Optimum real-time reconstruction of Gamma events for high resolution Anger camera with the use of GPGPU

Stefano Pedemonte, Alberto Gola, Andrea Abba, Carlo Fiorini, *Member, IEEE,*

***Abstract*—Aim of the HICAM project is to build a new, compact Gamma-ray imager with a submillimeter spatial resolution, based on the Anger Camera principle. The system is composed of a detection module with frontend ASICs, an acquisition board and a host PC, where real-time data processing and image reconstruction is implemented. The detector is based on an array of 100 Silicon Drift Detectors (SDDs) of $1\ cm^2$ each in a 10×10 cm^2 format, coupled to a single scintillator crystal. Position of the interaction with the crystal and energy of the incident radiation are obtained through the use of a MLE algorithm, that optimally exploits the information obtained from the detectors. Moreover it is possible to modify the algorithm in order to determine the depth of interaction of the Gamma photon inside the crystal. The MLE algorithm, on the other hand, requires a large amount of calculations per event. In order to process the events in real-time we have implemented the MLE algorithm on a GPGPU, obtaining a processing rate of $150000\ events/second$, considering a FOV of $512 \times 512 \times 10$ points (calculation of the z coordinate is performed). In the paper we discuss the derivation of the algorithm, its performance for what concerns spatial resolution and distortion, and the speed of its implementation on the GPGPU.**

***Index Terms*—Anger Camera, optimum reconstruction, maximum likelihood, CUDA.**

I. Introduction

THE aim of the HICAM project is to build a compact Gamma ray imager with submillimeter spatial resolution, based on an array of Silicon Drift Detectors coupled to a scintillator crystal [3]. Its potential applications are in the field of human imaging, in applications where high resolution along with probe compactness is needed. These include, e.g. intra operative probes and parathyroid imaging. Moreover, thanks to the high resolution, the probe finds application in in-vivo imaging of small animals, in both planar and tomographic applications.

The detector is designed according to the Anger Camera principle: a monolithic scintillator covers the active area of the light sensitive array and, when a Gamma event is detected, the scintillation light is shared among several pixels. The analog signals form the photodetector are amplified and filtered in order to maximize the S/N ratio in each cell.

Manuscript received November 13, 2009. Acknowledgement of support from the EC contract number LSH-CT-2006-037737 (HI-CAM)

S. Pedemonte was with Politecnico di Milano, Milano, Italy. He is now with the Department of Medical Physics, University College London, London, NW1 9EE, United Kingdom (telephone: 00447807544119, e-mail: s.pedemonte@cs.ucl.ac.uk).

A. Gola is with Politecnico Di Milano, Dipartimento di Elettronica e Informazione, Milano, 20133 Italy and INFN sez. di Milano (e-mail: gola@elet.polimi.it).

A. Abba is with Politecnico Di Milano, Dipartimento di Elettronica e Informazione, Milano, 20133 Italy (e-mail: abba@elet.polimi.it).

C. Fiorini is with Politecnico Di Milano, Dipartimento di Elettronica e Informazione, Milano, 20133 Italy and INFN sez. di Milano (e-mail: fiorini@elet.polimi.it).

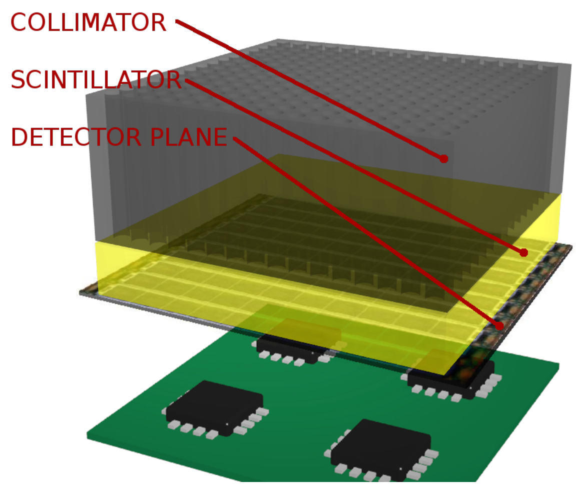

Fig. 1. Drawing of the HiCAM Gamma Camera.

In the original Anger camera design the analog signals are processed by means of a centroid algorithm which can be implemented by a net of resistors [3]. Many commercial cameras implement reconstruction by the centroid method as it allows for real time imaging. However this method does not optimally exploit the information from the camera, which is a serious drawback as it has direct effect on the dose of radiation given to the patient.

Several solutions have been proposed for digital processing of event data, including DSP and FPGA based ad-hoc hardware, in order to address the computational cost of maximum likelihood. In our implementation the signals are digitized and then acquired by a PC, where a suitable algorithm locates the position of interaction of the Gamma-photon with the crystal, by processing the information carried by the distribution of amplitude signals generated in the cells of the photodetector array. We exploit the flexibility of a PC based solution along with the processing power of a General Porpouse Graphics Processing Unit (GPGPU), aiming at real-time, high resolution, optimum coordinate and energy reconstruction.

In the paper we describe a Maximum Likelihood Estimator (MLE) [1] that provides optimal reconstruction of the coordinates and energy of the incident radiation, i.e. it exploits all the information that is carried by the signals from the array of detectors. In fact, the amount of light collected in the cells of the array constitutes a realization of a random process described by the joint probability functions associated to the light intensity measured from each cell, that are functions of the coordinates of interaction of the Gamma-photon with the crystal. Since the joint probability density functions are known from analytical model or characterization of the device, the problem is that of optimally estimating the point of interaction, given one realization of the random process. The problem can be solved by an algorithm based on the Maximum Likelihood Estimator (MLE).

From simulation of the Hicam Gamma camera, the use of a MLE based algorithm instead of a method based on the centroid method increases the resolution (FWHM spread is reduced from 3 mm to 0.5 mm), eliminates barrel distortion and allows for 3D recontruction of the interaction, i.e. we can obtain also information on the depth of interaction of the photon with the crystal. The z-axis information is especially useful when the anger-camera is coupled to a pin-hole collimator, infact, in this case, simple projection of the point of interaction on the x-y plane can lead to non negligible errors in the image. The actual performance of the algorithm with the real Gamma Camera strtongly depends on the precise knowledge of the system response and is actually under assesment.

We have conceived an algorithm based on MLE, considering several optimization methods for quick convergence and have finally implemented the algorithm for efficient reconstruction on a x86 CPU. However, in order to preserve the high intrinsic resolution of the Gamma Camera, we were not able to reduce the number of operations (multiply and add - MAD) per event to less than 60000. This yields a real-time rate of reconstruction of 10000 events per second, which can be limiting for the Gamma Camera.

In order to overcome this performance limitation we have evaluated other hardware computing resources and we have finally implemented the MLE algorithm for a general purpose graphics processing unit (GPGPU) [2], that represents a consumer off-the-shelf solution for high performance computing. The MLE algorithm and the optimization method have been modified for efficient execution on the architecture of the GPGPU.

II. Optimum reconstruction

The task of reconstruction is the task of associating the electrical signals produced by each Gamma event to the coordinates of the point of interaction and to the energy of the incident radiation. We will assume that each photodetector and its amplifing electronics are perfectly linear, that means that the ouput electrical signal is proportional to light intesity being absorbed; that they all have the same gain (which is achieved by proper calibration) and offset (non zero output when zero light intensity - which is also achieved by calibration).

Assuming that light is emitted randomly in every direction in the scintillator from a point source where the Gamma photon interacts with the crystal, the mean number of photons that are absorbed by each photodetector depends on the coordinates of interaction. As a first approximation each photodetector absorbs an amount of light proportional to the solid angle it subtends to the point of interaction. The information about the point of interaction is contained in the intensities read by the photodetectors; different points of interaction may be discriminated because they produce different patterns of light intensity on the photodetectors. However, given a point of interaction, only the *mean* of the number of photons collected by each photodetector is deterministic, as it is subject to fluctuations of the number of secondary photons generated, of radiation being emitted per unit solid angle and to fluctuation associated to quantum efficiency. The electrical signals (and finally the digitized values) present additional fluctuation due to noise sources in the electronics, and sampling noise. However we discard noise from electronics as a careful design can reduce it to a negligible effect, while Poisson statistics associated to the scintillation process are intrinsic to the Anger camera. The aim of the optimal estimator is to calculate the best estimate of the point of interaction, minimizing the uncertainty due to the noise sources of the system.

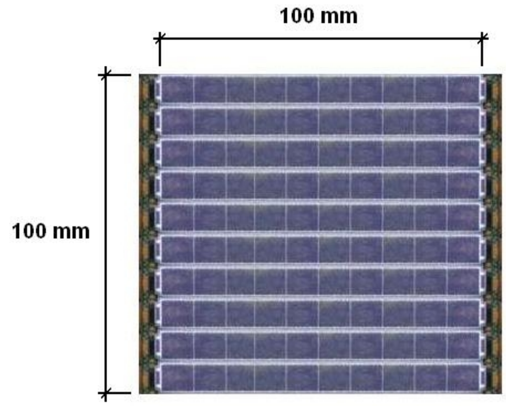

Fig. 2. Array of Silicon Drift Detectors.

The amplitude signals from the photodetectors are amplified and filtered by the readout electronics, sampled and digitized by the digital acquisition system; the digital processing unit deals with an array of numerical values, where each value is proportional to the light absorbed by one of the photodetectors. We will refer, in the following, to the array of electrical values, or vector of electrical values, or event data, which will be indicated by

$$x = (x_1, x_2, x_3, .., x_n) \tag{1}$$

where n is the number of photodetectors.

If one sets up an experiment where Gamma rays are forced

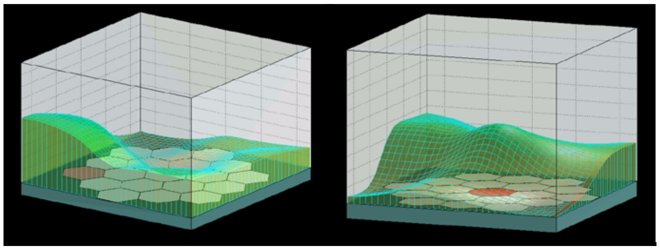

Fig. 3. Likelihood functions for a Gamma Camera based on an hexagonal array of Silicon Drift Detectors.

to interact always at the same spatial coordinate, and a number of interactions are considered, then a frequency distribution function for each signal may be obtained. For a big number of interactions the frequency distribution function associated to each signal will represent the probability density function for that signal, given the interaction point. The set of all the probability density functions (one for each photodetector), which is a function of the coordinate of interaction, represents the most complete model of the detector along with its electronics.

A geometrical/optical model may approximate the mean value of each probability distribution fuction. If noise from the electronics is negligible, then the probability distribution function associated to each signal is a Poisson distribution.

The problem of reconstruction of coordinates of interaction, given the electrical signals from the photodetector array is an estimation problem. Given the joint probability density function of the signals from the photodetectors, function of the point of interaction, the problem is that of estimating the point of interaction given one single realization of the random process described by the joint probability density function.

Let's consider the generic event data vector $X = (X_1, X_2, X_3, .., X_n)$. The probability distribution of X depends on an unknown parameter θ, that is a vector of coordinates of interaction. $X_1, X_2, X_3, .., X_n$ have the following joint probability density function that depends on the parameter θ.

$$f_\theta(X_1, X_2, X_3, .., X_n) = f(X_1, X_2, X_3, .., X_n|\theta) \quad (2)$$

Given the observed values $X_1, X_2, X_3, .., X_n = x_1, x_2, x_3, .., x_n$, the question is: which is the most likely value of the parameter θ that yields the given observation? $x_1, x_2, x_3, .., x_n$, for more clarity, are the aplitude signals from the photodetectors, and constitute one realization of the random process $X_1, X_2, X_3, .., X_n$, that is parameterized in θ.

Since the distribution of $X_1, X_2, X_3, .., X_n$ is known, as function of θ, the likelihood function is obtained by reversing the roles of X and θ. The likelihood function is $L(\theta|x)$:

$$L(\theta|x) = f(X|\theta) = f(x_1, x_2, x_3, .., x_n|\theta) \quad (3)$$

in other words $L(\theta|x)$ is the probability of observing the given data set $x = x_1, x_2, x_3, .., x_n$ as a function of θ.

The *maximum likelihood estimate* of θ is that value of θ that maximizes $L(\theta|x)$.

If X_i are *independent and identically distributed* random variables, then the likelihood $L(\theta|x)$ is

$$L(\theta|x) = \prod_{i=1}^{n} f(x_i|\theta) \quad (4)$$

Being log an always increasing function, maximizing $L(\theta|x)$ is equivalent to maximize the log likelihood function

$$l(\theta|x) = \sum_{i=1}^{n} log(f(x_i|\theta)) \quad (5)$$

As an approximation, we can assume that the signals from the photodetectors $X_1, X_2, X_3, .., X_n$ are statistically independent and identically distributed, so the joint probability distribution function is the product of the *pdf* of each signal.

Each signal is described by its *pdf*, function of θ:

$$P(x_i|\theta) = \frac{\lambda_{i,\theta}^{x_i} \cdot e^{-\lambda_{i,\theta}}}{x_i!} \quad (6)$$

Substituting into 5, the logarithmic likelihood function is then

$$l(\lambda_\theta|x) = \sum_{i=1}^{n} log\frac{\lambda_{i,\theta}^{x_i} \cdot e^{-\lambda_{i,\theta}}}{x_i!} \quad (7)$$

$$l(\lambda_\theta|x) = \sum_{i=1}^{n} x_i log(\lambda_{i,\theta}) - \lambda_{i,\theta} - log(x_i!) \quad (8)$$

The parameter that maximizes the liklelihood function is the vector $\lambda_\theta = \lambda_{1,\theta}, \lambda_{2,\theta}, \lambda_{3,\theta}, .., \lambda_{n,\theta}$, that is a function of θ, but the *invariance property* of maximum likelihood ensures the value of λ_θ that determines maximum $l(\lambda_\theta|x)$ is for θ that gives maximum $l(\theta|x)$, being $\lambda_\theta = \lambda_{1,\theta}, \lambda_{2,\theta}, \lambda_{3,\theta}, .., \lambda_{n,\theta}$ a deterministic function of θ.

The expression found in 8 represents an *estimator* of θ, given one realization $x = x_1, x_2, x_3, ..x_n$ of the random process $X = X_1, X_2, X_3, ..X_n$ that is determined by parameter θ. If the mean squared error is considered as the measure of precision of the estimation, the MLE is the optimum estimator as it achieves the Cramer-Rao lower bound when the sample size tends to infinity [1].

III. Discretization

In order to solve for the maximum likelihood in the digital domain, the problem needs to be discretized, the solution domain being the volume of the scintillator. The solution of reconstruction by the centroid method is 2-dimensional, while MLE is solved by definition in 3 dimensions. It is possible, however, to solve for MLE on a 2D grid, where the expected values at each node are chosen to account for all expected values corresponding to the same z coordinate. This

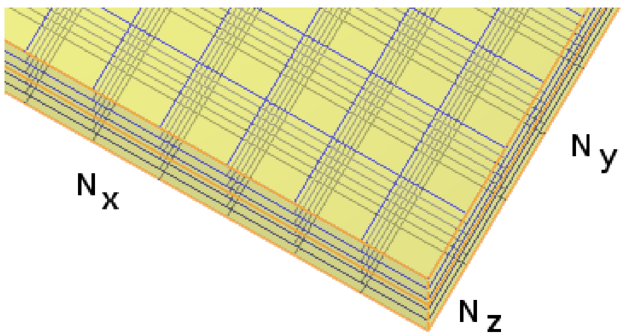

Fig. 4. Grid domain for computation of likelihood function

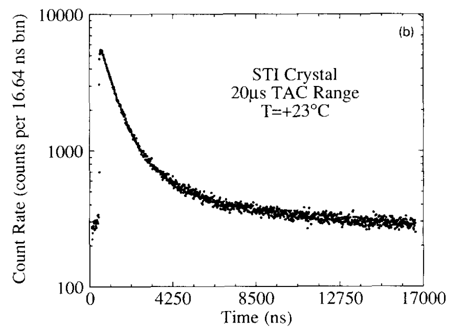

Fig. 5. Luminescence timing of Gamma excited CsI(Tl) Crystal (picture from [5])

can be done for example, if expected values are obtained by simulation, averaging expected values over z or computing them for an average value of z. In case expected values are observed from experimental data, it is straightforward to obtain the 2-dimensional values, but algorithms that extract 3D expected values have been developed. While tipically MLE is solved in the 2D domain in order to reduce the computational cost, this introduces an error in the estimation of the coordinates of interaction, that has been shown to be particularly cumbersome in case of pin-hole collimator where it introduces artefacts in the image. In order to obtain optimum reconstruction we compute the likelihood function in 3 dimensions.

The distance between points in the grid has to be small enough as to not limit resolution, in other words it has to be at some degree smaller than the variance of the reconstructed coordinates introduced by statistical fluctuations in the measured data. In order to quantify the computational cost we refer to the Hicam [3] high resolution Gamma camera, that has a crystal of size 100 $mm \times 100\ mm \times 5\ mm$ and an intrinsic resolution of about $0.5\ mm$. For this high resolution camera a reasonable discretization grid has $N_x = 512, N_x = 512, N_x = 10$ nodes. We will use these values in the next section.

The problem of reconstruction is also a problem of determining a model of the detector and representing the model in way that is practicable. A model that is represented by all the probability distribution functions, point by point, from every detector for all the possible points of interaction, cannot be stored in a memory as it is. In case of independent and identically distributed variables with Poisson pdf it is only necessary to store the expected value for the data vector, for each point of the domain. We will refer in the following section to vectors of expected data, and in our implementation we store the logarithm of the expected vectors in a lookup table.

When computing the log likelihood function, the third term in equation 8 can be discarded because it does not depend on θ. In our application we store in a lookup table the logarithm of the expected vectors $log(\lambda_{i,\theta})$, and the sum of each expected vector (the second term in equation 8), so evaluation of the log likelihood function reduces to a vector multiplication and one sum. Hence computing one point of the likelihood function costs approximately n sums of products (where n is the number of detectors).

IV. Real time reconstruction

Real-time reconstruction of the events implies that the processing rate is equal or graeater than the rate of interaction events. The rate of events that an Anger Gamma camera can detect is limited upwards by the characteristics of the scintillator chrystal [5]. Scintillation light is characterized by a fast rise time and a decay time that is tipically of a few μs (see 5). Being the emission of Gamma radiation a Poisson process with mean rate r, the everage number of events occurring in a period T is $\lambda = rT$. The probability of n events to occur in a period T is a Poisson function:

$$p_n = \frac{(rT)^n e^{-rT}}{n!} = \frac{\lambda^n e^{-\lambda}}{n!}$$

The probability of multiple events to occur in a period T is equal to $1 - p_0 - p_1$ (see 6-left for $T = 5\tau$ and $\tau = 2.5\mu s$):

$$p_{multiple} = 1 - p_0 - p_1 = 1 - e^{-\lambda}(1 + \lambda)$$

The rate of events that within a period T happen singularly is $R = rp_1$:

$$R = rp_1 = r\lambda e^{-\lambda}$$

We consider the reconstruction impossible if the event partially overlaps with another event; given the scintillation luminescence is approximately decaying exponentially (see figure 5), it is reasonable to consider the events non overlapping if $T = 5\tau$, where τ is the time constant of the decay. It this case 0.67% (e^{-5}) of the signal from the previous event constitutes noise for the current event.

$$R = 5r^2\tau e^{-5r\tau}$$

R is plotted in figure 6 for $\tau = 2.5\mu s$, which is an average value for an inorganic scintillator. For $\tau = 2.5\mu s$ the maximum rate of processable events is about 70000 events/second. A reconstruction rate of this order of magnitude can be considered realtime.

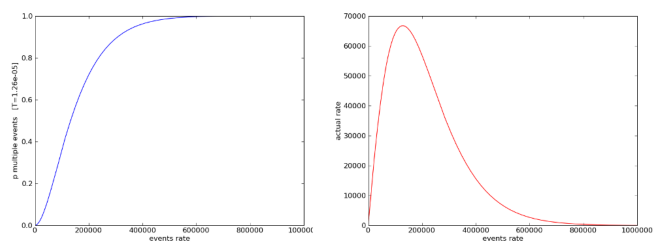

Fig. 6. Left: Probability of multiple events occurring in $T = 5\tau$,for scintillator decay constant $\tau = 2.5\mu s$. Right: Rate of non overlapping events observed in $T = 5\tau$,for scintillator decay constant $\tau = 2.5\mu s$.

V. GPGPU ACCELERATION

The recontruction of an event consists in the computation of the maximum of the likelihood function, over the solution domain, that is the 3D volume of the crystal. The naive solution is to compute the function in each point and find the greatest value. The evaluation of one point of the MLE function costs $n\ Flops$, with $n = 100$ for the Hicam camera.

Another approach is to find the maximum of the likelihood function by an iterative optimization algorithm, computing only the points that are needed. We did not explore this solution as it does not seem to fit GPGPU. However one must take into account that the likelihood function presents local maxima.

In our approach we sample the likelihood function on a grid of dicreasing size as schematized in figure 8-right. At each step the likelihood function is sampled on a grid, then at next step a more dense grid is created around the point that has maximum likelihood. Ad-hoc parameters for this algorithm (number of steps and size of the grids) allow to find the absolute maximum of the likelihood function, given the noise characteristics of the Gamma camera and the displacement of the detectors.

Table I summarizes the cost in terms of floating point operations, to solve for one event, with the mentioned approaches, with the parameters we used.

Given the high computational cost of MLE based recon-

TABLE I

COMPUTATIONAL COST FOR RECONSTRUCTION OF SINGLE EVENT WITH MLE BASED ALGORITHMS - DETECTOR SIZE $n = 100$

One point	100	$100Flop$
Naive maximization	$512 \times 512 \times 10 \times 100$	$262MFlop$
Iterative sampling	$(32 \times 32 + 32 \times 32 \times 10) \times 100$	$1.2MFlop$

TABLE II

COMPUTATIONAL COST FOR RECONSTRUCTIO OF 100000 EVENTS WITH MLE BASED ALGORITHMS - DETECTOR SIZE $n = 100$

Naive maximization	$26 \times 10^{12} Flop$
Iterative sampling	$112GFlop$

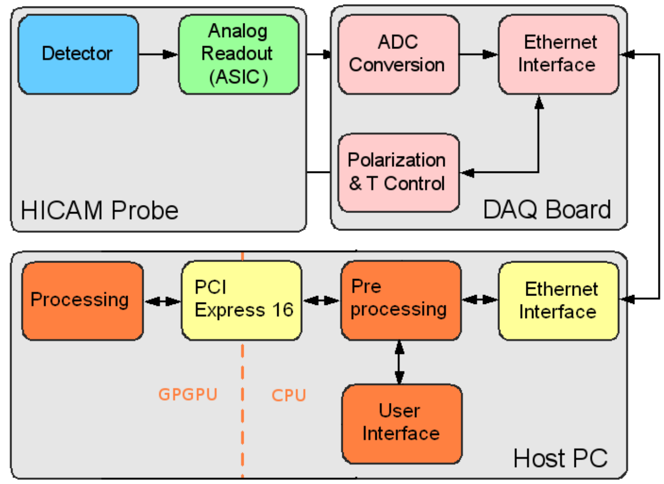

Fig. 7. Acquisition system of the HiCAM Gamma Camera.

struction, in order to meet real-time requirements, we have developed an implementation of the algorithm for General Purpose Graphics Processing Units (GPGPU). Our approach for acceleration of the reconstruction allows for good fitting of the algorithm to the characteristics of GPGPUs. Generally when implementing an algorithm for GPGPU the two essential aspects that affect performance are the amount of parallelism of the algorithm, that reflects on concurrent usage of multiple processors (nVidia defines *load* the percentage of processors in use by a kernel) and data locality, that allows for efficient use of shared memory. While this second aspect is often overlooked, it can impact enormously the performace.

A. Maximizing device load

Reconstruction of events by maximum likelihood is intrinsically prone to efficient implementation on a GPGPU. Let's consider first the case where the likelihood function is maximized by means of an iterative optimization algorithm. Multiple threads can compute single multiplication of a vector component with a component of the log of the expected vector. With our acceleration algorithm, at each step the event vector is multiplied by several vectors (all the expected vectors in the grid at that step); parallelism of the algorithm is increased. If the grid at each step is of size 32×32, and the vector is of lenght 100, one can launch $32 * 32 * 100 = 102400$ threads. Then the results from groups of 100 threads have to be summed, this can be accomplished by adding two components and killing half of the threads until one thread is left.
However we follow another approach as memory bandwidth actually limits the performance in this case, for any given possible arrangement of thread and block sizes in the GPU. This is due to low arithmetic intensity, infact, since each multiplication requires unique data from the lookup table, for each operation there is one access to memory.

B. Overcoming memory bandwidth bottleneck

The GPGPU we have used for the project (nVidia GTX260) has a bandwidth to device RAM memory of 140 GB/sec

but a latency of 600 $clockcycles$; on the contrary the shared memory is accessed in 4 $clockcycles$, but it is limited in size to 16 $Kbytes$ per multiprocessor. If one processes one vector at the time, each vector multiplication requires a unique vector (the logarithm of expectation vector for one node). With any arrangement of block and grid size, the number of operations per second is limited by the interface to device RAM memory.

The only way to exploit the bandwidth of the shared memory is to process multiple events at the same time, sharing data between the events. The problem is that, in order to reduce the computational cost, one tends to implement an acceleration algorithm, for the single event, that finds the maximum iterating over the likelihood function. Each event being independent by the others, their solutions follow different trajectories. Our acceleration algorithm allows for concurrent use of expected vectors by processing k events concurrently in the following way:

1) A GPU kernel function computes the likelihood function on a rough grid for the k events concurrently
2) A GPU kernel function finds the maximum of all the k likelihood functions
3) Events are clustered on the CPU: events with maximum likelihood in the same node are joined
4) One GPU kernel is launched for each node of the rough grid: this function computes the likelihood function on a fine grid that is centered around a specific node of the rough grid
5) One GPU kernel function that finds the maximum value of the sampled likelihood function is launched after each kernel launch at point 4.

Picture 8-right exemplifies this process. We implemented 2 steps only as it provides enough memory locality to overcome the bandwidth bottleneck. The two steps consist in the evaluation of a matrix multiplication and a function that computes per-row maxima of the resultimg matrix. In the first step the k events all undergo vector multiplication with the expected vectors at each node of the rough grid (this is a matrix multiplication); in step 2 each of the events that have been associated to a same node undergo vector multiplication with the expected vectors of the fine grid associated to that node. It is possible to tune the algorithm to any size of the rough and fine grid (and to include z axis) as each matrix multiplication is again partitioned in blocks in order to fit the shared memory efficiently, as explained in the nVidia Cuda programming guide [6].

VI. Results

We have discretized the crystal volume with a grid of $512 \times 512 \times 10$. By means of a Monte Carlo simulator that simulates our gamma camera we precompute the expected values for each node of the grid and store it to disk. The lookup with the expected values is loaded to the GPGPU device memory and events are streamed to the GPU in blocks of a few thousand events to be processed. Table III summarises the performance that we obtained. The potential of the GPGPU is fully exploited as the number of floating point operations per second approaches peak performance of the graphics card.

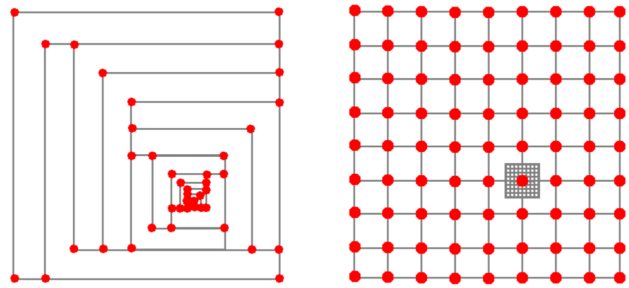

Fig. 8. GPGPU based acceleration algorithm.

TABLE III
Reconstruction rate with and without reconstruction of depth of interaction on nVidia GTX260

FOV	Number of detectors	**Reconstruction rate**
$512 \times 512 \times 10$	100	**150000 ev/sec**
$512 \times 512 \times 1$	100	**500000 ev/sec**
$256 \times 256 \times 1$	19	**1200000 ev/sec**

VII. Conclusions

The HiCAM Gamma Camera provides very high interinsic resolution but, in order to exploit the potential gain in resolution, it is necessary to implement more precise reconstruction algorithms. Reconstruction by maximum likelihood estimation provides minimum variance of the coordinates of interaction, however it is computationally expensive. Moreover, the higher the resolution of the camera, the more the algorithm is expensive, as it requires more dense solution domain.

Reconstruction by maximum likelihood moreover allows estimation of the depth of interaction, that increases the resolution and removes artifacts particularly in case of pin-hole collimators. However reconstructing the depth of interaction multiplies the computational cost.

We have focused on the implementation of a real-time optimum reconstruction by maximum likelihood, including depth of interaction, for high resolution gamma cameras and, in order to process the events in real time, we have adopted a GPGPU.

The key to the speedup of our reconstruction algorithm is that acceleration of reconstruction for one event allows

Fig. 9. Image reconstructed in real-time with Hicam Gamma camera coupled to a π shaped collimator and iodine source.

concurrent reconstruction of multiple events. This increases memory locality, essential for high performance computing on GPGPU.

References

[1] H. Cramer. *Mathematical Methods of Statistics.* Princeton University Press, 1957

[2] Lei Pan, Lixu Gu, and Jianrong Xu. *Implementation of medical image segmentation in cuda.* Technology and Applications in Biomedicine, May 2008, pages 82-85.

[3] C. Fiorini, A. Gola, R. Peloso, A. Geraci, A. Longoni,G. Padovini,P. Lechner,L. Struder,B. Hutton,K. Erlandsson,S. Mahmood,P. Van Mullekom,A. Pedretti,R. Moretti,G. Poli,G. Lucignani *Hicam: Development of a high-resolution anger camera for nuclear medicine.* Nuclear Science Symposium Conference Record, 2008. NSS 2008. IEEE, pages 3961-3964

[4] A. Abba, A. Manenti, A. Suardi, C. Fiorini, A. Geraci. *An effcient algorithm for spatial localization of multiple events from detector arrays in fpga devices.* Nuclear Science Symposium Conference Record, 2008. NSS 2008. IEEE, pages 1661-1664.

[5] J. Valentine, W. Moses, S. Derenzo, D. Wehe, G. Knoll, *Temperature dependence of CsI(Tl) Gamma-ray excited scintillation characteristics*, Nuclear Instruments and Methods in Physics Research A325 1993, pages 147-157.

[6] NVidia Cuda Compute Unified Device Architecture Programming Guide, Version 1.1, October 2007