

# **Undergraduate Topics in Computer Science**

Undergraduate Topics in Computer Science (UTiCS) delivers high-quality instructional content for undergraduates studying in all areas of computing and information science. From core foundational and theoretical material to final-year topics and applications, UTiCS books take a fresh, concise, and modern approach and are ideal for self-study or for a one- or two-semester course. The texts are all authored by established experts in their fields, reviewed by an international advisory board, and contain numerous examples and problems. Many include fully worked solutions.

For further volumes:

<http://www.springer.com/series/7592>

Yosi Ben-Asher

# Multicore Programming Using the ParC Language

 Springer

Yosi Ben-Asher  
Department of Computer Science  
University of Haifa  
Haifa, Israel

*Series editor*  
Ian Mackie

*Advisory board*

Samson Abramsky, University of Oxford, Oxford, UK  
Karin Breitman, Pontifical Catholic University of Rio de Janeiro, Rio de Janeiro, Brazil  
Chris Hankin, Imperial College London, London, UK  
Dexter Kozen, Cornell University, Ithaca, USA  
Andrew Pitts, University of Cambridge, Cambridge, UK  
Hanne Riis Nielson, Technical University of Denmark, Kongens Lyngby, Denmark  
Steven Skiena, Stony Brook University, Stony Brook, USA  
Iain Stewart, University of Durham, Durham, UK

ISSN 1863-7310 Undergraduate Topics in Computer Science  
ISBN 978-1-4471-2163-3 ISBN 978-1-4471-2164-0 (eBook)  
DOI 10.1007/978-1-4471-2164-0  
Springer London Heidelberg New York Dordrecht

Library of Congress Control Number: 2012940236

© Springer-Verlag London 2012

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. Exempted from this legal reservation are brief excerpts in connection with reviews or scholarly analysis or material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work. Duplication of this publication or parts thereof is permitted only under the provisions of the Copyright Law of the Publisher's location, in its current version, and permission for use must always be obtained from Springer. Permissions for use may be obtained through RightsLink at the Copyright Clearance Center. Violations are liable to prosecution under the respective Copyright Law.

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.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Printed on acid-free paper

Springer is part of Springer Science+Business Media ([www.springer.com](http://www.springer.com))

# Preface

This book discusses principles of practical parallel programming using shared memory on multicore machines. It uses a simple yet powerful parallel dialect of C called ParC as the basic programming language. The book contains a mixture of research directions combined with elementary material and basic concepts of parallel processing. As such it is intended to serve both as a text-book for a course on shared memory programming and as a way to improve the research background of specialists working on practical implementations of parallel programming systems over multicore machines. The objective of this book is to provide a firm basis for the “delicate art” of creating efficient parallel programs. The students can exercise parallel programming using a simulation software, which is portable on PC/Unix Multicore computers. Hence no special hardware is needed in order to gain experience in parallel programming using this book. Apart from the simulator *ParC* can be directly executed on Multicore machines. In the first four chapters elementary and advance concepts of parallel programming are carefully introduced via *ParC* examples. There is a special effort to present parallel programs as partial order over a set of assignments. Partial orders thus form the basis for determining elementary concepts such as the execution time of parallel programs scheduling, atomicity and threads.

The remaining chapters cover issues in parallel operating systems and compilation techniques that are relevant for shared memory and multicore machines. This reflects our belief that knowing how a parallel system works is necessary to master parallel programming. It describes the principles of scheduling parallel programs and the way ParC’s constructs are implemented. A separate chapter presents a set of ParC programs that forms a benchmark for measuring basic constants of the underlying system, e.g., the cost of creating a thread. A set of relevant compilation techniques of ParC programs is described. A special treatment is given to stack management, showing that essentially the stacks of different threads can be mapped to the regular stack of each core.

This book targets multicore machines which are now available as regular personal computers, laptops and servers. The shared memory in multicore machines is implemented using complex algorithms maintaining cache-consistency of the dif-

ferent cores. The hidden costs involved with the use of cache-consistency can easily damage the expected benefit of using shared memory parallel programming in multicore machines. Thus we carefully study the issue of cache consistency in multicore machines and consider several programming techniques to optimize the cache usage in ParC programs.

The usefulness of ParC stems from two sources. On the one hand, it is similar in flavor to the pseudo-languages that are used by theoreticians for describing their parallel algorithms. Indeed, parallel algorithms can be directly coded in ParC. On the other hand, sequential code segments can be easily parallelized in ParC. These two directions of coding parallel algorithms and parallelizing sequential code form the basis for developing practical parallel programs and are thus widely studied in this book. The ParC language promotes an understanding of the problems hidden in parallel code, and allows various approaches for their solutions. Though there are several popular ways of programming multicore machines such as OpenMP and Java-multithreading, we prefer ParC for learning parallel programming:

- It is significantly simpler than OpenMP and is naturally learned and used by C programmers.
- In some sense, it is more powerful than OpenMP and Java due to its extended scoping rules and support for optimizing locality of shared memory references.

As parallel hardware becomes more and more popular the demand to develop methodologies for teaching parallel programming becomes more evident. This book promotes the idea that parallel programming is best taught through the issue of **efficiency**. The claim is, that parallel programming capability is best gained through a systematic learning of programming techniques used to derive better efficient versions of parallel programs. Rather than surveying, parallel architectures, programming styles and a verity of applications, as is the case in many of the existing books in this field.

Finally, the only motive in developing and using parallel machines and languages is to speedup performances. This calls for a highly skilled careful programmer, who can create efficient programs, which exploit all potential advantages of the parallel hardware. A parallel algorithm can be programmed in many ways, yielding different performances. Lack of experience and careless programming can easily lead to extreme situations, namely when the performances of a parallel program executed on a given parallel machine, are actually worse than its sequential version.

In this book the method used to evaluate the expected performances of programs is **analytical**. The source code of the program is analyzed via a formula to determine the execution time of a program based on derived upper and a lower bounds. Next a speedup formula is derived and is separated to the different factors that may block the program from achieving a good speedup. Following this procedure it is possible to identify bottlenecks areas in the program and correct them. We shortly consider a second method to locate limiting factors that is based on the ParC simulator. The program is executed on a simulator, and its time measurements are used to spot performance hot-spots. Two types of time statistics for a parallel execution of a program are defined. Ideal-times are execution times that reflect optimal execution

of the program without any practical considerations (such as number of cores and overheads), while the Effective-times reflect limiting factors. An efficient program is obtained (using a simulator) when the gap between the ideal times and the effective times is reduced.

# Acknowledgements

*ParC* was first developed by Yosi Ben-Asher and Marc Snir at NYU the first implementation run on the Ultra Computer at NYU. At a later stage *ParC* was implemented on the Makbilan research project led by Larry Rudolph in the Hebrew University. The compiler and simulator were written by Yosi Ben-Asher. A new run-time environment for *ParC* on the Makbilan was implemented by Dror Feitelson, Moshe Ben Ezra, and Lior Picherski. Gudy Habber modified the *ParC* simulator from SUN workstation to PC-DOS systems.



# Contents

<b>1</b>	<b>Basic Concepts in Parallel Algorithms and Parallel Programming . . .</b>	<b>1</b>
1.1	Parallel Machines . . . . .	1
1.2	Basic Concepts of Parallel Programs . . . . .	7
1.2.1	Partial Orders as Basic Parallel Programs . . . . .	9
1.3	Time Diagrams . . . . .	14
1.4	Execution Times of Parallel Programs . . . . .	17
1.5	Moving to a General Parallel Programming Language . . . . .	22
1.6	Developing Practical Parallel Algorithms . . . . .	29
1.6.1	Some Basic Properties of Efficient Parallel Algorithms . . . . .	30
1.6.2	Reducing Synchronization in Practical Parallel Algorithms . . . . .	33
1.6.3	Some General Methods for Implementing the Scalability of Parallel Algorithms . . . . .	36
1.7	Exercises . . . . .	39
1.8	Bibliographic Notes . . . . .	42
	References . . . . .	43
<b>2</b>	<b>Principles of Shared Memory Parallel Programming Using <i>ParC</i> . . .</b>	<b>45</b>
2.1	Introduction . . . . .	45
2.1.1	System Structure . . . . .	47
2.2	Parallel Constructs . . . . .	48
2.3	Scoping Rules . . . . .	55
2.3.1	Using Static Variables . . . . .	60
2.3.2	Bypassing the Scoping Rules . . . . .	61
2.4	Effect of Execution Order and Atomicity . . . . .	62
2.5	Parallel Algorithms for Finding the Minimum . . . . .	63
2.5.1	Divide and Conquer Min . . . . .	64
2.5.2	Crash Min . . . . .	64
2.5.3	Random Min . . . . .	66
2.5.4	Divide and Crush . . . . .	67
2.5.5	Asynchronous Min . . . . .	71

2.6	Recursion and Parallel Loops: The Parallel Prefix Example . . . .	72
2.7	Minimum Spanning Tree . . . . .	77
2.8	Exercises . . . . .	82
2.8.1	Questions About Summing . . . . .	84
2.8.2	Questions About Min Algorithms . . . . .	86
2.8.3	Questions About Snake Sorting . . . . .	86
2.9	Bibliographic Notes . . . . .	88
	References . . . . .	90
<b>3</b>	<b>Locality and Synchronization . . . . .</b>	<b>93</b>
3.1	Locality of Memory Operations via the Scoping Rules . . . . .	94
3.2	Locality of Memory Operations Through Partitioned Arrays . . . .	96
3.3	Synchronization . . . . .	102
3.4	Synchronization Through Shared Variables . . . . .	103
3.5	Fetch and Add Synchronization . . . . .	105
3.6	The Sync Instruction . . . . .	110
3.7	Controlling the Execution by Using the Yield and Switch Instructions . . . . .	114
3.8	Fairness Related Issues . . . . .	117
3.9	Scheduling Policy . . . . .	118
3.10	Forced Termination . . . . .	118
3.11	Exercises . . . . .	122
3.11.1	Replacing Sync with f&a() . . . . .	122
3.11.2	Draw Program . . . . .	123
3.12	Bibliographic Notes . . . . .	124
	References . . . . .	126
<b>4</b>	<b>Multicore Machines . . . . .</b>	<b>129</b>
4.1	Caches . . . . .	130
4.2	Basic Mode of Operation of Multicore Machines . . . . .	132
4.3	The MESI Protocol for Memory Coherence in Multicore Machines . . . . .	134
4.4	Counting MESI's Transitions and Overhead . . . . .	138
4.5	The MESI Protocol Preserves Sequential Consistency . . . . .	140
4.6	The MOESI Extension . . . . .	142
4.7	Preserving the Locality Principle of ParC . . . . .	143
4.8	False Sharing of Cache Lines . . . . .	147
4.9	Controlling Cache Line Sharing at Different Times . . . . .	149
4.10	Prefetching Data . . . . .	151
4.11	Exercises . . . . .	154
4.12	Bibliographic Notes . . . . .	156
	References . . . . .	157
<b>5</b>	<b>Improving the Performance of Parallel Programs: The Analytical Approach . . . . .</b>	<b>159</b>
5.1	Introduction . . . . .	159

- 5.2 A Simple Model of a Virtual Machine . . . . . 160
- 5.3 The *ParC* Virtual Machine Model . . . . . 161
- 5.4 Speedup Notion for Programs . . . . . 167
- 5.5 Using the Speedup Factors . . . . . 170
- 5.6 The Effect of Scheduling on  $T(R)$  . . . . . 175
- 5.7 Accounting for Memory References for Multicore Machines . . . . 183
- 5.8 On the Usage of a Simulator . . . . . 185
- 5.9 Conclusions . . . . . 187
- 5.10 Exercises . . . . . 188
  - 5.10.1 Optimizing Speedup Factors . . . . . 188
  - 5.10.2 Amdhal’s Law . . . . . 189
  - 5.10.3 Scheduling . . . . . 190
  - 5.10.4 Two-Processor Machines . . . . . 191
  - 5.10.5 Sync-Based Exercises . . . . . 192
  - 5.10.6 The Effect of Sync on the Execution Time . . . . . 193
- 5.11 Bibliographic Notes . . . . . 195
- References . . . . . 196
  
- 6 Compilation Techniques . . . . . 197**
  - 6.1 Introduction . . . . . 197
  - 6.2 Inserting Explicit Context Switch Instructions . . . . . 200
  - 6.3 Using Function Calls to Spawn Threads . . . . . 202
  - 6.4 Scheduling and Supporting Context-Switch . . . . . 207
  - 6.5 Memory Management of Stacks . . . . . 211
  - 6.6 Bibliographical Notes . . . . . 215
  - References . . . . . 216
  
- 7 Working with Sequential Versions . . . . . 219**
  - 7.1 The Transformation Set . . . . . 222
    - 7.1.1 *parfor* Transformation . . . . . 223
    - 7.1.2 *parblock* Transformation . . . . . 223
    - 7.1.3 Privatization Transformation . . . . . 223
    - 7.1.4 Chunk Transformations . . . . . 224
    - 7.1.5 Loop Reorder Transformation . . . . . 225
    - 7.1.6 Asynchronous Transformation . . . . . 226
    - 7.1.7 Localization Transformation . . . . . 227
  - 7.2 Loop Transformations . . . . . 228
    - 7.2.1 Loop Interchange . . . . . 228
    - 7.2.2 Loop Distribution and Loop Fusion . . . . . 228
  - 7.3 Case Study: Transitive Closure . . . . . 230
  - 7.4 Case Study: Matrix Multiplication . . . . . 236
  - 7.5 Case Study: PDE—Partial Differential Equations . . . . . 238
  - 7.6 Case Study: Dynamic Chunking . . . . . 241
  - 7.7 Case Study: Using Pointer Arrays . . . . . 244
  - 7.8 Parallelization of Sequential Loops . . . . . 245
  - 7.9 Lamport’s Diagonalization . . . . . 246

- 7.10 Exercises . . . . . 252
  - 7.10.1 Q1 . . . . . 252
  - 7.10.2 Q2 . . . . . 255
- 7.11 Bibliographic Notes . . . . . 256
- References . . . . . 257
  
- 8 Performance and Overhead Measurements . . . . . 259**
  - 8.1 Introduction . . . . . 259
  - 8.2 Processing Rates . . . . . 260
  - 8.3 The Effect of MESI/MOESI Bus Transactions on the Processing Rates . . . . . 261
  - 8.4 The Effect of Hardware Threads . . . . . 263
  - 8.5 Initial Measurements of MESI . . . . . 266
  - 8.6 Remote Access Penalty . . . . . 267
  - 8.7 Wrapping MESI Overhead by Local Accesses . . . . . 269
  - 8.8 Experimenting with the MESI's Protocol States . . . . . 272
  - 8.9 Spawning New Threads . . . . . 273
  - 8.10 Bibliographic Notes . . . . . 275
  - References . . . . . 276