A Holistic Evaluation of Computing and the Environment

STS 4500 Prospectus

Department of Engineering

Computer Science

The University of Virginia, Charlottesville

Name: Edward Lue

Technical Advisor:

STS Advisor: Alice Fox

Projected Graduation Date: May 2024

Submission Date: 05/09/2024

A Holistic Evaluation of Computing and the Environment

Overview:

Processing power drives the world. Without it, countless aspects of our lives would be significantly changed. Hardware optimizations allow our processors to keep up with the increasing processing demands of data centers and growing AI models. However, processing also comes at an enormous energy and environmental cost. I want to understand the positive and negative relationships between people, computation, and the environment using the lens of Actor Network Theory. From this, we can learn how computing affects us all and what side effects of computing society may be ignoring.

Positionality:

Throughout my life, I have always been interested in challenging problems. I initially discovered this trait by playing games against my two brothers. In this period of my life, I learned that many simple games can be both challenging and also rewarding to solve. This thinking process led me to enjoy math and science throughout my middle and high school years. In high school, the primary topic I studied was Biology, which is fairly uncommon among Computer Science majors. Over my first few years in college, my interests shifted away from Biology and into computer science. The reason for this is due to a conflict in motivation between me and my peers in BME. I felt that most of them did not care to learn any Biology and instead only wanted to get good grades (to get into med school). I found that it was very important for me to not only solve problems, but also to seek out others from whom I can learn from.

Overall, my background is built upon an unyielding desire to solve challenging problems. I like spending time solving problems, learning to solve problems, and finding other people who

want to solve problems. Currently my main focuses are in competing in algorithmic competitions and researching computer architecture. In both of these areas, I am always able to find new problems to solve and others who want to work to solve them with me. Altogether, my passion for solving challenging problems will contribute to my project where I want to delve deeper into computer architecture to find processor optimizations.

Problematization:

Computers are one of the most prominent technologies today and they provide many benefits to everyone on a daily basis. However, an unfortunate side-effect of computers is that they are also one of the largest power consuming sectors and thus also have a significant environmental impact. I want to investigate as many interactions between computers, users, and the environment as possible to understand how we can make better decisions when designing and deploying computers.

Guiding Question:

In what ways do our interactions with computers impact the environment?

Projected Outcomes:

My research will give a holistic analysis of the interactions between computers, their users, and the environment. After investigating these interactions, I want to draw conclusions about what kinds of computing applications are sustainable and equitable. Theoretically, the downstream benefits of this research would help everyone, since it would bring about the best possible balance between computing and environmental sustainability. More directly, this research would affect computing hardware and software producers and large-scale computing centers.

Technical Project Description:

My technical project will likely focus on computer hardware optimizations. My current research explores speculative optimizations within computer processors. When a processor is reading a stream of instructions, it may require other information to continue computation, for instance if the program branches into two paths based on a computed value. As long as the prediction is relatively accurate, we will observe significant speedups from this addition. This process is called branch prediction, and modern branch prediction schemes have very high accuracy. More recently, a more complex processor optimization called value prediction has arisen. In value prediction, the processor predicts values that haven't been accessed from memory yet. Once again, the speedup of this optimization depends on the accuracy, but as with branch prediction, lots of research has investigated value prediction schemes and modern schemes have high accuracy.

I plan to target speculative optimizations based on value prediction. Predictions that are produced can be propagated to future instructions. In some cases, knowledge of a previous value can actually eliminate multiple future instructions and thus speed up execution. My lab has already looked at some of these optimizations, but the areas I would like to look at are in vector processing units and floating point computation. These areas require additional considerations when propagating predicted values that have not been explored in prior research.

Preliminary Literature Review & Findings:

There has been lots of research that attempts to quantify various environmental impacts of different aspects of large-scale computing. In general, the data indicates that environmental impacts of computing continue to grow in an unsustainable way (Andrae, 2019; Kline et al.,

2016). However, there are a few applications where computing does not follow this general trend. For instance, there is an interesting interaction between internet usage and carbon emissions where the impact is dependent on the amount of human capital (Wang & Xu, 2021). There also are lots of positive impacts of computation that are easy to overlook such as the decreased need for transportation when information can be instantly transported through the internet (Chowdhury, 2010) or environmental calculators (Al-Janabi et al., 2020; Lacoste et al., 2019; Lannelongue et al., 2021). From what I've found so far, most of the research in this area has been performed from a more technical viewpoint. The one exception to this is research on e-waste, the unregulated recycling of computer components (Perkins et al., 2014; Widmer et al., 2005). I found significant research that focuses both on the ethics and inequity of e-waste and the policies related with its regulation. There also is research on regulation of cryptocurrency, however, this research tends to be less focused on the environmental impacts of the regulation (Badea & Mungiu-Pupăzan, 2021). Cryptocurrency also has seen recent technological developments that are focused on sustainability (De Vries, 2023), which shows that these interactions are constantly evolving.

One benefit of my research is that I am using STS techniques to evaluate the topic of computing, which seems to be a fairly uncommon evaluation. In addition, current research seems to focus on more narrow aspects of computing, which is not sufficient for evaluating the overall state of computing and the environment.

STS Project Proposal:

STS explores the interaction between science and society. This interaction occurs in both directions. Science and technology can shape how we live in society, but our actions and thoughts can also affect the direction of science. Therefore, the evolution of both science and

society should be studied together. My project is an STS project because I investigate one of the many ways in which society and technology interact. Computers are arguably the most prominent technology today and environmental sustainability becomes a more and more pressing social issue. Understanding the interaction between computers and the environment will help us make better sociotechnical decisions.

My project will focus on understanding computing from the perspective of environmental sustainability. Because computers are so prevalent today, an important aspect of this research will be covering as many different computing applications as possible. This holistic analysis is a missing piece in sustainable computing research today. A contributing factor to this is that most research in the field is technical, which lends itself to a narrow computing application. Because of this, I will draw from prior research from a variety of authors. There certainly are landmark papers in some of the "subfields" that I'm researching, at least based on citations. Many of the ones I found are in my annotated bibliography.

I chose Actor Network Theory (ANT) as my analytic framework for my project. ANT is a natural choice for my project because it is directly related with defining interactions between actors. The actors I immediately would choose for my project would be the environment, computers, and the general public. As I research, I will likely define more fine-grained actors, but the general idea for the network would stay the same. Another reason for choosing ANT is that it doesn't require as much detail in the interactions as other frameworks. This is good for my project for two reasons. First, I want my project to be very broad, so requiring details may be infeasible for the scale of my project. Secondly, many of the interactions are indirect and passive, so it is harder to define all the details of some of the interactions. For instance, machine learning

provides many small efficiency gains to people on a daily basis. However, the extent of this interaction and the actual "impact" of the interaction is hard to describe.

Currently my chosen method for research will be literature review. I think this would be the most feasible choice for a semester-long project, especially considering I am not experienced in ANT. In addition, I think a literature review is sufficient for the broad analysis I want to do for my project. If I complete my literature review and want to add more analysis, I would probably try to look at the interactions more closely perhaps using narrative analysis and discourse theory to understand how individuals are affected by the environment or certain kinds of computing.

Barriers & Boons

Currently, I realize that the research I have done does not fully connect back to the general public. I have mainly looked at the interactions between computing and the environment, which is arguably the central interaction in my project. Since my background is in Computer Science, especially in hardware, I have significantly less experience with the state of environmental research. Thus, my understandings and conclusions of the environmental impacts on society may fall into some common pitfalls. I can address this by reading more background on the environment and possibly discussing with professors at UVA whose research is related with the environment. Defining the interaction between computing and users may also be challenging. Many of the benefits of computing assist with our daily lives in fairly benign ways, so it will be easy to miss many of these applications. I believe the best way to deal with this would be to simply take more time to review a broader range of computing applications and their environmental impacts. Lastly my background is not in STS, so using Actor Network Theory will also be a challenge. However, I will have assistance with my project through STS 4600.

References

Al-Janabi, S., Mohammad, M., & Al-Sultan, A. (2020). A new method for prediction of air pollution based on intelligent computation. *Soft Computing*, *24*(1), 661–680. https://doi.org/10.1007/s00500-019-04495-1

Andrae, A. (2019). *Prediction Studies of Electricity Use of Global Computing in 2030*. *8*, 27–33.

Badea, L., & Mungiu-Pupăzan, M. C. (2021). The Economic and Environmental Impact of Bitcoin. *IEEE Access*, *9*, 48091–48104. https://doi.org/10.1109/ACCESS.2021.3068636

Chowdhury, G. (2010). Carbon footprint of the knowledge sector: What's the future? *Journal of Documentation*, *66*(6), 934–946. https://doi.org/10.1108/00220411011087878

De Vries, A. (2023). Cryptocurrencies on the road to sustainability: Ethereum paving the way for Bitcoin. *Patterns*, *4*(1), 100633. https://doi.org/10.1016/j.patter.2022.100633

Kline, D., Parshook, N., Ge, X., Brunvand, E., Melhem, R., Chrysanthis, P. K., & Jones, A. K. (2016). Holistically evaluating the environmental impacts in modern computing systems. *2016 Seventh International Green and Sustainable Computing Conference (IGSC)*, 1–8. https://doi.org/10.1109/IGCC.2016.7892605

Lacoste, A., Luccioni, A., Schmidt, V., & Dandres, T. (2019). *Quantifying the Carbon Emissions of Machine Learning* (arXiv:1910.09700). arXiv. https://doi.org/10.48550/arXiv.1910.09700

Lannelongue, L., Grealey, J., & Inouye, M. (2021). Green Algorithms: Quantifying the Carbon Footprint of Computation. *Advanced Science*, *8*(12), 2100707. https://doi.org/10.1002/advs.202100707

Perkins, D. N., Brune Drisse, M.-N., Nxele, T., & Sly, P. D. (2014). E-Waste: A Global Hazard. *Annals of Global Health*, *80*(4), 286–295. https://doi.org/10.1016/j.aogh.2014.10.001

Wang, J., & Xu, Y. (2021). Internet Usage, Human Capital and CO2 Emissions: A Global Perspective. *Sustainability*, *13*(15), Article 15. https://doi.org/10.3390/su13158268

Widmer, R., Oswald-Krapf, H., Sinha-Khetriwal, D., Schnellmann, M., & Böni, H. (2005). Global perspectives on e-waste. *Environmental Impact Assessment Review*, *25*(5), 436–458. https://doi.org/10.1016/j.eiar.2005.04.001