MITIGATING CACHE ASSOCIATIVITY AND COHERENCE SCALABILITY CONSTRAINTS FOR MANY-CORE CHIP MULTIPROCESSORS

Thesis By

Malik Al-Manasia

In Partial Fulfilment of the Requirements for the Degree of

Doctor of Philosophy in Computer Engineering

School of Computing and Communications

Faculty of Engineering and Information Technology

University of Technology Sydney

CERTIFICATE OF ORIGINAL AUTHORSHIP

I certify that the work in this thesis has not previously been submitted for a degree nor has it been submitted as part of requirements for a degree except as part of the collaborative doctoral degree and/or fully acknowledged within the text.

I also certify that the thesis has been written by me. Any help that I have received in my research work and the preparation of the thesis itself has been acknowledged. In addition, I certify that all information sources and literature used are indicated in the thesis.

Signature of Student:

Date:

Acknowledgments

O Allah, being able to thank You is the greatest blessing You have blessed me with, You are the one who gave me the capability to accomplish this significant task of my life.

I would like to express my deep and sincere gratitude to my supervisor, Dr. Zenon Chaczko, for his persistent guidance, encouragement, and support throughout my period of study at the University of Technology, Sydney. His advice was critical in helping me lay the right foundation for my research, and thereby develop the right perspective in moving forward.

There are no words that could express my true gratitude to my beloved parents, wife, relatives and friends for their constant motivation, endless support and understanding. This thesis would have been impossible without every one of them. I married my wife half-way into my Ph.D. canditature, for which I appreciate her programming assistance, emotional support and unfaltering understanding. I extend a special gratitude towards my parents, the two people who raised me, who offered their full support, loved me and made the person I am today, for which I express my deepest gratitude.

Abstract

Chip Multi-Processor (CMP) designs have become dominant in the processor market. The evaluation and development of CMPs is essential for product improvement. Up to date, CMPs have presented many challenges for system designers, including cache memory system scalability. My research aims to implement a highly scalable CMP cache memory system using an associative cache, with enhanced replacement policy and a scalable cache coherent protocol.

This thesis establishes a novel Adaptive Hashing and Replacement Cache (AHRC) design, which can maintain high associativity with an advanced method of replacement policy. The AHRC design can improve associativity and keep the possible number of locations of each block (or ways) to a minimum. For the AHRC, the Adaptive Reuse Interval Prediction (ARIP) replacement policy was used because of its ability to resist both scan and thrash.

This research involved simulating several workloads on a large-scale CMP with AHRC as the last-level cache. The results demonstrated that AHRC has better energy efficiency and higher performance than conventional caches. Additionally, larger caches that utilise AHRC are the most suitable in many-core CMPs, as they support scalability as opposed to smaller caches. Scalable cache coherence protocols are essential for CMPs systems, in order to satisfy the requirement for more dominant high-performance chips with shared memory. However, the limited size of the directory cache, associated with larger systems, may result in recurrent directory entries, evictions and invalidations of cached blocks thus compromising system performance.

This thesis proposes the Private/Shared, Read-Only/Read-Write, Invalid/Valid scalable coherence protocol called PROI. This novel protocol implements a slight modification on the caches' tags, allowing it to differentiate between the private and shared data on a block granularity level. Also, PROI employs a dynamic writing policy with self-invalidation and self-downgrade for each L1 cache and can sustain system coherence and performance, scale with the raised number of cores and reduce area, energy, and performance associated costs with the coherence mechanism. The result indicates that PROI can reduce various variables, including the miss ratio of the private L1 cache by 17%, the network traffic, application runtime of approximately 6%, and energy consumption by about 35%. Therefore, utilising AHRC, ARIP, and PROI can mitigate the cache scalability constraints significantly and maintain the performance level while enhancing energy consumption of the CMP cache.

Table of Contents

List of Figures

List of Tables

List of Abbreviations

AHRC	Adaptive Hashing and Replacement Cache
ARIP	Adaptive Reuse Interval Prediction
BIP	Bimodal Insertion Policy
BSD	Berkeley-Style Open Source
CC-NUMA	Cache Coherent Non-Uniform Memory Architecture
CMP	Chip Multi-Processors
CPI	Cycles Per Instruction
CPU	Central Processing Unit
CSA	Computer System Architecture
DIP	Dynamic Insertion Policy
DMA	Direct Memory Access
DPIIP	Dynamic Promotion with Interpolated Increments Policy
DRAM	Dynamic Random Access Memory
DRF	Data-Race-Free
DSM	Distributed Shared Memory
DSP	Digital Signal Processor
DVFS	Dynamic Voltage and Frequency Scaling
ECC	Error Correction Codes
EEN	Explicit Eviction Notification
FIFO	First in First Out
FS/A	Full-System / Application-Level

ICEmon	In-Cache Estimation Monitor
IIC	Indirect Index Cache
IPC	Instruction Per Cycle
IPSEL	Insertion Policy Selection
ISA	Instruction Set Architectures
KIPS	Kilo Instructions Per Second
L1	Cache Level 1
L2	Cache Level 2
L3	Cache Level 3
LIFO	Last in First Out
LIP	LRU Insertion Policy
LLC	Last Level Cache
LoD	Level of Details
LRU	Least Recently Used
MLP	Memory Level Parallelism
MPKI	Misses Per 1000 Instructions
MRU	Most Recently Used
NACK	Negative Acknowledgment
NRU	Not Recently Used
NUCA	Non-Uniform Cache Access
NUMA	Non-Uniform Memory Architecture
OPT	Optimal Replacement Policy
OS	Operating System
PC	Personal Computer
PIPP	Promotion/Insertion Pseudo-Partitioning
PPSEL	Promotion Policy Selection

PROI	Private/Shared, Read-Only/Read-Write, Invalid/Valid
PSEL	Policy Selection
PTE	Page Table Entry
QoS	Quality-of-Service
RO	Read Only
ROI	Region of Interest
RRIP	Re-Reference Interval Prediction
RRPV	Re-Reference Prediction Value
RSWEL	Reconstituted SWEL
RW	Read Write
SDM	Set Dueling Monitor
SIPP	Single-step Incremental Promotion Policy
SMP	Symmetric Multi-Processor
SMT	Symmetric Multi-threaded
SRAM	Static Random Access Memory
SWEL	Protocol states are Shared, Written, Exclusivity. Level
SWMR	Single-Writer/Multiple-Readers
TADIP	Thread-Aware Dynamic Insertion Policy
TLB	Table Lookaside Buffer
TMA	Trap-based Memory Architecture
UCP	Utility-based Cache Partitioning
UIUC/NCSA	University of Illinois/NCSA Open Source License
VIPS	Valid/Invalid Private/Shared
W	watt