MEDEA Workshop

MEmory performance:
DEaling with Applications, systems and architecture

held in conjunction with PACT 2007 Conference

http://garga.iet.unipi.it/medea07/

"*Medea (memory) helps Jason (processor)…*
…will Medea kill Jason's sons?"

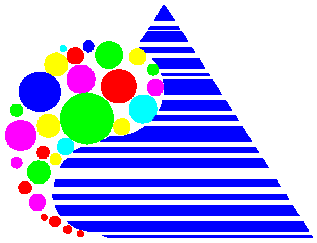

PACT-2007
Brasov, Romania
Sept 15-19, 2007

ACM

Call for Papers

MEDEA aims to continue the high level of interest of the previous editions held with PACT Conference since 2000.

Due to the ever-increasing gap between CPU and memory speed, there is always great interest in evaluating and proposing processor, multiprocessor, CMP, multi-core, SoC and system architectures dealing with the "memory wall" and wire-delay problems. At the same time, a modular high-level design is becoming more and more attracting in order to reduce design costs.

In this scenario, design solutions and their corresponding performance are shaped by the combined pressure of a) technological opportunities and limitations, b) features and organization of system architecture and c) critical requirements of specific application domains. Evaluating and controlling the effects on the memory subsystem (e.g. caches, interconnection, bus, memory, coherence) of any architectural proposal is extremely important both from the performance (e.g. bandwidth, latency, predictability) and power (e.g. static, dynamic, manageability) points of view.

In particular, the emerging trend of single-chip multi-core solutions, will push towards new design principles for memory hierarchy and interconnection networks, especially when the design is aimed to build systems with a high number of cores (many-core instead of multi-core systems), which aim to scale performance and power efficiency in a variety of application domains.

From a slightly different point of view, the mutual interaction between the application behavior and the system on which it executes, is responsible of the figures of merit of the memory subsystem and, therefore, pushes towards specific solutions. In addition, it can suggest specific compile/link time tunings for adapting the application to the features of the target architecture.

In the overall picture, power consumption requirements are increasingly important cross-cutting issues and raise specific challenges.

Typical architectural choices of interest include, but are not limited to, single processor, chip and board multiprocessors, SoC, traditional and tiled/clustered architectures, multithreaded or VLIW architectures with emphasis on single-chip design, massive parallelism designs, heterogeneous architectures, architectures equipped with application-domain accelerators as well as endowed with reconfigurable modules. Application domains encompass embedded (e.g. multimedia, mobile, automotive, automation, medical), commercial (e.g. Web, DB, multimedia), networking applications, security, etc. The emerging network on chip infrastructure may suggest new solutions and issues.

MEDEA Workshop wants to continue to be a forum for academic and industrial people to meet, discuss and exchange their ideas, experience and solutions in the design and evaluation of architectures for embedded, commercial and general/special purpose systems taking into account memory issues, both directly and indirectly.

Proceedings of the Workshop will be published under ACM ISBN, and appear also in the ACM Digital Library.
Accepted papers will be considered for publication, in an extended version, for the March 2008 special issue of Transactions on HIPEAC (www.hipeac.net/journal).

Organizing and steering Committee

Sandro Bartolini, University of Siena, Italy
Pierfrancesco Foglia, University of Pisa, Italy
Roberto Giorgi, University of Siena, Italy
Cosimo Antonio Prete, University of Pisa, Italy

Program Committee

Erik Altman, IBM T.J. Watson Res. Center, NY,USA
Alessio Bechini, University of Pisa, Italy
Koen Bertels, TU-DELFT, Delft, NL
John Cavazos,University of Edinburgh,Edinburgh,UK
Björn Franke, University of Edinburgh, Edinburgh, UK
Georgi Gaydadjiev, TU-DELFT, Delft, NL
Ali Hurson, Penn. State University, PA, USA
Timothy Jones,University of Edinburgh,Edinburgh,UK
David Kaeli, Northeastern University, MA, USA
Krishna Kavi, University of North Texas, TX, USA
Stefanos Kaxiras, University of Patras, Greece
Stephen Keckler,University of Texas at Austin,TX,USA
Hiroaki Kobayashi, Tohoku University, Sendai, Japan
David M. Koppelman, LSU, Baton Rouge, LA, USA
Enrico Martinelli, University of Siena, Italy
Avi Mendelson, Intel, Haifa, Israel
Aleksandar Milenković, Un. of Alabama, Huntsville, AL
Yale Patt, University of Texas at Austin, TX, USA
Alex Ramirez, UPC, Barcelona, Spain
Toshinori Sato, Kyushu Institute of Tech.,Iizuka, Japan
Per Stenström, Chalmers Univ. of Technol., Sweden
Theo Ungerer, University of Augsburg, DE, EU
Mateo Valero, UPC, Barcelona, Spain
Wei Zhang, South. Illinois Univ.,Carbondale, IL, USA

Important Dates

Jul. 16	Abstract submission (not mandatory)
Jul. 23	**Paper submission deadline**
Aug. 14	Notification of acceptance
Aug. 25	Final papers due
Sep. 16	Start of MEDEA-2007 Workshop in Brazov

Information for Authors

The papers should be **6-8 pages** in length. Abstracts and papers should be submitted in either postscript or PDF format (NO Hard copy, postal, submissions) by e-mail to Pierfrancesco Foglia and Sandro Bartolini.
Paper should be written in standard ACM SIG Proceedings Template (see MEDEA web site for details).
To speed-up the reviewing process, we encourage (it is not mandatory) the abstract submission by July 16, 2007.

Topics of interest include (but are not limited to):

- Memory hierarchy design, analysis, tuning for embedded, general and special purpose systems
- On-chip Multiprocessors and System On Chip architectures, development tools and applications
  - Issues in memory hierarchy design of scalable single chip systems
  - Memory hierarchy issues for heterogeneous and accelerator-based systems
  - Solutions for embedded, DSP, commercial, scientific and technical workloads
  - Inter-Chip and Intra-Chip memory latency tolerant and reduction techniques
- Cache coherence and memory models
- Compile/link time optimization techniques
- Bus/Interconnection Architectures
- Network On Chip
- Low-Power/Wire Delay design of memory hierarchies
- Processor and System Architectures
- Exploitation of application parallelism (e.g.: ILP, TLP, DLP)
- Academic/industrial experience in high performance, embedded systems and memory design