casper is more heterogeneous

Casper

Feature	Importance
last_state_change_time	20.50%
resources_available.mem	19.18%
resources_available.ncpus	18.73%
last_used_time	16.28%
resources_available.vmem	8.72%
resources_assigned.mem	7.60%
resources_assigned.ncpus	5.43%
resources_available.ngpus	3.56%

Derecho

Feature	Importance
last_state_change_time	43.23%
last_used_time	30.23%
resources_assigned.ncpus	9.62%
resources_assigned_mem	9.11%
resources_avail_mem	2.76%
resources_avaiable_vmem	2.43%
resources_available.ncpus	1.65%
resources_available.ngpus	0.97%

Heat Maps

Casper:

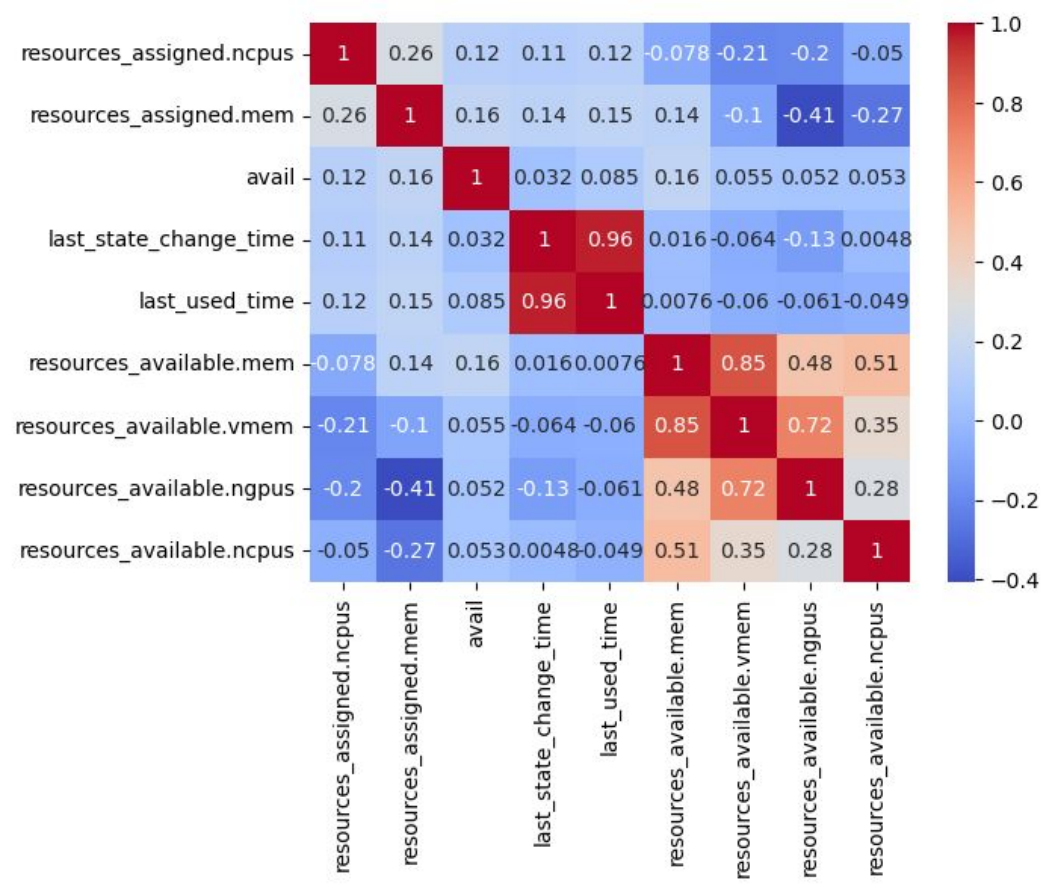

Derecho:

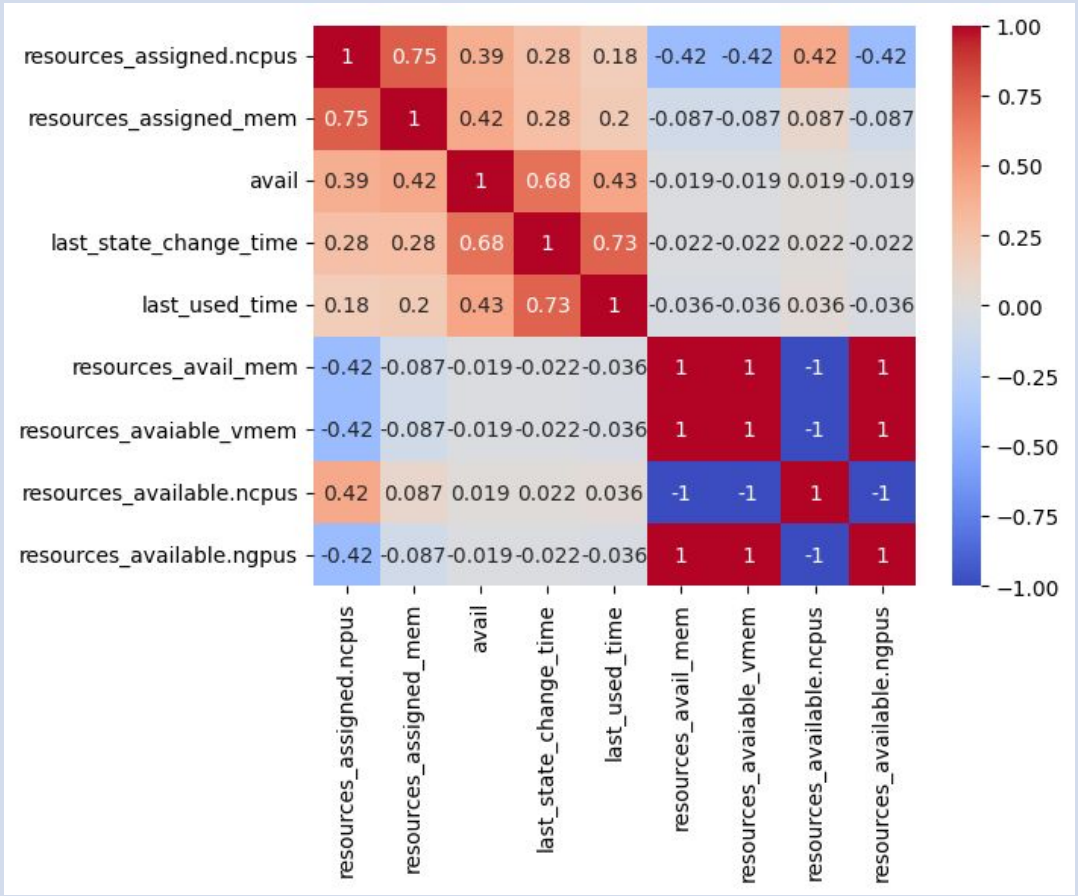

Neural Networks

Casper Model:

Parameter	Value
Window size	50.0
Batch size	16.0
Epochs	50.0
Mean Absolute Error	0.006460301480735768
Mean Squared Error	0.006460301480735768
Root Mean Squared Error	0.08037600065153633
R-squared	0.6207334316575936

Derecho Model:

Parameter	Value
Window size	50.0
Batch size	16.0
Epochs	50.0
Mean Absolute Error	0.003933516897455215
Mean Squared Error	0.003933516897455215
Root Mean Squared Error	0.06271775583879907
R-squared	0.7869194775815369

Confusion Matrices

Casper Model:

Derecho Model:

Future Work

- Use more heat maps
- Group nodes by specific type
- Create my own database for indexing
  - Combining data from other tables was challenging because of not matching times

select * from pbs_no... where tim... Enter a SQL expression to filter re...

	time	cluster	dc	host
1	2024-07-23 21:14:01.000	de	NWSC	desched1
2	2024-07-23 21:14:01.000	de	NWSC	desched1
3	2024-07-23 21:14:01.000	de	NWSC	desched1
4	2024-07-23 21:14:01.000	de	NWSC	desched1
5	2024-07-23 21:14:01.000	de	NWSC	desched1
6	2024-07-23 21:14:01.000	de	NWSC	desched1
7	2024-07-23 21:14:01.000	de	NWSC	desched1
8	2024-07-23 21:14:01.000	de	NWSC	desched1
9	2024-07-23 21:14:01.000	de	NWSC	desched1
10	2024-07-23 21:14:01.000	de	NWSC	desched1

select * from cdu where time > '20... Enter a SQL expressior...

	time	chassis	host
1	2024-07-23 21:14:02.000	c1	x1000c1
2	2024-07-23 21:14:02.000	c1	x1100c1
3	2024-07-23 21:14:02.000	c1	x1007c1
4	2024-07-23 21:13:02.000	c1	x1000c1
5	2024-07-23 21:13:02.000	c1	x1100c1
6	2024-07-23 21:13:02.000	c1	x1007c1
7	2024-07-23 21:12:02.000	c1	x1000c1
8	2024-07-23 21:12:02.000	c1	x1007c1
9	2024-07-23 21:12:02.000	c1	x1100c1
10	2024-07-23 21:11:02.000	c1	x1100c1

Thank you!

- My mentors Ben Matthews and Jenett Tillotson and the rest of the HPCRD team
- Virginia Do, Jerry Cyccone, Jessica Wang, Eva Sosoo, Ben Fellman and the rest of the CISL team
- All of the SIParCS interns for a fun summer!