

1. Record Nr.	UNINA9910131743203321
Titolo	Kant et les Lumieres europeennes : actes du 7e Congres de la Societe d'etudes kantiennes de langue francaise organise a Naples, 20-22 octobre 2005 // par le Dipartimento di filosofia e politica de l'Universita degli studi di Napoli l'Orientale avec le concours de la Societa italiana di studi sul secolo XVIII et de l'Istituto italiano per gli studi filosofici ; publies sous la direction de Lorenzo Bianchi, Jean Ferrari et Alberto Postigliola ; [publies par l'Association des societes de philosophie de langue francaise et la Societe d'etudes kantiennes de langue francaise]
Pubbl/distr/stampa	Napoli, : Liguori Paris, : Librairie philosophique J. Vrin, 2009
Altri autori (Persone)	BianchiLorenzo <1949-> FerrariJean PostigliolaAlberto <1942->
Soggetti	Enlightenment - Europe Philosophy, Modern - 18th century
Lingua di pubblicazione	Francese
Formato	Materiale a stampa
Livello bibliografico	Monografia

2. Record Nr.	UNINA9910349261003321
Titolo	Accelerator Programming Using Directives : 4th International Workshop, WACCPD 2017, Held in Conjunction with the International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2017, Denver, CO, USA, November 13, 2017, Proceedings // edited by Sunita Chandrasekaran, Guido Juckeland
Pubbl/distr/stampa	Cham : , : Springer International Publishing : , : Imprint : Springer, , 2018
ISBN	9783319748962 3319748963
Edizione	[1st ed. 2018.]
Descrizione fisica	1 online resource (IX, 183 p. 59 illus.)
Collana	Programming and Software Engineering, , 2945-9168 ; ; 10732
Disciplina	004.3
Soggetti	Compilers (Computer programs) Logic design Operating systems (Computers) Computer programming Computer engineering Computer networks Computer science Compilers and Interpreters Logic Design Operating Systems Programming Techniques Computer Engineering and Networks Models of Computation
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Nota di contenuto	Intro -- Preface -- Organization -- Contents -- Applications -- An Example of Porting PETSc Applications to Heterogeneous Platforms with OpenACC -- Abstract -- 1 Introduction -- 2 Workflow and System Description -- 2.1 Workflow -- 2.2 System -- 3 Results and Discussion -- 3.1 Profiling with Score-P -- 3.2 The Most Expensive Kernel:

MatMult_SeqAIJ -- 3.3 Four Steps Toward the Final Version of OpenACC Kernel -- 4 Speedups and Strong Scaling -- 5 Conclusion -- Acknowledgement -- References -- Hybrid Fortran: High Productivity GPU Porting Framework Applied to Japanese Weather Prediction Model -- 1 Introduction -- 1.1 ASUCA on GPU -- 1.2 Parallelization Granularity -- 1.3 Memory Layout -- 1.4 Related Work -- 1.5 Problem Summary -- 2 Hybrid Fortran Language Extension and Code Transformation -- 2.1 Parallel Loop Abstraction -- 2.2 Compile-Time Defined Memory Layout and Device Data Region -- 2.3 Transformed Code -- 3 Code Transformation Method -- 4 Productivity- and Performance Results -- 5 Conclusion and Future Work -- References -- Implicit Low-Order Unstructured Finite-Element Multiple Simulation Enhanced by Dense Computation Using OpenACC -- 1 Introduction -- 2 Finite-Element Earthquake Simulation Designed for the K Computer -- 3 Proposed Solver for GPUs Using OpenACC -- 3.1 Modification of Algorithm for GPUs -- 3.2 Introduction of OpenACC -- 4 Performance Measurements -- 5 Application Example -- 6 Concluding Remarks -- References -- Runtime Environments -- The Design and Implementation of OpenMP 4.5 and OpenACC Backends for the RAJA C++ Performance Portability Layer -- 1 Introduction -- 2 RAJA -- 2.1 Basic Execution Policies -- 2.2 RAJA::NestedPolicy and Loop Transformations -- 3 Embedding Directives in the C++ Type System -- 3.1 Defining Policy Tags for a Backend -- 3.2 Constructing Explicit Execution Policy Types.

3.3 Implement for all Specializations -- 4 Case Study: OpenMP 4.5 -- 5 Case Study: OpenACC -- 6 Evaluation -- 6.1 Test Set -- 6.2 Goals and Non-Goals -- 6.3 Compilation Overhead -- 6.4 Runtime Overhead -- 7 Future Work and Conclusion -- References -- Enabling GPU Support for the COMPSs-Mobile Framework -- 1 Introduction -- 2 Related Work -- 3 Programming Model -- 3.1 Extension for GPU Support -- 4 Runtime Support Implementation -- 4.1 COMPSs-Mobile Runtime Architecture -- 4.2 OpenCL Platform -- 5 Performance Evaluation -- 5.1 OpenCL Platform Performance -- 5.2 Load Balancing Policies -- 6 Conclusions and Future Work -- References -- Concurrent Parallel Processing on Graphics and Multicore Processors with OpenACC and OpenMP -- Abstract -- 1 Introduction -- 2 MBFLO3 Application -- 2.1 Mathematical Formulation -- 2.2 Numerical Method -- 3 Heterogeneous Multiblock Computing Strategy -- 3.1 Multicore Host Parallelism -- 3.2 Manycore Accelerator Parallelism -- 3.3 Heterogeneous Host-Device Parallelism -- 4 Performance Results and Analysis -- 5 Conclusions -- Acknowledgements -- References -- Program Evaluation -- Exploration of Supervised Machine Learning Techniques for Runtime Selection of CPU vs. GPU Execution in Java Programs -- 1 Introduction -- 2 Motivation -- 3 Compiling Java to GPUs -- 3.1 Java Parallel Stream API -- 3.2 JIT Compilation for GPUs -- 4 Exploring Supervised Machine Learning Algorithms -- 4.1 Supervised Machine Learning -- 4.2 Generating Subsets of Features -- 4.3 Constructing Prediction Models -- 4.4 Integrating Prediction Models -- 5 Experimental Results -- 5.1 Experimental Protocol -- 5.2 Overall Summary -- 5.3 Accuracies on the Full Set of Features -- 5.4 Exploring ML Algorithms by Feature Subsetting -- 5.5 Lessons Learned -- 6 Related Work -- 6.1 GPU Code Generation from High-Level Languages -- 6.2 Offline Model Construction.

7 Conclusions -- A Appendix -- References -- Automatic Testing of OpenACC Applications -- 1 Introduction -- 2 Testing a GPU Port of a Numerical Application -- 3 Autocompare with OpenACC -- 4 Autocompare Implementation -- 5 Experiments -- 6 Related Work -- 7 Future Work -- 8 Conclusion -- References -- Evaluation of

Asynchronous Offloading Capabilities of Accelerator Programming Models for Multiple Devices -- 1 Introduction -- 2 Related Work -- 3 Accelerator Programming Models -- 3.1 CUDA -- 3.2 OpenCL -- 3.3 OpenACC -- 3.4 OpenMP -- 4 Implementing the Conjugate Gradient Method -- 5 Performance Results on NVIDIA GPUs -- 5.1 Data Transfers with the Host -- 5.2 Single Device -- 5.3 Two Devices -- 6 Performance Results on Intel Xeon Phi Coprocessors -- 6.1 Single Device -- 6.2 Two Devices -- 7 Summary -- References -- Author Index.

Sommario/riassunto

This book constitutes the refereed post-conference proceedings of the 4th International Workshop on Accelerator Programming Using Directives, WACCPD 2017, held in Denver, CO, USA, in November 2017. The 9 full papers presented have been carefully reviewed and selected from 14 submissions. The papers share knowledge and experiences to program emerging complex parallel computing systems. They are organized in the following three sections: applications; environments; and program evaluation.
