

# Fast GPU-based KLT feature point tracking using CUDA

Hannes Fassold, Jakub Rosner, Peter Schallauer and Werner Bailer

## Introduction

- Automatic selection and tracking of feature points is a basic task for many CV algorithms (e.g. structure from motion, object tracking)
- Very popular due to good performance: KLT algorithm [Kanade, Lucas and Tomasi]
- Optimized OpenCV implementation available, but is not realtime capable (> 30 fps) for High Definition (HD) video
- GPUs provide a lot of computation power and can be used for general purpose computation (CUDA)
- Goal: Port KLT feature point tracker to CUDA to achieve realtime performance for HD video

## Feature point selection

- Algorithm
  - Calculate cornerness for each pixel
  - Enforce min. quality (5–10 % of max. cornerness) and do non-maxima suppression
  - Enforce minimum distance between all feature points
- First two steps are done on the GPU
- Step 3 is inherently serial, therefore done on CPU
- Alternative method for step 3 was implemented (uses a mask image, much faster for many feature points)

## Feature point tracking

- Algorithm
  - Calculate image pyramids
  - For each feature on every pyramid level (from coarse to fine), calculate the optical flow iteratively, using Gauss-Newton method
- All steps are done on the GPU (each feature point is one thread)
- For a low number of feature points (< 1000) modern GPUs are under-utilized

## Evaluation

- Tracker quality (subjective assessment)
  - Most of the feature points and their optical flow are the same for CPU and GPU implementation
  - Differences mainly for feature points which seem to be not correctly tracked by both implementations
- Tracker runtime
  - CPU (OpenCV) routine uses IPP & OpenMP internally, runs on Intel Xeon QuadCore 2.4 Ghz
  - GPU routine runs on Geforce GX280
  - Achieves overall **5–10 times speedup**
  - Higher speedup for larger images and more feature points

## Conclusion

- Goal of HD realtime tracking achieved
- Key for successfully porting an algorithm to the GPU
  - Parallelization of the algorithm into a large number of loosely coupled threads possible
  - Knowledge of the GPU architecture (e.g. the different memory types and their properties)



Figure 1: KLT feature points (green: GPU implementation, red: CPU implementation, yellow: both)

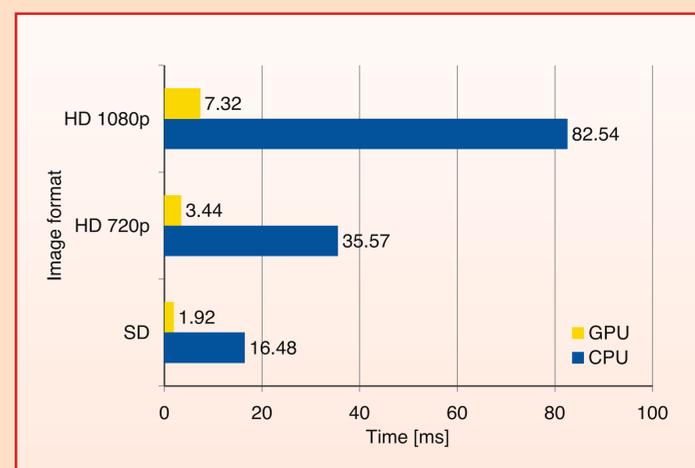


Figure 2: Runtime for feature point selection for different image sizes (without minimum distance enforcement)

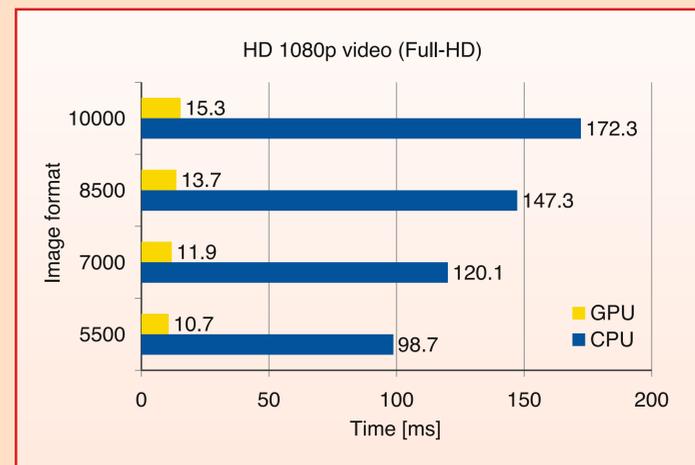


Figure 3: Runtime for feature point tracking for Full-HD images for different numbers of feature points

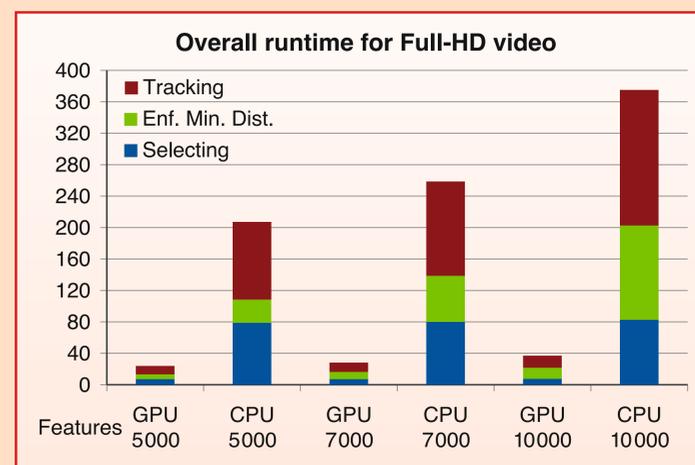


Figure 4: Overall runtime of the KLT feature point tracker for Full-HD (1920x1080) video