



University of
Zurich ^{UZH}

Department of Informatics

University of Zürich
Department of Informatics
Binzmühlestr. 14
CH-8050 Zürich
Phone. +41 44 635 43 11
Fax +41 44 635 68 09
www.ifi.uzh.ch/dbtg

UZH, Dept. of Informatics, Binzmühlestr. 14, CH-8050 Zürich

Johann Schwabe

Prof. Dr. Michael Böhlen
Professor
Phone +41 44 635 43 33
Fax +41 44 635 68 09
boehlen@ifi.uzh.ch

Zürich, January 24, 2020

Bachelor's Thesis **Database Technology**

Topic: A GPU-enabled Single-Point Incremental Fourier Transform

The Fourier transform is an important technique in signal and image processing. In radio astronomy, the Fourier transform is applied to frequency data to compute the sky image. The common practice is to apply the Fourier transform after the observation stage has been completed and all the data has been acquired. The *Single-Point Incremental Fourier Transform* (SPIFT) [1] computes the sky image incrementally during the observation stage and continuously refines the image as new observations arrive.

The goal of this Bachelor's thesis is to leverage the processing capabilities of Graphical Processing Units (GPUs) to speed up the Single-Point Incremental Fourier Transform. GPUs are specialized processors that are used for computation-heavy and parallel computations. They are suitable for data-parallel processing tasks where data elements can be mapped to parallel processing threads.

- **Task 1 - CUDA-based implementation of matrix-vector product**

Study the architecture of heterogeneous memory hierarchies with GPUs. Familiarize yourself with the programming model and the API of CUDA for parallel code execution on GPUs [2]. Use C++ and CUDA to implement matrix-vector products on GPUs. Evaluate the performance of the matrix-vector product computation on architectures with and without GPUs. Describe your solution and empirical results in a short document.

- **Task 2 - Implementation of SPIFT algorithm with CUDA**

Study the SPIFT algorithm [1] and implement it with CUDA. Start out with a naive CUDA implementation of SPIFT that serves as a baseline. Study GPU optimization strategies [4] and improve the performance of your baseline implementation.



- **Task 3 - Integrate the CUDA implementation of SPIFT into Apache Flink**

Familiarize yourself with the architecture of the Apache Flink stream processing platform and the basic concepts in stream processing [3]. Use JCUDA [5], which are JAVA bindings for CUDA libraries, to integrate the SPIFT CUDA implementation into an Apache Flink streaming pipeline that computes the Single-Point Incremental Fourier Transform over streams.

- **Task 4 - Experimental Evaluation**

Analyze the runtime performance of your solution empirically. Run experiments on the IfI Minion cluster that compare the performance of the SPIFT streaming pipeline for GPU and non-GPU architectures.

- **Task 5 - Write a Bachelor's thesis**

Write the Bachelor thesis that describes your solution, analyzes it analytically, and evaluates it empirically. Give a 20 minutes presentation at a DBTG meeting.

References

1. M. Saad, A. Bernstein, M. H. Böhlen, D. Dell'Aglio, Single Point Incremental Fourier Transform in Apache Flink (working draft), 2020.
2. NVIDIA Corporation, CUDA C Programming Guide, https://docs.nvidia.com/cuda/pdf/CUDA_C_Programming_Guide.pdf
3. P. Carbone, A. Katsifodimos, S. Ewen, V. Markl, S. Haridi, and K. Tzoumas. Apache flink: Stream and batch processing in a single engine. Bulletin of the IEEE Computer Society Technical Committee on Data Engineering, 36(4), 2015.
4. K. Swirydowicz, Basic GPU optimization strategies, 2018, <https://www.paranumal.com/single-post/2018/02/26/Basic-GPU-optimization-strategies>
5. <http://www.jcuda.org>

Supervisor: Muhammad Saad (saad@ifi.uzh.ch)

Start Date: January 26, 2020

End Date: June 26, 2020

Presentation Date: TBD

University of Zurich
Department of Informatics

Prof. Dr. Michael Böhlen
Professor