

1. Record Nr.	UNINA9910141439003321
Autore	Zomaya Albert Y
Titolo	Energy efficient distributed computing systems / / edited by Albert Y. Zomaya, Young Choon Lee
Pubbl/distr/stampa	Hoboken, New Jersey : , : Wiley, , c2012 [Piscataqay, New Jersey] : , : IEEE Xplore, , [2012]
ISBN	1-283-54601-9 9786613858467 1-118-34201-1 1-118-34200-3 1-118-34198-8
Descrizione fisica	1 online resource (855 p.)
Collana	Wiley series on parallel and distributed computing ; ; 88 Wiley series on parallel and distributed computing
Classificazione	COM043000
Altri autori (Persone)	ZomayaAlbert Y LeeYoung-Choon <1973->
Disciplina	004/.36
Soggetti	Computer networks - Energy conservation Electronic data processing - Distributed processing - Energy conservation Green technology
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Note generali	Description based upon print version of record.
Nota di bibliografia	Includes bibliographical references and index.
Nota di contenuto	PREFACE xxix -- ACKNOWLEDGMENTS xxxi -- CONTRIBUTORS xxxiii -- 1 POWER ALLOCATION AND TASK SCHEDULING ON MULTIPROCESSOR COMPUTERS WITH ENERGY AND TIME CONSTRAINTS 1 / Keqin Li -- 1.1 Introduction 1 -- 1.1.1 Energy Consumption 1 -- 1.1.2 Power Reduction 2 -- 1.1.3 Dynamic Power Management 3 -- 1.1.4 Task Scheduling with Energy and Time Constraints 4 -- 1.1.5 Chapter Outline 5 -- 1.2 Preliminaries 5 -- 1.2.1 Power Consumption Model 5 -- 1.2.2 Problem Definitions 6 -- 1.2.3 Task Models 7 -- 1.2.4 Processor Models 8 -- 1.2.5 Scheduling Models 9 -- 1.2.6 Problem Decomposition 9 -- 1.2.7 Types of Algorithms 10 -- 1.3 Problem Analysis 10 -- 1.3.1 Schedule Length Minimization 10 -- 1.3.1.1 Uniprocessor computers 10 -- 1.3.1.2 Multiprocessor computers 11 -- 1.3.2 Energy Consumption Minimization 12 -- 1.3.2.1

Uniprocessor computers 12 -- 1.3.2.2 Multiprocessor computers 13 --  
1.3.3 Strong NP-Hardness 14 -- 1.3.4 Lower Bounds 14 -- 1.3.5  
Energy-Delay Trade-off 15 -- 1.4 Pre-Power-Determination  
Algorithms 16 -- 1.4.1 Overview 16 -- 1.4.2 Performance Measures 17  
-- 1.4.3 Equal-Time Algorithms and Analysis 18 -- 1.4.3.1 Schedule  
length minimization 18 -- 1.4.3.2 Energy consumption minimization  
19 -- 1.4.4 Equal-Energy Algorithms and Analysis 19 -- 1.4.4.1  
Schedule length minimization 19 -- 1.4.4.2 Energy consumption  
minimization 21 -- 1.4.5 Equal-Speed Algorithms and Analysis 22 --  
1.4.5.1 Schedule length minimization 22 -- 1.4.5.2 Energy  
consumption minimization 23 -- 1.4.6 Numerical Data 24 -- 1.4.7  
Simulation Results 25 -- 1.5 Post-Power-Determination Algorithms 28  
-- 1.5.1 Overview 28 -- 1.5.2 Analysis of List Scheduling Algorithms  
29 -- 1.5.2.1 Analysis of algorithm LS 29 -- 1.5.2.2 Analysis of  
algorithm LRF 30 -- 1.5.3 Application to Schedule Length Minimization  
30 -- 1.5.4 Application to Energy Consumption Minimization 31 --  
1.5.5 Numerical Data 32 -- 1.5.6 Simulation Results 32 -- 1.6  
Summary and Further Research 33 -- References 34 -- 2 POWER-  
AWARE HIGH PERFORMANCE COMPUTING 39 / Rong Ge and Kirk W.  
Cameron.  
2.1 Introduction 39 -- 2.2 Background 41 -- 2.2.1 Current Hardware  
Technology and Power Consumption 41 -- 2.2.1.1 Processor power 41  
-- 2.2.1.2 Memory subsystem power 42 -- 2.2.2 Performance 43 --  
2.2.3 Energy Efficiency 44 -- 2.3 Related Work 45 -- 2.3.1 Power  
Profiling 45 -- 2.3.1.1 Simulator-based power estimation 45 -- 2.3.1.2  
Direct measurements 46 -- 2.3.1.3 Event-based estimation 46 -- 2.3.2  
Performance Scalability on Power-Aware Systems 46 -- 2.3.3 Adaptive  
Power Allocation for Energy-Efficient Computing 47 -- 2.4 PowerPack:  
Fine-Grain Energy Profiling of HPC Applications 48 -- 2.4.1 Design and  
Implementation of PowerPack 48 -- 2.4.1.1 Overview 48 -- 2.4.1.2  
Fine-grain systematic power measurement 50 -- 2.4.1.3 Automatic  
power profiling and code synchronization 51 -- 2.4.2 Power Profiles of  
HPC Applications and Systems 53 -- 2.4.2.1 Power distribution over  
components 53 -- 2.4.2.2 Power dynamics of applications 54 --  
2.4.2.3 Power bounds on HPC systems 55 -- 2.4.2.4 Power versus  
dynamic voltage and frequency scaling 57 -- 2.5 Power-Aware Speedup  
Model 59 -- 2.5.1 Power-Aware Speedup 59 -- 2.5.1.1 Sequential  
execution time for a single workload  $T_1(w, f)$  60 -- 2.5.1.2 Sequential  
execution time for an ON-chip/OFF-chip workload 60 -- 2.5.1.3  
Parallel execution time on N processors for an ON-/OFF-chip workload  
with  $DOP = i$  61 -- 2.5.1.4 Power-aware speedup for DOP and ON-  
/OFF-chip workloads 62 -- 2.5.2 Model Parametrization and Validation  
63 -- 2.5.2.1 Coarse-grain parametrization and validation 64 --  
2.5.2.2 Fine-grain parametrization and validation 66 -- 2.6 Model  
Usages 69 -- 2.6.1 Identification of Optimal System Configurations 70  
-- 2.6.2 PAS-Directed Energy-Driven Runtime Frequency Scaling 71 --  
2.7 Conclusion 73 -- References 75 -- 3 ENERGY EFFICIENCY IN HPC  
SYSTEMS 81 / Ivan Rodero and Manish Parashar -- 3.1 Introduction 81  
-- 3.2 Background and Related Work 83 -- 3.2.1 CPU Power  
Management 83 -- 3.2.1.1 OS-level CPU power management 83 --  
3.2.1.2 Workload-level CPU power management 84.  
3.2.1.3 Cluster-level CPU power management 84 -- 3.2.2 Component-  
Based Power Management 85 -- 3.2.2.1 Memory subsystem 85 --  
3.2.2.2 Storage subsystem 86 -- 3.2.3 Thermal-Conscious Power  
Management 87 -- 3.2.4 Power Management in Virtualized Datacenters  
87 -- 3.3 Proactive, Component-Based Power Management 88 -- 3.3.1  
Job Allocation Policies 88 -- 3.3.2 Workload Profiling 90 -- 3.4  
Quantifying Energy Saving Possibilities 91 -- 3.4.1 Methodology 92 --

3.4.2 Component-Level Power Requirements 92 -- 3.4.3 Energy Savings 94 -- 3.5 Evaluation of the Proposed Strategies 95 -- 3.5.1 Methodology 96 -- 3.5.2 Workloads 96 -- 3.5.3 Metrics 97 -- 3.6 Results 97 -- 3.7 Concluding Remarks 102 -- 3.8 Summary 103 -- References 104 -- 4 A STOCHASTIC FRAMEWORK FOR HIERARCHICAL SYSTEM-LEVEL POWERMANAGEMENT 109 / Peng Rong and Massoud Pedram -- 4.1 Introduction 109 -- 4.2 Related Work 111 -- 4.3 A Hierarchical DPM Architecture 113 -- 4.4 Modeling 114 -- 4.4.1 Model of the Application Pool 114 -- 4.4.2 Model of the Service Flow Control 118 -- 4.4.3 Model of the Simulated Service Provider 119 -- 4.4.4 Modeling Dependencies between SPs 120 -- 4.5 Policy Optimization 122 -- 4.5.1 Mathematical Formulation 122 -- 4.5.2 Optimal Time-Out Policy for Local Power Manager 123 -- 4.6 Experimental Results 125 -- 4.7 Conclusion 130 -- References 130 -- 5 ENERGY-EFFICIENT RESERVATION INFRASTRUCTURE FOR GRIDS,CLOUDS, AND NETWORKS 133 / Anne-Ce'cile Orgerie and Laurent Lefevre -- 5.1 Introduction 133 -- 5.2 Related Works 134 -- 5.2.1 Server and Data Center Power Management 135 -- 5.2.2 Node Optimizations 135 -- 5.2.3 Virtualization to Improve Energy Efficiency 136 -- 5.2.4 Energy Awareness in Wired Networking Equipment 136 -- 5.2.5 Synthesis 137 -- 5.3 ERIDIS: Energy-Efficient Reservation Infrastructure for Large-Scale Distributed Systems 138 -- 5.3.1 ERIDIS Architecture 138 -- 5.3.2 Management of the Resource Reservations 141 -- 5.3.3 Resource Management and On/Off Algorithms 145 -- 5.3.4 Energy-Consumption Estimates 146. 5.3.5 Prediction Algorithms 146 -- 5.4 EARl: Energy-Aware Reservation Infrastructure for DataCenters and Grids 147 -- 5.4.1 EARl's Architecture 147 -- 5.4.2 Validation of EARl on Experimental Grid Traces 147 -- 5.5 GOC: Green Open Cloud 149 -- 5.5.1 GOC's Resource Manager Architecture 150 -- 5.5.2 Validation of the GOC Framework 152 -- 5.6 HERMES: High Level Energy-Aware Model for BandwidthReservation in End-To-End Networks 152 -- 5.6.1 HERMES' Architecture 154 -- 5.6.2 The Reservation Process of HERMES 155 -- 5.6.3 Discussion 157 -- 5.7 Summary 158 -- References 158 -- 6 ENERGY-EFFICIENT JOB PLACEMENT ON CLUSTERS, GRIDS, AND CLOUDS 163 / Damien Borgetto, Henri Casanova, Georges Da Costa, and Jean-Marc Pierson -- 6.1 Problem and Motivation 163 -- 6.1.1 Context 163 -- 6.1.2 Chapter Roadmap 164 -- 6.2 Energy-Aware Infrastructures 164 -- 6.2.1 Buildings 165 -- 6.2.2 Context-Aware Buildings 165 -- 6.2.3 Cooling 166 -- 6.3 Current Resource Management Practices 167 -- 6.3.1 Widely Used Resource Management Systems 167 -- 6.3.2 Job Requirement Description 169 -- 6.4 Scientific and Technical Challenges 170 -- 6.4.1 Theoretical Difficulties 170 -- 6.4.2 Technical Difficulties 170 -- 6.4.3 Controlling and Tuning Jobs 171 -- 6.5 Energy-Aware Job Placement Algorithms 172 -- 6.5.1 State of the Art 172 -- 6.5.2 Detailing One Approach 174 -- 6.6 Discussion 180 -- 6.6.1 Open Issues and Opportunities 180 -- 6.6.2 Obstacles for Adoption in Production 182 -- 6.7 Conclusion 183 -- References 184 -- 7 COMPARISON AND ANALYSIS OF GREEDY ENERGY-EFFICIENT SCHEDULING ALGORITHMS FOR COMPUTATIONAL GRIDS 189 / Peder Lindberg, James Leingang, Daniel Lysaker, Kashif Bilal, Samee Ullah Khan, Pascal Bouvry, Nasir Ghani, NasroMin-Allah, and Juan Li -- 7.1 Introduction 189 -- 7.2 Problem Formulation 191 -- 7.2.1 The System Model 191 -- 7.2.1.1 PEs 191 -- 7.2.1.2 DVS 191 -- 7.2.1.3 Tasks 192 -- 7.2.1.4 Preliminaries 192 -- 7.2.2 Formulating the Energy-Makespan Minimization Problem 192 -- 7.3 Proposed Algorithms 193. 7.3.1 Greedy Heuristics 194 -- 7.3.1.1 Greedy heuristic scheduling

algorithm 196 -- 7.3.1.2 Greedy-min 197 -- 7.3.1.3 Greedy-deadline 198 -- 7.3.1.4 Greedy-max 198 -- 7.3.1.5 MaxMin 199 -- 7.3.1.6 ObFun 199 -- 7.3.1.7 MinMin StdDev 202 -- 7.3.1.8 MinMax StdDev 202 -- 7.4 Simulations, Results, and Discussion 203 -- 7.4.1 Workload 203 -- 7.4.2 Comparative Results 204 -- 7.4.2.1 Small-size problems 204 -- 7.4.2.2 Large-size problems 206 -- 7.5 Related Works 211 -- 7.6 Conclusion 211 -- References 212 -- 8 TOWARD ENERGY-AWARE SCHEDULING USING MACHINE LEARNING215 / Josep LL. Berral, Inigo Goiri, Ramon Nou, Ferran Julia` , Josep O. Fito' , Jordi Guitart, Ricard Gavalda', and Jordi Torres -- 8.1 Introduction 215 -- 8.1.1 Energetic Impact of the Cloud 216 -- 8.1.2 An Intelligent Way to Manage Data Centers 216 -- 8.1.3 Current Autonomic Computing Techniques 217 -- 8.1.4 Power-Aware Autonomic Computing 217 -- 8.1.5 State of the Art and Case Study 218 -- 8.2 Intelligent Self-Management 218 -- 8.2.1 Classical AI Approaches 219 -- 8.2.1.1 Heuristic algorithms 219 -- 8.2.1.2 AI planning 219 -- 8.2.1.3 Semantic techniques 219 -- 8.2.1.4 Expert systems and genetic algorithms 220 -- 8.2.2 Machine Learning Approaches 220 -- 8.2.2.1 Instance-based learning 221 -- 8.2.2.2 Reinforcement learning 222 -- 8.2.2.3 Feature and example selection 225 -- 8.3 Introducing Power-Aware Approaches 225 -- 8.3.1 Use of Virtualization 226 -- 8.3.2 Turning On and Off Machines 228 -- 8.3.3 Dynamic Voltage and Frequency Scaling 229 -- 8.3.4 Hybrid Nodes and Data Centers 230 -- 8.4 Experiences of Applying ML on Power-Aware Self-Management230 -- 8.4.1 Case Study Approach 231 -- 8.4.2 Scheduling and Power Trade-Off 231 -- 8.4.3 Experimenting with Power-Aware Techniques 233 -- 8.4.4 Applying Machine Learning 236 -- 8.4.5 Conclusions from the Experiments 238 -- 8.5 Conclusions on Intelligent Power-Aware Self-Management238 -- References 240 -- 9 ENERGY EFFICIENCY METRICS FOR DATA CENTERS 245 / Javid Taheri and Albert Y. Zomaya.  
9.1 Introduction 245 -- 9.1.1 Background 245 -- 9.1.2 Data Center Energy Use 246 -- 9.1.3 Data Center Characteristics 246 -- 9.1.3.1 Electric power 247 -- 9.1.3.2 Heat removal 249 -- 9.1.4 Energy Efficiency 250 -- 9.2 Fundamentals of Metrics 250 -- 9.2.1 Demand and Constraints on Data Center Operators 250 -- 9.2.2 Metrics 251 -- 9.2.2.1 Criteria for good metrics 251 -- 9.2.2.2 Methodology 252 -- 9.2.2.3 Stability of metrics 252 -- 9.3 Data Center Energy Efficiency 252 -- 9.3.1 Holistic IT Efficiency Metrics 252 -- 9.3.1.1 Fixed versus proportional overheads 254 -- 9.3.1.2 Power versus energy 254 -- 9.3.1.3 Performance versus productivity 255 -- 9.3.2 Code of Conduct 256 -- 9.3.2.1 Environmental statement 256 -- 9.3.2.2 Problem statement 256 -- 9.3.2.3 Scope of the CoC 257 -- 9.3.2.4 Aims and objectives of CoC 258 -- 9.3.3 Power Use in Data Centers 259 -- 9.3.3.1 Data center IT power to utility power relationship259 -- 9.3.3.2 Chiller efficiency and external temperature 260 -- 9.4 Available Metrics 260 -- 9.4.1 The Green Grid 261 -- 9.4.1.1 Power usage effectiveness (PUE) 261 -- 9.4.1.2 Data center efficiency (DCE) 262 -- 9.4.1.3 Data center infrastructure efficiency (DCiE) 262 -- 9.4.1.4 Data center productivity (DCP) 263 -- 9.4.2 McKinsey 263 -- 9.4.3 Uptime Institute 264 -- 9.4.3.1 Site infrastructure power overhead multiplier (SI-POM) 265 -- 9.4.3.2 IT hardware power overhead multiplier (H-POM) 266 -- 9.4.3.3 DC hardware compute load per unit of computing work done266 -- 9.4.3.4 Deployed hardware utilization ratio (DH-UR) 266 -- 9.4.3.5 Deployed hardware utilization efficiency (DH-UE) 267 -- 9.5 Harmonizing Global Metrics for Data Center Energy Efficiency267 -- References 268 -- 10 AUTONOMIC GREEN COMPUTING IN LARGE-SCALE DATA CENTERS271 / Haoting Luo, Bithika Khargharia, Salim Hariri, and Youssif Al-Nashif / -- 10.1 Introduction 271 -- 10.2 Related

Technologies and Techniques 272 -- 10.2.1 Power Optimization Techniques in Data Centers 272 -- 10.2.2 Design Model 273. 10.2.3 Networks 274 -- 10.2.4 Data Center Power Distribution 275 -- 10.2.5 Data Center Power-Efficient Metrics 276 -- 10.2.6 Modeling Prototype and Testbed 277 -- 10.2.7 Green Computing 278 -- 10.2.8 Energy Proportional Computing 280 -- 10.2.9 Hardware Virtualization Technology 281 -- 10.2.10 Autonomic Computing 282 -- 10.3 Autonomic Green Computing: A Case Study 283 -- 10.3.1 Autonomic Management Platform 285 -- 10.3.1.1 Platform architecture 285 -- 10.3.1.2 DEVS-based modeling and simulation platform 285 -- 10.3.1.3 Workload generator 287 -- 10.3.2 Model Parameter Evaluation 288 -- 10.3.2.1 State transitioning overhead 288 -- 10.3.2.2 VM template evaluation 289 -- 10.3.2.3 Scalability analysis 291 -- 10.3.3 Autonomic Power Efficiency Management Algorithm(Performance Per Watt) 291 -- 10.3.4 Simulation Results and Evaluation 293 -- 10.3.4.1 Analysis of energy and performance trade-offs 296 -- 10.4 Conclusion and Future Directions 297 -- References 298 -- 11 ENERGY AND THERMAL AWARE SCHEDULING IN DATA CENTERS 301 / Gaurav Dhiman, Raid Ayoub, and Tajana S. Rosing -- 11.1 Introduction 301 -- 11.2 Related Work 302 -- 11.3 Intermachine Scheduling 305 -- 11.3.1 Performance and Power Profile of VMs 305 -- 11.3.2 Architecture 309 -- 11.3.2.1 vgnode 309 -- 11.3.2.2 vgxe 310 -- 11.3.2.3 vgdom 312 -- 11.3.2.4 vgserve 312 -- 11.4 Intramachine Scheduling 315 -- 11.4.1 Air-Forced Thermal Modeling and Cost 316 -- 11.4.2 Cooling Aware Dynamic Workload Scheduling 317 -- 11.4.3 Scheduling Mechanism 318 -- 11.4.4 Cooling Costs Predictor 319 -- 11.5 Evaluation 321 -- 11.5.1 Intermachine Scheduler (vGreen) 321 -- 11.5.2 Heterogeneous Workloads 323 -- 11.5.2.1 Comparison with DVFS policies 325 -- 11.5.2.2 Homogeneous workloads 328 -- 11.5.3 Intramachine Scheduler (Cool and Save) 328 -- 11.5.3.1 Results 331 -- 11.5.3.2 Overhead of CAS 333 -- 11.6 Conclusion 333 -- References 334 -- 12 QOS-AWARE POWER MANAGEMENT IN DATA CENTERS 339 / Jiayu Gong and Cheng-Zhong Xu -- 12.1 Introduction 339. 12.2 Problem Classification 340 -- 12.2.1 Objective and Constraint 340 -- 12.2.2 Scope and Time Granularities 340 -- 12.2.3 Methodology 341 -- 12.2.4 Power Management Mechanism 342 -- 12.3 Energy Efficiency 344 -- 12.3.1 Energy-Efficiency Metrics 344 -- 12.3.2 Improving Energy Efficiency 346 -- 12.3.2.1 Energy minimization with performance guarantee 346 -- 12.3.2.2 Performance maximization under power budget 348 -- 12.3.2.3 Trade-off between power and performance 348 -- 12.3.3 Energy-Proportional Computing 350 -- 12.4 Power Capping 351 -- 12.5 Conclusion 353 -- References 356 -- 13 ENERGY-EFFICIENT STORAGE SYSTEMS FOR DATA CENTERS 361 / Sudhanva Gurumurthi and Anand Sivasubramaniam -- 13.1 Introduction 361 -- 13.2 Disk Drive Operation and Disk Power 362 -- 13.2.1 An Overview of Disk Drives 362 -- 13.2.2 Sources of Disk Power Consumption 363 -- 13.2.3 Disk Activity and Power Consumption 365 -- 13.3 Disk and Storage Power Reduction Techniques 366 -- 13.3.1 Exploiting the STANDBY State 368 -- 13.3.2 Reducing Seek Activity 369 -- 13.3.3 Achieving Energy Proportionality 369 -- 13.3.3.1 Hardware approaches 369 -- 13.3.3.2 Software approaches 370 -- 13.4 Using Nonvolatile Memory and Solid-State Disks 371 -- 13.5 Conclusions 372 -- References 373 -- 14 AUTONOMIC ENERGY/PERFORMANCE OPTIMIZATIONS FOR MEMORY INSERVERS 377 / Bithika Khargharia and Mazin Yousif -- 14.1 Introduction 378 -- 14.2 Classifications of Dynamic Power Management Techniques 380 -- 14.2.1 Heuristic and Predictive Techniques 380 -- 14.2.2 QoS and Energy Trade-Offs 381 -- 14.3 Applications of Dynamic Power Management (DPM) 382 --

14.3.1 Power Management of System Components in Isolation 382 --  
14.3.2 Joint Power Management of System Components 383 -- 14.3.3 Holistic System-Level Power Management 383 -- 14.4 Autonomic Power and Performance Optimization of Memory Subsystems in Server Platforms 384 -- 14.4.1 Adaptive Memory Interleaving Technique for Power and Performance Management 384 -- 14.4.1.1 Formulating the optimization problem 386.  
14.4.1.2 Memory appflow 389 -- 14.4.2 Industry Techniques 389 -- 14.4.2.1 Enhancements in memory hardware design 390 -- 14.4.2.2 Adding more operating states 390 -- 14.4.2.3 Faster transition to and from low power states 390 -- 14.4.2.4 Memory consolidation 390 -- 14.5 Conclusion 391 -- References 391 -- 15 ROD: A PRACTICAL APPROACH TO IMPROVING RELIABILITY OF ENERGY-EFFICIENT PARALLEL DISK SYSTEMS 395 / Shu Yin, Xiaojun Ruan, Adam Manzanares, and XiaoQin -- 15.1 Introduction 395 -- 15.2 Modeling Reliability of Energy-Efficient Parallel Disks 396 -- 15.2.1 The MINT Model 396 -- 15.2.1.1 Disk utilization 398 -- 15.2.1.2 Temperature 398 -- 15.2.1.3 Power-state transition frequency 399 -- 15.2.1.4 Single disk reliability model 399 -- 15.2.2 MAID, Massive Arrays of Idle Disks 400 -- 15.3 Improving Reliability of MAID via Disk Swapping 401 -- 15.3.1 Improving Reliability of Cache Disks in MAID 401 -- 15.3.2 Swapping Disks Multiple Times 404 -- 15.4 Experimental Results and Evaluation 405 -- 15.4.1 Experimental Setup 405 -- 15.4.2 Disk Utilization 406 -- 15.4.3 The Single Disk Swapping Strategy 406 -- 15.4.4 The Multiple Disk Swapping Strategy 409 -- 15.5 Related Work 411 -- 15.6 Conclusions 412 -- References 413 -- 16 EMBRACING THE MEMORY AND I/O WALLS FOR ENERGY-EFFICIENT SCIENTIFIC COMPUTING 417 / Chung-Hsing Hsu and Wu-Chun Feng -- 16.1 Introduction 417 -- 16.2 Background and Related Work 420 -- 16.2.1 DVFS-Enabled Processors 420 -- 16.2.2 DVFS Scheduling Algorithms 421 -- 16.2.3 Memory-Aware, Interval-Based Algorithms 422 -- 16.3 -Adaptation: A New DVFS Algorithm 423 -- 16.3.1 The Compute-Boundedness Metric, 423 -- 16.3.2 The Frequency Calculating Formula,  $f \geq \frac{C}{P}$ ; 424 -- 16.3.3 The Online Estimation 425 -- 16.3.4 Putting It All Together 427 -- 16.4 Algorithm Effectiveness 429 -- 16.4.1 A Comparison to Other DVFS Algorithms 429 -- 16.4.2 Frequency Emulation 432 -- 16.4.3 The Minimum Dependence to the PMU 436 -- 16.5 Conclusions and Future Work 438 -- References 439.  
17 MULTIPLE FREQUENCY SELECTION IN DVFS-ENABLED PROCESSORS TO MINIMIZE ENERGY CONSUMPTION 443 / Nikzad Babaii Rizvandi, Albert Y. Zomaya, Young Choon Lee, Ali Javadzadeh Boloori, and Javid Taheri -- 17.1 Introduction 443 -- 17.2 Energy Efficiency in HPC Systems 444 -- 17.3 Exploitation of Dynamic Voltage-Frequency Scaling 446 -- 17.3.1 Independent Slack Reclamation 446 -- 17.3.2 Integrated Schedule Generation 447 -- 17.4 Preliminaries 448 -- 17.4.1 System and Application Models 448 -- 17.4.2 Energy Model 448 -- 17.5 Energy-Aware Scheduling via DVFS 450 -- 17.5.1 Optimum Continuous Frequency 450 -- 17.5.2 Reference Dynamic Voltage-Frequency Scaling (RDVFS) 451 -- 17.5.3 Maximum-Minimum-Frequency for Dynamic Voltage-Frequency Scaling (MMF-DVFS) 452 -- 17.5.4 Multiple Frequency Selection for Dynamic Voltage-Frequency Scaling (MFS-DVFS) 453 -- 17.5.4.1 Task eligibility 454 -- 17.6 Experimental Results 456 -- 17.6.1 Simulation Settings 456 -- 17.6.2 Results 458 -- 17.7 Conclusion 461 -- References 461 -- 18 THE PARAMOUNTCY OF RECONFIGURABLE COMPUTING 465 / Reiner Hartenstein -- 18.1 Introduction 465 -- 18.2 Why Computers are Important 466 -- 18.2.1 Computing for a Sustainable Environment 470 -- 18.3 Performance Progress Stalled 472 -- 18.3.1 Unaffordable

Energy Consumption of Computing 473 -- 18.3.2 Crashing into the Programming Wall 475 -- 18.4 The Tail is Wagging the Dog (Accelerators) 488 -- 18.4.1 Hardwired Accelerators 489 -- 18.4.2 Programmable Accelerators 490 -- 18.5 Reconfigurable Computing 494 -- 18.5.1 Speedup Factors by FPGAs 498 -- 18.5.2 The Reconfigurable Computing Paradox 501 -- 18.5.3 Saving Energy by Reconfigurable Computing 505 -- 18.5.3.1 Traditional green computing 506 -- 18.5.3.2 The role of graphics processors 507 -- 18.5.3.3 Wintel versus ARM 508 -- 18.5.4 Reconfigurable Computing is the Silver Bullet 511 -- 18.5.4.1 A new world model of computing 511 -- 18.5.5 The Twin-Paradigm Approach to Tear Down the Wall 514 -- 18.5.6 A Mass Movement Needed as Soon as Possible 517. 18.5.6.1 Legacy software from the mainframe age 518 -- 18.5.7 How to Reinvent Computing 519 18.6 Conclusions 526 -- References 529 -- 19 WORKLOAD CLUSTERING FOR INCREASING ENERGY SAVINGS ONEMBEDDED MPSOCS 549 / Ozcan Ozturk, Mahmut Kandemir, and Sri Hari Krishna Narayanan -- 19.1 Introduction 549 -- 19.2 Embedded MPSoC Architecture, Execution Model, and RelatedWork 550 -- 19.3 Our Approach 551 -- 19.3.1 Overview 551 -- 19.3.2 Technical Details and Problem Formulation 553 -- 19.3.2.1 System and job model 553 -- 19.3.2.2 Mathematical programing model 554 -- 19.3.2.3 Example 557 -- 19.4 Experimental Evaluation 560 -- 19.5 Conclusions 564 -- References 565 -- 20 ENERGY-EFFICIENT INTERNET INFRASTRUCTURE 567 / Weirong Jiang and Viktor K. Prasanna -- 20.1 Introduction 567 -- 20.1.1 Performance Challenges 568 -- 20.1.2 Existing Packet Forwarding Approaches 570 -- 20.1.2.1 Software approaches 570 -- 20.1.2.2 Hardware approaches 571 -- 20.2 SRAM-Based Pipelined IP Lookup Architectures: Alternativeto TCAMs 571 -- 20.3 Data Structure Optimization for Power Efficiency 573 -- 20.3.1 Problem Formulation 574 -- 20.3.1.1 Non-pipelined and pipelined engines 574 -- 20.3.1.2 Power function of SRAM 575 -- 20.3.2 Special Case: Uniform Stride 576 -- 20.3.3 Dynamic Programming 576 -- 20.3.4 Performance Evaluation 577 -- 20.3.4.1 Results for non-pipelined architecture 578 -- 20.3.4.2 Results for pipelined architecture 578 -- 20.4 Architectural Optimization to Reduce Dynamic PowerDissipation 580 -- 20.4.1 Analysis and Motivation 581 -- 20.4.1.1 Traffic locality 582 -- 20.4.1.2 Traffic rate variation 582 -- 20.4.1.3 Access frequency on different stages 583 -- 20.4.2 Architecture-Specific Techniques 583 -- 20.4.2.1 Inherent caching 584 -- 20.4.2.2 Local clocking 584 -- 20.4.2.3 Fine-grained memory enabling 585 -- 20.4.3 Performance Evaluation 585 -- 20.5 Related Work 588 -- 20.6 Summary 589 -- References 589 -- 21 DEMAND RESPONSE IN THE SMART GRID: A DISTRIBUTED COMPUTING PERSPECTIVE 593 / Chen Wang and Martin De Groot. 21.1 Introduction 593 -- 21.2 Demand Response 595 -- 21.2.1 Existing Demand Response Programs 595 -- 21.2.2 Demand Response Supported by the Smart Grid 597 -- 21.3 Demand Response as a Distributed System 600 -- 21.3.1 An Overlay Network for Demand Response 600 -- 21.3.2 Event Driven Demand Response 602 -- 21.3.3 Cost Driven Demand Response 604 -- 21.3.4 A Decentralized Demand Response Framework 609 -- 21.3.5 Accountability of Coordination Decision Making 610 -- 21.4 Summary 611 -- References 611 -- 22 RESOURCE MANAGEMENT FOR DISTRIBUTED MOBILE COMPUTING615 / Jong-Kook Kim -- 22.1 Introduction 615 -- 22.2 Single-Hop Energy-Constrained Environment 617 -- 22.2.1 System Model 617 -- 22.2.2 Related Work 620 -- 22.2.3 Heuristic Descriptions 621 -- 22.2.3.1 Mapping event 621 -- 22.2.3.2 Scheduling communications 621 -- 22.2.3.3 Opportunistic load balancing and minimum energy greedyheuristics 622 -- 22.2.3.4 ME-MC heuristic 622 -- 22.2.3.5 ME-

ME heuristic 624 -- 22.2.3.6 CRME heuristic 625 -- 22.2.3.7 Originator and random 626 -- 22.2.3.8 Upper bound 626 -- 22.2.4 Simulation Model 628 -- 22.2.5 Results 630 -- 22.2.6 Summary 634 -- 22.3 Multihop Distributed Mobile Computing Environment 635 -- 22.3.1 The Multihop System Model 635 -- 22.3.2 Energy-Aware Routing Protocol 636 -- 22.3.2.1 Overview 636 -- 22.3.2.2 DSDV 637 -- 22.3.2.3 DSDV remaining energy 637 -- 22.3.2.4 DSDV-energy consumption per remaining energy 637 -- 22.3.3 Heuristic Description 638 -- 22.3.3.1 Random 638 -- 22.3.3.2 Estimated minimum total energy (EMTE) 638 -- 22.3.3.3 K-percent-speed (KPS) and K-percent-energy (KPE) 639 -- 22.3.3.4 Energy ratio and distance (ERD) 639 -- 22.3.3.5 ETC and distance (ETCD) 640 -- 22.3.3.6 Minimum execution time (MET) 640 -- 22.3.3.7 Minimum completion time (MCT) and minimum completion time with DVS (MCT-DVS) 640 -- 22.3.3.8 Switching algorithm (SA) 640 -- 22.3.4 Simulation Model 641 -- 22.3.5 Results 643 -- 22.3.5.1 Distributed resource management 643 -- 22.3.5.2 Energy-aware protocol 644.  
22.3.6 Summary 644 -- 22.4 Future Work 647 -- References 647 -- 23 AN ENERGY-AWARE FRAMEWORK FOR MOBILE DATA MINING 653 / Carmela Comito, Domenico Talia, and Paolo Trunfio -- 23.1 Introduction 653 -- 23.2 System Architecture 654 -- 23.3 Mobile Device Components 657 -- 23.4 Energy Model 659 -- 23.5 Clustering Scheme 664 -- 23.5.1 Clustering the M2M Architecture 666 -- 23.6 Conclusion 670 -- References 670 -- 24 ENERGY AWARENESS AND EFFICIENCY IN WIRELESS SENORN NETWORKS: FROM PHYSICAL DEVICES TO THE COMMUNICATION LINK 673 / Fla' via C. Delicato and Paulo F. Pires -- 24.1 Introduction 673 -- 24.2 WSN and Power Dissipation Models 676 -- 24.2.1 Network and Node Architecture 676 -- 24.2.2 Sources of Power Dissipation in WSNs 679 -- 24.3 Strategies for Energy Optimization 683 -- 24.3.1 Intranode Level 684 -- 24.3.1.1 Duty cycling 685 -- 24.3.1.2 Adaptive sensing 691 -- 24.3.1.3 Dynamic voltage scale (DVS) 693 -- 24.3.1.4 OS task scheduling 694 -- 24.3.2 Internode Level 695 -- 24.3.2.1 Transmission power control 695 -- 24.3.2.2 Dynamic modulation scaling 696 -- 24.3.2.3 Link layer optimizations 698 -- 24.4 Final Remarks 701 -- References 702 -- 25 NETWORK-WIDE STRATEGIES FOR ENERGY EFFICIENCY IN WIRELESS SENSOR NETWORKS 709 / Fla' via C. Delicato and Paulo F. Pires -- 25.1 Introduction 709 -- 25.2 Data Link Layer 711 -- 25.2.1 Topology Control Protocols 712 -- 25.2.2 Energy-Efficient MAC Protocols 714 -- 25.2.2.1 Scheduled MAC protocols in WSNs 716 -- 25.2.2.2 Contention-based MAC protocols 717 -- 25.3 Network Layer 719 -- 25.3.1 Flat and Hierarchical Protocols 722 -- 25.4 Transport Layer 725 -- 25.5 Application Layer 729 -- 25.5.1 Task Scheduling 729 -- 25.5.2 Data Aggregation and Data Fusion in WSNs 733 -- 25.5.2.1 Approaches of data fusion for energy efficiency 735 -- 25.5.2.2 Data aggregation strategies 736 -- 25.6 Final Remarks 740 -- References 741 -- 26 ENERGY MANAGEMENT IN HETEROGENEOUS WIRELESS HEALTH CARE NETWORKS 751 / Nima Nikzad, Priti Aghera, Piero Zappi, and Tajana S.Rosing.  
26.1 Introduction 751 -- 26.2 System Model 753 -- 26.2.1 Health Monitoring Task Model 753 -- 26.3 Collaborative Distributed Environmental Sensing 755 -- 26.3.1 Node Neighborhood and Localization Rate 757 -- 26.3.2 Energy Ratio and Sensing Rate 758 -- 26.3.3 Duty Cycling and Prediction 759 -- 26.4 Task Assignment in a Body Area Network 760 -- 26.4.1 Optimal Task Assignment 760 -- 26.4.2 Dynamic Task Assignment 762 -- 26.4.2.1 DynAGreen algorithm 763 -- 26.4.2.2 DynAGreenLife algorithm 768 -- 26.5 Results 771 -- 26.5.1 Collaborative Sensing 771 -- 26.5.1.1 Results

Sommario/riassunto

"The energy consumption issue in distributed computing systems raises various monetary, environmental and system performance concerns. Electricity consumption in the US doubled from 2000 to 2005. From a financial and environmental standpoint, reducing the consumption of electricity is important, yet these reforms must not lead to performance degradation of the computing systems. These contradicting constraints create a suite of complex problems that need to be resolved in order to lead to 'greener' distributed computing systems. This book brings together a group of outstanding researchers that investigate the different facets of green and energy efficient distributed computing.Key features: One of the first books of its kind Features latest research findings on emerging topics by well-known scientists Valuable research for grad students, postdocs, and researchers Research will greatly feed into other technologies and application domains"--