Reducing the Environmental Impact of Data Centers

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science
University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science, School of Engineering

Nikita Saxena

Spring 2022

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Advisor

Joshua Earle, Department of Engineering and Society

Reducing the Environmental Impact of Data Centers

How do interest groups, including environmentalists, policymakers, and technology companies, strive to reduce the environmental impact of data centers?

Background

Why the Data Center?

Data centers are the backbone of the worlds computing infrastructure, supporting everything from the business sector to government and academic operations. As the economy has become increasingly dependent on the digital medium, data centers have become critical in turn. Data centers are large scale computing infrastructures containing equipment for data storage, processing, and communication, collectively known as "Information Technology" (IT) equipment, as well as supporting power generation, cooling, and environmental control equipment.

There is great variety in the structure of data centers. There exist both large scale data centers located in specialized buildings and small-scale data centers within general purpose commercial buildings. While some institutions manage their own data centers, others rely on consolidated data centers managed by outside companies that cater to multiple institutions. Some larger companies can manage upwards of hundreds of data centers, whether for themselves or to partition out to others. (Shehabi et al., 2016)

Introduction

Increasing Computing Workloads

Computing workloads going through data centers have become more intensive within the last two decades due to a flood of data from new technologies. Online platforms, social media, "smart" infrastructure, artificial intelligence, amongst other existing and budding technologies demand the processing, storage, and communication of data. Computing workloads increased by an estimated 550% between 2010 and 2018, with data center traffic having an annual growth rate of 23 percent. Increases in computing workload has naturally caused the average power density, or power required per rack, to increase in tandem. (Siddik et al., 2021) This landscape continues to shift in the 2020s. As of the 2021 State of the Data Center Report, 58% of survey respondents saw a shift of computing workload from the cloud onto on-site or colocation data centers (centers that can be rented out to individual enterprises), while 62% of respondents reported a continued increase in rack density to account for computation intensive workloads. (Kleyman et al., 2021)

Resource Consumption of Data Centers

Maintaining this data is similarly intensive, not only in energy consumption, but in greenhouse gas emissions and water use. Data centers account for 0.3 percent of global carbon emissions and about 2 percent of global electricity consumption, with demand increasing at an annual rate of about 5 percent. (Dayarathna et al., 2016; Jones, 2018). A benchmarking study by Andrae et al. (2015) predicted the IT ecosystem could use up to 20 percent of global electricity use by 2030, with data centers accounting for a third.

Water is used in data centers to dissipate heat produced by IT equipment in operation. A typical data center uses about as much water as a city with a population of about 30,000 to 40,000 people. Electricity is also used indirectly in wastewater plants supplying these centers to treat water for use, and to clean it after use. (Calma, 2021) Cooling infrastructure is particularly energy intensive, with high potential for waste. (Jones, 2018). The power grid, another large polluter and the second heaviest water consumer, supplies most of this electricity. (Siddik et al., 2021) A comprehensive view of these systems reveals the potential for environmental impact.

Electricity use was essentially steady in data centers from 2016 – 2018 despite the after mentioned increase in workloads. In this time, data center electricity consumption only rose 6% in this time even given the high annual growth rate of traffic, implying the centers were able to keep up with increased demand. This trend usually attributed to industry-wide improvements in energy efficiency and storage-drive density. The movement towards modernized hyperscale data centers and colocation data centers which are almost 6 times more water efficient over internal data centers has also heavily contributed. (Jones, 2018; Siddik et al., 2021) However, once workload is moved to more efficient hyperscale data centers further optimization is limited. "The trend is good right now, but it's questionable what it's going to look like in 5–10 years," says Dale Sartor, who oversees the Federal Energy Management Program's Center of Expertise for Data Centers. (Jones, 2018; LBNL, n.d.a). Thus, we must look beyond only technological developments to maintain this trend.

Methods

A review of existing technology initiatives, policy, and activism as follows will highlight techniques that allowed for improvements in efficiency in the past decade on both the local and global level. This review focuses primarily on the sociotechnical systems rather than technological developments as those systems shape the progress of large-scale sustainability measures in data centers. The primary participants in the system surrounding data centers are the technology companies responsible for the centers, the governmental agencies regulating them, and the communities who live around them. Primary sources were used to explore the perspectives of these participants. This included press releases regarding sustainability practices, policy documents and summaries, and first-hand accounts from activists. Secondary news reports, analysis, and academic journal articles were used to provide context to and supplement these perspectives as well as explore the effectiveness of various approaches.

Technology Companies

Complex Sustainability Plans

The technology companies that build and maintain data centers play a key role in sustainability. Many "tech giants" outline self-imposed sustainability plans to modernize technology and source green energy. (Joppa & Walsh, 2021) These plans are usually multipronged, having to tackle the issue at multiple levels due to the scale of operations, and implementation is often capital intensive. (Calma, 2021; Walsh, 2022) Taking Microsoft as a case study, it plans to build 50 to 100 new data centers per year to keep up with their cloud demand. They have also pledged to cut water usage in their data centers by up to 95% by 2024 and become "carbon negative" (store more carbon than they generate) by the same deadline - in the long term they similarly aim to become "water positive" (replace more water than they use). (Calma, 2021) This involves a lot of varied and intensive research and development - whether it be for less carbon-intensive building materials, carbon tracking tools, new cooling methodology, or simply researching how servers perform under higher temperatures. (Walsh, 2022)

Special attention must be given to data centers in hot and dry climates - the same techniques provide less benefits. "You can become water positive as a company, but it's not useful if you save all the water in one location, as an extreme case, and you create bigger deficits elsewhere." remarked Venkatesh Uddameri, the director of the Water Resources Center at Texas Tech University. Managing these data centers on a local level must happen on a case-to-case basis, and centers in hot climates should be prioritized due to the relative scarcity of water. (Calma, 2021) For example, Companies such as Microsoft have promoted recent efforts to build “sustainable data centers”, the center built cooling technologies that are stated to reduce impact on water supplies by working with locals near their Arizona plant (Walsh, 2021)

The potential positives of these changes are great, but the undertaking of changing these supply chains and designs is itself a challenge. (Calma, 2021) However, these initiatives can be waylaid even despite the economic incentive of lowered energy use. In the big picture, Mills et. Al. (2007) found that most smaller companies prioritize capital use for short-term gain despite long term profits from improved energy efficiency in high tech facilities such as data centers

Open Source Information

Promoting data center modernization must be a group effort on the corporate scale. Modernization is supported by open data access to originally propriety designs and research which allows efficient designs to become ubiquitous and be iterated on. This idea is espoused by key technology companies such as Microsoft (Walsh, 2022) and Facebook. (Facebook, n.d.) Microsoft’s Vice President of Cloud Ops and Innovation Noelle Walsh has stated, “to get there, we need to share our learnings and progress, and create new tools and solutions to benchmark where we are today, measure our progress and make them widely available” in relation to the release of the Microsoft Cloud for Sustainability tools which help track and manage carbon emissions. (Walsh, 2022)

Modernization is also buoyed by organizations such as Open Compute by Facebook that promote free access, redesigned hardware, and research into making computing more energy efficient. This project aims to make important information more available by publishing standards for green data centers and encouraging “major technology innovators” to make their work public. This allows for more hands to take an attempt at boosting data center efficiency and enables these more efficient designs to become more ubiquitous. (Facebook, n.d.) Both

Facebook and Google have openly published server design advancements under this framework. Sales of hardware based on Open Compute Project designs reached $1.2 billion in 2017. (Sayer, 2018)

Effective Policy

Research and Outreach

Policymakers have historically also promoted data center energy efficiency. The U.S. Department of Energy spearheads promotion efforts by establishing publicly accessible data center design guidelines, publishing case studies, and funding research for energy efficient technologies. All such research and guidelines are publicly accessible. (Office, n.d.). In particular, the Department of Energy partnered with the Environmental Protection Agency to create the National Data Center Energy Efficiency Information Program in 2011. This program was notably voluntary, rather than a requirement or law. It standardized efficiency metrics, developed tools to assist data center managers, and provided labeling and recognition for well preforming server technology and data centers. (DOE & EPA, 2008) These programs aimed to ease the path towards greener data centers by making the information easier to acquire and interpret.

Initiatives also partnered with data center efficiency laboratories which provide training directly to private stakeholders, closing the gap between federal guidelines and the private enterprise that run data centers. These specialized laboratories also provide technical support, advising, and training programs directly to private stakeholders regarding good data center practice. Some of these laboratories were also established via government initiative, such as Center of Expertise for Energy Efficiency in Data Centers which was established by the Federal Energy Management Program (LBNL, n.d.b). Government backed research initiatives also allows for development and demonstration of new technology outside of commercialization, which is vital in areas where incentive for privately funded research is low. In particular, public funding over private funding may be especially necessary as chip performance limits approach, lowering the incentive for privately funded research. (Koronen et al. 2019)

Incentives and Requirements in Policy

States offer tailored economic incentives for data centers as they are a significant taxable investment that provide many employment opportunities. Incentives include credits, tax exemptions, and lowered electricity rates. Michael F. Kaestner (2014) discusses how these incentives are often incongruous with energy policies in the same states. Data centers take up a disproportionate amount of electricity, often subsidized by the locals, for a service that isn't always buoying the local stores. Policymakers could be more mindful of energy availability and environmental impact of data centers by limiting incentives to facilities that have proven to be committed to sustainability and conservation efforts. (Kaestner, 2014) General reviews of the technological challenges, environmental impact, and ethical concerns of large-scale data usage have also drawn similar comparisons between environmental and economic policy (Song et al., 2019; Lucivero, 2019)

Policy promoting the use of sustainable technology in private companies can take the form of requirements codified into law, sustainability programs that have voluntary participation, or a mix of both – to varying results. For example, the European Code of Conduct for Data Centre Energy Efficiency initiative (CoC) was a successful non-regulatory/voluntary policy that aimed account for the diversity in data center facilities - from age of IT equipment to style of core infrastructure. It provided a platform for stakeholders to agree on acceptable actions towards greener data centers and requested both an actionable plan and regular audits to track efficiency over time. All participating sites implemented at least the minimum expected best practices guidelines outlined within the program, and the number of applications for the next round of the program increased. This was one of the most successful data center policies in the EU, and shows a data backed non-regulatory/voluntary policy. As such Avgerinou et al. (2017) concluded voluntary agreements were favored over mandatory regulations when dealing with the private companies as they account for diversity in design, goals, and ownership.

In contrast, the "Ecodesign Directive" established in 2009 sets minimum efficiency requirements for servers and storage devices often used in data centers as of 2020. They also set reporting standards for manufacturers selling this equipment. While Ecodesign could at least prevent the least energy efficient servers coming into circulation, this area is difficult to regulate compared to domestic appliances due to the rapid pace of development and the complexity of technology, The resulting standards are both harder to define well and harder to measure the effectiveness of over that of more comprehensive, but voluntary, programs such as the CoC. (Koronen et al. 2019)

Local Level

Tensions over Water

Traditional wisdom states that communitys welcome data centers due to the lack of visible environmental detriments and the internet sector's "sleek" reputation. (Avgerinou et al., 2017). Even so, tensions have risen in certain local communities regarding high water and energy usage. An example is the previously referenced Microsoft data centers in Arizona where Microsoft promoted efforts to build more sustainable centers working with locals. (Walsh, 2021) However, Vice Mayor Jenn Duff of the Mesa City Council, Arizona echoed local concerns and voted against a new data center development - citing an "an irresponsible use of our water" in the "driest 12 months in 126 years".

This is not an out of the norm case. Data centers tend to be located close to infrastructure and customers, and where the cost of land, local tax incentives, and electricity prices are low. As such, they are often drawn to arid, water-starved areas in the west such as Arizona where solar and wind energy is plentiful, but usual energy saving techniques are less effective. Virginia Tech researchers estimated that one-fifth of data centers draw water from moderately or highly stressed watersheds in 2021. (Solon, 2021) Managing the impact of data centers must be a local effort from a technological perspective. (Calma, 2021; Walsh, 2022)

There is a trend in such cases were environmental concerns butt heads with economic benefit. In 2017, Google requested a permit to draw 1.5 million gallons of water per day from an aquifer in Berkeley County, South Georgia to cool its local data center as it expanded. Both residents and local conservation groups were concerned about this increased water usage in an already depleted aquifer, where groundwater supplies were running low. Google maintained that it had met with local stakeholders and experts that their water use was in fact sustainable in the area. This two-year push and pull with the South Carolina Coastal Conservation League culminated with an agreement for limited use of local groundwater, with google using alternative water sources instead. (Solon, 2021) A review from Gilmore & Troutman (2021) demonstrates how communities can respond to concerns and how concerns gaining traction, in this case framing the news as a contest between the potential threat to public health and the economic input of Google, with Google presenting themselves as also being a member of the community. The usage of natural resources generally draws public attention to the environmental impact of data centers and embroils companies in local politics. (Gilmore & Troutman, 2021)

Grassroots and The Power Plant

These issues don’t receive as much press in less populated areas, with the tax revenue providing incentive to maintain the data centers, despite negative effects on the local populace. (Glanz, 2012) A case study is of Microsoft’s data centers in Quincy, Washington, attracted by incentivized electricity rates and state level economic incentives in the farming town. However, soon after its 2007 completion, it was a target of a local citizen’s group protesting the use of over 40 diesel generators used for consistent backup power. Microsoft was able to use its status and economic leverage to avoid fines and accelerate construction as local officials attempted to address the issue. In contrast, Microsoft came under scrutiny for its use of diesel generators in 2008 and 2009. Their Bay Area data center was considered one of the largest diesel polluters and the increased emissions prompted review from the State of California.

Local concerns about carbon emissions and power usage can still be drivers of change especially regarding power generation structures. “Data centers don't usually freak people out. But power plants do.” (Miller, 2014). “Not in My Backyard” (NIMBY) regional grassroot organizations who organize local community and environmental groups can build momentum. An example was the collapse of a planned data center in Newark, Deleware, due to its innovation of an onsite power plant. NIMBY organizations organized community leaders, consistently communicated with their base, gained support from both local authorities and environmental groups, and eventually built enough momentum to file a lawsuit against the project. Miller (2014) called out the need for clear communication of future energy consumption and clarity in power operations to avoid such issues in the future.

Similar cases of grassroot organization have been recorded for other gas-powered high energy computing endeavors, such as Bitcoin mining by Greenridge Generation being opposed by citizens of the town of Torrey on Seneca lake. (Jordan, 2021) Another similar phenomenon occurred in Plattsburgh, New York where the cheap electricity drew bitcoin miners who regularly overran the city's quota of cheap power, causing electricity prices for locals to increase by almost 50%. Attention drawn by these organizations can spur local lawmakers to enact policy to mitigate heat output, noise, and energy drain from local data farms and even prevent future operations. (D'Ambrosio, 2019).

This is not to say that these controversies are entirely local - a similar push and pull can be observed at the country level. In Ireland, new data center construction was blocked for several months due to potential blackouts as the data centers eat up a great portion of electricity in the area. This was compounded by environmental concerns that the constructions of new data centers would impact carbon neutrality goals. (Fox, 2022) In response, a business lobby group was created by tech giants such as Google, Apple, and Amazon opposing those country-wide data center construction limits. (Beesley, 2021).

Conclusion

On the policy level, it is shown repeatedly that effective communication with private stakeholders is key in meeting efficiency goals. This should mostly take the form of voluntary programs such as the CoC and the laboratory partnerships as they create demonstratable increases in data center efficiency not only in these larger companies. These must be paired with research and outreach programs to create solutions in a way that is publicly available and independent from economic incentive. (Koronen et al. 2019) The free sharing of information is tantamount to the success of green data centers. (Facebook, n.d.).

-

At the local level, data centers benefit from the reputation of the internet and the lack of visible environmental effect. (Avgerinou et al., 2017) However, there are many cases where use of natural resources affects communities surrounding these data centers. While big technology companies have created effective, comprehensive plans to bolster both their cadre of data centers and their surrounding power generation structures, special attention must be given to the local community circumstances where data centers are built. Community pushback can draw attention to these issues to promote change, while clarity and better communication from companies would help assess data center impact beyond its economic incentives. (D'Ambrosio, 2019; Glanz, 2012; Miller, 2014) The environmental security of local areas must involve advocacy on the local level, both from residents and the companies wishing to station there.

A multidisciplinary approach will be required to manage their energy requirements as data centers only become more essential to computing infrastructure. Keeping the environmental footprint of data centers in check will require a mix of policy implementation, local efforts, and a large knowledge base as efficiency gains from chip optimization and data relocation cap out.

References

Andrae, A., & Edler, T. (2015). On global electricity usage of communication technology: Trends to 2030. *Challenges*, *6*(1), 117–157. https://doi.org/10.3390/challe6010117

Avgerinou, M., Bertoldi, P., & Castellazzi, L. (2017). Trends in Data Centre energy consumption under the European Code of Conduct for Data Centre Energy Efficiency. *Energies*, *10*(10), 1470. https://doi.org/10.3390/en10101470

Beesley, A. (2021, July 12). *Big Tech lobbying coalition against Curbing Data Centres*. The Irish Times. Retrieved November 1, 2021, from https://www.irishtimes.com/news/ireland/irish-news/big-tech-lobbying-coalition-against-curbing-data-centres-1.4617306

Calma, J. (2021, October 27). *Microsoft ramps up plans to make its data centers less thirsty*. The Verge. Retrieved February 20, 2022, from https://www.theverge.com/2021/10/27/22747394/microsoft-data-centers-water-drought-climate-change-energy-emissions

D'Ambrosio, D. (2019, March 19). *Plattsburgh turns back invasion of Bitcoin Miners*. Forbes. Retrieved November 1, 2021, from https://www.forbes.com/sites/danieldambrosio/2018/10/31/plattsburgh-turns-back-invasion-of-bitcoin-miners/?sh=13f065264b5b

Dayarathna, M., Wen, Y., & Fan, R. (2016). Data Center Energy Consumption Modeling: A survey. *IEEE Communications Surveys & Tutorials*, *18*(1), 732–794. https://doi.org/10.1109/comst.2015.2481183

DOE & EPA. (2008, March 19). *Fact sheet on national data center energy efficiency ...* Energy Star. Retrieved February 9, 2022, from https://www.energystar.gov/ia/partners/prod_development/downloads/NDCFactSheet.pdf?18d4-b6e9

Facebook. (n.d.). *About*. Open Compute Project. Retrieved November 1, 2021, from https://www.opencompute.org/about

Fox, K. (2022, January 23). *Ireland's data centers are an economic lifeline. but environmentalists say they're wrecking the planet*. CNN. Retrieved February 9, 2022, from https://www.cnn.com/2022/01/23/tech/ireland-data-centers-climate-intl-cmd/index.html

Gilmore, J. N., & Troutman, B. (2020). Articulating Infrastructure To Water: Agri-culture and google's South Carolina Data Center. *International Journal of Cultural Studies*, *23*(6), 916–931. https://doi.org/10.1177/1367877920913044

Glanz, J. (2012, September 24). *Data Barns in a farm town, gobbling power and flexing muscle*. The New York Times. Retrieved November 1, 2021, from https://www.nytimes.com/2012/09/24/technology/data-centers-in-rural-washington-state-gobble-power.html?ref=technology

Jones, N. (2018). How to stop data centres from gobbling up the world's electricity. *Nature*, *561*(7722), 163–166. https://doi.org/10.1038/d41586-018-06610-y

Joppa, L., & Walsh, N. (2021, July 16). *Made to measure: Sustainability Commitment Progress and updates*. The Official Microsoft Blog. Retrieved November 1, 2021, from https://blogs.microsoft.com/blog/2021/07/14/made-to-measure-sustainability-commitment-progress-and-updates/.

Jordan, J. (2021, July 19). *Environmentalists rally against Greenidge expansion, state lawmakers take notice*. WRFI Community Radio. Retrieved November 1, 2021, from https://www.wrfi.org/2021/04/19/environmentalists-rally-against-greenidge-expansion-state-lawmakers-take-notice/

Kaestner, M. F. (2014). SENSIBLE BYTES: STATES NEED A NEW APPROACH TO JUSTIFY THEIR RECRUITMENT OF INTERNET DATA CENTERS. *William & Mary Environmental Law & Policy Review*, *38*(3), 733–766.

Kleyman, B., Gillooly, B., & Letourneau, K. (2021, May 10). The 2021 State of the Data Center Report. AFCOM.

Kolbert, E. (2021, April 22). *Why bitcoin is bad for the environment*. The New Yorker. Retrieved November 1, 2021, from https://www.newyorker.com/news/daily-comment/why-bitcoin-is-bad-for-the-environment

Koronen, C., Åhman, M., & Nilsson, L. J. (2019). Data Centres in future European Energy Systems—Energy Efficiency, integration and policy. *Energy Efficiency*, *13*(1), 129–144. https://doi.org/10.1007/s12053-019-09833-8

Lawrence Berkeley National Laboratory. (n.d.). *Dale Sartor*. Energy Technologies Area. Retrieved November 1, 2021, from https://eta.lbl.gov/people/dale-sartor

Lawrence Berkeley National Laboratory. (n.d.). *Who we are*. Center of Expertise for Energy Efficiency in Data Centers. Retrieved November 1, 2021, from https://datacenters.lbl.gov/who-we-are

Lucivero, F. (2019). Big Data, big waste? A reflection on the environmental sustainability of Big Data Initiatives. *Science and Engineering Ethics*, *26*(2), 1009–1030. https://doi.org/10.1007/s11948-019-00171-7

Miller, R. (2014, July 22). *Nimby and the Data Center: Lessons from the battle of Newark*. Data Center Knowledge. Retrieved November 1, 2021, from https://www.datacenterknowledge.com/archives/2014/07/22/lessons-of-the-newark-data-center-cogeneration-project-fiasco

Mills, E., Shamshoian, G., Blazek, M., Naughton, P., Seese, R. S., Tschudi, W., & Sartor, D. (2007). The business case for energy management in high-tech industries. *Energy Efficiency*, *1*(1), 5–20. https://doi.org/10.1007/s12053-007-9000-8

Noelle Walsh. (2021, June 1). *Expanding cloud services: Microsoft launches its sustainable Datacenter Region in Arizona*. Azure Blog and Updates | Microsoft Azure. Retrieved November 1, 2021, from https://azure.microsoft.com/en-us/blog/expanding-cloud-services-microsoft-launches-its-sustainable-datacenter-region-in-arizona/

Office (n.d.). Office of Energy Efficiency & Renewable Energy. Data Centers and servers. Energy.gov. https://www.energy.gov/eere/buildings/data-centers-and-servers

Sephton, C. (n.d.). *How many cryptocurrencies are there?* How Many Cryptocurrencies Are There? | Currency.com. Retrieved November 2, 2021, from https://currency.com/how-many-cryptocurrencies-are-there

Shehabi, A., Smith, S., Sartor, D., Brown, R., Herrlin, M., Koomey, J., Masanet, E., Horner, N., Azevedo, I., & Lintner, W. (2016, June 1). *United States Data Center Energy Usage Report*. United States Data Center Energy Usage Report (Technical Report) | OSTI.GOV. Retrieved February 19, 2022, from https://www.osti.gov/servlets/purl/1372902

Siddik, M. A. B., Shehabi, A., & Marston, L. (2021). The environmental footprint of data centers in the United States. *Environmental Research Letters*, *16*(6), 064017. https://doi.org/10.1088/1748-9326/abfba1

Solon, O. (2021, June 19). *Drought-stricken communities push back against data centers*. NBCNews.com. Retrieved February 20, 2022, from https://www.nbcnews.com/tech/internet/drought-stricken-communities-push-back-against-data-centers-n1271344

Song, M., Fisher, R., & Kwoh, Y. (2019). Technological challenges of Green Innovation and Sustainable Resource Management with large scale data. *Technological Forecasting and Social Change*, *144*, 361–368. https://doi.org/10.1016/j.techfore.2018.07.055

Walsh, N. (2022, February 11). *Supporting our customers on the path to net zero: The microsoft cloud and decarbonization*. The Official Microsoft Blog. Retrieved February 20, 2022, from https://blogs.microsoft.com/blog/2021/10/27/supporting-our-customers-on-the-path-to-net-zero-the-microsoft-cloud-and-decarbonization/