TASK MIGRATION OPTIMIZATION FOR IMPROVED DARK SILICON MANY-CORE SYSTEMS PERFORMANCE UNDER THERMAL CONSTRAINT

MOHAMMED SULTAN AHMED MOHAMMED

A thesis submitted in fulfilment of the
requirements for the award of the degree of
Doctor of Philosophy

School of Electrical Engineering
Faculty of Engineering
Universiti Teknologi Malaysia

JANUARY 2022

DEDICATION

To my parents, my beloved wife, my lovely kids (Firas, Alaa, and Heba), my brothers and sisters, my friends, and my beloved country "Yemen".

ACKNOWLEDGEMENT

First and foremost, all praise and thanks are due to Allah, and peace and blessings be upon His Messenger, Muhammad (Peace Be Upon Him). I would like to thank Allah SWT for providing me with health, strength, and patience and guiding me through all the difficulties of carrying out this thesis.

I also would like to acknowledge the tremendous support of my supervisor, Assoc. Prof. Dr. Muhammad Nadzir Marsono. I especially thank him for his advice, guidance, patience, and encouragement throughout my PhD journey. He spent a lot of time revising my thesis, helping with setting the research direction, and offering great ideas when I was confused. Without his help, this thesis would not be possible. I am also very thankful to my co-supervisors, Dr. Norlina Paraman and Dr. Ab Al-hadi Ab Rahman, for their precious advice, comments, and encouragement to complete this work.

I would like to thank Hodeidah University and Universiti Teknologi Malaysia (UTM) for their financial support. I am also thankful to the School of Electrical Engineering staff for their assistance during my study at UTM. My sincere appreciation also extends to all my UTM friends and my VeCAD lab colleagues for their friendly environment, moral support, valuable opinions, and assistance on various occasions.

Finally, I wish to express my deepest gratitude and love to my beloved family members, especially my wife, children, mother, father, brothers, and sisters, for their utmost support, patience, and understanding throughout my PhD journey. I also extend my heartfelt thanks to all those who helped me during my doctoral study.

ABSTRACT

Contemporary thermally-constrained techniques for optimizing dark silicon many-core system performance do not use dynamic thermal management efficiently and do not consider the wake-up latency of dark cores. This thesis proposes two improved techniques to overcome these limitations. First is a dynamic thermal-aware performance optimization (DTaPO) technique for dark silicon many-core systems. DTaPO optimizes dark silicon many-core system performance under thermal constraint. The proposed technique utilizes both task migration and dynamic voltage frequency scaling (DVFS) to optimize the performance of a many-core system while keeping the system temperature at a safe operating limit. Task migration puts hot cores in low-power states and moves tasks to cool dark cores to aggressively reduce chip temperature while maintaining high overall system performance. To reduce task migration cold start overhead during task migration, source cores keep their level-2 cache content accessible to the destination cores. Moreover, task migration is limited among cores sharing the last level cache. In the case where task migration cannot be used due to no cool dark core destinations being available, DVFS is used to gradually cool the hot cores by reducing their frequency. Second, a prediction-based early wake-up (PEW) technique for dark cores is proposed to reduce the impact of dark core wake-up latency during the task migration process. An online sliding window-based ridge regression is used as the prediction model. In real-time, the prediction model uses the previous thermal, power, and core status (i.e., active or dark) to predict the subsequent temperature of each core . If task migration is expected to be used in the next control period, PEW puts the dark cores in a power state with low wake-up latency. Thus, it reduces the time needed by the dark cores to start running the migrating tasks, which improves the many-core system's overall performance. Experimental results show that DTaPO improves the system's performance by up to 80% compared to the Optimal Sprinting Patterns technique and reduces the temperature by up to 13.6 °C. Moreover, the comparison results show that the proposed PEW reduces the application execution time by up to 7.9% and 4.1% compared to DTaPO and the fixed-threshold wake-up (FEW) technique, respectively. It also shows that the proposed PEW increases the energy-efficiency by up to 5.5% and 2.3% MIPS/W over DTaPO and FEW, respectively.

ABSTRAK

Teknik kekangan-haba kontemporari untuk mengoptimumkan prestasi sistem banyak-teras silikon gelap tidak menggunakan pengurusan haba dinamik dengan cekap serta tidak mengambil kira kependaman bangun teras gelap. Tesis ini mencadangkan dua teknik yang ditambah baik untuk mengatasi batasan ini. Pertama ialah teknik pengoptimuman prestasi sedar-haba dinamik (DTaPO) bagi sistem banyak-teras silikon gelap. DTaPO mengoptimumkan prestasi sistem banyak-teras silikon gelap di bawah kekangan haba. Teknik yang dicadangkan menggunakan kedua-dua penghijrahan tugas dan penskalaan frekuensi voltan dinamik (DVFS) untuk mengoptimumkan prestasi sistem banyak-teras sambil mengekalkan suhu sistem pada had operasi yang selamat. Penghijrahan tugas meletakkan teras panas dalam keadaan kuasa rendah dan memindahkan tugas kepada teras gelap sejuk bagi mengurangkan suhu cip secara agresif sambil mengekalkan prestasi sistem keseluruhan yang tinggi. Untuk mengurangkan overhed permulaan sejuk semasa pemindahan tugas, teras sumber memastikan kandungan cache tahap-2 mereka boleh diakses oleh teras destinasi. Selain itu, pemindahan tugas adalah terhad di kalangan teras yang berkongsi cache tahap terakhir. Jika pemindahan tugas tidak boleh dilaksanakan kerana tiada destinasi teras gelap sejuk tersedia, DVFS digunakan untuk menyejukkan teras panas secara beransur-ansur dengan menurunkan frekuensi teras panas. Kedua, teknik bangun awal berasaskan ramalan (PEW) untuk teras gelap dicadangkan untuk mengurangkan kesan kependaman bangun teras gelap semasa pemindahan tugas. Regresi rabung berasaskan tingkap gelongsor dalam talian digunakan sebagai model ramalan. Dalam masa nyata, model ramalan menggunakan bacaan haba, kuasa dan status teras (iaitu, aktif atau gelap) untuk meramalkan suhu teras seterusnya. Jika pemindahan tugas dijangka akan digunakan dalam tempoh kawalan seterusnya, PEW meletakkan keadaan kuasa teras gelap dalam keadaan kuasa dengan kependaman bangun yang rendah. Oleh itu, ia dapat mengurangkan masa yang diperlukan oleh teras gelap untuk memulakan tugas yang dipindahkan, yang dapat meningkatkan prestasi keseluruhan sistem banyak-teras. Keputusan eksperimen menunjukkan bahawa DTaPO meningkatkan prestasi sistem sehingga 80% berbanding dengan teknik Pola Pecutan Optimum dan mengurangkan suhu sehingga 13.6 °C. Selain itu, hasil perbandingan menunjukkan bahawa PEW yang dicadangkan mengurangkan masa pelakuan aplikasi masing-masing sehingga 7.9% dan 4.1% berbanding dengan DTaPO dan teknik bangun ambang tetap (FEW). Ia juga menunjukkan bahawa PEW yang dicadangkan meningkatkan kecekapan tenaga masing-masing sehingga 5.5% dan 2.3% MIPS/W berbanding DTaPO dan FEW.

TABLE OF CONTENTS

	TITLE	PAGE
	DECLARATION	**iii**
	DEDICATION	**iv**
	ACKNOWLEDGEMENT	**v**
	ABSTRACT	**vii**
	ABSTRAK	**viii**
	TABLE OF CONTENTS	**ix**
	LIST OF TABLES	**xii**
	LIST OF FIGURES	**xiii**
	LIST OF ABBREVIATIONS	**xvi**
	LIST OF SYMBOLS	**xviii**
CHAPTER 1	**INTRODUCTION**	**1**
1.1	Dark Silicon Many-Core System and Its Thermal Constraints	2
1.2	Problem Statement	3
1.3	Objectives	5
1.4	Research Scope	5
1.5	Thesis Organization	6
CHAPTER 2	**BACKGROUND AND LITERATURE REVIEW**	**9**
2.1	System Performance Background	9
	2.1.1 Performance in Single-Core Era	9
	2.1.2 Performance in Multi/Many-Core Era	12
	2.1.3 Dark Silicon Problem	17
2.2	Dark Silicon Many-Core Systems Optimization Techniques	18
	2.2.1 Architectural Heterogeneity	18
	2.2.2 Resource Management Techniques	20
2.3	Related Work on Dark Silicon Performance Optimization under Thermal Constraint	26

2.4 Related Work on Reducing Dark Silicon Wake-up Latency 29
2.5 Chapter summary 32

CHAPTER 3 METHODOLOGY 35
3.1 Proposed Methodology 35
3.1.1 Dynamic Thermal-Aware Performance Optimization Technique (DTaPO) 37
3.1.2 Prediction-Based Early Wake-Up (PEW) 37
3.2 Research Approach 39
3.3 Design Environment and Simulation Tools 40
3.3.1 LifeSim 41
3.3.2 Benchmarks 44
3.3.3 Performance Validation Setup 49
3.3.4 Performance Evaluation Metrics 50
3.4 Chapter Summary 52

CHAPTER 4 DYNAMIC THERMAL-AWARE PERFORMANCE OPTIMIZATION FOR DARK SILICON MANY-CORE SYSTEMS 53
4.1 Proposed DTaPO Methodology 53
4.1.1 Proposed System Model 54
4.1.2 Proposed DTaPO Algorithm 56
4.2 Experimental Setup 59
4.3 Experimental Results and Discussion 62
4.3.1 Preliminary Results 62
4.3.2 Comprehensive Results 64
4.4 Chapter Summary 70

CHAPTER 5 PREDICTION-BASED EARLY DARK CORES WAKE-UP USING ONLINE RIDGE REGRESSION FOR MANY-CORE SYSTEMS 73
5.1 Proposed PEW Methodology 73
5.1.1 System Model and Problem Definition 74

5.1.2 Prediction Model: Online Ridge Regression 75
5.1.3 Proposed EW Algorithm 77
5.1.4 Complexity analysis 78
5.2 Experimental Setup 79
5.3 Comparative Results and Analysis 83
5.3.1 Preliminary Results 83
5.3.2 Comprehensive Results 84
5.3.3 Significance Test 89
5.4 Chapter Summary 91

CHAPTER 6 CONCLUSION AND FUTURE WORKS 93
6.1 Research Summary 93
6.2 Research Contributions 94
6.3 Directions for Future Works 95

REFERENCES 97
LIST OF PUBLICATIONS 111

LIST OF TABLES

TABLE NO.	TITLE	PAGE
Table 2.1	Power density scaling during and post-Dennard scaling.	12
Table 2.2	A summary of related works for optimizing the dark silicon many-core system performance under thermal constraints.	28
Table 2.3	A summary of related works that used an early wake-up concept.	31
Table 3.1	HotSpot thermal configuration.	45
Table 3.2	A summary of the SPLASH-2 benchmarks used in this work with their input size, application domain, and application description.	47
Table 3.3	A summary of the PARSEC benchmarks used in this work with their input size, application domain, and application description.	48
Table 4.1	The description of DTaPO symbols.	57
Table 4.2	A summary of the system configuration.	61
Table 4.3	Applications characteristics for the comprehensive study.	61
Table 4.4	The average number of times task migration and DVFS have been used in DTaPO at different thermal thresholds.	66
Table 5.1	The definition of the EW algorithm's symbols.	77
Table 5.2	A summary of the system setup.	79
Table 5.3	The MAE for the studied applications at different values of λ.	82
Table 5.4	Average number of task migrations and wake-up accuracy.	84
Table 5.5	A significance test (t-test) for the proposed PEW against FEW and DTaPO.	91

LIST OF FIGURES

FIGURE NO.	TITLE	PAGE
Figure 1.1	An illustration of the technology node's impact on the dark silicon percentage.	2
Figure 1.2	An illustration of transient temperatures and steady-state temperatures.	3
Figure 2.1	Microprocessor trend data over the past five decades.	10
Figure 2.2	Execution speedup according to Amdahl's law of four applications with 80%, 90%, 95%, and 99% parallel parts on a various number of cores.	14
Figure 2.4	Epiphany-V interconnection overview.	15
Figure 2.3	High bandwidth memory architecture.	15
Figure 2.5	Scaling projection of CPU core computation throughput at maximum clock frequency and thermally constrained average frequency.	16
Figure 2.6	Dark silicon estimation for different technology nodes.	17
Figure 2.7	Dark silicon optimization techniques.	18
Figure 2.8	The effect of the core power states on the cache content and wake-up latency.	23
Figure 2.9	An illustration of utilizing all active and dark cores when using task migration.	24
Figure 2.10	Illustration of using an early wake-up of dark cores.	31
Figure 3.1	Overview of the proposed system model.	36
Figure 3.2	The simulated many-core system floorplan.	36
Figure 3.3	The DTaPO technique: (a) performs task migration when temperatures exceed the threshold temperature; (b) performs DVFS when not all dark cores are cool.	38
Figure 3.4	Sniper simulator interval model.	42
Figure 3.5	An overview of McPAT framework.	43
Figure 3.6	Package layers of a ceramic ball grid array (CBGA).	44
Figure 3.7	HotSpot RC thermal model.	45

Figure 3.8	An illustration of a multi-threaded application.	46
Figure 3.9	An overview of the experimental setup of the proposed work.	49
Figure 3.10	An illustration of how Sniper maps tasks into the many-core system at the beginning and how task migration moves the tasks.	50
Figure 4.1	Overview of the system model of the proposed DTaPO.	54
Figure 4.2	An illustration of how DTaPO performs: (a) task migration; (b) and DVFS.	56
Figure 4.3	Experimental setup of the proposed DTaPO.	59
Figure 4.4	Illustration of cores utilization when using task migration.	63
Figure 4.5	The transient temperature of running 16 cores with different DTM techniques: (a) without DTM; and (b) DVFS only.	64
Figure 4.5	The transient temperature of running 16 cores with different DTM techniques (continued): (c) Task migration only; (d) DTaPO.	65
Figure 4.6	Execution time comparison using different DTM techniques.	65
Figure 4.7	Normalized memory accesses with and without clustering at a threshold temperature of 60°C.	66
Figure 4.8	The active and dark core patterns with their transient temperatures: (a) chessboard pattern; (b) contiguous pattern; (c) transient temperatures for chessboard pattern; and (d) transient temperatures for contiguous pattern.	67
Figure 4.9	The transient temperatures of running four applications on 64 cores with different DTM techniques: (a) without DTM; (b) DVFS; and (c) DTaPO.	68
Figure 4.10	The performance slowdown as well as maximum, average, and minimum temperatures using different DTM techniques.	69

Figure 4.11	Comparative results of running ten applications from SPLASH-2 and PARSEC: (a) a comparison of computational efficiency in terms of completion time; (b) a comparison of thermal efficiency in terms of peak temperature.	71
Figure 5.1	Integrating the proposed PEW into the system model.	74
Figure 5.2	An illustration of waiting time for a task when: (a) no early wake-up is used; (b) an early wake-up is used.	75
Figure 5.3	Experimental setup of the proposed PEW.	80
Figure 5.4	Relative completion time of the proposed PEW at different accuracy levels, FEW at different wake-up thresholds, and DTaPO.	83
Figure 5.5	Actual and predicted temperature of core 0 for all studied applications at 65 °C temperature threshold.	85
Figure 5.6	Actual and predicted temperature of cores 1-3 for *Blackscholes* (a-c) and *FFT* (d-f) applications at 65 °C temperature threshold.	86
Figure 5.7	The ratio of serial and parallel phases of the studied applications.	87
Figure 5.8	Relative performance in terms of completion time.	87
Figure 5.9	Relative performance in terms of completion time when the dark state wake-up latency is 261.77 ms.	88
Figure 5.10	Relative performance in terms of MIPS/W.	89
Figure 5.11	The average, maximum, and minimum of: (a) temperature variation between the coldest and hottest cores; (b) transient temperatures of all cores.	90

LIST OF ABBREVIATIONS

ACPI	–	Advanced Configuration and Power Interface
CMGA	–	Ceramic Ball Grid Array
DPM	–	Dynamic Power Management
DRAM	–	Dynamic Random Access Memory
DTaPO	–	Dynamic Thermal-aware Performance Optimization
DTM	–	Dynamic Thermal Management
DVFS	–	Dynamic Voltage Frequency Scaling
EW	–	Early Wake-up
FEW	–	Fixed-threshold Early Wake-up
HBM	–	High Bandwidth Memory
ILP	–	Instruction-Level Parallelism
IPC	–	Instructions Per Cycle
ISA	–	Instruction Set Architecture
LLC	–	Last Level Cache
MAE	–	Mean Absolute Error
MIPS/W	–	Million Instructions Per Second per Watt
NoC	–	Network-on-Chip
OSP	–	Optimal Sprinting Patterns
RMSE	–	Root Mean Square Error
RR	–	Ridge Regression
PEW	–	Prediction-based Early Wake-up
QoS	–	Quality-of-Service
SA	–	Simulated Annealing

TDP	–	Thermal Design Power
TLP	–	Thread-Level Parallelism
TSP	–	Thermal Safe Power
VLSI	–	Very Large Scale Integration

LIST OF SYMBOLS

$\boldsymbol{A}$	–	Set of all active cores
a_i	–	Active core $i \in \boldsymbol{A}$
$\boldsymbol{\beta}$	–	Regression coefficients
$\boldsymbol{D}$	–	Set of all dark cores
d_i	–	Dark core $i \in \boldsymbol{D}$
ζ	–	Frequency level step
ϵ	–	Random errors
$\mathcal{F}_{thr}$	–	Threshold frequency
λ	–	Regression regularization parameter
t_m	–	Makespan time
n	–	Number of samples
p	–	number of features
$\mathbf{T}$	–	Set of the transient temperature of all cores
ε	–	Safe margin value
$\boldsymbol{T}_p$	–	A set of predicted transient temperature of all cores
T_{thr}	–	Threshold temperature
$\boldsymbol{H}$	–	A set of all tasks on cores that exceeded the threshold temperature
t_i	–	The task on core i, $t_i \in \boldsymbol{H}$
w	–	Sliding window size
t_w	–	Task waiting time
$\boldsymbol{X}$	–	Matrix of independent variable
$\boldsymbol{Y}$	–	Dependent variable

CHAPTER 1

INTRODUCTION

The evolution of electronic components has been continuing since the transistor was invented. Moore's law [1] predicted that the number of transistors on a chip would double every two years, while Dennard scaling [2] predicted that power downscaling is proportional to technology size. These two laws were the key concepts for increasing processor performance. As the size of fabrication technology decreases, it becomes more difficult to scale down the supply voltage as it approaches the threshold voltage. Thus, further increases in frequency are infeasible due to increasing power densities that directly contribute to increasing chip temperature. As a solution, more cores on a single chip are integrated to improve processing performance. According to the international technology roadmap for semiconductors (ITRS) [3], the number of cores in future many-core systems will increase to hundreds in mobile devices and thousands in servers.

Although the many-core system is a promising solution for improving processing performance, further reducing technology size without downscaling the supply voltage would increase the many-core system power density, leading to increase chip temperature. To ensure a safe chip operating temperature, only some cores can be active (i.e., turned on) while others should be dark (i.e., turned off). Dynamic Thermal Management (DTM) manages active cores to run at different voltage/frequency levels. Consequently, turning some cores off will prevent a many-core system from fully utilizing a large number of cores for improved processing performance. This problem is called the *dark silicon* problem [4]. It is expected to be significant in future many-core systems [5].

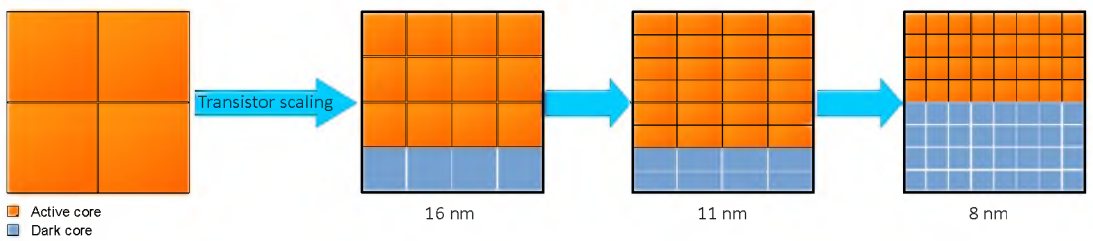

Figure 1.1 An illustration of the technology node's impact on the dark silicon percentage.

1.1 Dark Silicon Many-Core System and Its Thermal Constraints

The dark silicon in modern many-core systems is considered the most significant performance limitation because it prevents many-core systems from utilizing and gaining improved performance from a large number of processing cores. Increasing the number of cores increases the dark silicon ratio, which represents the portion of a chip that cannot be used. Figure 1.1 illustrates the impact of the technology node on the dark silicon ratio. Reducing technology size allows the integration of more cores on a chip. However, integrating more cores means more heat due to increasing power density. Studies in [4, 6] predicted that for the 8 nm technology node, more than half of the cores on a chip would be dark cores. This prediction has prompted researchers to find techniques to maximize multi/many-core system performance for dark silicon while maintaining safe thermal operations.

Thermal constraints are the most significant bottlenecks to maximizing performance, especially in modern chips with extremely high power densities. A chip generates heat as a result of power consumption. However, temperature changes do not occur instantaneously with changes in power consumption due to the thermal capacitance of chip elements [7]. The temperature reaches a steady state when sufficient time has passed with no changes in power. Before reaching the steady-state temperature, the intermediate temperatures are called *transient temperatures*, as shown in Figure 1.2. However, the power consumption in a many-core system is highly changeable with time. It is critical to keep the chip's temperature under a specific critical value called the *threshold temperature*. Otherwise, a permanent failure of the chip may occur due to the high temperatures. DTM techniques should be applied to keep the chip at a safe operating temperature.

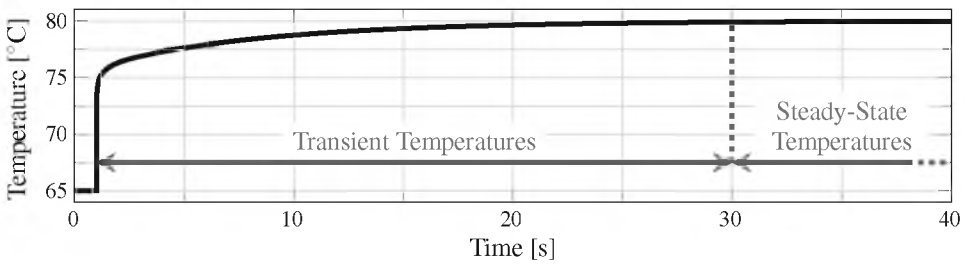

Figure 1.2 An illustration of transient temperatures and steady-state temperatures [8].

1.2 Problem Statement

DTM is an efficient technique for optimizing cores' performance under thermal constraints [9]. Task migration and dynamic voltage frequency scaling (DVFS) are the most commonly used DTM techniques for run-time thermal management. The task migration technique moves tasks from a hot core to a cool core to reduce system temperature and balance core processing loads such that all cores can operate at their maximum frequency under safe thermal constraints. Migrating tasks to dark cores can improve many-core system performance because dark cores are cool and can run at maximum frequency. Moreover, tasks are moved only in one direction after activating the dark core, i.e., a task is moved from an active core to the dark core.

On the other hand, using DVFS can guarantee that the average temperature is not higher than the critical core temperature by reducing the voltage/frequency level, which reduces the power consumption and chip temperature. However, using task migration and DVFS may cause performance degradation due to task migration overhead and downscaling the voltage/frequency level to avoid thermal violations. Thus, the resource management needs to address the task migration overhead due to the cold start cache misses and wake-up latency of dark cores. Additionally, downscaling the voltage/frequency should only be used when no cool cores are available.

Some previous thermal constraint optimization techniques use complex mapping and pattern mechanisms unsuitable for run-time thermal management [5, 10]. Other techniques use a computation sprinting mechanism, which increases cores' frequencies for a short period using DVFS [11–19]. However, sprinting techniques may decrease the chip lifetime due to high peak temperatures. Some techniques that

use DTM, i.e., task migration/DVFS, avoid migration to dark cores due to cold start cache misses overhead [20–22]. However, in modern many-core systems, the core goes into multiple low-power states before it completely shuts off, as implemented in the Intel Xeon Phi [23]. During the first low-power state, its L2 cache stays active, by which the destination core can access data from it rather than from the shared L3 cache or the main memory.

Task migration is widely used for controlling the temperature and improving the utilization of many-core systems. However, a large wake-up latency is required to activate the dark cores, which degrades the overall performance. Some studies used dark cores to migrate the tasks [24–27]. However, all these studies did not provide a solution to the wake-up latency of dark cores due to task migration. Waking up dark cores early just before performing the task migration can improve the overall system performance. Several previous studies proposed an early wake-up of dark cores [28, 29]. However, these studies depend on a fixed threshold to switch the dark cores to an idle state. Switching dark cores to idle mode makes the chip heats up. This results in the DTM being used more frequently, which further degrades the system performance. Moreover, using a fixed wake-up threshold may not suit high thermal fluctuating applications, such as *Fluidanimate* (see Section 5.3.2). Instead of using a fixed wake-up threshold, a simple predictive model can be used to determine when to wake up the dark cores at run-time.

In summary, dark silicon many-core system performance can be improved by addressing the following problems:

1. The lack of efficiency in using DTM techniques to improve dark silicon many-core system performance while keeping system temperature at a safe operating limit.
2. The large wake-up latency of the dark cores when waking from a dark state.

1.3 Objectives

The main aim of this thesis is to improve the overall dark silicon many-core system performance under thermal constraints by utilizing task migration. In specific terms, the objectives of this thesis are as follows.

1. To propose a dynamic thermal-aware performance optimization technique for dark silicon many-core systems. The proposed technique utilizes task migration to aggressively reduce system temperature and maintain a high overall many-core system performance. If task migration cannot be used due to very high core temperatures, DVFS is used to gradually reduce only the hot core frequencies to maintain the system performance while keeping the system temperature within a safe operating limit.
2. To propose a prediction-based early dark cores wake-up technique to reduce the impact of dark cores wake-up latency during the task migration. The proposed technique utilizes a prediction model to predict the future temperatures of cores and an early wake-up algorithm to put the dark cores in a power state with low wake-up latency based on the predicted temperatures.

1.4 Research Scope

This section is an outline of the assumptions and restrictions regarding the work presented in this thesis.

1. The optimization goal of this work is the performance in terms of completion time, while the temperature is used as a thermal constraint.
2. This work focuses on improving many-core performance from the computation perspective. The communication perspective is out of the scope of this work. As many-core system task mapping requires placement consideration, application mapping is not considered in this work.
3. Many-core architecture:

(a) A many-core system with shared memory was used to evaluate the proposed work. The simulated cores have a homogeneous microarchitecture, i.e., they have the same instruction set architecture (ISA), and a heterogeneous frequency, i.e., each core can run at a different frequency.

(b) Many-core system supports multiple power states.

(c) The many-core system supports preemptable tasks that can be stopped and moved to another core to continue the execution.

(d) A mesh network-on-chip (NoC) is used as a communication medium in a many-core system.

4. Simulation environment:

(a) Sniper simulation [30] is used to simulate a many-core system and generate performance traces.

(b) McPAT power model [31] is used to extract the power-related information of the applications.

(c) HotSpot thermal simulator [32] is used to generate temperature traces.

(d) Compute- and memory-intensive applications from SPLASH-2 [33] and PARSEC [34] benchmark suites are used to evaluate the efficiency of the proposed work.

5. Completion time, temperature, mean absolute error (MAE), root mean square error (RMSE), and a million instructions per second per Watt (MIPS/W) are used to evaluate the proposed work.

1.5 Thesis Organization

The rest of the thesis is structured as follows.

Chapter 2 provides the theoretical background and an overview of system performance. It presents a brief introduction to the dark silicon problem. This chapter also reviews different types of dark silicon optimization techniques. It also presents

related works on optimizing the performance of dark silicon many-core systems under thermal constraints.

Chapter 3 describes the proposed methodology for the work done in this thesis. This includes a general overview of the proposed techniques, the step-by-step research approach, the design environment and simulation tools used to validate the proposed work, and the performance metrics used to evaluate and measure the proposed work.

Chapter 4 proposes a dynamic thermal-aware performance optimization (DTaPO) technique for dark silicon many-core systems. The end of this chapter presents the performance of the proposed method. It describes the proposed DTaPO methodology, including the system model and the proposed algorithm. The experimental setup and performance evaluation are presented at the end of this chapter.

Chapter 5 proposes a prediction-based early wake-up (PEW) technique for the dark cores technique that utilizes an online sliding window-based ridge regression (RR) to reduce the wake-up latency of dark cores during the task migration. It describes the proposed PEW methodology, including the online ridge regression prediction model and the early wake-up algorithm. The experimental setup and performance evaluation are presented at the end of this chapter.

Chapter 6 summarizes the research work, highlighting the effectiveness of the proposed work and outlining future research directions.

REFERENCES

1. Moore, G. E. Cramming more components onto integrated circuits. *Electronics*, 1965. 38(8): 114–117.

2. Dennard, R. H., Gaensslen, F. H., Rideout, V. L., Bassous, E. and LeBlanc, A. R. Design of ion-implanted MOSFET's with very small physical dimensions. *IEEE Journal of Solid-State Circuits*, 1974. 9(5): 256–268.

3. ITRS. International technology roadmap for semiconductors 2.0. URL `https://www.semiconductors.org/wp-content/uploads/2018/06/0_2015-ITRS-2.0-Executive-Report-1.pdf`, accessed October 16, 2021.

4. Esmaeilzadeh, H., Blem, E., Amant, R. S., Sankaralingam, K. and Burger, D. Dark silicon and the end of multicore scaling. *Proceedings of the 38th Annual International Symposium on Computer Architecture (ISCA)*. San Jose, CA, USA. 2011. 365–376.

5. Khdr, H., Pagani, S., Shafique, M. and Henkel, J. Thermal constrained resource management for mixed ILP-TLP workloads in dark silicon chips. *Proceedings of the 52nd ACM/EDAC/IEEE Design Automation Conference (DAC)*. San Francisco, CA, USA. 2015. 1–6.

6. Henkel, J., Khdr, H., Pagani, S. and Shafique, M. New trends in dark silicon. *Proceedings of the 52nd ACM/EDAC/IEEE Design Automation Conference (DAC)*. San Francisco, CA, USA. 2015. 1–6.

7. Pagani, S., Manoj, P. S., Jantsch, A. and Henkel, J. Machine learning for power, energy, and thermal management on multicore processors: A survey. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2020. 39(1): 101–116.

8. Pagani, S., Chen, J.-J., Shafique, M. and Henkel, J. *Advanced Techniques for Power, Energy, and Thermal Management for Clustered Manycores*. Springer. 2018.

9. Donald, J. and Martonosi, M. Techniques for multicore thermal management: classification and new exploration. *Proceedings of the 33rd International Symposium on Computer Architecture (ISCA)*. Boston, MA, USA. 2006. 78–88.

10. Wang, J., Chen, Z., Guo, J., Li, Y. and Lu, Z. ACO-based thermal-aware thread-to-core mapping for dark-silicon-constrained CMPs. *IEEE Transactions on Electron Devices*, 2017. 64(3): 930–937.

11. Raghavan, A., Luo, Y., Chandawalla, A., Papaefthymiou, M., Pipe, K. P., Wenisch, T. F. and Martin, M. M. Computational sprinting. *Proceedings of the 18th International Symposium on High-Performance Computer Architecture (HPCA)*. New Orleans, LA, USA. 2012. 1–12.

12. Raghavan, A., Emurian, L., Shao, L., Papaefthymiou, M., Pipe, K. P., Wenisch, T. F. and Martin, M. M. Utilizing dark silicon to save energy with computational sprinting. *IEEE Micro*, 2013. 33(5): 20–28.

13. Zhan, J., Xie, Y. and Sun, G. NoC-sprinting: Interconnect for fine-grained sprinting in the dark silicon era. *Proceedings of the 51st Annual Design Automation Conference (DAC)*. Francisco, CA, USA. 2014. 1–6.

14. Shao, L., Raghavan, A., Emurian, L., Papaefthymiou, M. C., Wenisch, T. F., Martin, M. M. and Pipe, K. P. On-chip phase change heat sinks designed for computational sprinting. *Proceedings of the 2014 Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM)*. San Jose, CA, USA. 2014. 29–34.

15. Kaplan, F. and Coskun, A. K. Adaptive sprinting: How to get the most out of Phase Change based passive cooling. *Proceedings of the 2015 IEEE/ACM International Symposium on Low Power Electronics and Design (ISLPED)*. Rome, Italy. 2015. 37–42.

16. Rezaei, A., Zhao, D., Daneshtalab, M. and Wu, H. Shift sprinting: fine-grained temperature-aware NoC-based MCSoC architecture in dark silicon age. *Proceedings of the 53rd Annual Design Automation Conference (DAC)*. Austin, TX, USA. 2016. 1–6.

17. Morris, N., Stewart, C., Birke, R., Chen, L. and Kelley, J. Early work on modeling computational sprinting. *Proceedings of the 2017 Symposium on Cloud Computing*. Santa Clara, CA, USA. 2017. 661–661.

18. Morris, N., Stewart, C., Chen, L., Birke, R. and Kelley, J. Model-driven computational sprinting. *Proceedings of the 13th EuroSys Conference (EuroSys)*. Porto, Portugal. 2018. 1–13.

19. Wang, J., Chen, Z., Guo, S., Li, Y.-b. and Lu, Z. Optimal Sprinting Pattern in Thermal Constrained CMPs. *IEEE Transactions on Emerging Topics in Computing*, 2021. 9(1): 484–495.

20. Shafique, M., Gnad, D., Garg, S. and Henkel, J. Variability-aware dark silicon management in on-chip many-core systems. *Proceedings of the 2015 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. Grenoble, France. 2015. 387–392.

21. Wang, H., Ma, J., Tan, S. X.-D., Zhang, C., Tang, H., Huang, K. and Zhang, Z. Hierarchical dynamic thermal management method for high-performance many-core microprocessors. *ACM Transactions on Design Automation of Electronic Systems*, 2016. 22(1): 1–21.

22. Wang, H., Zhang, M., Tan, S. X.-D., Zhang, C., Yuan, Y., Huang, K. and Zhang, Z. New power budgeting and thermal management scheme for multi-core systems in dark silicon. *Proceedings of the 29th IEEE International System-on-Chip Conference (SOCC)*. Seattle, WA, USA. 2016. 344–349.

23. Intel. Intel®Xeon Phi Processor x200 Product Family Datasheet, Vol. 1. URL `https://www.intel.com/content/dam/www/public/us/en/documents/datasheets/xeon-phi-processor-x200-product-family-datasheet.pdf`, accessed October 16, 2021.

24. Wang, X., Singh, A. K. and Wen, S. Exploiting dark cores for performance optimization via patterning for many-core chips in the dark silicon era. *Proceedings of the 12th IEEE/ACM International Symposium on Networks-on-Chip (NOCS)*. Turin, Italy. 2018. 1–8.

25. Wang, X., Singh, A., Li, B., Yang, Y., Li, H. and Mak, T. Bubble budgeting: throughput optimization for dynamic workloads by exploiting dark cores in many core systems. *IEEE Transactions on Computers*, 2018. 67(2): 178–192.

26. Wen, S., Wang, X., Singh, A. K., Jiang, Y. and Yang, M. Performance optimization of many-core systems by exploiting task migration and dark core allocation. *IEEE Transactions on Computers*, 2022. 71(1): 92–106.

27. Huang, X., Wang, X., Jiang, Y., Singh, A. K. and Yang, M. Dynamic Allocation/Reallocation of Dark Cores in Many-Core Systems for Improved System Performance. *IEEE Access*, 2020. 8: 165693–165707.

28. Bashir, Q., Shehzad, M. N., Awais, M. N., Farooq, U., Hamayun, M. T. and Ali, I. A scheduling based energy-aware core switching technique to avoid thermal threshold values in multi-core processing systems. *Microprocessors and Microsystems*, 2018. 61: 296–305.

29. Bashir, Q., Shehzad, M. N., Awais, M. N., Baig, S., Dogar, M. G. and Rashid, A. An online temperature-aware scheduling technique to avoid thermal emergencies in multiprocessor systems. *Computers & Electrical Engineering*, 2018. 70: 83–98.

30. Carlson, T. E., Heirman, W. and Eeckhout, L. Sniper: Exploring the level of abstraction for scalable and accurate parallel multi-core simulation. *Proceedings of the 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*. Seatle, WA, USA. 2011. 1–12.

31. Li, S., Ahn, J. H., Strong, R. D., Brockman, J. B., Tullsen, D. M. and Jouppi, N. P. McPAT: an integrated power, area, and timing modeling framework for multicore and manycore architectures. *Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture*. New York, NY, USA. 2009. 469–480.

32. Huang, W., Ghosh, S., Velusamy, S., Sankaranarayanan, K., Skadron, K. and Stan, M. R. HotSpot: A compact thermal modeling methodology for early-stage VLSI design. *IEEE Transactions on Very Large Scale Integration (VLSI) systems*, 2006. 14(5): 501–513.

33. Woo, S. C., Ohara, M., Torrie, E., Singh, J. P. and Gupta, A. The SPLASH-2 programs: Characterization and methodological considerations. *Proceedings of the 22nd Annual International Symposium on Computer Architecture*. Santa Margherita Ligure, Italy. 1995. 24–36.

34. Bienia, C., Kumar, S., Singh, J. P. and Li, K. The PARSEC benchmark suite: Characterization and architectural implications. *Proceedings of the 2008 International Conference on Parallel Architectures and Compilation Techniques (PACT)*. Toronto, ON, Canada. 2008. 72–81.

35. Faggin, F., Hoff, M., Mazor, S. and Shima, M. The history of the 4004. *IEEE Micro*, 1996. 16(6): 10–20.

36. Rupp, K. Microprocessor Trend Data. URL `https://github.com/karlrupp/microprocessor-trend-data`, accessed Oct 16, 2021.

37. Mudge, T. Power: A first-class architectural design constraint. *Computer*, 2001. 34(4): 52–58.

38. Borkar, S. Getting gigascale chips: Challenges and opportunities in continuing moore's law. *Queue*, 2003. 1(7): 26–33.

39. Liao, W., Basile, J. M. and He, L. Leakage power modeling and reduction with data retention. *Proceedings of the 2002 IEEE/ACM International Conference on Computer-Aided Design*. San Jose, California, USA. 2002. 714–719.

40. Taylor, M. B. A landscape of the new dark silicon design regime. *IEEE Micro*, 2013. 33(5): 8–19.

41. Borkar, S. Thousand core chips: a technology perspective. *Proceedings of the 44th Annual Design Automation Conference*. San Diego, California, USA. 2007. 746–749.

42. Lotfi-Kamran, P. and Sarbazi-Azad, H. Chapter One - Dark Silicon and the History of Computing. In: Hurson, A. R. and Sarbazi-Azad, H., eds. *Dark Silicon and Future On-chip Systems*. Elsevier, *Advances in Computers*, vol. 110. 1–33. 2018.

43. Amdahl, G. M. Validity of the single processor approach to achieving large scale computing capabilities. *Proceedings of the Spring Joint Computer Conference*. New York, NY, USA. 1967. 483–485.

44. Mohammed, M. S. and Abandah, G. A. Communication characteristics of parallel shared-memory multicore applications. *Proceedings of the 2015 IEEE Jordan Conference on Applied Electrical Engineering and Computing Technologies (AEECT)*. Amman, Jordan. 2015. 1–6.

45. ITRS. Overall Roadmap Technology Characteristics. URL `https://www.semiconductors.org/wp-content/uploads/2018/08/2003Overall-Roadmap-Technology-Characteristics.pdf`, accessed October 16, 2021.

46. Joint Electron Device Engineering Council (JEDEC). Standard High Bandwidth Memory (HBM) DRAM Specification, JESD235D, 2021. URL `https://www.jedec.org/standards-documents/docs/jesd235a`, accessed October 16, 2021.

47. Dally, W. J. and Towles, B. Route packets, not wires: on-chip inteconnection networks. *Proceedings of the 38th Annual Design Automation Conference*. Las Vegas, Nevada, USA. 2001. 684–689.

48. Kumar, S., Jantsch, A., Soininen, J.-P., Forsell, M., Millberg, M., Oberg, J., Tiensyrja, K. and Hemani, A. A network on chip architecture and design methodology. *Proceedings IEEE Computer Society Annual Symposium on VLSI. New Paradigms for VLSI Systems Design. ISVLSI*. Pittsburgh, PA, USA. 2002. 117–124.

49. Jun, H., Cho, J., Lee, K., Son, H.-Y., Kim, K., Jin, H. and Kim, K. HBM (high bandwidth memory) dram technology and architecture. *Proceedings of the 2017 IEEE International Memory Workshop (IMW)*. Monterey, CA, USA. 2017. 1–4.

50. Olofsson, A. Epiphany-v: A 1024 processor 64-bit risc system-on-chip. *arXiv preprint arXiv:1610.01832*, 2016.

51. Lee, D. U., Kim, K. W., Kim, K. W., Lee, K. S., Byeon, S. J., Kim, J. H., Cho, J. H., Lee, J. and Chun, J. H. A 1.2 V 8 Gb 8-channel 128 GB/s high-bandwidth memory (HBM) stacked DRAM with effective I/O test circuits. *IEEE Journal of Solid-State Circuits*, 2014. 50(1): 191–203.

52. Isci, C., Buyuktosunoglu, A., Cher, C.-Y., Bose, P. and Martonosi, M. An analysis of efficient multi-core global power management policies: Maximizing performance for a given power budget. *Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. Orlando, FL, USA. 2006. 347–358.

53. Ganapathy, D. and Warner, E. Defining thermal design power based on real-world usage models. *Proceedings of the 2008 11th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems*. Orlando, FL, USA. 2008. 1242–1246.

54. Goyal, H. and Agrawal, V. D. Characterizing processors for energy and performance management. *Proceedings of the 16th International Workshop on Microprocessor and SOC Test and Verification (MTV)*. 2015. 67–72.

55. Pagani, S., Chen, J.-J., Shafique, M. and Henkel, J. MatEx: Efficient transient and peak temperature computation for compact thermal models. *Proceedings of the 2015 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. renoble, France. 2015. 1515–1520.

56. IRDS. International Roadmap for Devices and Systems (IRDS) 2020 Edition (More Moore). URL `https://irds.ieee.org/editions/2020`, accessed October 16, 2021.

57. Shafique, M. and Garg, S. Computing in the Dark Silicon Era: Current Trends and Research Challenges. *IEEE Design & Test*, 2017. 34(2): 8–23.

58. Venkatesh, G., Sampson, J., Goulding, N., Garcia, S., Bryksin, V., Lugo-Martinez, J., Swanson, S. and Taylor, M. B. Conservation cores: reducing the energy of mature computations. *ACM Sigplan Notices*, 2010. 45(3): 205–218.

59. Goulding-Hotta, N., Sampson, J., Zheng, Q., Bhatt, V., Auricchio, J., Swanson, S. and Taylor, M. B. Greendroid: An architecture for the dark silicon age. *Proceedings of the 17th Asia and South Pacific Design Automation Conference*. Sydney, NSW, Australia. 2012. 100–105.

60. Venkatesh, G., Sampson, J., Goulding-Hotta, N., Venkata, S. K., Taylor, M. B. and Swanson, S. QsCores: Trading dark silicon for scalable energy efficiency with quasi-specific cores. *Proceedings of the 44th Annual*

IEEE/ACM International Symposium on Microarchitecture (MICRO). Porto Alegre, Brazil. 2011. 163–174.

61. Turakhia, Y., Raghunathan, B., Garg, S. and Marculescu, D. HaDeS: Architectural synthesis for heterogeneous dark silicon chip multi-processors. *Proceedings of the 50th ACM/EDAC/IEEE Design Automation Conference (DAC)*. Austin, TX, USA. 2013. 1–7.

62. Yang, L., Liu, W., Jiang, W., Li, M., Chen, P. and Sha, E. H.-M. FoToNoC: A folded torus-like network-on-chip based many-core systems-on-chip in the dark silicon era. *IEEE Transactions on Parallel and Distributed Systems*, 2017. 28(7): 1905–1918.

63. Yang, L., Liu, W., Jiang, W., Chen, C., Li, M., Chen, P. and Edwin, H. Hardware-software collaboration for dark silicon heterogeneous many-core systems. *Future Generation Computer Systems*, 2017. 68: 234–247.

64. Greenhalgh, P. Big. Little Processing with Arm Cortex-A15 & Cortex-A7. URL `https://www.eetimes.com/big-little-processing-with-arm-cortex-a15-cortex-a7`, accessed October 16, 2021.

65. Intel. Intel® Xeon® Platinum 8380 Processor. URL `https://www.intel.com/content/www/us/en/products/sku/212287/intel-xeon-platinum-8380-processor-60m-cache-2-30-ghz/specifications.html`, accessed October 16, 2021.

66. Muthukaruppan, T. S., Pricopi, M., Venkataramani, V., Mitra, T. and Vishin, S. Hierarchical power management for asymmetric multi-core in dark silicon era. *Proceedings of the 50th Annual Design Automation Conference*. Austin, TX, USA. 2013. 1–9.

67. Raghunathan, B., Turakhia, Y., Garg, S. and Marculescu, D. Cherry-picking: exploiting process variations in dark-silicon homogeneous chip multi-processors. *Proceedings of the 2013 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. Grenoble, France. 2013. 39–44.

68. Raghunathan, B. and Garg, S. Job arrival rate aware scheduling for asymmetric multi-core servers in the dark silicon era. *Proceedings of the*

2014 International Conference on Hardware/Software Codesign and System Synthesis (CODES). New Delhi, India. 2014. 1–9.

69. Pagani, S., Khdr, H., Chen, J.-J., Shafique, M., Li, M. and Henkel, J. Thermal safe power (TSP): Efficient power budgeting for heterogeneous manycore systems in dark silicon. *IEEE Transactions on Computers*, 2016. 66(1): 147–162.

70. Kanduri, A., Haghbayan, M.-H., Rahmani, A.-M., Liljeberg, P., Jantsch, A. and Tenhunen, H. Dark Silicon Aware Runtime Mapping for Many-core Systems: A Patterning Approach. *Proceedings of the 33rd IEEE International Conference on Computer Design (ICCD)*. New York, NY, USA. 2015. 573–580.

71. Rahmani, A. M., Haghbayan, M.-H., Miele, A., Liljeberg, P., Jantsch, A. and Tenhunen, H. Reliability-aware runtime power management for many-core systems in the dark silicon era. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 2016. 25(2): 427–440.

72. Rapp, M., Sagi, M., Pathania, A., Herkersdorf, A. and Henkel, J. Power-and cache-aware task mapping with dynamic power budgeting for many-cores. *IEEE Transactions on Computers*, 2019. 69(1): 1–13.

73. Shafique, M., Garg, S., Henkel, J. and Marculescu, D. The EDA challenges in the dark silicon era: Temperature, reliability, and variability perspectives. *Proceedings of the 51st Annual Design Automation Conference*. Francisco, CA, USA. 2014. 1–6.

74. Unified Extensible Firmware Interface Forum. Advanced Configuration and Power Interface (ACPI) Specification. URL `https://www.uefi.org/specifications`, accessed October 16, 2021.

75. Zhang, Y., Chakrabarty, K. and Swaminathan, V. Energy-aware fault tolerance in fixed-priority real-time embedded systems. *Proceedings of the 2003 International Conference on Computer Aided Design*. San Jose, CA, USA. 2003. 209–213.

76. Van Craeynest, K., Akram, S., Heirman, W., Jaleel, A. and Eeckhout, L. Fairness-aware scheduling on single-ISA heterogeneous multi-cores.

Proceedings of the 22nd international conference on Parallel architectures and compilation techniques. Edinburgh, UK. 2013. 177–187.

77. Dietrich, B., Nunna, S., Goswami, D., Chakraborty, S. and Gries, M. LMS-based low-complexity game workload prediction for DVFS. *Proceedings of the 2010 IEEE International Conference on Computer Design*. Amsterdam, Netherlands. 2010. 417–424.

78. Shen, H., Lu, J. and Qiu, Q. Learning based DVFS for simultaneous temperature, performance and energy management. *Proceedings of 13th International Symposium on Quality Electronic Design (ISQED)*. Santa Clara, CA, USA. 2012. 747–754.

79. Yang, S., Shafik, R. A., Merrett, G. V., Stott, E., Levine, J. M., Davis, J. and Al-Hashimi, B. M. Adaptive energy minimization of embedded heterogeneous systems using regression-based learning. *Proceedings of the 25th International Workshop on Power and Timing Modeling, Optimization and Simulation (PATMOS)*. Salvador, Brazil. 2015. 103–110.

80. Mahbub ul Islam, F. M. and Lin, M. A framework for learning based DVFS technique selection and frequency scaling for multi-core real-time systems. *Proceedings of the 17th International Conference on High Performance Computing and Communications, 7th International Symposium on Cyberspace Safety and Security, and 12th International Conference on Embedded Software and Systems*. New York, NY, USA. 2015. 721–726.

81. Cochran, R. and Reda, S. Consistent runtime thermal prediction and control through workload phase detection. *Proceedings of the Design Automation Conference*. Anaheim, CA, USA. 2010. 62–67.

82. Khanna, R., John, J. and Rangarajan, T. Phase-aware predictive thermal modeling for proactive load-balancing of compute clusters. *Proceedings of the 2012 International Conference on Energy Aware Computing*. Guzelyurt, Northern Cyprus. 2012. 1–6.

83. Bartolini, A., Cacciari, M., Tilli, A. and Benini, L. Thermal and energy management of high-performance multicores: Distributed and self-calibrating model-predictive controller. *IEEE Transactions on Parallel and Distributed Systems*, 2012. 24(1): 170–183.

84. Kanduri, A., Haghbayan, M.-H., Rahmani, A. M., Shafique, M., Jantsch, A. and Liljeberg, P. adBoost: Thermal aware performance boosting through dark silicon patterning. *IEEE Transactions on Computers*, 2018. (8): 1062–1077.

85. Hanumaiah, V., Vrudhula, S. and Chatha, K. S. Performance optimal online DVFS and task migration techniques for thermally constrained multi-core processors. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2011. 30(11): 1677–1690.

86. Fang, Z., Hallnor, E., Li, B., Leddige, M., Dai, D., Lee, S. E., Makineni, S. and Iyer, R. Boomerang: reducing power consumption of response packets in NoCs with minimal performance impact. *IEEE Computer Architecture Letters*, 2010. 9(2): 49–52.

87. Matsutani, H., Koibuchi, M., Ikebuchi, D., Usami, K., Nakamura, H. and Amano, H. Ultra fine-grained run-time power gating of on-chip routers for cmps. *Proceedings of 4th ACM/IEEE International Symposium on Networks-on-Chip*. Grenoble, France. 2010. 61–68.

88. Matsutani, H., Koibuchi, M., Ikebuchi, D., Usami, K., Nakamura, H. and Amano, H. Performance, area, and power evaluations of ultrafine-grained run-time power-gating routers for CMPs. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2011. 30(4): 520–533.

89. Chen, L., Zhu, D., Pedram, M. and Pinkston, T. M. Power punch: Towards non-blocking power-gating of noc routers. *Proceedings of the 21st International Symposium on High Performance Computer Architecture (HPCA)*. Burlingame, CA, USA. 2015. 378–389.

90. Neelkamal, Yadav, S. and Kapoor, H. K. Lightweight Message Encoding of Power-Gating Controller for On-Time Wakeup of Gated Router in Network-on-Chip. *Proceedings of the 9th International Symposium on Embedded Computing and System Design (ISED)*. Kollam, India. 2019. 1–6.

91. Ikebuchi, D., Seki, N., Kojima, Y., Kamata, M., Zhao, L., Amano, H., Shirai, T., Koyama, S., Hashida, T., Umahashi, Y., Masuda, H., Usami, K., Takeda, S., Nakamura, H., Namiki, M. and Kondo, M. Geyser-1: A MIPS R3000 CPU core with fine grain runtime power gating. *Proceedings of the 2009 IEEE Asian Solid-State Circuits Conference*. Taipei, Taiwan. 2009. 281–284.

92. Roy, S., Ranganathan, N. and Katkoori, S. A framework for power-gating functional units in embedded microprocessors. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 2009. 17(11): 1640–1649.

93. Yeh, C.-C., Chang, K.-C., Chen, T.-F. and Yeh, C. Maintaining performance on power gating of microprocessor functional units by using a predictive pre-wakeup strategy. *ACM Transactions on Architecture and Code Optimization (TACO)*, 2011. 8(3): 1–27.

94. Hoerl, A. E. and Kennard, R. W. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 1970. 12(1): 55–67.

95. Farrar, D. E. and Glauber, R. R. Multicollinearity in regression analysis: the problem revisited. *The Review of Economic and Statistics*, 1967: 92–107.

96. Tibshirani, R. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 1996. 58(1): 267–288.

97. Rohith, R., Rathore, V., Chaturvedi, V., Singh, A. K., Thambipillai, S. and Lam, S.-K. LifeSim: A lifetime reliability simulator for manycore systems. *Proceedings of the 8th Annual Computing and Communication Workshop and Conference (CCWC)*. Las Vegas, NV, USA. 2018. 375–381.

98. Genbrugge, D., Eyerman, S. and Eeckhout, L. Interval simulation: Raising the level of abstraction in architectural simulation. *Proceedings of the 16th International Symposium on High-Performance Computer Architecture (HPCA)*. Bangalore, India. 2010. 1–12.

99. Skadron, K., Stan, M. R., Huang, W., Velusamy, S., Sankaranarayanan, K. and Tarjan, D. Temperature-aware computer systems: Opportunities and challenges. *IEEE Micro*, 2003. 23(6): 52–61.

100. Dubey, P. Recognition, mining and synthesis moves computers to the era of tera. *Technology@ Intel Magazine*, 2005. 9(2): 1–10.

101. Bienia, C., Kumar, S. and Li, K. PARSEC vs. SPLASH-2: A quantitative comparison of two multithreaded benchmark suites on chip-multiprocessors. *Proceedings of the 2008 IEEE International Symposium on Workload Characterization*. Seattle, WA, USA. 2008. 47–56.

102. Silberschatz, A., Galvin, P. B. and Gagne, G. *Operating System Concepts*. 9th ed. Wiley Publishing. 2012. ISBN 1118063333.

103. Rothberg, E. and Gupta, A. An efficient block-oriented approach to parallel sparse Cholesky factorization. *SIAM Journal on Scientific Computing*, 1994. 15(6): 1413–1439.

104. Bailey, D. H. FFTs in external or hierarchical memory. *The Journal of Supercomputing*, 1990. 4(1): 23–35.

105. Singh, J. P. and Hennessy, J. L. Finding and exploiting parallelism in an ocean simulation program: Experience, results, and implications. *Journal of Parallel and Distributed Computing*, 1992. 15(1): 27–48.

106. Blelloch, G. E., Leiserson, C. E., Maggs, B. M., Plaxton, C. G., Smith, S. J. and Zagha, M. A comparison of sorting algorithms for the connection machine CM-2. *Proceedings of the 3rd annual ACM symposium on Parallel algorithms and architectures*. Hilton Head, South Carolina, USA. 1991. 3–16.

107. Singh, J. P., Gupta, A. and Levoy, M. Parallel visualization algorithms: Performance and architectural implications. *Computer*, 1994. 27(7): 45–55.

108. Black, F. and Scholes, M. The pricing of options and corporate liabilities. *Journal of Political Economy*, 1974. 81(3): 637–654.

109. Deutscher, J. and Reid, I. Articulated body motion capture by stochastic search. *International Journal of Computer Vision*, 2005. 61(2): 185–205.

110. Banerjee, P. *Parallel algorithms for VLSI Computer-Aided Design*. Prentice-Hall, Inc. 1994.

111. Quinlan, S. and Dorward, S. Venti: A new approach to archival storage. *Proceedings of the 2002 File and Storage Technologies (FAST)*. Monterey, California, USA. 2002. 89–102.

112. Müller, M., Charypar, D. and Gross, M. H. Particle-based fluid simulation for interactive applications. *Symposium on Computer Animation*. San Diego, California, USA. 2003. 154–159.

113. Heath, D., Jarrow, R. and Morton, A. Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica: Journal of the Econometric Society*, 1992: 77–105.

114. Grochowski, E. and Annavaram, M. Energy per instruction trends in Intel microprocessors. *Technology@ Intel Magazine*, 2006. 4(3): 1–8.

115. Chittamuru, S. V. R., Thakkar, I. G. and Pasricha, S. LIBRA: Thermal and process variation aware reliability management in photonic networks-on-chip. *IEEE Transactions on Multi-Scale Computing Systems*, 2018. 4(4): 758–772.

116. Schöne, R., Molka, D. and Werner, M. Wake-up latencies for processor idle states on current x86 processors. *Computer Science-Research and Development*, 2015. 30(2): 219–227.

117. Yoon, C., Shim, J. H., Moon, B. and Kong, J. 3D die-stacked DRAM thermal management via task allocation and core pipeline control. *IEICE Electronics Express*, 2018. 15(3): 1–12.

118. Hastie, T. and Tibshirani, R. Efficient quadratic regularization for expression arrays. *Biostatistics*, 2004. 5(3): 329–340.

119. Intel. Linux's Intel® Driver. URL `https://github.com/intel/intel-vaapi-driver/blob/master/src/intel_driver.c`, accessed October 16, 2021.

120. Mohammed, M. S., Al-Dhamari, A. K., Rahman, A. A.-H. A., Paraman, N., Al-Kubati, A. A. M. and Marsono, M. N. Temperature-Aware Task Scheduling for Dark Silicon Many-Core System-on-Chip. *Proceedings of the 8th International Conference on Modeling Simulation and Applied Optimization (ICMSAO)*. Manama, Bahrain. 2019. 1–5.

121. Mohammed, M. S., Al-Kubati, A. A., Paraman, N., Ab Rahman, A. A.-H. and Marsono, M. N. DTaPO: Dynamic Thermal-Aware Performance Optimization for Dark Silicon Many-Core Systems. *Electronics*, 2020. 9(11): 1980.

122. Mohammed, M. S., Paraman, N., Ab Rahman, A. A.-H., Ghaleb, F. A., Al-Dhamari, A. and Marsono, M. N. PEW: Prediction-Based Early Dark Cores Wake-up Using Online Ridge Regression for Many-Core Systems. *IEEE Access*, 2021. 9: 124087–124099.

LIST OF PUBLICATIONS

Journal with Impact Factor

1. Mohammed, M. S., Paraman, N., Ab Rahman, A. A.-H., Ghaleb, F. A., Al-Dhamari, A. and Marsono, M. N. PEW: Prediction-Based Early Dark Cores Wake-up Using Online Ridge Regression for Many-Core Systems. *IEEE Access*, 2021. 9: 124087–124099.
2. Mohammed, M. S., Al-Kubati, A. A., Paraman, N., Ab Rahman, A. A.-H. and Marsono, M. N. DTaPO: Dynamic Thermal-Aware Performance Optimization for Dark Silicon Many-Core Systems. *Electronics*, 2020. 9(11): 1980. 122.

Indexed Conference Proceedings

1. Mohammed, M. S., Al-Dhamari, A. K., Ab Rahman, A. A. H., Paraman, N., Al-Kubati, A. A., and Marsono, M. N. Temperature-Aware Task Scheduling for Dark Silicon Many-Core System-on-Chip. *In Proceedings of the 8th International Conference on Modeling Simulation and Applied Optimization* (ICMSAO). Manama, Bahrain: IEEE. 2019. 1–5.
2. Mohammed, M. S., Tang, J. W., Ab Rahman, A. A. H., Paraman, N., and Marsono, M. N. Rapid Prototyping of NoC-based MPSoC based on Dataflow Modeling of Real-World Applications. *In Proceedings of the 9th IEEE Control and System Graduate Research Colloquium* (ICSGRC). Shah Alam, Malaysia: IEEE. 2019. 217-222.