
Parallel and Distributed Computing Handbook

Albert Y. Zomaya, Editor

McGraw-Hill

New York San Francisco Washington, D.C. Auckland Bogotá
Caracas Lisbon London Madrid Mexico City Milan
Montreal New Delhi San Juan Singapore
Sydney Tokyo Toronto

Contents

Foreword	xix
Preface	xxi
Acknowledgments	xxiii
List of contributors	xxv

Part I Theory

Foundations

Chapter 1. Parallel and Distributed Computing: The Scene, the Props, the Players 5

Albert Y. Zomaya

1.1	A Perspective	
1.2	Parallel Processing Paradigms	7
1.3	Modeling and Characterizing Parallel Algorithms	11
1.4	Cost vs. Performance Evaluation	13
1.5	Software and General-Purpose PDC	15
1.6	A Brief Outline of the Handbook	16
1.7	Recommended Reading	19
1.8	References	21

Chapter 2. Semantics of Concurrent Programming 24

J. Desharnais, A. Mili, R. Mili, J. Mullins, and Y. Slimani

2.1	Models of Concurrent Programming	25
2.2	Semantic Definitions	27

2.3	Axiomatic Semantic Definitions	30
2.4	Denotational Semantic Definitions	36
2.5	Operational Semantic Definitions	54
2.6	Summary and Prospects	57
2.7	References	57
Chapter 3. Formal Methods: A Petri Nets Based Approach		59
<i>Giorgio De Michelis, Lucia Pomello, Eugenio Battiston, Fiorella De Cindio, and Carla Simone</i>		
3.1	Process Algebras	61
3.2	PETRI Nets	68
3.3	High-Level Net Models	78
3.4	Conclusions	84
3.5	References	86
Chapter 4. Complexity Issues in Parallel and Distributed Computing		89
<i>E. V. Krishnamurthy</i>		
4.1	Introduction	89
4.2	Turing Machine as the Basis, and Consequences	93
4.3	Complexity Measures for Parallelism	101
4.4	Parallel Complexity Models and Resulting Classes	103
4.5	VLSI Computational Complexity	121
4.6	Complexity Measures for Distributed Systems	121
4.7	Neural Networks and Complexity Issues	121
4.8	Other Complexity Theories	123
4.9	Concluding Remarks	124
4.10	References	125
Chapter 5. Distributed Computing Theory		127
<i>Hagit Attiya</i>		
5.1	The Computation Model	129
5.2	A Simple Example	131
5.3	Leader Election	132
5.4	Sparse Network Covers and Their Applications	138
5.5	Ordering of Events	142
5.6	Resource Allocation	146
5.7	Tolerating Processor Failures in Synchronous Systems	146
5.8	Tolerating Processor Failures in Asynchronous Systems	151
5.9	Other Types of Failures	154
5.10	Wait-Free Implementations of Shared Objects	156
5.11	Final Comments	157
5.12	References	158

Models

Chapter 6. PRAM MODELS	163
<i>Lydia I. Kronsjö</i>	
6.1 Introduction	163
6.2 Techniques for the Design of Parallel Algorithms	165
6.3 The PRAM Model	168
6.4 Optimality and Efficiency of Parallel Algorithms	171
6.5 Basic PRAM Algorithms	175
6.6 The NC-Class	180
6.7 P-Completeness: Hardly Parallelizable Problems	180
6.8 Randomized Algorithms and Parallelism	181
6.9 List Ranking Revisited: Optimal $O(\log n)$ Deterministic List Ranking	184
6.10 Taxonomy of Parallel Algorithms	186
6.11 Deficiencies of the PRAM Model	186
6.12 Summary	189
6.13 References	189
Chapter 7. Broadcasting with Selective Reduction: A Powerful Model of Parallel Computation	192
<i>Selim G. Akl and Ivan Stojmenović</i>	
7.1 Introduction	192
7.2 A Generalized BSR Model	197
7.3 One Criterion BSR Algorithms	200
7.4 Two Criteria BSR Algorithms	212
7.5 Three Criteria BSR Algorithms	215
7.6 Multiple Criteria BSR Algorithms	218
7.7 Conclusions and future Work	220
7.8 References	221
Chapter 8. Dataflow Models	223
<i>R. Jagannathan</i>	
8.1 Kinds of Dataflow	224
8.2 Data-Driven Dataflow Computing Models	225
8.3 Demand-driven Dataflow Computing Models	230
8.4 Unifying Data-Driven and Demand-Driven	234
8.5 Lessons Learned and Future Trends	235
8.6 Summary	236
8.7 References	237
Chapter 9. Partitioning and Scheduling	239
<i>Hesham El-Rewini</i>	
9.1 Program Partitioning	241
9.2 Task Scheduling	243

x Parallel and Distributed Computing

9.3	Scheduling System Model	244
9.4	Communication Models	248
9.5	Optimal Scheduling Algorithms	249
9.6	Scheduling Heuristic Algorithms	255
9.7	Scheduling Nondeterministic Task Graphs	262
9.8	Scheduling Tools	266
9.9	Task Allocation	267
9.10	Heterogeneous Environments	268
9.11	Summary and Concluding Remarks	270
9.12	References	272

Chapter 10. Checkpointing in Parallel and Distributed Systems 274

Avi Ziv and Jehoshua Bruck

10.1	Introduction	274
10.2	Checkpointing Using Task Duplication	276
10.3	Techniques for Consistent Checkpointing	288
10.4	Conclusions and Future Directions	299
10.5	References	300

Chapter 11. Architecture for Open Distributed Software Systems 303

Kazi Farooqui and Luigi Logrippo

11.1	Introduction to Open Distributed Systems Architecture	303
11.2	Computational Model	307
11.3	Engineering Model	315
11.4	ODP Application	324
11.5	Conclusion and Directions for Future Research	327
11.6	References	328

Algorithms

Chapter 12. Fundamentals of Parallel Algorithms 333

Joseph F. Jájá

12.1	Introduction	333
12.2	Models of Parallel Computation	334
12.3	Balanced Trees	339
12.4	Divide and Conquer	342
12.5	Partitioning	345
12.6	Combining	350
12.7	Conclusions and Future Trends	353
12.8	Acknowledgment	354
12.9	References	354

Chapter 13. Parallel Graph Algorithms	355
<i>Stephan Olariu</i>	
13.1 Graph-Theoretic Concepts and Notation	356
13.2 Tree Algorithms	358
13.3 Algorithms for General Graphs	372
13.4 Algorithms for Particular Classes of Graphs	384
13.5 Concluding Remarks	401
13.6 References	402
Chapter 14. Parallel Computational Geometry	404
<i>Mikhail J. Atallah</i>	
14.1 Parallel CG: Why New Techniques Are Needed	405
14.2 Basic Subproblems	407
14.3 CG on the PRAM	409
14.4 CG on the Mesh	416
14.5 CG on the Hypercube	419
14.6 Other Parallel Models	420
14.7 Conclusions and Future Work	422
14.8 References	423
Chapter 15. Data Structures for Parallel Processing	429
<i>Sajal K. Das and Kwang-Bae Min</i>	
15.1 Arrays and Balanced Binary Trees	430
15.2 Linked Lists	432
15.3 Trees and Euler Tour	434
15.4 General Trees and Binarized Trees	435
15.5 Euler Tour vs. Parentheses String	436
15.6 Stacks	440
15.7 Queues	445
15.8 Priority Queues (Heaps)	448
15.9 Search Trees/Dictionaries	455
15.10 Conclusions	463
15.11 References	464
Chapter 16. Data Parallel Algorithms	466
<i>Howard Jay Siegel, Lee Wang, John John E. So, and Muthucumar Maheswaran</i>	
16.1 Chapter Overview	466
16.2 Machine Model	467
16.3 Impact of Data Distribution	469
16.4 CU/PE Overlap	476
16.5 Parallel Reduction Operations	480
16.6 Matrix and Vector Operations	487

16.7	Mapping Algorithms onto Partitionable Machines	489
16.8	Achieving Scalability Using a Set of Algorithms	492
16.9	Conclusions and Future Directions	494
16.10	References	497
Chapter 17. Systolic and VLSI Processor Arrays for Matrix Algorithms		500
<i>D. J. Evans and M. Gusev</i>		
17.1	Processor Array Implementations	500
17.2	VLSI Processor Arrays	501
17.3	Systolic Array Algorithms	508
17.4	Mathematical Methods in DSP	510
17.5	Implementation of Systolic Algorithms in DSP	517
17.6	Conjugate Gradient Method	530
17.7	Summary	535
17.8	References	536
Chapter 18. Direct Interconnection Networks		537
<i>Ivan Stojmenović</i>		
18.1	Topological Properties of Interconnection Networks	537
18.2	Hypercube Networks	547
18.3	Routing and Broadcasting	555
18.4	Conclusions	563
18.5	References	564
18.6	Suggested Readings	565
Chapter 19. Parallel and Communication Algorithms on Hypercube Multiprocessors		568
<i>Afonso Ferreira</i>		
19.1	Topological Aspects	569
19.2	Communication Issues	573
19.3	Useful Algorithmic Tools	577
19.4	Solving Problems	581
19.5	Conclusions and Future Directions	587
19.6	References	588

Part II Architectures and Technologies

Architectures

Chapter 20. RISC Architectures		595
<i>Manolis Katevenis</i>		
20.1	What is RISC?	596
20.2	Pipelining and Bypassing	599

20.3	Dependences and Parallelism in CISC and in RISC	604
20.4	Instruction Alignment, Size, and Format	609
20.5	Implementation Disadvantages of CISC	615
20.6	History, Perspective, and Conclusions	617
20.7	References	619
Chapter 21. Superscalar and VLIW Processors		621
<i>Thomas M. Conte</i>		
21.1	Superscalar Processors	622
21.2	VLIW Processors	634
21.3	Superscalar vs. VLIW: Which Is Better?	645
21.4	Bibliography	647
Chapter 22. SIMD-Processing: Concepts and Systems		649
<i>Michael Jurczyk and Thomas Schwederski</i>		
22.1	Basic Concepts	649
22.2	SIMD Machine Components	654
22.3	Associative Processing	660
22.4	Case Studies of SIMD Systems	662
22.5	Applications and Algorithms	669
22.6	Languages and Programming	673
22.7	Conclusions	677
22.8	References	677
Chapter 23. MIMD Architectures: Shared and Distributed Memory Designs		680
<i>Ralph Duncan</i>		
23.1	Proliferation of MIMD Designs	681
23.2	Shared Memory Architectures	682
23.3	Distributed Memory Architectures	689
23.4	Hybrid Shared/Distributed Memory Architectures	695
23.5	Conclusion	696
23.6	References	697
Chapter 24. Memory Models		699
<i>Leonidas J. Kontothanassis and Michael L. Scott</i>		
24.1	Memory Hardware Technology	700
24.2	Memory System Architecture	702
24.3	User-Level Memory Models	707
24.4	Memory Consistency Models	711
24.5	Implementation and Performance of Memory Consistency Models	714
24.6	Conclusions and Trends	718
24.7	References	719

Technologies

Chapter 25. Heterogeneous Computing	725
<i>Howard Jay Siegel, John K. Antonio, Richard C. Metzger, Min Tan, and Yan Alexander Li</i>	
25.1 Introduction	725
25.2 Mixed-Mode Systems	727
25.3 Examples of Existing Mixed-Machine HC Systems	733
25.4 Examples of Software Tools for Mixed-Machine HC Systems	735
25.5 A Conceptual Model for Automatic Mixed-Machine HC	739
25.6 Task Profiling and Analytical Benchmarking	741
25.7 Matching and Scheduling for Mixed-Machine HC Systems	747
25.8 Conclusions and Future Directions	756
25.9 References	758
Chapter 26. Cluster Computing	762
<i>Louis Turcotte</i>	
26.1 Technological Evolution	763
26.2 Overview of Clustering	767
26.3 Distinct Uses of Clusters	771
26.4 Open Issues	777
26.5 References	779
Chapter 27. Massively Parallel Processing with Optical Interconnections	780
<i>Eugen Schenfeld</i>	
27.1 Parallel Processing Motivations	782
27.2 General-Purpose Parallel Computers	785
27.3 How Much Interconnection?	793
27.4 Considerations in Choosing the Interconnection Topology	795
27.5 Optical Communication: Free-Space Interconnection	796
27.6 Conclusions and Future Work	808
27.7 References	808
Chapter 28. ATM-Based Parallel and Distributed Computing	811
<i>Salim Hariri and Bei Lu</i>	
28.1 Introduction	811
28.2 Broadband Integrated Service Data Network (B-ISDN)	812
28.3 ATM Protocols	814
28.4 ATM Switches	824
28.5 Host-to-Network Interfaces	831
28.6 Parallel and Distributed Computing Environment Over ATM	835
28.7 Conclusions and Future Directions	836
28.8 References	837

Part III Tools and Applications

Development Tools

Chapter 29. Parallel Languages	843
<i>R. H. Perrott</i>	
29.1 Introduction	843
29.2 Language Categories	844
29.3 Programming Languages	846
29.4 Summary	861
29.5 References	862
Chapter 30. Tools for Portable High-Performance Parallel Computing	865
<i>Doreen Y. Cheng</i>	
30.1 Introduction	865
30.2 Criteria for Evaluating Portability Support	867
30.3 Portable Message-Passing Libraries	871
30.4 Language-Centered Tools	878
30.5 Parallelizing Compilers and Preprocessors	886
30.6 Conclusion	892
30.7 References	893
Chapter 31. Visualization of Parallel and Distributed Systems	897
<i>Michael T. Heath</i>	
31.1 Performance Monitoring	898
31.2 Performance Visualization	899
31.3 Example	910
31.4 Future Directions	913
31.5 References	915
Chapter 32. Constructing Numerical Software Libraries for High-Performance Computer Environments	917
<i>Jack J. Dongarra and David W. Walker</i>	
32.1 Introduction	917
32.2 The BLAS as the Key to Portability	924
32.3 Block Algorithms and Their Derivation	925
32.4 LU Factorization	930
32.5 Data Distribution	932
32.6 Parallel Implementation	935
32.7 Optimization, Tuning, and Trade-Offs	941
32.8 Conclusions and Future Research Directions	948
32.9 References	951

Chapter 33. Testing of Distributed Programs	955
<i>K. C. Tai and Richard H. Carver</i>	
33.1 SYN-Sequences of Distributed Programs	956
33.2 Definitions of Correctness and Faults for Distributed Programs	961
33.3 Approaches to Testing Distributed Programs	963
33.4 Test Generation for Distributed Programs	968
33.5 Analysis and Replay of Program Executions	973
33.6 Building Testing Tools for Distributed Programs	975
33.7 Conclusions and Future Work	976
33.8 References	977

Applications

Chapter 34. Scientific Computation	981
<i>Timothy G. Mattson</i>	
34.1 Programming Models for Parallel Computing	982
34.2 Algorithms for Parallel Scientific Computing	983
34.3 Case Studies: Molecular Modeling	989
34.4 Trends	997
34.5 Further Reading	999
34.6 Conclusion	1000
34.7 References	1001
Chapter 35. Parallel and Distributed Simulation of Discrete Event Systems	1003
<i>Alois Ferscha</i>	
35.1 Simulation Principles	1003
35.2 "Classical" LP Simulation Protocols	1009
35.3 Conservative vs. Optimistic Protocols	1037
35.4 Conclusions and Outlook	1037
35.5 References	1039
Chapter 36. Parallelism for Image Understanding	1042
<i>Viktor K. Prasanna and Cho-Li Wang</i>	
36.1 Vision Tasks	1046
36.2 A Model of CM-5	1050
36.3 Scalable Parallel Algorithms	1051
36.4 Implementation Details and Experimental Results	1060
36.5 Concluding Remarks	1068
36.6 References	1069

Chapter 37. Parallel Computation in Biomedicine: Genetic and Protein Sequence Analysis	1071
<i>Tieng K. Yap, Ophir Frieder, and Robert L. Martino</i>	
37.1 The Origin of Genetic and Protein Sequence Data	1072
37.2 An Example Database: GenBank	1074
37.3 Residue Substitution Scoring Matrices	1077
37.4 Sequence Comparison Algorithms	1080
37.5 Parallel Techniques for Sequence Similarity Searching	1083
37.6 Performance	1089
37.7 Discussion and Conclusions	1093
37.8 Future Work	1095
37.9 References	1095
Chapter 38. Parallel Algorithms for Solving Stochastic Linear Programs	1097
<i>Amal De Silva and David Abramson</i>	
38.1 Stochastic Linear Programming	1098
38.2 Techniques for Solving Stochastic Linear Programs	1104
38.3 Comparison of Methods	1113
38.4 Conclusion and Future Directions	1115
38.5 References	1115
Chapter 39. Parallel Genetic Algorithms	1118
<i>Andrew Chipperfield and Peter Fleming</i>	
39.1 What Are Genetic Algorithms?	1118
39.2 Major Elements of the Genetic Algorithm	1121
39.3 Parallel GAs	1130
39.4 Conclusions and Future Trends	1140
39.5 References	1141
Chapter 40. Parallel Processing for Robotic Computations: A Review	1144
<i>Tarek M. Nabhan and Albert Y. Zomaya</i>	
40.1 Overview of Robotic Systems	1144
40.2 The Task Planner	1145
40.3 Sensing	1147
40.4 Robot Control	1149
40.5 Applications of Advanced Architectures for Robot Kinematics and Dynamics	1150
40.6 Summary, Conclusions, and Future Directions	1154
40.7 References	1155

Chapter 41. Distributed Flight Simulation: A Challenge for Software Architecture	1160
<i>Rick Kazman</i>	
41.1 The Challenge of Distributed Flight Simulation	1160
41.2 A Generic Flight Simulator	1162
41.3 Introduction to Software Architecture	1164
41.4 Structural Modeling	1166
41.5 Motivations for Structural Modeling	1167
41.6 Flight Simulator Software Architecture: Overview	1168
41.7 Flight Simulator Software Architecture: Base Types	1169
41.8 A Simplified Software Structure	1173
41.9 Requirements of Flight Simulation	1173
41.10 Lessons Learned/Future Directions	1176
41.11 Summary	1176
41.12 References	1176
Index	1179