

1. Record Nr.	UNISA996475771803316
Titolo	Accelerator programming using directives : 7th international workshop, WACCPD 2020, virtual event, November 20, 2020, proceedings // edited by Sridutt Bhalachandra, Christopher Daley, and Veronica Melesse Vergara
Pubbl/distr/stampa	Cham, Switzerland : , : Springer, , [2022] ©2022
ISBN	3-030-97759-5
Descrizione fisica	1 online resource (157 pages)
Collana	Lecture Notes in Computer Science ; ; v.13194
Disciplina	005.13
Soggetti	High performance computing Microprogramming Computer programming Càlcul intensiu (Informàtica) Programació (Ordinadors) Congressos Llibres electrònics
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 -- Directive Alternatives -- Can Fortran's `do concurrent' Replace Directives for Accelerated Computing?*-8pt -- 1 Introduction -- 2 Code and Test Description -- 2.1 Code Description -- 2.2 Test Description -- 2.3 Computational Environment -- 2.4 Baseline Performance Results -- 3 Implementation -- 3.1 The Fortran do concurrent construct -- 3.2 Code Versions -- 3.3 Compiler Flag Options -- 4 Results -- 4.1 Results Using nvfortran -- 4.2 Results Using gfortran -- 4.3 Results Using ifort -- 4.4 Experimental Results -- 5 Discussion -- 6 Artifact Availability Statement -- References -- Achieving Near-Native Runtime Performance and Cross-Platform Performance Portability for Random Number Generation Through SYCL Interoperability -- 1 Introduction -- 1.1 Contribution -- 2 Related Work -- 2.1 Parallel Programming Frameworks -- 2.2 Linear Algebra Libraries -- 2.3 The Proposed

Approach -- 3 SYCL Overview -- 4 SYCL-Based RNG Implementations of NVIDIA and AMD GPUs in oneMKL -- 4.1 Technical Aspects -- 4.2 Native cuRAND and hipRAND flow -- 4.3 Implementation of cuRAND and hipRAND in oneMKL -- 5 Benchmark Applications -- 5.1 Random Number Generation Burner -- 5.2 FastCaloSim -- 6 Performance Evaluation -- 6.1 Performance Portability Metrics -- 6.2 Hardware Specifications -- 6.3 Software Specifications -- 7 Results -- 8 Conclusions and Future Work -- References -- Directive Extensions -- Extending OpenMP for Machine Learning-Driven Adaptation -- 1 Introduction -- 2 A Motivating Example -- 3 A Vision -- 4 The declare adaptation Directive -- 4.1 Syntax and Semantics of declare adaptation -- 4.2 Examples Using metadirective -- 5 Implementation -- 5.1 Compiler Support -- 5.2 Runtime Support -- 6 Evaluation -- 6.1 Software and Hardware Configurations -- 6.2 Performance Results -- 6.3 Accuracy of Prediction Models. 6.4 Overhead Analysis -- 7 Related Work -- 8 Conclusion -- References -- Directive Case Studies -- GPU Porting of Scalable Implicit Solver with Green's Function-Based Neural Networks by OpenACC -- 1 Introduction -- 2 Solver with Green's Function-Based NN Preconditioner -- 2.1 Target Problem -- 2.2 GF-Based NN Predictor -- 2.3 Scalable Solver Algorithm Using GF-Based NN Predictor -- 3 GPU Porting of Solver with Green's Function-Based NN Preconditioner Using OpenACC -- 4 Performance Measurement -- 4.1 Problem Used for Measurement -- 4.2 Performance Measurement Environment -- 4.3 Solver Performance on GPU-Based System -- 4.4 Weak Scaling on GPU-Based System -- 5 Closing Remarks -- References -- Challenges Porting a C++ Template-Metaprogramming Abstraction Layer to Directive-Based Offloading -- 1 Introduction -- 2 Related Work -- 3 Methods and APIs -- 3.1 Alpaka and PIconGPU -- 3.2 Review of OpenACC and OpenMP Target -- 3.3 Experimental Setup -- 4 Porting Alpaka -- 4.1 Final Touches: PIconGPU -- 5 Major Hurdles and Discussion -- 5.1 Standards Issues -- 5.2 Compiler and Runtime Issues -- 5.3 Preliminary Results -- 6 Conclusions and Outlook -- References -- Accelerating Quantum Many-Body Configuration Interaction with Directives -- 1 Introduction -- 2 Computational Motifs in Configuration Interaction Code MFDn -- 2.1 Matrix Sparsity Determination -- 2.2 Parallel Prefix Sum -- 2.3 Filling Shared Arrays -- 2.4 Array Reductions -- 3 Conclusion and Outlook -- References -- GPU Offloading of a Large-Scale Gyrokinetic Particle-in-Cell Fortran Code on Summit: From OpenACC to OpenMP -- 1 Introduction -- 2 Software and Experimental Setup -- 2.1 Experimental Setup -- 2.2 OpenMP GPU Offloading -- 3 The Structure of GEM -- 4 Results and Analysis -- 4.1 Speedup Performance and Roofline Analysis for Single Node -- 4.2 Scalability Analysis. 4.3 Investigation of Hardware Threads -- 5 Discussion -- 6 Summary -- References -- Author Index.

---