

1. Record Nr.	UNINA9910484010103321
Titolo	Network and Parallel Computing : 13th IFIP WG 10.3 International Conference, NPC 2016, Xi'an, China, October 28-29, 2016, Proceedings // edited by Guang R. Gao, Depei Qian, Xinbo Gao, Barbara Chapman, Wenguang Chen
Pubbl/distr/stampa	Cham : , : Springer International Publishing : , : Imprint : Springer, , 2016
ISBN	3-319-47099-X
Edizione	[1st ed. 2016.]
Descrizione fisica	1 online resource (X, 213 p. 94 illus.)
Collana	Theoretical Computer Science and General Issues, , 2512-2029 ; ; 9966
Disciplina	004.6
Soggetti	Computer engineering Computer networks Software engineering Algorithms Application software Computer Engineering and Networks Software Engineering Computer and Information Systems Applications
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Nota di contenuto	Intro -- Preface -- Organization -- Contents -- Memory: Non-Volatile, Solid State Drives, Hybrid Systems -- VIOS: A Variation-Aware I/O Scheduler for Flash-Based Storage Systems -- 1 Introduction -- 2 Background and Related Work -- 2.1 SSD Organization -- 2.2 Process Variation of Flash Memory -- 2.3 I/O Scheduler for Flash-Based SSDs -- 3 Details of VIOS -- 3.1 Block Management -- 3.2 Global Chip-State Vector -- 3.3 Conflict Optimized Scheduling Mechanism -- 4 Experimental Results -- 4.1 Performance Analysis of VIOS -- 4.2 Sensitivity Analysis of VIOS -- 5 Conclusion -- References -- Exploiting Cross-Layer Hotness Identification to Improve Flash Memory System Performance -- 1 Introduction -- 2 Background and Related Work -- 2.1 Tradeoff Between Flash Cell Wearing and Read Latency -- 2.2 Related Work -- 3 Exploiting Cross-Layer Hotness Identification to

Improve Read and Endurance Performance (HIRE) -- 3.1 Cross-Layer Study for Hotness Identifier -- 3.2 Voltage Controller in HIRE -- 3.3 Overhead Analysis -- 4 Performance Evaluation -- 4.1 Methodology -- 4.2 Experiment Results -- 5 Conclusion -- References -- Efficient Management for Hybrid Memory in Managed Language Runtime -- 1 Introduction -- 2 Managed Runtime Background -- 3 Related Work -- 4 Hybrid Memory Management Scheme -- 4.1 Overview -- 4.2 HMprof Offline Performance Instrumentation -- 4.3 Hot Object Marking -- 5 Experimental Methodology -- 6 Evaluation Results -- 6.1 Heap Partition Placement -- 6.2 Hot Object Allocation -- 7 Conclusion -- References -- Resilience and Reliability -- Application-Based Coarse-Grained Incremental Checkpointing Based on Non-volatile Memory -- Abstract -- 1 Introduction -- 2 Background and Motivation -- 2.1 Non-volatile Memory -- 2.2 Incremental Checkpointing with Non-volatile Memory -- 2.3 The Problem with Current Checkpoint -- 3 Design and Implementation.

3.1 Contiguous Memory Regions to be Visited -- 3.2 Application-Based Coarse-Grained Checkpoint: Loose Monitoring Granularity for "Hot" Applications -- 3.3 Huge Page Support -- 4 Experiments -- 4.1 Experimental Setup and Benchmarks -- 4.2 Performance Metrics and Corresponding Results -- 5 Related Work -- 6 Conclusion -- Acknowledgments -- References -- DASM: A Dynamic Adaptive Forward Assembly Area Method to Accelerate Restore Speed for Deduplication-Based Backup Systems -- 1 Introduction -- 2 Background and Motivation -- 2.1 Data Deduplication -- 2.2 Fragmentation Problem and Restoration Speed -- 2.3 Forward Assembly Area -- 2.4 Our Motivation -- 3 Design of DASM -- 3.1 Dynamic Orthotics -- 3.2 Near-Optimal Cache Policy -- 4 Performance Evaluation -- 4.1 Evaluation Methodology -- 4.2 Experimental Results -- 4.3 Analysis -- 5 Conclusions -- References -- Scheduling and Load-Balancing -- A Statistics Based Prediction Method for Rendering Application -- 1 Introduction -- 2 Related Works -- 3 Proposed Strategy -- 3.1 System Design -- 3.2 Parsing of Parameters -- 3.3 Strategy Description -- 4 Experimental Results -- 5 Conclusion and Future Work -- References -- IBB: Improved K-Resource Aware Backfill Balanced Scheduling for HTCondor -- Abstract -- 1 Introduction -- 2 Related Work -- 3 Improved Backfill Balance Scheduling -- 3.1 Thought of IBB -- 3.2 Apply Backfilling to HTCondor -- 4 Experiments and Evaluations -- 4.1 Simulation Experiments -- 4.2 Experiments Conducted on HTCondor -- 4.3 Discussion -- 5 Conclusion and Future Work -- Acknowledgments -- References -- Multipath Load Balancing in SDN/OSPF Hybrid Network -- Abstract -- 1 Introduction -- 2 Problem Formulation -- 2.1 Hybrid Network Scenario -- 2.2 Optimization Problem Definition -- 3 Load Balancing in Hybrid SDN Network -- 3.1 Disjoint Multipath Calculation -- 3.2 LRU: A New FPTAS Algorithm.

4 Experiments and Evaluation -- 5 Conclusion -- Acknowledgments -- References -- Heterogeneous Systems -- A Study of Overflow Vulnerabilities on GPUs -- 1 Introduction -- 2 Background on CUDA Architecture -- 3 Empirical Evaluation of GPU Vulnerabilities -- 3.1 Experiment Setup -- 3.2 Stack Overflow -- 3.3 Heap Overflow -- 3.4 Other Vulnerabilities -- 4 Discussions and Countermeasures -- 5 Related Work -- 6 Conclusion -- References -- Streaming Applications on Heterogeneous Platforms -- 1 Introduction -- 2 Related Work -- 3 A Statistical View -- 3.1 Benchmarks and Datasets -- 3.2 Experimental Platforms -- 3.3 Measurement Methodology -- 3.4 Results and Analysis -- 4 Our Streaming Approach -- 4.1 Categorization -- 4.2 Code Streaming -- 5 Experimental Results -- 6 Conclusion --

References -- Data Processing and Big Data -- DSS: A Scalable and Efficient Stratified Sampling Algorithm for Large-Scale Datasets -- 1 Introduction -- 2 Related Work -- 3 Stratified Sampling Queries -- 3.1 A Single Stratified Sampling Query -- 3.2 Sequential Answering Process -- 4 Distributed Sampling Design -- 4.1 Sampling Representativeness in Distributed Environment -- 4.2 Distributed Algorithm Spark-SQE -- 4.3 Improved Distributed Algorithm DSS -- 5 Experimental Evaluation -- 5.1 Experimental Setup -- 5.2 Results -- 6 Conclusion -- References -- A Fast and Better Hybrid Recommender System Based on Spark -- 1 Introduction -- 1.1 Hybrid Recommender Systems -- 1.2 Work of Paper -- 2 Related Work -- 2.1 Recommender Systems -- 2.2 Weight Analysis -- 2.3 Spark -- 3 Design Overview -- 3.1 Objective Function -- 3.2 Weight Calculation -- 4 Implementation -- 4.1 Modules -- 4.2 Discussion -- 5 Performance -- 5.1 Evaluation Index -- 5.2 Experimental Setup -- 5.3 Performance Comparison -- 6 Conclusion and Future Work -- References.

Discovering Trip Patterns from Incomplete Passenger Trajectories for Inter-zonal Bus Line Planning -- Abstract -- 1 Introduction -- 2 Related Work -- 3 Problem Definition -- 4 Mining Frequent Trip Patterns for Inter-zonal Bus Lines -- 4.1 Mining Commuters -- 4.2 Mining FBPT4BL Patterns -- 4.3 Time Complexity Analysis -- 5 Evaluation -- 5.1 Sensitivity Analysis of MinPts and -- 5.2 Effective and Efficient Analysis of MCFTP Approach -- 6 Conclusion -- Acknowledgments -- References -- FCM: A Fine-Grained Crowdsourcing Model Based on Ontology in Crowd-Sensing -- Abstract -- 1 Introduction -- 2 Characteristic Ontology Triple -- 3 Recommendation Method -- 3.1 Recommendation Index -- 3.2 Service Decision Tree -- 4 Experiment -- 5 Conclusion -- Acknowledgments -- References -- QIM: Quantifying Hyperparameter Importance for Deep Learning -- 1 Introduction -- 2 Background -- 2.1 Deep Learning (DL) -- 2.2 PB Design -- 3 QIM -- 3.1 Overview -- 3.2 Identifying the Value Range for Each Hyperparameter -- 3.3 QIM -- 4 Experimental Setup -- 5 Evaluation -- 5.1 Supervised Learning -- 5.2 Unsupervised Pre-training -- 5.3 Time Cost -- 6 Related Work -- 7 Conclusion -- References -- Algorithms and Computational Models -- Toward a Parallel Turing Machine Model -- 1 Introduction and Motivation -- 2 Existing Work on Parallel Turing Machine --- A Disappointing Status Report -- 2.1 Existing Parallel Turing Machine Proposals -- 2.2 The Disappointing Status on Parallel Turing Machine Studies -- 3 A Parallel Turing Machine Model --- Our Proposal -- 3.1 The Concept of Codelets and Codelet Graphs (CDGs) -- 3.2 The Parallel Turing Machine Model -- 3.3 An Example to Illustrate How the PTM Works -- 3.4 Determinacy Property of Our Proposed PTM -- 4 Related Work -- 4.1 Parallel Turing Machine -- 4.2 Memory Consistency Models -- 4.3 The Codelet Model. 4.4 Work on Parallel Computation Models -- 5 Conclusion and Future Work -- References -- On Determination of Balance Ratio for Some Tree Structures -- 1 Introduction -- 2 The Basic Properties and Algorithms -- 3 The Improvement of Time Complexity -- 4 Concluding Remarks -- References -- Author Index.

Sommario/riassunto

This book constitutes the proceedings of the 13th IFIP WG 10.3 International Conference on Network and Parallel Computing, NPC 2016, held in Xi'an, China, in October 2016. The 17 full papers presented were carefully reviewed and selected from 99 submissions. They are organized in the following topical sections; memory: non-volatile, solid state drives, hybrid systems; resilience and reliability; scheduling and load-balancing; heterogeneous systems; data processing and big data; and algorithms and computational models.
