Energy Efficiency and Sustainability in Data Centers

Simon Harrison-Michaels
University of Minnesota Morris
Morris, MN, USA

Studies to be focused on

This presentation will focus two different products designed to improve energy efficiency and sustainability in data centers:

FootPrinter was designed in 2024 at the University of Vrije in Amsterdam.

PACT (Per Application Turbo Controller) was designed in 2020 in a collaboration between researchers at Stanford University and Google

Sacheendra Talluri (Vrije)

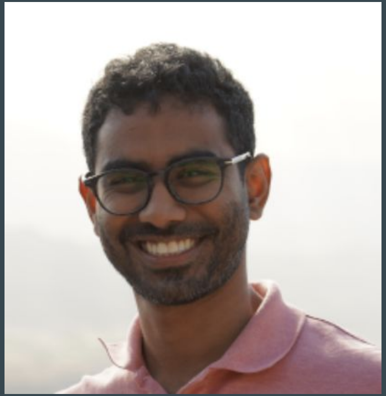

Christos Kozyrakis (Stanford)

Functions of a Data Center

- The most basic function of a data center is to store information.
- Data centers contain servers that store and manipulate data.
- Similarly, data centers require cooling systems to keep the servers at a productive temperature.

Aisle in a data center

Energy Expenditures in a Data Center

Energy expenditures in data centers can be divided into two categories:

1. *Energy required for IT equipment*
2. *Energy required for the cooling system*

IT equipment and cooling both individually require about 40% of the total energy requirement of the data center. The other 20% goes towards things like storage drives and lighting.

Typical energy consumption of a data center:

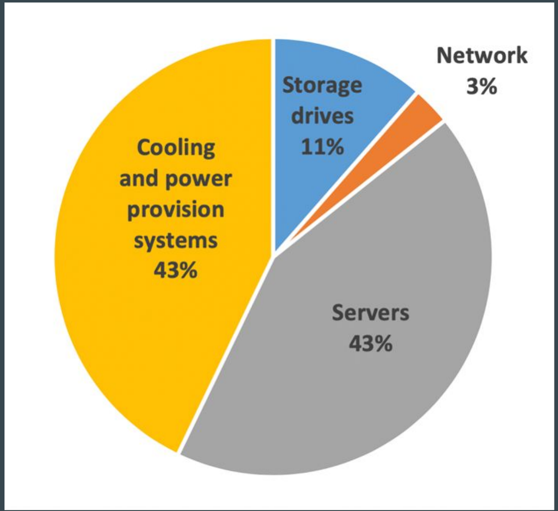

Cooling Systems (the basics)

- Cooling systems use HVAC/air conditioning or liquid cooling. Often they are a combination of both.
- Liquid cooling systems are more efficient than HVAC/air cooling systems
- Two types of liquid cooling systems - direct to chip cooling systems and immersion cooling

Direct to chip cooling

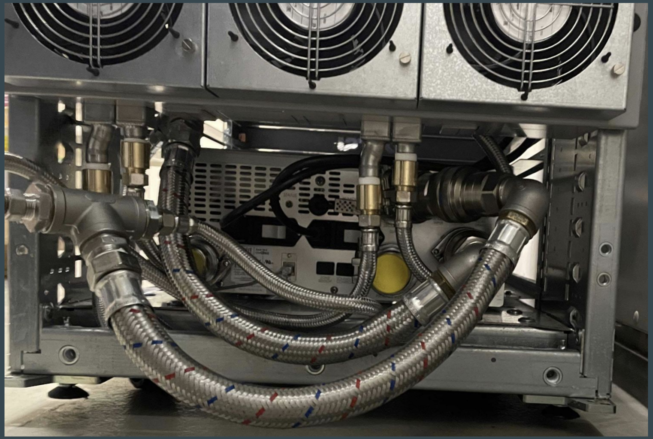

Immersion Cooling

Energy Efficiency Measurement Tools:

PUE: Power Usage Effectiveness

Power Usage Effectiveness is a common measurement to evaluate what percentage of energy delivered to a data center is also delivered to its IT equipment. The formula for PUE is (total facility energy usage) / IT Equipment Energy Usage. A PUE value of 1.2 or less is considered to be good.

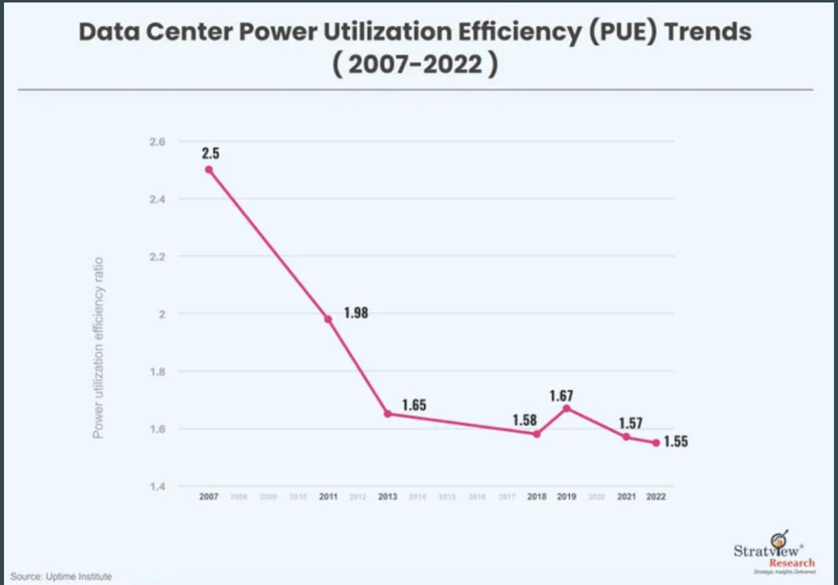

Average PUE has been declining since 2007

Software Carbon Intensity

- Unlike power usage effectiveness, software carbon intensity takes into account the energy's supply chain - the source of the energy itself
- Or in other words, a data center that uses coal for energy would have a higher carbon intensity than a data center that uses solar

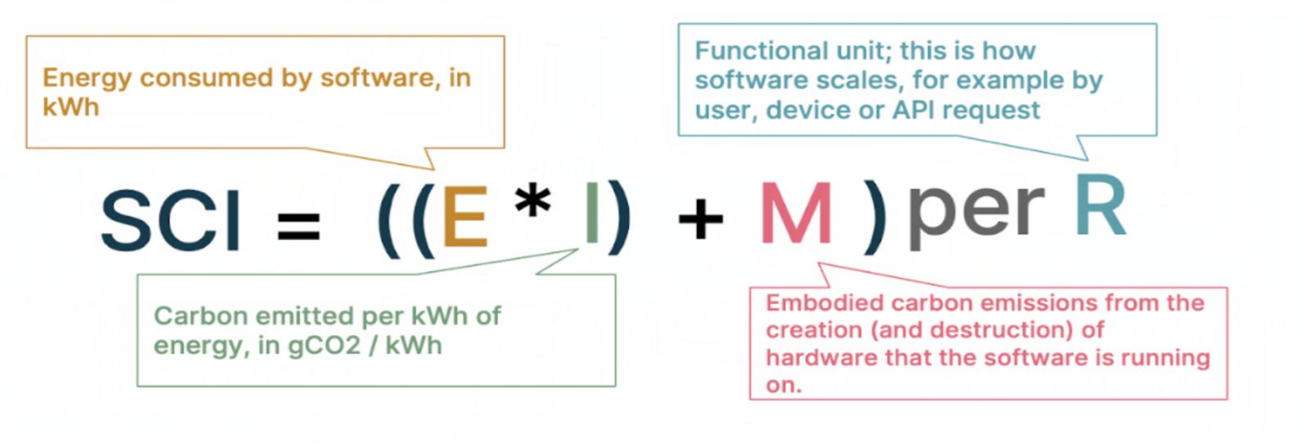

Problem Statement of FootPrinter

- PUE and carbon intensity metrics are not useful when it comes to simulating future energy efficiency. There is not a standardized way to simulate energy demand into the future.
- Because of this, a data center that wants to implement big changes - like a new cooling system or a new server configuration - will often be stuck using a trial and error method to determine energy efficiency. This is inefficient from every perspective.
- To combat these problems, researchers at the University of Vrije in Amsterdam designed FootPrinter. FootPrinter can broadly be described as a parameter-based energy efficiency simulator for data centers.

Parameters of FootPrinter

1: Workload Traces. Workload traces in a data center are similar to what a process is on an operating system. When specific tasks are done in a data center, data about each task is stored - timestamps of start/end time, and resources usage (CPU, memory).

2: Hardware and Environment Specifications.. Hardware and environment specifications include things like the location of the data center, the type of cooling system used, the hardware in the IT equipment itself. Similar to carbon intensity, hardware and environment specifications take into account the supply chain of the energy.

3: Operational Techniques. Operational techniques involve things like scheduling and resource allocation policies. Operational techniques could be described as all the human-made constraints made on the hardware/software of a data center.

FootPrinter Functionality Diagram

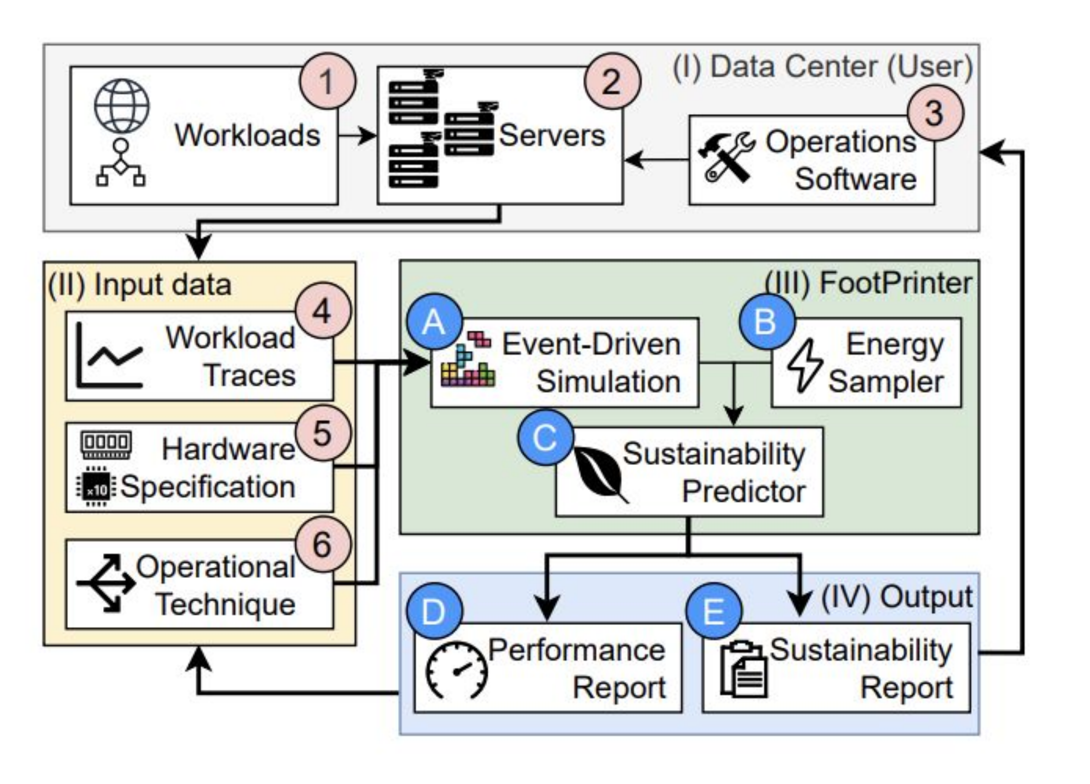

- Data from workload traces, hardware specifications, and operational techniques is combined with “energy sampler”.
- The energy sampler is the carbon intensity value of the incoming energy
- This information is combined together to produce two reports - a performance report and a sustainability report

FootPrinter Analysis

- This is a graph of the performance report power consumption simulation overlaid with the real-time power consumption data.
- The performance report was able to simulate the real power consumption within 1 kW of accuracy
- Although 1 kW is quite a bit - enough to run an average vacuum cleaner for an hour - that is nonetheless an impressive feat

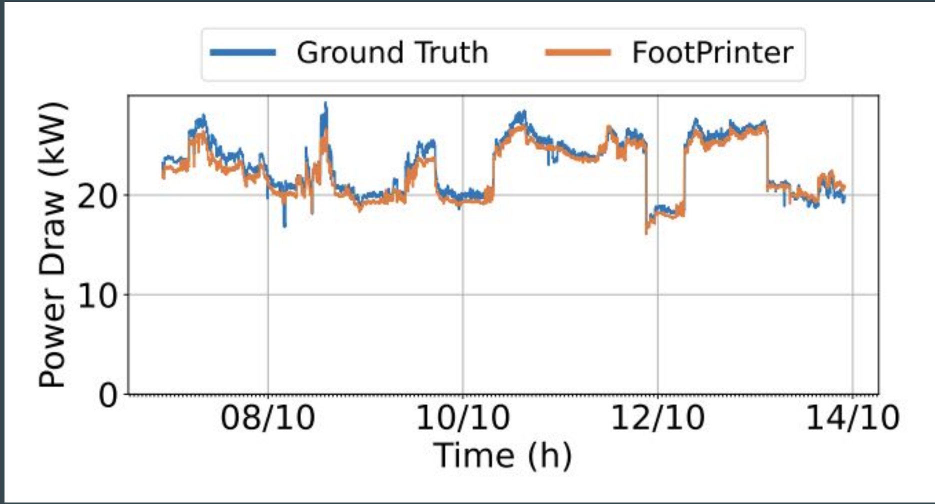

Demand Based Switching (DBS)

- The purpose of demand based switching is to minimize the number of idling CPU cores
- An important idea to understand DBS is the idea of idling CPUs. Data centers need to keep a certain percentage of available CPUs free (or idling) in order to be able to handle incoming requests
- DBS systems manipulate the clock speed and voltage provided to each CPU in accordance with the necessary workload at a specific time.

Multiplicative Saving

- Multiplicative Saving is the idea that saving power at one level of the data center can have a *multiplicative* effect and save power in the data center as a whole.
- For example, CPU requires less energy to do a task, then the cooling system will subsequently need less power.
- Saving 1 watt of power at the CPU level can save around 2.5 watts in total power consumption of the data center.

Multiplicative Saving is similar to a domino effect

Latency Sensitive vs Best Effort Tasks

- Latency sensitive tasks deal with information where delays would create safety, security, or monetary issues in the real world.
- An easy example of latency sensitive tasks are self-driving cars. Self-driving cars communicate with data centers to receive real-time info about the road beyond the radar of the self-driving car.

- Best effort tasks are any tasks that are not latency sensitive tasks. An example of a best effort task is sending an email. A sent email is allowed to have a slight delay in delivery.

PACT (Per Application Turbo Controller)

- In search of a more efficient DBS software system, researchers at Stanford University and Google collaborated to make PACT (Per Application Turbo Controller).
- PACT is made up of two components - Turbo Controller and CPUJailing.
- The ultimate goal of PACT is to minimize CPU energy consumption
- Minimizing CPU energy consumption will maximize the effect of multiplicative saving
- This will reduce energy consumption of the data center as a whole

How PACT works

Turbo Control:

- Turbo Control is an advanced Demand Based Switching (DBS) system.
- When demand is high, PACT will initiate Turbo Control on a certain amount of CPU cores. This means that the clock speed and voltage of the CPU will be altered. Then, more tasks can be undertaken.
- This is important because Turbo Control reduces the potential for latency issues, which means that less CPUs need to idling. This reduces power consumption at the CPU level, which will result in increased multiplicative saving

How PACT works (cont.)

CPUJailing

- Turbo Control means you don't require as many idling CPUs. This is where CPUJailing comes in.
- CPUJailing does more or less it's name - puts certain CPUs in a "jail" (i.e. reduces their energy).
- However, CPUJailing identifies CPUs by dividing tasks into two categories - Latency Sensitive (LS) and Best Effort (BE). It prioritizes Latency Sensitive tasks over Best Effort tasks, and therefore will implement *more* demand based switching on CPU cores running Best Effort tasks vs Latency Sensitive Tasks.

1 CPUJailing Controller

```
procedure CPUJailing :
    ls_jail ← ∅
    be_jail ← ∅
    for socket in sockets do
        ls_util ← ceil(get_ls_util(socket))
        ls_util ← ls_util + ls_buf
        be_util ← ceil(get_be_util(socket))
        num_cores ← get_num_cores(socket)
        if be_util + ls_util ≤ num_cores then
            ls_socket_mask ←
                get_N_CPUs(socket, ls_util)
            be_socket_mask ←
                socket_mask \ ls_socket_mask
        else
            ls_socket_mask ← socket_mask
            be_socket_mask ← socket_mask
        ls_jail ←
            ls_jail ∪ ls_socket_mask
        be_jail ←
            be_jail ∪ be_socket_mask
    for task in ls_tasks do
        task.cpumask ←
            task.cpumask ∩ ls_jail
    for task in be_tasks do
        task.cpumask ←
            task.cpumask ∩ be_jail
```

PACT analysis

- The two separate components of PACT - Turbo Control and CPUJailing - both individually reduced power compared to the baseline, and more so when combined
- Cluster A and Cluster B each have more than 10,000 machines of different hardware platforms

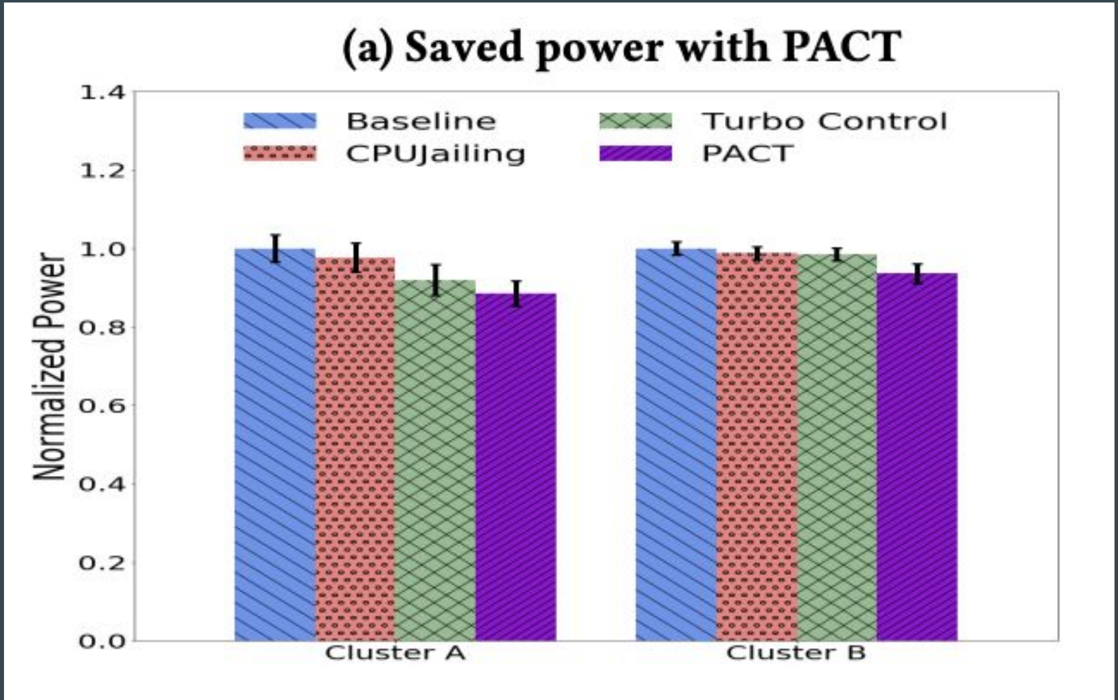

Total Power Consumption (watts) of Clusters A + B

Figure 3: Cluster characteristics.

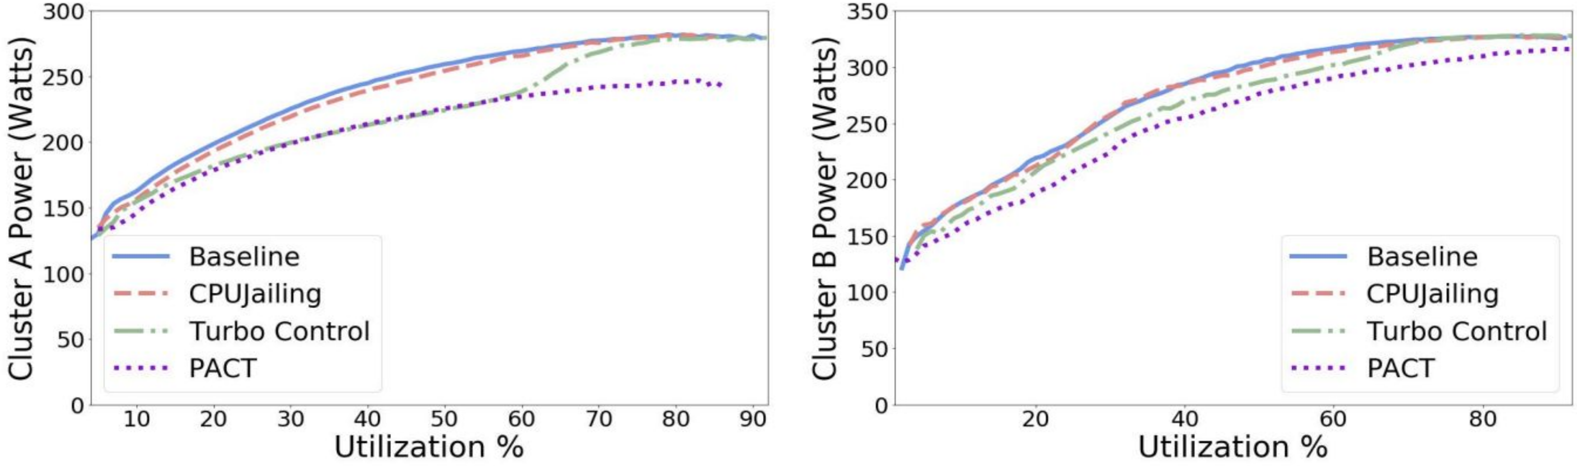

(a) Cluster A

(b) Cluster B

Real world connection

One reason why this topic is generating interest currently is because of increased demand for data centers in the future. In 2022, the IEA reported data centers consumed 460 TWh of electricity (about 2% of the world's energy usage).

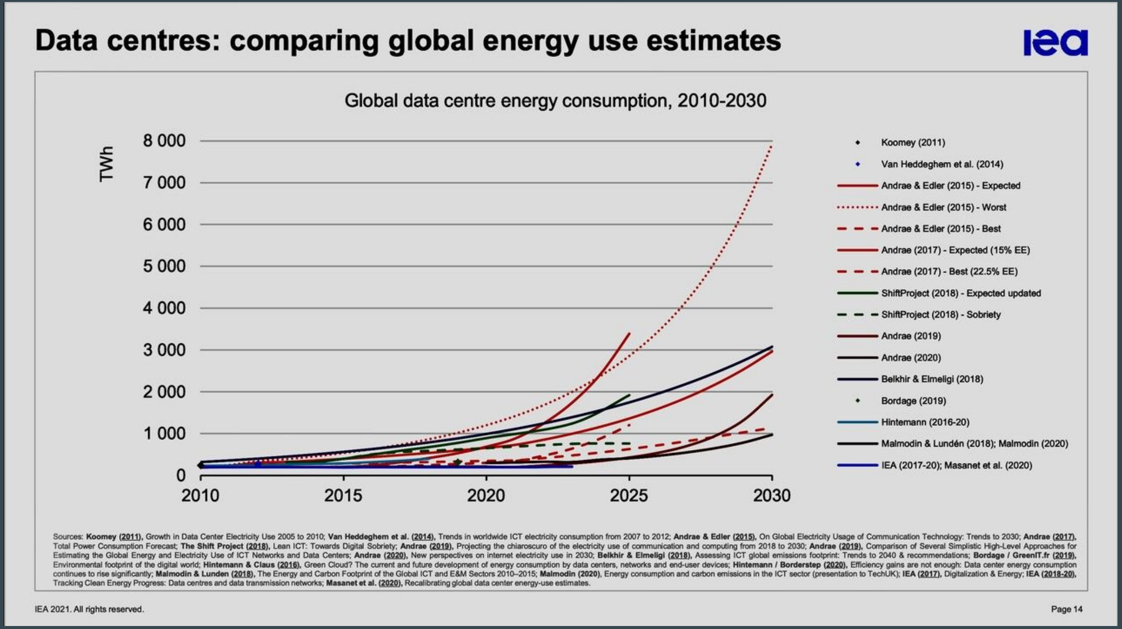

Different predictions of future global data center energy consumption layered on same graph

Citations

[1] Kaffes, K., Sbirlea, D., Lin, Y., Lo, D., and Kozyrakis, C. (2020). Leveraging application classes to save power in highly-utilized data centers. In Proceedings of the 11th ACM Symposium on Cloud Computing, SoCC ’20, page 134–149, New York, NY, USA. Association for Computing Machinery.

[2] Khan, W., De Chiara, D., Kor, A.-L., and Chinnici, M. (2022). Exploratory data analysis for data center energy management. In Proceedings of the Thirteenth ACM International Conference on Future Energy Systems, -Energy ’22, page 571–580, New York, NY, USA. Association for Computing Machinery.

[3] Niewenhuis, D., Talluri, S., Iosup, A., and De Matteis, T. (2024). Footprinter: Quantifying data center carbon footprint. In Companion of the Figure 4. Projected Data Center Electricity Usage (taken from) 15th ACM/SPEC International Conference on Performance Engineering, ICPE ’24 Companion, page 189–195, New York, NY, USA. Association for Computing Machinery.

[4] Piatek, W., Oleksiak, A., and vor dem Berge, M. (2015). Modeling impact of power- and thermal-aware fans management on data center energy consumption. In Proceedings of the 2015 ACM Sixth International Conference on Future Energy Systems, e-Energy ’15, page 253–258, New York, NY, USA. Association for Computing Machinery.

[5] Raje, S., Maan, H., Ganguly, S., Singh, T., Jayaram, N., Ghatikar, G., Greenberg, S., Kumar, S., and Sartor, D. (2015). Data center energy efficiency standards in india. In Proceedings of the 2015 ACM Sixth International Conference on Future Energy Systems, e-Energy ’15, page 233–240, New York, NY, USA. Association for Computing Machinery.

Questions...