Real-time direct georeferencing of images from airborne line scan cameras

Trym V. Haavardsholm, Ingebrigt Winjum, Thomas Opsahl

FFI, P.O. Box 25, NO-2027 Kjeller, Norway
trym.haavardsholm@ffi.no

Introduction

The Norwegian Defense Research Establishment (FFI) is developing a technology demonstrator for airborne real-time hyperspectral target detection. The system includes two nadir-pointing line scan cameras with a 17º field of view:

- Hyperspectral: 1600 pixels, 100 Hz (0.16 Mpixel/s)
- Panchromatic: 8192 pixels, 1000 Hz (8.19 Mpixel/s)

The line scanned images are georeferenced in real-time by intersecting rays cast from the cameras with a 3D model of the terrain underneath. The georeferenced images may then easily be ortho-rectified (e.g by using texture mapping in OpenGL) and overlaid digital maps. This task is very well suited for implementation on a GPU as it is highly parallelizable.

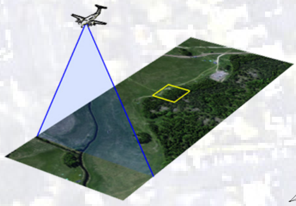

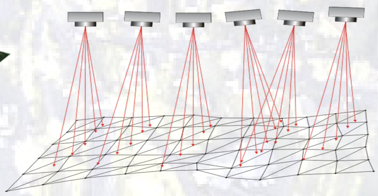

Georeferencing

In order to perform the georeferencing, we have to know the location and orientation of the cameras for each frame (which for line scan cameras is one single line). We obtain this information by combining Kalman filtered data from an onboard GPS and Inertial Measurement Unit (IMU) with camera models that describe the line of sight of each sensor element relative to the position and orientation of the IMU.

The ground surface is described by a digital elevation model (DEM). The DEM used in the hyperspectral detection system is typically based on the DTED level 2 format, in which the ground resolution is approximately 30 meters.

The points where each ray intersects the DEM gives us the geographical position for each corresponding pixel. To avoid having to check each ray with each triangle in the DEM, we apply an accelerating structure based on a 2D uniform grid in the local tangent plane. The DEM is divided into smaller volume elements, referred to as voxels, that are aligned with the vertices in the DEM, and which may be traversed very efficiently. We only check for ray-triangle intersection in the voxels that each ray enters.

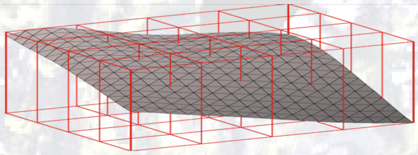

CUDA implementation

The CUDA implementation of the georeferencing method is split into two kernels. A block of navigation data for a chosen amount of frames adapted to an acceptable latency is loaded onto the device, and the first kernel finds the position and direction of each ray using the camera model and navigation data. The second kernel does the ray casting, where each thread in the kernel follows one individual ray until it intersects the DEM, which resides in texture memory. The resulting points of intersection is then finally returned to the host.

Performance results

In this test we have constructed flat "worst case" DEMs where we have manipulated the accelerating structure in such a way that we ensure that all rays start in the same voxel so that they have to traverse the maximum distance before reaching the correct voxel (the ceiling of the voxels is set to aircraft altitude). The simulated flight is a 60 s straight and level trajectory at 1000 m above the ground. The performance was tested on a GeForce 8800 graphics card.

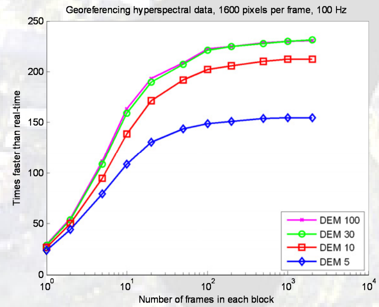

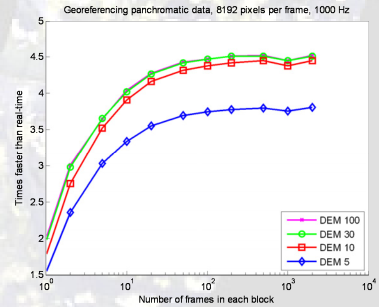

The figures above show that we for full resolution frames are able georeference faster than real-time. Optimal performance is reached for around 1 s large blocks, with a processing rate of about 35 Mpixel/s on a DEM corresponding to DTED level 2.

Since the ground resolution of the cameras is much higher than for the DEM, we do not have to ray cast each pixel. If we instead ray cast with twice the resolution of the DEM for each frame, we can interpolate the georeference along the line for the rest of the pixels. With a 17º field of view, 1000 m elevation and a DEM with 30 m resolution this corresponds to ray casting only 20 pixels per frame.

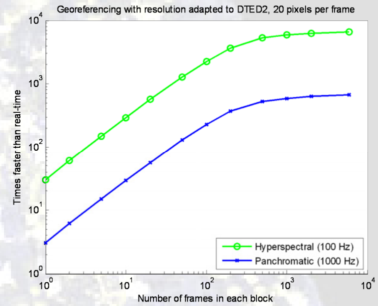

The figure above shows the performance when ray casting only 20 pixels per frame. For 1 s blocks, the processing rate for the panchromatic camera is about 600x faster than real-time. This is more than 100x faster than an optimized matlab implementation of the same method.

The images above show an example of applying georeferencing in rectifying a line scanned image severely disturbed by aircraft movement. Interpolation along the line is performed "automatically" in the texture mapping, which is "instantaneous".

FFI Forsvarets forskningsinstitutt

Norwegian Defence Research Establishment