Real-Time ML Integration with CFS to Improve Power Efficiency in Non-Hybrid Linux Systems

Dissanayake M.D.

(MS24035026)

A THESIS

SUBMITTED TO

SRI LANKA INSTITUTE OF INFORMATION TECHNOLOGY

IN PARTIAL FULFILMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

MASTER OF SCIENCE IN INFORMATION TECHNOLOGY

December 2025

I certify that I have read this thesis and that in my opinion it is fully adequate, in scope and in quality, as a thesis for the degree of Master of Science.

4/1/2026

Prof. Samantha Rajapaksha

Approved for MSc. Research Project:

MSc in IT Programme Co-ordinator, SLIIT

Approved for MSc:

Head of Graduate Studies, FoC, SLIIT

DECLARATION

This is to certify that the work is entirely my own and not of any other person, unless explicitly acknowledged (including citation of published and unpublished sources). The work has not previously been submitted in any form to the Sri Lanka Institute of Information Technology or to any other institution for assessment for any other purpose.

Sign: 

Dissanayake M.D.

Date: 30/12/2025

ABSTRACT

Real-Time ML Integration with CFS to Improve Power Efficiency in Non-Hybrid Linux Systems

Dissanayake M.D.

MSc. in Information Technology

Supervisor: Prof. Samantha Rajapaksha

December 2025

Modern non-hybrid Linux systems rely on the Completely Fair Scheduler (CFS) and Dynamic Voltage and Frequency Scaling (DVFS) to balance performance and energy efficiency. However, this default approach has key limitations: CFS distributes tasks evenly across cores without distinguishing between workload types, while DVFS adjusts frequency based on aggregate load. As a result, CPU-bound and I/O-bound tasks often share the same cores, leading to unnecessary frequency boosts and wasted energy for tasks that don't require high frequencies. This thesis proposes a machine learning-enhanced scheduling framework that integrates task-type awareness into non-hybrid Linux systems. A decision tree classifier predicts whether tasks are CPU-bound or I/O-bound using lightweight runtime features such as execution time, waiting time, context switches, priority, and niceness values. Classified CPU-bound tasks are consolidated onto selected cores running at higher frequencies, while I/O-bound tasks are grouped on separate cores maintained at lower frequencies. To prevent thermal hotspots and performance degradation, periodic rotation of CPU-bound and I/O-bound core groups is implemented. The scheduler is evaluated against the default CFS using controlled CPU-intensive, I/O-intensive, and mixed workloads generated with Sysbench and FIO. Energy consumption, performance, and Energy-Delay Product (EDP) are measured using perf and turbostat. Experimental results show that the ML-enhanced scheduler reduces energy consumption and improves EDP significantly, while maintaining performance levels comparable to the default scheduler. These findings highlight the potential of ML-based task classification to improve DVFS utilization and deliver more energy-efficient scheduling in general-purpose, non-hybrid Linux systems.

ACKNOWLEDGEMENT

I wish to extend my deepest gratitude to my supervisor for their invaluable guidance, constructive feedback, and continuous encouragement throughout the course of this research. Their expertise and support have been instrumental in shaping the direction and outcome of this work.

I am also grateful to the faculty and staff of my university for providing the academic environment and resources necessary for the successful completion of this study.

Finally, I would like to acknowledge the unwavering support and understanding of my family and friends, whose encouragement has been a constant source of strength during this research journey.

TABLE OF CONTENTS

List of Figures

List of Tables