

1. Record Nr.	UNISA996465608603316
Titolo	Languages and Compilers for Parallel Computing [[electronic resource] ] : 22nd International Workshop, LCPC 2009, Newark, DE, USA, October 8-10, 2009, Revised Selected Papers // edited by Guang R. Gao, Lori Pollock, John Cavazos, Xiaoming Li
Pubbl/distr/stampa	Berlin, Heidelberg : , : Springer Berlin Heidelberg : , : Imprint : Springer, , 2010
ISBN	1-280-38685-1 9786613564771 3-642-13374-6
Edizione	[1st ed. 2010.]
Descrizione fisica	1 online resource (XI, 426 p. 186 illus.)
Collana	Theoretical Computer Science and General Issues, , 2512-2029 ; ; 5898
Disciplina	005.13
Soggetti	Compilers (Computer programs) Computer programming Computer networks Artificial intelligence—Data processing Artificial intelligence Computer science Compilers and Interpreters Programming Techniques Computer Communication Networks Data Science Artificial Intelligence Models of Computation
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Note generali	Bibliographic Level Mode of Issuance: Monograph
Nota di bibliografia	Includes bibliographical references and index.
Nota di contenuto	A Communication Framework for Fault-Tolerant Parallel Execution -- The STAPL pList -- Hardware Support for OpenMP Collective Operations -- Loop Transformation Recipes for Code Generation and Auto-Tuning -- MIMD Interpretation on a GPU -- TL-DAE: Thread-Level Decoupled Access/Execution for OpenMP on the Cyclops-64 Many-Core Processor -- Mapping Streaming Languages to General Purpose Processors

through Vectorization -- A Balanced Approach to Application Performance Tuning -- Automatically Tuning Parallel and Parallelized Programs -- DFT Performance Prediction in FFTW -- Safe and Familiar Multi-core Programming by Means of a Hybrid Functional and Imperative Language -- Hierarchical Place Trees: A Portable Abstraction for Task Parallelism and Data Movement -- OSCAR API for Real-Time Low-Power Multicores and Its Performance on Multicores and SMP Servers -- Programming with Intervals -- Adaptive and Speculative Memory Consistency Support for Multi-core Architectures with On-Chip Local Memories -- Synchronization-Free Automatic Parallelization: Beyond Affine Iteration-Space Slicing -- Automatic Data Distribution for Improving Data Locality on the Cell BE Architecture -- Automatic Restructuring of Linked Data Structures -- Using the Meeting Graph Framework to Minimise Kernel Loop Unrolling for Scheduled Loops -- Efficient Tiled Loop Generation: D-Tiling -- Effective Source-to-Source Outlining to Support Whole Program Empirical Optimization -- Speculative Optimizations for Parallel Programs on Multicores -- Fastpath Speculative Parallelization -- PSnAP: Accurate Synthetic Address Streams through Memory Profiles -- Enforcing Textual Alignment of Collectives Using Dynamic Checks -- A Code Generation Approach for Auto-Vectorization in the Spade Compiler -- Portable Just-in-Time Specialization of Dynamically Typed Scripting Languages -- Reducing Training Time in a One-Shot Machine Learning-Based Compiler -- Optimizing Local Memory Allocation and Assignment through a Decoupled Approach -- Unrolling Loops Containing Task Parallelism.

---

## Sommario/riassunto

It is our pleasure to present the papers accepted for the 22nd International Workshop on Languages and Compilers for Parallel Computing held during October 8–10 2009 in Newark Delaware, USA. Since 1986, LCPC has become a valuable venue for researchers to report on work in the general area of parallel computing, high-performance computer architecture and compilers. LCPC 2009 continued this tradition and in particular extended the area of interest to new parallel computing accelerators such as the IBM Cell Processor and Graphic Processing Unit (GPU). This year we received 52 submissions from 15 countries. Each submission received at least three reviews and most had four. The PC also sought additional external reviews for contentious papers. The PC held a all-day phone conference on August 24 to discuss the papers. PC members who had a conflict of interest were asked to leave the call temporarily when the corresponding papers were discussed. From the 52 submissions, the PC selected 25 full papers and 5 short papers to be included in the workshop proceeding, representing a 58% acceptance rate. We were fortunate to have three keynote speeches, a panel discussion and a tutorial in this year's workshop. First, Thomas Sterling, Professor of Computer Science at Louisiana State University, gave a keynote talk titled "HPC in Phase Change: Towards a New Parallel Execution Model." Sterling argued that a new multi-dimensional research thrust was required to realize the design goals with regard to power, complexity, clock rate and reliability in the new parallel computer systems. ParalleX, an exploratory execution model developed by Sterling's group was introduced to guide the co-design of new architectures, programming methods and system software.

---