Dynamic Cache Partitioning and Adaptive Cache Replacement Schemes for Chip Multiprocessors

Norfadila Mahrom

B.Eng.(Computer Engineering)
M.Sc.(Intelligent System)

A THESIS SUBMITTED IN FULFILMENT OF THE
REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

SCHOOL OF ELECTRICAL AND ELECTRONIC ENGINEERING
THE UNIVERSITY OF ADELAIDE
AUSTRALIA

2014

CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ALGORITHMS

ABSTRACT

One of the dominant approaches towards implementing fast and high performance computer architectures is the Chip Multi Processor (CMP), in which the design of the memory hierarchy has a critical effect on performance. Performance can be improved by the use of a shared cache on the chip, but it is a matter of ongoing research as to how each processor can gain the greatest advantage from the cache without affecting the performance of other processors. Moreover, power is a critical issue in CMP design.

Cache replacement policies and cache partitioning schemes have been investigated and proven able to enhance shared cache management. However, it is still desirable to have an optimal replacement policy that can retain useful data as long as possible to minimise miss rate and not degrade performance in a partitioned shared cache. Many of the metrics that have led to innovations in various partitioning schemes have increased the complexity of the partitioning strategies and the hardware overhead. There is scope for more work in achieving the right balance between power consumption and performance improvement in the CMP.

This thesis investigates the effects of the cache replacement policy in a partitioned shared cache. The goal is to quantify whether a better power/performance trade-off can be achieved by using less complex replacement strategies. A Middle Insertion 2 Positions Promotion (MI2PP) policy is proposed to eliminate cache misses that could adversely affect the access patterns and the throughput of the processors in the system. The insertion, promotion and eviction strategies of the replacement policy are investigated and modified to improve shared cache utilisation by the competing processors. The MI2PP policy employs a static predefined insertion point, near distance promotion and the concept of ownership in the eviction policy to avoid resource stealing among the processors. With these strategies, the performance of the shared cache and the overall system were enhanced and the miss rate showed significant improvement over the Least Recently Used (LRU) policy.

While existing cache partitioning schemes use a variety of performance metrics to allocate the cache for each competing processor, most of the schemes focus only on one metric in their partitioning algorithm. An Adaptive Cycles per Instruction (CPI)-based Cache Partitioning (ACCP) scheme is introduced to investigate the efficiency of using two metrics to optimise partitioning decisions and to study the trade-offs between the complexity of using more performance metrics in partition decision-making and additional hardware cost. The analysis performed on ACCP showed that the performance of the processors was improved compared to the existing CPI-based partitioning scheme introduced by Muralidhara et al. [2010], which uses only one of the performance metrics employed in ACCP. Evaluation on a more complex scheme, namely the Utility Cache Partitioning (UCP) scheme demonstrated that the ACCP on average achieved similar performance although the ACCP is simpler to implement. The low hardware overhead incurred by ACCP showed that it is superior to UCP. ACCP demonstrates that the complexity of the partitioning mechanism and hardware cost could be reduced without degrading the overall system performance.

DECLARATION OF ORIGINALITY

This work contains no material that has been accepted for the award of any other degree or diploma in any university or other tertiary institution to Norfadila Mahrom and, to the best of my knowledge and belief, contains no material previously published written by another person, except where due reference has been made in the text.

I give consent to this copy of the thesis, when deposited in the University Library, being available for loan, photocopying, and dissemination through the library digital thesis collection, subject to the provisions of the Copyright Act 1968.

I also give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library catalogue, the Australasian Digital Thesis Program (ADTP) and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

Signed	Date

ACKNOWLEDGEMENTS

In the Name of Allah, the Most Gracious, the Most Merciful.

All praise and thanks to Allah.

My very great appreciation and deep gratitude to my principal supervisor, Associate Professor Michael J. Liebelt for his patient guidance, assistance, encouragement and useful criticism of my research work. His willingness to allocate his time so generously to support me in many aspects and his trust in me has been very much appreciated. I also would like to express my sincere thanks to Dr. Braden J. Phillips for his valuable and constructive suggestions during the planning and development of this research work.

The generosity of the School of Electrical and Electronic Engineering, The University of Adelaide in funding extensive access to the Tensilica tools under Tensilica Inc. University Program is highly appreciated. This work would not have been possible without the support of David Bowler in all issues related to technical resources. My research candidature was sponsored by the Ministry of Higher Education, Malaysia and Universiti Malaysia Perlis, and supported by Prof. Dr. R. Badlishah Ahmad, Prof. Dr. Ali Yeon Md Shakaff and Prof. Dato' Dr. Zul Azhar Zahid Jamal.

My grateful thanks are extended for the members of CHiPTec, particularly for everyone in the Journal Club for the valuable meetings that had introduced me to the best platform in expending my knowledge and research work. Not to forget, million thanks to Noorfazila

Kamal, Dr. Puteh Saad, Deborah Coleman-George, Ivana Rebellato and Chris Andrew Haerberli for their motivation and continuous moral support.

Finally, to my family. Thank you all for being you. You know how much you have done to make me come this far.

Norfadila Mahrom

LIST OF ABBREVIATIONS

ACCP	Adaptive CPI-based Cache Partitioning
ACCP_S	ACCP with Sampled sets
API	Application Programming Interface
ART 2	Adaptive Resonance Theory 2
BIP	Bimodal Insertion Policy
CLMP	Clustered Multi Processor
CMP	Chip Multi Processor
CPI	Cycles per Instruction
CPU	Central Processing Unit
DIP	Dynamic Insertion Policy
DPP	Dynamic Promotion Policy
DSS	Dynamic Set Sampling
EQ	Equally Partitioned
FIFO	First-In First-Out
FPGA	Field-Programmable Gate Array
HU	High Utility
I/O	Input/Output
ID	Identification
IPC	Instructions per Cycle
ISA	Instruction Set Architecture

ISS	Instruction Set Simulator
ISSCC	IEEE international Solid-State Circuits Conference
L1	Level one
L2	Level two
L3	Level three
LFU	Least Frequently Used
LIP	LRU Insertion Policy
LLC	Last Level Cache
LRU	Least Recently Used
LU	Low Utility
MI2PP	Middle Insertion 2 Positions Promotion
MMU	Memory Management Unit
MPKI	Misses per 1000 Instructions
MPP	MRU Promotion Policy
MRU	Most Recently Used
MU	Marginal Utility
PIF	Processor Interface
PIPP	Promotion/Insertion Pseudo-Partitioning
PRP	Partition-based Replacement Policy
PSEL	Policy Selector
RAM	Random Access Memory
RISC	Reduced Instruction Set Computer
ROM	Read Only Memory
SIPP	Single-step Incremental Promotion Policy
SMP	Symmetric Multi Processor

SPEC	Standard Performance Evaluation Corporation
SU	Saturating Utility
TADIP	Thread-Aware Dynamic Insertion Policy
TADPP	Thread-Aware Dynamic Promotion Policy
TIE	Tensilica Instruction Extension
UCP	Utility-Cache Partitioning
UMON	Utility Monitor
UMON_DSS	Utility Monitor with Dynamic Set Sampling
VPR	Versatile Place and Route
XCC	Xtensa C/C++ Compiler
XPG	Xtensa Processor Generator
XTMP	Xtensa Modeling Protocol