

# Task Scheduling for Multi-core and Parallel Architectures

Quan Chen · Minyi Guo

# Task Scheduling for Multi-core and Parallel Architectures

Challenges, Solutions and Perspectives

Quan Chen  
Shanghai Jiao Tong University  
Shanghai  
China

Minyi Guo  
Shanghai Jiao Tong University  
Shanghai  
China

ISBN 978-981-10-6237-7      ISBN 978-981-10-6238-4 (eBook)  
<https://doi.org/10.1007/978-981-10-6238-4>

Library of Congress Control Number: 2017956334

© Springer Nature Singapore Pte Ltd. 2017

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Printed on acid-free paper

This Springer imprint is published by Springer Nature

The registered company is Springer Nature Singapore Pte Ltd.

The registered company address is: 152 Beach Road, #21-01/04 Gateway East, Singapore 189721, Singapore

# Preface

Computers play an indispensable role in almost all scientific fields and disciplines such as biomedicine, physical simulations, computational chemistry, and astronautics. In order to fulfill the urgent requirement for high computational capacity (for example, huge amount of computational capacity is required in big data analysis), technologies of parallel computing and resource management increasingly grow. Nowadays, multi-core processors have become mainstream in both research and real-world settings, from warehouse-scale datacenter, personal desktops, laptops, to smartphones, since they demonstrate the superior performance per watt and the larger computational capacity compared to the traditional single-core processors. For these widely used parallel systems, a key problem is how to schedule the tasks to efficiently utilize the hardware, improve the performance, and guarantee the Quality-of-Service. Because different types of architectures have totally different features, there is not universal scheduling technique that can work the best across all these architectures and different applications. Scheduling techniques often need to be altered to match the various architectures (e.g., multi-core, datacenter, and distributed system) accordingly. For example, for a big data processing system that runs on large-scale distributed system, task scheduling should focus on load balancing and data locality; for a datacenter, the scheduling should aim to guarantee the quality-of-service of the customer-facing applications (e.g., web search).

However, a book that elaborates task scheduling techniques for emerging complex parallel architectures (e.g., multi-socket architecture, heterogeneous multi-core architecture, and cloud/big data processing platform) is missing. Previous published works mainly discuss traditional task scheduling techniques for generalized parallel systems mathematically. But these techniques suffer from low performance on the emerging complex parallel architectures. To this end, this book discusses the state-of-the-art task scheduling techniques that are optimized against different architectures, and these techniques can be applied in real parallel systems directly.

In this book, we will mainly introduce these task scheduling techniques in the scenarios of emerging parallel architectures, including the multi-core architecture,

cloud platform, and accelerator. The book will examine the current challenges in this topic, and present detailed solutions including algorithms, methods, and perspectives. It is well noticed that parallel architectures are becoming more and more complex. Instances are multi-socket multi-core architecture, asymmetric multi-core architecture, various parallel accelerators, distributed parallel architecture, etc. For different types of parallel architectures, in order to utilize the hardware efficiently and maximize the performance, different techniques have been introduced and integrated into the traditional task scheduling policy. To this end, we elaborate these techniques and demonstrate that they can be implemented efficiently for emerging parallel architectures. This book consists of three main parts: *background*, *task scheduling techniques*, and *perspectives*.

## Part I: Background

In this part, we mainly introduce the background of this book, including the emerging widely used parallel architectures and classic task scheduling techniques. This part includes two chapters.

In the first chapter, we do a survey on the emerging parallel architectures including multi-core architectures, NUMA-enabled multi-core architectures, asymmetric multi-core architectures, accelerators, cloud platform, and so on. Besides these parallel architectures, vendors are now producing other architectures. For instance, Google releases Tensor Processing Unit (TPU), which is a parallel architecture as well recently.

In the second chapter, we introduce the classic task scheduling policies and parallel programming environments. Work-sharing and work-stealing are the two most classic task scheduling policies. In addition, we introduce many parallel programming environments, such as Apache Hadoop, Spark, MIT Cilk, TBB, X10, and so on. These task scheduling policies can be applied in various parallel architectures but may suffer from low performance. In the next part, we introduce the techniques that optimize the parallel applications on various parallel architectures.

## Part II: Task Scheduling for Various Parallel Architectures

In this part, we introduce techniques that can be used to improve the performance of applications on the emerging widely used parallel architectures.

In the third chapter, targeting the multi-socket architecture, we will introduce cache-aware task scheduling policy, which can improve shared cache utilization in different sockets.

In the fourth chapter, on the NUMA (Non-Uniform Memory Access)-enabled architecture, the NUMA-aware task scheduling policy will be introduced, which

can reduce remote memory accesses and improve the performance of applications in consequence.

In the fifth chapter, we introduce workload-aware task scheduling policy, which is proposed for asymmetric multi-core architecture. On this kind of architecture, where different cores operate at different speeds, partitioning by the workload can truly improve the performance.

In the sixth chapter, we introduce asymptotic technique to allocate workload between CPU and GPU for CPU+GPU heterogeneous parallel architecture. On this kind of architecture, different applications have different speedup ratios on the GPU compared with CPU, because the applications have various characteristics. It is not trivial to find an optimal workload partition between CPU and GPU.

Nowadays, big data analysis requires tremendous amount of computers to process the data in parallel. In the seventh chapter, we will introduce several featured dynamic task scheduling policies that can significantly improve the performance of big data processing on heterogeneous cloud platforms.

The eighth chapter contains the quality-of-service aware task scheduling policy for accelerators. Using this policy, it can improve the accelerator utilization. Moreover, it also guarantees the quality-of-service of latency-critical applications.

### **Part III: Summary and Perspectives**

In this part, we summarize all the previously introduced task scheduling solutions for parallel architectures, provide our perspectives, and discuss the possibilities of designing new dynamic task scheduling policies for more other future parallel architectures. Especially, we give several guidelines of designing new efficient and effective task scheduling techniques for those newly released parallel architectures.

After reading this book, we expect the readers will have an overview on the recent progress of task scheduling policies in parallel architectures. And we also hope the book can help the readers to quickly master the focused issues and opening problems if they tend to work in this field. In order to understand this book, the readers are suggested to have some basic knowledge on computer architecture, multi-core, and parallel processing.

Shanghai, China  
September 2017

Quan Chen  
Minyi Guo

# Acknowledgements

This book was partially sponsored by the National Basic Research 973 Program of China under grant 2015CB352403, the National Natural Science Foundation of China (NSFC) (61602301). We are grateful for the editor of this book, Dr. XiaoLan Yao at Springer for her patience and support to make this book possible.

Shanghai, China  
September 2017

Quan Chen  
Minyi Guo

# Contents

## Part I Background

<b>1 Emerging Parallel Architectures</b> . . . . .	3
1.1 Parallel Architecture is Dominating the World . . . . .	3
1.2 Shared Memory Parallel Architecture . . . . .	4
1.2.1 Multi-core Architecture . . . . .	4
1.2.2 Multi-socket Multi-core Architecture . . . . .	5
1.2.3 Asymmetric Multi-core Architecture . . . . .	6
1.3 Distributed Memory Parallel Architecture . . . . .	6
1.3.1 Tight-Coupled Distributed Memory Architecture . . . . .	7
1.3.2 Loose-Coupled Distributed Memory Architecture . . . . .	7
1.4 Accelerator . . . . .	8
1.4.1 GPGPU . . . . .	9
1.4.2 Intel Xeon Phi . . . . .	10
1.5 Heterogeneous Parallel Architecture . . . . .	11
1.6 Chapter Highlights . . . . .	11
References . . . . .	12
<b>2 Conventional Task Scheduling Policies</b> . . . . .	13
2.1 Manual Task Scheduling Policies . . . . .	14
2.1.1 Message Passing . . . . .	14
2.1.2 Multi-threading . . . . .	15
2.2 Automatic Task Scheduling Policies . . . . .	15
2.2.1 Task Scheduling Policies for Data Parallelism . . . . .	15
2.2.2 Task Scheduling Policies for Task Parallelism . . . . .	16
2.3 Parallel Programming Environments . . . . .	19
2.3.1 Programming Environments for Data Parallelism . . . . .	19
2.3.2 Programming Environments for Task Parallelism . . . . .	22



2.4	Problems in Existing Task Scheduling Systems . . . . .	24
2.5	Chapter Highlights . . . . .	24
	References . . . . .	25

## **Part II Optimized Task Scheduling for Parallel Architectures**

<b>3</b>	<b>Work-Stealing for Multi-socket Architecture . . . . .</b>	<b>29</b>
3.1	Background and Existing Problems . . . . .	29
3.1.1	The TRICI Problem . . . . .	30
3.2	Prior Solutions . . . . .	32
3.2.1	Scalable Locality-Aware Adaptive Work-Stealing (SLAW) . . . . .	32
3.2.2	Multi-Threaded Shepherds (MTS) . . . . .	33
3.2.3	Probability Work-Stealing (PWS) . . . . .	35
3.2.4	Hierarchical Work-Stealing (HWS) . . . . .	35
3.2.5	CONTROLLED-PDF . . . . .	36
3.3	Cache-Aware Bi-tier Work-Stealing . . . . .	37
3.3.1	Solution Overview . . . . .	37
3.3.2	Design Overview . . . . .	38
3.4	Cache-Aware Task Graph Partition Policy . . . . .	40
3.4.1	Full Tree Oriented Partition Policy . . . . .	40
3.4.2	General Tree Oriented Partition Policy . . . . .	43
3.5	Bi-tier Work-Stealing Scheduling Policy . . . . .	47
3.5.1	Work Stealing Algorithm . . . . .	48
3.5.2	Task Generating Algorithm . . . . .	49
3.6	Theoretical Time and Space Bounds . . . . .	51
3.6.1	Theoretical Bounds for Random Work-Stealing . . . . .	51
3.6.2	Theoretical Bounds for CAB . . . . .	52
3.7	Implementation Methodology . . . . .	54
3.7.1	Compiler Support . . . . .	54
3.7.2	Runtime Support . . . . .	55
3.8	Evaluation of CAB . . . . .	55
3.8.1	Performance of CAB-FTO . . . . .	57
3.8.2	Performance of CAB-GTO . . . . .	61
3.9	Summary . . . . .	69
3.9.1	Chapter Highlights . . . . .	70
	References . . . . .	71
<b>4</b>	<b>Work-Stealing for NUMA-enabled Architecture . . . . .</b>	<b>73</b>
4.1	Chapter Organization . . . . .	73
4.2	Background and Existing Problems . . . . .	73
4.3	Prior Solutions . . . . .	75
4.3.1	Random Pushing . . . . .	76
4.3.2	Cluster-Aware Hierarchical Stealing (CHS) . . . . .	77

4.3.3	Cluster-Aware Load-Based Stealing (CLS) . . . . .	77
4.3.4	Cluster-Aware Random Stealing (CRS) . . . . .	79
4.3.5	TATL . . . . .	80
4.3.6	NUMALB . . . . .	82
4.3.7	Offline Technique for Unstructured Parallelism . . . . .	84
4.4	Design of Locality-Aware Work-Stealing . . . . .	87
4.5	Load-Balanced Task Allocator . . . . .	88
4.6	Cache-Friendly Task Graph Partitioner . . . . .	91
4.6.1	Decide the Initial Partitioning . . . . .	91
4.6.2	Search for the Optimal Partitioning . . . . .	92
4.7	Triple-Level Work-Stealing Policy . . . . .	94
4.8	Theoretical Validation . . . . .	96
4.9	Implementation Methodology . . . . .	97
4.10	Performance Evaluation of LAWS . . . . .	98
4.10.1	Experimental Platforms . . . . .	98
4.10.2	Performance of LAWS . . . . .	100
4.10.3	Effectiveness of Cache-Friendly Task Graph Partitioner . . . . .	103
4.10.4	Scalability of LAWS . . . . .	104
4.10.5	Overhead of LAWS . . . . .	107
4.10.6	Applicability of LAWS . . . . .	108
4.11	Summary . . . . .	109
4.11.1	Chapter Highlights . . . . .	109
	References . . . . .	110
<b>5</b>	<b>Dynamic Load Balancing for Asymmetric Multi-core</b>	
	<b>Architecture . . . . .</b>	<b>113</b>
5.1	Chapter Organization . . . . .	113
5.2	Problem Formulation . . . . .	114
5.3	Existing Solutions . . . . .	115
5.3.1	Task Snatching Technique . . . . .	115
5.3.2	CAMP . . . . .	117
5.3.3	Bias Scheduling . . . . .	119
5.3.4	Age-Based Scheduling . . . . .	121
5.3.5	Speed-Based Balancing . . . . .	123
5.3.6	Scheduling on AMC with Hardware Support . . . . .	126
5.4	Theoretical Ideal Task Scheduling . . . . .	126
5.5	A Practical Polynomial Time Solution . . . . .	127
5.6	Design of Asymmetric-Aware Task Scheduling . . . . .	129
5.6.1	Processing Flow of AATS . . . . .	130
5.7	History-Based Task Allocation . . . . .	131
5.7.1	Build Task Classes . . . . .	132
5.7.2	Allocate Task Classes to C-Groups . . . . .	134

5.8	Preference-Based Work-Stealing . . . . .	136
5.8.1	Scheduling Within a C-Group . . . . .	136
5.8.2	Scheduling Among C-Groups . . . . .	137
5.9	Implementation Methodology of AATS . . . . .	139
5.10	Performance of AATS . . . . .	140
5.10.1	Experimental Configurations . . . . .	140
5.10.2	Performance on Emulated Platform . . . . .	143
5.10.3	Effectiveness of the Preference-Based Work-Stealing . . . . .	145
5.10.4	Scalability of AATS . . . . .	146
5.10.5	Integrating Task-Snatching in AATS . . . . .	148
5.11	Summary . . . . .	149
5.11.1	Chapter Highlights . . . . .	150
	References . . . . .	150
<b>6</b>	<b>Load Balancing for Heterogeneous Parallel Architecture . . . . .</b>	<b>153</b>
6.1	Background and Existing Problems . . . . .	153
6.2	Prior Solutions . . . . .	155
6.2.1	Static Scheduling . . . . .	155
6.2.2	Quick Scheduling . . . . .	156
6.2.3	Split Scheduling . . . . .	158
6.2.4	FinePar . . . . .	159
6.3	Heterogeneous-Aware Task Scheduling . . . . .	161
6.4	Comparison of the Scheduling Policies . . . . .	162
6.5	Performance of Dynamic Scheduling Policies . . . . .	164
6.5.1	Experimental Setup . . . . .	164
6.5.2	Performance . . . . .	165
6.5.3	Effectiveness of Balancing Workload . . . . .	167
6.5.4	Effectiveness of Predicting the Performance of GPU . . . . .	168
6.5.5	Impact of Profiling Granularity . . . . .	168
6.6	Summary . . . . .	169
6.6.1	Chapter Highlights . . . . .	169
	References . . . . .	170
<b>7</b>	<b>MapReduce for Cloud Computing . . . . .</b>	<b>173</b>
7.1	Introduction to MapReduce . . . . .	173
7.1.1	Scheduling Policy in MapReduce . . . . .	174
7.1.2	Adapting to Other Platforms . . . . .	175
7.1.3	Variations of MapReduce . . . . .	176
7.1.4	Existing Problem in Heterogeneous Environment . . . . .	176
7.2	Prior Solutions . . . . .	177
7.2.1	Least Progress Policy . . . . .	177
7.2.2	Longest Approximate Time to End Policy . . . . .	178

7.2.3	Calculating Progress Score . . . . .	179
7.2.4	Problems in Existing Solutions . . . . .	180
7.2.5	Tarazu . . . . .	181
7.3	Self-adaptive MapReduce Scheduling . . . . .	184
7.3.1	Overview of SAMR . . . . .	184
7.3.2	Tuning Phase Weights . . . . .	185
7.3.3	Calculating Progress Score . . . . .	185
7.3.4	Identifying Straggler Task. . . . .	186
7.3.5	Identifying Slow Node . . . . .	187
7.3.6	Boosting Straggler Task . . . . .	188
7.4	Implementation of SAMR . . . . .	189
7.5	Performance Evaluation . . . . .	190
7.5.1	Experimental Setup . . . . .	190
7.5.2	Performance. . . . .	191
7.5.3	Effectiveness of Speculative Execution and Weight Tuning . . . . .	192
7.5.4	Parameter Selection in SAMR. . . . .	193
7.6	Summary . . . . .	196
7.6.1	Chapter Highlights . . . . .	196
	References . . . . .	197
<b>8</b>	<b>QoS-Aware Task Reordering for Accelerators . . . . .</b>	<b>199</b>
8.1	Background and Existing Problems . . . . .	199
8.2	Prior Work on Handling Accelerator Co-location . . . . .	200
8.2.1	TimeGraph . . . . .	201
8.2.2	GPU-EvR . . . . .	202
8.2.3	Simultaneous Multi-kernel (SMK). . . . .	204
8.2.4	GPU Thread Preemption. . . . .	206
8.3	Real System Investigation on Accelerator Co-location. . . . .	206
8.4	Investigation on Priority-Based Scheduling Policy . . . . .	208
8.5	Design of Task Scheduling Mechanism on Accelerators. . . . .	209
8.6	Case Study: QoS-Aware Task Scheduling on Accelerator . . . . .	210
8.6.1	Root Causes of Long Tail Latency at Co-location . . . . .	210
8.6.2	Design of Baymax . . . . .	211
8.7	Task Duration Modeling in Baymax . . . . .	212
8.7.1	Task Duration Predictor . . . . .	212
8.7.2	Selecting Representative Features . . . . .	213
8.7.3	Low Overhead Prediction Models . . . . .	214
8.7.4	Minimizing Prediction Error . . . . .	215
8.7.5	Prediction Accuracy . . . . .	215
8.8	Scheduling Hand-Written Kernels and Library Calls . . . . .	218
8.8.1	Breaking down the End-to-end Latency . . . . .	218
8.8.2	Scheduling Policy . . . . .	219

8.9	Scheduling Data Transfer Tasks . . . . .	222
8.9.1	Characterizing PCI-e Bandwidth Contention . . . . .	222
8.9.2	Scheduling Policy . . . . .	223
8.10	Performance of Baymax . . . . .	224
8.10.1	Experimental Configuration . . . . .	224
8.10.2	QoS and Throughput . . . . .	225
8.10.3	Scheduling Data Transfer Tasks . . . . .	226
8.10.4	Beyond Pair-Wise Co-locations . . . . .	227
8.11	Summary . . . . .	228
8.11.1	Chapter Highlights . . . . .	229
	References . . . . .	230
 <b>Part III Summary and Discussion</b>		
<b>9</b>	<b>Summary and Discussion . . . . .</b>	<b>235</b>
9.1	Guideline of Scheduling Technique Design . . . . .	235
9.2	Multi-socket Architecture . . . . .	236
9.3	NUMA-Enabled Multi-socket Architecture . . . . .	236
9.4	Asymmetric Multi-core Architecture . . . . .	237
9.5	Heterogeneous CPU+GPU Architecture . . . . .	238
9.6	Heterogeneous Cloud Platform . . . . .	238
9.7	Non-preemptive Accelerator Architecture . . . . .	239
	<b>Glossary . . . . .</b>	<b>241</b>

# Acronyms

AATS	Asymmetric-Aware Task Scheduling
AMC	Asymmetric Multi-Core
ANN	Approximate Nearest Neighbor
BIOS	Basic Input/Output System
CAB	Cache-Aware Bi-tier Work-Stealing
CF	Cache-Friendly
c-group	Core Group
CMPI	Cache Misses Per Instruction
CPU	Central Processing Unit
D&C	Divide-and-Conquer
DAG	Directed Acyclic Graph
DDR	Dual Data Rate
DRAM	Dynamic Random Access Memory
DT	Duration Table
DVFS	Dynamic Voltage and Frequency Scaling
EDC	MCDRAM Controller
FIFO	First-In-First-Out
FPGA	Field-Programmable Gate Array
FTO	Full Tree Oriented
GFS	Google File System
GPGPU	General-Purpose Graph Processing Unit
GTO	General Tree Oriented
HATS	Heterogeneous-Aware Task Scheduler
HDFS	Hadoop Distributed File System
HTT	Hyper-Threading Technology
IC	Inter-Connect
IPA	Intelligent Personal Assistant
ISA	Instruction Set Architecture
KNL	Intel Xeon Phi (codenamed Knights Landing)
KNN	K-Nearest Neighbor

LATE	Longest Approximate Time to End
LAWS	Locality-Aware Work-Stealing
LR	Linear Regression
MCDRAM	Multi-Channel Dynamic Random Access Memory
MIMD	Multiple-Instruction-Multiple-Data
MPI	Message Passing Interface
MPS	Multi-Process Service
MSMC	Multi-Socket Multi-core
NUMA	Non-Uniform Memory Access
PFWS	Parent-First Work-Stealing
PMC	Performance Monitoring Counter
QoS	Quality-of-Service
RDD	Resilient Distributed Dataset, defined in Apache Spark
RPC	Remote Procedure Call
SAMR	Self-Adaptive Map-Reduce
SIMD	Single-Instruction-Multiple-Data
SM	Streaming Multiprocessor
SMC	Symmetric Multi-core
SOID	Size Of Involved Data
SVM	Support Vector Machines
TBB	Intel Thread Building Blocks
TCO	Total Cost of Ownership
TRICI	Task Relocation Incurred Cache Interference problem