



33rd NORDUnet 2026 · SUSTAINABILITY SESSION


GREENCODE Constitution:

Agent Principles for Optimizing Code and LLM Inference Energy Use


Esther Hahyeon Kim

DSSS research group · Department of Computer Science, Aalborg University


Pippi




Pippi, beloved across the Nordics



Pippi had a bag of gold coins by Father



Pippi, learned LLM!




✦ Gemini


🌀 ChatGPT

✦ Claude


Pippi, built stock-trading model using LLM!



LLM cost > Earning




LLM cost + Sustainability cost > Earning




Pippi want to built good stock-trading model



Pippi is using LLM agent to upgrade model



LLM, Code and energy



Count generation and repeated execution

Development

- LLM calls
- Code generation and revision
- Builds, tests and experiments

Operation

- Execution of the deployed program
- Its dependencies and I/O
- Its actual workload and lifetime

Compare the complete task at a declared measurement boundary.

Who we are

Research group of Distributed Sustainable Software and Systems (DSSS)
Department of Computer Science, Aalborg University, Denmark

**Software choices shape
physical energy demand.**

Who we are

Research group of Distributed Sustainable Software and Systems (DSSS)
Department of Computer Science, Aalborg University, Denmark



Energy benchmarking

How software design choices matters to real power draw.



Energy-consumption analysis

Finding energy hotspots - where energy goes, and which components dominate.




Process-level measurement tools

Building instruments such as Metrion + RAPL to attribute energy down to the process.

Energy gains must preserve correctness

Benchmarking LLM-generated code across tasks, languages and prompting strategies.



Correct-solution ratio

Prompt engineering can improve the energy efficiency of generated code by up to 91.9% for several models..The best strategy depends on the model.


Up to 91.9%

Reported energy-efficiency improvement in selected cases

Energy principles can guide a coding agent

Detect energy anti-patterns, propose refactorings, then measure the result.

Reported energy reduction (%)



4.7%


GPU energy reduction in
llama.cpp
(3.7% total energy)

**An LLM can improve the
software that runs LLMs.**

Architecture is an energy design choice

Bridging the Gap Between Code Efficiency and Sustainable Design

LLMs recover architectural intent from legacy code.




Evaluate architectural alternatives against the workload.

Esther et al. · Green Architectural Reconstruction · ICSE 2026 Software Architecture BoF

Reconstruction can reduce package power

Green Architectural Reconstruction and Petstore Use Case

Normalized package power



Normalized from the reported maximum reduction.

Up to 37.7%

lower CPU package power
Reported with Intel RAPL


Java EE 7 → asynchronous services
OpenAPI 3.0 intermediate
representation
Up to 1,000 concurrent users

Our research agenda



Data-center cooling & thermals

Optimizing cooling and heat dissipation to cut the energy overhead of running compute at scale.



Hardware-level scheduling

Scheduling and allocating hardware resources to squeeze more useful work from every watt.



Process-level optimization

Profiling processes to locate energy hotspots — then restructuring software to remove them.



Formal verification of green code

Formal-method checks that LLM-generated programs actually meet green-coding criteria.



Let's build greener software together!
If you have any question, want collaboration, please contact me.

Esther Hahyeon Kim

Postdoc

hki@cs.aau.dk

Aalborg University, Denmark