GPU IMPLEMENTATION OF A POISSON MAP-MRF SEGMENTATION ALGORITHM FOR 3D TRACKING OF LOW LIGHT LEVEL IMAGERY

Danny Crookes, Damian McCaughey, Hugh Gribben and Paul Miller

The Institute of Electronics, Communications and Information Technology (ECIT)
Queen's University Belfast, Belfast, UK. BT3 9DT
{d.crookes, dmccaughey03, hgribben01, p.miller}@qub.ac.uk Tel. 028 90971724

1 Introduction

In this work, we report on the implementation, on an NVIDIA GeForce 8800GTX, of the segmentation phase of a tracking system for 4D low-light level imagery obtained from recently developed Electron Multiplying Charge Coupled Device (EMCCD) sensors.

The high sensitivity of these sensors provides us with the ability to acquire multidimensional datasets at high temporal sampling rates under low-light level conditions. This allows, for the first time, the study of the dynamics of biomedical bodies such as cells, and their intra-cellular components. Novel statistical distributions are present within the data due to both the initial low-light levels at which the data is acquired, and the avalanche multiplication to which the data is subjected via the EMCCD. The size of these datasets, coupled with the statistical distributions, present a new challenge within the area of biomedical image processing. Related work in this area has included the segmentation of medical images using an automated volumetric segmentation system [1], and the segmentation of 3D MRI brain images [2]. While these works have adopted a Gaussian MAP-MRF approach, we instead assume low-light level Poisson statistics for the optimisation of the novel L3CCD data; and apply the resulting technique.

2 Segmentation

In many segmentation applications, an initial baseline estimate is obtained using Mixture Modelling, assuming a mixture density of Gaussian distributions. At low light levels (up to approx. 50 photons per pixel per integration time), however, we know that photon counts are distributed according to the Poisson distribution. Using this knowledge, we model the intensity distribution of the EMCCD dataset as a mixture of Poisson distributions. The algorithm used in practice to find the mixture of distributions that best model the dataset is the EM algorithm, first introduced by Dempster. This is an iterative algorithm, which estimates the parameters via a Maximum Likelihood (ML) criterion. The mixture density returned can then be used to associate pixel observations with a Poisson density k using a simple ML estimation. This estimate is then refined to consider spatial correlations using MRF Modelling. In practice, we use of the Iterated Conditional Modes (ICM) algorithm, originally introduced by Besag. This is an iterative algorithm that begins with the observed scene Y, and the initial estimate of the true scene X from mixture modelling. By considering each pixel site in turn, it iteratively provides a new estimate of the true scene, until convergence is reached.

The technique has been applied to a 4D EMCCD dataset of dimension 208*235*100, with 30 time frames. The scene consists of Green Fluorescent Protein labelled telomeres in a cell nucleus. These are a component of cells which act as a buffer at the end of chromosomes to prevent the loss of genetic information needed to sustain its activities; with implications in anti-aging and anti-cancer therapy. Fig. 1 shows the qualitative segmentation of 1 3D volume. The results clearly show that the Poisson MAP-MRF technique provides a result closer to that obtained via the manual segmentation technique than the traditional Gaussian MAP-MRF technique. Quantitative results back up this finding. The reason is that Poisson statistics provide a better model of the true distributions present in the L3CCD datasets.

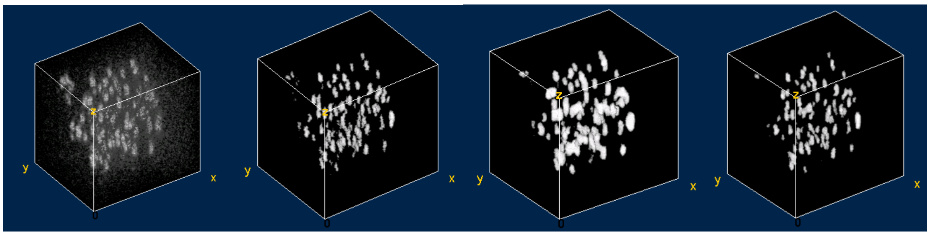

(a) (b) (c) (d)

Figure 1: *Segmentation of L3CCD Dataset* (a) Input Dataset
(b) Output from Manual Segmentation (c) Output from Gaussian MAP-MRF Technique
(d) Output from Poisson MAP-MRF Technique.

3 GPU Implementation of the Segmentation Algorithm

Our goal is to accelerate the complete tracking process using a low cost accelerator in the form of an NVIDIA GeForce 8800 GTX GPU card, using the Compute Unified Device Architecture (CUDA). The starting point of the implementation was:

(i) A mathematical specification of the segmentation algorithm
(ii) A Matlab implementation of the algorithm.

The original intention was to use the Matlab implementation as the model for the parallelization on the GPU. It was hoped that this would lead to a more general-purpose Matlab-based development environment which would be very advantageous for a range of applications. However, one clear result from the GPU implementation exercise was that the path from the mathematical specification the final GPU implementation is much clearer and more direct than starting from the Matlab version. It was found that the Matlab implementation of the algorithm had been forced to introduce many extra code variations and additions which a software tool would find difficult or impossible to disentangle. The much more direct correspondence between the mathematical algorithm specification and the GPU implementation has led us to propose an algebraic programming notation for 4D imaging. The Matlab implementation is being used to validate/evaluate the intermediate and final results of the GPU implementation, including the issue of floating point precision.

At the time of writing, the complete segmentation has been coded in CUDA and run in emulation mode. Reusable strategies for processing 3D and 2D datasets have been developed. Experimentation is underway to investigate the best partitioning strategies. The most heavily used operations tend to be neighbourhood (3x3x3) and image reduction (2D and 3D) operators. These are very amenable to efficient parallelization, given proper exploitation of the on-chip shared memory. Timings will be presented to indicate speed-up.

4 References

[1] D. Montgomery, A. Amira, & F. Murtagh, “An automated volumetric segmentation system combining multi-scale and statistical reasoning,” *Proceedings of the IEEE International Symposium on Circuits and Systems, Kobe, Japan*, 4, pp. 3789-3792, May 2005.

[2] Z. Peng, W. Wee & J. Lee, “MR brain imaging segmentation based on spatial Gaussian Mixture Model and Markov Random Field,” *IEEE International Conference on Image Processing*, 2005.