Published in final edited form as:

IEEE Nucl Sci Symp Conf Rec (1997). 2009 October 24; 2009: 4072. doi:10.1109/NSSMIC.2009.5402392.

Maximum Likelihood Event Estimation and List-mode Image Reconstruction on GPU Hardware

Luca Caucci, Lars R. Furenlid[Member, IEEE], and Harrison H. Barrett[Fellow, IEEE]

The authors are with the College of Optical Sciences, University of Arizona, 1630 E. University Blvd., Tucson, Arizona 85721 and also with the Center for Gamma-Ray Imaging, University of Arizona, 1609 N. Warren Ave., Tucson, Arizona 85719.

Abstract

The scintillation detectors commonly used in SPECT and PET imaging and in Compton cameras require estimation of the position and energy of each gamma ray interaction. Ideally, this process would yield images with no spatial distortion and the best possible spatial resolution. In addition, especially for Compton cameras, the computation must yield the best possible estimate of the energy of each interacting gamma ray. These goals can be achieved by use of maximum-likelihood (ML) estimation of the event parameters, but in the past the search for an ML estimate has not been computationally feasible. Now, however, graphics processing units (GPUs) make it possible to produce optimal, real-time estimates of position and energy, even from scintillation cameras with a large number of photodetectors. In addition, the mathematical properties of ML estimates make them very attractive for use as list entries in list-mode ML image reconstruction. This two-step ML process—using ML estimation once to get the list data and again to reconstruct the object—allows accurate modeling of the detector blur and, potentially, considerable improvement in reconstructed spatial resolution.

I. Introduction

In tomographic imaging of gamma-ray-emitting radiotracers with scintillation detectors, each gamma-ray photon that interacts with a detector produces a set of signals on photomultipliers (PMTs) or other photodetectors. From these signals the position and energy of each gamma-ray interaction are determined by some algorithm, known as the *event-estimation process* and then stored, either as a list of position and energy for each interaction events or as a binned array where each bin entry represents the number of events with position and energy within some interval. We refer to the first alternative as *list-mode* data storage and the second as *binned-mode* storage. In either case, we refer to the estimated position and energy as *attributes* of the event. Image reconstruction can then be performed either directly from the attribute list or from the binned data.

The current trend towards measuring more and more attributes of each event and building larger and larger 2-D and 3-D detectors causes the memory requirements for binned data to grow geometrically, making binned-data approaches challenging. For example, a Compton camera typically yields six attributes (coordinates x and y and energy E for each of two interactions) for every measured coincidence event. Similarly, a PET system using scintillation cameras with depth-of-interaction capability can yield up to ten attributes (x , y , z and E plus time of arrival on each detector) for each event.

Fortunately, list-mode (LM) maximum-likelihood (ML) expectation-maximization (EM) image reconstruction algorithms are well developed and often used for PET, Compton cameras and other applications in nuclear medicine. There are, however, two problems with

LMMLEM algorithms: they are computationally demanding, and for best performance they require an accurate probability model for the attributes. It is the purpose of this paper to show that both of these problems can be addressed effectively with graphics processing units (GPUs).

In Sec. II we present the basics of LMMLEM reconstruction, and in Sec. III we show in detail how to develop the probability model for a Compton camera based on thick, monolithic scintillation crystals; a similar model applies to monolithic-crystal PET systems. In Sec. IV we discuss the advantages of using ML methods to estimate the attributes used in the reconstruction algorithm, and we show how an important asymptotic property of ML estimates can be used to simplify the probability model. Thus we recommend a dual use of ML theory, first in the event-estimation step to give values for the attributes of each event and then in the reconstruction step to estimate the object distribution from the attribute list.

In Sec. V we discuss the computational requirements of this approach and present some preliminary results obtained by running a 3-D ML position estimation algorithm on a GPU supercomputer equipped with four NVIDIA GeForce 9800 GX2 computing processors. Our GPU implementation resulted in a $250\times$ speedup with respect to an implementation of the same algorithm on a conventional cluster machine. These promising results show the enormous potential of commodity GPU hardware for medical image reconstruction.

II. MLEM and LMMLEM algorithms

Given the data vector $\mathbf{g} = \{g_1, \dots, g_M\}$, where g_m represents the number of events stored in the m^{th} bin, the goal of an image reconstruction algorithm is to produce a vector $\hat{\mathbf{f}} = \{\hat{f}_1, \dots, \hat{f}_N\}$ approximating the radiotracer distribution in the object [1], [2].

A popular iterative reconstruction algorithm in the medical literature is the maximum-likelihood (ML) expectation-maximization (EM) algorithm [3], [4]. If the data vector $\mathbf{g}$ follows Poisson statistics, the MLEM algorithm takes the form:

$$\hat{\mathbf{f}}_n^{(k+1)} = \hat{\mathbf{f}}_n^{(k)} \left\{ \frac{1}{\sum_{m=1}^M H_{mn}} \sum_{m=1}^M \frac{g_m H_{mn}}{[H \hat{\mathbf{f}}^{(k)}]_m} \right\}, \quad (1)$$

where the matrix H (with components H_{mn}) represents a discretized version of the system operator [2]. This way of organizing the sensitivity data requires such a large amount of storage and computational capability [5] to make it impossible to record more than just a few attributes for each event.

Alternatively, the MLEM algorithm is applied to raw event estimates recorded in list mode [5]–[9], resulting in the list-mode (LM) MLEM algorithm, which we now detail for specific configurations of a Compton camera and a PET imager.

In the Compton camera of Fig. 1, two position-and-energy-sensitive thick detectors image the radiotracer distribution. The two detectors allow high sensitivity because collimation is accomplished in software rather than by placing physical collimators in front of the detectors. For the j^{th} event (with $j=1, \dots, J$), we can consider two 3-D positions of interaction estimates, $\hat{\mathbf{r}}_{j1}$ and $\hat{\mathbf{r}}_{j2}$ along with the estimate $\hat{\beta}_j$ of the cosine of the scatter angle

$\widehat{\alpha}_j$ (calculated from estimates $\widehat{E}_{j1}$ and $\widehat{E}_{j2}$ of the deposited energies). Given $\widehat{r}_{j1}$, $\widehat{r}_{j2}$, and $\widehat{\beta}_j$ we define the j^{th} attribute vector as $A_j = (\widehat{r}_{j1}, \widehat{\beta}_j, \widehat{r}_{j2})$.

In the PET setup shown in Fig. 2, two detectors are placed on opposite sides of the object. For the j^{th} event, two 3-D positions of interaction estimates are computed. Thus, the j^{th} attribute vector is $A_j = (\widehat{r}_{j1}, \widehat{r}_{j2})$. We will assume that coincidence windowing is used to pair events occurring in different detectors. We will also assume that, for each event, both photons have been collected. Energy estimation is not used as the vast majority of electron-positron decays result in two 511 keV gamma photons being emitted at an angle $\alpha \approx \pi$. For the sake of simplicity, we will assume $\alpha = \pi$.

In both cases, we can define the list $A = \{A_1, \dots, A_J\}$ of the J attribute vectors. Assume we subdivide the field of view in N voxels, centered at 3-D locations r_n , for $n = 1, \dots, N$. Let $f = \{f_1, \dots, f_N\}$ be the radiotracer distribution we want to estimate, where f_n is the number of photons emitted from voxel centered at r_n . Furthermore, let S_n be the probability that a photon emitted from voxel at r_n is detected. Thus, S_n is referred to as the sensitivity for the n^{th} voxel [7]. The probability that a detected photon was emitted from voxel at r_n is then written as:

$$\Pr(r_n | f) = \frac{S_n f_n}{\sum_{n'=1}^N S_{n'} f_{n'}}.$$

If we denote with $\text{pr}(A_j | r_n)$ the probability density function (PDF) for measuring the attribute vector A_j given that a photon is emitted from the n^{th} voxel, then:

$$\text{pr}(A_j | f) = \sum_{n=1}^N \text{pr}(A_j | r_n) \Pr(r_n | f).$$

The likelihood of f given the list A of J statistically independent attribute events is:

$$\mathcal{L}(f; A) = \prod_{j=1}^J \text{pr}(A_j | f).$$

As [3] and [7] show, the iterative expression for LMMLEM reconstruction is [cf. (1)]:

$$\widehat{f}_n^{(k+1)} = \widehat{f}_n^{(k)} \left\{ \frac{1}{T} \sum_{j=1}^J \frac{\text{pr}(A_j | r_n)}{\sum_{n'=1}^N \text{pr}(A_j | r_{n'}) S_{n'} \widehat{f}_{n'}^{(k)}} \right\}, \quad (2)$$

in which T is the total exposure time. The expression above emphasizes that the quantity $\text{pr}(A_j | r_n)$ plays a central role in the implementation of LMMLEM reconstruction algorithms.

It can be shown [7] that if the first estimate, $\widehat{f}^{(0)}$, satisfies positivity constraints, then all the subsequent estimates obtained using (2) satisfy positivity constraints as well and, under reasonable conditions [7], the algorithm has the desired property of converging to the global

maximum of the likelihood function $\mathcal{L}(f; A)$. Global convergence is proved by showing that the Hessian (with respect to f) of $\mathcal{L}(f; A)$ is negative definite [7].

III. Derivation of the PDF $\text{pr}(A_j|r_n)$

In this section, we provide a derivation of the expression for the probability density function $\text{pr}(A_j|r_n)$ for the Compton camera setup of Fig. 1. The same expression can easily be modified to accommodate the PET setup shown in Fig. 2.

Recall that $A_j = (\hat{r}_{j1}, \hat{\beta}_j, \hat{r}_{j2})$. By Bayes' theorem,

$$\text{pr}(A_j|r_n) = \int_{\text{det.1}} \int_{-1}^1 \int_{\text{det.2}} \text{pr}(r_{j1}, \beta_j, r_{j2}|r_n) \times \text{pr}(\hat{r}_{j1}, \hat{\beta}_j, \hat{r}_{j2}|r_{j1}, \beta_j, r_{j2}, r_n) d^3 r_{j1} d\beta_j d^3 r_{j2}, \quad (3)$$

where vectors r_{j1} and r_{j2} denote actual (i.e. not estimates) 3-D interaction within the detector. Similarly, β_j is the cosine of the actual scattering angle. Notations $\int_{\text{det.1}}$ and $\int_{\text{det.2}}$ denote integration over detectors 1 and 2, respectively.

Notice first that in $\text{pr}(\hat{r}_{j1}, \hat{\beta}_j, \hat{r}_{j2}|r_{j1}, \beta_j, r_{j2}, r_n)$ there is no need to condition on r_n because estimates depend only on the interaction locations and the scattering angle. Furthermore, estimates of $\hat{r}_{j1}$ and $\hat{r}_{j2}$ are statistically independent because they use two different sets of detector outputs. Thus:

$$\text{pr}(\hat{r}_{j1}, \hat{\beta}_j, \hat{r}_{j2}|r_{j1}, \beta_j, r_{j2}, r_n) = \text{pr}(\hat{r}_{j1}, \hat{\beta}_j|r_{j1}, \beta_j) \text{pr}(\hat{r}_{j2}|r_{j2}).$$

If the detector is made of homogeneous material, then r_{j1} and β_j are statistically independent. However, r_{j2} depends on the location of the first interaction and the scattering angle. These observations lead to:

$$\text{pr}(r_{j1}, \beta_j, r_{j2}|r_n) = \text{pr}(r_{j1}|r_n) \text{pr}(\beta_j) \text{pr}(r_{j2}|r_{j1}, \beta_j, r_n).$$

Thus, the probability density function $\text{pr}(A_j|r_n)$ in (3) assumes the form:

$$\text{pr}(A_j|r_n) = \int_{\text{det.1}} \int_{-1}^1 \int_{\text{det.2}} \text{pr}(r_{j1}|r_n) \text{pr}(\beta_j) \text{pr}(r_{j2}|r_{j1}, \beta_j, r_n) \times \text{pr}(\hat{r}_{j1}, \hat{\beta}_j|r_{j1}, \beta_j) \text{pr}(\hat{r}_{j2}|r_{j2}) d^3 r_{j1} d\beta_j d^3 r_{j2}. \quad (4)$$

To compute $\text{pr}(r_{j1}|r_n)$, consider a small volume element centered at r_{j1} . It is convenient to take this volume element as a cube of side ε with one face perpendicular to the line connecting r_n to r_{j1} . Then, the probability that a Compton interaction occurs in the cube is:

$$\text{Pr}(\text{interaction in cube at } r_{j1}|r_n) = \frac{\varepsilon^3 \mu_{C1}}{4\pi |r_{j1} - r_n|^2},$$

where μ_{C1} is the Compton attenuation coefficient in detector 1. From the result above, it follows:

$$\text{pr}(\mathbf{r}_{j1}|\mathbf{r}_n) = \frac{\mu_{c1}}{4\pi|\mathbf{r}_{j1} - \mathbf{r}_n|^2}. \quad (5)$$

To derive an expression for $\text{pr}(\mathbf{r}_{j2}|\mathbf{r}_{j1}, \beta_j, \mathbf{r}_n)$, consider a cube of side ε center at $\mathbf{r}_{j2}$ and having one face perpendicular to the line connecting point $\mathbf{r}_{j2}$ with $\mathbf{r}_{j1}$. Conditioning on $\mathbf{r}_{j1}$ and $\mathbf{r}_n$ fixes a direction $\mathbf{s}_{1n}^{(j)}$ from the origin of the event $\mathbf{r}_n$ to the interaction location $\mathbf{r}_{j1}$. Conditioning on β_j imposes that the scattered photon lie on a cone of infinitesimal with around the direction $\mathbf{s}_{12}^{(j)}$, defined by $\mathbf{r}_{j1}$ and $\mathbf{r}_{j2}$. By passing to the limit for $\varepsilon \rightarrow 0$, we get:

$$\text{pr}(\mathbf{r}_{j2}|\mathbf{r}_{j1}, \beta_j, \mathbf{r}_n) = \frac{\mu_{p2}}{2\pi|\mathbf{r}_{j1} - \mathbf{r}_{j2}|^2} \delta(\mathbf{s}_{1n}^{(j)} \cdot \mathbf{s}_{12}^{(j)} - \beta_j), \quad (6)$$

in which μ_{p2} is the photoelectric attenuation coefficient in detector 2, $\delta(\dots)$ denotes the delta function, and the unit vectors $\mathbf{s}_{1n}^{(j)}$ and $\mathbf{s}_{12}^{(j)}$ are given by:

$$\mathbf{s}_{1n}^{(j)} = \frac{\mathbf{r}_{j1} - \mathbf{r}_n}{|\mathbf{r}_{j1} - \mathbf{r}_n|}, \quad \mathbf{s}_{12}^{(j)} = \frac{\mathbf{r}_{j2} - \mathbf{r}_{j1}}{|\mathbf{r}_{j2} - \mathbf{r}_{j1}|}. \quad (7)$$

Substituting (5) and (6) into (4) and performing the integration over β_j gives:

$$\text{pr}(A_j|\mathbf{r}_n) = \int_{\text{det.1}} \int_{\text{det.2}} \frac{\mu_{c1} \mu_{p2}}{8\pi^2|\mathbf{r}_{j1} - \mathbf{r}_n|^2|\mathbf{r}_{j1} - \mathbf{r}_{j2}|^2} \times \text{pr}_{\beta}(\mathbf{s}_{1n}^{(j)} \cdot \mathbf{s}_{12}^{(j)}) \text{pr}(\widehat{\mathbf{r}}_{j1}, \widehat{\beta}_j|\mathbf{r}_{j1}, \mathbf{s}_{1n}^{(j)} \cdot \mathbf{s}_{12}^{(j)}) \times \text{pr}(\widehat{\mathbf{r}}_{j2}|\mathbf{r}_{j2}) d^3\mathbf{r}_{j1} d^3\mathbf{r}_{j2}.$$

So far, no approximations have been made. Assumptions that might be valid for silicon detectors are that position and energy estimates in the first detector are statistically independent, that spatial resolution does not depend of the amount of energy deposited, and, finally, that energy resolution is constant over the whole 3-D detector. With these assumptions:

$$\text{pr}(\widehat{\mathbf{r}}_{j1}, \widehat{\beta}_j|\mathbf{r}_{j1}, \beta_j) = \text{pr}(\widehat{\mathbf{r}}_{j1}|\mathbf{r}_{j1}) \text{pr}(\widehat{\beta}_j|\beta_j),$$

leading to:

$$\text{pr}(A_j|\mathbf{r}_n) = \int_{\text{det.1}} \int_{\text{det.2}} \frac{\mu_{c1} \mu_{p2}}{8\pi^2|\mathbf{r}_{j1} - \mathbf{r}_n|^2|\mathbf{r}_{j1} - \mathbf{r}_{j2}|^2} \times \text{pr}_{\beta}(\mathbf{s}_{1n}^{(j)} \cdot \mathbf{s}_{12}^{(j)}) \text{pr}(\widehat{\mathbf{r}}_{j1}|\mathbf{r}_{j1}) \text{pr}(\widehat{\beta}_j|\mathbf{s}_{1n}^{(j)} \cdot \mathbf{s}_{12}^{(j)}) \times \text{pr}(\widehat{\mathbf{r}}_{j2}|\mathbf{r}_{j2}) d^3\mathbf{r}_{j1} d^3\mathbf{r}_{j2}.$$

If we further assume that $\text{pr}(\widehat{\mathbf{r}}_{j1}|\mathbf{r}_{j1})$, $\text{pr}(\widehat{\mathbf{r}}_{j2}|\mathbf{r}_{j2})$ are sharply peaked about $\widehat{\mathbf{r}}_{j1}$ and $\widehat{\mathbf{r}}_{j2}$, respectively, then:

$$\text{pr}(A_j|\mathbf{r}_n) \approx \frac{\mu_{c1} \mu_{p2} \text{pr}_{\beta}(\mathbf{s}_{1n}^{(j)} \cdot \mathbf{s}_{12}^{(j)})}{8\pi^2|\widehat{\mathbf{r}}_{j1} - \mathbf{r}_n|^2|\widehat{\mathbf{r}}_{j1} - \widehat{\mathbf{r}}_{j2}|^2} \int_{\text{det.1}} \int_{\text{det.2}} \text{pr}(\widehat{\mathbf{r}}_{j1}|\mathbf{r}_{j1}) \text{pr}(\widehat{\beta}_j|\mathbf{s}_{1n}^{(j)} \cdot \mathbf{s}_{12}^{(j)}) \text{pr}(\widehat{\mathbf{r}}_{j2}|\mathbf{r}_{j2}) d^3\mathbf{r}_{j1} d^3\mathbf{r}_{j2}, \quad (8)$$

where [cf. (7)]:

$$\widehat{s}_{1n}^{(j)} = \frac{\widehat{r}_{j1} - r_n}{|\widehat{r}_{j1} - r_n|}, \quad \widehat{s}_{12}^{(j)} = \frac{\widehat{r}_{j2} - \widehat{r}_{j1}}{|\widehat{r}_{j2} - \widehat{r}_{j1}|}.$$

If instead of a Compton camera we have the PET system of Fig. 2, then (8) becomes:

$$\text{pr}(A_j|r_n) \approx \frac{\mu_p}{8\pi^2|\widehat{r}_{j1} - r_n|^2|\widehat{r}_{j1} - \widehat{r}_{j2}|^2} \times \int_{\text{det.1}} \int_{\text{det.2}} \text{pr}(\widehat{r}_{j1}|r_{j1}) \text{pr}(\widehat{r}_{j2}|r_{j2}) d^3r_{j1} d^3r_{j2}, \quad (8')$$

because $\beta_0 = \cos(\alpha_0) = -1$ for $\alpha_0 = \pi$. The delta function that appears in (6) and (8') is an angular delta function [2], defined as:

$$\int_{4\pi} \delta(s - s_0) t(s) d\Omega_s = t(s_0),$$

in which the integral is performed over the solid angle 4π , Ω_s is the infinitesimal solid angle element, and $t(s)$ is a test function.

The two expressions for $\text{pr}(A_j|r_n)$ derived above might require experimentally determining $\text{pr}(\beta_j)$, $\text{pr}(\widehat{\beta}_j|\beta_j)$, $\text{pr}(\widehat{r}_{j1}|r_{j1})$, and $\text{pr}(\widehat{r}_{j2}|r_{j2})$. Alternatively, assumptions can be made on the shape of those PDFs. For example, if we assume that $\text{pr}(\widehat{\beta}_j|\beta_j)$, $\text{pr}(\widehat{r}_{j1}|r_{j1})$, and $\text{pr}(\widehat{r}_{j2}|r_{j2})$ are all Gaussians, integrals in (8) and (8') can be evaluated analytically.

IV. ML estimation of position of interaction

The integrands in (8) and (8') contain quantities of the form $\text{pr}(\widehat{r}|\mathbf{r})$, which expresses how well $\mathbf{r}$ is estimated from the detector outputs. One way to calculate $\widehat{r}$ is by means of ML estimation [10], [11]. Assume that $\mathbf{G}(\mathbf{r})$ represents noisy (i.e. random) detector outputs corresponding to an event occurring at location $\mathbf{r}$ in the detector. For example, $\mathbf{G}(\mathbf{r})$ might be a vector with K random components, each of them corresponding to a PMT output in a scintillation camera. In such a case, the ML estimate of $\mathbf{r}$ can be mathematically defined as:

$$\widehat{r} = \underset{r_0}{\text{argmax}} \text{pr}[\mathbf{G}(\mathbf{r})|r_0]. \quad (9)$$

If $\widehat{r}$ satisfies the condition above then $\text{pr}(\widehat{r}|\mathbf{r})$ asymptotically approximates a multivariate normal PDF with mean $\mathbf{r}$ and covariance matrix given by the inverse of the Fisher information matrix $\mathbf{F}_{\widehat{r}|\mathbf{r}}$ for the estimation of $\mathbf{r}$ [2], [12], [13]:

$$\text{pr}(\widehat{r}|\mathbf{r}) \rightarrow \frac{\sqrt{\det(\mathbf{F}_{\widehat{r}|\mathbf{r}})}}{(2\pi)^{3/2}} \exp \left[-\frac{1}{2} (\widehat{r} - \mathbf{r}) \mathbf{F}_{\widehat{r}|\mathbf{r}} (\widehat{r} - \mathbf{r}) \right].$$

In a scintillation camera, multiple visible photons are produced for each gamma-ray event, and each visible photon leads to the generation of many photoelectrons. Thus, asymptotic properties of ML estimation hold true and the approximation above is a valid assumption.

The Fisher information matrix provides a way of quantifying an upper bound on how well information can be extracted from detector outputs [2]. If the noise in the components of

$\mathbf{G}(\mathbf{r})$ follows a Poisson distribution, the $(m, n)^{\text{th}}$ component of the Fisher information matrix is given by [2]:

$$\left[\mathbf{F}_{\hat{\mathbf{r}}|\mathbf{r}} \right]_{mm} = \sum_{k=1}^K \frac{1}{\bar{G}_k(\mathbf{r})} \frac{\partial \bar{G}_k(\mathbf{r})}{\partial x_m} \frac{\partial \bar{G}_k(\mathbf{r})}{\partial x_n} \quad (10)$$

where $\mathbf{r} = (x_1, x_2, x_3)$ and $\bar{G}_k(\mathbf{r})$ is the mean of the k^{th} detector output for an interaction event occurring at $\mathbf{r}$. The derivatives that appear in (10) can be evaluated analytically by first interpolating calibration data with spline functions [14]. Also, spline interpolation can be used to deliver continuous values for $\hat{\mathbf{r}}$ when (9) is used.

To summarize, the overall reconstruction algorithm that gives an estimate $\hat{\mathbf{f}}$ of the radiotracer distribution consists of two separate ML-estimation steps:

1. For each event, use (9) to estimate the position of interaction;
2. Given the list of estimates previously calculated, iterate (2) to get $\hat{\mathbf{f}}$.

It is worth noting that ML methods provide a rigorous and—in some sense—optimal way to extract information from detector data. The optimality of ML estimation is used in the two steps above, giving an estimate $\hat{\mathbf{f}}$ that most likely produced the collected detector data.

V. Hardware configuration and results

For this research, we used a GPU supercomputer, which we assembled using commodity hardware. The hardware configuration included four NVIDIA GeForce 9800 GX2 graphics cards. The software configuration included the Linux operating system and the NVIDIA's CUDA software development kit. For a total cost of less than \$3000, our GPU supercomputer is capable of about 4 TFLOPS.

The theory outlined in the previous sections requires the ability of performing 3-D ML estimation of event position from PMT data. A 3-D contracting-grid [10] algorithm was implemented for a CUDA-capable GPU device and run on our GPU machine. For the ℓ^{th} event occurring on a detector, a vector $\mathbf{G}_\ell = \{G_{\ell 1}, \dots, G_{\ell K}\}$ of K PMT outputs is collected. The elements of $\mathbf{G}_\ell$ are integer numbers. Assuming Poisson statistics, the ML estimation problem can now be mathematically formulated as [2], [10]:

$$\begin{aligned} \hat{m}_\ell &= \arg\max_{m \in \{1, \dots, M\}} \left\{ \prod_{k=1}^K \left[\bar{G}_k(m) \right]^{G_{\ell k}} \frac{1}{[G_{\ell k}]!} e^{-\bar{G}_k(m)} \right\} \\ &= \arg\max_{m \in \{1, \dots, M\}} \Lambda_{\ell m}, \end{aligned} \quad (11)$$

in which the vector $\bar{\mathbf{G}}(m)$ is the mean PMT output for an event occurring at pixel m [15]. The goal of the contracting-grid algorithm is to compute $\hat{m}_\ell$ by avoiding the exhaustive search above. In our case, we considered a $69 \times 69 \times 25$ -voxel 3-D detector and we assumed $K = 64$ PMT outputs for each event. The contracting-grid algorithm starts by considering a uniform $4 \times 4 \times 4$ coarse grid that covers the whole detector. The quantity $\Lambda_{\ell m}$ in (11) is evaluated at each point of the grid and the value of m that maximizes $\Lambda_{\ell m}$ on the grid is used as the center of the grid in the next iteration. The grid spacing is halved at each iteration. The algorithm terminates when the grid spacing is one voxel in each dimension.

In our GPU implementation, each event was assigned to a $4 \times 4 \times 4$ thread block. The threads in a thread block are responsible for loading into the shared memory the PMT outputs G_r . Once the PMT data are loaded, each thread computes the likelihood Λ_{lm} for a point on the grid. The 64 values of the likelihood are then searched for the maximum value. The center of the grid for the next iteration and the new grid spacing are calculated as outlined above, and shared with all the threads in the block using shared variables.

Each NVIDIA GeForce 9800 GX2 graphics card installed in our machine contains two CUDA-capable devices, which can be used independently, for a total of eight devices. We also implemented the same contracting-grid algorithm for a conventional cluster machine, which we used for our comparison. Results of the comparison are reported in Table I.

VI. Conclusions

In this paper the list-mode (LM) MLEM algorithm is applied to ML position and energy estimates for a Compton camera. An expression for the probability density function of the event attribute vector was given. The full implementation of the reconstruction algorithm requires two distinct estimation steps: the first is the 3-D ML estimation of event position from PMT data and the second consists in the reconstruction of the object distribution from the list of position estimates by using an iterative formula. We have written a GPU algorithm that implements the first estimation step and we ran our code on a GPU supercomputer assembled from commodity hardware. Our implementation could achieve a $250\times$ speedup with respect to an implementation on a conventional cluster machine. We are currently working on the GPU implementation of the iterative reconstruction step.

Acknowledgments

This work was supported by NIH grants 5 R37 EB000803-18 and 5 P41 EB002035-10.

References

1. Barrett HH. Objective assessment of image quality: Effects of quantum noise and object variability. *J. Opt. Soc. Am. A* Jul;1990 7(7):1266–1278. [PubMed: 2370589]
2. Barrett, HH.; Myers, KJ. *Foundations of Image Science*. Wiley-Interscience; Hoboken, NJ: 2004.
3. Dempster AP, Laird NM, Rubin DB. Maximum likelihood from incomplete data via the EM algorithm. *J. Roy. Stat. Soc. B* Met 1977;32(1):1–38.
4. Shepp L, Vardi Y. Maximum likelihood reconstruction for emission tomography. *IEEE Trans. Med. Imag* Oct.;1982 1(2):113–122.
5. Barrett HH, White T, Parra LC. List-mode likelihood. *J. Opt. Soc. Am. A* Nov.;1997 14(11):2914–2923.
6. Reader A, Manavaki R, Zhao S, Julyan P, Hastings D, Zweit J. Accelerated list-mode EM algorithm. *IEEE Trans. Nucl. Sci* Feb.;2002 49(1):42–49.
7. Parra LC, Barrett HH. List-mode likelihood: EM algorithm and image quality estimation demonstrated on 2-D PET. *IEEE Trans. Med. Imag* Apr.;1998 17(2):228–235.
8. Snyder DL, Polite DG. Image reconstruction from list-mode data in an emission tomography system having time-of-flight measurements. *IEEE Trans. Nucl. Sci* Jun.;1983 30(3):1843–1849.
9. Reader A, Erlandsson K, Ott R, Flower M. Attenuation and scatter correction of list-mode data driven iterative and analytic image reconstruction algorithms for rotating 3D PET systems. *IEEE Trans. Nucl. Sci* Dec.;1999 46(6):2218–2226.
10. Furenlid, LR.; Hesterman, JY.; Barrett, HH. Real-time data acquisition and maximum-likelihood estimation for gamma cameras. 14th IEEE-NPSS Real Time Conference; Jun. 2005; p. 498-501.

11. Hesterman, JY.; Caucci, L.; Kupinski, MA.; Barrett, HH.; Furenlid, LR. Maximum-likelihood estimation with a contracting-grid search algorithm. 2009. accepted for publication in *IEEE Trans. Nucl. Sci.*
12. Barrett HH, Hunter WCJ, Miller BW, Moore SK, Chen Y, Furenlid LR. Maximum-likelihood methods for processing signals from gamma-ray detectors. *IEEE Trans. Nucl. Sci* Jun.;2009 56(3): 725–735. [PubMed: 20107527]
13. Fisher RA. Theory of statistical estimation. *P. Camb. Philos. Soc* 1925;22:700–725.
14. Piegl, LA.; Tiller, W. The NURBS book. Springer; Berlin, Germany: 1997.
15. Hunter WCJ, Barrett HH, Furenlid LR. Calibration method for ML estimation of 3D interaction position in a thick gamma-ray detector. *IEEE Trans. Nucl. Sci* Feb.;2009 56(1):189–196. [PubMed: 20191099]

Fig. 1.
Setup for a Compton camera system

Fig. 2.
Setup for a PET system

TABLE I

Comparison results

Architecture	Events/s	Speedup
AMD Phenom 9850 2.5 GHz CPU	989.67	—
NVIDIA GeForce 9800 GX2, 1 device	34224.59	34.58
NVIDIA GeForce 9800 GX2, 8 devices	246003.92	248.57