

1. Record Nr.	UNISA996503564203316
Titolo	High performance computing : 9th Latin American conference, CARLA 2022, Porto Alegre, Brazil, September 26-30, 2022 : revised selected papers // edited by Philippe Navaux [and three others]
Pubbl/distr/stampa	Cham, Switzerland : , : Springer, , [2022] ©2022
ISBN	3-031-23821-4
Descrizione fisica	1 online resource (246 pages)
Collana	Communications in Computer and Information Science ; ; v.1660
Disciplina	004.11
Soggetti	High performance computing
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Nota di bibliografia	Includes bibliographical references and index.
Nota di contenuto	Intro -- Preface -- Organization -- Contents -- A Comparative Evaluation of Parallel Programming Python Tools for Particle-in-Cell on Symmetric Multiprocessors -- 1 Introduction -- 2 Background -- 2.1 Particle-in-Cell -- 2.2 Python Parallel Programming -- 2.3 Related Work -- 3 Implementation -- 3.1 Profiling -- 3.2 Code Transformation -- 4 Experimental Results -- 4.1 Setup -- 4.2 Experiments -- 5 Discussion -- 6 Final Remarks -- References -- Accelerating GNN Training on CPU+Multi-FPGA Heterogeneous Platform -- 1 Introduction -- 2 Background -- 2.1 GNN Models -- 2.2 Mini-Batch GNN Training -- 2.3 Related Work -- 3 GNN Training on CPU+Multi-FPGA Platform -- 4 Optimizations -- 4.1 Graph Partitioning and Workload Balancing -- 4.2 Optimized GNN Kernels -- 5 Experiments -- 5.1 Experimental Setup -- 5.2 Hardware Parameter Selection and Resource Utilization -- 5.3 Performance Metrics -- 5.4 Comparison with Multi-GPU Platform -- 5.5 Scalability -- 5.6 Impact of Optimizations -- 6 Conclusion -- References -- Implementing a GPU-Portable Field Line Tracing Application with OpenMP Offload -- 1 Introduction -- 2 Background -- 2.1 Directive-Based Programming for Accelerators with OpenMP -- 2.2 Simulating Plasma Confinement in Stellarator Devices -- 2.3 Related Work -- 3 Directive-Based GPU Offloading Implementation -- 3.1 Breakdown of the Execution Flow -- 3.2 Data Management for Offloading -- 3.3 Parallelism Implementation -- 4 Results -- 4.1

Experimental Setup -- 4.2 Baseline Comparison: Single CPU Node Versus Single GPU -- 4.3 Multi-GPU Scalability -- 4.4 Economic Analysis -- 5 Conclusions -- References -- Quantitative Characterization of Scientific Computing Clusters -- 1 Introduction -- 2 Related Work -- 3 Background -- 3.1 Cluster Overhead and Coupling -- 3.2 Cluster Performance Profile -- 4 Performance Evaluation -- 4.1 Experimental Setup. 4.2 Threats to Validity -- 4.3 Results -- 4.4 Clusters Performance Profiles -- 5 Discussion -- 6 Conclusion -- References -- Towards Parameter-Based Profiling for MARE2DEM Performance Modeling -- 1 Introduction -- 2 Dataset and Application Background -- 2.1 CSEM Data -- 2.2 MARE2DEM -- 2.3 Refinement Groups -- 3 Methodology and Experimental Context -- 4 Results -- 4.1 Performance Characterization of the Microkernels -- 4.2 Iterations and Refinement Groups -- 5 Conclusion -- References -- Time-Power-Energy Balance of BLAS Kernels in Modern FPGAs -- 1 Introduction -- 2 FPGAs and NLA -- 2.1 BLAS -- 2.2 FPGAs -- 3 Evaluated Kernels -- 3.1 Vitis Libraries -- 3.2 Matrix-Matrix Multiplication (MMM) -- 4 Experimental Evaluation -- 4.1 Setup -- 4.2 Experimental Results and Discussion -- 5 Conclusions -- References -- Improving Boundary Layer Predictions Using Parametric Physics-Aware Neural Networks -- 1 Introduction -- 2 Related Work -- 3 Methodology -- 3.1 Boundary Layer Problem -- 3.2 Architecture Design -- 4 Experimental Results -- 4.1 First Setting: Reaction-Diffusion Problem -- 4.2 Second Setting: Advection-Diffusion Problem -- 5 Summary and Outlook -- References -- Towards Fire Identification Model in Satellite Images Using HPC Embedded Systems and AI -- 1 Introduction -- 2 Related Works -- 2.1 Satellite Imagery Multiscale Rapid Detection With Windowed Networks -- 2.2 Lapped Convolutional Neural Networks for Embedded Systems -- 3 Workflow -- 3.1 Dataset Elaboration -- 3.2 Algorithm Selection -- 4 Results -- 4.1 Artificial Learning -- 4.2 Evaluation Metrics -- 5 Conclusion -- 6 Future Work -- References -- A Machine Learning-Based Missing Data Imputation with FHIR Interoperability Approach in Sepsis Prediction -- 1 Introduction -- 2 State of the Art -- 2.1 Machine Learning on Clinical Features for Sepsis Prediction. 2.2 Interoperability of Healthcare Information Systems -- 3 Materials and Methods -- 3.1 Study Design -- 3.2 Dataset Early Prediction of Sepsis from Clinical Data -- 3.3 Processing and Transformation of Clinical Data to the FHIR Standard -- 3.4 Data Distribution - Hospitals A and B -- 3.5 Preprocessing of Data -- 3.6 Experiment Dataset -- 3.7 Creation of Train Test -- 3.8 Implementation of Classifiers -- 4 Experiments and Results -- 4.1 Experiment Results -- 5 Conclusions -- References -- Understanding the Energy Consumption of HPC Scale Artificial Intelligence -- 1 Introduction -- 2 Related Work -- 2.1 AI and Climate Change -- 2.2 Energy-Aware AI -- 2.3 AI Benchmarks -- 2.4 Energy Measurement Tools -- 2.5 Positioning of This Paper -- 3 Background -- 4 Benchmark Tracker -- 5 Results -- 5.1 Experimental Setting -- 5.2 Experimental Results -- 6 Conclusion and Future Work -- 6.1 Future Work -- References -- Using Big Data and Serverless Architecture to Follow the Emotional Response to the COVID-19 Pandemic in Mexico -- 1 Introduction -- 2 Related Work -- 3 Method -- 3.1 General System Architecture -- 4 Experiments -- 5 Results -- 6 Conclusions -- References -- Multi-GPU 3-D Reverse Time Migration with Minimum I/O -- 1 Introduction -- 2 Reverse Time Migration -- 3 Computational Implementation and Optimizations -- 3.1 Classical Reverse Time Migration -- 3.2 Reverse Time Migration with Wavefield Reconstruction -- 3.3 Hybrid OpenACC/MPI Implementation -- 4 Numerical Experiments -- 5 Conclusions -- References -- ParsIRNA-

Seq: An Efficient and Scalable RNAseq Analysis Workflow for Studies of Differentiated Gene Expression -- 1 Introduction -- 2 Related Works -- 3 Background on Differential Gene Expression Analysis -- 4 ParsIRNA-Seq: Workflow for DGE Analysis -- 4.1 Improvements in the Previous Implementation of the Workflow. 4.2 Multithreading and Multiprocessing -- 4.3 The Current Implementation of the ParsIRNA-Seq Workflow -- 5 Methods and Infrastructure -- 5.1 Experiment Dataset -- 5.2 Experiment Setup -- 5.3 Computational Environment Setup -- 6 Experimental Results -- 6.1 Performance and Scalability Analyses -- 6.2 I/O Performance Results Using Darshan -- 6.3 Performance Results Using SSD -- 6.4 Biological Results of RNA-Seq Data -- 7 Conclusion -- References -- Refactoring an Electric-Market Simulation Software for Massively Parallel Computations -- 1 Introduction -- 2 The SimSEE and Previous Results -- 3 Proposal -- 3.1 Loading the Playrooms for Massively-Parallel Trajectories, naive -- 3.2 Improving the Playrooms Replication, base -- 3.3 Sharing References to Avoid Memory Allocations, RefCat -- 3.4 Enhancing the Access to Shared References in the Simulation, RefDicc -- 4 Experimental Evaluation -- 4.1 Test Cases -- 4.2 Runtime Environment -- 4.3 Experimental Results -- 5 Conclusion and Future Work -- References -- Nearly Quantum Computing by Simulation -- 1 Introduction -- 2 Quantum Computing Modelling -- 2.1 An Overview of Quantum Mechanics -- 2.2 Information Theory -- 2.3 Quantum Information Theory -- 3 Quantum Computing Parallelism and Simulation -- 3.1 Quantum Computing Simulators -- 3.2 Popular Open Source Quantum Computer Simulators -- 4 Discussion and Further Work -- References -- Functionality Testing in the Automation of Scientific Application Workflows in an HPC Environment -- 1 Introduction -- 2 Infrastructure Used -- 3 Tools Used -- 3.1 Slurm -- 3.2 Singularity -- 3.3 Snakemake -- 4 Design of the Processing Flow for Testing -- 5 Analysis of Possible Cases -- 5.1 Running Python Script -- 5.2 Running Python Script with SLURM and Singularity -- 5.3 Running a Singularity Container with Snakemake Using SLURM -- 5.4 Notes for Tables 2, 3 and 4. 6 Results of Executions -- 6.1 Case 1: -- 6.2 Case 2: -- 6.3 Case 3: -- 7 Discussion of Results and Conclusions -- 7.1 Testing Time -- 7.2 Duration of Tests -- 7.3 What Limitations There Were -- 7.4 Learning -- 7.5 Conclusions and Benefits -- References -- Author Index.
