

**WORKLOAD-AWARE ADAPTIVE SCHEDULING
ALGORITHMS FOR HETEROGENEOUS MULTI-CORE
SYSTEMS**

Upasakage Saneth Maduranga

229351T

Degree of Master of Science in Computer Science

Department of Computer Science and Engineering

University of Moratuwa
Sri Lanka

June 2024

**WORKLOAD-AWARE ADAPTIVE SCHEDULING
ALGORITHMS FOR HETEROGENEOUS MULTI-CORE
SYSTEMS**

Upasakage Saneth Maduranga

229351T

Thesis submitted in partial fulfillment of the requirements for the degree Master
of Science in Computer Science

Department of Computer Science and Engineering

University of Moratuwa
Sri Lanka

June 2024

DECLARATION OF THE CANDIDATE & SUPERVISOR

I declare that this is my own work and this thesis does not incorporate without acknowledgement any material previously submitted for a Degree or Diploma in any other University or institute of higher learning and to the best of my knowledge and belief it does not contain any material previously published or written by another person except where the acknowledgement is made in the text.

Also, I hereby grant to University of Moratuwa the non-exclusive right to reproduce and distribute my thesis, in whole or in part in print, electronic or other medium. I retain the right to use this content in whole or part in future works (such as articles or books).

Signature: U.S.Maduranga

Date: 30/06/2024

The above candidate has carried out research for the Masters thesis under my supervision.

Name of the supervisor: Dr.Chathuranga Hettiarachchi

Signature of the supervisor:

Date : 5th July 2024

ABSTRACT

WORKLOAD-AWARE ADAPTIVE SCHEDULING ALGORITHMS FOR HETEROGENEOUS MULTI-CORE SYSTEMS

Microprocessors have undergone a revolutionary transformation since their inception, with transistor count exploding from 2,300 in the Intel 4004 to a 90 billion in the Apple M3 Ultra. This surge in processing power coexisted with challenges like performance limitations and power constraints. Multi-core processors emerged as a solution, but the "dark silicon era" highlighted the inability to fully utilize all available cores due to thermal limitations. Heterogeneous Multi-Core Processors (AMPs) addressed this by incorporating cores with diverse capabilities, maximizing power efficiency. AMPs are now all-over in the latest smartphones, tablets, and laptops. Still, the software and scheduler support remain questionable on the AMPs. Existing schedulers, designed for homogeneous systems, struggle to manage the varied capabilities of AMP cores. This research tackles this challenge by proposing a Workload-Aware Adaptive Scheduling (WAAS) algorithm. This novel approach leverages machine learning to create a workload-aware scheduler specifically designed for heterogeneous multi-core systems, aiming to unlock the full potential of AMPs and bridge the gap between processor capabilities and scheduling efficiency.

Keywords: AMP, Heterogeneous Multi Cores, Workload Aware Scheduler

TABLE OF CONTENTS

DECLARATION OF THE CANDIDATE & SUPERVISOR.....	i
ABSTRACT.....	ii
TABLE OF CONTENTS.....	iii
LIST OF FIGURES.....	v
LIST OF TABLES.....	vi
LIST OF ABBREVIATIONS AND ACRONYMS.....	vii
LIST OF APPENDICES.....	ix
INTRODUCTION.....	1
1.1 Background.....	1
1.2 Limitations of uni-processors.....	2
1.3 Introduction to Multi-Core Processors.....	3
1.4 Key innovations in Intel multi-core micro-architecture.....	4
1.4.1 Intel Wide Dynamic Execution.....	4
1.4.2 Intel Intelligent Power Capability.....	4
1.4.3 Intel Advanced Smart Cache.....	5
1.4.4 Intel Smart Memory Access.....	5
1.4.5 Intel Advanced Digital Media Boost.....	5
1.5 The Dark Silicon Era.....	5
1.6 Heterogeneous Multi-Core Processors.....	7
1.6.1 Static Asymmetric Multi-Core.....	8
1.6.2 Dynamic Asymmetric Multi-Core.....	11
1.7 Software Support for Heterogeneous Multi-Core Processors.....	13
1.8 Scheduler Support for Heterogeneous Multi-Core Processors.....	14
1.9 Research Problem.....	15
1.10 Research Objectives.....	16
LITERATURE REVIEW.....	17
2.1 Utilization-Based Scheduling.....	17
2.1.1 Capacity Aware Scheduling.....	17
2.1.2 Energy Aware Scheduling.....	18
2.1.3 Global Task Scheduling.....	18
2.2 Bias Scheduling.....	19
2.3 Fairness Scheduling.....	20
2.4 Offline Profiling and Signature Matching (HAAS).....	21
2.5 Conclusion.....	22
PROBLEM DEFINITION AND PROPOSED SOLUTION.....	23
3.1 The Problem Definition.....	23
3.2 Task Stealing.....	24

3.3 Workload-Aware Adaptive Scheduling(WAAS).....	25
3.3.1 Machine Learning Based Task Allocation.....	25
3.3.2 Priority Based Task Stealing Policy.....	26
IMPLEMENTATION AND VALIDATION.....	28
4.1 The PMCSched Framework.....	28
4.1.1 The PMCTrack Monitoring Tool.....	29
4.1.2 Limitations Of The PMCSched Framework.....	30
4.2 Implementation Of The WAAS.....	30
4.2.1 Task Become Inactive (on_inactive_thread).....	31
4.2.2 Task Become Active (on_active_thread).....	31
4.2.3 Periodic Timer Action (sched_timer_periodic).....	31
4.2.4 Periodic Kthread Action (sched_kthread_periodic).....	32
4.2.5 Task Migration Track (on_migrate_thread).....	32
4.2.6 Task exit from the CPU (on_exit_thread).....	32
4.2.7 The Machine Learning Logic.....	33
4.3 Test Setup - System Specification.....	34
4.4 Selection Of Workloads.....	36
4.5 Evaluation of ML Prediction Logic.....	38
4.6 Evaluation Of WAAS.....	40
4.6.1 Execution Criteria.....	40
4.6.2 Completion Time.....	40
4.6.3 Instructions Per Cycle (IPC).....	42
4.6.4 L1 references, L2 references and LLC references.....	44
4.6.5 Handling Dynamic Workloads.....	46
4.6.6 Handling Gradual Workloads.....	50
CONCLUSIONS AND FUTURE WORK.....	53
5.1 Key Contributions.....	53
5.2 Limitations.....	53
5.2.1 Overhead of the ML integration.....	54
5.3 Future Works.....	54
REFERENCES.....	55
APPENDIX A.....	59
APPENDIX B.....	63
APPENDIX C.....	66
APPENDIX D.....	69
APPENDIX E.....	71

LIST OF FIGURES

Figure	Description	Page
Figure 1.1	Comparison between uni-core & multi-core processors.	4
Figure 1.2	Symmetric multiprocessing vs asymmetric multiprocessing.	7
Figure 1.3	Static and dynamic asymmetric multi-core architectures.	8
Figure 1.4	ARM big.LITTLE static asymmetric multi-core.	9
Figure 1.5	Power-Performance characteristics of A15 & A17 clusters.	10
Figure 1.6	Bahurupi dynamic heterogeneous multi-core.	13
Figure 3.1	Ideal allocation of unbalanced workloads.	23
Figure 3.2	A random allocation of unbalanced workloads.	23
Figure 3.3	Core groups with different clock speeds.	26
Figure 3.4	Cg1 , Cg2 , and Cg3 core groups.	27
Figure 4.1	PMCSched components(in green) inside PMCTrack's architecture.	29
Figure 4.2	ML logic with task active and task inactive flow.	34
Figure 4.3	Small workload generation via the Jmeter.	37
Figure 4.4	Prediction accuracy of the ML model.	39
Figure 4.5	Residual error of the ML model.	39
Figure 4.6	Average task completion time of schedulers.	42
Figure 4.7	IPC data of each scheduler.	44
Figure 4.8	Architecture of the Intel Alder Lake processor.	46
Figure 4.9	Dynamic workload.	47
Figure 4.10	Throughput comparison of schedulers for dynamic workload	48
Figure 4.11	Average throughput comparison of schedulers for dynamic workload	48
Figure 4.12	Aggregate report of WAAS for dynamic workload.	49
Figure 4.13	Gradual workload.	50
Figure 4.14	Throughput comparison of schedulers for gradual workload.	51
Figure 4.15	Staggered graph for throughput comparison of schedulers for gradual workload (CFS and RTMS are given with offsets of 500 TPS and 1000 TPS respectively, for clarity).	51
Figure 4.16	Aggregate report of WAAS for gradual workload.	52

LIST OF TABLES

Table	Description	Page
Table 3.1	Task Stealing Priority of Cg1, Cg2, and Cg3 Core Groups.	27
Table 4.1	Workload Completion Time of Each Scheduler.	41
Table 4.2	IPC Data of Each Scheduler.	43
Table 4.3	L2 and LLC References Data of Each Scheduler.	45

LIST OF ABBREVIATIONS AND ACRONYMS

Abbreviation	Description
AMP	Asymmetric Multi-core Processors
GHz	GigaHertz
KHz	KiloHertz
ILP	Instruction Level Parallelism
IPC	Instructions Per Cycle
TDP	Thermal Design Power
IPC	Instructions Per Clock
DRAM	Dynamic Random Access Memory
TLP	Thread Level Parallelism
L2	Level 2
SSE	Streaming SIMD Extensions
ITRS	International Technology Roadmap for Semiconductors
DVFS	Dynamic Voltage-Frequency Scaling
CMOS	Complementary Metal-Oxide-Semiconductor
I/O	Input/Output
WPU	Wearable Processing Unit
GPU	Graphics Processing Unit
FPGA	Field-Programmable Gate Array
API	Application Programming Interface
PMC	Performance Monitoring Counter
TD	Thread Director
WAAS	Workload-Aware Adaptive Scheduling
CFS	Completely Fair Scheduler
RTMS	Random Task Migration Scheduler
CPU	Central Processing Unit
MIPS	Million Instructions Per Second
EAS	Energy-Aware Scheduling
CAS	Capacity-Aware Scheduling
SMT	Simultaneous Multi Threading

GTS	Global Task Scheduler
PELT	Per-Entity Load Tracking
TLB	Translation Lookaside Buffer
ML	Machine Learning
GB	Giga Bytes
JWT	Json Web Token
ID token	Identifier Token
DB	DataBase
L1	Level 1
LLC	Least Level Cache
L1D	Level 1 Data
L1I	Level 1 Instructions
MLC	Multi-Level Cell

LIST OF APPENDICES

Appendix	Description	Page
Appendix A	Sample Kernel Log Extract During The Data Collection	59
Appendix B	Sample Extraction Of The Initial Data Set	63
Appendix C	Sample Extraction Of The Prepared Data Set	66
Appendix D	Shell Script For Workload Start	69
Appendix E	Sample Data Extract Of Waas For Small Workload Via PMCTrack	71