Strategies for Minimization of Energy Consumption in Data Centers

Friis, Simon; Larsen, Line Maria Pyndt; Ruepp, Sarah Renée

Published in:
Proceedings of Twenty-Second International Conference on Networks

Publication date:
2023

Document Version
Peer reviewed version

Citation (APA):
Friis, S., Larsen, L. M. P., & Ruepp, S. R. (2023). Strategies for Minimization of Energy Consumption in Data Centers. In *Proceedings of Twenty-Second International Conference on Networks* (pp. 17-22)

Strategies for Minimization of Energy Consumption in Data Centers

Simon Friis*, Line M. P. Larsen*† and Sarah Ruepp*
* Department of Electrical and Photonics Engineering, Technical University of Denmark, Kgs. Lyngby, Denmark
† TDC Net, Copenhagen, Denmark
Corresponding author: email: s175405@student.dtu.dk

***Abstract*—Data Centers (DCs) consume a significant amount of energy, making overall energy consumption a major concern. The scale of a DC affects its energy efficiency, with larger centers having more resources for energy-saving measures but at the same time different challenges than those faced by DCs of smaller scale. This work analyses and compares energy efficiency of small, medium, and large DCs, analysing factors such as resource and server utilization and design. Finally, the energy minimisation techniques are evaluated for their potential impact on DC energy consumption, as well as their contribution to DCs of different sizes.**

Keywords-Data center; cloud; green; low power; scalability; energy efficiency.

I. INTRODUCTION

Energy efficiency in Data Centers (DCs) is a crucial topic of modern DC operations, as it can help to reduce energy costs and environmental impact of the Information and Communications Technology (ICT) sector. DCs are energy-intensive facilities that consume a large amount of electricity to power servers, storage systems and cooling equipment. The energy consumption of DCs has become an increasing concern for the industry, as well as businesses and organizations that operate these facilities, as DCs are responsible for approximately 1.5% [1] of global carbon emissions. With the increase in data volume, DCs will consume more energy, thus; there is a need to find new and more efficient ways to run them.

Two types of DCs include private entreprise DCs and public cloud DCs. As illustrated in Fig. 1, the end-users gain access to the DCs to store and process data through a network of computers, wireless Access Points (APs), switches/routers, firewalls, and the Internet. The computers are connected to both data centers through a switch or router, which directs the data traffic between them. An enterprise DC is located inside the same local network as the users, while a cloud DC is located outside of the local network. Cloud DCs are typically managed and owned by third-party service providers and the services they provide are accessed through the Internet. Users can access both types of DCs by authenticating themselves and then proceed to transfer data through the nodes in the network. Fig. 1 gives an overview of a basic network architecture where users have access to both an enterprise DC and a cloud DC. The benefits of connecting to both a cloud DC and an enterprise DC for data access and exchange count:

- The cloud DC enables remote, on-demand access to data and application services from other providers through the Internet.
- The enterprise DC provides more secure, local access to other types of data and applications, which is beneficial for vulnerable data.

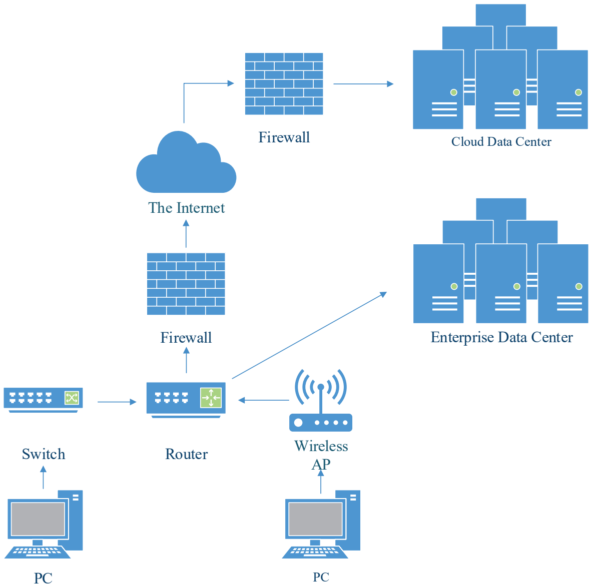

Figure 1. A basic network with computers connecting to both a cloud data center and an enterprise data center.

The architecture of a DC plays a crucial role in its overall energy efficiency. Several components make up a typical DC architecture, including [2]:

- *Server Racks:* The servers themselves, as well as the physical stations that house the servers in a DC and consume energy for processing and cooling. Server racks are designed to organize, store and manage numerous servers, while optimizing floor space at the same time.
- *Top of the Rack (ToR) Switches:* Switches connected to every server in a server rack and connects those to the network.
- *Aggregation Switches:* A centralized connection point for assigned ToR switches. Responsible for collecting data traffic from multiple servers and forward it.
- *Load Balancers:* Devices responsible for distributing network traffic evenly across several servers. Reduces the probability of network failures by lowering workload of overwhelmed servers.
- *Access Routers:* A secure connection point for external network traffic.

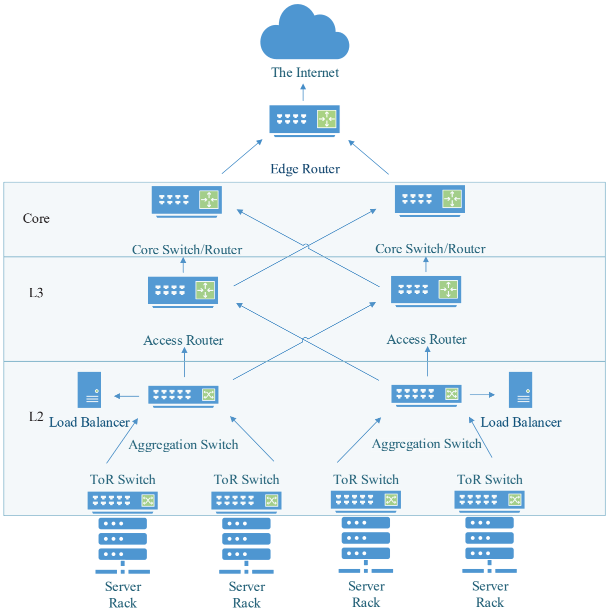

Figure 2. Key components of a basic data center architecture separated into different layers.

- *Core Switches/Routers:* Devices responsible for forwarding traffic at a high speed within nodes of a DC network.
- *Edge Routers:* Handles incoming and outgoing network traffic by routing data from and to the DC.

These devices cooperate to distribute, forward and transmit data traffic stored in a DC, and understanding the purpose and role of each device is key when optimizing energy efficiency in DCs. Fig. 2 illustrates the basic elements of a data center architecture, designed to efficiently process and manage data. It includes server racks connected to ToR switches which forwards traffic to the aggregation switch. The data is then directed to the load balancer which distributes the traffic between the servers. The access router controls access to the DC network and the core switch/router forwards to the edge routers which serves as the bridge between the internal DC network and the external network, being the Internet.

There are many different types of DCs and they will therefore be categorized into three different sizes in this paper ranging from small, medium and large.

- *Small DCs* are categorized as usually having less than 1,000 servers, as well as less complex infrastructure, limited storage and consume less power compared to larger DCs.
- *Mid-scale DCs* have a larger number of servers and usually range between 1,000 to 10,000 [3] with more complex infrastructure.
- *Large DCs* are defined as usually having more than 10,000 servers with an even more complex infrastructure [3]. Large DCs are typically used by large companies or governments.

This paper investigates methods for energy efficiency in DCs, with a focus on state-of-the-art technologies and techniques, as well as how and why these are beneficial for DCs of different scale. Section II presents other related papers and features our contribution to the topic. Section III goes in-depth with modern and commonly used strategies. Section VI discusses what and why some methods are most commonly used in DCs of certain sizes and their potential. Finally, the conclusion closes this paper.

II. RELATED WORK

There have been several studies and research conducted on energy efficiency in DCs in recent years. However, numerous surveys regarding energy efficiency in DCs tend to be older than 5 years, such as the work in [4], which presents an overview of energy-aware resource management approaches with focus on basic architecture of cloud DCs and virtualization technology. The survey in [5] investigates the green energy aware power management problem for Megawatt-scale DCs and classifies work that considers renewable energy and/or carbon emission. In [6], they discuss several state-of-the-art resource management techniques, that claim significant improvement in the energy efficiency and performance of ICT equipment and large-scale computing systems, such as DCs. In [7], they conduct an in-depth study of the existing literature on DC power modeling, covering more than 200 models.

Recent related work includes [8], where the approaches moving towards green computing are investigated and categorized to help researchers and specialists in cloud computing expand green cloud computing and improve the environment quality. The work in [9], gives a brief overview of the state-of-the-art in green cloud computing. They examine existing research in the area and categorize it into different themes. They also discuss the challenges and opportunities in the field, and provide insights into future directions for research. This paper provides valuable background information and a significant understanding of the current landscape of green cloud computing. In the survey [10], the authors discuss different mechanisms for lowering the power utilization in DCs. It provides in-depth details about the various mechanisms that can be employed at the hardware level so that the utilization of energy by component can be reduced. Table I lists relevant research in the field of energy improved DCs categorized into relevance regarding the different DC sizes.

A. Our contribution

Our work contributes to the field by exploring numerous strategies for improving energy efficiency in DCs, while taking the different sizes of DCs into consideration. We aim to provide an in-depth overview of key challenges, opportunities and methods for improving energy efficiency in all types of DCs. The contributions of this paper are not only an overview of current research directions but also an overview of how proposed technologies and techniques can be realised in modern DCs, including:

- Provide insights into effective ways to improve energy efficiency in DCs by synthesizing the state-of-the-art technologies and techniques for DCs of different sizes.

TABLE I. DC SIZE REFERENCES

DC Size	References
Large scale	[1] [2] [4] [5] [6] [7] [8] [10] [11] [12] [13] [14] [15] [16] [17] [18] [19]
Mid scale	[2] [7] [8] [10] [13] [14] [15] [16] [17] [19] [20]
Small scale	[2] [8] [10] [13] [14] [15] [16] [17] [19] [20] [21] [22]

- Analyzing previous surveys and papers on energy efficiency in DCs, and provide a overview of their key findings and limitations. Our work highlights that there has barely been performed any research on how effective certain energy efficiency strategies have been for different size DCs.

III. ENERGY MINIMIZATION METHODS

Energy minimization methods refer to the numerous strategies used to reduce the energy consumption of DCs with the goal of minimizing energy consumption while maintaining high levels of performance and reliability. Server utilization in DCs are found to be under 20% most of the time and with the servers still running fully, this results in very low energy efficiency since servers still consume a significant amount of energy even when not fully utilized [13]. A common tool for measuring energy efficiency in DCs are the Power Usage Effiectiveness (PUE) metric. It is calculated by dividing the total amount of energy used by the DC, including all systems and components, by the energy used by the IT equipment within the DC [23].

This section will give a brief overview of multiple state-of-the-art technologies and techniques as well as going in-depth with some subcategories of these strategies, being: sleep state methods and resource utilization in their own subsections. Fig. 3 illustrates where in a DC certain methods are utilized and what components are involved by highlighting the energy efficiency strategies with different colors: Green for Load Balancing and Scheduling, dark blue for cooling systems optimization, red for DCIM tools and yellow is for Sleep State methods.

A. Trending methodologies

Energy efficiency in DCs can be achieved through a variety of strategies. Examples of current research directions are:

- Advanced cooling systems
- Server virtualization
- Data Center Infrastructure Management (DCIM) tools
- Edge computing
- AI-driven DC Management
- Quantum computing

Advanced cooling systems are innovative technologies used for mainly cooling servers and can result in notable energy savings. Liquid cooling, free cooling and indirect cooling are some of the advanced types of cooling systems [22]. Another relevant method in this category is *heat re-use*, which refers to the process of utilizing waste heat generated from one process or system and using it for another purpose, rather than letting it go to waste. This results in energy savings and reduced carbon emissions [24].

Server virtualization can lower the number of needed servers in a DC by running multiple virtual servers on a single physical server, resulting in lower power consumption [14].

DCIM tools monitor, measure, manage and/or control data center utilization and energy consumption of DC equipment such as servers, storage systems and network switches/routers. This helps identify power-related issues and improve DC performance and energy efficiency [25].

Edge computing refers to a range of networks and devices at or near the users and enables processing data closer to where it is being generated. Edge computing can reduce the amount of data traffic that needs to be transmitted to a central DC, resulting in potential energy savings [15].

AI-driven DC Management is a method for automating control and monitoring of DC resources. By improving DC operations, energy efficiency improves as well [26].

Quantum computing have the potential to increase energy efficiency in DCs by solving complex problems at an incredible speed compared to traditional computing methods. However, it is a new technology and is still in its early stages [27].

B. Sleep states

Sleep states can be implemented to shut down several server components for a short period of time to reduce energy wasted on un-used server capacity. Fig. 3 illustrates what components in a DC that can be impacted by sleep state methods. When utilizing sleep state methods, the components that are being powered down are the Central Processing Unit (CPU), cores of the CPU, memory and storage devices [13]. The devices and nodes that are involved when utilizing this method are highlighted in Fig. 3, marked by yellow. Modern processors support multiple types of sleep states, primarily:

- Core C-states
- Package C-states
- P-states
- DRAM (Dynamic Random-Access Memory) power mode

Core C-states work by stopping executions on the core. They range from C1-C6 and the differences between those being the varied amounts of power savings and exit latency costs. C0 is the active state, with no CPU power savings. C1 is the state with the least power savings but with the shortest exit latency whereas C6 is having the longest exit latency at a $133\mu s$ transition time [13].

Package C-states are used when all cores are in state C1 to C6, hence; the entire CPU is idle. In this state a whole package of components turns off, such as shared caches, integrated Peripheral Component Interconnect Express (PCIe), memory controllers, and so on [16]. However, the concept is that additional power is saved compared to the power saved with the sub-components individually [16]. Package C-states can significantly reduce energy consumption but has the side effect of increasing the latency for cores going to or from low power states [13]. Furthermore, package C-states can be problematic because of high response times during re-activation when handling traffic spikes. Additionally, having the memory and/or storage of all servers to be available, even during times of light load, can be very beneficial. Lower latency can be achieved by AgileWatts (AW) which is a deep idle core power-state architecture that reduces the transition latency to/from very low power states. AW has been proven to result in up to 71% power savings per core with a less than 1% end-to-end performance decrease [13].

P-states changes the frequency and voltage of a part of the system. This being the cores or other components such as a shared L3 cache [16]. P-state is a module state affecting a collection of cores that share resources [16]. The concept of P-states is that a CPU running at lower frequencies requires lower performance and longer latency to complete a certain amount of work. Thus, under some circumstances, for example in low traffic periods, it is possible to complete a required amount of work with lower energy [16].

The DRAM power mode consists of two power-saving methods which are the Self-refresh function and the Clock Enable (CKE) mode. CKE sends a signal from the memory-controller (MC) to the DRAM device, and when this signal is no longer being sent, the DRAM is free to enter a low power state. The MC is also behind the Self-refresh function as it sends the refresh signal to DRAM to ensure that the data is valid. DRAM does have the ability to start the Self-refresh process itself which can reduce the power consumption in the MC [28].

C. Resource utilization

DCs' load rises when more requests are received, and these requests can be received seasonally. Thus, the workload demands of the servers are changing dynamically and are determined by a real-time workload status. By balancing the load on the servers carefully and properly, it is possible to increase the energy efficiency of components in a DC. Fig. 3 illustrates resource utilization techniques within a DC, highlighted with as green.

1) Load balancing: Dynamic Time Scale based Server Provisioning (DTSP) is a method which takes the variability of workloads into consideration when providing servers for workload demands. For DTSP to load balance properly, key information is gathered constantly so that DTSP can accurately estimate workload requirements on servers and specify the appropriate number of servers for the dynamic workloads [11]. Irregular arrivals of requests impact the accuracy of the expected workloads, so to increase the estimation, the gathered information of incoming requests is standardized before it is used in later calculations. When it comes to workload, the algorithm looks at the three factors; arrival rate of previous requests, the arrival rate of current requests and the mean service time of current requests. With these factors, the algorithm is able to figure out the intensity of previous workloads and reflect the available remaining capacity for the unfinished waiting workloads, as well as measure the intensity and time needed for current workloads to complete. These factors are also used when calculating the workload demand of incoming requests and to determine how many servers are needed to finish current and remaining workloads while satisfying the Quality of Service (QoS) requirements [17]. DTSP has been proven to be able to estimate the workload demands of servers in a DC. By periodically adjusting service resources to match workload demands, DTSP significantly improves and maintains the system energy efficiency under an acceptable QoS level [11].

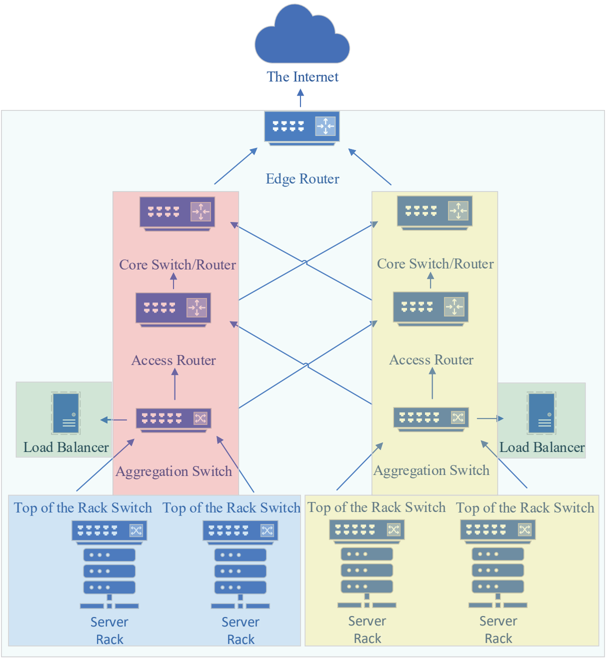

Figure 3. Resource utilization techniques highlighted in color for involved key components of a data center.

2) Scheduling: A cloud system uses virtualization technology to provide cloud resources such as CPU and memory to users in the form of virtual machines. Tasks and job requests are assigned on these VMs for execution. The technique known as job scheduling is a method used to assign a job to a VM based on classification. By allocating jobs based on types and availability, it is possible to increase energy efficiency by making better use of available resources. Minimizing the number of hosts used when allocating resources reduces energy consumption. The Energy Aware VM Available Time (EAVMAT) scheduling algorithm does exactly this [18]. By categorizing jobs into three types and then assigning jobs based on a predefined policy with the earliest available resource. Energy consumption is then reduced since less hosts are in a active state and resource utilization is higher. This method has been tested and was able to achieve up to 46% energy savings [18].

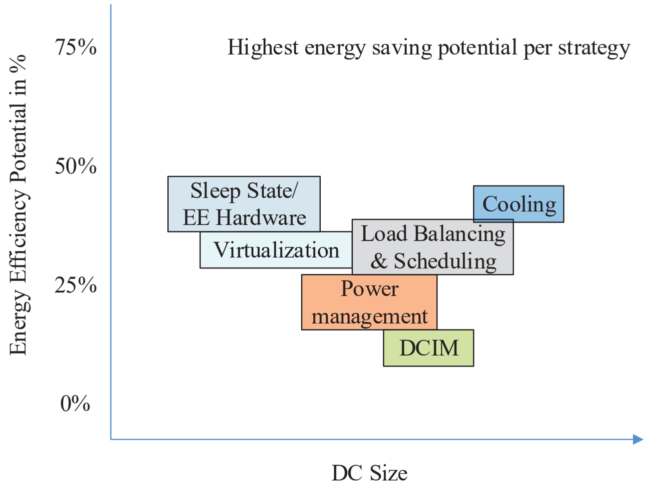

Figure 4. The graph compares the energy efficiency improvement percentages achieved through different energy efficiency strategies in DCs of different scale.

IV. DISCUSSION

DCs being responsible for approximately 1.5% of global carbon emission with an annual growth rate of 4.3% have become an area of focus within the last decade [20]. However, most attention has been directed towards the largest DCs, which are only responsible for a small portion of the overall energy consumption of DCs in general, since small-/mid-scale sized DCs being responsible for approximately 50% of the energy consumption [20]. Due to the increased attention, large-scale DCs have therefore advanced more than small-scale DCs and have numerous energy efficient methods implemented already. It is shown in [20], that energy efficient strategies such as virtualization are adopted less in smaller DCs compared to large DCs. Small DCs are in general behind on the energy efficiency front with around 43% of them not having energy efficiency objective in place at all [20]. The benefits of different strategies used for energy efficiency in DCs varies depending on the size of the DC. Below is recommended a set of guidelines for optimizing energy efficiency in DCs and evaluated based on three different sizes/categories. Techniques and technologies recommended for small-scale DCs are also excellent methods for larger DCs, whereas methods recommended for large-scale DCs are not always realistic/beneficial options for smaller DCs because of price and other circumstances. On the other hand, small-scale DCs might see greater gains when utilizing some of these strategies, since they are size-wise easier to manage, which can result in energy-efficient technologies and practices being adopted more easily. Large DCs managing thousands of servers and hundreds of server racks will likely achieve greater power savings by investing in advanced cooling systems than small-scale DCs managing less than hundred servers. Here, small-scale setups might see greater benefits investing in other technologies and techniques.

A. Strategies at Different Scale

Many different factors are decisive for how effective certain strategies are when it comes to the energy efficiency for data centers of various sizes. This makes it difficult to generalize the different methods as all DCs differ in relation to infrastructure, scale and utilization, environmental factors and what energy efficient technologies are already in place. Some energy efficiency strategies can provide the best results for smaller DCs compared to larger DCs, since small-scale DCs have fewer resources available as well as generally not even having implemented any energy efficiency strategies at all [20]. Fig. 4 shows potential power savings of different strategies for varying DC sizes.

Small-scale DCs benefit from energy efficiency strategies such as sleep state methods and power management tools, as well as virtualization, load balancing and energy efficient hardware. Additionally, small-scale DCs can also benefit from design optimization including efficient cooling systems and energy efficient infrastructure.

Large-scale DCs however, have access to more resources and can allocate those towards many different energy-saving measures, including advanced cooling systems, server virtualization, load balancing as well as renewable energy sources. Having access to additional resources opens up for other strategies such as AI-driven DC management and quantum computing as large-scale DCs also have more data traffic to handle. Modern energy efficiency strategies such as advanced cooling systems have proven to potentially achieve energy saving of up to 50% [22], virtualization has proven possibilities of 30% [14], sleep state methods can provide up to 34% energy savings [19], and resource utilization methods can reduce energy consumption by up to 46% [18]. All these strategies are beneficial for DCs of all sizes but can vary in potential energy savings depending on multiple different factors. DCIM and PUE are also excellent methods for working towards more energy efficient DCs and can provide beneficial tools for analysing DCs of all sizes. That being said, as well as being able to utilize and implement the technologies and techniques mentioned for smaller DCs, large-scale DCs does also have other possible methods for achieving greater energy efficiency. AI driven DC management and quantum computing are both methods which will most commonly be seen in large-scale DCs since the owners are able to provide sufficient resources for these technologies to be implemented and these methods are therefore recommended for large-scale DCs, along the methods mentioned for smaller DCs. Edge computing however, might prove to be most beneficial for smaller DCs. Large DCs are much more centralized and have a much greater power density which counteracts the whole principle of edge computing. For smaller DCs it's the complete opposite and operators of such facilities should therefore experience the implementation of this technology as less challenging. These technologies are new and are therefore still being researched, so it is not yet possible to provide potential power savings.

V. CONCLUSION

In conclusion, the scale of a DC can impact its energy efficiency. Small and medium-sized DCs can achieve notable energy-savings by improving design and infrastructure, as well as improving resource utilization, while large-scale DCs can make use of the greater amount of available resources to increase energy efficiency in the same and other ways. This is important since numerous factors impact how energy efficiency is achieved in each DC. This paper highlights the importance of considering DC-scale when estimating the potential impact of energy-saving strategies, as well as suggesting various methods for energy efficiency improvements for DCs of different sizes.

ACKNOWLEDGMENT

Support from Innovation fund Denmark, through grant no. 1045-00047B and the Nordic University Hub on Industrial IoT, Nordforsk grant agreement no. 86220, is gratefully acknowledged.

REFERENCES

[1] S. Mustafa *et al.*, "Performance evaluation of energy-aware best fit decreasing algorithms for cloud environments," *Proceedings - 2015 Ieee International Conference on Data Science and Data Intensive Systems; 8th Ieee International Conference Cyber, Physical and Social Computing; 11th Ieee International Conference on Green Computing and Communications and 8th Ieee International Conference on Internet of Things, Dsdis/cpscom/greencom/ithings 2015*, pp. 464–469, 2015.

[2] L. A. Barroso, U. Hölzle, and P. Ranganathan, "Data center basics: Building, power, and cooling," *Datacenter As a Computer*, pp. 75–98, 2019.

[3] S. Moss. (18/05/2022) In search of the world's largest data center. Accessed: 15/03/2023. [Online]. Available: https://www.datacenterdynamics.com/en/analysis/in-search-of-the-worlds-largest-data-center/

[4] X. Wang, X. Liu, L. Fan, and J. Huang, "Energy-aware resource management and green energy use for large-scale datacenters: A survey," *Advances in Intelligent Systems and Computing*, vol. 255, pp. 555–563, 2014.

[5] F. Kong and X. Liu, "A survey on green-energy-aware power management for datacenters," *Acm Computing Surveys*, vol. 47, no. 2, p. 2642708, 2014.

[6] M. Zakarya, "Energy, performance and cost efficient datacenters: A survey," *Renewable and Sustainable Energy Reviews*, vol. 94, pp. 363–385, 2018.

[7] M. Dayarathna, Y. Wen, and R. Fan, "Data center energy consumption modeling: A survey," *Ieee Communications Surveys and Tutorials*, vol. 18, no. 1, p. 7279063, 2016.

[8] L. R. Jahangard and A. Shirmarz, "Taxonomy of green cloud computing techniques with environment quality improvement considering: a survey," *International Journal of Energy and Environmental Engineering*, vol. 13, no. 4, pp. 1247–1269, 2022.

[9] M. H. M. Gavali, M. S. S. Patil, M. P. U. Patil, S. P. Mane, and M. K. N. Rode, "Green cloud computing," *International Journal for Research in Applied Science and Engineering Technology*, vol. 10, no. 4, pp. 581–583, 2022.

[10] A. Katal, S. Dahiya, and T. Choudhury, "Energy efficiency in cloud computing data center: a survey on hardware technologies," *Cluster Computing*, vol. 25, no. 1, pp. 675–705, 2022.

[11] C. Hu, Y. Guo, Y. Deng, and L. Lang, "Improve the energy efficiency of datacenters with the awareness of workload variability," *Ieee Transactions on Network and Service Management*, vol. 19, no. 2, pp. 1260–1273, 2022.

[12] M. Zakarya, "Energy, performance and cost efficient datacenters: A survey," *Renewable and Sustainable Energy Reviews*, vol. 94, pp. 363–385, 2018.

[13] G. Antoniou *et al.*, "Agilepkgc: An agile system idle state architecture for energy proportional datacenter servers," *Proceedings of the Annual International Symposium on Microarchitecture, Micro*, vol. 2022-, pp. 851–867, 2022.

[14] M. S. B. M. Desa *et al.*, "Energy efficient approach using server virtualization in cloud data center," *2018 Ieee 4th International Symposium in Robotics and Manufacturing Automation, Roma 2018*, p. 8986732, 2018.

[15] L. C. Yan, Y. Li, H. Song, H. D. Zou, and L. J. Wang, "Edge computing based data center monitoring," *Proceedings - Ieee International Conference on Edge Computing*, vol. 2021-, pp. 17–24, 2021.

[16] C. Gough, I. Steiner, and W. Saunders, *Energy efficient servers: Blueprints for data center optimization*. Apress Media LLC, 2015.

[17] C. Hu, Y. Deng, G. Min, P. Huang, and X. Qin, "Qos promotion in energy-efficient datacenters through peak load scheduling," *Ieee Transactions on Cloud Computing*, vol. 9, no. 2, pp. 777–792, 2021.

[18] S. Loganathan, R. D. Saravanan, and S. Mukherjee, "Energy aware resource management and job scheduling in cloud datacenter," *International Journal of Intelligent Engineering and Systems*, vol. 10, no. 4, pp. 175–184, 2017.

[19] V. Anagnostopoulou, S. Biswas, H. Saadeldeen, A. Savage, R. Bianchini, T. Yang, D. Franklin, and F. T. Chong, "Barely alive servers: Greener datacenters through memory-accessible, low-power states," *Design Technologies for Green and Sustainable Computing Systems*, pp. 149–178, 2013.

[20] T. L. Vasques, P. Moura, and A. de Almeida, "A review on energy efficiency and demand response with focus on small and medium data centers," *Energy Efficiency*, vol. 12, no. 5, pp. 1399–1428, 2019.

[21] B. Speitkamp and M. Bichler, "A mathematical programming approach for server consolidation problems in virtualized data centers," *Ieee Transactions on Services Computing*, vol. 3, no. 4, pp. 266–278, 2010.

[22] Y. Gong, F. Zhou, G. Ma, and S. Liu, "Advancements on mechanically driven two-phase cooling loop systems for data center free cooling," *International Journal of Refrigeration*, vol. 138, pp. 84–96, 2022.

[23] N. Horner and I. Azevedo, "Power usage effectiveness in data centers: Overloaded and underachieving," *Electricity Journal*, vol. 29, no. 4, pp. 61–69, 2016.

[24] L. M. P. Larsen, H. Christiansen, S. Ruepp, and M. Berger, "Towards greener 5g and beyond radio access networks - a survey," *IEEE Open journal of the Communications Society*, 2023.

[25] D. Huang, "Data center infrastructure management," *Data Center Handbook: Plan, Design, Build, and Operations of a Smart Data Center*, pp. 627–644, 2021.

[26] A. Garg and D. Shenkar, "Drive data center management and build better ai with it devices as sensors," *Data Center Handbook: Plan, Design, Build, and Operations of a Smart Data Center*, pp. 669–673, 2021.

[27] J. Liu, C. T. Hann, and L. Jiang, "Quantum data center: Theories and applications," p. 24, 2022, accessed: 15/03/2023. [Online]. Available: https://arxiv.org/abs/2207.14336

[28] J. Haj-Yahya, Y. Sazeides, M. Alser, E. Rotem, and O. Mutlu, "Techniques for reducing the connected-standby energy consumption of mobile devices," *Proceedings - 2020 Ieee International Symposium on High Performance Computer Architecture, Hpca 2020*, pp. 623–636, 2020.