Long Range, Passive Detection of Fissile Materials

L. Fabris, K.P. Ziock

February 2nd, 2005

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

Auspices Statement

This work was performed under the auspices of the U. S. Department of Energy (DOE) by the University of California, Lawrence Livermore National Laboratory (LLNL) under Contract No. W-7405-Eng-48. The project (03-ERD-048) was funded by the Laboratory Directed Research and Development Program at LLNL.

FY04 LDRD Final Report
Long Range, Passive Detection of Fissile Materials
LDRD Project Tracking Code: 03-ERD-048
Lorenzo Fabris, Principal Investigator

Abstract

We have recently completed a large-area, coded-aperture, gamma-ray imager for use in searching for radiation sources. The instrument was constructed to verify that weak point sources can be detected at considerable distances if one uses imaging to overcome fluctuations in the natural background. The instrument uses a rank-19, one-dimensional coded aperture to cast shadow patterns onto a 0.57 m² NaI(Tl) detector composed of 57 individual cubes each 10 cm on a side. These are arranged in a 19 x 3 array. The mask is composed of four-centimeter thick, one-meter high, 10-cm wide lead blocks. The instrument is mounted in the back of a small truck from which images are obtained as one drives through a region.

Introduction

Remote detection of small quantities of radioactive materials is of interest for a variety of applications. Although sufficient radiation can reach a large gamma-ray detector from a small source to make detection possible out to considerable distances, the naturally varying background makes such detections impossible unless it is known in advance. [1] Unfortunately, this is generally not the case. In fact, because the background can vary by order of itself, and is unknown, the standard signal-to-noise ratio (SNR) expression:

$$SNR = \frac{S}{\sqrt{B}} \propto \sqrt{At} \quad (1)$$

must be replaced by

$$SNR = \frac{S}{B} \propto \frac{At}{At} = \text{Constant} \quad (2)$$

Where S (B) is the signal (background) count rate, in a detector of area, A , measuring for a dwell time, t .

A physical interpretation of this problem is that variations in the background at modest distances have the same shape in an omni-directional detector as a point source at that distance. Hence, a concrete building at 20 m will look like the point source one wishes to find.

To overcome this effect, we have developed an imaging instrument [2] that makes a map of the radiation field as it traverses a region. The pixel size of the map is sufficiently small at distances out to 100 m to allow one to distinguish point from distributed objects. The background in neighboring pixels thereby provides a measure of the background, allowing one to apply the classic SNR expression of equation 2, increasing sensitivity by an order of magnitude or more over an omni-directional detector.

Instrument Description

I. Mechanical design

The imager (Fig. 1) is designed around an extant set of 57, 10 cm x 10 cm x 10 cm NaI(Tl) detectors. These are arranged in a 3 x 19 array to sit behind a one-dimensional, base-19, uniformly-redundant array, coded-aperture mask [3]. The details of the design choices are provided in [2]. Briefly, the number of pixels in the imager is determined primarily by the length of time that an object is in the instrument's field of view, given its one meter focal length. This size is selected as a compromise between vehicle width and the imager resolution at 100 m (10 m). The imager provides only one-dimensional images—a design choice based on its intended use in suburban or light-urban areas. Searches in areas with high-rise buildings beyond three stories are significantly complicated by the extra shielding from the concrete floors prevalent in their designs.

The mask comprises individual elements of linotype metal (85 % Pb, 11 % Sb, 3 % Sn, and 1 % other materials) that are 4 cm thick by 10.8 cm wide by 100 cm high. This material was selected for its lead-like stopping power with improved mechanical properties. The individual elements are supported in a steel frame that allows for rotation of sections of the mask around axis of anti-symmetry. Additional mask elements may be lowered into or out of the field of view of the detector (Fig. 2). These features allow us to exchange the mask for its inverse (anti-mask) in sequential passes past a source, providing an important means of noise control as explained later. The detectors are supported from a plate mounted in a steel frame similar to that used for the mask. The detectors are separated from the plate with a 0.5 cm thick sheet of elastomer to help protect them from mechanical and thermal shock. An additional 2.5 cm of Styrofoam is placed on the front, tops, sides and bottom of the detector array as thermal insulation.

Several pieces of aluminum channel connect the mask and detector frames to fix the one meter focal spacing and to restrict motion of one with respect to the other. The instrument relies on friction to keep it from moving within the truck. A safety strap ties the mask to the rear of the vehicle to protect the cab occupants in the event of a sudden stop. The instrument weighs 1585 kg with the 1225 kg mask comprising the bulk of this mass. The overall length of the mask is four meters.

Fig. 1. The mask and detector are fixed to an aluminum frame that maintains the imager focal length and alignment. The system is mounted to look through the curb side of a vehicle.

Fig. 2. The mask is converted to the anti-mask by rotating different parts about axis of anti-symmetry (A and B—B is hinged to reduce the swing space required) or by lowering a section out of the field of view (C).

II. Electronics System

Each of the 57 photomultiplier tubes is connected to a custom-built low-voltage/high voltage system. The low voltage supply is provided individually to each tube, while the high voltage is shared among three different tubes. Each tube has its own fine-adjustment for the high voltage value, in order to have individual gain control. The low-power electronics (120 Watts, full system) runs off the vehicle power system. It is based on a modular concept with one circuit board per detector (Fig. 3). The card supplies both a fast channel for triggering and a slow channel for spectroscopy. Signals above a lower-level threshold trigger the card which digitizes the output of the slow channel (a gated integrator was the choice for efficiently processing the analog information coming from our scintillators) using a low-power, 16-bit digitizer (decimated to 12 bits to achieve nuclear quality.) The logic also sets a signal to alert the overall system that the card has data. Up to 64 of these cards can populate a motherboard that provides communications to a centrally-located, digital-processing card (Fig. 4).

A field-programmable gate array (FPGA) mounted on this board reads the ADC values and forms a 32 bit digital word by adding the card number and whether upper and lower discriminators (not used in this experiment) have fired. The FPGA sends the data to a commercial PCI card [4] for processing by the data collection computer. In addition to servicing the individual cards, the FPGA card also maintains a system clock. The value of this clock is inserted at one-millisecond intervals into the data stream. Finally the FPGA card accepts events from a commercial, Doppler-radar-based "fifth-wheel" that provides location information. [5] This unit generates a frequency of 100 Hz for each mile per hour of the vehicle. The FPGA board counts the pulses from this unit and provides two sequential 32-bit events (ID plus 16 bits of the count per event) every 10 counts. If the vehicle is not moving, the location count is provided every second). Given the flexible nature of this design, the electronics system has been employed also in support of other ongoing LLNL projects.

Data Handling

Fig. 3. Single channel PC board with different regions identified.

Fig. 4. Electronics assembly in a standard 19 inch rack-mount chassis. The FPGA board (arrow) services the 64 individual cards via the mother board.

The data-collection computer is mounted with the other electronics in the back of the vehicle. A local area network is used to provide the vehicle passenger with control of the system via a laptop computer. The high-level, control-and-display software is written in Microsoft Visual Basic 6. [6] Data is collected from the National Instruments card using their NI DAQ function calls. As the data is collected, an event-mode file of the full data stream is recorded to disk for later analysis. Simultaneously, the data is parsed for image processing and display. On-line displays include a grayscale view of the 19 x 3 detector assembly, energy histograms of individual pixels or the entire detector array, a counts-per-meter histogram and the image of the track past the source.

The displays are populated from the incoming data using the following methods. First, the raw data is corrected for gain and offset based on stored calibration data from previous runs with ^{137}Cs and ^{60}Co sources. The linearized data is then checked against possible limits in both position and energy before it is added to an energy histogram maintained for each of the 57 detectors and a "short-term" detector. The latter is a 19-element array that is used to generate the images as described below.

I. Image Generation

As an image run proceeds, data is continuously summed to the short-term detector until a fifth-wheel event occurs. Upon such an event, the current location of the imager is checked against one of 11 possible pixel boundaries. These boundaries represent the edge of "world pixels" at different distances from the detector and are the key to understanding how an overall image such as that shown in Fig. 5 (top) is created. The simplest of these is used to generate the counts-versus-position histogram (Fig 5, bottom) and is explained first.

For the counts-versus-position histogram, the pixel boundaries occur every meter and the histogram simply keeps track of the number of counts that occur while the imager is within that one meter distance. Upon a fifth-wheel event, the current short term detector is added to a dedicated "distance-specific" detector for this histogram. A check is then made to see if the detector position exceeds the next pixel boundary. If not, the code returns without further changes. If the pixel

Fig. 5. Anti-mask track past 1 mCi ^{137}Cs source at ~ 5 km/h, range 46 m (top). The middle plot shows the histogram of counts vs position at the distance denoted by "<" (45 m) in the top graph. The bottom plot is the counts-per-meter histogram (converted to counts/sec for convenience.) Note that this does not cover as wide a position range as the imaged data.

boundary has been exceeded, then the counts in the distance-specific detector are normalized to the expected dwell-time in the pixel, summed and placed in the appropriate pixel of the counts-per-meter histogram. The distance-specific detector is then cleared, the time is recorded and the next pixel boundary for this range is calculated.

Normalization to the dwell time is performed by taking the actual time spent in the pixel (current time minus the start time for the pixel) divided by the expected dwell time at the estimated

velocity. The later number is selected at the start of the run and is the velocity that the driver attempts to maintain during the scan. Note that this number is arbitrary and the same relative image results although the numeric values differ if one uses a different assumed velocity. For this publication, the results have been further processed to give a counts-per-second value by multiplying by the assumed meters-per-second number.

The procedure to construct the overall image is only slightly more complicated. In this case we have ten different world-pixel maps to keep track of—one for each 10-meter range increment

from the imager. These pixel maps have different boundaries since the size of the pixels increase as the distance from the imager increases. The choice of ten, 10-meter pixel maps is arbitrary; one could just as easily have selected 100, one-meter pixel maps. However, the distance information is limited by the parallax as one passes a source and more pixels increase the processing required. No optimization of the number of distance pixels has been performed to date.

During the scan, (see Fig. 6) data is continuously accumulated into the “short-term” 19-element detector array. When a fifth-wheel event occurs, the code adds the data from this array to each of the 11, 19-element, “distance-specific” detector arrays and resets the short-term array. It then checks to see if any of the pixel boundaries at different distances have been crossed.

Consider the world-pixel map centered at 55 m that is meant to include sources from 50 to 60 m. For a one-meter focal length, the pixels are 5.94 m wide. The absolute edge of a pixel is selected in the code by assuming the pixels for all ranges line up at location “zero” halfway through the scan (again this is selected before the pass and is arbitrary.) When the code checks to see if the edge of the current 5.94 m wide pixel has been reached, it is using the edge determined at the start of the previous step. If the edge has not been reached, no further action is taken. However, if the edge of the pixel is exceeded, then the data in the distance-specific detector array is normalized to the dwell time and used to form an image based on the coded aperture deconvolution algorithm. [7] This provides a 19-element image that is

Fig. 6. Flow diagram used to generate world image.

added to the correct 19 elements of the world map at the 55 m range. (The center of the image maps to the current location.) The distance-specific array is erased, the next pixel boundary is computed and this is recorded together with the new start time for this source range. Finally, the grayscale histogram is updated and more data is processed.

II. Mask/Anti-Mask data handling

The imager can be run with the mask in either a “mask” or an “anti-mask” configuration. These different configurations are used to remove spurious background from the images as originally pointed out in [8]. Specifically, the coded aperture image process generates an image based on the variation in counts versus position across the face of the detector. These variations are supposed to be due to the shadow cast by the mask on the detector. However, variations are also possible due to radiation scattered off of nearby structures, radiation incident from other directions and even shielding of inside pixels by those on the edge of the detector array. A summation of equal-time exposures with the mask and its inverse should show no variation of counts versus position. Any variations that exist are extraneous and introduce artifacts in the image. The mask plus anti-mask data set can be subtracted from each of the separate runs before the image is generated to remove this noise. However, as shown in [9], this procedure is the same as adding mask-only data and subtracting anti-mask only data from the image. This is the approach taken here, so that if the mask is in the anti-mask configuration, the data accumulated in the short-term detector is subtracted from the distance-specific detector. The image generation multiplies these values by plus or minus one so that both positive and negative counts can occur in an image. In general the mask (anti-mask) only image is greater (less) than zero due to the non-modulated background in the detector (see Fig. 5.)

III. Error estimation

To first order, the statistical error in any single pixel of a coded aperture image is given by the square root of the total number of counts in the entire detector. [10] This follows since the image deconvolution process is linear. In this work, this simple observation is complicated by the fact that multiple images, each normalized to their own dwell time, are added to generate the final world image at any given distance. Therefore, the variance in any given pixel, dI , is:

$$(dI_n)^2 = \sum_j \frac{N_j}{D_j^2} \quad (3)$$

where N_j is the number of counts in the j^{th} image and D_j the appropriate dwell time for that image. To keep track of this, an additional summation is made whenever a pixel boundary is encountered and a world map of this sum is maintained to provide the variance at any given pixel of the final image.

Performance

The completed imager was used to track past a one millicurie ^{137}Cs source at a number of different velocities and distances. In all runs, only data in the 662 keV peak (590 - 730 keV) is used. Distances were fixed by the local parking lot and road structure and the requirements that the source remain inside its “assigned” building. For the first measurements, the source is mounted against the inside of the near wall of the building and the vehicle tracked past the source at a distance of 46 m at the

Fig. 7. Combined mask/anti-mask image of track past the source at 5 km/h at a range of 66 m (top). The multiple histogram in the middle shows the world map histograms at 65 m denoted by the "<" in the top image. The middle curve shows the histogram of the mask anti mask image; while the upper curve shows the mask pass (offset by + 1000 counts) and the bottom curve shows the anti-mask pass (offset by - 1000 counts.) The bottom histogram shows the non-imaged counts/sec detected per pass in the detector as a function of location. This plot shows no evidence of the source.

the far side of the building to obtain a track past the source at 83 m. This image (Fig. 8) was obtained traveling at 16 km/h. It also includes effects from all of the structural components of the near wall of the building. The counts histogram at 85 m is given in the middle panel of the figure. The peak represents a three-sigma

slowest velocity that could be maintained reliably (~ 5 km/h). The building itself is a tall corrugated metal-sided structure with apparent minimal wall thickness defined by the metal and insulation, both of unknown thickness.

The track past the source at 46 m with the imager in the anti-mask configuration is shown in Fig. 5. The top panel of the figure shows the complete grayscale image for the full set of distances on 10 m centers. The source is clearly visible in the center of the image. The middle panel shows the count histogram of the 45 m distance world map, clearly revealing a significant detection (± 162 counts, one sigma). The bottom plot shows the counts per meter histogram obtained without imaging. Although the source is also visible as the broad "bump" in the counts-per-meter histogram, as shown in Section V, such an assignment could not be made without imaging or a *priori* knowledge of the background.

In Fig. 7 we present the results for a track past the source at a range of 66 m, also traveling at 5 km/h. In the grayscale image at the top of the figure we show the combined mask/anti-mask data. This no longer suffers from the large offset in counts as observed in Fig. 5. The histogram of world pixels at 65 m is plotted below the image for the mask pass, the anti-mask pass and for the combined data as shown in the image. The offsetting nature of the two passes is clearly evident.

Finally, we moved the source to

departure from zero counts indicating that a detection has been made. The counts-per-meter histogram is shown in the bottom panel. It is hard to imagine any argument that would allow one to conclude that the source is present based on that

Fig. 8. Combined mask/anti-mask image of source at 16 km/h at a range of 83 m (top). The graph in the middle is the histogram of the image at 85 m denoted by the "<". The peak is three sigma above zero counts. The histogram at the bottom is the counts per meter in the detector

Fig. 9. Combined mask/anti-mask image of background at 16 km/h at a range of 83 m (top). The graph in the middle is the histogram of the image at 85 m denoted by the "<". The peak is three-sigma above zero counts. The histogram at the bottom is the counts per meter in the detector

data alone. The data in Fig. 9 shows a scan past the same location with the source removed.

Background Measurements

In addition to the tracks past the Cs source, a number of measurements were conducted at various locations in the Livermore area. A GPS receiver was used to record the tracks on the data collection computer and the resulting map showing some of the scan sites is shown in Fig. 10. [11] Note that the GPS data is not currently embedded in the data stream and is therefore not used for analysis. For each of these runs, the imager velocity was maintained at 16 km/h and both a mask and an anti-mask pass were made. As with the source data, alignment was

performed by sighting on objects next to the roadway. This provided a few meters uncertainty between the successive passes due to the higher speeds used in these runs. The images and non-imaged counts-per-meter histogram for a number of these runs are shown in Fig 11. Unlike the source tracks, this data represents the full spectral range (50 keV to 2.7 MeV) of the imager. (Scaled data for the Cs energy band is shown in the counts-per-meter histograms for comparison.)

Discussion

It is clear that the imager allows one to determine the presence of small sources out to distances beyond 80 m—a feat not possible with a non-imaging system. While the 3-sigma significance of the detection at the longest range is not sufficient to conduct a search for unknown sources, it does provide a proof of the value of imaging.

To verify the statistical confidence in the 83-meter detection, we played back the various background runs using the same energy cut (590 – 730 keV) used to generate Fig. 8. For each of the pixels at 85 m that were fully imaged in a pass, we calculated the number of sigma the pixel is from zero. The data is shown as a histogram in Fig. 12. A total of 435 pixels are included in the analysis. They have a mean value of 0.12. A Chi-square minimization fit to a Gaussian [12] generates the curve shown in the figure. This has a reduced chi-square of 1.29 and a one-sigma width of 1.00 ± 0.07 . This clearly indicates that the statistics have been handled correctly in the imaging, and that the sigma values are normally distributed.

It is interesting to note that the mean value of the image is zero. One clear interpretation of this is that strong sources are not present in the image. This is as expected since a building-sized source is necessarily a weak emitter per unit area—exactly the point of the imaging in the first place.

In contrast to the normally distributed background of the imaged data, Fig 13 shows that the background of the non-imaged data varies in a systematic fashion (i.e. the variation is not given by counting statistics.) This can be clearly seen by comparing the normally distributed, imaged data of Fig. 12 to the counts per second data in Fig. 13. The non-Gaussian nature of the latter is readily evident. Its effect on search sensitivity is seen by recognizing that the full width at half maximum of a normally distributed background based on counting statistics would be ~ 220 counts or just over one of the bin widths in the histogram!

Fig. 10. GPS track of imager at the west end of Livermore, CA. The arrows and numbers represent the rough locations and direction of travel during the different imaging background run. The imager is configured to look toward the near curb.

Fig. 11 Data collected around Livermore, California. Two passes were made, each at 16 km/h. On the left is the imaged data, on the right is the histogram of counts per second. The top trace of the counts-per-second histogram is the full-energy band from 50 to 2700 keV. The bottom trace is for the Cs-only energy band from 590 to 730 keV multiplied by a factor of ten to better match the range of the full-energy data in the plot. From top to bottom, the data correspond to run 1, 2, 4 and 5 of Fig. 10. Note that the bottom run passes under the freeway. The increases in count rate at 125 m and 200 m correspond to passing under the two four-lane bridges. It is important to note that these are the most obvious features found in any background run and that they do not show up in the corresponding image. The grayscale of each image is independent running from the minimum to the maximum value. In these runs a total of 61 (5) pixels are more than 3 (4) sigma above zero out of a total of 12,639 pixels.

To date, one of the most surprising results found is the uniformity of the

Fig. 12. Histogram of imaged background data reconstructed to a range of 85 m, using a Cs energy cut from 590-730 keV. The line is a fit to a Gaussian with a reduced Chi-square of 1.29. The one sigma width is 1.00 ± 0.07

Fig. 13. Histogram of Counts/Sec taken from several background runs. This is non-imaged data.

background. The absolute magnitude of its variance is an order of magnitude smaller than that predicted in [1] ($\pm 20\%$ not $\pm 200\%$). This does not invalidate the concept that systematic noise limits the range at which a point source can be detected. The standard deviation of the non-imaged background from its mean value of 8955 counts/sec is 1140 counts, an order of magnitude greater than would be expected from counting statistics alone. The effect of systematic variance on a source detection can be seen graphically by comparing the variation in the background in the Cs energy cut shown in Fig 11 (scaled by 10), to the variation seen in Fig. 5, the closest source pass at 46 m. The background increase due to the source in the non-imaged data is $\sim 20\%$. This is similar in size to other variations seen where no source is present. The ~ 1540 rise in counts seen in the anti-mask only image is ~ 9 sigma—a clear detection by any measure.

Acknowledgements

The authors wish to acknowledge the Department of Homeland Security for current support of this project. Mechanical assistance provided by Dennis Carr and Darrel Carter and electronics support by Marianne Ammendolia, all of LLNL, is also gratefully acknowledged.

References

- [1] K.P. Ziock, W. Goldstein, "The lost source, varying backgrounds and why bigger may not be better," in *Unattended Radiation Sensor Systems for Remote Applications*, J. Trombka, D. Spears, P. Solomon, editors, AIP Conf. Proc. vol. 632, pp. 60-70, 2002.
- [2] K.P. Ziock, W.W. Craig, L. Fabris, R.C. Lanza, S. Gallagher, B.K.P. Horn, N.W. Madden, "Large Area Imaging Detector for Long-Range, Passive Detection of Fissile Material," *IEEE Trans. Nucl. Sci.*, Vol. 51, no. 5, pp. 2238-2244, 2004.

- [3] S. Gottesman, E. Fenimore, "New family of binary arrays for coded aperture imaging," *App. Opt.*, vol. 28, No. 20, pp. 4344-4352, Oct. 1989.
- [4] National Instruments PCI-6534, <http://www.ni.com>
- [5] GMH Engineering, <http://www.gmheng.com/radar1.html>
- [6] Microsoft Corporation, <http://www.microsoft.com>
- [7] E.E. Fenimore, T.M. Cannon, "Coded aperture imaging with uniformly redundant arrays," *Applied Optics*, vol. 17, no. 3, pp. 337-347, Feb., 1978.
- [8] C. Brown, "Multiplex Imaging with multiple-pinhole cameras," *J. Appl. Phys.*, Vol. 45, no. 4, pp. 1806-1811, April, 1974.
- [9] K. Ziock, "Scintillation imaging: A technique to reduce coding noise in scanned, coded-aperture images," *Proc. of SPIE*, vol. 5540, pp. 225-234, 2004.
- [10] E.E. Fenimore, "Coded aperture imaging: predicted performance of uniformly redundant arrays," *Applied Optics*, Vol. 17, no. 22, pp. 3562-3570, Nov., 1978.
- [11] By permission Screen Capture from Microsoft "Streets & Trips 2005".
- [12] W.H. Press, S.A. Teukolsky, W. T. Vetterling, B. P. Flannery, Numerical Recipes in C, (Cambridge Univ. Press, New York, NY, 1995), p. 685.