

Can Large-Language Models Help us Better Understand and Teach the Development of Energy-Efficient Software?

Ryan Hasler, Konstantin Läufer,
George K. Thiruvathukal
Loyola University Chicago
Chicago, Illinois, USA
{rhasler,klauder,gthiruvathukal}@luc.edu

Huiyun Peng, Kyle Robinson, Kirsten Davis,
Yung-Hsiang Lu, James C. Davis
Purdue University
West Lafayette, Indiana, USA
{peng397,robin489,kad,yhl,davisjam}@purdue.edu

Abstract

Computing systems are consuming an increasing and unsustainable fraction of society’s energy footprint, notably in data centers. Meanwhile, energy-efficient software engineering techniques are often absent from undergraduate curricula. We propose to develop a learning module for energy-efficient software, suitable for incorporation into an undergraduate software engineering class. There is one major problem with such an endeavor: undergraduate curricula have limited space for mastering energy-related systems programming aspects. To address this problem, we propose to leverage the domain expertise afforded by large language models (LLMs). In our preliminary studies, we observe that LLMs can generate energy-efficient variations of basic linear algebra codes tailored to both ARM64 and AMD64 architectures, as well as unit tests and energy measurement harnesses. On toy examples suitable for classroom use, this approach reduces energy expenditure by 30–90%. These initial experiences give rise to our vision of LLM-based meta-compilers as a tool for students to transform high-level algorithms into efficient, hardware-specific implementations. Complementing this tooling, we will incorporate systems thinking concepts into the learning module so that students can reason both locally and globally about the effects of energy optimizations.

ACM Reference Format:

Ryan Hasler, Konstantin Läufer., George K. Thiruvathukal and Huiyun Peng, Kyle Robinson, Kirsten Davis, Yung-Hsiang Lu, James C. Davis. 2024. Can Large-Language Models Help us Better Understand and Teach the Development of Energy-Efficient Software?. In *Proceedings of (TBD)*. ACM, New York, NY, USA, 3 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

Global climate change poses a serious threat to societal well-being. Data centers, which account for an estimated 4% of annual energy consumption in the United States [10] and 3% in the European Union [6], contribute significantly to this issue. As data center energy consumption is expected to grow, immediate action is needed. Although advancements have been made in enhancing data center computing system performance along various metrics — such




Figure 1: Research overview: we are developing techniques and pedagogy to support a two-part learning module.

as reducing latency [9], increasing throughput [26], and improving parallelism [14] — knowledge of improving their energy efficiency remains limited [12]. *There is little literature on training software engineers to use data center resources with an eye toward energy efficiency* [15, 17]. While energy optimizations have been proposed for individual components [1, 18, 19, 23, 24] and entire systems [5, 8, 22, 25], they have largely focused on hardware. We have identified a need for improvement in software as well. Implementing energy-efficient software methods remains challenging, as current approaches rely on heavyweight design approaches [20, 21], pattern catalogs [11, 16], programming languages [7], or decision frameworks [13].

We argue that all software engineers should have access to training on how to write more energy-efficient software. They should be able to learn not only what lightweight tools are available to them, but also how systems-level thinking can drive energy-conscious design decisions. Our preliminary results suggest that large language models (LLMs) might be a useful tool to realize this vision. We are examining two educational research questions:

- (1) *RQ1*: How can we apply LLMs to help engineers learn to write energy-efficient software?
- (2) *RQ2*: Is systems thinking an effective approach for design-level reasoning about energy-efficient software engineering designs?

2 Initial Experiment: GPT-4 as Optimizer

Goal. Our initial objective was to evaluate whether the LLM ChatGPT-v4 (GPT) could assist in developing energy-efficient versions of common algorithms such as the vector dot product.

Methodology. We asked GPT to: (1) write a basic version of the algorithm in C; (2) make a parallelized version with OpenMP directives; (3) use SIMD mode on ARM64 processors; and finally (4) leverage both OpenMP and SIMD mode for greater performance and efficiency.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
TBD, ,

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM
<https://doi.org/XXXXXXX.XXXXXXX>

This sequence resembles a homework exercise. However, note the substantial background knowledge necessary to complete this activity without assistance: students would need to know C/C++ (not guaranteed in Python- or Java-first curricula); parallel programming libraries such as OpenMP; and specialized hardware instructions on ARM64. Such a homework would thus be unsuitable in many undergraduate curricula. While we do not discourage instructors from offering courses that teach these topics in depth, we believe that any technique that allows less sophisticated students to explore energy-efficient software programming would be valuable to the education community.

LLMs such as GPT provide just this capability. For the dot product task, GPT proposes a familiar function as a starting point:

```
1 double result = 0.0;
2 for(int i = 0; i < size; i++) {
3     result += a[i] * b[i];
4 }
```

To add OpenMP, GPT inserts a pragma and compilation notes:

```
1 #pragma omp parallel for reduction(+:result)
```

To target the ARM64 Neon SIMD architecture, it leverages the specialized instructions for 4-element float vectors.

```
1 for(int i = 0; i < size; i += 4) {
2     a_vec = vld1q_f32(a + i);
3     b_vec = vld1q_f32(b + i);
4     result_vec = vmlaq_f32(result_vec, a_vec, b_vec);
5 }
```

Evaluation. To evaluate the performance of the algorithms generated by GPT, we measured processing time. This decision was influenced by the standard model that reducing processing time reduces energy consumption [3, 4]. Performance testing was conducted on a Raspberry Pi v4 (ARM) processor. GPT’s recommendations resulted in improved latency and energy efficiency (Figure 2). The code and data for the performance tests are available on GitHub at <https://github.com/rhasler1/GPTEfficiencyMeasurements>.

With relatively simple prompts, GPT completed each of these tasks for all three algorithms. As a proof of our concept, we asked an undergraduate student to conduct this activity — he did not know OpenMP nor ARM64 instructions. GPT explained to him everything he needed to know to complete the task.




Figure 2: Preliminary results. GPT’s recommendations improved latency and energy (lower is better) on a Raspberry Pi v4.

3 Future Plans

Building on these preliminary results, we plan to develop an educational module to improve the training of software engineers to think about sustainability and energy efficiency. See Table 1 for an outline of the proposed learning module.

- (1) *Part One: Energy-Efficient Computing with LLMs:* We will describe energy-proportional computing [3, 4], examine the role of specialized hardware for energy efficiency, explain methods for measuring energy usage with industry-grade devices, and discuss energy benchmarking strategies. Additionally, we will demonstrate how LLMs can assist in optimizing algorithms to promote energy efficiency.
- (2) *Part Two: Systems Thinking as a Framework for Energy Efficiency:* We will introduce students to concepts from Systems Thinking [2] and apply these principles to energy-efficient software design. If successful, students will transition from reductionist thinking (e.g., considering only local optimizations) to reasoning about larger changes for greater software energy efficiency.

To disseminate our findings to university students and members of the data center programming workforce, we plan to host a regional workshop in partnership with Argonne National Laboratories. We will also share all materials over the Internet.

Table 1: Outline of proposed learning module.

Unit	Topic
1	Data center concepts
2	Energy usage and measure
3	Optimizations (incl. LLM)
4	Systems thinking for energy

4 Conclusion

The rising energy consumption of software in data centers must not be allowed to contribute to climate change. We propose to introduce a learning module to improve software engineers’ awareness of the energy implications of their code. We suggest that large language models (LLMs) can help undergraduates integrate the many kinds of knowledge needed to reason about energy. While our work is still in its early stages, our preliminary results indicate that these approaches have the potential to help students understand when and how to reduce software energy consumption, offering a promising direction for pedagogy. Given the limited educational literature on enhancing software energy efficiency (especially in data centers), these pedagogical materials could transform how both students and practitioners approach implementing energy-efficient software. Our approach looks to address some of the immediate energy challenges of data centers, and lay the groundwork for long-term sustainability.

Acknowledgments. This work was funded by NSF awards #2343595 and #2343596. This abstract describes part of the research agenda from those awards.

References

- [1] Vlasia Anagnostopoulou, Susmit Biswas, Heba Saadeldien, Alan Savage, Ricardo Bianchini, Tao Yang, Diana Franklin, and Frederic T. Chong. 2012. Barely Alive Memory Servers: Keeping Data Active in a Low-Power State. *ACM Journal on Emerging Technologies in Computing Systems* 8, 4 (2012), 1–20. <https://doi.org/10.1145/2367736.2367742>
- [2] Ross D. Arnold and Jon P. Wade. 2015. A Definition of Systems Thinking: A Systems Approach. *Procedia Computer Science* 44 (2015), 669–678. <https://doi.org/10.1016/j.procs.2015.03.050>
- [3] Luiz André Barroso and Urs Hölzle. 2007. The Case for Energy-Proportional Computing. *Computer* 40, 12 (2007), 33–37. <https://doi.org/10.1109/MC.2007.443>
- [4] Mark T. Bohr and Ian A. Young. 2017. CMOS Scaling Trends and Beyond. *IEEE Micro* 37, 6 (2017), 20–29. <https://doi.org/10.1109/MM.2017.4241347>
- [5] Ce Chi, Kaixuan Ji, Penglei Song, Avinab Marahatta, Shikui Zhang, Fa Zhang, Dehui Qiu, and Zhiyong Liu. 2021. Cooperatively Improving Data Center Energy Efficiency Based on Multi-Agent Deep Reinforcement Learning. *Energies* 14, 8 (2021), 2071. <https://doi.org/10.3390/en14082071>
- [6] European Commission. 2024. Commission adopts EU-wide scheme for rating sustainability of data centres. https://energy.ec.europa.eu/news/commission-adopts-eu-wide-scheme-rating-sustainability-data-centres-2024-03-15_en
- [7] Marco Couto, Rui Pereira, Francisco Ribeiro, Rui Rua, and João Saraiva. 2017. Towards a Green Ranking for Programming Languages. In *Proceedings of the 21st Brazilian Symposium on Programming Languages (Fortaleza, CE, Brazil) (SBLP '17)*. Association for Computing Machinery, New York, NY, USA, Article 7, 8 pages. <https://doi.org/10.1145/3125374.3125382>
- [8] Taliver Heath, Enrique V. Carrera, Eduardo Pinheiro, and Ricardo Bianchini. 2001. *Load Balancing and Unbalancing for Power and Performance in Cluster-Based Systems*. Technical Report. Rutgers University. <https://doi.org/10.7282/T3-AGFW-YT73>
- [9] Nan Hua, Zhizhen Zhong, and Xiaoping Zheng. 2017. Enabling low latency at large-scale data center and high-performance computing interconnect networks using fine-grained all-optical switching technology. In *2017 International Conference on Optical Network Design and Modeling (ONDM)*. 1–4. <https://doi.org/10.23919/ONDM.2017.7958532>
- [10] Electric Power Research Institute. 2024. Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption. <https://www.epri.com/research/products/3002028905>
- [11] Sepideh Maleki, Cuijiao Fu, Arun Banotra, and Ziliang Zong. 2017. Understanding the Impact of Object Oriented Programming and Design Patterns on Energy Efficiency. In *2017 Eighth International Green and Sustainable Computing Conference (IGSC)* (Orlando, FL, 2017-10). IEEE, 1–6. <https://doi.org/10.1109/IGCC.2017.8323605>
- [12] Irene Manotas, Christian Bird, Rui Zhang, David Shepherd, Ciera Jaspan, Caitlin Sadowski, Lori Pollock, and James Clause. 2016. An Empirical Study of Practitioners' Perspectives on Green Software Engineering. In *Proceedings of the 38th International Conference on Software Engineering* (Austin Texas, 2016-05-14). ACM, 237–248. <https://doi.org/10.1145/2884781.2884810>
- [13] Irene Manotas, Lori Pollock, and James Clause. 2014. SEEDS: A Software Engineer's Energy-Optimization Decision Support Framework. In *Proceedings of the 36th International Conference on Software Engineering* (Hyderabad India, 2014-05-31). ACM, 503–514. <https://doi.org/10.1145/2568225.2568297>
- [14] Koushik Mondal and Paramartha Dutta. 2015. Big data parallelism: Challenges in different computational paradigms. In *Proceedings of the 2015 Third International Conference on Computer, Communication, Control and Information Technology (C3IT)*. 1–5. <https://doi.org/10.1109/C3IT.2015.7060186>
- [15] Julia Mullen, Chansup Byun, Vijay Gadepally, Siddharth Samsi, Albert Reuther, and Jeremy Kepner. 2017. Learning by Doing, High Performance Computing Education in the MOOC Era. *J. Parallel and Distrib. Comput.* 105 (2017), 105–115. <https://doi.org/10.1016/j.jpdc.2017.01.015>
- [16] Gustavo Pinto, Kenan Liu, Fernando Castor, and Yu David Liu. 2016. A Comprehensive Study on the Energy Efficiency of Java's Thread-Safe Collections. In *2016 IEEE International Conference on Software Maintenance and Evolution (ICSME)* (Raleigh, NC, USA, 2016-10). IEEE, 20–31. <https://doi.org/10.1109/ICSME.2016.34>
- [17] Apan Qasem. 2019. A Gentle Introduction to Heterogeneous Computing for CS1 Students. In *2019 IEEE/ACM Workshop on Education for High-Performance Computing (EduHPC)* (Denver, CO, USA, 2019-11). IEEE, 10–16. <https://doi.org/10.1109/EduHPC49559.2019.00007>
- [18] Priya Sehgal, Vasily Tarasov, and Erez Zadok. 2010. Evaluating Performance and Energy in File System Server Workloads.. In *FAST*, Vol. 10. 253–266.
- [19] Priya Sehgal, Vasily Tarasov, and Erez Zadok. 2010. Optimizing Energy and Performance for Server-Class File System Workloads. *ACM Transactions on Storage* 6, 3 (2010), 1–31. <https://doi.org/10.1145/1837915.1837918>
- [20] Steven Te Brinke, Somayeh Malakuti, Christoph Bockisch, Lodewijk Bergmans, and Mehmet Akşit. 2013. A Design Method for Modular Energy-Aware Software. In *Proceedings of the 28th Annual ACM Symposium on Applied Computing* (Coimbra Portugal, 2013-03-18). ACM, 1180–1182. <https://doi.org/10.1145/2480362.2480584>
- [21] Steven Te Brinke, Somayeh Malakuti, Christoph Bockisch, Lodewijk Bergmans, Mehmet Akşit, and Shmuel Katz. 2014. A Tool-Supported Approach for Modular Design of Energy-Aware Software. In *Proceedings of the 29th Annual ACM Symposium on Applied Computing* (Gyeongju Republic of Korea, 2014-03-24). ACM, 1206–1212. <https://doi.org/10.1145/2554850.2554964>
- [22] Paul Townend, Stephen Clement, Dan Burdett, Renyu Yang, Joe Shaw, Brad Slater, and Jie Xu. 2019. Invited Paper: Improving Data Center Efficiency Through Holistic Scheduling In Kubernetes. In *2019 IEEE International Conference on Service-Oriented System Engineering (SOSE)* (San Francisco East Bay, CA, USA, 2019-04). IEEE, 156–15610. <https://doi.org/10.1109/SOSE.2019.00030>
- [23] Nedeljko Vasić, Martin Barisits, Vincent Salzgeber, and Dejan Kostic. 2009. Making Cluster Applications Energy-Aware. In *Proceedings of the 1st Workshop on Automated Control for Datacenters and Clouds* (Barcelona Spain, 2009-06-19). ACM, 37–42. <https://doi.org/10.1145/1555271.1555281>
- [24] Mark Weiser, Brent Welch, Alan Demers, and Scott Shenker. 1996. Scheduling for Reduced CPU Energy. In *Mobile Computing*, Tomasz Imielinski and Henry F. Korth (Eds.). Vol. 353. Springer US, 449–471. https://doi.org/10.1007/978-0-585-29603-6_17
- [25] Wei Wu, Wenjia Li, Deify Law, and Woonki Na. 2015. Improving Data Center Energy Efficiency Using a Cyber-physical Systems Approach: Integration of Building Information Modeling and Wireless Sensor Networks. *Procedia Engineering* 118 (2015), 1266–1273. <https://doi.org/10.1016/j.proeng.2015.08.481>
- [26] Jianfeng Zhan, Lixin Zhang, Ninghui Sun, Lei Wang, Zhen Jia, and Chunjie Luo. 2012. High Volume Throughput Computing: Identifying and Characterizing Throughput Oriented Workloads in Data Centers. In *2012 IEEE 26th International Parallel and Distributed Processing Symposium Workshops & PhD Forum*. 1712–1721. <https://doi.org/10.1109/IPDPSW.2012.213>