

# Energy Flow Graph: Modeling Software Energy Consumption

Saurabhsingh Rajput, Tushar Sharma  
Dalhousie University  
Halifax, Canada  
{saurabh,tushar}@dal.ca

## Abstract

The growing energy demands of computational systems necessitate a fundamental shift from performance-centric design to one that treats energy consumption as one of the primary design considerations. Current approaches treat energy consumption as an aggregate, deterministic property, overlooking the path-dependent nature of computation, where different execution paths through the same software consume dramatically different energy. We introduce the **Energy Flow Graph (EFG)**, a formal model that represents computational processes as state-transition systems with energy costs for both states and transitions. EFG enables various applications in software engineering, including static analysis of energy-optimal execution paths and a multiplicative cascade model that predicts combined optimization effects without exhaustive testing. Our early experiments demonstrate EFG’s versatility across domains: in software programs validated through 3.5 million executions, 15.6% of solutions exhibit high path-dependent variance ( $CV > 0.1$ ), while structural optimization reveals up to  $705\times$  energy reduction. In AI pipelines, the cascade model predicts optimization combinations within 5.1% error, enabling selection from 4.2 million possibilities using only 22 measurements. The EFG transforms energy optimization from trial-and-error to systematic analysis, providing a foundation for green software engineering across computational domains.

## Keywords

Energy Efficiency, Static Analysis, Green AI, Software Modeling

## 1 Introduction

The growing energy demands of computation—from mobile devices to massive AI training clusters—have made energy efficiency a primary design constraint [13, 23, 27]. The emerging fields of *green AI* and *green software engineering* aim to optimize software energy consumption [18, 21, 24]. However, current approaches—whether high-level profiling, hardware performance counters, or micro-benchmarking—share a critical limitation: they treat energy as an aggregate, deterministic attribute of the overall process. This perspective overlooks the intrinsic path-dependent nature of software execution, in which the actual energy consumed varies significantly with the specific sequence of state transitions determined by inputs, runtime conditions, and hardware states.

A fundamental disconnect exists in current methodology: while execution is dynamic and path-dependent, standard energy models remain static and aggregate. Current practices, such as annotating code segments with average measurements [7, 23], fail to distinguish between execution paths, mask state-transition costs, and provide no insight into how optimizations interact. While some approaches employ state machines for energy analysis [3, 11] or

trace-based profiling [14], they lack either the path-dependent optimization mechanisms or the thermodynamic foundations necessary for principled energy reasoning. For instance, database queries executing identical plans may consume radically different energy depending on join orders or data locality—nuances lost in aggregate annotations. Similarly, the energy footprints of machine learning (ML) models vary dramatically depending on input characteristics and batch sizes. Without path-level visibility, developers cannot identify energy-intensive execution patterns or predict how code changes affect energy behavior across workloads. This path-insensitive modeling creates a distinct but related challenge: the combinatorial explosion of optimization interactions. When  $n$  optimizations are available, practitioners must choose from  $2^n$  possible combinations. Measuring individual optimizations A and B at 20% and 30% savings provides no basis to predict their combined effect—whether multiplicative, additive, or antagonistic [22]—forcing either exhaustive testing, deployment of suboptimal single-optimization strategies or abandoned optimization efforts.

These practical challenges—*path insensitivity*, *hidden transition costs*, and *combinatorial uncertainty*—stem from a deeper gap: the lack of a formal model connecting abstract software states-transitions to thermodynamic costs. While Landauer’s principle [17] establishes that logical state transitions fundamentally drive physical energy dissipation, current software engineering (SE) treats energy as an abstract metric rather than a physical consequence of computation. Bridging this gap between logical flow and thermodynamic reality is essential for developing robust optimization frameworks.

To address these challenges, we introduce the **Energy Flow Graph (EFG)**, a computational model that bridges *graph theory*, *software engineering*, and *thermodynamic* principles. The EFG abstracts any finite computational process as a state-transition system with explicit energy annotations for both states and transitions. This unified representation enables energy analysis across the system lifecycle: *before* execution for design-time optimization planning, *during* runtime for adaptive decision-making, and *after* execution for post-mortem diagnosis of energy behavior. By modeling both potential paths (pre-execution) and activated paths (post-execution), EFG supports development of optimization strategies grounded in the underlying computational structure, including static analysis to identify energy-optimal paths, predictive models to forecast the combined effects of optimization strategies, and runtime adaptation based on path activation patterns.

In summary, this work makes the following **contributions**. First, we introduce and define EFG, creating a bridge between graph theory, software engineering, and thermodynamic principles. Second, we demonstrate EFG’s utility and versatility as an analytical tool for energy optimizations using two large-scale case studies. The replication package is available online [1].

## 2 Energy Flow Graph (EFG)

Energy Flow Graph (EFG) is a formal, domain-agnostic model that represents computational software processes as state–transition systems annotated with energy costs, enabling systematic reasoning about and optimizing their path-dependent energy consumption.

### 2.1 Definition

*Definition 2.1 (Energy Flow Graph).* An Energy Flow Graph is a directed graph formalized as a tuple  $\mathcal{G} = (V, E, C_s, C_t, \mathcal{P})$ , where:

- $V$  is finite set of **vertices** representing discrete computational states.
- $E \subseteq V \times V$  is a set of directed **edges** representing **transitions** between states.
- $C_s : V \rightarrow \mathbb{R}_{\geq 0}$  is the **state cost function**, assigning energy cost to execute each state.
- $C_t : E \rightarrow \mathbb{R}_{\geq 0}$  is the **transition cost function**, assigning energy overhead to each state transition.
- $\mathcal{P} : E \rightarrow [0, 1]$  is the **transition probability function** (optional, for stochastic analysis), where for each vertex  $v \in V$ ,  $\sum_{(v, v') \in E} \mathcal{P}(v, v') \leq 1$ .

A **valid execution path**  $\pi$  from vertex  $v_1$  to vertex  $v_k$  is a sequence of vertices  $\pi = (v_1, v_2, \dots, v_k)$  such that  $(v_i, v_{i+1}) \in E$  for all  $i \in \{1, \dots, k-1\}$ .

The **total energy** of path  $\pi$  is:

$$E_{\text{total}}(\pi) = \sum_{i=1}^k C_s(v_i) + \sum_{i=1}^{k-1} C_t(v_i, v_{i+1}) \quad (1)$$

This formulation explicitly captures the path-dependent nature of energy consumption: different execution paths through the same code consume different energy based on the states and transitions activated. **Due to space constraints, detailed derivations for equations are provided in the replication package [2].**

### 2.2 Energy Analysis Methods

The EFG framework supports two distinct modes of analysis depending on whether the execution path is fully deterministic or governed by probabilistic uncertainty.

**2.2.1 Deterministic Path Analysis (Bounds & Traces).** This mode applies when analyzing a specific execution trace (e.g., post-mortem profiling) or establishing safety bounds. Since energy costs  $C_s$  and  $C_t$  are non-negative, identifying energy limits is reducible to a shortest/longest path problem on a weighted graph. The Best-Case Energy Consumption (BCEC) and Worst-Case Energy Consumption (WCEC) are naturally defined as  $\text{BCEC} = \min_{\pi \in \Pi} E_{\text{total}}(\pi)$  and  $\text{WCEC} = \max_{\pi \in \Pi} E_{\text{total}}(\pi)$ , where  $\Pi$  is the set of all valid paths from entry to terminal states. For acyclic graphs, these bounds are computable using standard shortest-path algorithms [8]. These bounds are valuable for energy budgeting in resource-constrained, real-time systems.

**2.2.2 Stochastic Path Analysis.** When the precise execution path is not known a priori (e.g., input-dependent control flow), we model the system as a Markov Chain [3], where future states depend only on the current state and transition probabilities  $\mathcal{P}$ , not on execution history. We distinguish between *evaluating* a fixed system and *optimizing* a controllable one.

**Expected Energy (Fixed Policy):** For standard software analysis where transition probabilities are fixed by input distributions (e.g., profiled from representative workloads), the goal is to compute the *Expected Energy Consumption*  $E_{\text{exp}}(v)$  from state  $v$ . This is given by the Bellman equation [4] for a Markov Decision Process (MDP) [5]:

$$E_{\text{exp}}(v) = C_s(v) + \sum_{(v, v_j) \in E} \mathcal{P}(v, v_j) [C_t(v, v_j) + E_{\text{exp}}(v_j)] \quad (2)$$

with the boundary condition  $E_{\text{exp}}(t) = C_s(t)$  for terminal states. For finite state spaces, this yields a linear solvable system [20].

**Policy Optimization:** When the system can choose transitions at runtime (e.g., adaptive algorithm selection, approximate computing, compiler optimization choices), the objective becomes finding an energy-minimizing policy. Instead of probability-weighted sums (as Eq 2), we select the minimum-cost successor at each state, yielding the MDP formulation:

$$E_{\text{opt}}(v) = C_s(v) + \min_{(v, v_j) \in E} [C_t(v, v_j) + E_{\text{opt}}(v_j)] \quad (3)$$

The optimal policy  $\pi^*(v) = \arg \min_{(v, v_j) \in E} [C_t(v, v_j) + E_{\text{opt}}(v_j)]$  identifies the energy-minimizing transition at each state, computable via dynamic programming [5].

### 2.3 Energy Optimization Constructs

EFG enables systematic energy optimization through theoretical constructs grounded in graph theory and probability [4, 5].

**Node Optimization Impact:** Local optimizations to individual nodes create cascading effects throughout the graph. The global impact of optimizing node  $v_i$  from energy cost  $E_i$  to  $E'_i$  is the probability-weighted sum of per-path savings multiplied by execution frequency [26]:

$$\Delta E_{\text{global}} = \sum_{\pi \in \Pi(v_i)} \text{Pr}(\pi) \cdot (E_i - E'_i) \cdot \text{freq}(v_i, \pi) \quad (4)$$

where  $\Pi(v_i)$  is the set of paths containing  $v_i$ ,  $\text{Pr}(\pi)$  is the execution probability of path  $\pi$ , and  $\text{freq}(v_i, \pi)$  is the execution frequency of  $v_i$  within path  $\pi$ . Parameters are derived from representative profiling data or symbolic execution bounds. This formulation reveals that high-frequency nodes in frequently executed paths provide disproportionate optimization leverage.


**Multiplicative Cascade Model:** For orthogonal optimizations targeting independent computational resources (e.g., one reduces CPU cycles, another reduces memory bandwidth), energy savings compound multiplicatively rather than additively [2]:

$$S_{\text{combined}} = 1 - \prod_{i=1}^n (1 - s_i) \quad (5)$$

where  $s_i$  represents the fractional energy saving of optimization  $i$ . This model enables predicting combined optimization effects without exhaustive testing, addressing the combinatorial explosion problem in optimization space exploration.

**Hierarchical Composition:** EFG supports multi-scale analysis through compositional abstraction. A macro-state  $v_{\text{macro}}$  can represent an entire sub-graph  $\mathcal{G}_{\text{micro}} = (V_{\text{micro}}, E_{\text{micro}})$ , with the relationship:

$$C_s(v_{\text{macro}}) = \sum_{v \in V_{\text{micro}}} C_s(v) + \sum_{e \in E_{\text{micro}}} C_t(e) \quad (6)$$



**Figure 1: Energy variance analysis. (a) Path dependency: intra-solution CV distribution; 15.6% with  $CV > 0.1$  validate stochastic modeling. (b) Structural impact: 42% of problems offer  $> 90\%$  energy reduction via topology selection. (c) Thermodynamic variance: power vs. runtime (log scale);  $\sigma = 33.6W$  validates non-uniform state costs  $C_s$ .**

This compositional property ensures energy conservation across abstraction scales and enables modular analysis where optimizations at one level propagate predictably through the hierarchy.

### 3 Applications

To demonstrate EFG’s utility as a predictive tool across computational scales, we present two case studies that validate its core constructs through empirical analysis. These applications map directly to the formal capabilities defined in Section 2.

**Case Study 1** applies EFG to algorithmic software, demonstrating path-dependent energy variance and structural optimization potential through  $\approx 3.5$  million cycle-accurate simulation of 39744 algorithmic kernels. **Case Study 2** applies EFG to AI pipelines, validating the multiplicative cascade model for predicting optimization combinations of AI pipelines.

#### 3.1 Case Study 1: Algorithmic Energy Analysis

**Application Context.** Developers face a fundamental challenge: predicting which algorithmic implementations and input patterns will minimize energy consumption. EFG addresses this by modeling algorithms as state-transition systems where energy varies by execution path. We validate whether EFG’s constructs 2—transition probabilities  $\mathcal{P}$ , state costs  $C_s$ , and structural topology—accurately capture and explain observed energy behaviors.

**Dataset.** We use programs from Shypula *et al.* [25] with 3, 507, 435 verified executions across 1474 problems and 39, 744 unique solutions, each evaluated on an average of 88.3 test inputs. Records capture energy (Joules), instructions, cycles, power (Watts), and IPC. For each solution, we conceptually construct a statement-level EFG where vertices represent states and edges represent control flow; different inputs activate distinct paths  $\pi$ , leading to measurable energy differences  $E_{\text{total}}(\pi)$  (Eq. 1).

**Measurement.** We used Sniper cycle-accurate simulator [15] with McPAT power model [19], modeling AMD EPYC 9554P (Zen 4, 5nm, 1.5 GHz). Simulation ensures: (1) determinism ( $E(\pi)$  depends strictly on path  $\pi$ ), (2) microarchitectural fidelity (captures transition costs  $C_t$ ), and (3) reproducibility.

**Findings. Validating  $\mathcal{P}$  (Path Dependency)** Fig. 1(a) shows energy Coefficient of Variation ( $CV = \sigma/\mu$ ) across 39,744 solutions. The right-skewed distribution reveals a deterministic peak ( $CV \approx 0$ ) and heavy tail: 15.6% exhibit  $CV > 0.1$ , with extreme cases varying

10 $\times$  for identical code validating the need for stochastic analysis (Eq. 2); single “average energy” values mask order-of-magnitude variations driven by probabilistic path selection  $\mathcal{P}$ .

To confirm variance stems from path selection rather than data scaling, we analyzed energy variance ( $\sigma_E$ ) vs. instruction variance ( $\sigma_{\text{inst}}$ ). Table 1 shows high-CV solutions exhibit massive instruction variance (e.g.,  $\sigma_{\text{inst}} \approx 5.5 \times 10^6$ ), confirming inputs drive execution along different paths  $\pi$ . For Problem p00189, sparse graphs activate 50% of vertices while dense graphs force full traversal, explaining the 1.93 $\times$  energy variance.

**Validating Structural Topology (Optimization Potential)** Fig. 1(b) shows energy reduction potential across problems. Over 40% offer  $> 90\%$  reduction through topology selection. For Problem p02907, the worst solution (21.1J) vs. best (0.03J) shows 705 $\times$  reduction. This validates deterministic optimization (Sec. 2.2.1): identifying minimum-energy graph structures yields orders-of-magnitude gains.

**Table 1: Correlation of Energy and Inst. Variance (Top 5)**

Solution ID	Energy CV	$\sigma_E$ (J)	$\sigma_{\text{inst}}$
p03607_07bc...	10.10	0.0037	20,716
p02959_1ef7...	10.10	0.0025	77,836
p02614_6cae...	10.06	0.2583	5,520,196
p02357_8b1d...	10.00	0.0039	23,885
p00670_18d2...	10.00	0.0028	21,736

**Thermodynamic Validation: Non-Uniform State Costs** A common simplification in green software engineering assumes Energy  $\propto$  Time, implying constant power consumption across all computational operations. The EFG model explicitly rejects this by assigning distinct state costs  $C_s(v)$  based on thermodynamic properties of operations (e.g., vector instructions vs. memory stalls). This design choice validates the formal separation of state costs ( $C_s$ ) and transition costs ( $C_t$ ) in Definition 2.1. Fig. 1(c) shows power vs. runtime. Vertical dispersion demonstrates significant power variance ( $\sigma = 33.6W$ , 17% CV) even at fixed runtimes, proving distinct states dissipate energy at different rates. This validates non-uniform state costs  $C_s$  over time-based profiling.

**Theoretical Implications** The empirical evidence validates EFG tuple components (Section 2.1): (1) High CV values prove transition probabilities  $\mathcal{P}$  dominate for 15.6% of algorithms, validating stochastic analysis. (2) 17% power variance confirms states have distinct thermodynamic costs independent of time. (3) 705 $\times$  structural gaps validate topological optimization (Sec. 2.2.1) as critical.

#### 3.2 Case Study 2: AI Pipeline Optimization

**Application Context.** AI practitioners face a combinatorial explosion when selecting optimization combinations across pipeline stages. With 22 available optimizations yielding  $2^{22} \approx 4.2$  million possible combinations, exhaustive testing is prohibitive. EFG addresses this by modeling pipelines as hierarchical state-transition systems where the multiplicative cascade model (Eq. 5) predicts combined savings from just 22 individual measurements, eliminating the need for exhaustive testing.

**Experimental Setup.** We use data from Rajput et al. [22] fine-tuning ModernBERT-base [29] on BigVul [12], evaluating 22 single-knob variants and 8 multi-knob combinations monitored via CodeCarbon [10]. We model the pipeline as a macro-scale EFG with five states (*Data, Model, Training, System, Inference*), each with state costs  $C_s$  and transition costs  $C_t$ .

**Validation.** We apply Eq. 5 to predict savings for four multi-knob variants (V1–V4) combining optimizations across stages. Predicted values are obtained by multiplying individual savings measured in isolation; observed values are empirical measurements. This tests whether EFG forecasts energy without exhaustive testing (Table 2).

**Findings.** Table 2 confirms EFG’s predictive utility. V1 shows near-perfect agreement (-0.6%), confirming true orthogonality between inference engine, gradient checkpointing, LoRA, and FP16. V2 (+1.2%) and V3 (+2.4%) show strong alignment with minor noise, indicating slight resource interaction. V4’s larger delta (+5.1%) reveals non-linear hardware effects from compilation and precision reduction, likely thermal throttling or bandwidth saturation. These deviations are diagnostic: practitioners can identify optimization combinations requiring resource allocation adjustments. The near-zero delta in V1 confirms true orthogonality, providing a benchmark for ideal combinations.

**Table 2: Multiplicative Cascade Model Validation**

Variant (Combined strategies)	Predicted	Observed	$\Delta$
V1 (InfEngine+GradCkpt+LoRA+FP16)	86.0%	86.6%	-0.6%
V2 (Prune+Compile+FP16)	74.9%	73.7%	+1.2%
V3 (Prune+SeqLen+Compile)	83.2%	80.8%	+2.4%
V4 (Compile+FP16)	51.1%	46.0%	+5.1%

## 4 Threats to Validity

**Internal Validity:** Simulation abstracts thermal effects and OS noise. However, this is deliberate to isolate algorithmic energy costs from environmental noise. **External Validity:** Results are specific to the simulated Zen 4 architecture, though path-dependency is an architecture-agnostic property of control flow. **Construct Validity:** The dataset focuses on algorithmic kernels; production software may exhibit different I/O characteristics, a target for future work.

## 5 Related Work

Our work synthesizes formal behavioral modeling, worst-case analysis, and thermodynamic computing across three research streams.

**Energy-Annotated Behavioral Models.** Duarte *et al.* use Labeled Transition Systems (LTS) to analyze software [11], annotating transitions with action costs (e.g., method calls). Baier *et al.* apply Probabilistic Model Checking to MDPs with state–transition rewards to verify energy-utility trade-offs [3]. These approaches exhibit critical limitations that EFG addresses. First, their *additive cost models* ( $\sum \text{cost}_i$ ) treat energy as linearly accumulated, lacking mechanisms to predict optimization interactions. EFG introduces the *Multiplicative Cascade Model* to forecast combined effects of orthogonal optimizations without exhaustive testing. Second, prior work associates costs with either functional actions [11] or abstract rewards [3] without physical grounding. EFG explicitly separates

*State Cost* ( $C_s$ , computational residence) from *Transition Cost* ( $C_t$ , reconfiguration penalty)—critical for modern hardware where data movement and context switches often dominate computational costs. For instance, cache misses and branch mispredictions can incur substantial transition penalties beyond statement execution, which existing models fail to capture. Third, prior work targets specific domains (software traces [11], verification constraints [3]) and suffers state-space explosion for large systems [3]. EFG addresses both limitations through domain-agnostic formulation and *Hierarchical Composition* (Eq. 6), which abstracts sub-graphs as macro-states to enable scalable reasoning from algorithmic kernels to AI pipelines while preserving energy semantics.

**Worst-Case and Hardware-Aware Analysis.** Worst-Case Energy Consumption (WCEC) extends WCET using CFGs and ILP for safe upper bounds [16]. Wagemann *et al.* [28] advanced this with Power-State-Transition Graphs (PSTG), for hardware peripheral states. EFG generalizes from hardware device states to algorithmic states. Beyond WCEC’s deterministic bounds, EFG supports *dual-mode analysis*: deterministic for bounds (Eq. 1) and stochastic for expected values (through Bellman equations [4, 5]), bridging safety-critical and average-case optimization.

**Data-Driven Profiling & Thermodynamic Theory.** Tools like AMALTHEA [14] and CodeCarbon [9] employ trace-based regression for energy estimation. These “black-box” methods measure *what* happened but cannot structurally explain *why*, and fail to *extrapolate* to unseen execution paths. EFG provides a “white-box” structural framework for optimization. Landauer’s principle [17] and reversible computing establish physical energy bounds. EFG operationalizes these by modeling transition costs as thermodynamic expenses, combining empirical measurement with theoretical foundations for predictive optimization.

## 6 Implications

EFG offers key implications across the computing ecosystem. For **software engineers**, it bridges traditional program representations with energy semantics, enabling developers to reason about consumption using familiar graph-based abstractions while incorporating thermodynamic principles. This enables energy-aware programming beyond post-hoc profiling. For **AI practitioners**, the predictive optimization methodology addresses combinatorial explosion by leveraging multiplicative modeling—particularly valuable when optimization spaces grow exponentially. Forecasting optimization interactions without exhaustive testing enables strategic energy planning in resource-constrained environments. For **researchers and tool developers**, EFG establishes energy complexity alongside time and space complexity, opening research directions in energy-bounded computation program analysis. Its compositional nature supports modular analysis across system boundaries, from programs to ML pipeline optimization. The framework’s cross-domain applicability suggests broader utility for system architects building energy-transparent systems, enabling consistent reasoning from micro-scale operations to macro-scale workflows.

## 7 Future Plans

To advance EFG from its current foundation to a comprehensive methodology, we plan to pursue three directions. First, **toolchain**

**development** will create production-ready tooling including automated EFG construction from source code, IDE integration for real-time energy feedback, and compiler plugins for energy-aware code generation, making EFG practical for everyday development workflows. Second, **theoretical extensions** will develop formal methods to verify energy equivalence between program versions and composition rules, maintaining energy guarantees when building systems from components. Third, **empirical validation** will conduct large-scale studies across diverse hardware platforms and application domains, using hardware simulators [6, 15] for deterministic cost model calibration and extending validation to distributed systems and edge computing. These efforts will transform EFG from a conceptual framework into an ecosystem for energy-aware computing readily adoptable across computational domains.

## 8 Conclusions

This paper introduces **Energy Flow Graph (EFG)**, a novel computational model bridging graph theory and thermodynamics with software engineering to analyze and optimize energy consumption. EFG captures the path-dependent nature of energy use through state and transition costs, enabling static analysis of optimal paths and predictive modeling of optimization combinations. Our case studies demonstrate EFG's utility from software optimization to AI pipeline management, showing how energy savings propagate across system boundaries. The model establishes energy complexity as a fundamental consideration, enabling both theoretical foundations and practical methodologies for systematic energy optimization across the software, including AI-enabled software, development and deployment lifecycle.

## References

- [1] anon. 2025. Anonymous Github – anonymous.4open.science. <https://anonymous.4open.science/r/energy-flow-graph-case-study/README.md>. [Accessed 30-09-2025].
- [2] anon. 2025. Anonymous Github – anonymous.4open.science. <https://anonymous.4open.science/r/energy-flow-graph-case-study/derivations.md>. [Accessed 10-01-2026].
- [3] Christel Baier, Clemens Dubsloff, Joachim Klein, Sascha Klüppelholz, and Sascha Wunderlich. 2014. Probabilistic model checking for energy-utility analysis. In *Horizons of the Mind. A Tribute to Prakash Panangaden: Essays Dedicated to Prakash Panangaden on the Occasion of His 60th Birthday*. Springer, 96–123.
- [4] Richard Bellman. 1952. On the theory of dynamic programming. *Proceedings of the national Academy of Sciences* 38, 8 (1952), 716–719.
- [5] Richard Bellman. 1966. Dynamic programming. *science* 153, 3731 (1966), 34–37.
- [6] Nathan Binkert, Bradford Beckmann, Gabriel Black, Steven K Reinhardt, Ali Saidi, Arkaprava Basu, Joel Hestness, Derek R Hower, Tushar Krishna, Somayeh Sardashti, et al. 2011. The gem5 simulator. *ACM SIGARCH computer architecture news* 39, 2 (2011), 1–7.
- [7] Jae-Won Chung, Jiachen Liu, Jeff J. Ma, Ruofan Wu, Oh Jun Kweon, Yuxuan Xia, Zhiyu Wu, and Mosharaf Chowdhury. 2025. The MLENERGY Benchmark: Toward Automated Inference Energy Measurement and Optimization. arXiv:2505.06371 [cs.LG] <https://arxiv.org/abs/2505.06371>
- [8] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. 2022. *Introduction to algorithms*. MIT press.
- [9] Benoit Courty, Victor Schmidt, Goyal-Kamal, MarionCoutarel, Boris Feld, Jérémy Lecourt, LiamConnell, SabAmine, inimaz, supatomic, Mathilde Léval, Luis Blanche, Alexis Cruveiller, ouminasara, Franklin Zhao, Aditya Joshi, Alexis Bogroff, Amine Saboni, Hugues de Lavoreille, Niko Laskaris, Edoardo Abati, Douglas Blank, Ziyao Wang, Armin Catovic, alencon, Michal Stechly, Christian Bauer, Lucas-Otavio, JPW, and MinervaBooks. 2024. *mlco2/codecarbon: v2.4.1*. doi:10.5281/zenodo.11171501
- [10] Benoit Courty, Victor Schmidt, Sasha Luccioni, Goyal-Kamal, MarionCoutarel, Boris Feld, Jérémy Lecourt, LiamConnell, Amine Saboni, Inimaz, supatomic, Mathilde Léval, Luis Blanche, Alexis Cruveiller, ouminasara, Franklin Zhao, Aditya Joshi, Alexis Bogroff, Ziyao Wang, Armin Catovic, Marc Alencon, Michal Stechly, Christian Bauer, Lucas Otávio N. de Araújo, JPW, and MinervaBooks. 2024. *mlco2/codecarbon: v2.4.1*. doi:10.5281/zenodo.11171501
- [11] Lucio Mauro Duarte, Danilo da Silva Alves, Bruno Ramos Toresan, Paulo Henrique Maia, and Davi Silva. 2019. A model-based framework for the analysis of software energy consumption. In *Proceedings of the XXXIII Brazilian Symposium on Software Engineering*. 67–72.
- [12] Jiahao Fan, Yi Li, Shaohua Wang, and Tien N Nguyen. 2020. AC/C++ code vulnerability dataset with code changes and CVE summaries. In *Proceedings of the 17th international conference on mining software repositories*. 508–512.
- [13] Stefanos Georgiou, Maria Kechagia, Tushar Sharma, Federica Sarro, and Ying Zou. 2022. Green ai: Do deep learning frameworks have different costs?. In *Proceedings of the 44th International Conference on Software Engineering*. 1082–1094.
- [14] Ricardo Gomes, Tiago Carvalho, António Barros, and Luis Miguel Pinho. 2022. A model annotation approach for the support of software energy properties management using amalthea. In *2022 IEEE 5th International Conference on Industrial Cyber-Physical Systems (ICPS)*. IEEE, 01–08.
- [15] Wim Heirman, Trevor Carlson, and Lieven Eeckhout. 2012. Sniper: Scalable and accurate parallel multi-core simulation. In *8th International Summer School on Advanced Computer Architecture and Compilation for High-Performance and Embedded Systems (ACACES-2012)*. High-Performance and Embedded Architecture and Compilation Network of . . . , 91–94.
- [16] Ramkumar Jayaseelan, Tulika Mitra, and Xianfeng Li. 2006. Estimating the worst-case energy consumption of embedded software. In *12th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS'06)*. IEEE, 81–90.
- [17] Rolf Landauer. 1961. Irreversibility and heat generation in the computing process. *IBM journal of research and development* 5, 3 (1961), 183–191.
- [18] Sung Une Lee, Niroshinie Fernando, Kevin Lee, and Jean-Guy Schneider. 2024. A survey of energy concerns for software engineering. *Journal of Systems and Software* 210 (2024), 111944.
- [19] Sheng Li, Jung Ho Ahn, Richard D Strong, Jay B Brockman, Dean M Tullsen, and Norman P Jouppi. 2009. McPAT: An integrated power, area, and timing modeling framework for multicore and manycore architectures. In *Proceedings of the 42nd annual IEEE/ACM international symposium on microarchitecture*. 469–480.
- [20] Michael L Littman, Thomas L Dean, and Leslie Pack Kaelbling. 2013. On the complexity of solving Markov decision problems. *arXiv preprint arXiv:1302.4971* (2013).
- [21] 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*. 237–248.
- [22] Saurabhsingh Rajput, Mootez Saad, and Tushar Sharma. 5555. Tu(r)ning AI Green: Exploring Energy Efficiency Cascading with Orthogonal Optimizations. *IEEE Software* 01 (Dec. 5555), 1–7. doi:10.1109/MS.2025.3645090
- [23] Saurabhsingh Rajput, Tim Widmayer, Ziyuan Shang, Maria Kechagia, Federica Sarro, and Tushar Sharma. 2024. Enhancing energy-awareness in deep learning through fine-grained energy measurement. *ACM Transactions on Software Engineering and Methodology* 33, 8 (2024), 1–34.
- [24] Roy Schwartz, Jesse Dodge, Noah A Smith, and Oren Etzioni. 2020. Green ai. *Commun. ACM* 63, 12 (2020), 54–63.
- [25] Alexander Shyupla, Aman Madaan, Yimeng Zeng, Uri Alon, Jacob Gardner, Milad Hashemi, Graham Neubig, Parthasarathy Ranganathan, Osbert Bastani, and Amir Yazdanbakhsh. 2023. Learning performance-improving code edits. *arXiv preprint arXiv:2302.07867* (2023).
- [26] Vivek Tiwari, Sharad Malik, and Andrew Wolfe. 2002. Power analysis of embedded software: A first step towards software power minimization. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 2, 4 (2002), 437–445.
- [27] Roberto Verdecchia, June Sallou, and Luis Cruz. 2023. A systematic review of Green AI. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 13, 4 (2023), e1507.
- [28] Peter Wägemann, Christian Dietrich, Tobias Distler, Peter Ulbrich, and Wolfgang Schröder-Preikschat. 2018. Whole-system worst-case energy-consumption analysis for energy-constrained real-time systems. In *30th Euromicro Conference on Real-Time Systems (ECRTS 2018)*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 24–1.
- [29] Benjamin Warner, Antoine Chaffin, Benjamin Clavié, Orion Weller, Oskar Hallström, Said Taghadouini, Alexis Gallagher, Raja Biswas, Faisal Ladhak, Tom Aarsen, et al. 2024. Smarter, better, faster, longer: A modern bidirectional encoder for fast, memory efficient, and long context finetuning and inference. *arXiv preprint arXiv:2412.13663* (2024).