

CALL FOR PAPERS

Workshop on GPU Computing Held with PPAM 2015 -- 11th International Conference on Parallel Processing and Applied Mathematics

Krakow, Poland

September 6-9, 2015

<http://ppam.pl>

Overview:

GPU programming is now a much richer environment than it used to be a few years ago. On top of the two major programming languages, CUDA and OpenCL, libraries (e.g., cufft) and high level interfaces (e.g., thrust) have been developed that allow a fast access to the computing power of GPUs without detailed knowledge or programming of GPU hardware.

Annotation-based programming models (e.g., PGI Accelerator), GPU plug-ins for existing mathematical software (e.g., Jacket in Matlab), GPU script languages (e.g., PyOpenCL), and new data parallel languages (e.g., Copperhead) bring GPU programming to a new level.

A major decision for libraries and high level programming tools is the positioning within the triangle of performance, coding comfort and specialization. The spectrum ranges from high performance building blocks for common numeric or discrete transformations, to application domain specific libraries facilitating the solution of a certain class of problems, to general high level abstractions increasing the programmer's productivity.

By sharing the experiences of utilizing programming abstractions we hope that the participants of the workshop will gain a better understanding about which tools are good for which type of problem and which trade-offs between performance, coding comfort and specialization are available.

A major criticism of programming abstractions is that they look great on small examples but fail on practical problems. Therefore, this workshop invites, in particular, submissions that deal with practical applications that have successfully employed GPU libraries or high level programming tools. The focus may lie both on the development of the libraries or utilization of existing tools. Workshop topics include, but are not limited to:

- GPU applications coded with high level programming tools
- GPU library development and application
- Comparison of different programming abstractions on the same/similar applications
- Comparison of the same/similar programming abstractions on different applications
- Performance and coding effort of high level tools against hand-coded approaches on the GPU
- Performance and coding effort on multi-core CPUs against GPUs utilizing programming abstractions
- Classification of different programming abstractions with respect to their best application area

Submissions:

The rules of PPAM conference apply. In particular:

- Papers will be refereed and accepted on the basis of their scientific merit and relevance to the conference topics.
- Regular papers are not to exceed 10 pages (LNCS style).
- Abstracts of accepted papers will be available during the conference in form of a brochure.
- Only papers presented at PPAM 2015 will be included into the proceedings, which is planned to be published after the conference by Springer in the LNCS series.

Important dates:

Submission of Papers: ~~April 26th~~, **May 11th, 2015**

Notification of Acceptance: ~~June 5th~~, **June 20th, 2015**

Conference: September 6-9, 2015

Camera-Ready Papers: November 15th, 2015

Organizers:

José R. Herrero, Enrique S. Quintana-Orti, and Robert Strzodka

Program Committee:

All submitted papers will be reviewed by the workshop scientific committee members.

Special Issue:

The authors of the best papers presented at PPAM, selected by the Program Committee, will be invited to submit extended versions of their work to special issues of the following journals: International Journal of High Performance Computing Applications, and Concurrency and Computation: Practice and Experience.

Award:

The authors of the best paper presented at the Minisymposium will receive a GeForce GTX480 donated by NVIDIA.