Carbon Aware Software Deployment in Cloud

CS 7387 Independent Research Presentation

Dayuan Chen

Introduction

Overview

- Big tech companies have huge electricity usage in their data centers, which create greenhouse gas emissions from local grids. By analyzing the potential for using more renewable energy and shifting workloads, this project aims to reduce these emissions and promote sustainability in the tech industry.
- The carbon emission of various computing services increases rapidly but research in carbon-aware computing is still in its early stage.
- The research question here is how can carbon emissions in datacenter be reduced in a practical and cost-effective manner, while also promoting sustainability in the tech industry?

Introduction

Goal

- The main objective of this study is to explore the potential use of the Marginal Operating Emissions Rate (MOER) dataset as a means of identifying opportunities for reducing carbon emissions.
- Additionally, the study provides a comprehensive analysis of the regional and temporal patterns of MOER to gain a deeper understanding of the dataset's potential applications in carbon-aware computing.

Related Work
The "real" carbon-aware

- **Traditional approaches** to carbon-aware computing are primarily focused on reducing energy consumption, rather than reducing carbon emissions. For example using machine learning models to control the physical environment of data centers and scheduling air conditioners, or by selecting programming languages that use less energy.

- **Examples of carbon-aware** computing that aim to reduce emissions include Microsoft's carbon-aware Windows update and Apple's carbon-aware iPhone charging practices.

- **In this project**, we expand upon existing research by analyzing MOER patterns in multiple regions and exploring new approaches to reducing greenhouse gas emissions in the field of carbon-aware computing. The goal is to deepen understanding of MOER patterns and raise awareness of carbon-aware computing, inspiring action towards reducing greenhouse gas emissions.

Content

- MOER analysis
- Regional analysis
- Azure VM Dataset
- Time-shifting on Azure VM traces

MOER analysis

- The MOER dataset consists of a continuous measurement record taken every 5 minutes, covering the period from 2017 to the end of 2022 for 151 electric grid regions.

MOER analysis

MOER monthly/weekly average in different years

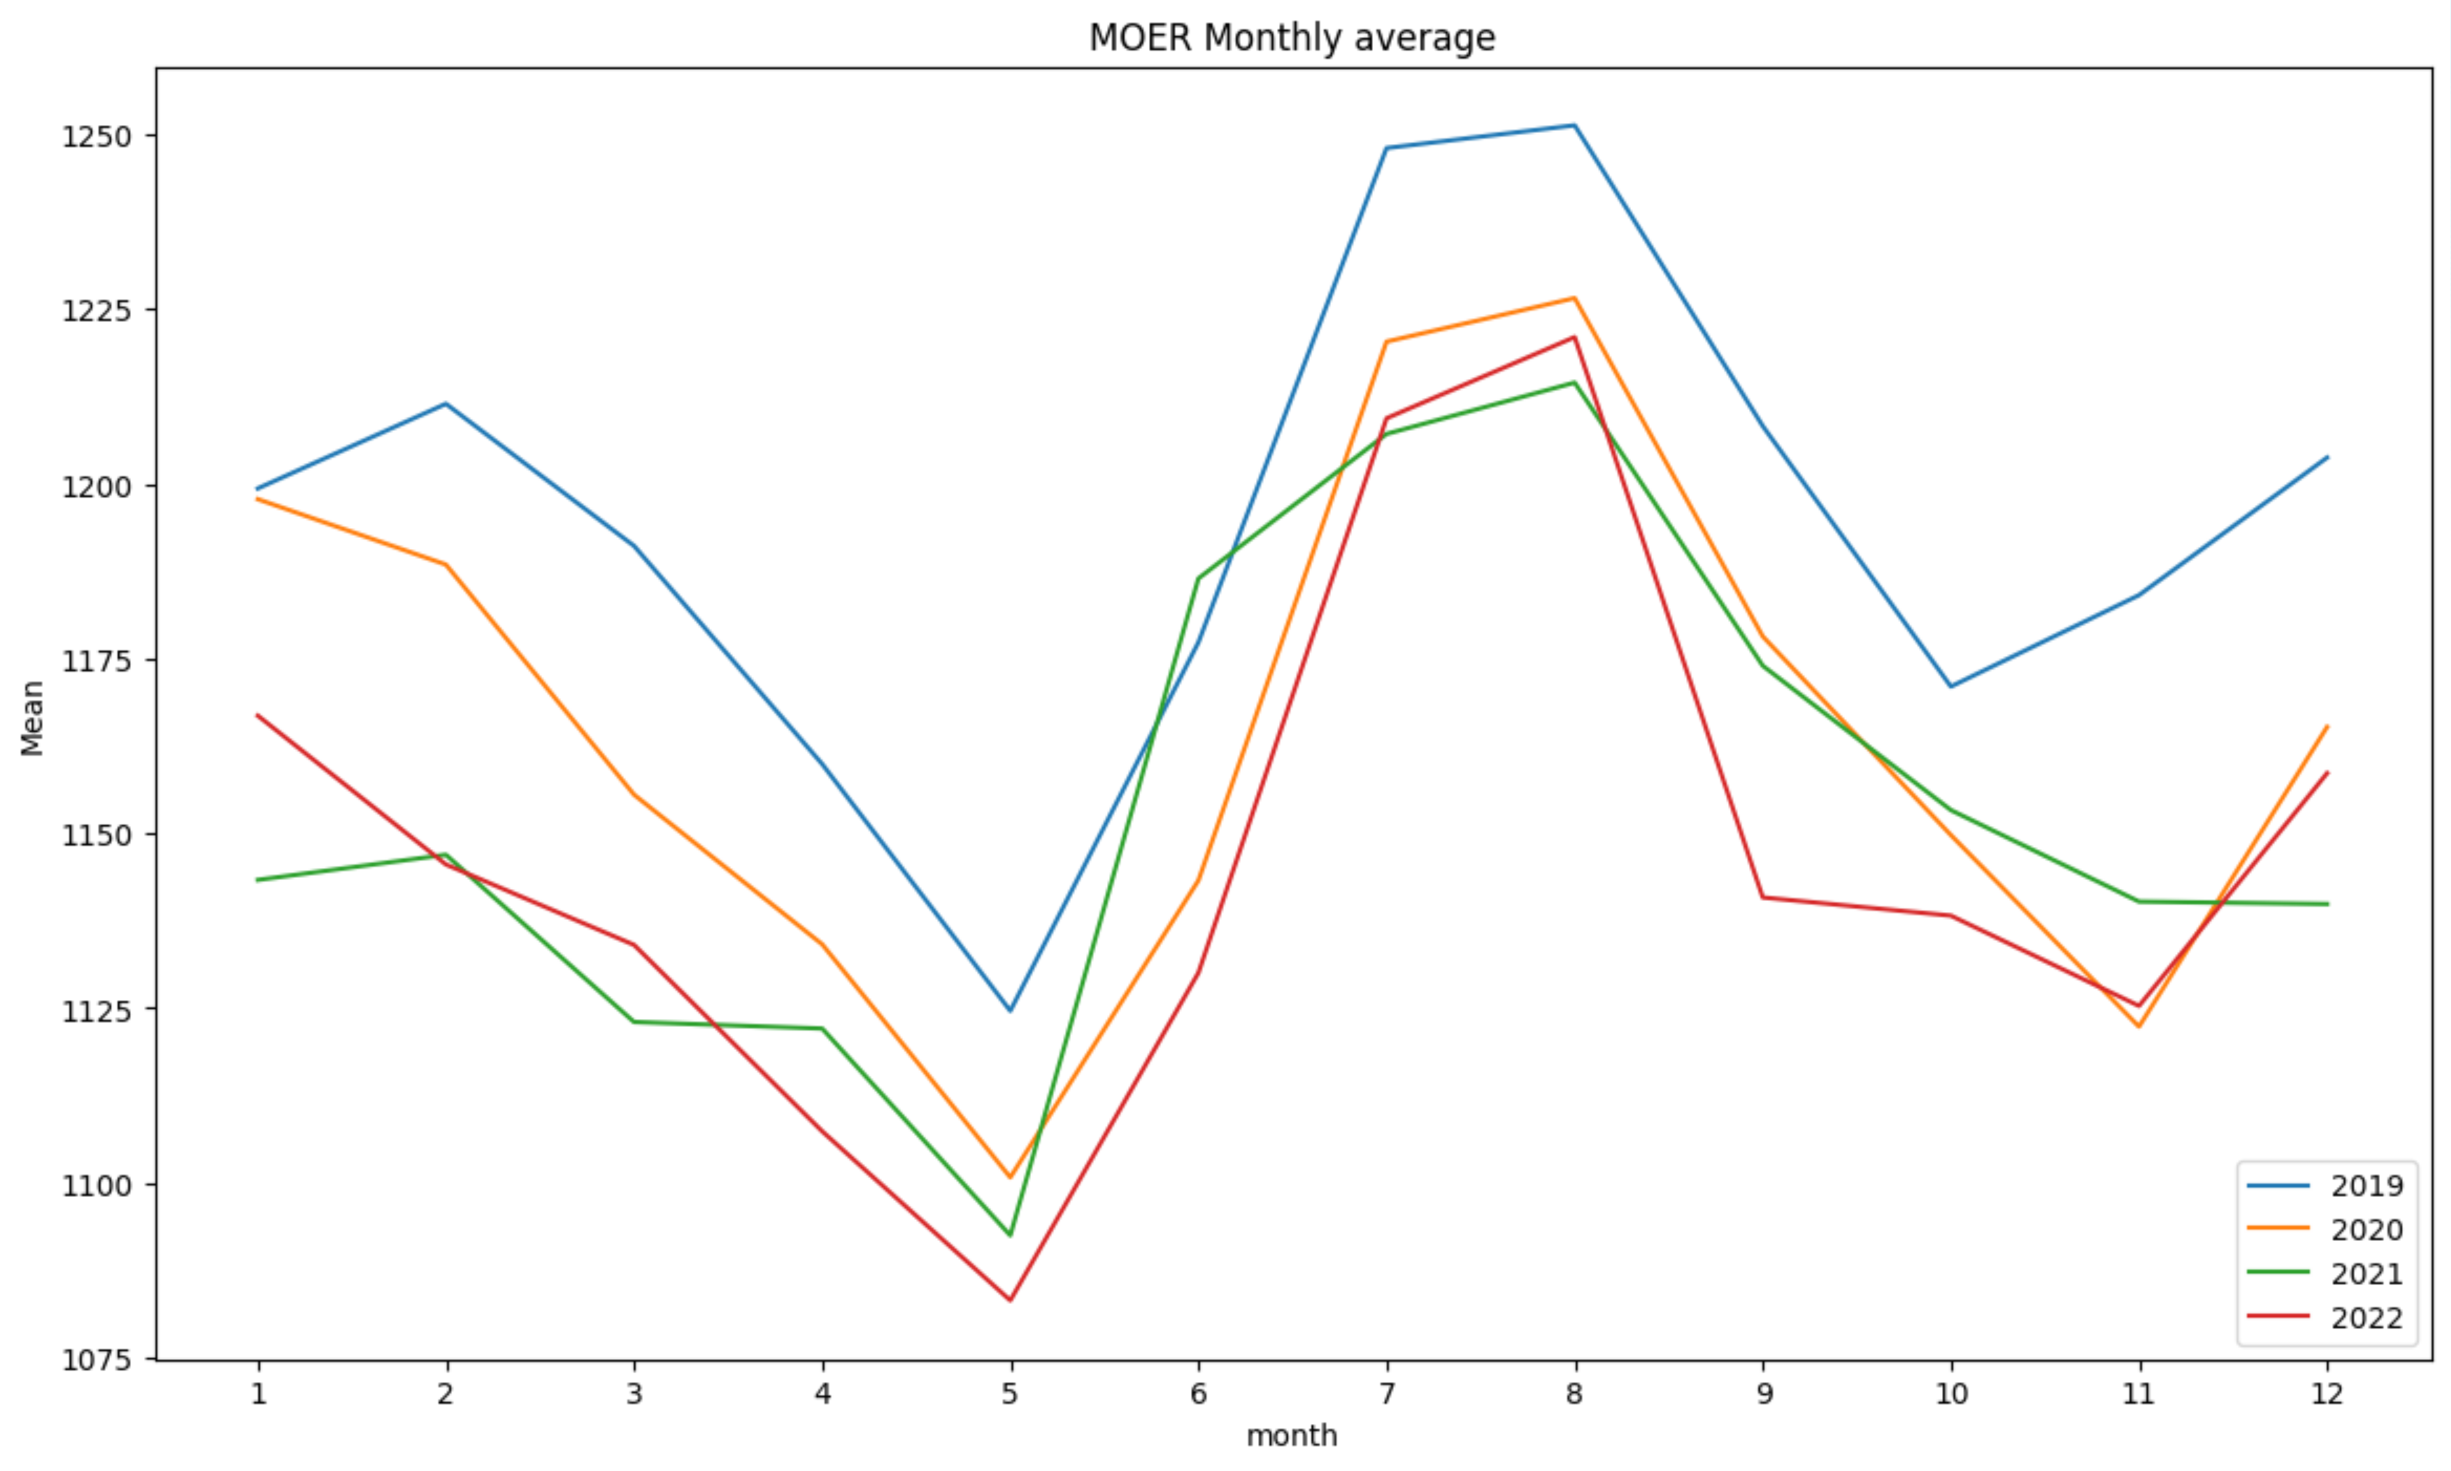

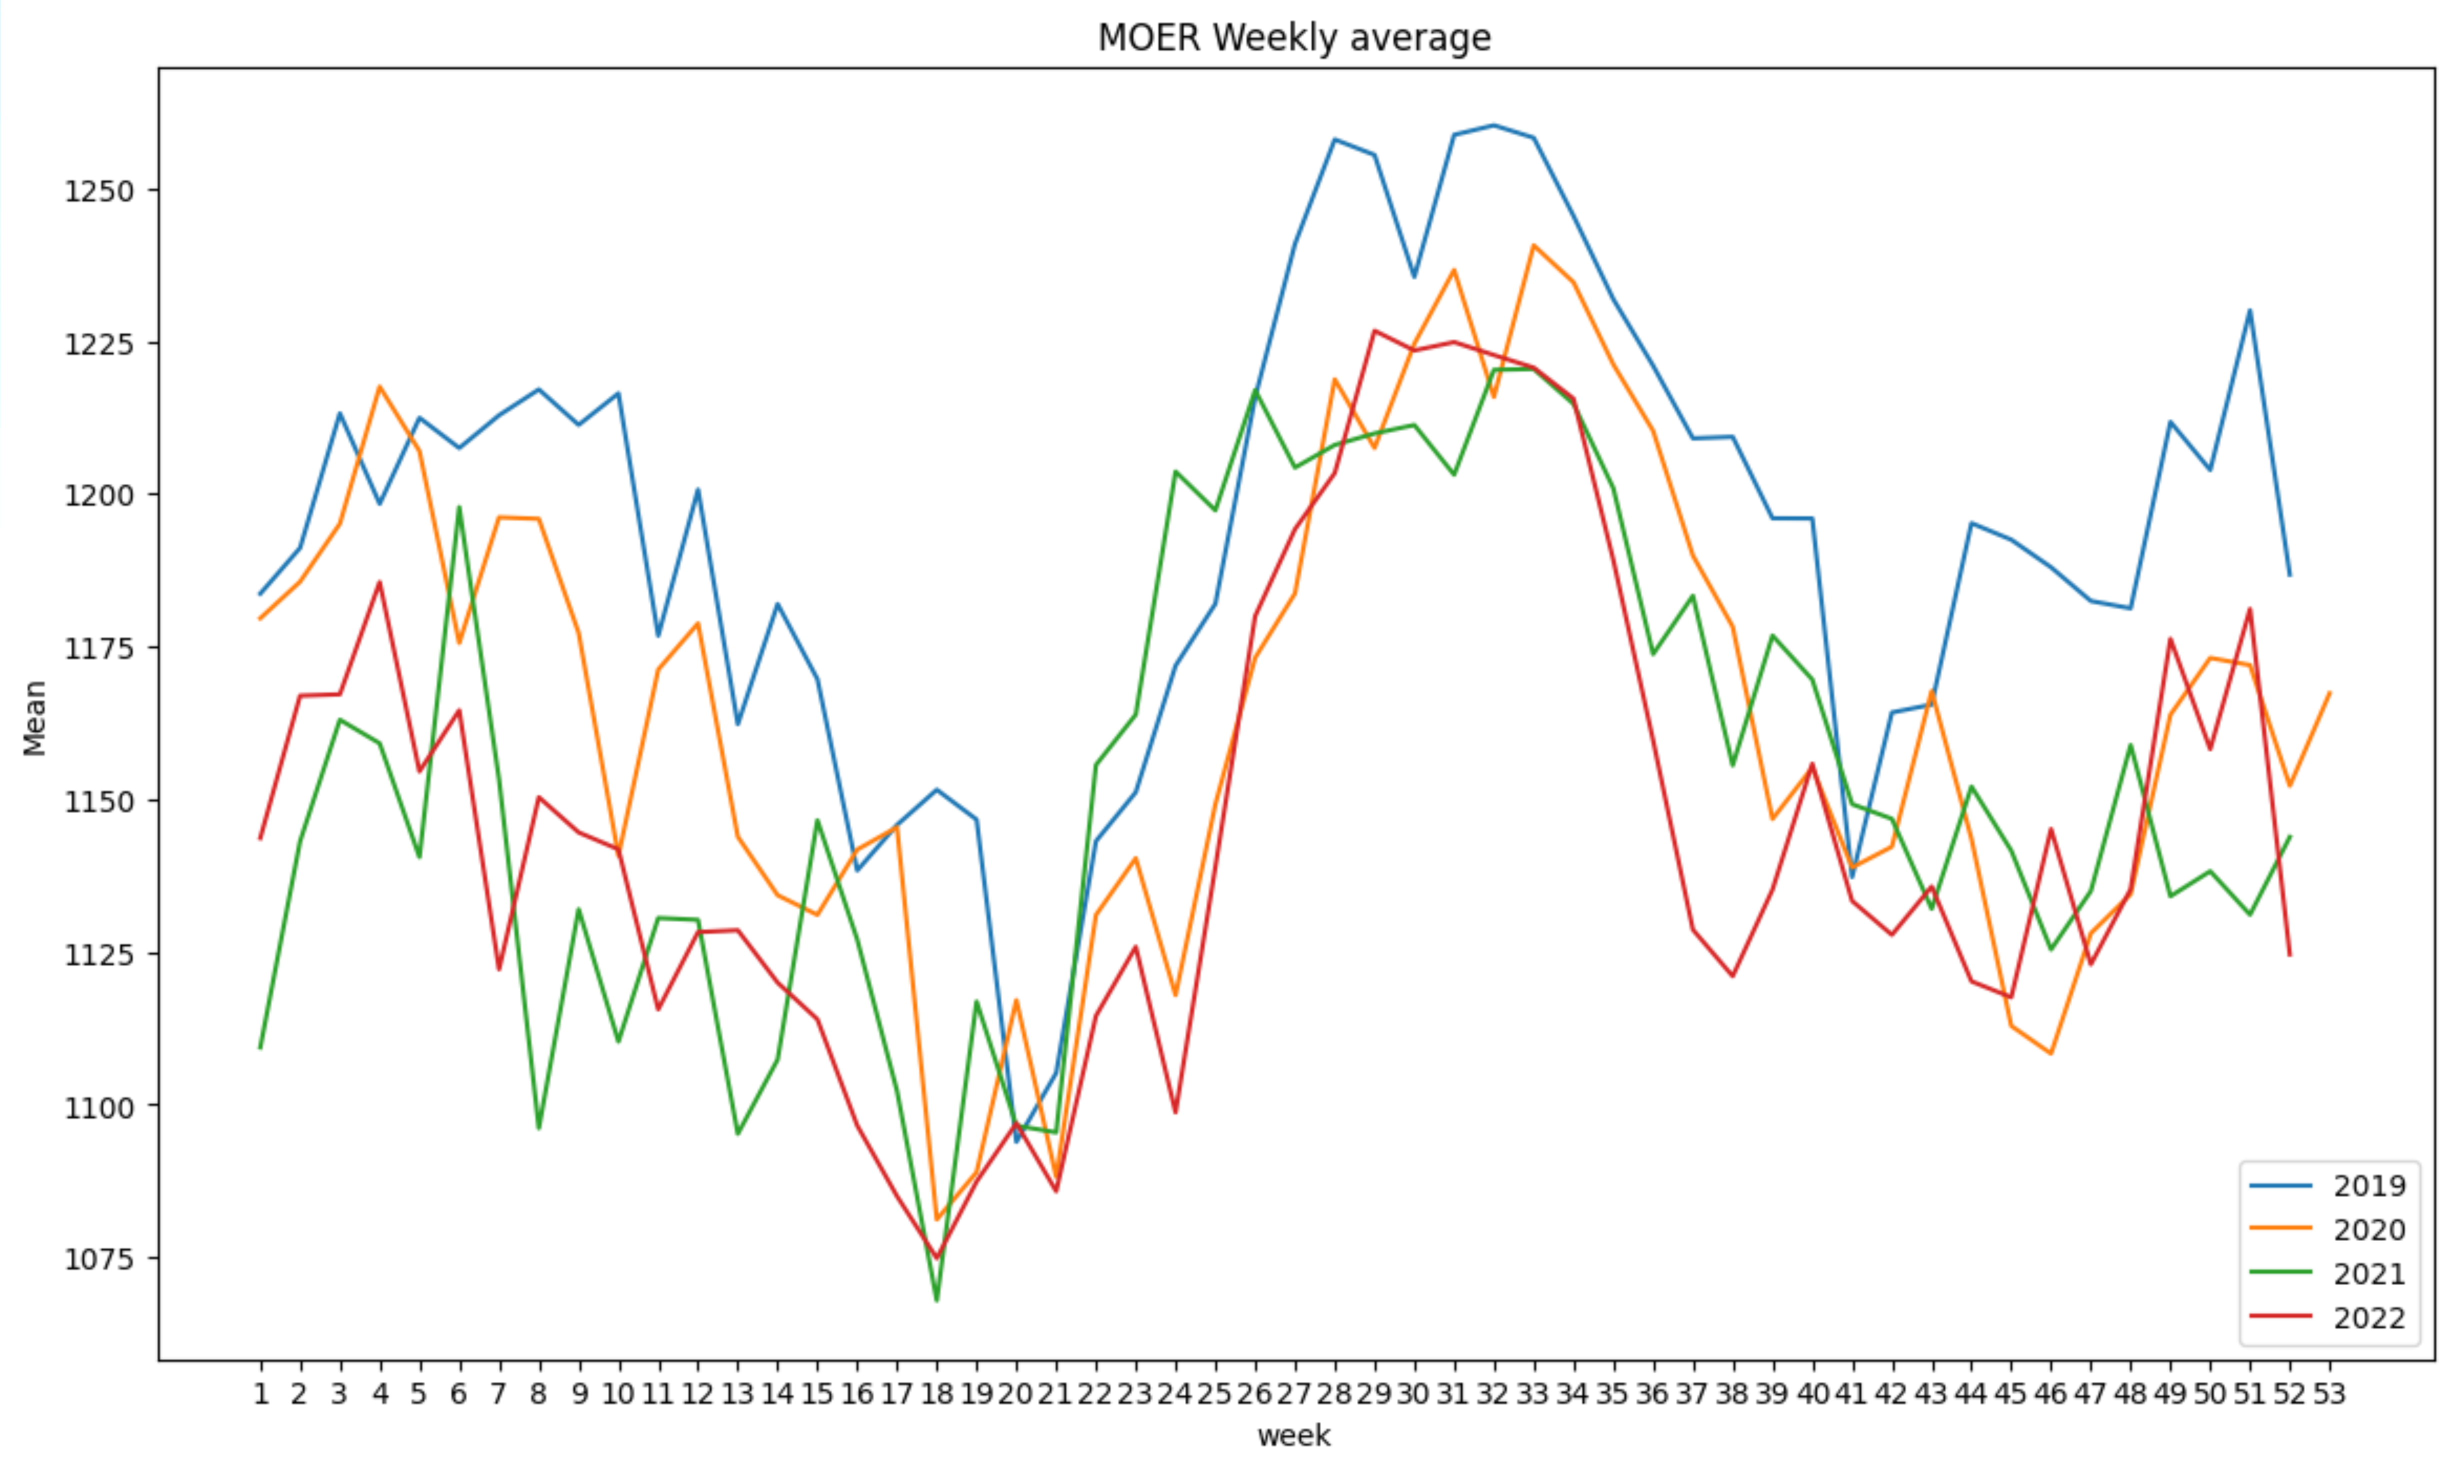

MOER analysis

Top 10 green/red regions

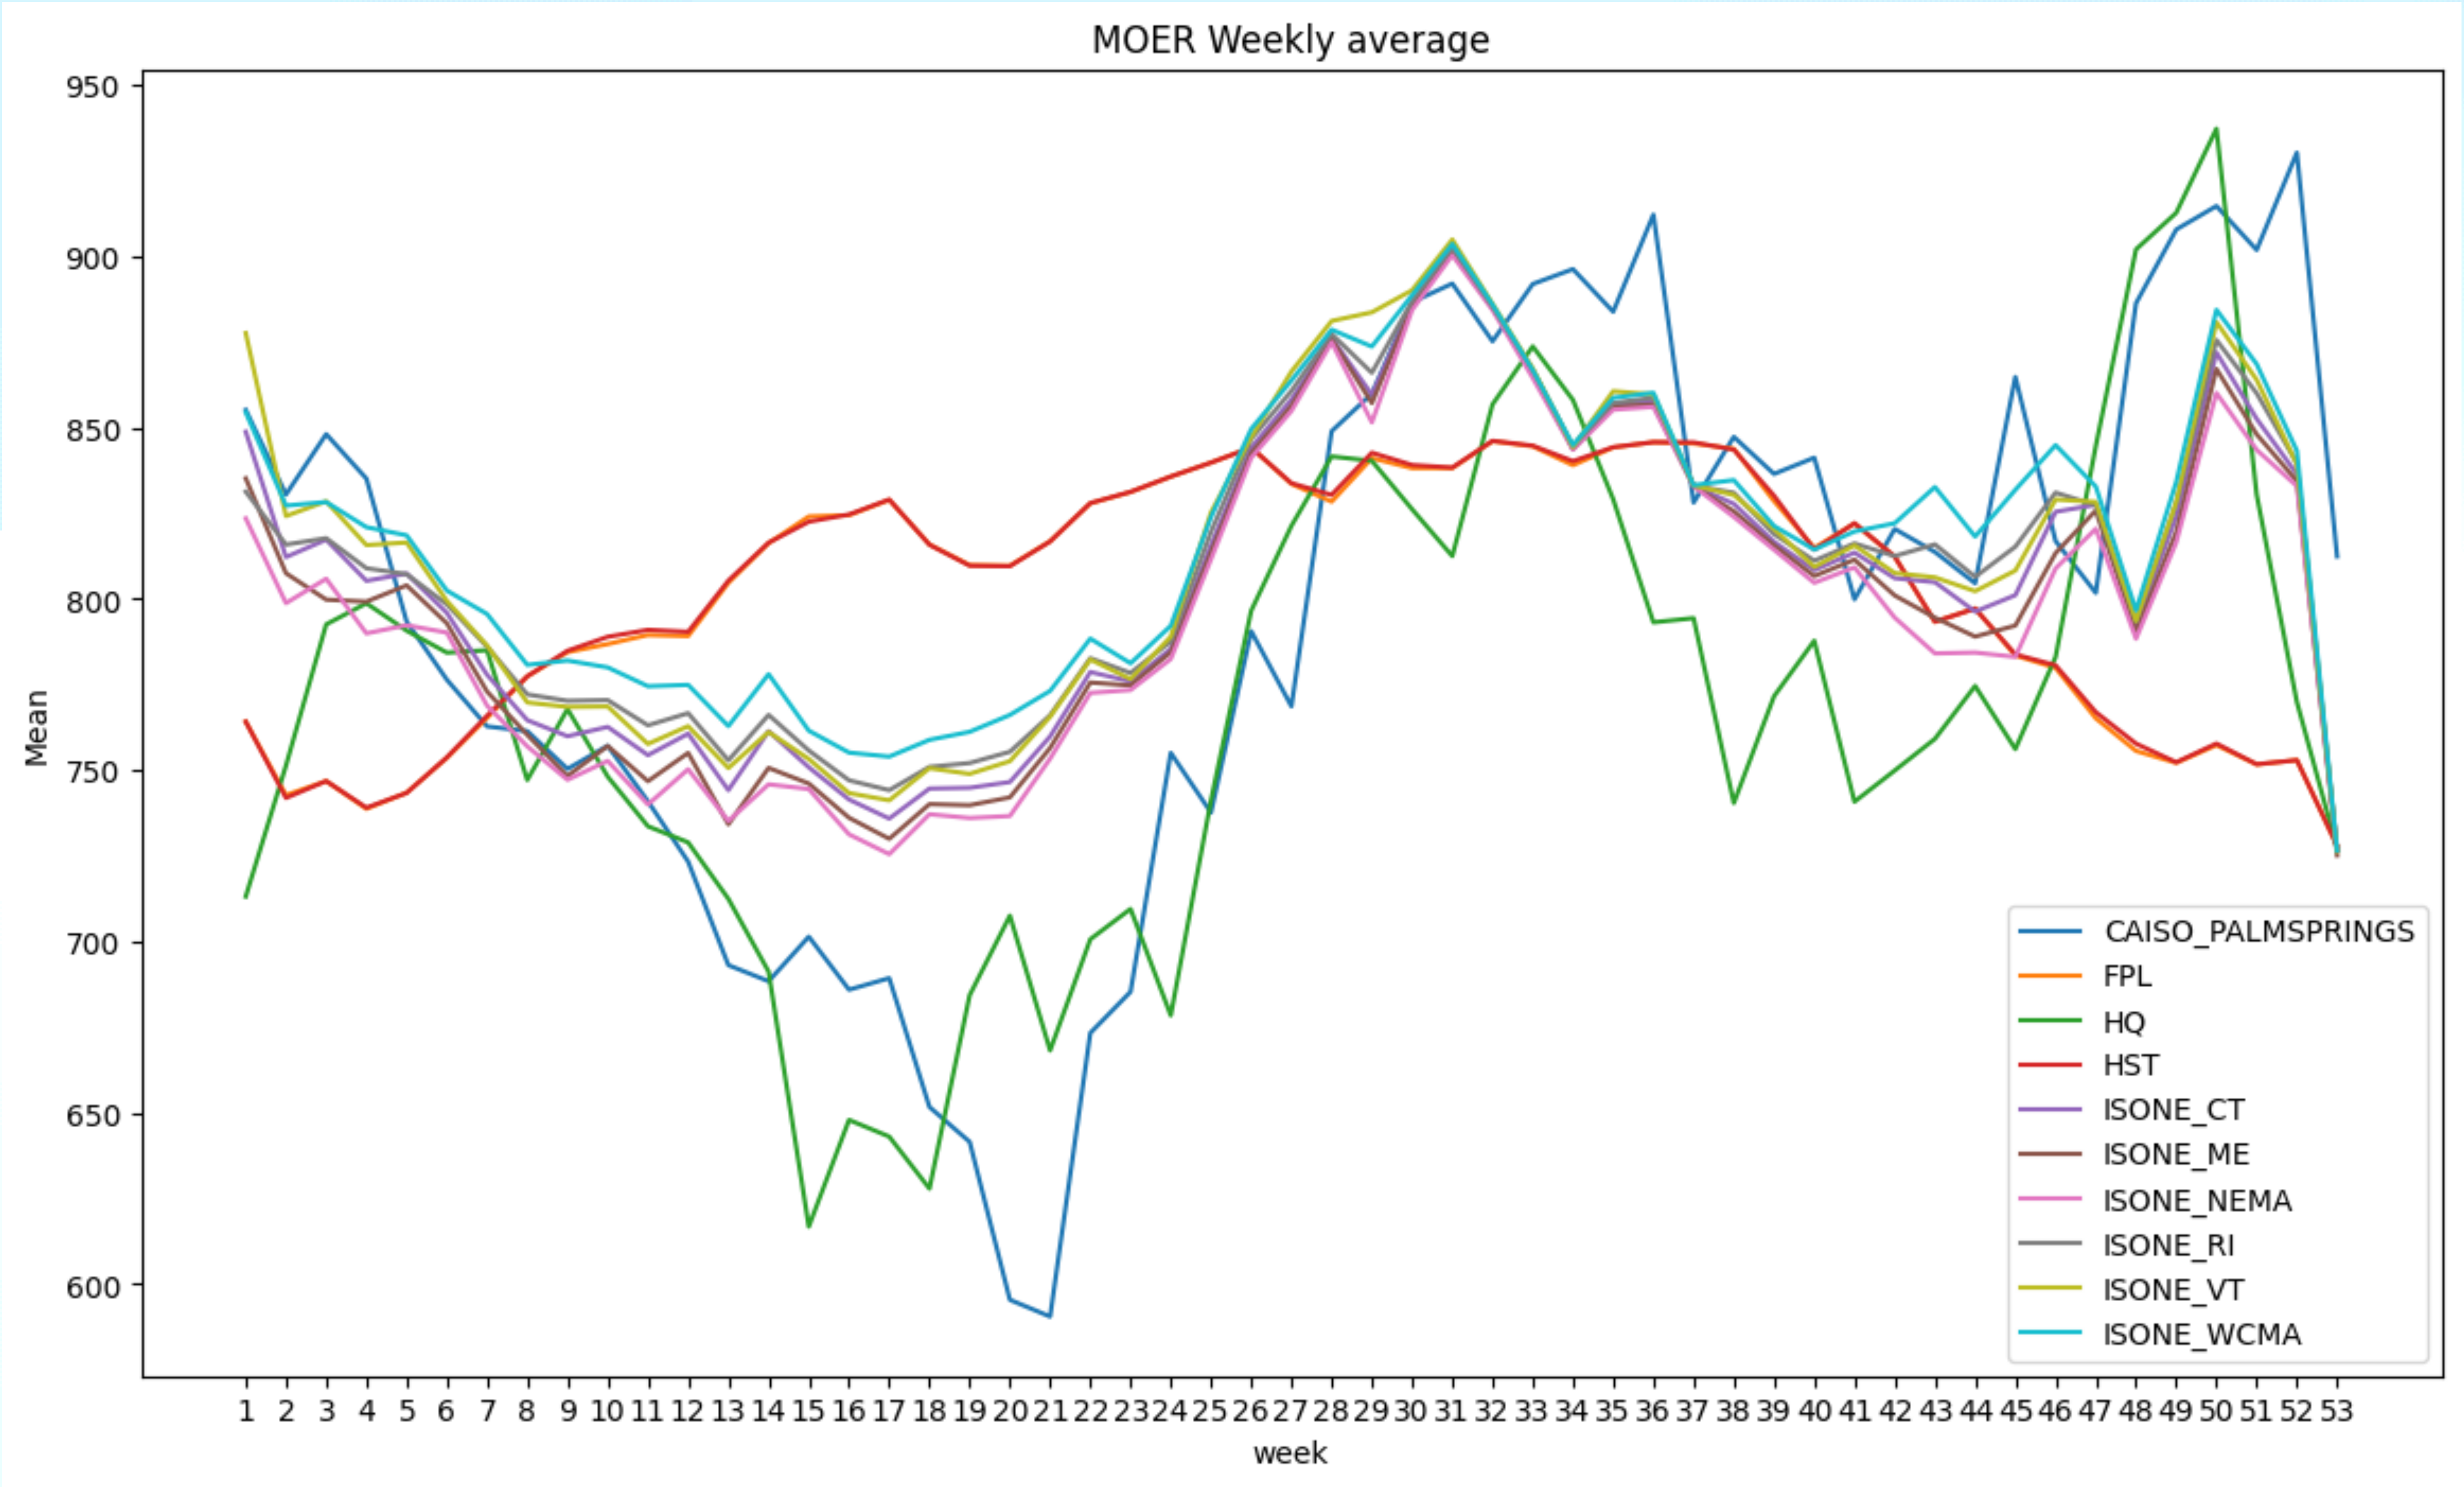

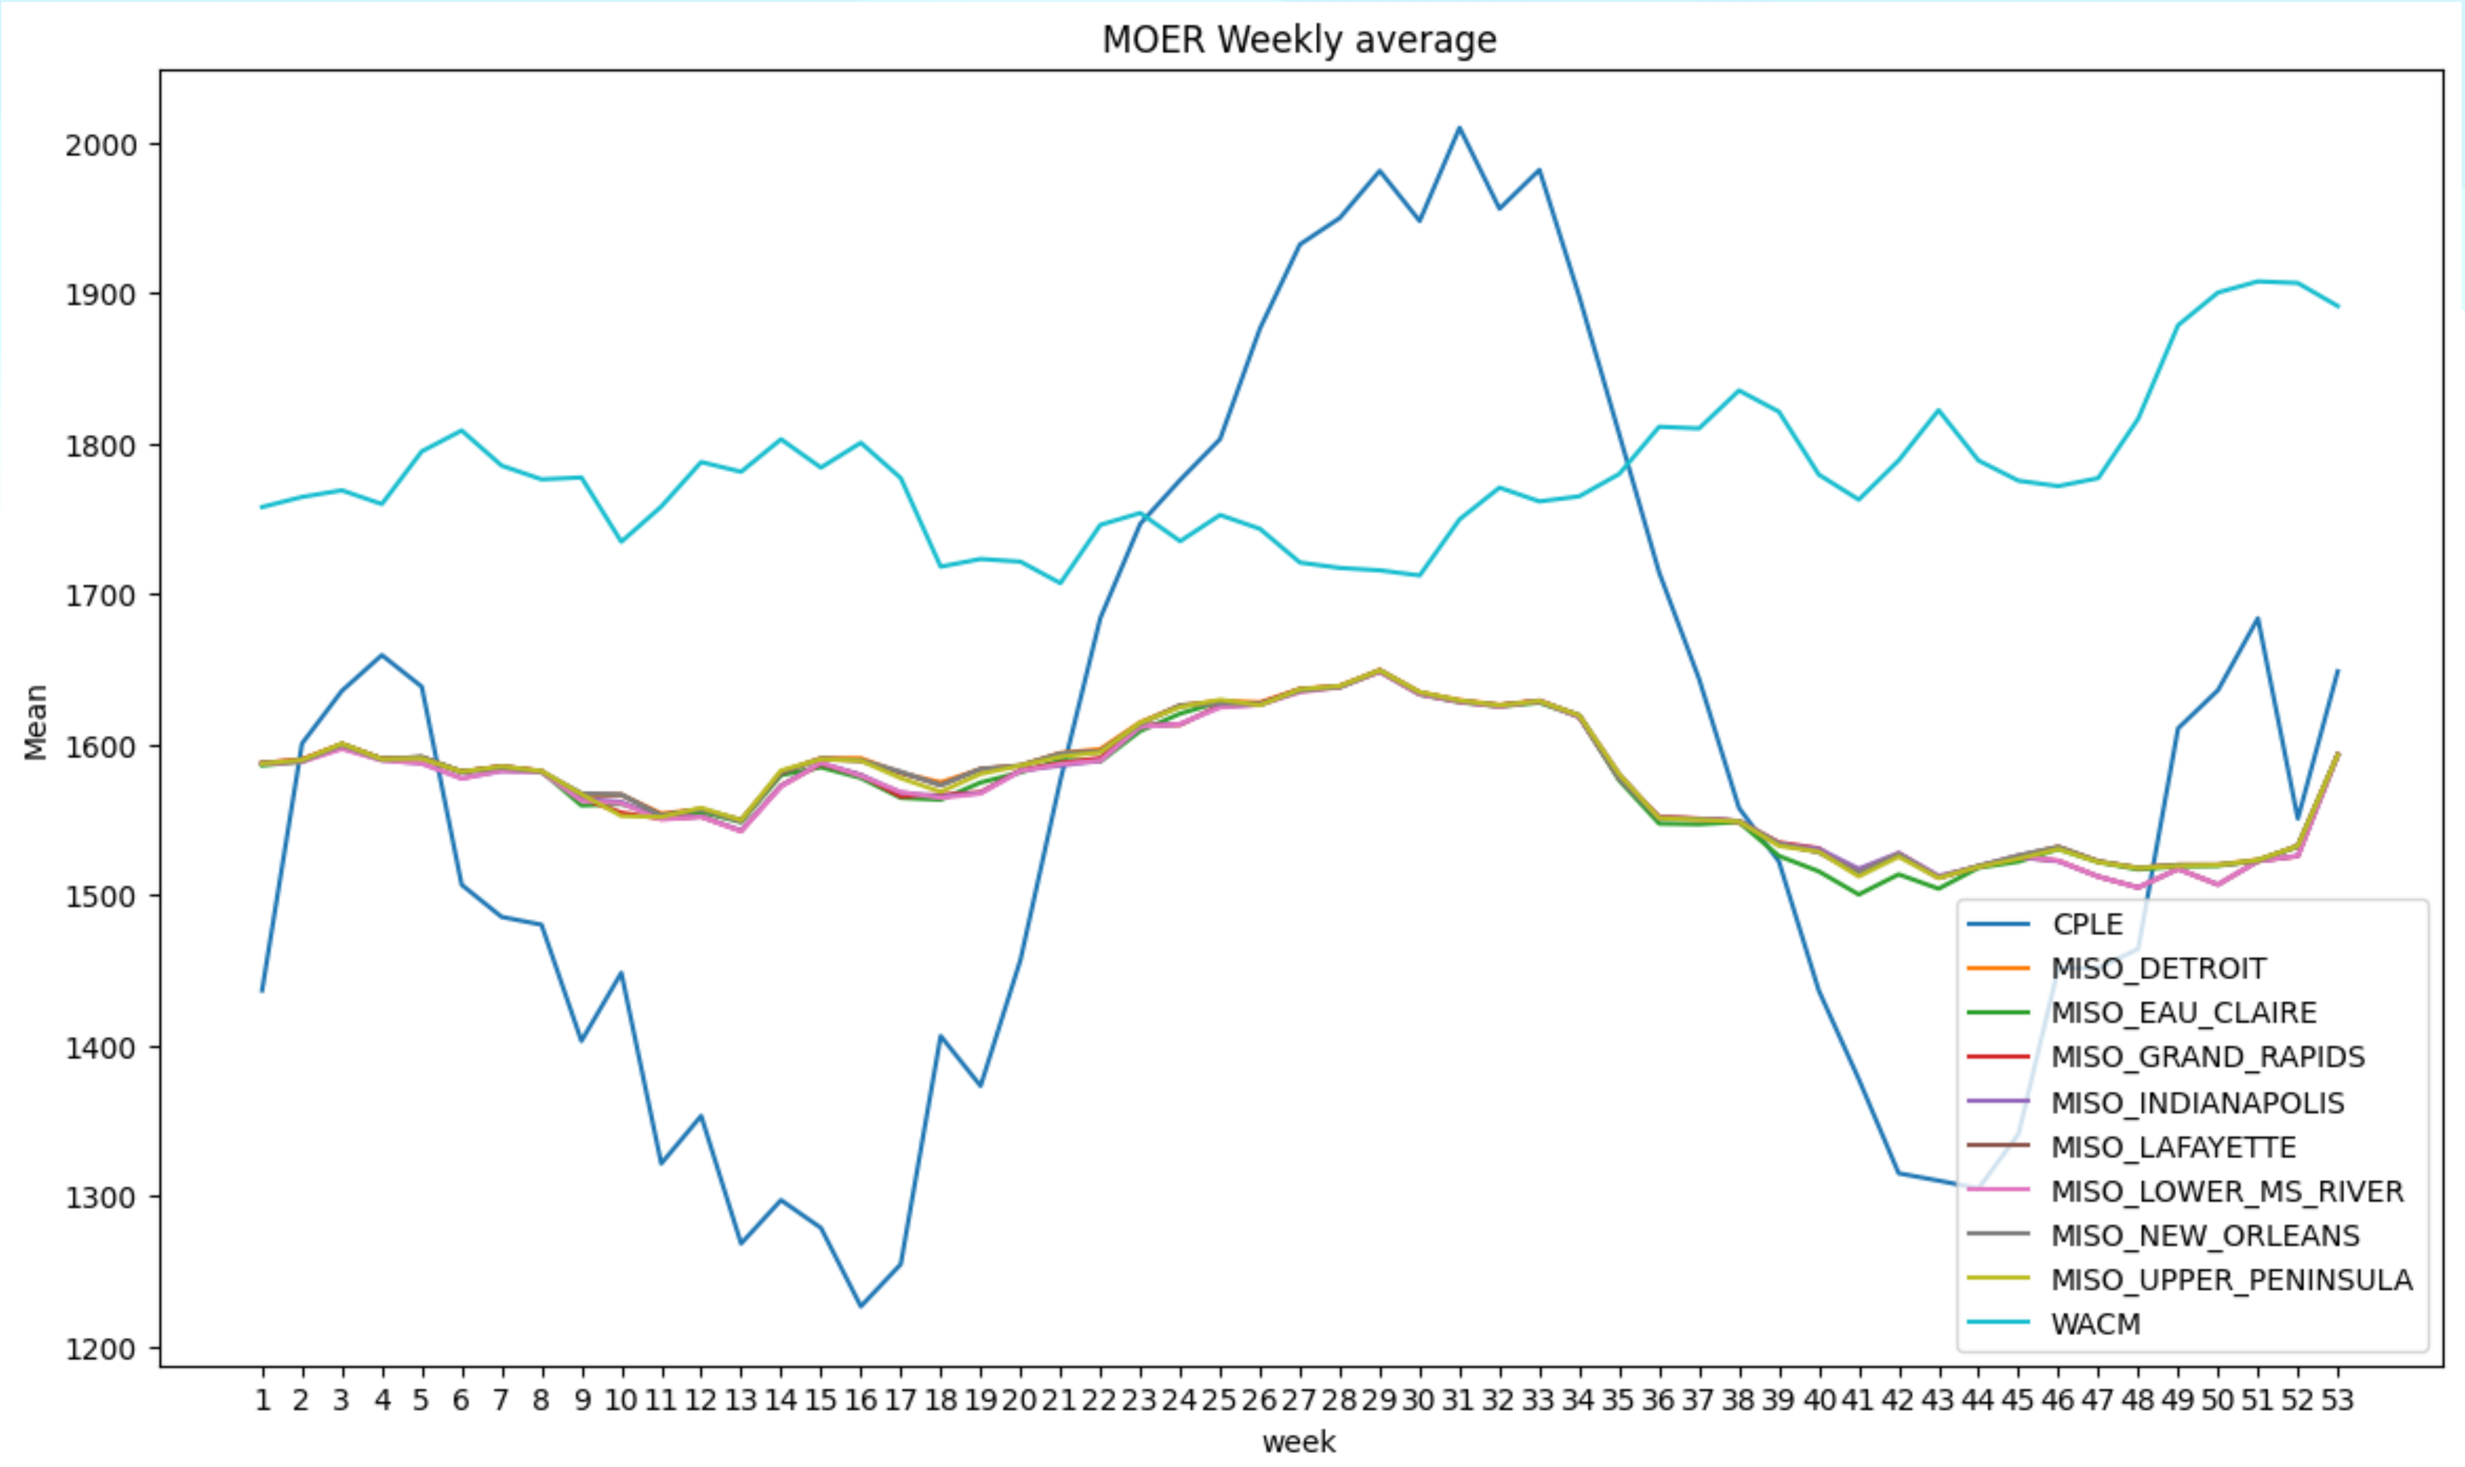

MOER analysis

MOER & Fuel Mix

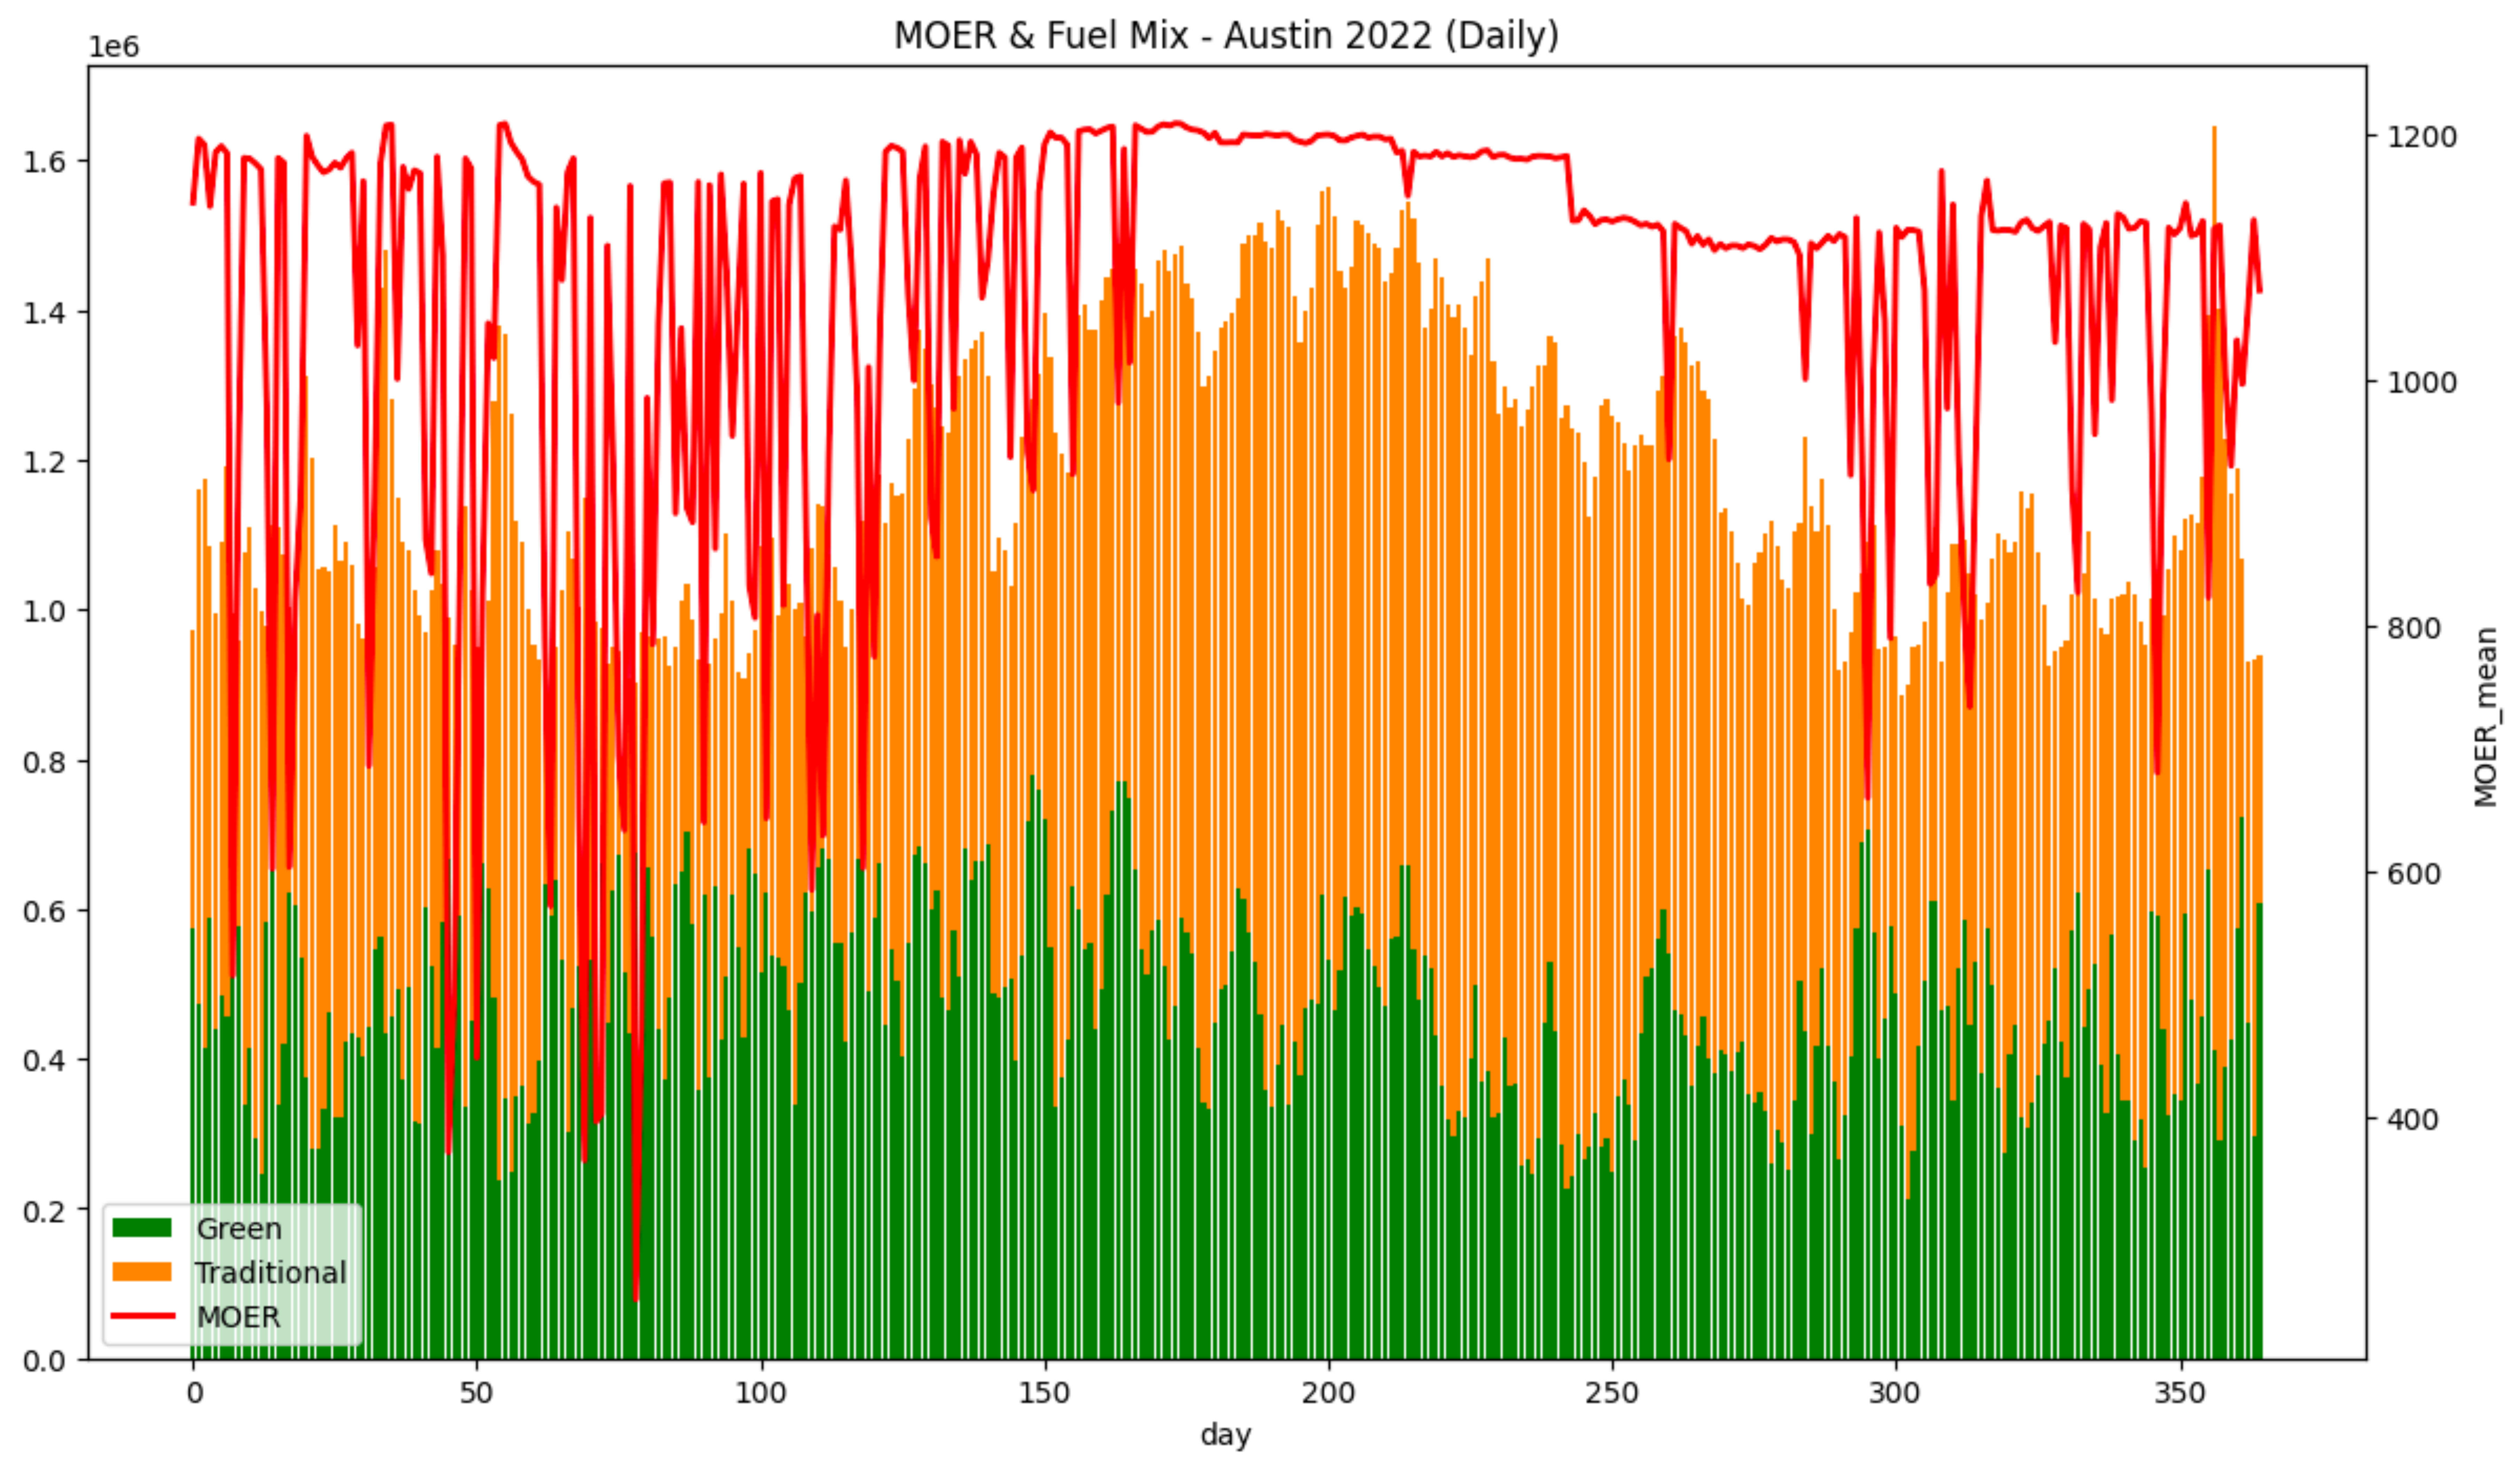

MOER analysis

MOER analysis

Regional analysis

Table 1: MOER & Data Centers

MOER	AWS	Google Cloud	Mircosoft Azure	Oracle Cloud	IBM Cloud	Alibaba Cloud	Tencent Cloud	TOTAL
600 - 900	5	6	9	6	-	2	2	30
900 - 1200	4	9	17	10	3	1	1	45
1200 - 1500	15	8	16	10	2	2	2	55
1500 - 2000	2	4	7	2	2	2	1	20

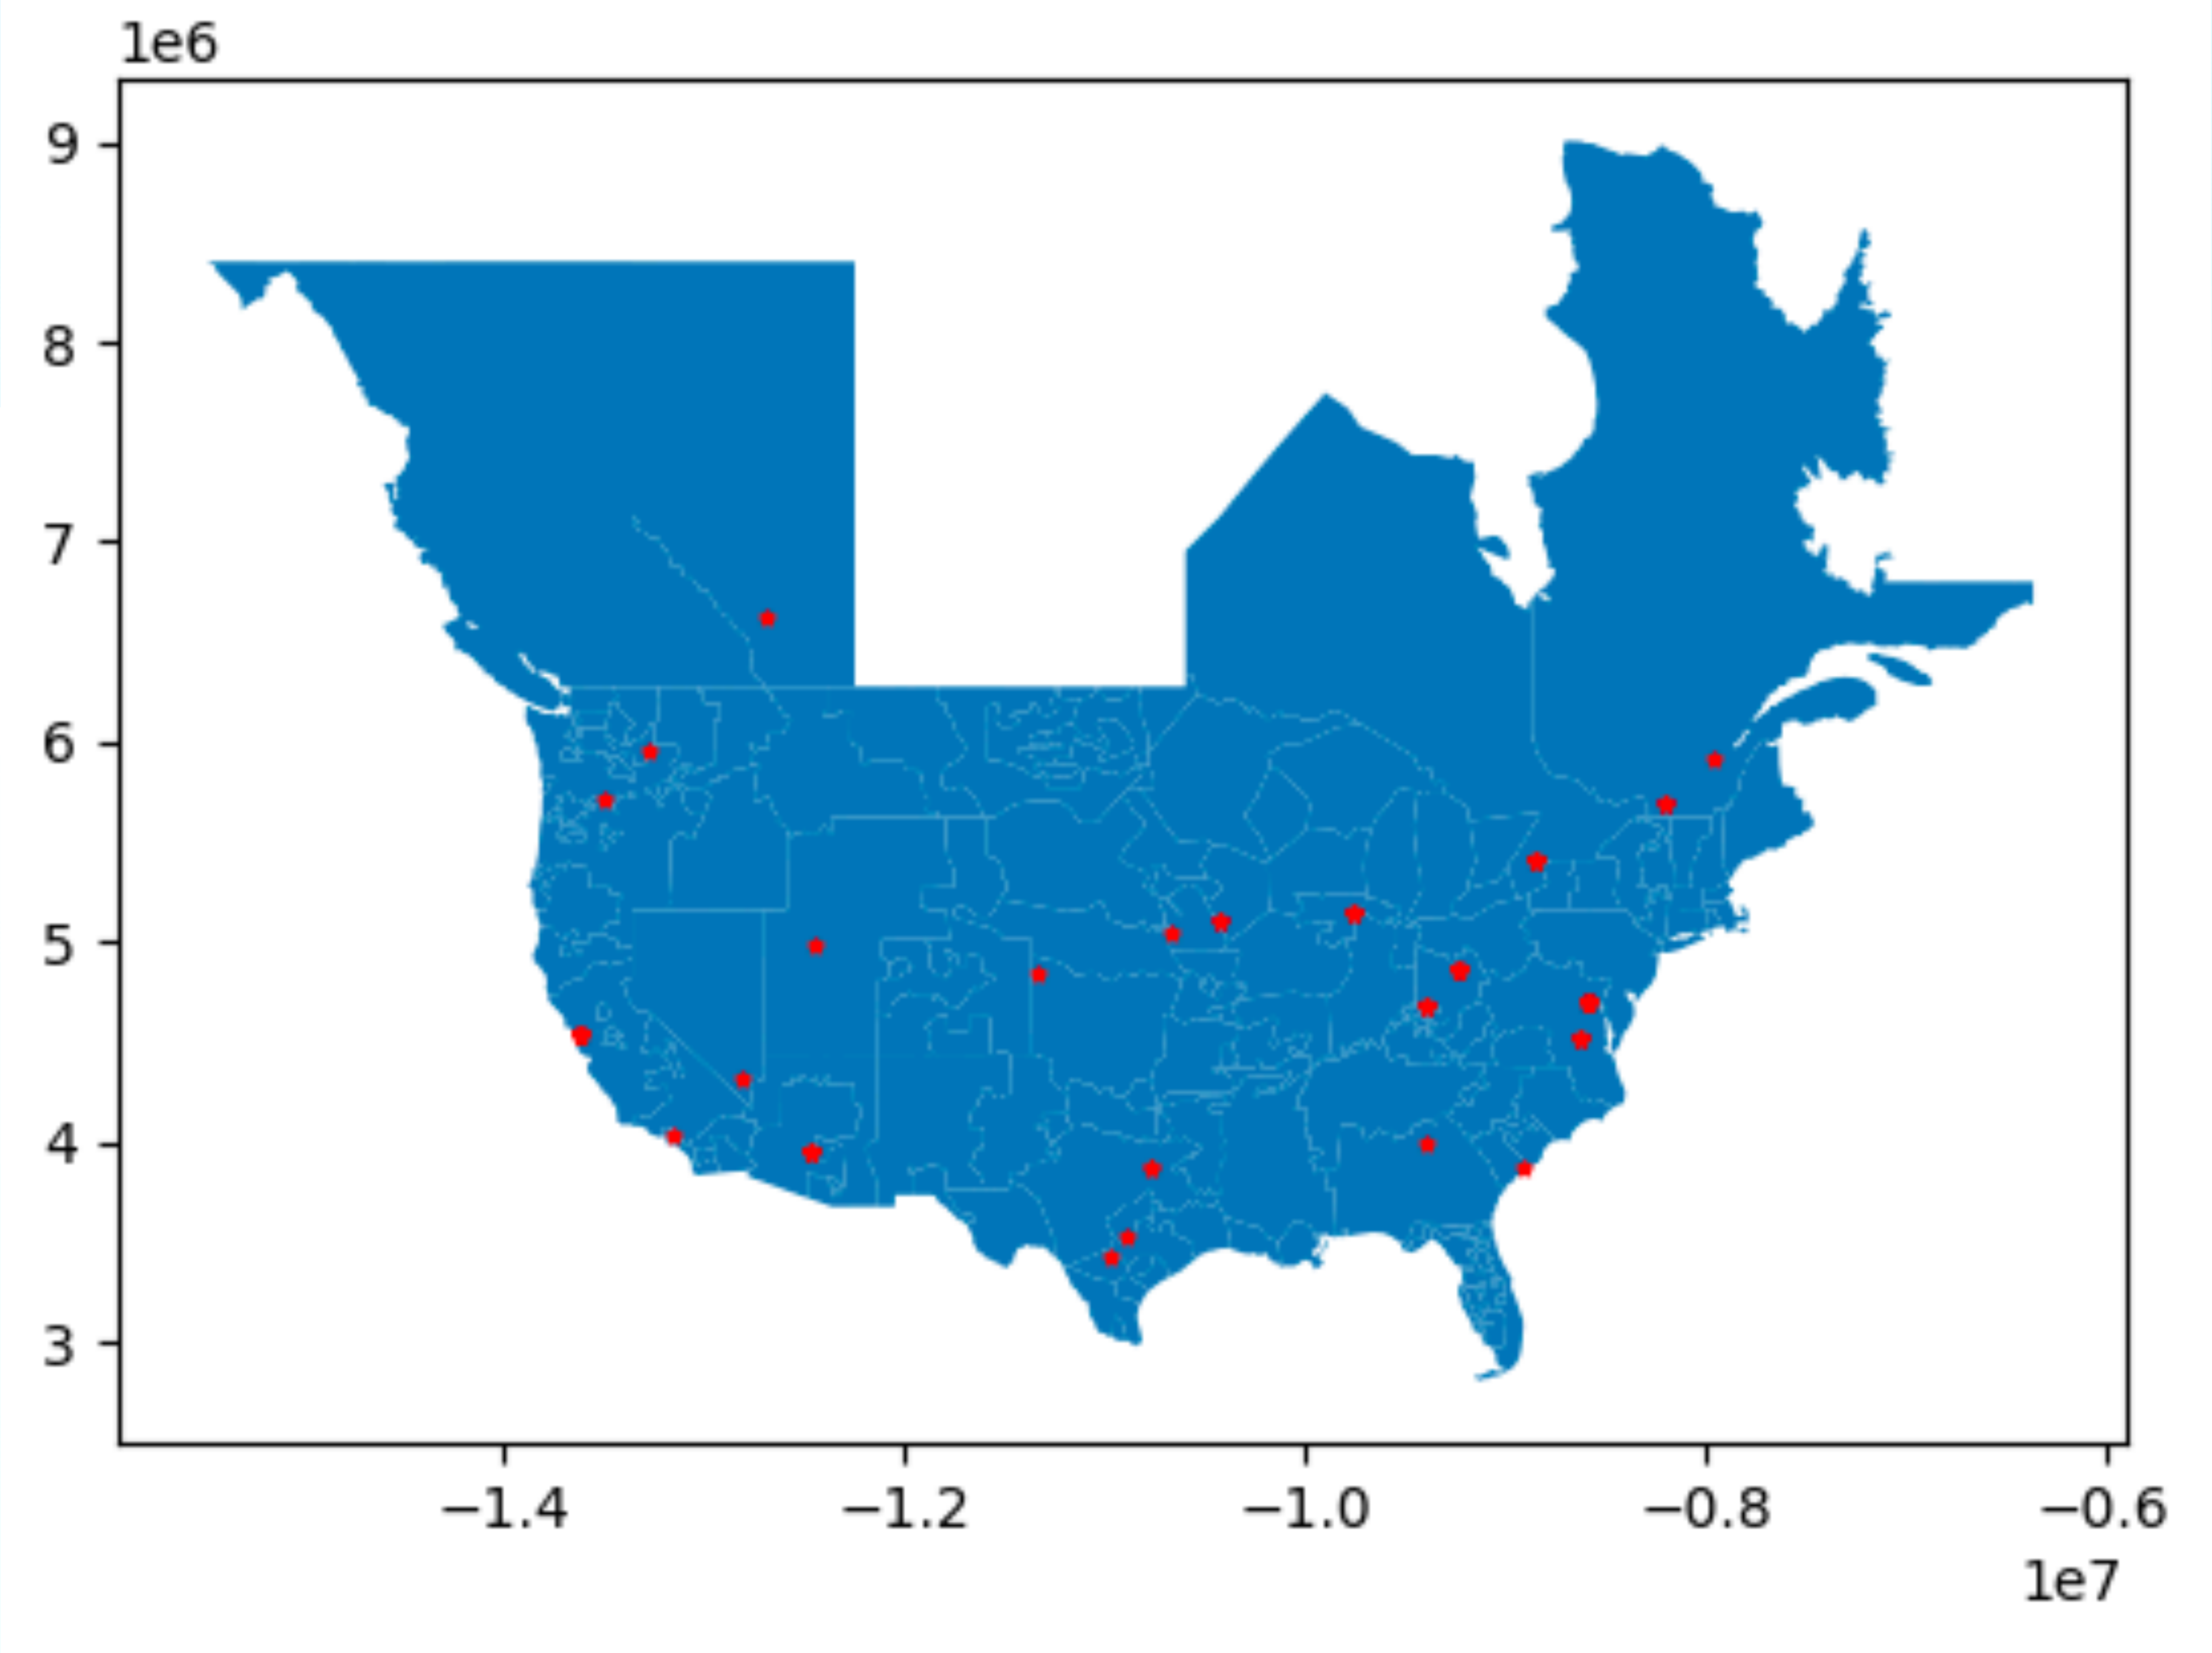

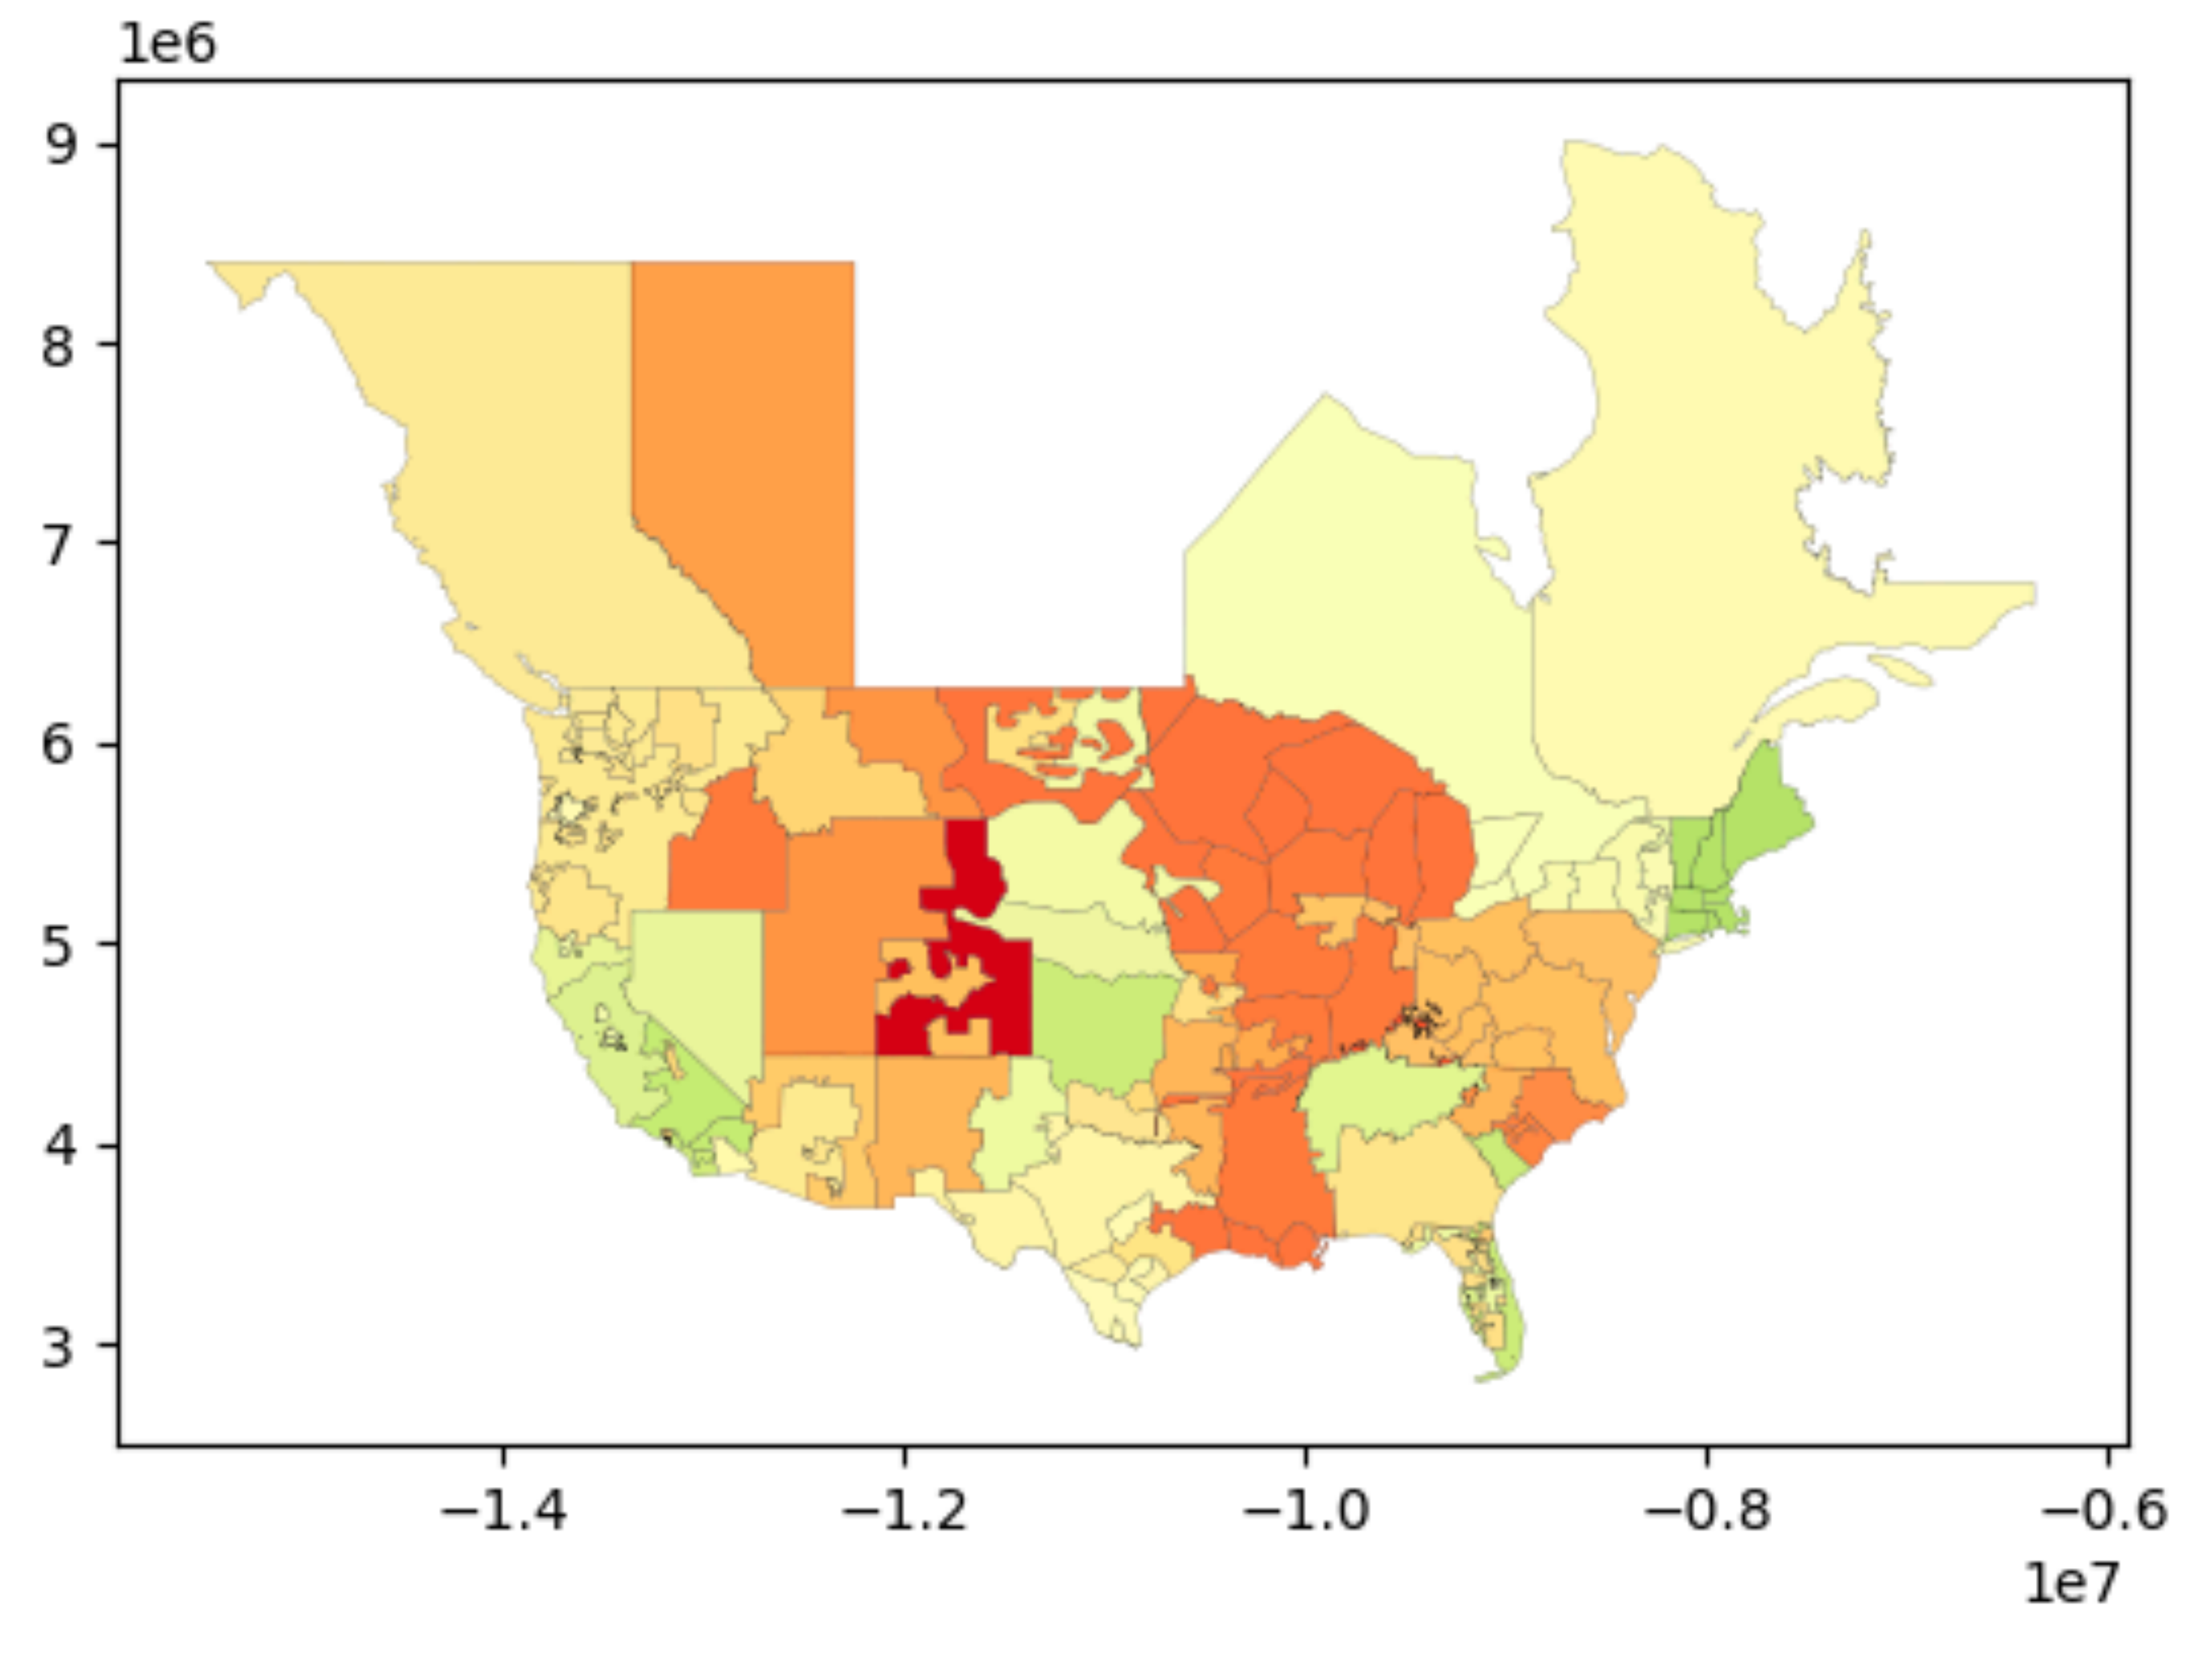

Regional analysis

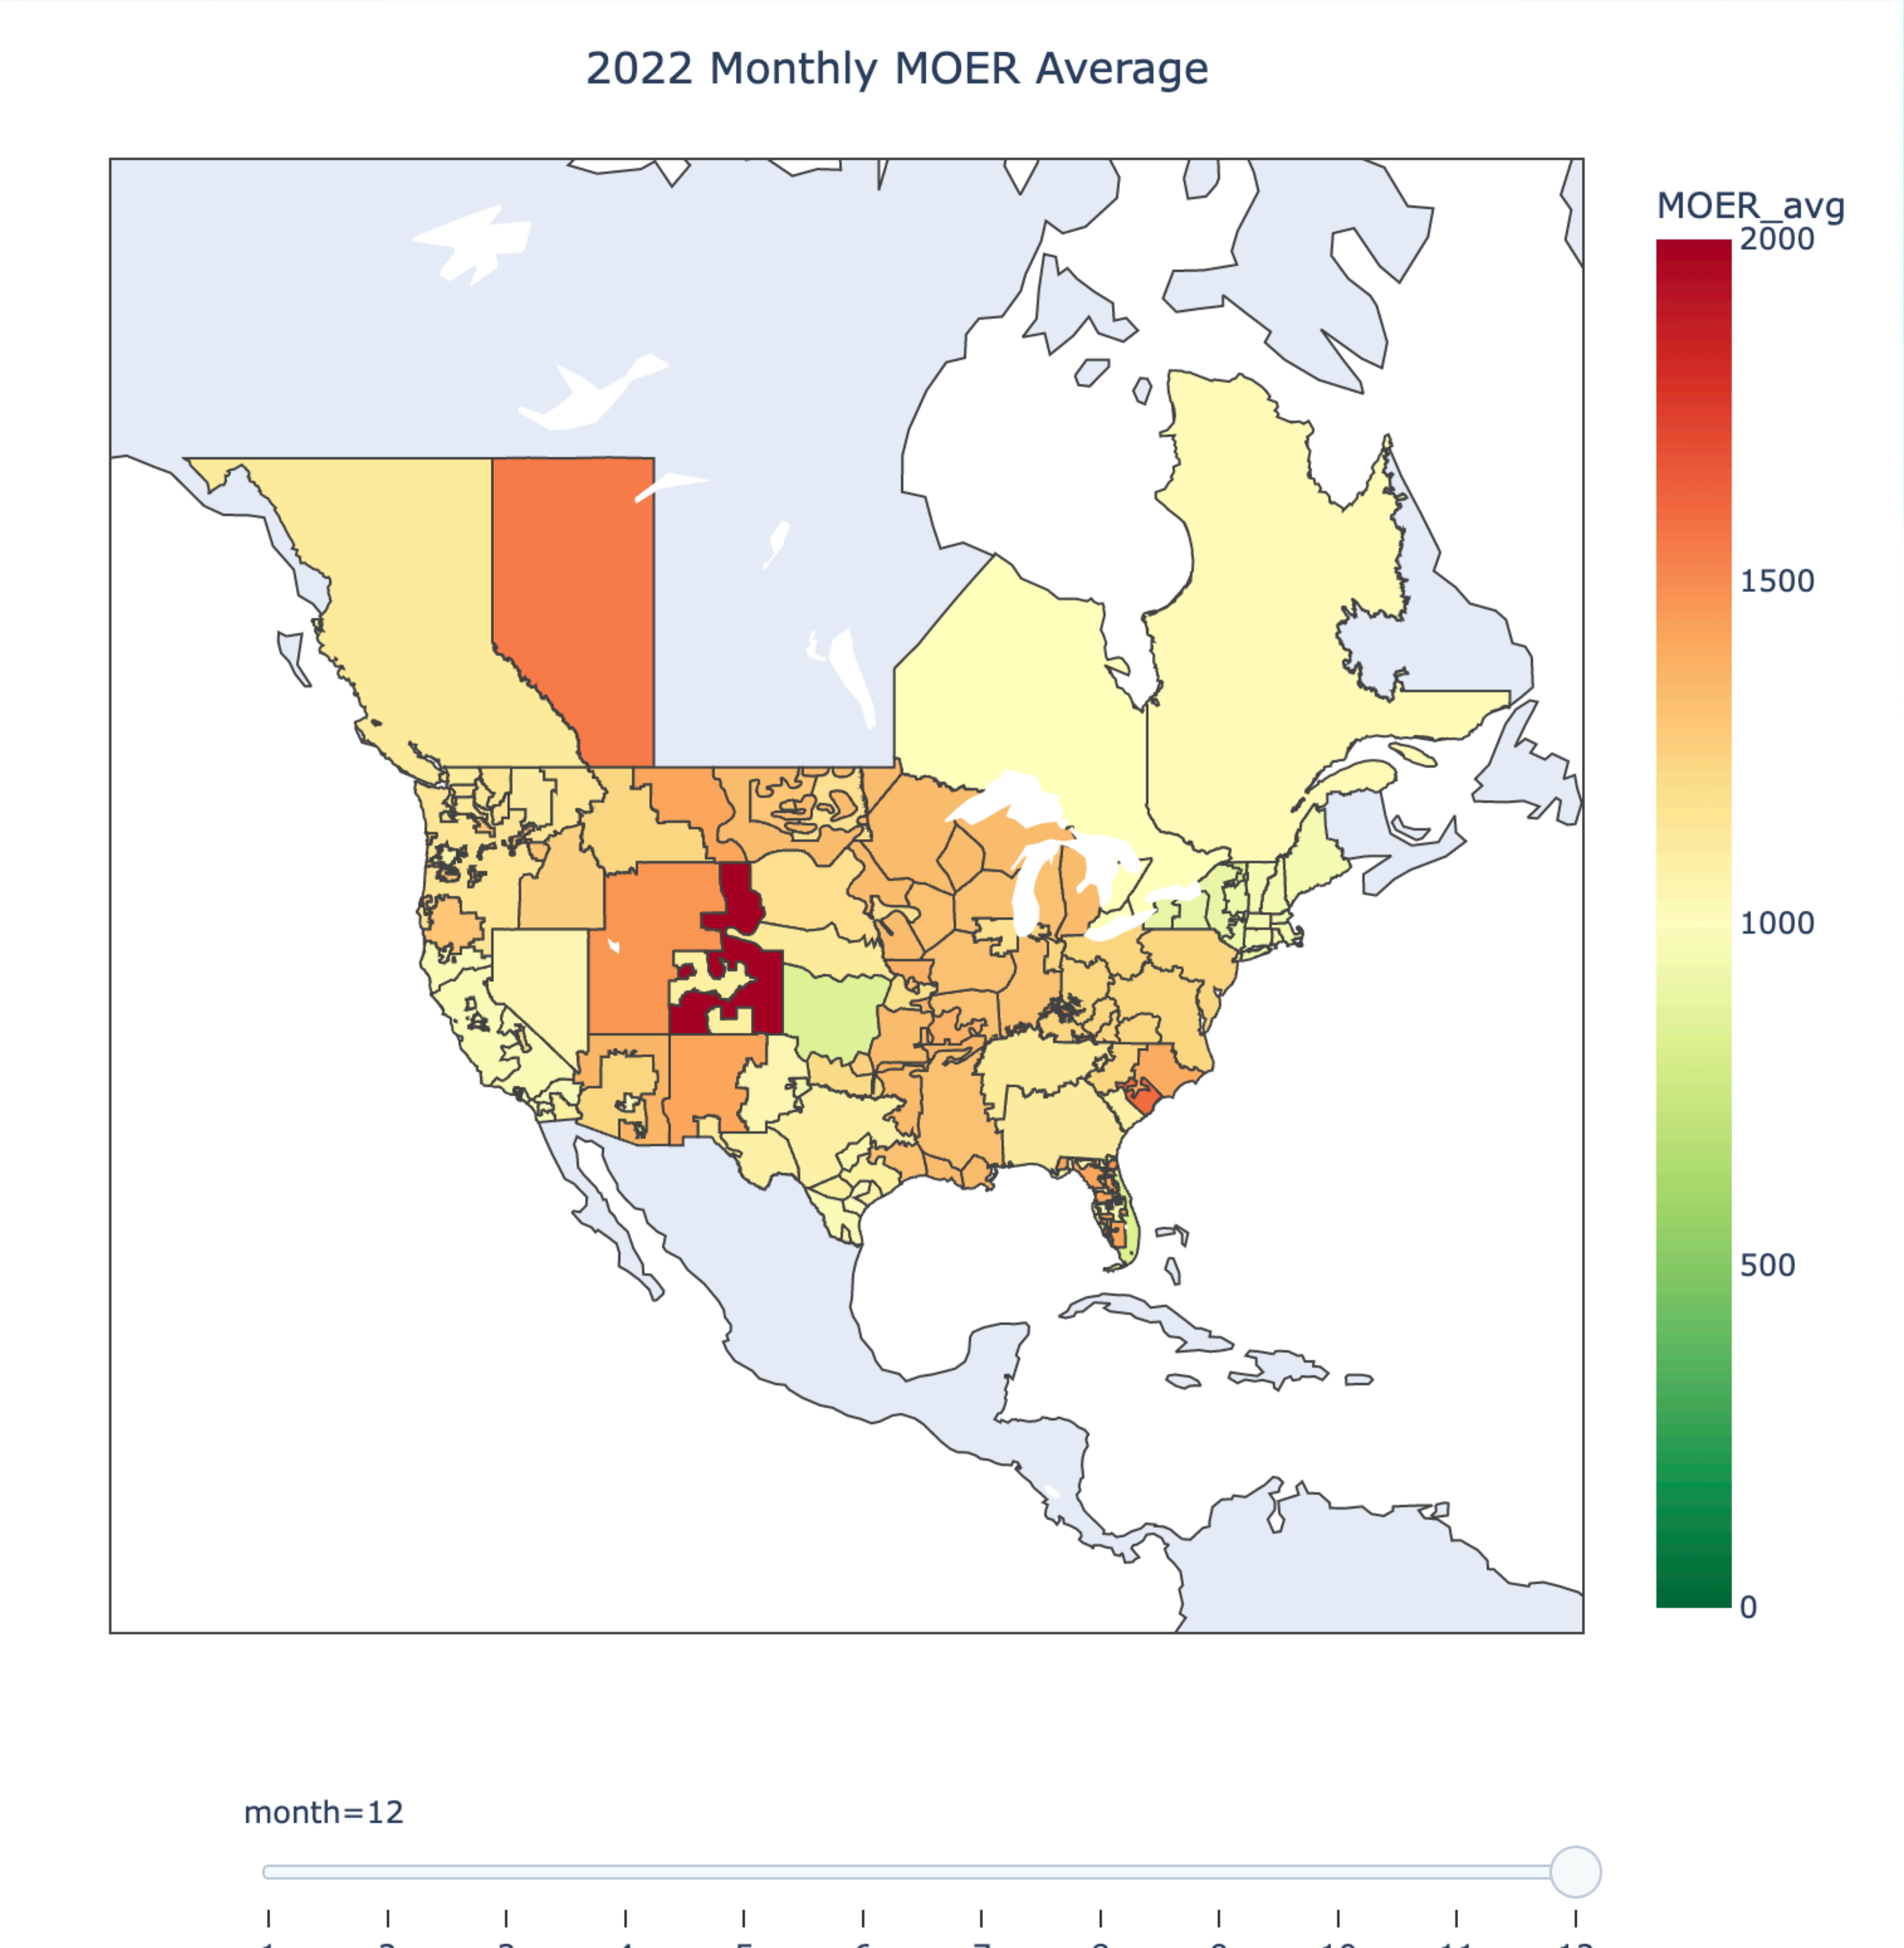

Azure VM Dataset

- The trace contains a representative subset of the first-party Azure VM workload in one geographical region.
- Duration: 30 consecutive days
- Total number of VMs: 2,695,548
- Total VM hours: 104,371,713
- Total virtual core hours: >380,000,000

Azure VM Dataset

AzurePublicDatasetV2 - VM Trace

1. Encrypted subscription id
2. Encrypted deployment id
3. Timestamp in seconds (starting from 0) when first VM created
4. Count VMs created
5. Deployment size (we define a “deployment” differently than Azure in our paper)
6. Encrypted VM id
7. **Timestamp VM created**
8. **Timestamp VM deleted**
9. Max CPU utilization
10. **Avg CPU utilization**
11. P95 of Max CPU utilization
12. VM category
13. **VM virtual core count bucket**
14. VM memory (GBs) bucket
15. Timestamp in seconds (every 5 minutes)
16. Min CPU utilization during the 5 minutes
17. Max CPU utilization during the 5 minutes
18. Avg CPU utilization during the 5 minutes
19. VM virtual core count bucket definition
20. VM memory (GBs) bucket definition

Azure VM Dataset

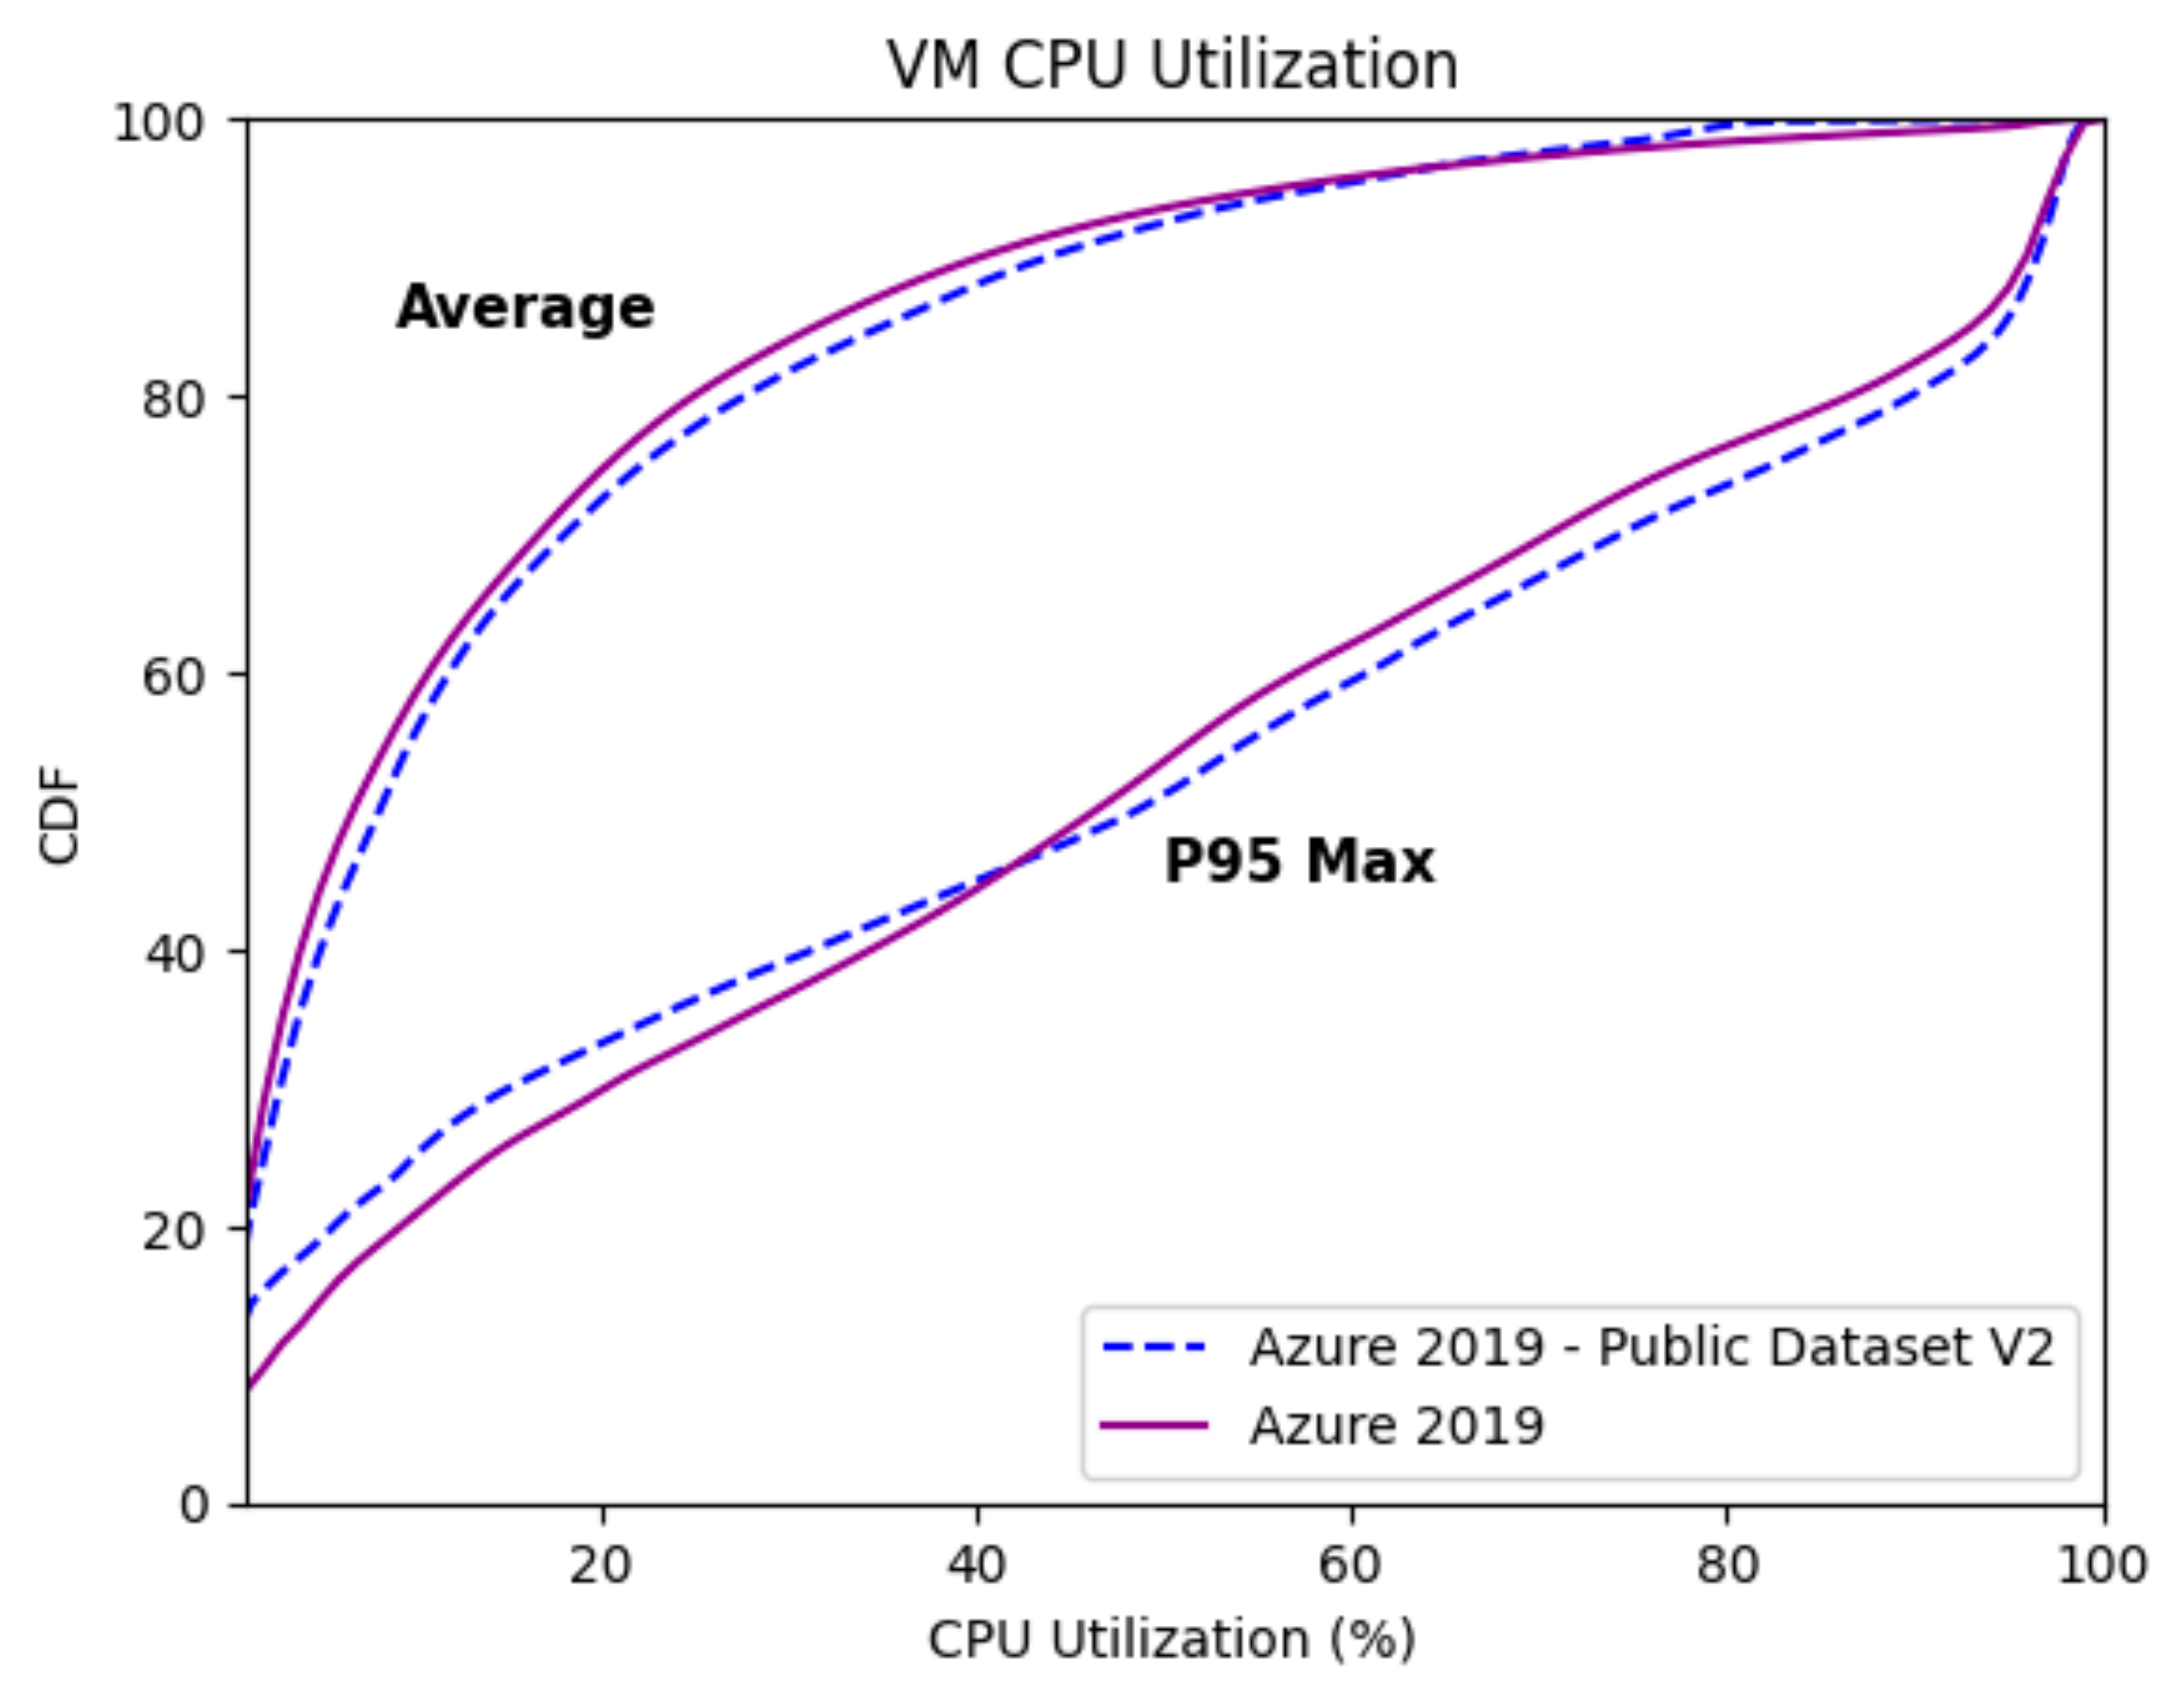

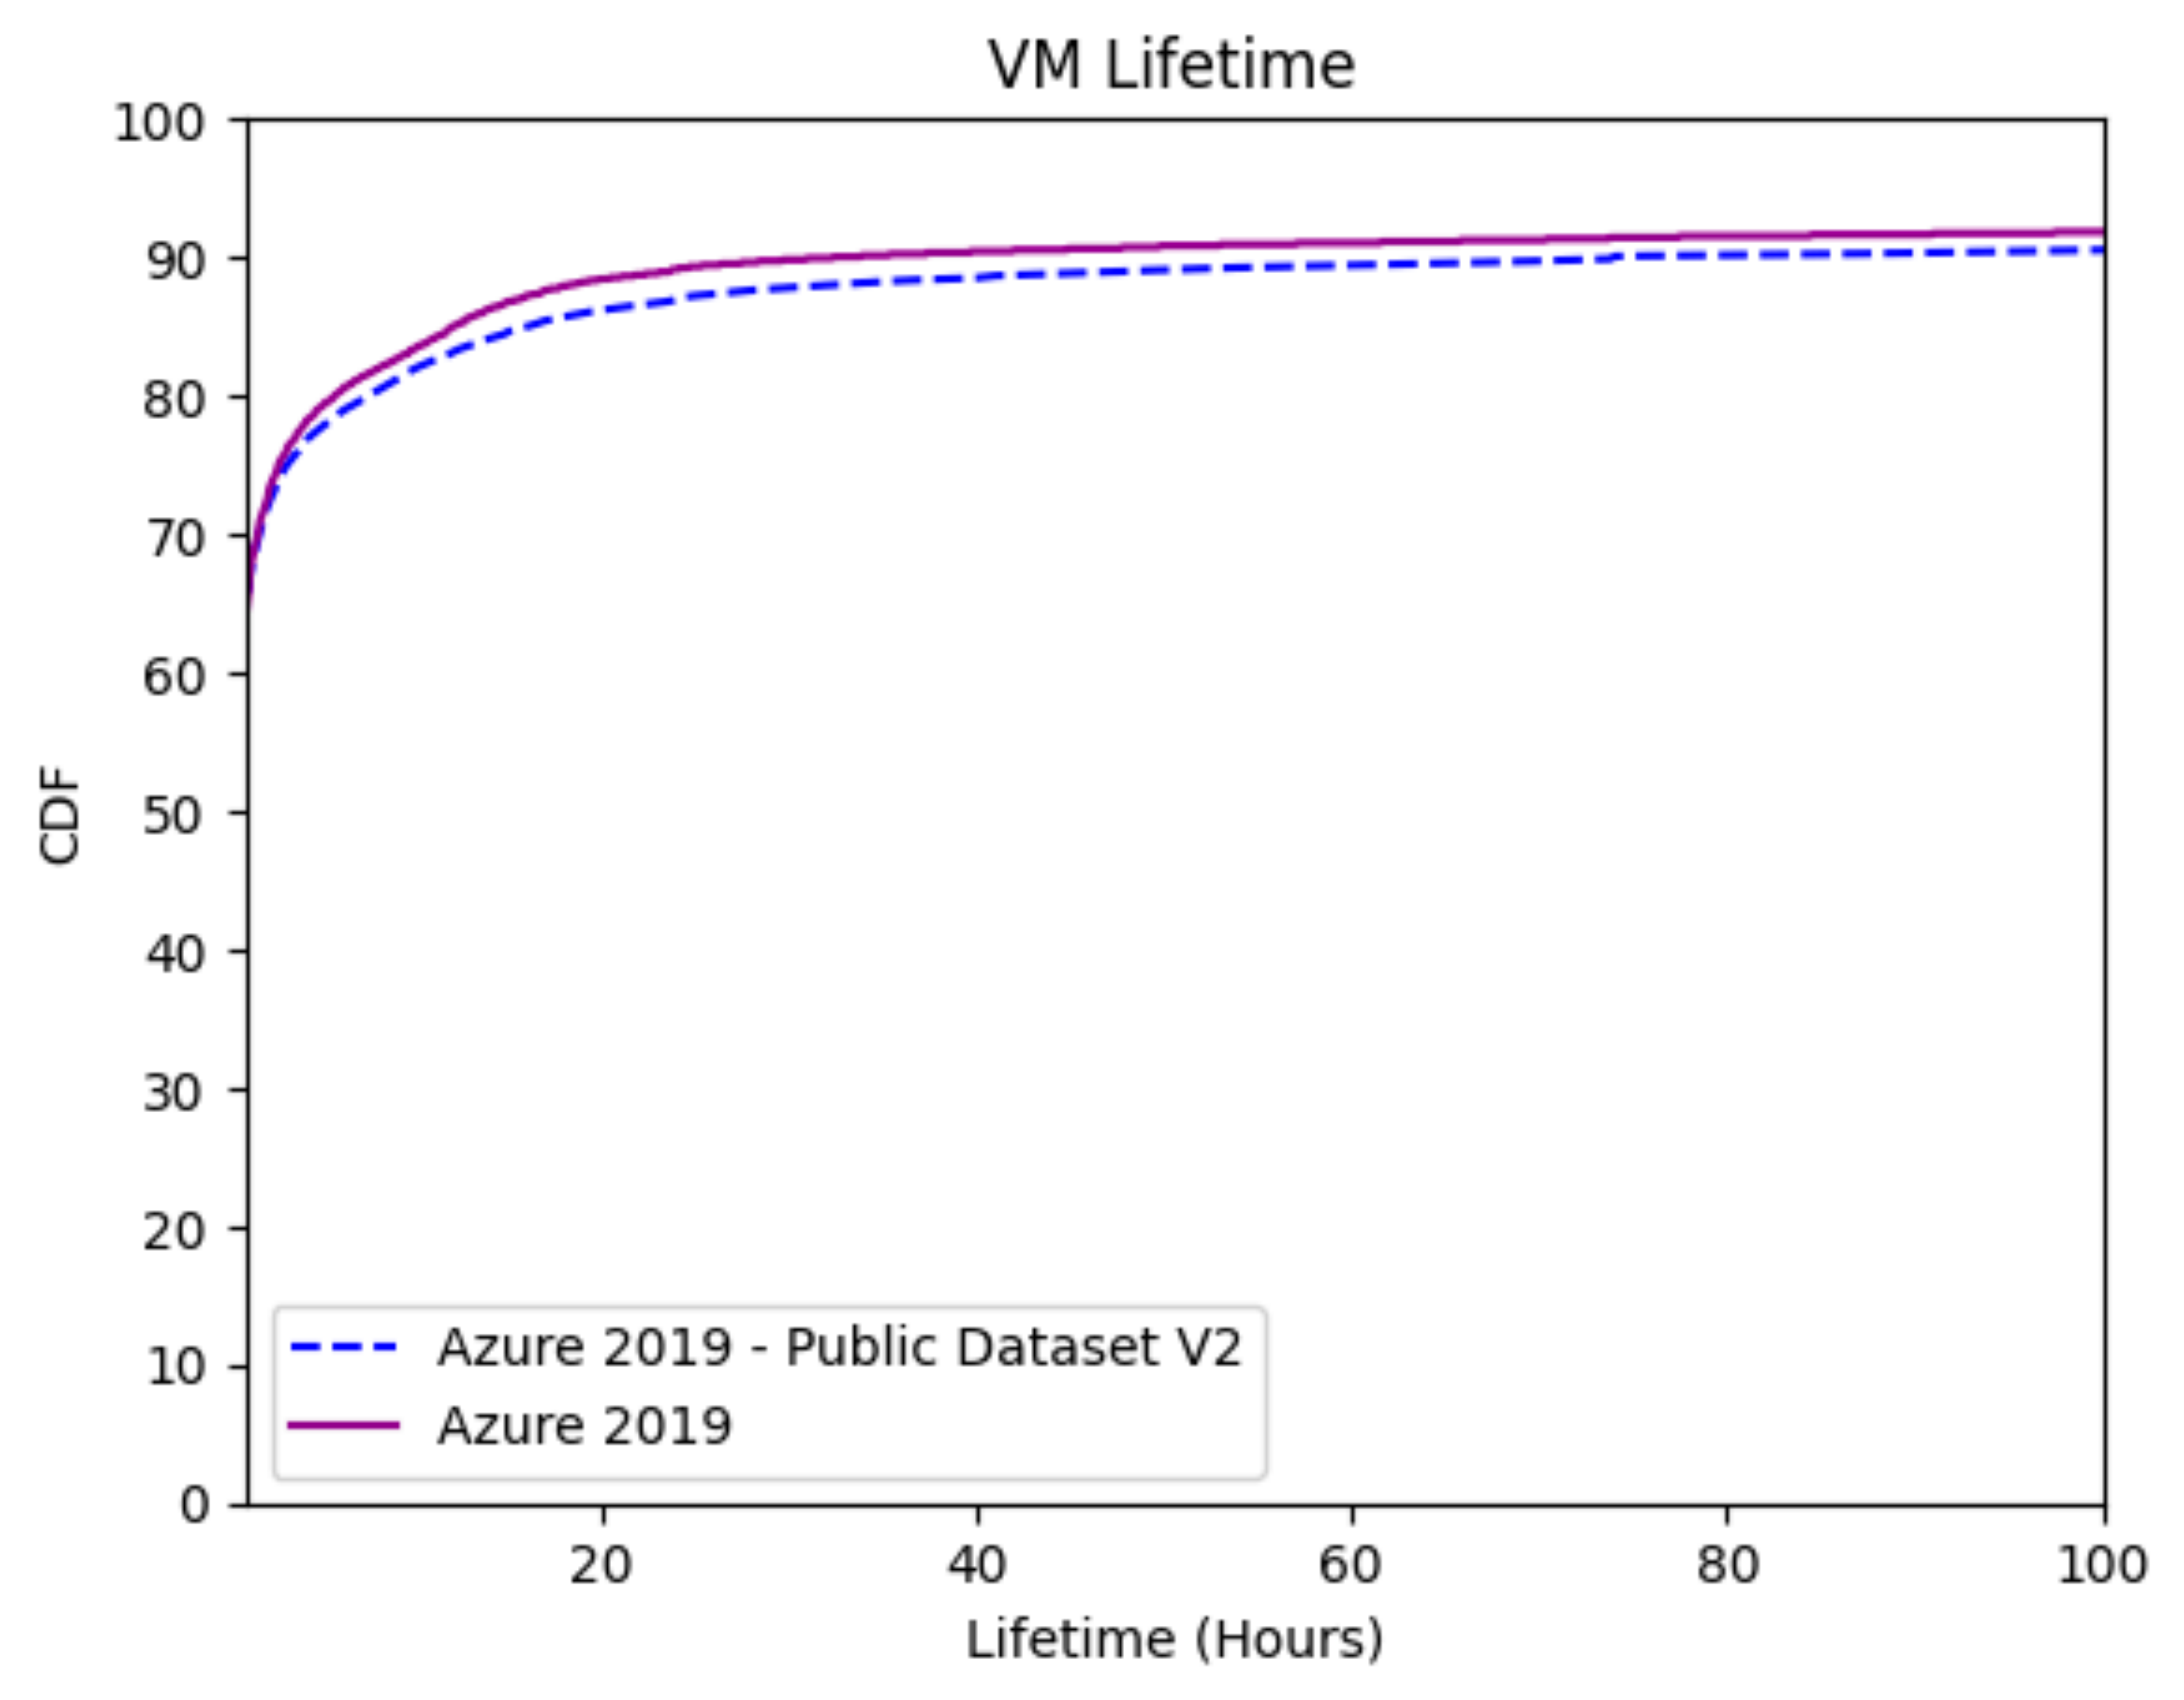

Azure VM Dataset

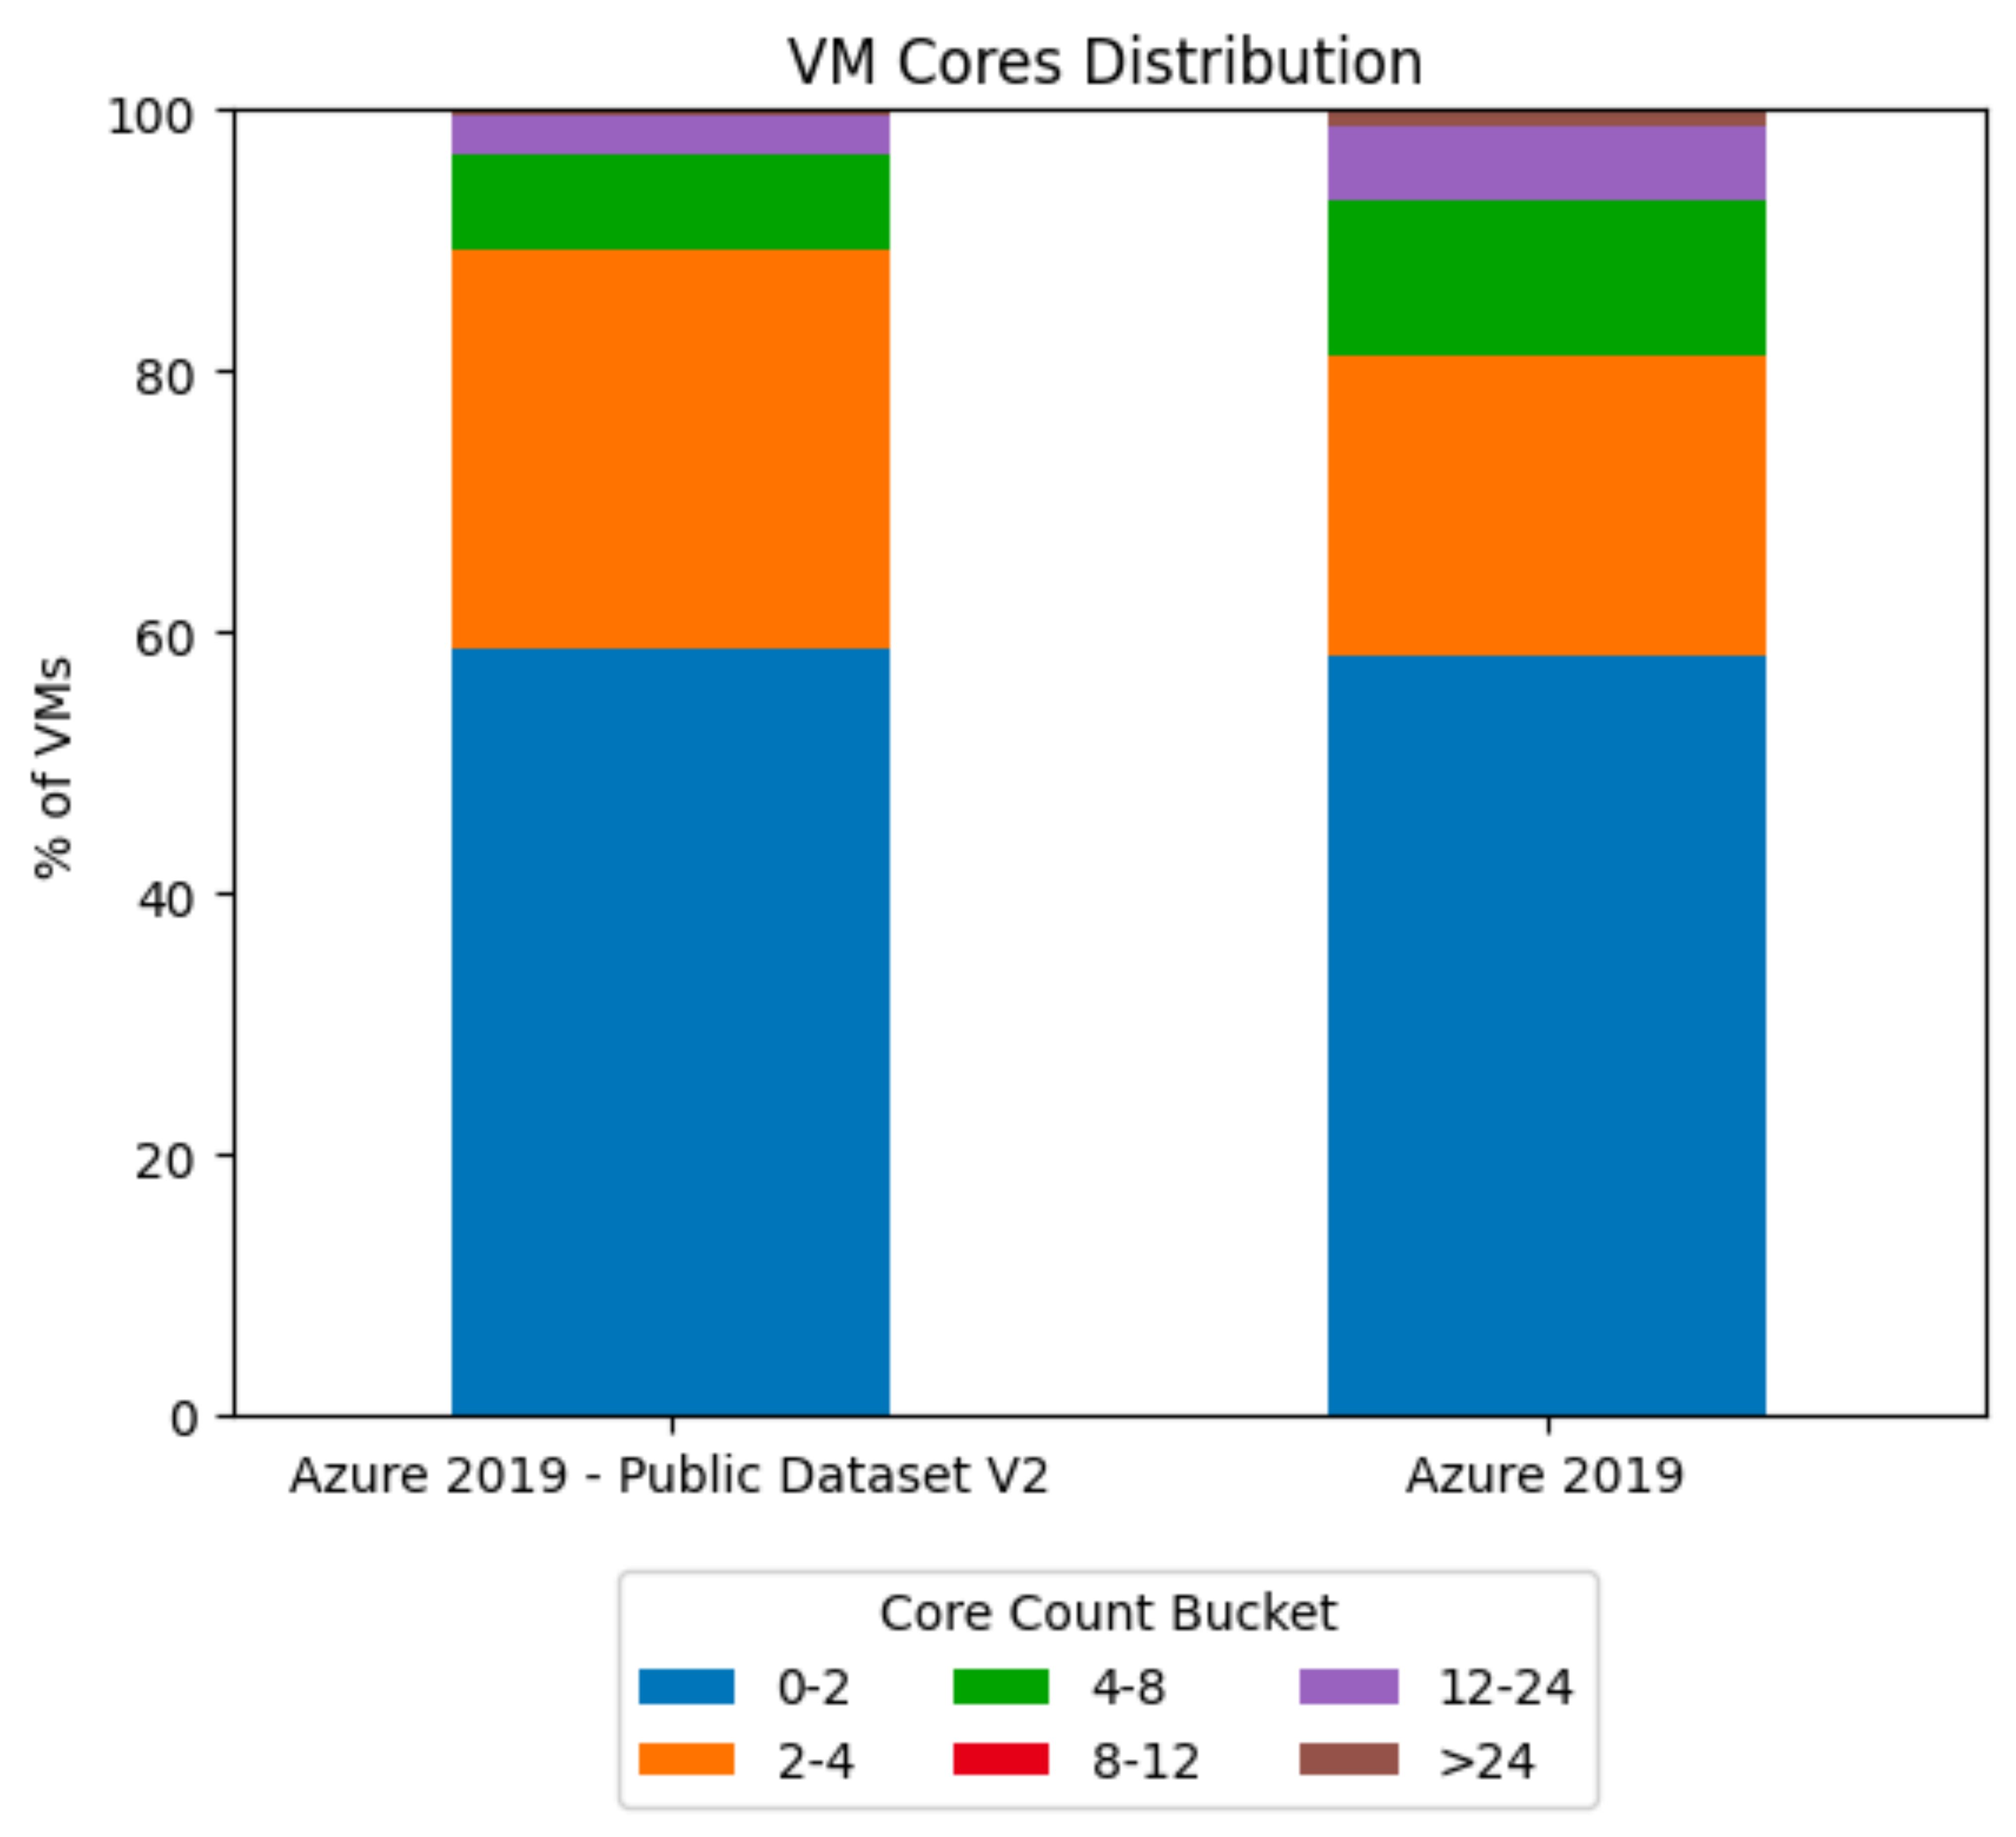

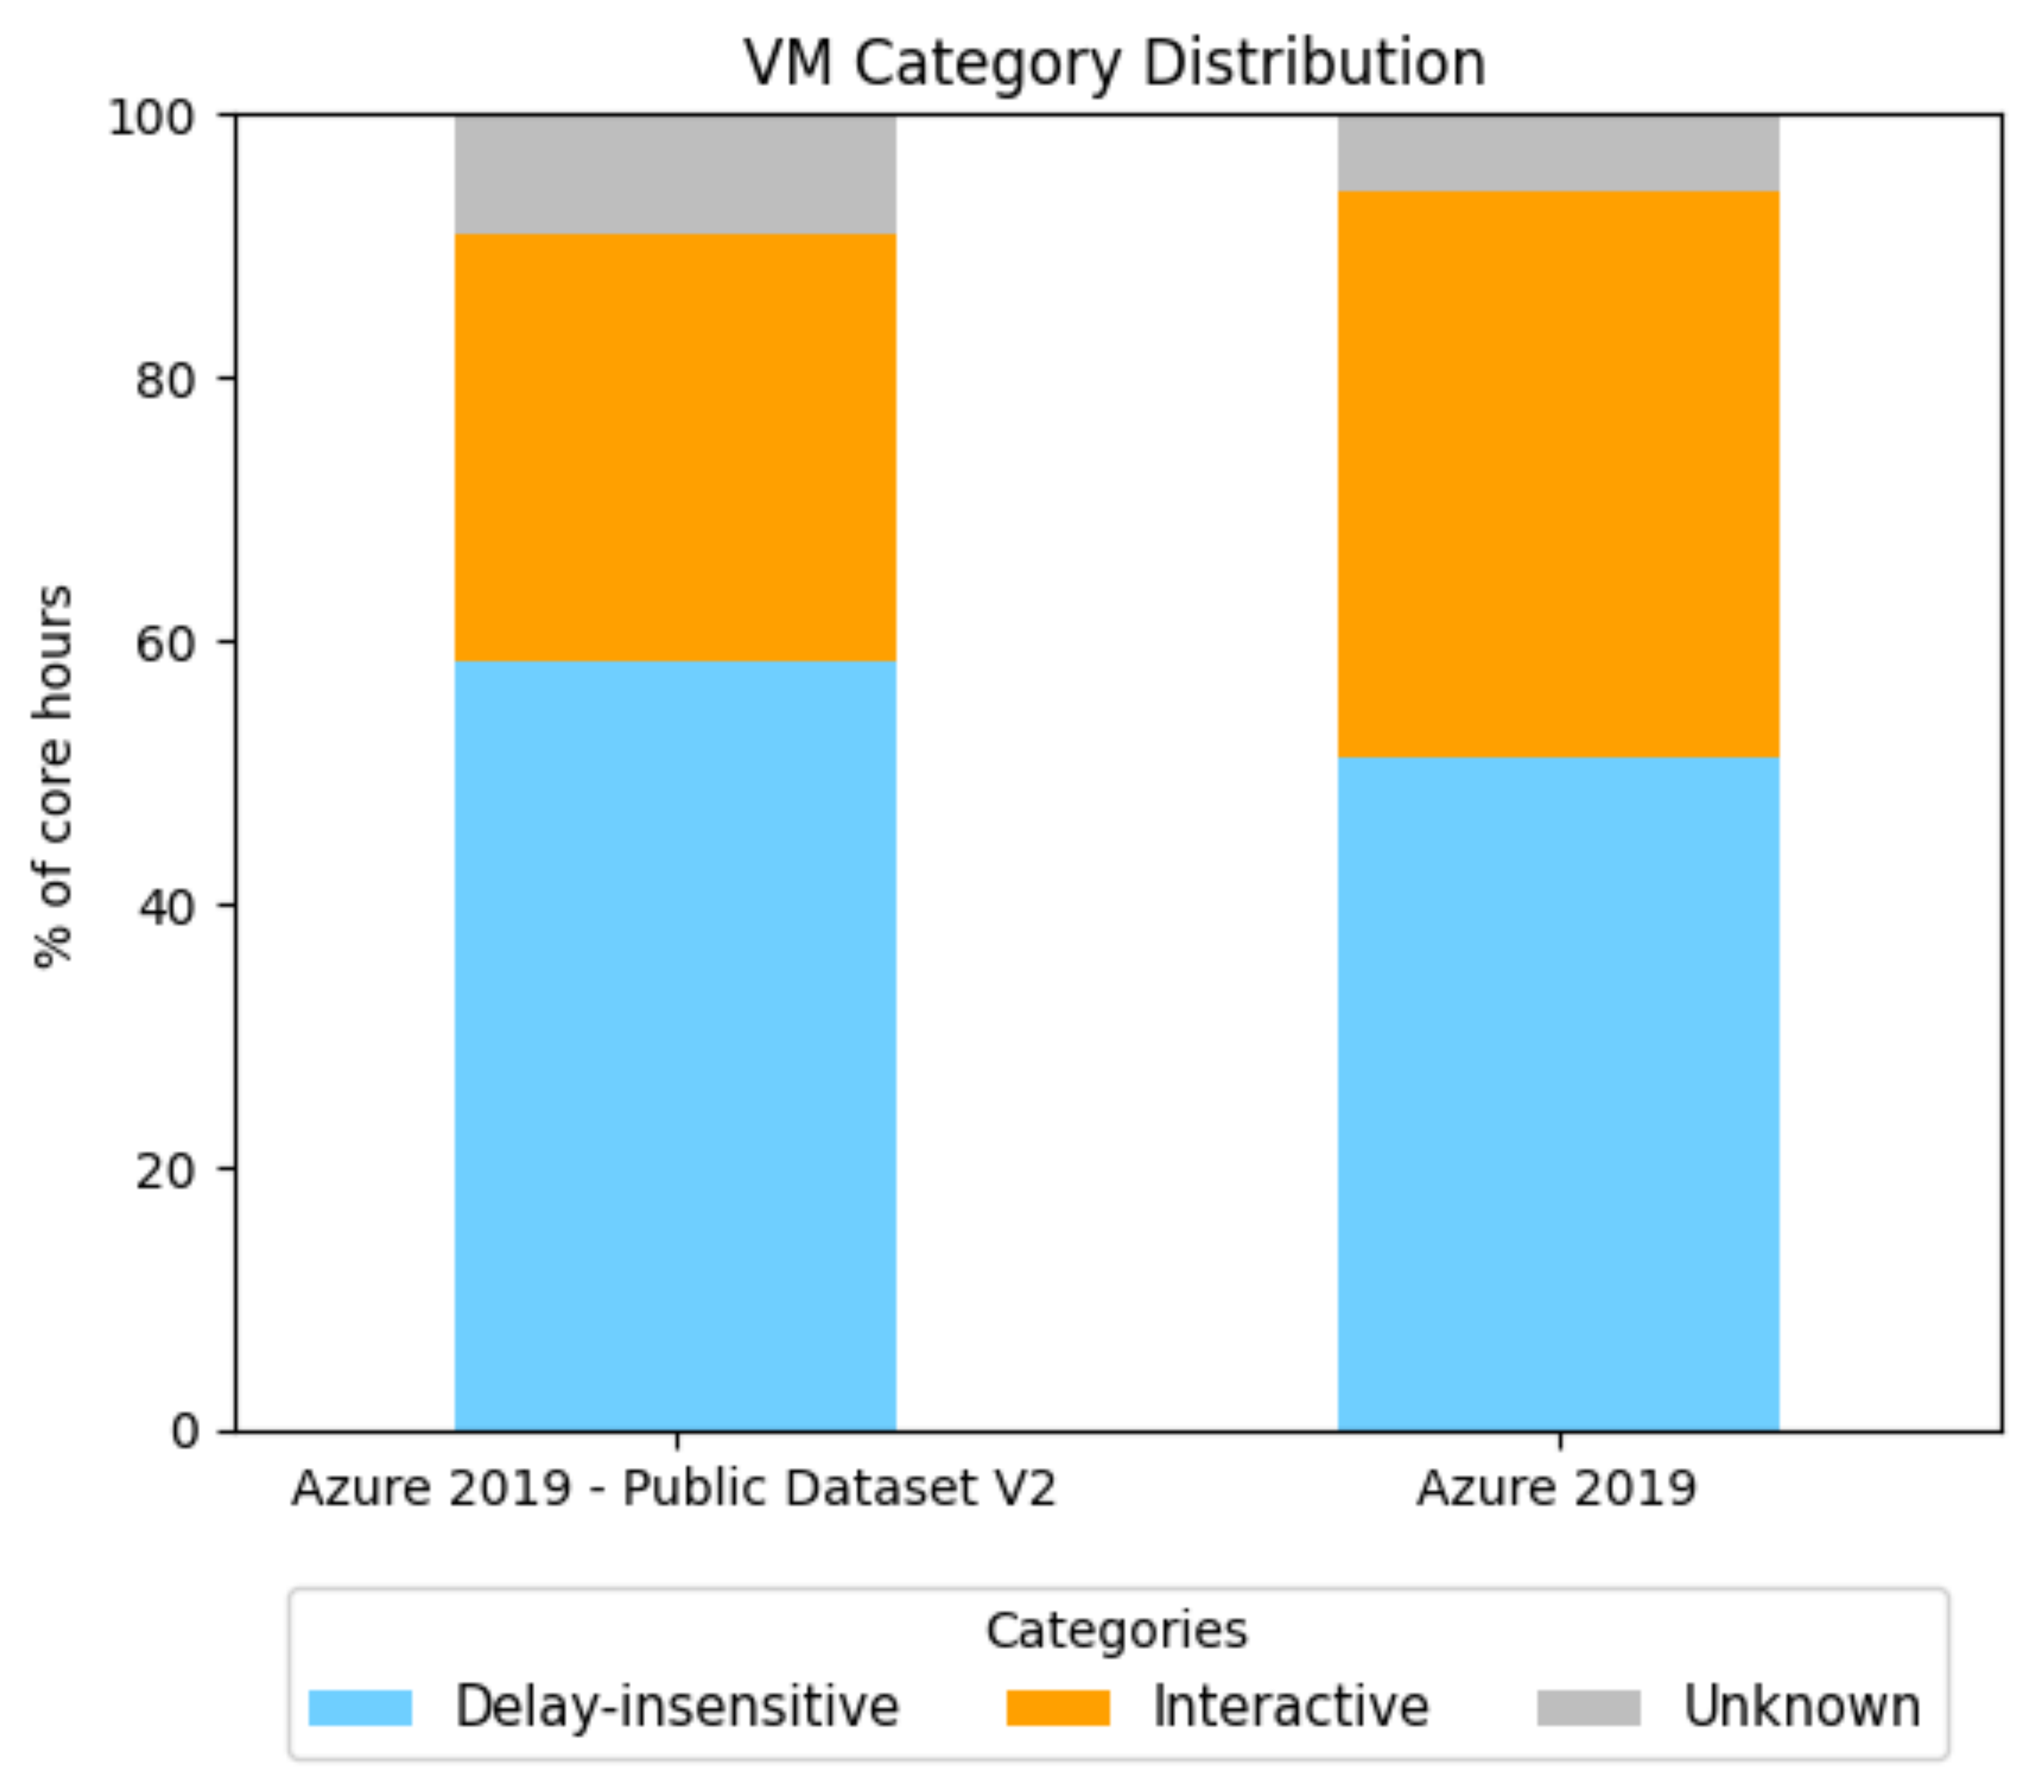

Time-shifting on Azure VM traces

Model: Total Relative Energy Cost (TREC)

$$avgMOER_i = \frac{1}{n} \sum_{starttime_i}^{endtime_i} MOER$$

$$TREC \propto \sum_{i}^{VMs} avgMOER_i \times lifetime_i \times avgCPU_i \times coreCount_i$$

$$avgMOER\_optimum_i = \min_{w \epsilon window} \frac{1}{n} \sum_{starttime_{i+w}}^{endtime_{i+w}} MOER$$

Time-shifting on Azure VM traces

Results

Start Date	City	Shift Window	Block Shift	Hourly Shift
2019-08-23	Chicago	24h	0.46%	1.23%
2019-10-19	WACM	24h	0.01%	0.64%

- Unit: (MOER) x hour
- e.g. TDP of Intel Xeon Platinum 8272CL: 195W

Future Work

- Scale the test on more regions and dates, and potentially larger window.
- Consider location shifting workload and combined shifting.

Thank you