Contents

PREFACE xv

1 INTRODUCTION 1

1.1 Parallel and Distributed Architectures 2

1.1.1 The Need for Parallel and Distributed Computation, 2
1.1.2 Parallel Computing Systems and their Classification, 3

1.2 Models, Complexity Measures, and Some Simple Algorithms 8

1.2.1 Models, 8
1.2.2 Complexity Measures, 10
1.2.3 Examples: Vector and Matrix Computations, 16
1.2.4 Parallelization of Iterative Methods, 19

1.3 Communication Aspects of Parallel and Distributed Systems 27

1.3.1 Communication Links, 30
1.3.2 Data Link Control, 33

1.3.3 Routing, 37
1.3.4 Network Topologies, 39
1.3.5 Concurrency and Communication Tradeoffs, 68
1.3.6 Examples of Matrix–Vector Calculations, 71

1.4 Synchronization Issues in Parallel and Distributed Algorithms 88

1.4.1 Synchronous Algorithms, 88
1.4.2 Asynchronous Algorithms and the Reduction of the Synchronization Penalty, 95

Part 1 Synchronous Algorithms **109**

2 ALGORITHMS FOR SYSTEMS OF LINEAR EQUATIONS AND MATRIX INVERSION ***109***

2.1 Parallel Algorithms for Linear Systems with Special Structure 110

2.1.1 Triangular Matrices and Back Substitution, 110
2.1.2 Tridiagonal Systems and Odd–Even Reduction, 113

2.2 Parallel Direct Methods for General Linear Equations 118

2.2.1 Gaussian Elimination, 119
2.2.2 Triangularization Using Givens Rotations, 124

2.3 A Fast Direct Matrix Inversion Algorithm 128

2.4 Classical Iterative Methods For Systems of Linear Equations 130

2.5 Parallel Implementation of Classical Iterative Methods 135

2.5.1 An Example: Poisson's Equation, 137
2.5.2 Multigrid Methods, 139

2.6 Convergence Analysis of Classical Iterative Methods 143

2.6.1 Background on Maximum Norms and Nonnegative Matrices, 144
2.6.2 Convergence Analysis Using Maximum Norms, 151
2.6.3 Convergence Analysis Using a Quadratic Cost Function, 153
2.6.4 The Poisson Equation Revisited, 155

2.7 The Conjugate Gradient Method 158

2.7.1 Description of the Algorithm, 160
2.7.2 Speed of Convergence, 162

2.7.3 Preconditioned Conjugate Gradient Method, 164
2.7.4 Parallel Implementation, 165

2.8 Computation of the Invariant Distribution of a Markov Chain 166

2.9 Fast Iterative Matrix Inversion 173

3 ITERATIVE METHODS FOR NONLINEAR PROBLEMS 180

3.1 Contraction Mappings 181

3.1.1 General Results, 182
3.1.2 Contractions Over Cartesian Product Sets, 185
3.1.3 Some Useful Contraction Mappings, 191

3.2 Unconstrained Optimization 198

3.2.1 The Main Algorithms, 198
3.2.2 Convergence Analysis Using the Descent Approach, 202
3.2.3 The Case of a Convex Cost Function, 206
3.2.4 Nonlinear Algorithms, 207

3.3 Constrained Optimization 210

3.3.1 Optimality Conditions and the Projection Theorem, 210
3.3.2 The Gradient Projection Algorithm, 212
3.3.3 Scaled Gradient Projection Algorithms, 215
3.3.4 The Case of a Product Constraint Set: Parallel Implementations, 217
3.3.5 Nonlinear Algorithms, 219

3.4 Parallelization and Decomposition of Optimization Problems 224

3.4.1 Quadratic Programming, 225
3.4.2 Separable Strictly Convex Programming, 229
3.4.3 The Proximal Minimization Algorithm, 232
3.4.4 Augmented Lagrangian Methods, 243

3.5 Variational Inequalities 264

3.5.1 Examples of Variational Inequality Problems, 264
3.5.2 Preliminaries, 267
3.5.3 The Projection Algorithm, 269
3.5.4 Linearized Algorithms, 273
3.5.5 The Cartesian Product Case: Parallel Implementations, 275
3.5.6 Nonlinear Algorithms, 278
3.5.7 Decomposition Methods for Variational Inequalities, 281

4 SHORTEST PATHS AND DYNAMIC PROGRAMMING **291**

4.1 The Shortest Path Problem 293

4.1.1 The Bellman–Ford Algorithm, 294
4.1.2 Other Parallel Shortest Path Methods, 302

4.2 Markov Chains with Transition Costs 308

4.3 Markovian Decision Problems 312

4.3.1 Discounted Problems, 316
4.3.2 Undiscounted Problems—Stochastic Shortest Paths, 317
4.3.3 Parallel Implementation of the Dynamic Programming Iteration, 323

5 NETWORK FLOW PROBLEMS **331**

5.1 The Linear Network Flow Problem and its Dual 332

5.2 The Relaxation Method 340

5.2.1 Application to the Shortest Path Problem, 343
5.2.2 Multiple Node Relaxation Method, 345

5.3 The ϵ – Relaxation Method 355

5.3.1 The Auction Algorithm for the Assignment Problem, 364
5.3.2 Parallel Versions of the ϵ–Relaxation and the Auction Algorithms, 371

5.4 Complexity Analysis of the ϵ–Relaxation Method and its Scaled Version 376

5.4.1 The Scaled Version of the Algorithm, 384
5.4.2 Application to the Assignment Problem, 386

5.5 Network Flow Problems with Strictly Convex Cost 390

5.5.1 The Relaxation Method, 397
5.5.2 Convergence Analysis, 398
5.5.3 The Problem without Arc Flow Bounds, 406
5.5.4 An Example: Constrained Matrix Problems, 408
5.5.5 Parallel Implementations of the Relaxation Method, 410

5.6 Nonlinear Multicommodity Flow Problems–Routing Applications 414

Part 2 Asynchronous Algorithms 425

6 TOTALLY ASYNCHRONOUS ITERATIVE ALGORITHMS 425

6.1 The Totally Asynchronous Algorithmic Model 426

6.2 A General Convergence Theorem 431

6.3 Applications to Problems Involving Maximum Norm Contraction Mappings 434

6.3.1 Solution of Linear Systems of Equations, 434
6.3.2 Unconstrained Optimization, 437
6.3.3 Constrained Optimization and Variational Inequalities, 440
6.3.4 Dynamic Programming, 440
6.3.5 Convergence Rate Comparison of Synchronous and Asynchronous Algorithms, 441

6.4 Applications to Monotone Mappings and the Shortest Path Problem 445

6.5 Linear Network Flow Problems 451

6.6 Nonlinear Network Flow Problems 457

6.7 Asynchronous Relaxation for Ordinary Differential Equations and Two-Point Boundary Value Problems 469

6.7.1 The Asynchronous Relaxation Algorithm, 470
6.7.2 Two–Point Boundary Value Problems, 477
6.7.3 The Discrete Time Case, 478

7 PARTIALLY ASYNCHRONOUS ITERATIVE METHODS 481

7.1 The Partially Asynchronous Algorithmic Model 483

7.2 Algorithms for Fixed Points of Nonexpansive Mappings 490

7.2.1 A Convergence Result, 490
7.2.2 Weakly Diagonally Dominant Systems of Linear Equations, 498
7.2.3 Strictly Convex Network Flow Problems, 501

7.3 Algorithms for Agreement and for Markov Chain Problems 508

7.3.1 The Agreement Algorithm, 508
7.3.2 An Asynchronous Algorithm for the Invariant Distribution of a Markov chain, 515

7.4 Load Balancing in a Computer Network 519

7.5 Gradient–Like Optimization Algorithms 527

7.5.1 The Algorithm and its Convergence, 527
7.5.2 The Role of the Various Parameters, 532
7.5.3 Block–Iterative Algorithms, 533
7.5.4 Gradient Projection Algorithms, 534

7.6 Distributed Asynchronous Routing in Data Networks 536

7.6.1 Problem Definition, 536
7.6.2 The Algorithm and its Convergence, 538
7.6.3 Discussion, 546

7.7 A Model in Which Several Processors Update the Same Variables 550

7.8 Stochastic Gradient Algorithms 556

7.8.1 Description of the Algorithm and Assumptions, 558
7.8.2 A Convergence Result, 559
7.8.3 Discussion and Extensions, 567

8 ORGANIZING AN ASYNCHRONOUS NETWORK OF PROCESSORS FOR DISTRIBUTED COMPUTATION 570

8.1 Detecting Termination of a Distributed Algorithm 571

8.2 Snapshots 579

8.3 Resource Scheduling 587

8.4 Synchronization Using Rollback: Asynchronous Simulation 592

8.5 Maintaining Communication with a Center 605

A LINEAR ALGEBRA AND ANALYSIS 619

B GRAPH THEORY 653

C DUALITY THEORY 659

D PROBABILITY THEORY AND MARKOV CHAINS 670

REFERENCES *680*

INDEX *707*