GPU-Accelerated Data Reconstruction on FastSPECT II

Introduction

FastSPECT II is a SPECT imaging system developed at CGRI for small-animal imaging. It consists of sixteen stationary (i.e., non-moving) modular gamma-ray cameras. Each camera has an input face measuring about 120×120 mm^2. Dedicate circuitry interfaces the cameras to a computer station for data acquisition of the camera's photomultiplier tubes (PMT) signals. The size of the field of view is approximatively $42 \times 42 \times 54$ mm^3.

GPU Implementation

GPU code was developed to speed up the pre-processing of the acquired data and the actual reconstruction via the expectation maximization maximum likelihood (MLEM) algorithm. Pre-processing of the data consisted in 2D maximum likelihood estimation of position of interaction for each detected gamma-ray photon.

Our algorithm for estimation of position of interaction uses a contracting-grid approach. We took full advantage of the texture mapping unit of the GPU devices to implement on-the-fly interpolation of the PMT mean detector response function (MDRF) of each camera. This interpolation step is automatically performed by the GPU hardware, resulting in code that is extremely fast, compact, and easy to understand and to maintain.

Our reconstruction code implements the MLEM algorithm in the following form:

$$\hat{f}_n^{(k+1)} = \hat{f}_n^{(k)} \frac{1}{s_n} \sum_{m=1}^{M} \frac{g_m h_{m,n}}{\sum_{n'=1}^{N} h_{m,n'} \hat{f}_{n'}^{(k)}},$$

in which $\hat{f}_n^{(k)}$ is the estimated activity for the n-th voxel at the k-th iteration, s_n is the sensitivity for voxel n, g_m is the bin count for the m-th bin (hence m both encodes the camera and the pixel on the camera's face), and $h_{m,n}$ is the (m,n)-th component of the system matrix. A plot of the sensitivity is shown in Figure 1. Our parallel implementation of the MLEM algorithm takes advantage of the capabilities of modern GPU devices. For example, dynamic parallelism was used to avoid useless calculations for the case $h_{m,n} = 0$. Our code can be adapted to handle list-mode MLEM reconstruction as well.

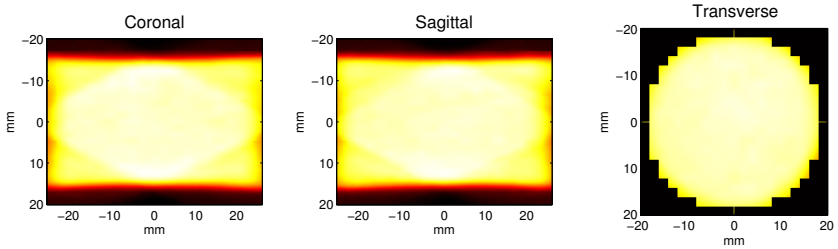

Figure 1: Plot of the sensitivity across three planes, showing almost uniform sensitivity over a large portion of the system's field of view. The volume for which the sensitivity exceeds 5% its maximum value is about 56.60 cm^3

Reconstruction Results

We performed a simple simulation study to assess the correctness and performance of our algorithm. Starting from a simulated mouse, we generated sample PMT signals which were used to estimate 2D event position for each camera. These data were reconstructed with our GPU-accelerated implementation of the MLEM algorithm. The data were reconstructed on a $208 \times 161 \times 161$ voxel grid. The size of each voxel was $0.25 \times 0.25 \times 0.25$ mm^3. We ran 30 iterations of the MLEM algorithm. With an 8-Tesla K40 GPU machine, the total processing time was less than six minutes. We finally generated a volume rendering (shown on the right) of the reconstructed data.

Click on image to play clip