CONTACT NAME
Hermann Furntratt: hermann.fuerntratt@joanneum.at

JOANNEUM RESEARCH DIGITAL

GPU-accelerated SIFT descriptor matching

Hermann Furntratt, Jakub Rosner, Harald Stiegler, Hannes Fassold

JOANNEUM RESEARCH
Forschungsgesellschaft mbH

DIGITAL
Institute for Information and
Communication Technologies

Steyrergasse 17
8010 Graz, Austria

Phone +43 316 876 -1199
Fax +43 316 876-1191

hermann.fuerntratt@joanneum.at
www.joanneum.at/digital

Introduction

Brand visibility in broadcast content is an important information for companies in order to steer their advertising plan. Furthermore, automatic detection of logo occurrences (size, duration) within a video is desired. Therefore we use SIFT descriptors to detect such logos and have implemented parts of the workflow on GPU.

SIFT descriptor matching algorithm

- Given N templates with K descriptors, and the current frame with M descriptors
- For each template descriptor, find best matching frame descriptor
- Compute the distance matrix between template descriptors (rows) and frame descriptors (columns)
- For each row in distance matrix, find the index with the minimum distance (or multiple indices if the L best matches are requested)

GPU implementation

- Handles different descriptor sizes
- Distance matrix calculation
  - One thread block (32x32 threads) calculates the distance values for a 32x32 block of the matrix
  - Buffering of parts of template and frame descriptors in shared memory during distance calculation (resembles strategy for matrix multiplication)
- Calculation of best match(es)
  - Each thread handles one row and finds the minimum element there
  - Templatized implementation for finding the L best matches, specialized implementation for L = 1 (significantly faster than CUDPP and Thrust routines for segmented scan)
  - Reorganization of matrix sub-block in shared memory for achieving perfectly coalesced accesses to global memory

Evaluation

- Evaluation on a Windows 7 system with a Xeon 2.67 GHz Quadcore CPU and a Tesla K20 GPU
- CPU implementation of SIFT descriptor matching uses FLANN (fast approximate nearest neighbor) algorithm from OpenCV library
- Speedup factor of 6 – 10 of GPU implementation compared to multi-threaded CPU implementation

Conclusion

- GPU implementation significantly faster than CPU implementation, even without employing any acceleration data structure
- GPU-accelerated matching integrated in BrandDetector application for automatic brand monitoring in broadcast content

http://www.joanneum.at/en/digital/products-solutions/branddetector.html

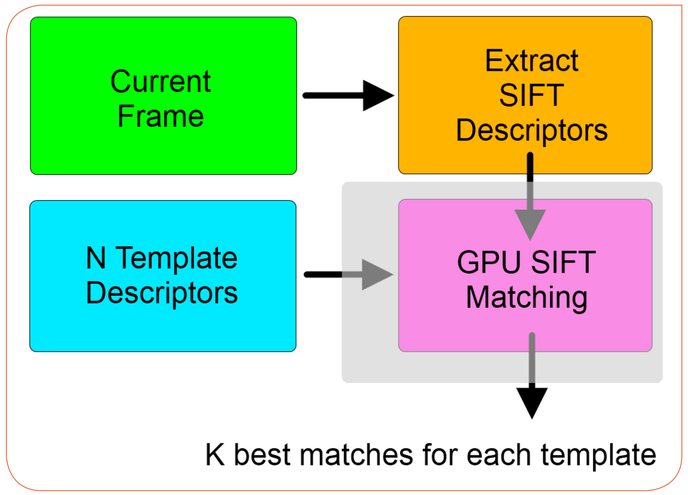

Figure 1: Schematic workflow of the algorithm

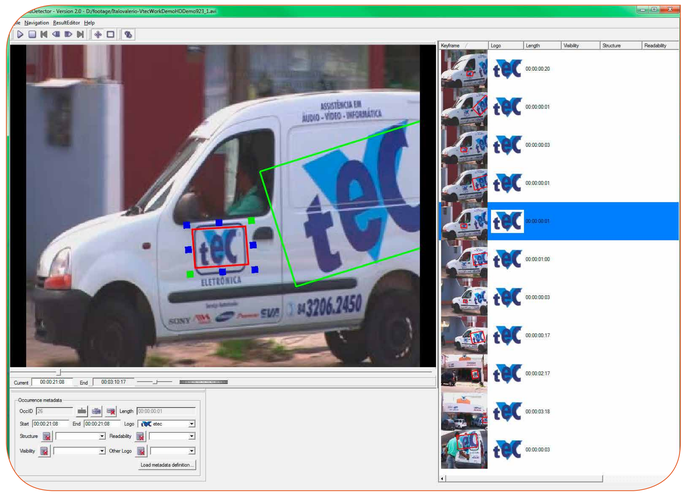

Figure 2: BrandDetector application

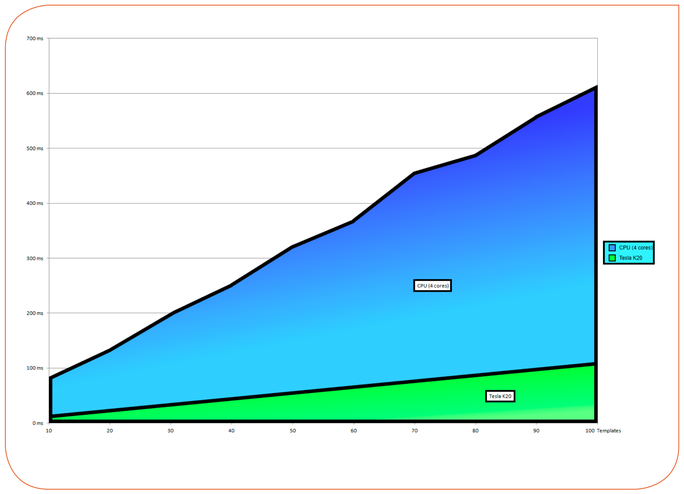

Figure 3: Runtime comparison of GPU and multi-threaded CPU implementation, with 500 descriptors per template and 1500 descriptors per frame

SEVENTH FRAMEWORK PROGRAMME

TOSCA-MP

FFG

KIRAS Sicherheitsforschung

This work was supported by the European Commission under the contract FP7-287532 „TOSCA-MP“ http://www.tosca-mp.eu and the Austrian Research Promotion Agency (FFG) under the contract 836278 “QuOIMA”

Jakub Rosner is a scholarship holder in project SWIFT POKL.08.02.01-24-005/10 and in project UDA-POKL.04.01.01-00-106/09 both supported by the European Union from the European Social Fund.