universität wien

Faculty of Computer Science

Research Group Scientific Computing

Prof. Dr. Siegfried Benkner
Währingerstrasse 29
A- 1090 Wien

T+43-1-4277-784 10
siegfried.benkner@univie.ac.at

Vienna, 2020-02-28

Review of habilitation thesis
"GPU-Accelerated Methods for Content-based Retrieval"
by Dr. Martin Krulis

Computer architectures have undergone very significant changes over the past 15 years. On the one hand there has been a shift towards multicore processors, which are now ubiquitous, on the other hand, heterogeneous architectures, combining different types of processors like CPUs and GPUs have become increasingly important. A major driving force behind these developments is energy efficiency, which has become a priority in the design of processor architectures and computer systems. However, while we witness tremendous developments in parallel processor architectures, major challenges remain in developing efficient parallel software and applications that can exploit those architectures to their full potential. This is mainly due to the architectural complexity of parallel systems and the lack of appropriate higher-level parallel programming interfaces. Thus, developing efficient algorithms for a parallel system requires to take into consideration a multitude of low-level architectural details with respect to the partitioning, granularity and degree of parallelism, data management, and synchronization. These challenges become even more drastic when targeting hybrid/heterogeneous architectures, where, e.g., some parts of an algorithm need to run on a CPU and others on a GPU. Another major issue is performance portability of codes across different types and generations of architectures, which is often not achievable without re-tuning the code for each specific architecture.

The research of Dr. Krulis done in the context of his habilitation thesis exactly addresses these challenges. More specifically, the presented thesis deals with the parallelization of important algorithms and methods for content-based retrieval in multi-media databases in order to improve their efficiency on modern GPU-based parallel systems.

The presented thesis is an accumulative thesis that summarizes the key scientific contributions of seven papers that have been published in top international journals and conference proceedings. In the following, I briefly review some of the main scientific contributions of this thesis for advancing the state-of-the art in efficient methods for content-based retrieval in multi-media databases.

Dr. Krulis has developed a world-wide first implementation of the Signature Quadratic Form Distance (SQFD) method for GPUs, which is key for similarity search in multimedia databases. For this work, he developed a generic SIMD variant of a GPU-accelerated Wagner-Fisher algorithm with recursive partitioning and special techniques for optimizing the usage of GPU registers. An important property of this method is that the recursive partitioning, facilitates performance portability across different types of parallel architectures. Such portability across CPUs, Xeon Phi co-processors, and GPUs has been successfully demonstrated, and will be very important for targeting future processor architectures.

Another significant contribution is the design and realization of an efficient parallel method for metric indexing. This work focused on the development of a parallel Linear Approximating Eliminating Search Algorithm (LAESA) for GPUs and the demonstration of its near optimal performance. Key findings of this work comprise new methods for selecting effective block sizes for a given similarity model as well as the development of novel heuristics for range estimation.

Original research work was also performed for the development of hybrid algorithms that simultaneously use CPU and GPU resources. In particular, a hybrid CPU/GPU algorithm for permutation-based indexing for approximative similarity search was developed, which significantly outperformed other existing algorithms. The data management techniques developed for this algorithm should be applicable for the optimization of similar data-intensive algorithms on CPU/GPU systems.

Additional scientific contributions are the design and implementation of a GPU algorithm for the extraction of feature signatures, again pioneering work, since it is the first such algorithm for GPUs, and a novel GPU algorithm for image denoising.

Although the major scientific contributions of this thesis are in the area of algorithms and methods for multi-media databases, most of the presented research is of much broader relevance.

Overall Assessment

The papers on which the habilitation thesis are based are well-written, in excellent English, and technically sound. All papers have been published in high-quality, peer-reviewed international journals and conferences, and all of them contain new scientific contributions. The proposed methods and algorithms have been rigorously specified, theoretically analyzed, implemented, and experimentally evaluated on state-of-the art parallel systems. Related work has been covered appropriately.

Concluding, the methodological quality of Dr. Krulis' habilitation thesis and the presented research papers are of impeccable methodological quality and contain new academic insights. Dr. Krulis has clearly demonstrated his command of the habilitation subject as well as his ability for advancing it. Thus, I strongly recommend to accept the presented thesis as a habilitation thesis.

Siegfried Benkner