Spring 5-27-2024

Analysis of Green Data Center Efforts and Energy Usage

Dillon J. Goicoechea

Analysis of Green Data Center Efforts and Energy Usage

by

Dillon Goicoechea

FACULTY MENTORS:

Ange Kakpo

HONORS PROGRAM DIRECTOR:

DR. Joshua Tom

A project submitted in partial fulfillment of the requirements
for the Bachelor of Arts degree in Honors Liberal Arts
Seattle Pacific University
2024

Presented at the SPU Honors Research Symposium
May 18th, 2024

Analysis of Green Data Center Efforts and Energy Usage

Dillon Goicoechea

***Abstract*—This paper is an undergraduate level literature review and analysis of research surrounding the Green Data Center phenomenon. Review of work covering energy usage, data usage, usage predictions, and strategies for decreasing energy requirements is the main analysis of this work. The analysis shows that while data centers are becoming greener, the increase in usage of their capacities is negating those efficiency increases. The increase in the energy efficiency of data centers is crucial, however, there must be made efforts to lower computational and data usage to help achieve lower energy usage of data centers.**

***Index Terms*—PUE, Data Center, Green Data Center, DCIE, PEMC, EUED, PDD, distributed computing**

I. Introduction

MODERN data centers (DC) have become an integral part of business operations over the last 2 decades, especially with the mass adoption of cloud-based work programs such as Microsoft Teams and Slack [1]. Due to the workload of modern DC, they are some of the biggest singular place energy users on the planet, with their current estimate being about 2.4% of all electricity globally [2]. The energy requirement of modern DC has become a global environmental concern, as in 2008 they accounted for 2% of global carbon emissions, with their models predicting a strong increase in energy usage [3]. These trends have led to the concept of green DC, which are more energy efficient and renewable based DC, to lower the ever-growing energy usage of modern DC. This movement began in the early 2000s, with a private effort to create these as businesses wanted to save money on resource usage. This eventually grew into regulation, particularly starting with the EU. The academic body of work surrounding green DC is primarily positing new engineering solutions or algorithmic enhancements to create greener solutions. Data centers are composed of 3 main components: servers, cooling technology, and networks.

DC began as government owned and operated computing centers in the mid-1900s, with commercial usage seeing its birth during the 1970s. The term data center did not come about until the 1990s, when it described server computers [4]. Since then, corporate integration and use of DC has skyrocketed, especially with the creation of hosting services and cloud data storage, both of which utilize most of DC computation [5].

II. Measurement Units

Modern DCs have various measurements to gauge effectiveness, but the most common is PUE, or Power Usage Effectiveness, which measures a DC's Total Facility Power divided by IT Power, which tells how much of the facility's power usage went to computing power [6][7]. The best this metric can ever be is a 1.0 by the nature of this equation, and many companies have been able to threshold into a sub 1.2 PUE for their DC. This metric is simple, which means it holds much stronger when comparing DC. DCiE, or Data Center Infrastructure Efficiency, is the inverse of PUE [6]. They effectively measure the same thing, just with different phrasing. HVAC is a metric for measuring the energy usage of cooling technologies within DC [6].

In recent years, new metrics such as EUED, PDD, and PEMC have been suggested as ways to get a better understanding of the mathematical relation of energy usage factors of DC. EUED is the total DC energy with variation per year divided by the energy directly consumed by IT equipment per year. This is important to understand PDD, as it is EUED altered to prioritize efficiency. PDD, or Perfect Design DC index, is 1 divided by the EUED of a DC [8]. This version of the equation allows for stronger linear correlation to be made. Additionally, there is PEMC, or Power Efficiency Measurement Calculator, which is a proposed algorithm to do a series of calculations with the PUE and DCiE of a DC to create tiered measurement of the energy efficiency of a DC [9].

All these metrics are for energy efficiency, redefined based on different explanatory variables.

III. Data Consumption

The birth of hyper-scale DC leans toward the conclusion of drastically increased consumer demand for cloud computing resources, with the total number of global internet users reaching 3.9 billion in 2018 [1]. This is supported by the increase in the number of applications that require an internet connection, as applications and services are now hosted on DC, creating a need for significantly more and faster cloud computing resources. US DC handled about 300 million tb of information in 2016, with data usage of internet resources predicted to hit 175 ZB by 2025, supporting an increased demand for cloud computing from WFH products, streaming services, online games, etc. [1][10]. This has correlated to both increased computing usage of existing DC and a need to build DC in more diverse locations to maintain low latency communications for all users. Low latency communications drive up energy usage, as DC cannot reliably sleep without hindering desired user experience and communication needs. This has led to Time of Use pricing models for computation to surge, AWS being a great example of these phenomena [1].

IV. Energy Consumption

DC energy usage is predicted to increase going to 2030 in all scenarios, with the best case being at least double its energy usage in 2020 [11]. DC have previously accounted for up to 10% of global electricity usage, with their usage to represent about 3–4% of total energy usage by 2030. US based DC consumed about 70 billion kWh of energy, about 100 times the amount used by a similar sized office spaces [12]. Hyper-scale DCs use as much as 25,000 households worth of energy in the US, with cooling systems accounting for up to two thirds of energy consumption [3]. This directly relates to consumer demand data usage, as US DC often utilizes 28 kWh per Tb of data used. This correlates to about 35 kg of CO_2 waste per tb of processed information in 2016. The predicted CO_2 emissions of DC are projected to hit 720 million tons by 2030 [10].

DC have become more efficient with introduction of hyper-scale solutions, as FPE has significantly lowered since 2010, but energy metrics show that energy usage is still at an all-time for data centers. Global electricity usage is of communication technology is expected to grow to about 1000 TWh, with data centers predicted to grow in energy usage between 15–20% each year [11].

V. Emerging Strategies

The biggest targets for creating Green DC are energy sourcing, hardware, structural, and algorithmic improvements. Modern designs of Green DC and ideas for improving existing DC focus on these areas primarily for reducing energy consumption, however, energy sourcing is the exception. Energy sourcing is important as the push to create greener DC is rooted in environmental concerns around carbon emissions related to energy consumption (additionally with financial incentives for private owners). If a DC sources all of its energy from non-emissions sources, the issue of emissions concerns is gone. The other 3 targets are by-products of general advancements in technology.

A. Energy Sourcing Example

Under ANR DATAZERO project, researchers have proposed designs of a fully green power source DC. This model works with 2 overarching concepts of power, with 4 separate sources. The two big concepts are renewable resource farms and energy storage systems. The individual sources proposed are Solar Panels, Wind Turbines, Batteries from the ESS system, and hydrogen powered fuel cells fed by an electrolyzer. In the power architecture, there would be 3 decision-making modules: The IT decision module, the negotiation module, and the power decision module. These 3 modules would conceptually debate and choose a power plan for the DC, given by measured factors, and implement the power plan for a set period of time. The power plan is implemented by the power decision module who chooses which sources to use, and the IT decision module operates the DC with the anticipated power plan, whether higher or lower energy usage. Experimentation of the proposed model under various conditions showed success with some caveats to be addressed, such as the feasibility of power plans and scale of measurements [13].

B. Structural Improvement Example

Network architecture is a critical part of DC efficiency, and the equipment used to build out the architecture heavily affects the energy efficiency and usage of a DC. In 2008, a special instance of Clos topology named fat-tree, organized with a k-ary approach. The topology has key advantages, such as creating a structure where all switching elements are equivalent, and are rearrangeably non-blocking, essentially allowing paths to utilize maximum bandwidth. This structure requires a strong traffic diffusion to achieve ideal bandwidth, where it then falls to routing protocols and algorithms to complete tasks. This structure allows for usage of less total networking equipment and less taxing network equipment, leading to a reduction in maximum power usage by 56.6%. The effects of the fat-tree architecture address both scalability and bottle necking issues relevant for modern delivery of cloud-based applications and decrease energy consumption compared to hierarchical traditional architecture [14].

C. Hardware Improvement Example

The costliest part of DC operations is the cooling of internal computer equipment to increase efficiency and lifespan. Cooling solutions of DC primarily fall into 2 categories: air-based and liquid-based. While liquid-based systems are typically more effective, they are more costly energy wise. Research around better utilization of air-based technologies has been a key part, especially on how to solve the problem of exhaust air (EA) recirculation. Current methodology uses outdoor air (OA) primarily for cooling of systems, to maintain a lowered air temperature. When EA mixes back in with the OA, it causes a general rise in temperature, lowering the effectiveness of the cooling efforts. Usage of air-side economizers, both direct and indirect, is a method to prevent EA from mixing back into the system. Direct air-side economizers work by pushing EA to the outside directly, whereas indirect air-side economizers create wind paths forcibly preventing recirculation. Both types of air-side economizers consumed less energy than the existing water-chilled system in the study [15].

D. Algorithmic Improvement Example

Proposed in 2023, the GEECO DC framework suggests a 3 headed approach to task scheduling. This module utilizes 3 task categorizing algorithms to schedule tasks based on deadline and resource requirements. The system additionally supports dynamic rescheduling. The 3 categorization algorithms utilized to schedule tasks are Longest Process Time (LPT), Shortest Processing Time (SPT), and Longest Setup Times First (LSTF). LPT schedules the most process intensive tasks first, allowing the DC to complete simpler tasks more simultaneously once more resources are free. SPT is the opposite of LPT, putting simpler tasks first to optimize task throughput, creating a more responsive and efficient task completion system. LSTF considers the preparation required for the task queue, and puts the task with more complex preparation first, reducing idle time between tasks. After each task is categorized, the resource estimation module performs cost analysis on each task, allowing the system to perform more cognizant load balancing and scaling, leading to less downtime [2].

VI. Conclusion

DC are essential for the modern Internet of Things, and understanding efforts to improve them is crucial for planning and predicting resource usage. It can be seen by the literature gathered and reviewed in this paper that DCs are becoming more energy efficient and using more green energy. However, the increasing computation and data usage in the modern environment leads to more energy being used in DC than ever before, even with improvements being implemented and designed, such as the examples from the previous section. This paper is not as comprehensive as other literature analysis in this topic and could most definitely benefit from an increase in strategies discussed as well as a further dive into the specific results of each tested strategy. These conclusions lead to the topic of discussing how to effectively lower the usage of data and DC in the modern environment, as the growth of consumption of information leads to a significant increase in the energy requirements.

Appendix A
4th Annual SPU Honors Research Symposium Script, Presented May 18th, 2024, Panel 7:Eating the Internet: Mass Consumption of Media, Data, and Energy

Good Afternoon all, and welcome to my presentation on Energy Efficiency of Green Data Center Efforts. To begin the presentation, let us start with establishing some baseline knowledge. In the Modern Information Era, Data Centers are the backbone of the entire internet. Any piece of data that is transmitted, received, sent, or saved on the cloud depends on data centers in some capacity. Modern data centers are large internet infrastructure sites filled with immense data storage and computational power combined with low-latency communications technology. They began in the mid-20th century as government computation sites, and began their transformation into their modern incarnation in the 1990s, when the terminology switched from business server sites to Data centers.

Green Data centers are data center that are meant to consume less or no “brown” energy, or energy via fossil fuels. This effort came into spotlight due to a private push to waste less money on resources for data center and was further enacted into law starting in Europe. The key part of this issue is the “brown energy” aspect when it comes to impact on our direct lives. While the brown energy usage of data centers is measured, it is not how we judge energy efficiency with this infrastructure sites.

Measuring the energy efficiency of any system, let alone one as complex as a data center, is hard at its easiest. The most common way to measure is PUE, Power Usage Effectiveness, which is a facilities total power consumption over its IT power consumption, telling how much power went to the computing tasks. The best score is a 1 on this metric, with it being worse as it gets higher. There is additionally an algorithmic assessment, PEMC, proposed to create tiered measurements of data center energy efficiency utilizing multiple metrics in a series of calculations.

Measurements of energy efficiency of data centers often leave out the most complex part of the equation: Data is energy. Any and all transactions involving data requires energy, from viewing to copying. As Data usage of the internet is expected to become 175 ZB by 2025. This is most definitely related to the increased demand of content streaming, such as video streaming, social media, WFH products, cloud access data, and even online gaming(although significantly smaller). This increase in data correlates with a strong increase in energy consumption of data centers, with their usage to represent 3-4% of total global electricity by 2030. US based data centers consumed about 100 times more energy than comparable commercial office spaces, and Hyperscale data centers (such as azure regional centers) consume as much as 25,000 households. US data centers consumed about 28 kWh per Tb of data transactions, leading to about 35 kg of CO_2 per Tb of processed information. To put into context, HD video streaming uses about 1.5 to 2.5 GBs of data per hour. 4k HDR video streaming uses between 7-10 GBs of data per hour. Online gaming uses about 10 GBs of data per month. Say you watch a 10 episode season of a new show, where each episode is an hour. 10 hours of HDR streaming, you would use 0.1 TB of data, or 3.5 kg of CO_2.

Barring the issue data, data centers are now more energy efficient than ever, they are becoming significantly more energy efficient, especially over the last decade. The biggest targets for creating Green data centers are energy sourcing, hardware, structural, and algorithmic improvements. Modern designs of Green data centers and ideas for improving existing data centers focus on these areas primarily for reducing energy consumption, however, energy sourcing is the exception. Energy sourcing is important as the push to create greener data centers is rooted in environmental concerns around carbon emissions related to energy consumption (additionally with financial incentives for private owners). If a data centers sources all of its energy from non-emissions sources, the issue of emissions concerns is gone. The other 3 targets are by-products of general advancements in technology.In an example of sourcing improvements, Under ANR DATAZERO project, researchers have proposed designs of a fully green power source data centers. This model works with 2 overarching concepts of power, with 4 separate sources. The two big concepts are renewable resource farms and energy storage systems. The individual sources proposed are Solar Panels, Wind Turbines, Batteries from the ESS system, and hydrogen powered fuel cells fed by an electrolyzer. In the power architecture, there would be 3 decision-making modules: The IT decision module, the negotiation module, and the power decision module. These 3 modules would conceptually debate and choose a power plan for the data centers, given by measured factors, and implement the power plan for a set period of time. The power plan is implemented by the power decision module who chooses which sources to use, and the IT decision module operates the data centers with the anticipated power plan, whether higher or lower energy usage. Experimentation of the proposed model under various conditions showed success with some caveats to be addressed, such as the feasibility of power plans and scale of measurements.

Networking is costly to improve, but it makes all the difference. Network architecture is a critical part of data centers efficiency, and the equipment used to build out the architecture heavily affects the energy efficiency and usage of a data centers. In 2008, a special instance of Clos topology named fat-tree, organized with a k-ary approach was designed and implemented. The topology has key advantages, such as creating a structure where all switching elements are equivalent, and are rearrangeably non-blocking, essentially allowing paths to utilize maximum bandwidth. This structure requires a strong traffic diffusion to achieve ideal bandwidth, where it then falls to routing protocols and algorithms to complete tasks. This allows for usage of less total networking equipment and less taxing network equipment, leading to a reduction in maximum power usage by 56.6%. The effects of the fat-tree architecture address both scalability and bottle necking issues relevant for modern delivery of cloud-based applications and decrease energy consumption compared to hierarchical traditional architecture.

Hardware of computation is always improving, but stagnation in surronding technologies is strong. The costliest part of data centers operations is the cooling of internal computer equipment to increase efficiency and lifespan. Cooling solutions of data centers primarily fall into 2 categories: air-based and liquid-based. While liquid-based systems are typically more effective, they are more costly energy wise. Research around better utilization of air-based technologies has been a key part, especially on how to solve the problem of exhaust air (EA) recirculation. Current methodology uses outdoor air (OA) primarily for cooling of systems, to maintain a lowered air temperature. When EA mixes back in with the OA, it causes a general rise in temperature, lowering the effectiveness of the cooling efforts. Usage of air-side economizers, both direct and indirect, is a method to prevent EA from mixing back into the system. Direct air-side economizers work by pushing EA to the outside directly, whereas indirect air-side economizers create wind paths forcibly preventing recirculation. Both types of air-side economizers consumed less energy than the existing water-chilled system in the study.

Algorithmic improvements are rare, and harder to truly test. Proposed in 2023, the GEECO data centers framework suggests a 3 headed approach to task scheduling. This module utilizes 3 task categorizing algorithms to schedule tasks based on deadline and resource requirements. The system additionally supports dynamic rescheduling. The 3 categorization algorithms utilized to schedule tasks are Longest Process Time (LPT), Shortest Processing Time (SPT), and Longest Setup Times First (LSTF). LPT schedules the most process intensive tasks first, allowing the DATA CENTER to complete simpler tasks more simultaneously once more resources are free. SPT is the opposite of LPT, putting simpler tasks first to optimize task throughput, creating a more responsive and efficient task completion system. LSTF considers the preparation required for the task queue, and puts the task with more complex preparation first, reducing idle time between tasks. After each task is categorized, the resource estimation module performs cost analysis on each task, allowing the system to perform more cognizant load balancing and scaling, leading to less downtime.

data centers are essential for the modern Internet of Things, and understanding efforts to improve them is crucial for planning and predicting resource usage. It can be seen by the literature gathered and reviewed in this project that data centers are becoming more energy efficient and using more green energy. However, the increasing computation and data usage in the modern environment leads to more energy being used in data centers than ever before, even with improvements being implemented and designed, such as the examples from the previous section. These conclusions lead to the topic of discussing how to effectively lower the usage of data and data centers in the modern environment, as the growth of consumption of information leads to a significant increase in the energy requirements.

Essentially all usage of the internet in the modern day in age, and especially the uses studied by my fellow panelists, require data centers. As content and media become more abundant and more consumed, so does the energy required to store and access this data. Now increasing energy consumption is a problem of itself enough, but even more important when put into the context of the CO2 generated as a byproduct.

There are many ways to reduce data centers consumption, but the heavy hitter is reduce brown energy consumption and data consumption. This simply means smaller things being transported across the internet. Now this is not an issue on the consumer, as you do not control the standard used in the modern technology business. It also means pushing your government for more green energy production and access.

While reduction of data size is not a thing problem put onto the shoulders of the consumers, there are some best practices that can help YOU both financially and energy wise, as some ISPs charge by data usage AND speed: 1) Cut down on the streams of media you have ongoing. I know that sometimes it helps have multiple streams going to generate background noise, but be conscious of what you are using. Instead of tv, try music, audiobooks, or podcasts, all very low data usage. 2) Reduce the size of files you send via zip or even just lower the quality a little. 3) Be smart about downloading items. Try not to install and uninstall then reinstall a 200gb game often. 4) Sometimes on streaming services, you can opt to use the quality of streaming that is actually correlated to your device(I will note most streaming services say they do this automatically now) so that you aren't wasting your data.

The good news is that the internet and using it is not going to end the world anytime faster than all the other stuff we do. However, just as cars and energy production, we can do them smarter and cleaner to create a healthier environment both now and in the future.

References

[1] M. T. Takcı, T. Gözel, and M. H. Hocaoğlu, "Quantitative evaluation of data centers' participation in demand side management," *IEEE Access*, vol. 9, p. 14883–14896, 2021.

[2] S. Mondal, F. B. Faruk, D. Rajbongshi, M. M. K. Efaz, and M. M. Islam, "Geeco: Green data centers for energy optimization and carbon footprint reduction," *Sustainability*, vol. 15, no. 21, p. 15249, 2023.

[3] T. Bhattacharya, t. l. w. o. i. a. n. t. Link to external site, M. Rahgouy, X. Peng, T. Takreeti, T. Cao, J. Mao, A. Das, X. Qin, and A. Sinha, "Capping carbon emission from green data centers," *International Journal of Energy and Environmental Engineering*, vol. 14, no. 4, p. 627–641, Dec. 2023.

[4] 2018. [Online]. Available: https://web.archive.org/web/20181024232055 /https://blog.rackspace.com/datacenter-evolution-1960-to-2000

[5] C. Harvey, "What is a data center? — history, design, cooling & types," 2017. [Online]. Available: https://www.datamation.com/data-center/what-is-data-center/

[6] L. Wang and S. Khan, "Review of performance metrics for green data centers: a taxonomy study," *Journal of Supercomputing*, vol. 63, no. 3, p. 639–656, Mar. 2013.

[7] C. C. Zhou, F. Zhang, Z. Yuan, L. M. Zhou, F. M. Wang, W. Li, and J. H. Yang, "Study on load forecasting to data centers of high power density based on power usage effectiveness," *IOP Conference Series: Earth and Environmental Science*, vol. 40, no. 1, p. 012053, Aug. 2016.

[8] A. F. Santos, P. D. Gaspar, and H. J. L. d. Souza, "New data center performance index: Perfect design data center—pdd," *Climate*, vol. 8, no. 1010, p. 110, Oct. 2020.

[9] A. Shaikh, M. Uddin, M. A. Elmagzoub, and A. Alghamdi, "Pemc: Power efficiency measurement calculator to compute power efficiency and co_2 emissions in cloud data centers," *IEEE Access*, vol. 8, p. 195216–195228, 2020.

[10] K. M. U. Ahmed, M. H. J. Bollen, and M. Alvarez, "A review of data centers energy consumption and reliability modeling," *IEEE Access*, vol. 9, p. 152536–152563, 2021.

[11] A. S. G. Andrae and T. Edler, "On global electricity usage of communication technology: Trends to 2030," *Challenges*, vol. 6, no. 11, p. 117–157, Jun. 2015.

[12] Y. Liu, X. Wei, J. Xiao, Z. Liu, Y. Xu, and Y. Tian, "Energy consumption and emission mitigation prediction based on data center traffic and pue for global data centers," *Global Energy Interconnection*, vol. 3, no. 3, p. 272–282, Jun. 2020.

[13] M. Haddad, J.-M. Nicod, M.-C. Péra, and C. Varnier, "Stand-alone renewable power system scheduling for a green data center using integer linear programming," *Journal of Scheduling*, vol. 24, no. 5, p. 523–541, Oct. 2021.

[14] M. Al-Fares, A. Loukissas, and A. Vahdat, "A scalable, commodity data center network architecture," in *Proceedings of the ACM SIGCOMM 2008 conference on Data communication*. Seattle WA USA: ACM, Aug. 2008, p. 63–74. [Online]. Available: https://dl.acm.org/doi/10.1145/1402958.1402967

[15] S. Park and J. Seo, "Analysis of air-side economizers in terms of cooling-energy performance in a data center considering exhaust air recirculation," *Energies*, vol. 11, no. 22, p. 444, Feb. 2018.