Undergraduate Thesis Prospectus

Lancium Front-End Redesign

(technical research project in Computer Science)

The Lack of Green Data Centers

(STS research project)

by

Samuel McBroom

November 17, 2020

technical project collaborators:

Courtney Jacobs

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

signed: _____ date: Nov 30, 2020
Samuel McBroom

signed: _____ date: _____
S. Travis Elliot, Department of Engineering and Society

signed: _____ date: Nov. 24, 2020
Andrew Grimshaw, Department of Computer Science

Introduction

"Data center carbon emissions are a growing global concern" (Bouley, D., 2010). My STS research paper will use the Social Construction of Technology framework to explore the ways data centers could be made more environmentally friendly and why these technological developments are not mainstream practice among large data center providers. With the continual development of more robust hardware and the recent rise in renewable energy sources, data centers have the potential to achieve lower carbon emissions, cut the release of harmful refrigerant gases into the atmosphere, and reduce the amount of electronic waste going to landfills. However, many large technology companies, such as Google, Facebook, and Amazon, continue to operate heavily air conditioned data centers powered primarily by fossil fuels and cutting edge hardware.

My technical project will be to improve the accessibility of Lancium, a green data center provider, by enhancing the capabilities of the web interface. Currently, Lancium customers can either use a command line interface (CLI)-based application programming interface (API) or web interface to create, schedule, run, and analyze computation jobs. The web interface uses a backend that parallels the basic functionality of the API, but is underused due to its lack of functionality compared to the CLI. I will be working with Courtney Jacobs, another UVA Computer Science student, to reimplement the web interface to use the API endpoints and refine the design of the front-end to make the use of Lancium more accessible. This will hopefully lead to an increased usage of the carbon-neutral distributed computation services Lancium provides (Lancium, 2020).

Technical Problem

Most major data center providers give users three primary ways to manage their computation jobs and resources, an API that can be accessed via most major programming languages, a CLI, and a web interface that provides a visual representation of the API. While "power" or more experienced users might feel more comfortable with the reproducibility and flexibility of a CLI or programming API, many new users will find these systems hard to learn. A visual representation, such as a web interface, can be a helpful tool that allows newer or smaller scale users to take advantage of the services offered by a data center provider. For Lancium, this is especially important since they are competing with data center giants like Google Cloud and AWS. Thus, having an easy-to-use web interface is a great way to attract new customers away from Amazon Web Services (AWS) or Google Cloud to a more sustainable option.

The primary technical problem of this project is that Lancium's web interface has only a fraction of the CLI functionality. This excludes most of Lancium's capabilities from users who might prefer or require the used of a web interface. The CLI functions by communicating with an underlying API, which exposes the full functionality of the Lancium computing grid. Our project will be to expand the functionality of the web interface by reconstructing the back-end to use the API endpoints and to design and build new pages to expose the added functionality.

To accomplish this, we will be using a variety of tools and technologies. The current front-end system uses Ruby on Rails, a Ruby-based web application framework, for the back-end. The Rails back-end replicates the functionality of the API, leading to duplicated code that is harder to maintain. We will rewrite the current web interface as a progressive web app (PWA) using Vue.js, a Javascript-based web app framework, making the web interface feel more like

using a native desktop application than a website. The Rails controller will be updated to present the Vue app, which will interact with the API endpoints instead of replicating the functionality. This will not only reduce the complexity of the code, making the system more maintainable, but will allow for the exposure of the entire API if desired. The result will be a web interface that feels native, without the hassle of installing the CLI, and will provide the entire capability of the API in a widely-accessible visual wrapper. One major benefit of using the CLI is the ability to write scripts to automatically create and run jobs, which will not be available using the web interface. Earlier this year we presented mock-ups of our new website design and have begun work on integrating a basic Vue app with the API. We aim to first replicate the existing web functionality, then improve the performance using caching, and finally to finish the aesthetics of the website near the end of 2020.

STS Problem

In his book *Grow a Greener Data Center* (Alger, D., 2009), Douglas Alger defines a green data center as "a computing environment that uses resources in a more efficient manner and has less impact upon people and the environment." As the number and processor density of data centers grow, the impact of their carbon, refrigerant, and technological waste will continue to adversely affect the environment. However, many studies and some existing data centers show that this outcome is not inevitable. Lancium, a provider of Clean Compute Centers, demonstrates a number of successful solutions to the rising environmental impact of data centers (About Lancium, 2020). Why then, are larger data center providers not also adopting these greener measures?

To analyze this question, three harmful impacts of traditional data centers will be considered: carbon emission cost of power, refrigerant use, and technological waste produced. For each area of impact, the economic effects will be considered along with the environmental consequences. In addition, the feasibility of the solutions will discussed for a smaller scale data center, such as Lancium, and massive data center providers, such as AWS. The Social Construction of Technology (SCOT) framework will be used to analyze these impacts and the limited use of greener solutions. A core principle of SCOT is that theories do not succeed because they are true but because they are socially supported. I will apply this principle to understand why the use of greener solutions is rare in the data center industry and how their use might become more prominent in the future.

Lancium is a carbon-negative data center, which it accomplishes by "co-locat[ing] directly at renewable facilities and draw load when power prices are low (or when power would otherwise be wasted) and drop load when power is unavailable or expensive" (About Lancium). They are able to drop load by moving, suspending, and resuming compute jobs when needed. This leads Lancium to be "up to 90% cheaper vs. the leading cloud providers" (About Lancium). However, while massive data center providers would benefit even more from the power cost savings, Lancium's solution limits its market to customers who do not need "always on" power. One solution would be to store the intermittent power when consumption is low and production is high using "energy storage systems such as pump hydro, compressed air, molten salt, flywheel, etc." (Kao, W., 2015). However, these options would create an added cost and might offset the energy savings of using renewables. SCOT clearly demonstrates a link between a short-term switch to greener, intermittent power and potential backlash from customers and shareholders

which might explain why large data center providers choose to source power from always-on, yet harmful fossil fuel sources. Thus, while there is a clear environmental benefit, large data center providers might wait to update their energy technology to renewables due to the added challenge and cost of handling intermittent power availability.

Another change Lancium makes to the traditional data center design is to use moving airflow to cool the racks rather than an air conditioning system. Hydrofluorocarbons (HFCs) are the primary molecule used as cooling agents in refrigeration and air conditioning. While they were developed as a greener alternative to chlorofluorocarbons (CFCs) (Benhadid-Dib, S., Benzaoui, A., 2011), "HFCs are potent greenhouse gases with a global warming potential up to 12400 times that of CO2 per mass unit" (Purohit, P., Höglund-Isaksson, L., Dulac, J. *et.al*, 2020). Thus, it is important to limit the use of these molecules as much as possible. Most data centers are kept at temperatures around 60ºC, in accordance with manufacturer recommendations. This provides some cost protection if the technology fails as manufacturer warranties still apply. However, technology has continued to become more resilient, and studies have found that there is little impact on failure rates from a higher ambient temperature. Google has been experimenting with ambient temperatures in the range of 80ºC to 90ºC. Lancium uses older computers which are cheaper than new machines, which allows them to forgo air conditioning all together. Lancium has data centers located in Houston, Texas, and has measured ambient temperatures of up to 108ºC in their data centers since they began operating. While it would be more environmentally beneficial for large data center providers to adopt a similar system, this solution may not be feasible for all providers. Lancium runs on renewable energy, and when the cost of energy is too high they are able to shut down or move running processes. When

temperatures are high, air conditioning use also rises which raises the cost of electricity. Thus, while Lancium continues to operate some machines in high temperatures, many of their processors remain off or idle during high temperatures, which reduces wear on the machines. Large data centers do not have this flexibly and therefore might still need to use air conditioning systems, but as Google and Facebook have begun to discover, they may be able to keep the temperatures higher than previously though. In addition, data centers have a high cost per square foot, so larger companies will use newer processors with denser computation power. This makes it difficult to cool using even air conditioning, and thus other forms such as water cooling must be used. Switching to older machines, which having a positive impact on the planet and reducing cooling costs might be more expensive due to the larger data centers required. Social factors, as shown by SCOT, can also impact the design of data centers. Most large data center providers originated in the United States, where the use of air conditioning is much more commonplace and viewed necessary compared to other parts of the world. This may have influenced the initial design of data centers and might still play a role in their construction today.

The Clean Compute Centers Lancium operates are not only more sustainable because of the reduced greenhouse gas emissions, but also because they help to reduce technological waste being sent to landfills. “Yearly e-waste generation relays on the average lifespan, no. of total units in check, and mass of the electronic device. Three years is the common lifespan of computers” (Arkam, R., Natasha, Fahad, S. *et al*, 2019). Lancium purchases and uses machines produced years ago instead of purchasing cutting edge processors. Moore’s Law “is the observation that the number of transistors in a dense integrated circuit doubles every two years” (Mollick, E. 2006). However, in recent years there has been a decline in the performance

improvement of integrated circuits. An article titled *The End of Moore's Law* (Theis, T. N., Wong, H.-S. Philip, 2016) "shows a very important exponential trend in information technology that already shows a sharp slowing of progress." The trend being discussed is the switching energy, which is a form of quantitatively measuring the power of different processors. Thus, chips produced today are, on average, only about 25% to 30% more efficient than chips produced three or four years ago. Lancium purchases the older chips, which reduces both e-waste and operating costs because the older computers are priced significantly cheaper than newer machines. Larger data center providers have likely not adopted this strategy due to the increase size of the data center structures, lower computation power density, and possibly public view. Marketing the use of "cutting edge" technology might appeal to more customers and make for better press releases than using chips that are a few years old. However, if larger data center providers were to switch to using older machines, they would help reduce e-waste, and could also take advantage of Lanciums cooling solution described above due to the lower energy density and cheaper processor cost. There are certainly technical factors that influence the choice of processor used, but social factors can also play a major role. SCOT will be helpful in analyzing how the views of customers and shareholders might also play a role in the choice to use cutting-edge technology instead of recycling older machines.

While greener data center designs exist and have been proven effective by companies like Lancium, larger data center providers continue to use traditional designs which continue to pollute our environment. The Social Construction of Technology (SCOT) framework will be useful in analyzing why these greener solutions have not already become wide-spread across the industry. I will apply one of the core principles of SCOT, that technologies are adopted, not

because they are objectively better, but because they are socially supported, to analyze three core areas of improvement data center providers could take to build more sustainable options: the carbon output of power providers, amount and impact of refrigerants used, and e-waste generated.

Conclusion

The technical report will describe the improvements our team makes to the Lancium web interface. It will specifically focus on API integration, PWA implementation, and visual design.

The STS research paper will explore the effectiveness of the green data center design Lancium uses, and how the sustainable solutions could be applied to data center provider giants like Amazon and Google.

The technical report and research paper will help to show how green data center design can be used to help the environment without sacrificing usability and the capabilities required by large data center providers.

References

Lancium. (n.d.). Retrieved October 17, 2020, from https://lancium.com

About Lancium. (n.d.). Retrieved October 17, 2020, from https://portal.lancium.com/about/company

Pinch, T. J., Bijker, W. E. (1984). The Social Construction of Facts and Artifacts: Or How the Sociology of Science and the Sociology of Technology Might Benefit Each Other. Retrieved October 17, 2020, from https://www.jstor.org/stable/285355

Bouley, D. (2010). Estimating a Data Center's Electrical Carbon Footprint. Retrieved October 17, 2020, from https://sg.insight.com/content/dam/insight/en_US/pdfs/apc/apc-estimating-data-centers-carbon-footprint.pdf

Alger, D. (2009, August). Grow a Greener Data Center. Retrieved October 17, 2020, from https://learning.oreilly.com/library/view/grow-a-greener/9781587059919/

Kao, W. (2014, October 24). Renewable and Clean Energy for Data Centers. Retrieved October 17, 2020, from https://doi.org/10.1002/9781118937563.ch30

Benhadid-Dib, S., Benzaoui, A. (2011). Refrigerants and their impact in the environment. Use of the solar energy as the source of energy. Retrieved October 17, 2020, from https://doi.org/10.1016/j.egypro.2011.05.040

Purohit, P., Höglund-Isaksson, L., Dulac, J., Shah, N., Wei, M., Rafaj, P., and Schöpp, W. (2020, October 5). Electricity savings and greenhouse gas emission reductions from global phase-down of hydrofluorocarbons. Retrieved October 17, 2020, from https://doi.org/10.5194/acp-20-11305-2020

Akram, R., Natasha, Fahad, S., Hashmi, M. Z., Wahid, A., Adnan, M., Mubeen, M., Khan, N., Asif Rehmani, M. I., Awais, M., Abbas, M., Shahzad, K., Ahmad, S., Hammad, H. M., Nasim, W. (2019, April 25). Trends of electronic waste pollution and its impact on the global environment and ecosystem. Retrieved October 17, 2020, from https://doi-org.proxy01.its.virginia.edu/10.1007/s11356-019-04998-2

Mollick, E. (2006, August 21). Establishing Moore's Law. Retrieved October 17, 2020, from https://doi.org/10.1109/MAHC.2006.45

Theis, T. N., Wong, H.-S. Philip. (2016, May/June). The End of Moore's Law. Retrieved October 17, 2020, from https://doi.org/10.1109/MCSE.2017.29