

Introduction to Parallel Processing

Algorithms and Architectures

Behrooz Parhami

*University of California at Santa Barbara
Santa Barbara, California*

Plenum Press • New York and London

Contents

Part I. Fundamental Concepts	1
1. Introduction to Parallelism	3
1.1. Why Parallel Processing?	5
1.2. A Motivating Example	8
1.3. Parallel Processing Ups and Downs	13
1.4. Types of Parallelism: A Taxonomy	15
1.5. Roadblocks to Parallel Processing	16
1.6. Effectiveness of Parallel Processing	19
Problems	21
References and Suggested Reading	23
2. A Taste of Parallel Algorithms	25
2.1. Some Simple Computations	27
2.2. Some Simple Architectures	28
2.3. Algorithms for a Linear Array	30
2.4. Algorithms for a Binary Tree	34
2.5. Algorithms for a 2D Mesh	39
2.6. Algorithms with Shared Variables	40
Problems	41
References and Suggested Reading	43
3. Parallel Algorithm Complexity	45
3.1. Asymptotic Complexity	47
3.2. Algorithm Optimality and Efficiency	50
3.3. Complexity Classes	53
3.4. Parallelizable Tasks and the NC Class	55
3.5. Parallel Programming Paradigms	56
3.6. Solving Recurrences	58

Problems	61
References and Suggested Reading	63
4. Models of Parallel Processing	65
4.1. Development of Early Models	67
4.2. SIMD versus MIMD Architectures	69
4.3. Global versus Distributed Memory	71
4.4. The PRAM Shared-Memory Model	74
4.5. Distributed-Memory or Graph Models	77
4.6. Circuit Model and Physical Realizations	80
Problems	82
References and Suggested Reading	85
Part II. Extreme Models	87
5. PRAM and Basic Algorithms	89
5.1. PRAM Submodels and Assumptions	91
5.2. Data Broadcasting	93
5.3. Semigroup or Fan-In Computation	96
5.4. Parallel Prefix Computation	98
5.5. Ranking the Elements of a Linked List	99
5.6. Matrix Multiplication	102
Problems	105
References and Suggested Reading	108
6. More Shared-Memory Algorithms	109
6.1. Sequential Rank-Based Selection	111
6.2. A Parallel Selection Algorithm	113
6.3. A Selection-Based Sorting Algorithm	114
6.4. Alternative Sorting Algorithms	117
6.5. Convex Hull of a 2D Point Set	118
6.6. Some Implementation Aspects	121
Problems	125
References and Suggested Reading	127
7. Sorting and Selection Networks	129
7.1. What Is a Sorting Network	131
7.2. Figures of Merit for Sorting Networks	133
7.3. Design of Sorting Networks	135
7.4. Batcher Sorting Networks	136
7.5. Other Classes of Sorting Networks	141
7.6. Selection Networks	142
Problems	144
References and Suggested Reading	147

8. Other Circuit-Level Examples	149
8.1. Searching and Dictionary Operations	151
8.2. A Tree-Structured Dictionary Machine	152
8.3. Parallel Prefix Computation	156
8.4. Parallel Prefix Networks	157
8.5. The Discrete Fourier Transform	161
8.6. Parallel Architectures for FFT	163
Problems	165
References and Suggested Reading	168
Part III. Mesh-Based Architectures	169
9. Sorting on a 2D Mesh or Torus	171
9.1. Mesh-Connected Computers	173
9.2. The Shearsort Algorithm	176
9.3. Variants of Simple Shearsort	179
9.4. Recursive Sorting Algorithms	180
9.5. A Nontrivial Lower Bound	183
9.6. Achieving the Lower Bound	186
Problems	187
References and Suggested Reading	190
10. Routing on a 2D Mesh or Torus	191
10.1. Types of Data Routing Operations	193
10.2. Useful Elementary Operations	195
10.3. Data Routing on a 2D Array	197
10.4. Greedy Routing Algorithms	199
10.5. Other Classes of Routing Algorithms	202
10.6. Wormhole Routing	204
Problems	208
References and Suggested Reading	210
11. Numerical 2D Mesh Algorithms	211
11.1. Matrix Multiplication	213
11.2. Triangular System of Equations	215
11.3. Tridiagonal System of Linear Equations	218
11.4. Arbitrary System of Linear Equations	221
11.5. Graph Algorithms	225
11.6. Image-Processing Algorithms	228
Problems	231
References and Suggested Reading	233
12. Other Mesh-Related Architectures	235
12.1. Three or More Dimensions	237

12.2. Stronger and Weaker Connectivities	240
12.3. Meshes Augmented with Nonlocal Links	242
12.4. Meshes with Dynamic Links	245
12.5. Pyramid and Multigrid Systems	246
12.6. Meshes of Trees	248
Problems	253
References and Suggested Reading	256
Part IV. Low-Diameter Architectures	257
13. Hypercubes and Their Algorithms	259
13.1. Definition and Main Properties	261
13.2. Embeddings and Their Usefulness	263
13.3. Embedding of Arrays and Trees	264
13.4. A Few Simple Algorithms	269
13.5. Matrix Multiplication	272
13.6. Inverting a Lower Triangular Matrix	274
Problems	275
References and Suggested Reading	278
14. Sorting and Routing on Hypercubes	279
14.1. Defining the Sorting Problem	281
14.2. Bitonic Sorting on a Hypercube	284
14.3. Routing Problems on a Hypercube	285
14.4. Dimension-Order Routing	288
14.5. Broadcasting on a Hypercube	292
14.6. Adaptive and Fault-Tolerant Routing	294
Problems	295
References and Suggested Reading	298
15. Other Hypercubic Architectures	301
15.1. Modified and Generalized Hypercubes	303
15.2. Butterfly and Permutation Networks	305
15.3. Plus-or-Minus-2 ⁱ Network	309
15.4. The Cube-Connected Cycles Network	310
15.5. Shuffle and Shuffle-Exchange Networks	313
15.6. That's Not All, Folks!	316
Problems	317
References and Suggested Reading	320
16. A Sampler of Other Networks	321
16.1. Performance Parameters for Networks	323
16.2. Star and Pancake Networks	326
16.3. Ring-Based Networks	329

16.4. Composite or Hybrid Networks	335
16.5. Hierarchical (Multilevel) Networks	337
16.6. Multistage Interconnection Networks	338
Problems	340
References and Suggested Reading	343
Part V. Some Broad Topics	345
17. Emulation and Scheduling	347
17.1. Emulations among Architectures	349
17.2. Distributed Shared Memory	351
17.3. The Task Scheduling Problem	355
17.4. A Class of Scheduling Algorithms	357
17.5. Some Useful Bounds for Scheduling	360
17.6. Load Balancing and Dataflow Systems	362
Problems	364
References and Suggested Reading	367
18. Data Storage, Input, and Output	369
18.1. Data Access Problems and Caching	371
18.2. Cache Coherence Protocols	374
18.3. Multithreading and Latency Hiding	377
18.4. Parallel I/O Technology	379
18.5. Redundant Disk Arrays	382
18.6. Interfaces and Standards	384
Problems	386
References and Suggested Reading	388
19. Reliable Parallel Processing	391
19.1. Defects, Faults, Failures	393
19.2. Defect-Level Methods	396
19.3. Fault-Level Methods	399
19.4. Error-Level Methods	402
19.5. Malfunction-Level Methods	404
19.6. Degradation-Level Methods	407
Problems	410
References and Suggested Reading	413
20. System and Software Issues	415
20.1. Coordination and Synchronization	417
20.2. Parallel Programming	421
20.3. Software Portability and Standards	425
20.4. Parallel Operating Systems	427
20.5. Parallel File Systems	430

20.6. Hardware/Software Interaction	431
Problems	433
References and Suggested Reading	435
Part VI. Implementation Aspects	437
21. Shared-Memory MIMD Machines	439
21.1. Variations in Shared Memory	441
21.2. MIN-Based BBN Butterfly	444
21.3. Vector-Parallel Cray Y-MP	445
21.4. Latency-Tolerant Tera MTA	448
21.5. CC-NUMA Stanford DASH	450
21.6. SCI-Based Sequent NUMA-Q	452
Problems	455
References and Suggested Reading	457
22. Message-Passing MIMD Machines	459
22.1. Mechanisms for Message Passing	461
22.2. Reliable Bus-Based Tandem NonStop	464
22.3. Hypercube-Based nCUBE3	466
22.4. Fat-Tree-Based Connection Machine 5	469
22.5. Omega-Network-Based IBM SP2	471
22.6. Commodity-Driven Berkeley NOW	473
Problems	475
References and Suggested Reading	477
23. Data-Parallel SIMD Machines	479
23.1. Where Have All the SIMDs Gone?	481
23.2. The First Supercomputer: ILLIAC IV	484
23.3. Massively Parallel Goodyear MPP	485
23.4. Distributed Array Processor (DAP)	488
23.5. Hypercubic Connection Machine 2	490
23.6. Multiconnected MasPar MP-2	492
Problems	495
References and Suggested Reading	497
24. Past, Present, and Future	499
24.1. Milestones in Parallel Processing	501
24.2. Current Status, Issues, and Debates	503
24.3. TFLOPS, PFLOPS, and Beyond	506
24.4. Processor and Memory Technologies	508
24.5. Interconnection Technologies	510

24.6. The Future of Parallel Processing	513
Problems	515
References and Suggested Reading	517
Index	519