

GPU Computing and Applications

Yiyu Cai • Simon See
Editors

GPU Computing and Applications

 Springer

Editors

Yiyu Cai
Nanyang Technological University
Singapore

Simon See
Nvidia
Singapore

ISBN 978-981-287-133-6

ISBN 978-981-287-134-3 (eBook)

DOI 10.1007/978-981-287-134-3

Springer Singapore Heidelberg New York Dordrecht London

Library of Congress Control Number: 2014955661

© Springer Science+Business Media Singapore 2015

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)

Preface

Graphics Processing Unit (GPU) technology is playing an important role in computing today. This book presents a collection of state-of-the-art research on GPU computing and their applications. The major part of this book is selected from the work submitted to the 2013 Symposium on GPU Computing and Applications jointly organized by the Institute for Media Innovation of Nanyang Technological University, and NVIDIA Corporation (South East Asia).

The book addresses the fundamental issues in GPU computing with a focus on big data processing. Three major domains of GPU applications are covered in the book including (1) Engineering design and simulation; (2) Biomedical Sciences; and (3) Interactive & Digital Media.

This book has 17 chapters. Each chapter is carefully selected to present ideas and techniques with insight in a specific area. In Chap. 1, Panpan Cai et al. will report a GPU-enabled parallel genetic algorithm for path planning. In Chap. 2, Alexandre Kaspar and Bailin Deng will introduce real-time deformation of constrained meshes using GPU. In Chap. 3, Yanlin Luo et al. will investigate GPU-based real-time volume interaction for scientific visualization education. In Chap. 4, Petros Papanikolaou and George Papagiannakis will illustrate real-time separable sub-surface scattering for animated virtual characters. In Chap. 5, Yusha Li et al. will describe adaptive NURBS tessellation on GPU. In Chap. 6, Huagen Wan et al. will discuss a graphics native approach to identifying surface atoms of macromolecules. In Chap. 7, Farhoosh Alghabi et al. will explain their scalable software framework for stateful stream data processing on multiple GPUs. In Chap. 8, Tananan Pattanangkur et al. will share their solution for high performance mobile medical imaging. In Chap. 9, [David Mainzer](#) and Gabriel Zachmann will showcase their collision detection based on fuzzy scene subdivision. In Chap. 10, Philip Boyer et al. will present the smoothed particle hydrodynamics applied to cartilage deformation. In Chap. 11, Kyrlyo Shegeda and Pierre Boulanger will describe a GPU-based real-time algorithm for virtual viewpoint rendering from multi-video. In Chap. 12, Ettikan K. Karuppiyah et al. will illustrate a middleware framework for programmable multi-GPU based big data applications. In Chap. 13, Byungjoon Chang et al. will talk on the efficient implementation of a real-time Kd-tree

construction algorithm. In Chap. 14, [Niko Lukac](#) and [Borut Zalik](#) will discuss fast approximate k-nearest neighbors search using GPGPU. In Chap. 15, Shafaatunnur Hasan et al. will share their soft computing methods for big data problems. In Chap. 16, Martin Němec and Lumír Janošek will show a numerical solution of BVP on GPU with application to path-planning. And in Chap. 17, Amirul Abdullah et al. will investigate fast multi-keyword range search in GPGPU.

Readers will benefit from this book which is contributed by experienced GPU researchers and educators. The book may also motivate researchers and developers to develop new possible applications of GPU technology in various areas.

Singapore
Singapore

Yiyu Cai
Simon See

Contents

1	A GPU-Enabled Parallel Genetic Algorithm for Path Planning of Robotic Operators	1
	Panpan Cai, Yiyu Cai, Indhumathi Chandrasekaran, and Jianmin Zheng	
1.1	Introduction	1
1.1.1	Motivation	1
1.1.2	Objectives	2
1.1.3	Organization of the Chapter	2
1.2	Prior Arts	2
1.2.1	GA Study	2
1.2.2	Parallel GA (PGA) Study	3
1.3	GPU-Enabled PGA	4
1.3.1	GPU Fundamentals	4
1.3.2	GPU PGA Design and Implementation	5
1.4	GPU-Enabled PGA Application for Collision-Free Path Planning	8
1.4.1	Collision Avoidance Path Planning	8
1.4.2	Terminal Condition	9
1.4.3	Global Optimization	9
1.4.4	Performance	11
	Conclusions	12
	References	12
2	Real-Time Deformation of Constrained Meshes Using GPU	15
	Alexandre Kaspar and Bailin Deng	
2.1	Introduction	15
2.1.1	Related Work	16
2.1.2	Overview	17
2.2	Overview of the Method	18
2.2.1	Problem Formulation	18
2.2.2	Numerical Solution	20

2.3	General Implementation Strategies	22
2.4	CUDA Implementation Details	24
2.4.1	Kernels	25
2.4.2	Sparse Linear Algebra	28
2.5	Results	28
2.6	Limitation and Future Work	32
	Conclusion	32
	References	33
3	GPU-Based Real-Time Volume Interaction for Scientific Visualization Education	35
	Yanlin Luo, Zhongke Wu, Zuying Luo, and Yanhong Luo	
3.1	Introduction	35
3.2	Related Work	36
3.2.1	GPU-Accelerated Volume Rendering	36
3.2.2	Volume Interaction	37
3.2.3	Volume Illustration	38
3.3	The Proposed Method	38
3.3.1	Transfer Function Design	38
3.3.2	Plane Tools	40
3.3.3	Superquadric Tools	42
3.3.4	Virtual Lenses	44
3.3.5	User Interaction and Implementation	45
3.4	Results	47
	Conclusions	48
	References	50
4	Real-Time Separable Subsurface Scattering for Animated Virtual Characters	53
	P. Papanikolaou and G. Papagiannakis	
4.1	Introduction	53
4.2	Previous Work	54
4.3	Separable Subsurface Scattering for Dynamic Surfaces	55
4.4	Implementation	56
4.4.1	Light's Transmission Through Thin Skin	59
4.4.2	Subsurface Scattering Implementation Novelties	62
4.5	Comparison with Ground Truth	63
	Conclusions	65
	References	66
5	Adaptive NURBS Tessellation on GPU	69
	Yusha Li, Xingjiang Lu, Wenjing Zhang, and Guozhao Wang	
5.1	Introduction	69
5.2	Estimating the Tessellation Intervals	71
5.2.1	Tessellation Intervals for Rational Bézier Curves	71
5.2.2	Tessellation Intervals for Rational Bézier Surfaces	75

5.3	Creating Transition Regions	76
5.3.1	Extracting Bézier Patches	76
5.3.2	Filling the Gaps	77
5.4	Implementations on GPU	77
5.5	Experiment Results	80
5.5.1	Comparisons to Zheng and Sederberg's Method	80
5.5.2	Run Time on CPU and GPU	81
	Conclusions	83
	References	84
6	Graphics Native Approach to Identifying Surface Atoms of Macromolecules	85
	Huagen Wan, Yunqing Guan, and Yiyu Cai	
6.1	Introduction	85
6.2	Prior Work	87
6.3	Algorithm Overview and Implementation	88
6.3.1	Algorithm Overview	88
6.3.2	Implementation	89
6.3.3	Improvements	91
6.4	Experimental Results and Discussions	92
	Conclusions	93
	References	95
7	A Scalable Software Framework for Stateful Stream Data Processing on Multiple GPUs and Applications	99
	Farhoosh Alghabi, Ulrich Schipper, and Andreas Kolb	
7.1	Introduction	100
7.2	Related Work	102
7.3	The Framework	103
7.3.1	Basic Concepts	104
7.3.2	Distributed Graph	105
7.3.3	Multiple Graph Instantiation	106
7.4	Experimental Evaluation	108
7.4.1	Comparison of Preliminary Implementations	109
7.4.2	Scalability and Feedback	112
7.5	Applications	113
7.5.1	Information Security Using Crypto- and Steganography	113
7.5.2	Crystallography Using a pnCCD Camera	114
	Conclusion	117
	References	117

8	The Design of SkyPACS: A High-Performance Mobile Medical Imaging Solution	119
	Tananan Pattanangkur, Sikana Tanupabrungson, Katchaguy Areekijseree, Sarunya Pumma, and Tiranee Achalaku	
8.1	Introduction	120
8.2	Imagery Procedure	121
8.3	Features of SkyPACS	122
8.4	Software Design	123
8.5	Implementation and Deployment	127
8.6	Product Comparisons	128
	Conclusion	130
	Appendix A: Screenshots	131
	References	132
9	Collision Detection Based on Fuzzy Scene Subdivision	135
	David Mainzer and Gabriel Zachmann	
9.1	Introduction	135
9.1.1	Our Contributions	136
9.2	Previous Work	137
9.2.1	Approaches Using Bounding Volume Hierarchies	137
9.2.2	GPU-Based Collision Detection	138
9.3	Sweep-Plane Technique Using PCA for Collision Detection	138
9.3.1	Thread Management	140
9.4	Object Subdivision Using Fuzzy C-Means	141
9.5	GPU-Based Collision Detection	143
9.5.1	Accuracy and Limitations	145
9.6	Results	146
9.6.1	Benchmarking	146
	Conclusions and Future Work	148
	References	149
10	Smoothed Particle Hydrodynamics Applied to Cartilage Deformation	151
	Philip Boyer, Sean LeBlanc, and Chris Joslin	
10.1	Introduction and Background	151
10.2	Materials and Methods	154
10.2.1	Elastic Solid Forces	154
10.2.2	Rigid Boundary Collision Handling	157
10.2.3	Implementation	158
10.3	Results	159
10.3.1	Solid Rod and Falling Wedge Tests	159
10.3.2	Cartilage Simulation	162
	Conclusions and Future Work	163
	References	163

- 11 A GPU-Based Real-Time Algorithm for Virtual Viewpoint Rendering from Multi-video** 167
 Kyrlo Shegeda and Pierre Boulanger
- 11.1 Introduction 167
- 11.2 Common Plane Sweeping Algorithm 169
 - 11.2.1 Depth-Map Estimation Algorithm 169
 - 11.2.2 Pixel Similarity Function 171
 - 11.2.3 Projective Block Matching 171
 - 11.2.4 Virtual Viewpoint Rendering 172
 - 11.2.5 GPU-Accelerated Algorithm and Its Implementation 174
 - 11.2.6 Constructing OpenGL Model-View and Projection Matrices 179
- 11.3 Experimental Results 180
- Conclusion 183
- References 184

- 12 A Middleware Framework for Programmable Multi-GPU-Based Big Data Applications** 187
 Ettikan K. Karuppiyah, Yong Keh Kok, and Keeratpal Singh
- 12.1 Introduction 187
- 12.2 Related Work 190
- 12.3 Middleware Framework Design 192
 - 12.3.1 Big Data Needs 192
 - 12.3.2 Presentation Layer 194
 - 12.3.3 Interface Layer 194
 - 12.3.4 Middleware Layer 194
 - 12.3.5 Orchestration Engine (with Example of Use Case) 195
 - 12.3.6 Storage 196
 - 12.3.7 Mi-AccLib and Analytics Component 197
- 12.4 Implementation 199
- 12.5 Results 201
- Conclusions 204
- References 205

- 13 On the Efficient Implementation of a Real-Time Kd-Tree Construction Algorithm** 207
 Byungjoon Chang, Woong Seo, and Insung Ihm
- 13.1 Background and Our Contribution 207
- 13.2 Optimizations for the Large-Node Stage 208
 - 13.2.1 Triangle Sorting with Respect to Splitting Planes 209
 - 13.2.2 AABB Computations for Active Large Nodes 212
- 13.3 Optimizations for the Small-Node Stage 213
- 13.4 Experimental Results 213
- Concluding Remarks 216

Appendix: A Single-Kernel Implementation for the Triangle-Sorting Process (Sect. 13.2.1.2)	217
References	219
14 Fast Approximate k-Nearest Neighbours Search Using GPGPU . . .	221
Niko Lukač and Borut Žalik	
14.1 Introduction	221
14.2 Related Work	223
14.3 Parallel Multi-probe LSH	224
14.3.1 Locality-Sensitive Hashing	224
14.3.2 Parallel MLSH Using CUDA	227
14.4 Results	230
Conclusion	232
References	233
15 Soft Computing Methods for Big Data Problems	235
Shafaatunnur Hasan, Siti Mariyam Shamsuddin, and Noel Lopes	
15.1 Introduction	236
15.2 Related Work	237
15.3 GPU Machine Learning Library Implementation	238
15.3.1 Parallel Multiple Back-Propagation	238
15.3.2 Parallel Self-Organizing Map	239
15.4 Experimental Setup	239
15.4.1 Dataset Preparation	239
15.4.2 Performance Measurement	240
15.5 Experimental Result and Analysis	242
15.5.1 Speed Analysis	243
15.5.2 Classification Analysis	245
Conclusion	246
References	246
16 Numerical Solution of BVP on GPU with Application to Path Planning	249
Lumír Janošek, Martin Němec, and Radoslav Fasuga	
16.1 Introduction	249
16.2 Harmonic Potential Field	250
16.3 Iterative Methods	251
16.4 Implementation	252
16.5 Results	254
Conclusion	256
References	256
17 Fast Multi-Keyword Range Search Using GPGPU	259
Amirul Abdullah, Amril Nazir, Mohanavelu Senapan Soo Saw Meng, and Ettikan Karuppiah	
17.1 Introduction	259
17.2 Background	261

- 17.2.1 Keyword Search 261
- 17.2.2 Binary Search 262
- 17.2.3 Multi-keyword Search (P-ary Search) 262
- 17.3 Implementation 264
 - 17.3.1 Data Packing 265
 - 17.3.2 Memory Coalescing 266
 - 17.3.3 Shared Memory 267
- 17.4 Experimental Evaluation 268
 - 17.4.1 Response Time 268
 - 17.4.2 Speed-Ups 271
 - 17.4.3 Throughput 272
- Conclusions and Future Work 272
- References 273
- Index 275**

List of Contributors

Amirul Abdullah MIMOS Berhad, Kuala Lumpur, Malaysia

Tiranee Achalakul Computer Engineering, King Mongkut's University of Technology Thonburi, Bangkok, Thailand

Farhoosh Alghabi Institute for Vision and Graphics, University of Siegen, Siegen, Germany

Katchaguy Areekijseree Computer Engineering, King Mongkut's University of Technology Thonburi, Bangkok, Thailand

Pierre Boulanger Computing Science Department, University of Alberta, Edmonton, AB, Canada

Philip Boyer Department of Systems and Computer Engineering, Carleton University, Ottawa, ON, Canada

Panpan Cai School of Mechanical & Aerospace Engineering, Nanyang Technological University, Singapore, Singapore

Yiyu Cai Nanyang Technological University, Singapore, Singapore

Institute for Media Innovation, Nanyang Technological University, Singapore, Singapore

Indhumathi Chandrasekaran School of Mechanical & Aerospace Engineering, Nanyang Technological University, Singapore, Singapore

Byungjoon Chang Digital Media & Communications R&D Center, Samsung Electronics, Suwon-si, Gyeonggi-do, South Korea

Bailin Deng Computer Graphics and Geometry Laboratory, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland

Radoslav Fasuga Department of Computer Science, VŠB-Technical University, Ostrava, Czech Republic

Yunqing Guan Institute for Media Innovation, Nanyang Technological University, Singapore, Singapore

Shafaatunnur Hasan Soft Computing Research Group, Faculty of Computing, Universiti Teknologi Malaysia, Skudai, Johor, Malaysia

Wan Huagen State Key Lab of CAD&CG, Zhejiang University, Hangzhou, China

Insung Ihm Department of Computer Science and Engineering, Sogang University, Seoul, South Korea

Lumír Janošek Department of Computer Science, VŠB-Technical University, Ostrava, Czech Republic

Chris Joslin Department of Systems and Computer Engineering, Carleton University, Ottawa, ON, Canada

Ettikan K. Karupiah MIMOS Berhad, Kuala Lumpur, Malaysia

Alexandre Kaspar Computer Graphics and Geometry Laboratory, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland

Yong Keh Kok MIMOS Berhad, Kuala Lumpur, Malaysia

Andreas Kolb Institute for Vision and Graphics, University of Siegen, Siegen, Germany

Sean LeBlanc Department of Systems and Computer Engineering, Carleton University, Ottawa, ON, Canada

Yusha Li School of Computer Engineering, Nanyang Technological University, Singapore, Singapore

Noel Lopes UDI, Institute of Guarda, Guarda, Portugal

CISUC, University of Coimbra, Coimbra, Portugal

Xingjiang Lu Department of Mathematics, Zhejiang University, Hangzhou, China

Niko Lukač Faculty of Electrical Engineering and Computer Science, University of Maribor, Maribor, Slovenia

Yanlin Luo The College of Information Science and Technology, Beijing Normal University, Beijing, China

Zuying Luo The College of Information Science and Technology, Beijing Normal University, Beijing, China

Yanhong Luo The College of Electrical Engineering, Northwest University for Nationalities, Lanzhou, China

The School of Nuclear Science and Technology, Lanzhou University, Lanzhou, China

David Mainzer Clausthal University, Clausthal-Zellerfeld, Germany

Soo Saw Meng MIMOS Berhad, Kuala Lumpur, Malaysia

Amril Nazir MIMOS Berhad, Kuala Lumpur, Malaysia

Martin Němec Department of Computer Science, VŠB-Technical University, Ostrava, Czech Republic

G. Papagiannakis Computer Science Department, University of Crete, Heraklion, Greece

Foundation for Research and Technology Hellas, Heraklion, Greece

P. Papanikolaou Computer Science Department, University of Crete, Heraklion, Greece

Foundation for Research and Technology Hellas, Heraklion, Greece

Tananan Pattanangkur Computer Engineering, King Mongkut's University of Technology Thonburi, Bangkok, Thailand

Sarunya Pumma Computer Engineering, King Mongkut's University of Technology Thonburi, Bangkok, Thailand

Ulrich Schipper Institute for Vision and Graphics, University of Siegen, Siegen, Germany

Mohanavelu Senapan MIMOS Berhad, Kuala Lumpur, Malaysia

Woong Seo Department of Computer Science and Engineering, Sogang University, Seoul, South Korea

Siti Mariyam Shamsuddin UDI, Institute of Guarda, Guarda, Portugal

Kyrylo Shegeda Computing Science Department, University of Alberta, Edmonton, AB, Canada

Keeratpal Singh MIMOS Berhad, Kuala Lumpur, Malaysia

Sikana Tanupabrungron Computer Engineering, King Mongkut's University of Technology Thonburi, Bangkok, Thailand

Guozhao Wang Department of Mathematics, Zhejiang University, Hangzhou, China

Zhongke Wu The College of Information Science and Technology, Beijing Normal University, Beijing, China

Gabriel Zachmann University of Bremen, Bremen, Germany

Borut Žalik Faculty of Electrical Engineering and Computer Science, University of Maribor, Maribor, Slovenia

Wenjing Zhang School of Computer Engineering, Nanyang Technological University, Singapore, Singapore

Jianmin Zheng School of Computer Engineering, Nanyang Technological University, Singapore, Singapore