

Table of Contents

Third International Conference on High-Performance and Embedded Architectures and Compilers (HiPEAC)

Dynamic Cache Partitioning Based on the MLP of Cache Misses	3
<i>Miquel Moreto, Francisco J. Cazorla, Alex Ramirez, and Mateo Valero</i>	
Cache Sensitive Code Arrangement for Virtual Machine	24
<i>Chun-Chieh Lin and Chuen-Liang Chen</i>	
Data Layout for Cache Performance on a Multithreaded Architecture . . .	43
<i>Subhradyuti Sarkar and Dean M. Tullsen</i>	
Improving Branch Prediction by Considering Affectors and Affectees Correlations	69
<i>Yiannakis Sazeides, Andreas Moustakas, Kypros Constantinides, and Marios Kleanthous</i>	

Eighth MEDEA Workshop (Selected Papers)

Introduction	91
<i>Sandro Bartolini, Pierfrancesco Foglia, and Cosimo Antonia Prete</i>	
Exploring the Architecture of a Stream Register-Based Snoop Filter . . .	93
<i>Matthias Blumrich, Valentina Salapura, and Alan Gara</i>	
CROB: Implementing a Large Instruction Window through Compression	115
<i>Fernando Latorre, Grigorios Magklis, Jose González, Pedro Chaparro, and Antonio González</i>	
Power-Aware Dynamic Cache Partitioning for CMPs	135
<i>Isao Kotera, Kenta Abe, Ryusuke Egawa, Hiroyuki Takizawa, and Hiroaki Kobayashi</i>	
A Multithreaded Multicore System for Embedded Media Processing . . .	154
<i>Jan Hoogerbrugge and Andrei Terechko</i>	

Regular Papers

Parallelization Schemes for Memory Optimization on the Cell Processor: A Case Study on the Harris Corner Detector	177
<i>Tarik Saidani, Lionel Lacassagne, Joel Falcou, Claude Tadonki, and Samir Bouaziz</i>	
Constructing Application-Specific Memory Hierarchies on FPGAs	201
<i>Harald Devos, Jan Van Campenhout, Ingrid Verbauwhede, and Dirk Stroobandt</i>	
First Workshop on Programmability Issues for Multi-core Computers (MULTIPROG)	
autopin – Automated Optimization of Thread-to-Core Pinning on Multicore Systems	219
<i>Tobias Klug, Michael Ott, Josef Weidendorfer, and Carsten Trinitis</i>	
Robust Adaptation to Available Parallelism in Transactional Memory Applications	236
<i>Mohammad Ansari, Mikel Luján, Christos Kotselidis, Kim Jarvis, Chris Kirkham, and Ian Watson</i>	
Efficient Partial Roll-Backing Mechanism for Transactional Memory Systems	256
<i>M.M. Waliullah</i>	
Software-Level Instruction-Cache Leakage Reduction Using Value-Dependence of SRAM Leakage in Nanometer Technologies	275
<i>Maziar Goudarzi, Tohru Ishihara, and Hamid Noori</i>	
Author Index	301