




Robert Hegner

Graduate Candidate	Robert Hegner
Examiner	Prof. Dr. Guido Schuster
Co-Examiner	Assoc. Prof. Dr. Ivar Austvoll, Assoc. Prof. Dr. Tom Ryen
Master Research Unit	Sensor, Actuator and Communication Systems
Project Partner	University of Stavanger, Norway


Efficient Implementation and Evaluation of Methods for the Estimation of Motion in Image Sequences



Introduction: Optical flow estimation (the estimation of the apparent motion of objects in an image sequence) is used in many applications, including video compression, object detection and tracking, robot navigation, and so on. This project focussed on one specific optical flow estimation algorithm, which uses directional filters and an AM-FM demodulation algorithm to estimate the velocities.

Objective: The main goals of this project were

- implementing the algorithm in CUDA (parallel computing architecture developed by NVIDIA) to make use of the huge parallel computing power of modern GPUs (graphic processing units).
- extending the algorithm to a multiresolution scheme to allow the estimation of higher speeds (Pyramid Algorithm).
- integrating the algorithms into an (existing) Matlab GUI which allows comparison of the new algorithm with other optical flow estimation algorithms.



Result: The speedup of the CUDA implementation (running on a GeForce GTX 260 with 216 parallel cores) compared to an existing Matlab implementation (running on an Intel Core 2 Quad 2.4GHz) is several thousand (depending on the dimensions of the image sequence). For most of the image sequences used for evaluation, the accuracy of the Pyramid Algorithm is better than or comparable to the accuracy of the OpenCV implementation of the famous Lucas-Kanade algorithm. Due to the large spatial support of the directional filters, the algorithm has some problems handling motion discontinuities, particularly in the border regions of an image sequence. The estimation of high speeds was achieved by computing a rough estimate of the (high) speeds in a downsampled image sequence. The motion in the full-resolution image sequence can then be compensated before estimating the speeds on the full-resolution image. These new estimates of the compensated speeds can finally be used to refine the first rough estimates. This procedure can be applied recursively over several levels (Pyramid Algorithm). One problem of the pyramid implementation of the algorithm is that unreliable estimates from downsampled levels can have a negative impact on the estimations at the full-resolution level.

