

Roadside Tracker Portal-less Portal Monitor

July 2013

**Prepared by
Klaus-Peter Ziock**

DOCUMENT AVAILABILITY

Reports produced after January 1, 1996, are generally available free via the U.S. Department of Energy (DOE) Information Bridge

Web site <http://www.osti.gov/bridge>

Reports produced before January 1, 1996, may be purchased by members of the public from the following source.

National Technical Information Service

5285 Port Royal Road

Springfield, VA 22161

Telephone 703-605-6000 (1-800-553-6847)

TDD 703-487-4639

Fax 703-605-6900

E-mail info@ntis.gov

Web site <http://www.ntis.gov/support/ordernowabout.htm>

Reports are available to DOE employees, DOE contractors, Energy Technology Data Exchange (ETDE) representatives, and International Nuclear Information System (INIS) representatives from the following source.

Office of Scientific and Technical Information

P.O. Box 62

Oak Ridge, TN 37831

Telephone 865-576-8401

Fax 865-576-5728

E-mail reports@osti.gov

Web site <http://www.osti.gov/contact.html>

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, 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 any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Nuclear Security and Isotope Technology Division

ROADSIDE TRACKER PORTAL-LESS PORTAL MONITOR

Klaus-Peter Ziock

Date Published: November 2013

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, Tennessee 37831-6283
managed by
UT-BATTELLE, LLC
for the
U.S. DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

Table of Contents

Table of Contents	iv
List of Figures	v
List of Tables	vi
1. Introduction	1
2. Background	2
3. Imaging Solution.....	4
4. POC Instrument	6
Gamma-Imager Design	7
Gamma Imager Optimization.....	7
Gamma-Data Acquisition system and Data handling	9
Gamma-Ray Image Generation	12
Video system design	14
Camera System	15
Video data handling	16
Processing Pipeline	23
Acquisition.....	17
Alignment camera transformation	18
Stereo Vision Reconstruction	19
TAT Output.....	19
Dual-road Fusion track clustering.....	20
Vehicle detection, Monitoring Estimation, and Video Event Generation	20
High-Resolution Images for Vehicle ID	21
Rig drift monitor	22
5. Project Evolution	22
Phase 2: POC design	24
Phase 3: POC Construction.....	24
Phase 4: Integration and Test and Evaluation	25
Phase 5: Additional Development.....	27
Independent Characterization	33
Summary InterAgency Agreement History	33
6. Conclusion	36
7. References.....	37
Appendix A: List of Acronyms.....	39
Appendix B: POF Report.....	40

List of Figures

Figure	Description	Page
1	Operational Roadside Tracker Portal-less Portal Monitor	1
2	Classic portal monitor installation at a border crossing	5
3	Large Area Imager for stand-off-detection search instrument.	6
4	Roadside Tracker results for a moving vehicle with two sources.	7
5	Graphical representation of the process of single frame vehicle image capture.	8
6	Block diagram of the overall data acquisition system.	12
7	Block diagram of s single channel of gamma-ray data acquisition.	13
8	Schematic diagram of gamma-ray buffer data structure.	14
9	Illustration of vehicle slices.	15
10	Progress of video hardware during the project.	18
11	Bayer image format.	18
12	Block diagram of video processing data flow.	19
13	Targets on the trailer roof.	20
14	Stereo point search approach.	21
15	Standard and high resolution video images.	24
16	Project timeline.	25
17	Major events during design phase.	26
18	Major events during POC construction phase.	26
19	Major events during phase 4.	27
20	RST standard 2-lane development site.	28
21	RST development site on main ORNL thoroughfare.	28
22	First offsite data collection site.	29
23	Major events during improvement, test and evaluation phase.	29
24	Major events during the limited use exercise.	30
25	Site at SRNL where high-speed tests were conducted.	32
26	RST performance vs velocity.	33
27	RST deployed near Belmont raceway.	34
28	Improved gamma performance after removal of video frame offset.	35

List of Tables

Table	Description	Page
1	Gamma imager properties.	9
2	InterAgency Agreements	35

1. INTRODUCTION

This report documents the full development cycle of the Roadside Tracker (RST) Portal-less Portal monitor (Fig. 1) funded by DHS DNDO. The project started with development of a proof-of-feasibility prototype, proceeded through design and construction of a proof-of-concept (POC) prototype, a test-and-evaluation phase, participation in a Limited Use Exercise that included the Standoff Radiation Detections Systems developed under an Advanced Technology Demonstration and concluded with participation in a Characterization Study conducted by DNDO.

The RST represents a new concept in radiation detection portal monitors, specifically providing the ability to scan passing traffic for illicit radiation sources without impacting the flow of that traffic. To achieve this it makes use of a radiation imaging approach that relies on an automated, Target Acquisition and Tracking (TAT) software engine applied to visible light images. Output from the TAT engine provides the information necessary to remove motion blur, and generate a one-to-one link, of the gamma-ray images of each vehicle with a visible light image of that vehicle. In turn, these can be used to identify the vehicle to response personnel positioned with, or just downstream of, the RST. In addition to providing a one-to-one link between threats and vehicles, the system can also separate point (threat) from distributed (non-threat) sources, provide the approximate 2D-location of any detected sources within a vehicle, and generate background-subtracted gamma-ray spectra of any sources that are detected (for use with isotope ID software). Finally, the approach also provides first-order immunity from systematic sensitivity issues that plague classic radiation portal monitors, specifically, baseline suppression, varying backgrounds, and source confusion.

Following this introduction, the remainder of the report is organized to provide the reader with the information required to understand the issues encountered in the instrument's development, including a general introduction to how and why the imaging approach was adopted, and a technical description of the instrument and the reasoning behind selection of different components. This is followed by an overview of the project timeline and a phase-by-phase description of issues, achievements, and lessons learned. Finally, the report concludes with recommendations for additional development to extend the capabilities of the current POC instrument as well as additional application scenarios and modifications such scenarios would require.

Roadside Tracker Portal-less Portal Monitor

CFP 06 – TA-02-OR18

ORNL DNDO Project Manager:

Richard A. Turner

Prepared by:

Klaus Ziock, Principle Investigator

Oak Ridge National Laboratory

P.O. Box 2008, MS-6010

Oak Ridge, TN 37831

zockk@ornl.gov

Team:

K.P. Ziock, A. Cheriyyadat, E. C. Bradley, M. F. Cunningham, L. Fabris, J. S. Goddard, D. E. Hornback, T. P. Karnowski, R. A. Kerekes, Jason Newby



Fig. 1. Operational Roadside Tracker Portal-less Portal Monitor deployed on the DC beltway.

2. BACKGROUND

By their very nature, illicit nuclear materials are radioactive, emitting radiation that can be intercepted by remote detectors requiring contact with neither the materials nor the containers that hold them. Although the passive radiation emissions can be shielded to reduce or eliminate their detection, such shielding increases the mass and volume of the materials, making in-transit searches for such materials one place where shielding might be minimal or non-existent. Classic radiation portal monitors were developed to exploit this vulnerability in a broader, layered approach to control nuclear materials. The devices comprise large radiation detectors typically deployed on either side of a traffic lane and vehicles are required to move slowly between them Fig. 2.



Fig. 2. Classic portal monitor installation at a border cross-

An occupancy monitor is used to determine when a vehicle is present and to separate background measurements (no vehicle) from vehicle measurements. In general the monitors use a threshold alarming approach, setting an alarm if the counts when a vehicle is present are statistically higher than those when no vehicle is present. The devices were originally developed for use around known nuclear facilities [1, 2] and have since expanded to this more general application in society at large. While a detailed description of standard portal monitors is beyond the scope of this report, there are several important aspects of the performance of these devices and their initial application that informed the decisions made in development of the RST and these are briefly discussed below.

In their initial application, at the periphery of nuclear facilities, radiation portal monitors (RPM) were designed to detect specific materials. The low level of traffic in and out of the facilities meant that occasional false alarms were tolerable due to the risks associated with the loss of materials known to be present. In addition, access to such facilities is carefully controlled and personnel understand that occasional delays are a necessary inconvenience to avoid loss of material. Thus, when an alarm is triggered, traffic can be stopped until *all* vehicles in a queue can be cleared individually. With the migration of RPMs to more wide spread use due to fears of nuclear terrorism, their short comings have become more problematic. When used at complex multilane facilities such as border crossings, there are issues associated with a single strong source causing alarms on multiple lanes even before the source vehicle reaches the head of the queue in its own lane [3]. Such an occurrence can overwhelm the secondary response capabilities of the facility, leaving operators two undesirable options, slow all traffic to a rate that can be sustained by secondary inspection, or allow vehicles with an alarm to proceed without further inspection.

Another issue associated with standard portal monitors is background suppression. A vehicle in a portal shields the detectors from background radiation, lowering the count rate seen by the detectors [4]. This means that an alarm level set based on the no-vehicle background is not as sensitive as one might think from first principles. In essence, a source in the vehicle must first return the rate at the detector to the background level and only the extra radiation above this level counts toward tripping the alarm. While algorithms exist to try and compensate for this effect, the problem is that the degree of shielding varies

from vehicle to vehicle so that full compensation is not possible. Further, once a detector is large enough to see such a depression, making it larger does not improve the sensitivity since the depression effects are larger than the counting statistics variations normally thought to limit sensitivity [5]. Finally, many normally occurring radioactive materials (NORM), such as terracotta tiles, cat litter, and residual gas in propane trucks, are distributed throughout a vehicle while sources of concern are likely to be point-like in nature. Hence, alarms based on overall count rate will lead to needless alarms on such materials.

3. IMAGING SOLUTION

To overcome the various issues discussed above, this project sought to apply radiation imaging to the problem of scanning traffic for illicit radiation sources. Imaging was deemed suitable because it can distinguish point from distributed sources, because it can localize a source to a given vehicle, and because, as shown in the Large Area Imager stand-off-detection search instrument (see Fig. 3) [6], the approach overcomes the systematic issues associated with baseline depression and other causes of systematic variation in the radiation levels (change in weather such as rain, other sources in the vicinity, and shielding effects due to other vehicles in the area.) To this end, we elected to use the coded-aperture imaging approach [7] that generates images by using a shadow mask between a position-sensitive detector and the region to be imaged to encode the image as a spatial variation in counts across the face of the detector.



Fig. 3. Large Area Imager search instrument is housed in a 16 foot trailer (right) and the instrument itself as seen from the back doors of the trailer (left). A detector array runs down the center of the system with a shadow mask on either side.

The technique was developed by the astrophysical community and has the advantage that the uncertainty associated with each pixel of an image is known from first principles [8]. This means that search sensitivities are well characterized and can be quantified at the surface of the target vehicle based on the radiation image obtained by the system. (As with any passive system, the effects of shielding within a vehicle remain unknowable and are a reason that a layered approach must be used.) This allows one to knowledgeably trade sensitivity for false alarms as might be desirable if intelligence were to indicate that a source may be in a specific region.

An additional advantage of the coded aperture technique is that it provides background subtracted spectra for each pixel of the image, independent of any other sources in the image or in the vicinity of the system. In essence the system is continuously measuring the background with the detector region that is behind the closed mask pixels, while measuring the background plus foreground (source) radiation with the pixels behind the open mask elements. This is demonstrated graphically by the data shown in Fig. 4 where the RST results for a vehicle with two sources is shown. The spectra from each source region shows the correct spectrum for that source, while the spectrum from off-source pixels show no net spectral content!

In a classic RPM each source in the vehicle acts as a source of systematic background for the other and for no-source regions of the vehicle.




Fig. 4. Results for a vehicle with two sources. The spectra for the three cursor locations are shown on the right. [9].

To successfully apply radiation imaging to moving vehicles requires that one compensate for motion blur as a source moves through the imager's field of view. This process is shown graphically in Fig. 5. One can collect the gamma-ray data into a series of short-exposure images that “freeze” the motion of a passing vehicle. Each of these images can be thought of as a single frame in a movie that shows the passing source. A sufficiently strong source will appear as a hot spot that changes location from frame to frame. To maximize the signal from a source and enhance its detectability, one would line up the peaks from the separate images before adding the images together (i.e. one compensates for the motion blur with a digital offset of the images.) In this way, weak detection in each of the single frames becomes a much stronger signal in the single, motion-compensated image.

While the procedure is straight forward, the problem is that most vehicles will not emit enough radiation to show up as a source in any of the frames so that one does not know how much the radiation image must be shifted to compensate for the motion. This problem is solved by adding a visible-light-based target-acquisition-and-tracking (TAT) system that automatically finds passing vehicles, calculates their location, and passes this information to the gamma-ray imaging system. The gamma-ray imager uses this information to generate and appropriately shift the gamma-ray images to compensate for motion blur.

Detailed descriptions of the various subcomponents of the system and how they function to generate the motion compensated images have been published in the open literature [9-14] and are summarized in the following section. The basic idea is that we have coaligned video and gamma-ray imaging systems and the data from the video system is used to select the gamma-ray data used to generate each movie frame, and to determine how much that frame must be shifted before adding it to the final composite gamma-ray image of the vehicle. This composite image is then searched for pixels that exceed a selected threshold level (typically 5-sigma or one alarm in approximately two million pixels searched.) Greater technical detail will only be given as appropriate to aid the discussions on how the team responded to various technical challenges as the project developed.




Fig. 5. A source in passing vehicle proceeds from left to right in a series of short-exposure gamma-ray images (left.) By shifting the images so that the peaks line up (center) one can add the data to get a single image with a much higher signal to noise ratio (bottom right).

4. POC INSTRUMENT

Development of the POC instrument started with a design phase. This included a survey of extant, turn-key, TAT engines that might be applied to the problem. Two classes of commercial systems existed, those that simply counted traffic passing a given point, and those that provided tracking of a specific vehicle manually selected from the stream of traffic by an operator. None of the extant solutions both automatically identified individual the vehicles and provided their locations as they tracked through the field of view. The design work also included establishing the basic problem that would be addressed by the POC system. Two concepts were explored: developing a rapid-response instrument for monitoring freely flowing traffic from the side of a roadway or a system designed for use at classic multilane choke points. Based on consultations with DNDO, the former was selected as an instrument that would provide the most utility to DHS. The work was naturally broken into the visible-light tracking design and the design of the gamma-ray imaging hardware. It is interesting to note that while the hardware development was clearly dominated by the latter, the intellectual challenges remained with the former.

Once the basic problem to be addressed (rapid-deployment, roadside detection) was selected, design of the gamma-ray imaging system was highly leveraged by prior work on the LAI, since the geometry of the two concepts was largely identical (a source passing to one side of the imager in a field of view orthogonal to the relative motion between the two). This meant that the bulk of the gamma-ray system design (detectors, electronics, basic data acquisition software) were complete and we had only to select the requisite detector area, the number and detail of the imaging modules, the system field of view, and shielding. These were straight forward tasks and we were able to utilize the design and simulation tools from the LAI and Proof of Feasibility (POF) prototype development.

For the Video system the bulk of the work was in the TAT software and for most of the project we adopted and maintained a two-strand approach that included an offline concept testing strand and an instrument strand. The former was used to test ideas in a rapid turn-around software environment, while the latter developed those ideas found to be viable into the code used by the system. This approach was developed

during the POF phase of the project. We also carried forward fundamental tracking approaches validated in the POF prototype: using a point-based segmentation approach and definition of vehicle bounding boxes and computation of detailed tracks only after a vehicle left the instrument's field of view. Key mechanical video design issues for the POC included the cameras to be used, the design of a stable stereo imaging system (once stereo vision was selected as an important observable), and the height at which the video cameras should be deployed to minimize occlusion of small vehicles in the far lanes by tall vehicles in the near lanes.

Finally, we addressed the question of the overall system architecture that would be used to control and combine data from the video and gamma-ray instruments into a single output stream. One important component in the system was the ability to collect all of the data in its most basic form, i.e. as it came in from the sensors. This approach was taken because it would allow conceptual and software development to be continued by replaying the saved data without having to have the system in the field collecting data. In the end, this approach saved both money and time since data collection requires significant manpower. This was particularly beneficial in the early days as the reliability of the system was not always as high as desired. However, the large amounts of data collected due to the raw video streams from a total of six video cameras did prove to be a challenge. It contributed to some of the growing pains in the early days of data collection where a good fraction of the time was spent keeping the system operating, as opposed to collecting data. Some of this was necessary software development that had to be done in the field, but even a significant fraction of the software debugging was done in the replay mode.

Gamma-Imager Design

As mentioned earlier, the gamma-ray imager design heavily leveraged the work on the LAI. Most of the information on that system has been published in the technical literature [6, 15] and will only be summarized below with the bulk of the discussion focusing on the most significant differences between the LAI and the RST gamma-ray systems: the choice of mask pattern and field of view, the need to synchronize the gamma-ray data with the visible-light cameras, and the specifics of data handling required to generate images based on information from the video system.

Gamma Imager Optimization

There were a number of instrumental parameters that were explored to design the final gamma-ray imager system. For a system deployed by the side of a roadway, the primary system attribute is the field of view. This determines how long data is integrated during a vehicle transit past the system and it is easily seen that there is an optimum value. Since all radiation detectors receive a background flux, the goal is to minimize the background collected while maximizing the signal from a source. Since flux from a source only reaches the detector when a source is near the instrument, if the field of view is too large, then one collects only background during much of the integration time. On the other hand, if the field of view is too small, then one does not collect as much signal flux as possible. At either extreme the signal-to-noise ratio for the measurement is less than optimum.

To optimize the field of view and to determine the expected performance, we relied on software that simulated vehicle motion past a stationary imager. The vehicle (with or without a source) motion was simulated using very short time steps (1 ms). At each step the code used Poisson statistics [16], to determine the number of counts expected at the gamma-ray imager based on the vehicle's instantaneous location, the strength of the source in the vehicle and the background rate. The source strength reaching the detector

was adjusted to account for atmospheric attenuation.¹ Once the expected number of counts from the source was calculated, that number of events was randomly distributed across the detector array and mapped back to the source location within the target vehicle. The line of site from the detector to the source was used to determine if the radiation would encounter a closed or open element of the coded-aperture mask. Only those events reaching the detector through open mask elements were recorded as hits by the system. Similarly, the expected number of background events was also distributed across the detector array, however all such counts were retained. The millisecond step interval is the same time resolution as provided by the gamma-ray imager and is also much smaller than the dwell time in a single gamma-ray image pixel (GRIP) by a vehicle at even the highest velocities. This provided a simulated data stream that closely mimicked that expected from the actual hardware and was used to develop the data handling codes at the same time that system optimization was performed.

The field of view optimization is essentially similar to that performed for the LAI and results have been reported in the literature [6]. In short, we found that independent of background rate, if we focused on situations where a source was reliably detectable above the background (~ 5 sigma) the SNR as a function of field of view has a broad maximum around 90 degrees and the optimum field of view increases only slowly with range to the target. The SNR is also found to fall off more slowly for wider fields of view than for those that are too small. Thus, we conservatively selected to use a field of view suitable for the greatest distances expected to be encountered. That distance was selected to include 5-lanes of travel and the width of the road shoulders or ~ 72 feet.

In addition to determining the optimum system parameters, one of the key benefits of the simulation code is that it provided a data stream of gamma-ray events that was very similar to that expected from the final instrument. This allowed us to develop and test much of the gamma-ray imaging code early in the project, greatly easing integration of that part of the system.

The gamma-ray imager field of view is closely related to several other design parameters, including the focal length (spacing between the mask and the detector), the size of the gamma-ray pixels, and the number of pixels in the detector array. These parameters are largely governed by the size of the detector pixels used. This was set early in the project at the same size as the detector elements in the LAI, since these had demonstrated excellent performance in that application. While minor changes to the detectors might have improved the overall performance of the system somewhat, we eschewed this secondary goal to focus on the primary project research: motion compensation of the gamma-ray images. Once the detector size was fixed, the mask element size was largely set as well, leaving only the number of elements in the base coded aperture array and the focal length. These were determined by the desired sensitivity (total active area) and the minimum spatial resolution for a vehicle in the furthest lane. The former was set by the goal of detecting an unshielded 1-mCi source in a vehicle travelling in any of 5-lanes of traffic at speeds up to 70 mph, and the overall cost of the detectors, while the latter was set to give at least two resolution elements on a sub-compact vehicle in the furthest lane.

¹ Most of the optimization work used the 662-keV emission line from a 1-mCi ^{137}Cs source.

In addition to the simulation code used to set imaging parameters, we also used standard radiation transport codes (MCNP [17] and GEANT [18]) to explore the stopping power of the detectors, mask elements, and to design shielding. For the first two we looked at the response of the system to radiation at a series of energies from 100 keV through 1500 keV. We validated that our detector sizes and mask element thickness (4 cm) were suitable for the entire range of interest. For the shielding design, we generated background from a thick disk of material below the system populated with appropriate levels of radio-isotopes found

in nature. The leakage flux reaching the detectors through different shielding thicknesses and configurations were then explored. Because the background varies from location to location, a precise background model was not required. Rather, we looked at the fractional reduction in background in different energy bands covering the same overall range. From these simulations, different thicknesses for shielding below, behind, on the sides, and above the system were selected. The final imager properties are given in Table 1.

Gamma-Data Acquisition system and Data handling

In addition to reusing the basic design of the gamma-ray detectors from the LAI, we kept the overall design of the data acquisition system as well. This allowed us to leverage all of the work that had gone into that instrument as well as the bulk of the data collection software. A block diagram of the overall system is provided in Figs. 6 and a detail of one gamma ray channel in Fig. 7.

Table 1.
Gamma Imager Properties

Detector size	4.2 x 4.2 x 40 cm ³
Detector material	CsI(Na)
# detectors per trailer	27
# trailers in system	2
Event time resolution	1 ms
Coded aperture pattern	Rank 23 MURA [16]
Focal Length	63 cm
Nominal Energy resolution (at 662 keV)	9% FWHM full array
Mask Elements	4.6 x 4.6 x 100 cm ³
Mask/Shielding Material	Pb
Detector back shielding Back	2 cm
Detector side shielding	1 cm
Detector bottom shielding	1 cm
Mask-detector gap shielding (top and bottom)	1 cm

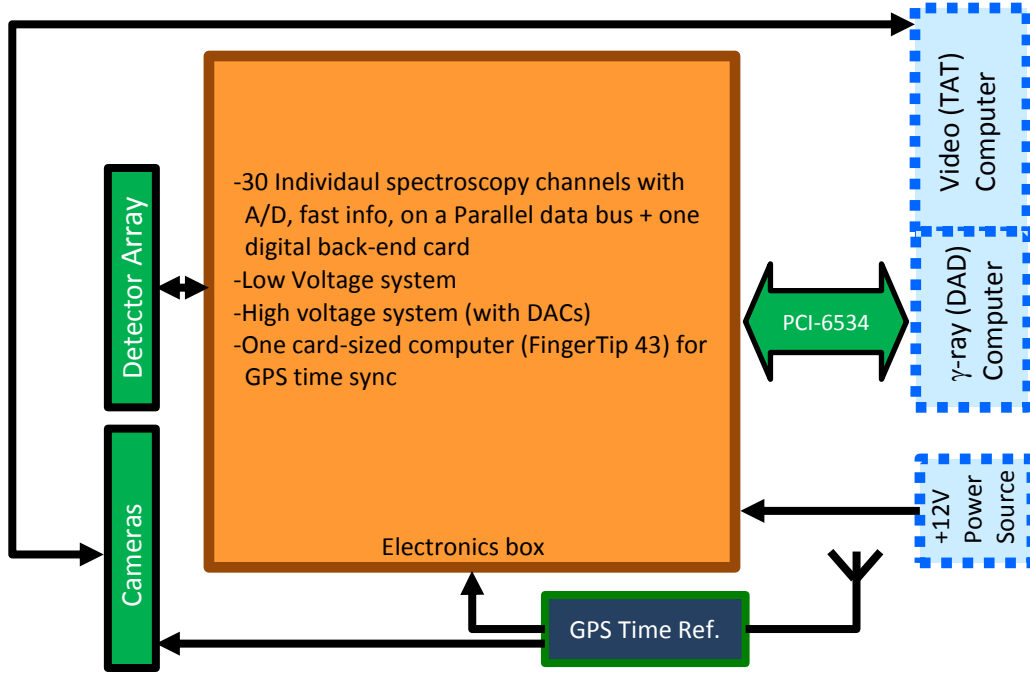


Fig. 6. Gamma-ray system block diagram

Each of the 27 gamma-ray detectors in a trailer comprises a $4.2 \times 4.2 \times 40 \text{ cm}^3$ CsI(Na) crystal read out by single Hamamatsu R980 photomultiplier tube (PMT). A preamplifier in the base for each tube generates a high level signal that is sent to a dedicated custom acquisition card for that tube. The card includes a full nuclear processing electronics chain from the front end all the way through digitization. Both a fast (triggering) and a slow (spectroscopy) channel are included in the card. Each card also includes a small FPGA used to control various parameters (such as digitally variable gain, trigger threshold, etc.). The cards are connected to the readout system through a mother-board (each card has its own digitally addressable slot) that has a parallel bus connecting the cards to a readout card. That card includes an FPGA that reads out the cards and processes the events in a time ordered fashion into blocks of events that occur within the same millisecond. Each event includes an ID (detector number) and the energy of the event. These are transmitted to the gamma-ray data collection computer (referred to as the “Data Acquisition Device” or DAD, based on earlier development from the LAI) via a National Instruments 6534 card [20]. More details on the system can be found in Ziock, et al. [6, 9, 15]. In addition to the gamma-ray events, the system accepts digital inputs and these are used to entrain the clock pulses used to synchronize the gamma-ray and video systems as described in detail later. Each non-gamma-ray digital input type is given its own unique readout ID, with the energy data replaced by the scalar values of the number of pulses received since the last reset of the scalar associated with that line. Associated with the clock pulses are communications via an RS-232 interface to a small computer running a real-time version of Linux that provides periodic information on the absolute time and Global Positioning System (GPS) location. These are also entrained in the data stream as different event types.

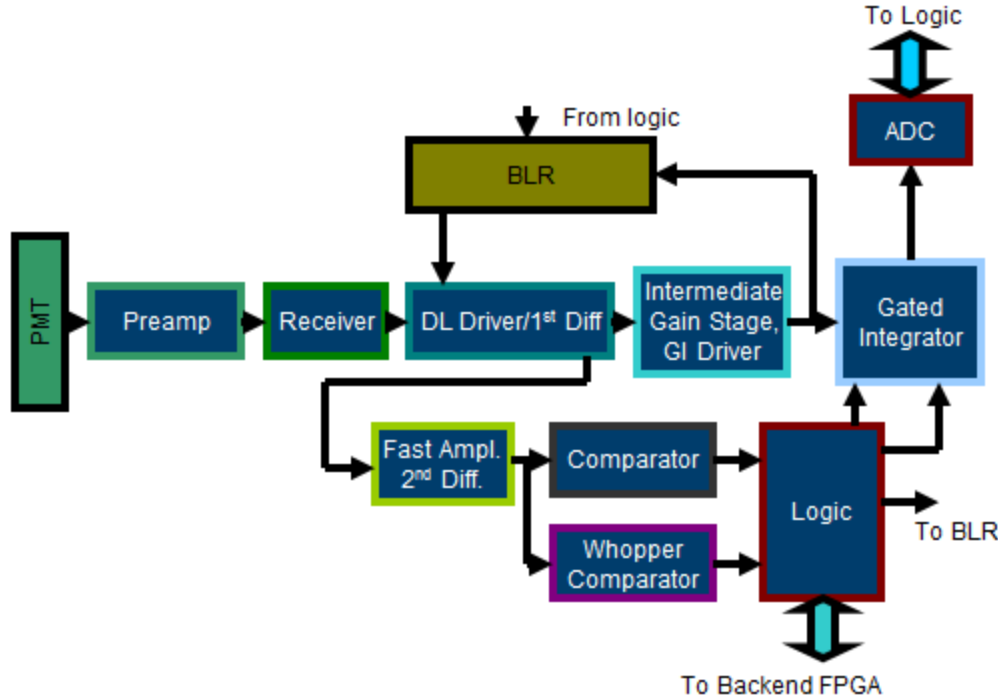


Fig. 7. Single channel of the Gamma-Data Acquisition System. The output from each PMT goes through a preamplifier (at the PMT base) and is buffered at the dedicated card for that tube (see text). On the card are slow shaping components including a baseline restorer (BLR) and differentiation (diff.). The output of this channel goes to a gated integrator. The integrator is triggered by a fast electronics chain that includes additional differentiation (diff.) and a check for very large signals (Whopper) such as those from cosmic muons that overload the channel and result in significant deadtime. These are flagged by the logic.

Data received by the computer are immediately written to disk without processing, allowing playback at a later time for continued system development without requiring system deployment. The software then parses the data stream into a buffer structure that maintains the time-ordered nature of the events. The structure forms a ring buffer that is large enough to hold several minutes of data (see Fig. 8) and is organized around the time structure of the data stream. The data arrives in millisecond blocks and as it arrives, the data is parsed into individual events that are saved in an array with all of the other events from that millisecond. The millisecond arrays are then saved in collections of 50, one for each millisecond in the interval between visible-light triggers.

When a visible light trigger event is received, the code starts a new collection of millisecond arrays. In this way the time of an event is provided by the data structure itself, a feature that is used to generate the time-based, vehicle-specific, gamma-ray images. Gamma-ray images are generated based on video events received from the Master computer. The events are generated based on the times (in units of interpolated fractional video frames) when the front of a vehicle crosses from one gamma-ray image pixel to the next. Upon receipt of a video event, the fractional frame number is converted to milliseconds and the gamma-ray computer searches the ring buffer for the starting (when the vehicle enters a GRIP) and ending (when the vehicle leaves the GRIP) frame numbers. All events between these two times are then used to gener-

ate an instantaneous gamma-image that is shifted to account for the location of the GRIP, and added to the vehicle specific image.

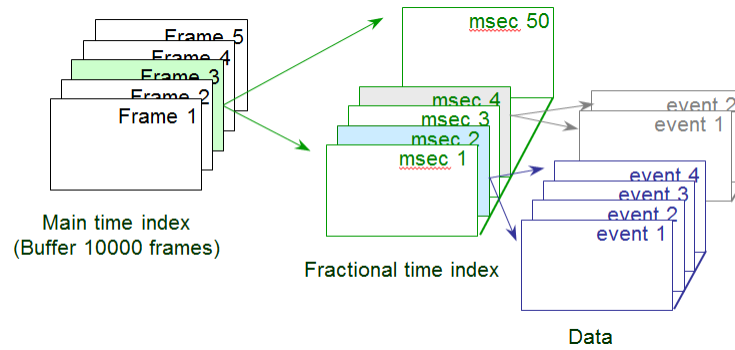


Fig. 8. Data structure for the gamma-ray events. The events are grouped by frame, each frame having 50 possible millisecond buffers.

Gamma-Ray Image Generation

The data handling was conceptually developed for the POC instrument based on the proof-of-feasibility work. For clarity and ease of presentation, the following description includes the steps of the final image generation routines developed for the system. This presentation avoids confusion that might result as changes were made to the code, instead significant differences will be pointed out in the different stages/tasks of the project based on the final solution.

The primary complication of image generation is that it must account for the different angular velocities of sources at different distances (ranges) from the detector (the source location within a vehicle is unknown). Consequently, the video system breaks the vehicle into different range slices (Fig. 9) and generates consistent video events for each slice. The slice widths are the same as those of the many “lanes” (not the width of travel lanes) into which the roadway is divided by the video system. The video events from different slices are treated separately by the gamma-ray software, so that a separate gamma image is formed for each slice. The intensity of a source will be greatest in the slice that correctly matches the angular velocity of the source through the imager field of view and this information provides the location of the source within the vehicle perpendicular to the direction of travel. To account for the fact that vehicles can change lanes, each video event gives not only the current lane of a slice, but also the furthest lane that that particular slice enters during its transit. In processing, the instantaneous gamma-ray images are generated for the slice at the correct distance, but then immediately rebinned to match the resolution of an image at the furthest lane that slice travels. In this way imaged data is always moved to the lowest resolution plane, based on the concept that this will provide the highest and sharpest signal (i.e. if the images were moved to closer lanes then a point source would have to be distributed to many pixels since its location is not known with more precision.)

Image generation is performed using precalculated response functions (event images) of the imager to an event in each of the individual detectors. Because the magnification of the mask on the detector array is a function of source distance, a response image is calculated for each of the different roadway slices. To generate these, a single event is placed into one of the detectors and then rebinned into an ideal detector array that oversamples the mask pattern by a factor of two, based on the distance to the source lane. This

rebinning detector array is then used to create the event image using standard cross-correlation techniques [7]. From each such image, additional event images are formed by geometrically rebinning the image to all further lanes used by the system. These additional event images are used by the system to handle lane changes.




Fig. 9. Transit of a vehicle through the system. The vehicle is broken into slices the width of the RST lanes (blue). These are distinct from the actual roadway lanes (yellow). A track is generated for each slice. The track for the top slice of the vehicle (green) is shown in light green. The image from the lower trailer will be binned to the resolution of furthest lane the slice reaches (red arrow).

At the same time that the event images are generated, a similar procedure is used to calculate the uncertainty coefficients associated with the single gamma-ray in a detector. For this, the squared fractional contribution of each of the single photons to each of the final response image pixels is calculated as well. This is used to calculate the variance of that photon to the overall variance in a given image. This information is used to determine the significance of the signal in each pixel of the image.

During a vehicle transit, the gamma-ray system generates the instantaneous gamma-ray image associated with a video event by taking all of the events from the buffer associated with the appropriate time interval. The events are then projected singly onto the composite image for the video event by selecting the appropriate event image based on the detector associated with that event (i.e. the appropriate event image from the set of rebinned event images generated from the event image at the current lane the vehicle slice is in.) Note that each event is added to two images for the vehicle, a composite image including all events, and a separate image for the appropriate spectral bin of the event. This approach provides the ability to rapidly generate spectral images and is one of the primary reasons for using precomputed event images and an event-by-event projection approach [21]. Finally, the system also generates a variance image using the uncertainty coefficients.

When a vehicle leaves the system field of view (i.e. receives a retirement video event,) the final gamma-ray images from each of the trailers are generated and sent to the Master computer where an additional composite image is formed. To combine the different resolution images for each vehicle slice from opposite sides of the roadway, the same geometric rebinning approach (always generating the image at the

lowest resolution from the system) is used. Here the data are combined and the instrument searches the composite images of all of a vehicle's slices and finds the one with the greatest statistical significance. The significance for each pixel is calculated by dividing the counts in the image pixel by the square-root of the corresponding variance image pixel. An alarm is generated if the significance exceeds five. Note that the availability of the full data cube (i.e. an image for each spectral bin) means that images associated with different spectral regions can be generated by simply summing the appropriate spectral bins. This capability means that automated alarms can be configured for different isotopes. Although such alarms were not enabled during the project, the user interface did allow for selection of manual regions of interest to explore sensitivity versus energy windows.

Late in the project, the LLNL Radio Nuclide Analysis Kit (RNAK) isotope ID software was added to the software suite. This code uses a principle component analysis approach, coupled with the known response function of the detector, to determine isotopes that are present in a given spectrum. This analysis is available manually to the user. To apply it, the user selects the pixel of the gamma-ray image for which the analysis is applied and "pushes" an on-screen button. The spectrum from that pixel is sent to the algorithm.

Video system design

Mechanically, the design of the video system was driven by several factors. The most obvious aspect of the video system is the tall mast that holds the cameras. This was selected because of concerns about vehicle occlusion where a large vehicle in a lane close to the instrument, blocks a smaller vehicle in a lane further out. By mounting the cameras at the top of a mast, they can see over vehicles in near lanes. However simple geometry indicated that to completely avoid all cases of occlusion would have required a 75-foot-tall mast. Although such a deployable mast could have been purchased, it would have been both considerably more expensive, and a larger deviation from the ~25-foot-tall mast used in the POF prototype instrument. As a compromise, a 50-foot-tall mast was selected, providing some view of almost all vehicle combinations, with only a few cases where a short vehicle in the center lane could be completely obscured by trucks on either side of it. In some cases calculations showed that some of the roof of the shorter vehicle would be visible beyond the taller one, and in these cases there was the issue of how to determine where each vehicle ended.

The second influence on the video system came from the use of point-based segmentation carried over from the POF prototype. This was preferred as it did not require obtaining vehicle-free background images and was able to handle rapid changes in lighting as occurred when clouds obscured the sun. In the POF, the locations of the points used for segmentation were projected to the plane of the road surface and then clustered to form vehicles and subsequent tracking. With the possibility of the tops of neighboring vehicles merging (as mentioned above) and general issues associated with projecting a point from a 2D image to a 3D surface, it was decided to use stereo video data. In this way the 3D coordinates of each point used for segmentation can be determined; theoretically improving both segmentation into different vehicles, and the accuracy of the tracking. To generate the stereo images, a two-camera video system was needed. Given the range to the road surface from the top of the mast a 2-m baseline was initially selected for the stereo pair. This was based on an initial uncertainty goal of $1/10^{\text{th}}$ of a gamma-ray pixel. (Near the end of FY 2010, a redesign of the video camera system was performed to include weatherization; at that point a smaller 1-m baseline was selected with one of the advantages being to reduce the difficulty in mechanical handling.)

Finally, one of the requirements for a fieldable prototype was rapid deployment and minimal disruption at the site of the deployment. This meant that alignment of the video and gamma imagers, a process that requires the use of a gamma-ray source in the overlapping fields of view of the two systems (i.e. in the roadway) could not be performed *in situ*. To achieve this, a careful alignment between the two systems was performed off-line using a source in both fields of view simultaneously. A third video camera on the video rig is used to reestablish the alignment in the field. The extra camera was optimized to view a series of visible targets mounted to the top of the trailer. As the alignment between the targets and the gamma-ray imager are fixed, the targets provide the fiducials required to reestablish alignment in the field. They were also used to compensate for mast motion on a frame-by-frame basis.

We note that there are actually four independent coordinate systems that must be coaligned for the complete system; those of the two gamma-ray imagers and those of the two stereo imagers. While off-line procedures and the addition of roof targets allowed alignment of the gamma-ray and video systems within each trailer, there is also the additional requirement to reconcile the coordinate systems from the two trailers and this must be done at the deployment site. To accomplish this requires only that the transformation from one video coordinate system to the other be determined; so that no radiation sources are required. This is accomplished by finding simultaneous points in the images from the two sides of the roadway. While the process was automated, at the conclusion of the project in FY 2012, the method of choice was manual selection of points. The process of aligning the two sides of the road is referred to as the REMLOC (remote-to-local) procedure, and also includes specification of several other run-time parameters such as defining the extent of the roadway.

Camera System

The different video systems used over the life of the project are shown in Fig. 10. The final video stereo rig has three Gevicam [22] 1032 x 779 pixel color cameras. One camera – the alignment camera - uses a 12-mm focal length lens for a relatively small field of view and is mounted to image the rooftop of the trailer; the other two cameras use 3.5-mm focal length lenses and are mounted 1 m apart to generate stereo data. All three cameras are triggered by a common signal at 20 Hz and also feature a frame counter which is reset every second to allow an extra synchronization check. The alignment camera is used to determine the instantaneous position of the stereo rig as the mast shifts and sways and this is used to transform the coordinate system of the video imager to that of the gamma-ray imager. The cameras all obtain color data in the Bayer format (Fig. 11) where each single byte pixel is filtered with a red, green, or blue filter. A standard 24-bit, RGB image can be generated for each pixel by interpolation of the color at the pixel in combination with values from surrounding pixels. The 8-bit Bayer images are obtained from the cameras and stored to disk in this form to save transfer and write times. For subsequent processing, however, the values are converted to 24-bit RGB or 8-bit gray scale values.




Fig. 10. Progress in the video hardware. Left shows the camera used in the POP mounted at the top of its 25 foot tall mast. Center shows the original 2 m system mounted at the top of the 50 foot tall mast. The cameras are under the white rectangular sun shields with the stereo pair at either end of the 2 m pipe. Right shows the final 1 m baseline system with the cameras mounted in weatherized enclosures.

Video data handling

The video data handling was initially designed for the POF instrument as a pipelined process where different processing threads ran in parallel and handled data frames one by one, passing them to the next stage of processing through shared data buffers that were synchronized between threads. This same architecture was used throughout the project, as it has the advantages of being easily parallelized and scalable, can easily be adapted and extended to new processing needs and requirements, and will map well to other computational platforms for future development efforts.




Fig. 11. A segment of a Bayer color sensor showing the filter pattern.

As in the gamma data handling description, here we focus on the steps of the final video handling developed for the system, and we avoid confusion as code changes were introduced by only pointing out significant differences over the course of the project.

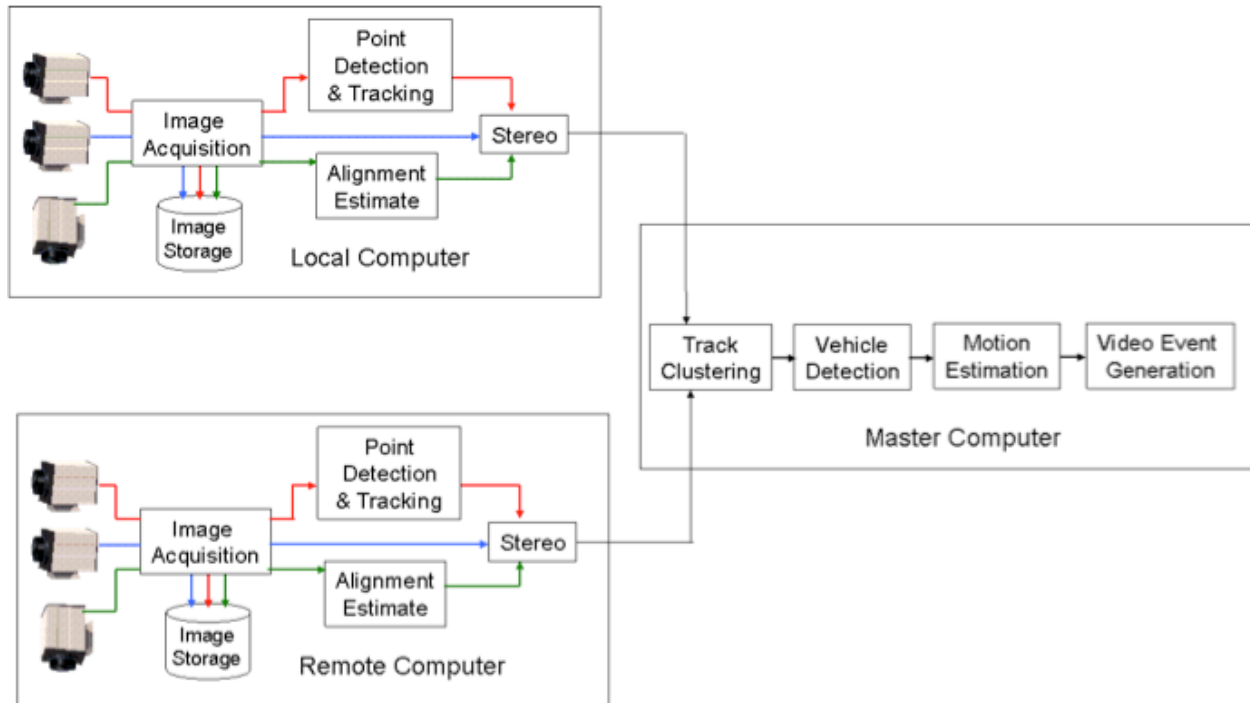


Fig. 12. Data flow for the overall video processing.

Processing Pipeline

The individual steps of the processing pipeline are shown in Fig. 12. The video system starts when the computers boot up and after initial camera connection testing, an Ethernet link is established to the “Master” computer and communications are started. When processing begins, the TAT software starts a set of processing threads that undergo initialization and then prepare to acquire video data. The initialization procedure includes transmission and storage / usage of “scene-specific” parameters, such as the transformation between local and remote coordinate systems and roadway location that are generated by the REMLOC procedure. When acquisition and processing are running, the data is buffered between pipeline steps using shared memory areas. Due to the computer operating system (Windows XP), the code is only quasi-really real-time, in that each and every frame is not guaranteed to be processed in the 50 ms available between frames; rather, the average or composite processing is less than 50 ms per frame, meaning *on the average* the system is operating in real-time.

Acquisition

To allow for overall time synchronization, data is not acquired until the gamma system begins triggering the cameras. Once the trigger is started, data is read by the Acquisition thread which collects data from all three cameras and checks that frame numbers are consistent. It also monitors auxiliary data such as temperature. The thread bundles the data together for disk storage as a “raw video file” (.RWV) that is created uniquely for each processing run. Because disk write speeds over long periods of time (>20 minutes) was a limiting factor throughout the project, the data is written to disk asynchronously in large memory batches of 11 frames. Nevertheless, the performance of the SATA disks used during the project degraded to the point they were unusable when they were approximately half full. Consequently, the thread also monitors the rate at which data are written to disk to forestall issues during a run. Finally, the acquisition thread sends the different images from each camera to the next three threads in the pipeline: the alignment

image goes to the “Alignment Camera Transformation” thread, the reference image goes to the “Point Detection and Low-Level Tracking” thread, and the stereo image goes to the “Rectification” thread. These are described below.



Fig. 13. Image of rooftop from alignment camera showing bulls-eye targets and corner fiducials

Alignment camera transformation

The Alignment Camera Transformation thread processes the alignment camera images, starting by locating the two “bulls-eye” targets at the corners of the trailer roof (Fig. 13). The bulls-eye pattern was selected because it features unique properties such as high space-frequency bandwidth, high contrast, and rotational invariance making it easy to identify and locate. In the first frame of a run, the targets are found using a wide-area normalized correlation search. However, this search is time consuming, taking much longer than the 50 ms allowed per frame, so in subsequent frames restricted windows centered on the pattern locations found in the first frame are used. In the restricted windows, the search is much faster than real-time and the processing quickly catches up to the real-time video sequence after a few dozen frames (a few seconds of data acquisition). After the bulls-eye target coordinates are obtained, the locations of a set of corner targets are found based on their known relative positions to a line drawn between the bulls-eyes. These targets were included in the roof-top design because their locations can be localized with high accuracy. The target locations are then used to calculate the overall transformation between the rooftop coordinate system—which has a fixed relationship to the gamma imager—and the reference camera. This information is then passed to the Stereo Vision thread.

Point detection and low-level tracking

The Point Detection and Low-Level Tracking Thread operates on the reference image, using the “good features to track” algorithm [23] to find points that can be identified with the same points in subsequent images. (The points are similar to ‘corners’.) The “candidate” points are tracked in subsequent frames using the pyramidal Lucas-Kanade algorithm [24]. Note that, due to constraints on the processing time, while existing points are tracked every frame, a search for new points is limited by a parameter currently set to every 4th frame. For each frame, the thread outputs a list of point coordinates along with a reference to the index of the matching point in the previous frame to comprise a ‘track’.

In the initial POC implementation, the tracks were also clustered into groups in this step using heuristic rules. Although these clusters are still being generated, they are now used only for informational purposes.

es, since track clustering is now performed later in the processing pipeline on tracks obtained from both sides of the roadway.

Stereo Vision Reconstruction

Stereo vision includes three main processing steps (“Stereo Rectification”, “Initial Stereo”, and “Stereo Fusion”) to combine the reference image, detected points, stereo image, and alignment transform into real-world coordinates. Stereo Rectification performs a specialized image warping operation on each image from the stereo camera. This operation removes lens distortions and warps the image so that it appears to be obtained from the same plane as the reference image. In early experiments under very sparse traffic conditions, frames with a vehicle present would often require more than 50 ms to process. Although this would cause the system to lag behind the real-time data stream, it would “catch up” when frames with no points were processed after the vehicle left the field of view.




Fig. 14. The image from the stereo camera (right) is only searched in a small rectangular region to find the point matching to the original point found in the reference image (left).

As the system matured and was used in more challenging environments, the need to process frames with multiple vehicles was encountered more frequently, and to improve temporal performance, the stereo vision was split into three threads, one for each processing step. Consequently, the stereo rectification step now runs in parallel with the point detection performed on the reference image. The next thread performs image rectification on the reference image, but only if there are detected points in that frame. The final thread operates on each point in a frame to find a match between the point in the reference image and the same point in the stereo image. This is used to compute a disparity value (3D distance.) To match points the system uses a normalized correlation algorithm [25] to search a region of the stereo image, based on the point in the reference image. To minimize processing time, we leverage the point tracking information for those points belonging to a track to reduce the size of the search region (see Fig. 14). Note that from time to time some points cannot be matched and these are saved and flagged for possible later use if there is a gap in an otherwise continuous track. Under some conditions, if they are needed to repair the gap, these points can be replaced in later processing with a point calculated using linear interpolation on matched points in the surrounding frames.

TAT Output

A final thread on the video computers takes the real-world results and track information on a frame-by-frame basis and bundles this data into “segmentation results” which are simultaneously stored to disk (as the “Segmentation Results” or SGR file) and transmitted via Ethernet to the Master computer. Each frame of the segmentation results file contains a low-resolution, compressed image of the scene, along with the

list of detected points, and the real-world coordinates of those points. The data also include the alignment camera transformation and the auxiliary frame information (frame number, time stamps, temperature data, etc.). The SGR files are much more efficient with respect to disk storage than the raw video files, with a compressed file generally requiring only half the disk space. These files also proved very useful as a resource for debugging problems in the code because they allowed us to troubleshoot the generation of video event files, remote-to-local trailer alignment, and temporal issues such as starting time offsets, without the overhead and complexity of processing the video images themselves.

Dual-road Fusion track clustering

The output of the TAT computers is received by the Master computer where the data from the two sides of the roadway are merged to form a single unified view for tracking purposes. The two systems (Remote and Local) are designed to start on 15-second boundaries (i.e., at 1:05:00, or 1:05:15, etc.) synchronized by the GPS-based timing systems in each trailer. During development we discovered that if both of the GPS units had not locked, there could be an offset of 1-3 frames (50-150 ms) between the Local and Remote data. This offset could be removed by comparing the time stamps on the initial frames received from all cameras. Once the GPS units have locked, the two units start simultaneously.

Merging the data is handled by the Dual-Road Fusion (DRF) processing thread. When data is received, the first step is to put the data into a single coordinate system. To this end the remote SGR data are transformed to the coordinate system of the Local trailer and the points from the two sides of the roadway are combined. (Combination does not mean that the same points are found in both data sets, but rather that a union of the two sets of points is formed.) The module then performs track clustering using the coherent motion region (CMR) algorithm [26] on the x-y coordinates of the tracks. This step finds groups of points that maintain their relative spacing with respect to each other. Each such CMR is identified, labeled, and associated with CMRs found in the previous batch. These CMR sform the basis of vehicle detection as described below.

Vehicle detection, Motion Estimation and Video Event Generation

A vehicle is identified by the CMR algorithm when a sufficient number of tracks have been clustered. The software then assigns the vehicle a unique identification number and monitors it as long as it continues to be found by the CMR code. When it is no longer updated by a CMR batch, meaning all tracks have ended (i.e. the vehicle has left the field of view, or remained stationary for a long time) the vehicle is given a “retired” designation.

Once retired, the data associated with a vehicle are analyzed using a variety of filtering operations to estimate the motion of the vehicle through the field of view. (Currently this processing step is not parallelized, but it should be simple to start a separate thread to operate on the data for each vehicle.) The first step of the motion estimation is to determine the size and mean location of a 3D bounding box for the vehicle in each frame. The different bounding box sizes are used to find an average bounding box to be used in all frames. The average bounding box is then positioned in space based on the location of the mean of the original bounding boxes in each frame, but only after the locations have been filtered to ensure smooth motion from frame to frame. Using the box locations, the time that the front of the bounding box crosses the boundary between gamma-ray image pixels (GRIP) is then calculated. This information is used to trigger a video event that is sent to the gamma-ray imager. As described previously in Gamma-Ray Image Formation, the vehicle is actually split into a series of slices and video events are independent-

ly sent for each slice. Because the GRIP boundaries of the two trailers do not overlap, different video events are generated for both instruments and are based on their own internal coordinate systems. For the Remote trailer, this is done by taking the vehicle motion estimation values and transforming them back to the Remote imager's frame of reference.

The video event includes primary information for use by the gamma-ray systems such as: the time of the event (given in fractional frame numbers since the start of a given acquisition), the vehicle ID, the slice number, the number of slices, the current lane of the slice, the maximum lane the slice enters during its transit through the field of view, the width of the vehicle in GRIPS, and the type of the event. There are three types of events defined by the system, a new vehicle event, meaning a new vehicle has been identified and entered the field of view, a GRIP-crossing event that tells when the front of a slice crosses a GRIP boundary, and a retire event, that says the vehicle has left the system field of view. Additional diagnostic information is passed in the events as well, including: the times of entry and exit into and from the field of view, the vehicle velocity, its location, and various information on the number of points used in the processing.

High Resolution Images for Vehicle ID

One of the issues with the system operation encountered late in the program was how to identify vehicles to emergency response assets positioned with the instrument. While each frame of the SGR file includes a compressed, low-resolution image of the vehicle, these images are only suitable for their planned diagnostic purposes—primarily to show the bounding box on the vehicle to the operator in real time. They are insufficient to identify and interdict any vehicles that cause a radiation alarm (Fig. 15). While the original high-resolution images exist, and are most suitable for identification, getting them to the Local trailer on a continuous basis would overwhelm the system throughput. To address this, an independent thread is run in parallel with the rest of the TAT processing pipeline. It is used to form a high-resolution compressed JPEG image of every sixth (chosen heuristically) frame that is written to the TAT machine's hard drive. In addition, the last 200 images are maintained in a memory buffer by each TAT machine. The images are maintained so that if the Master computer detects a gamma-ray alarm, it can request a particular frame from the TAT machines which would then search the stored images to find the image closest in time to the requested image and transmit it to the Master. For a deployed system this process would be automated and the high-resolution image would be displayed to the operators. However, because the need for high-resolution images only became important towards the end of the project when the system was used with local law enforcement during both the Limited Use Exercise (LUE) and the characterization effort, this level of automation was not implemented. Instead, the set of images on the local hard drive were manually displayed when needed.



Fig. 15. A high resolution image for vehicle ID (left) and “standard” resolution image of the same vehicle (right).

Rig drift monitor

A final thread, the “rig drift monitor,” is included in the TAT code. It computes the real-world coordinates of a set of pre-defined points on the road surface found in a stereo-reference image pair. The process is repeated every 2 seconds (40 frames) and the results are stored in the SGR file data. The thread was introduced as a means of assessing the stability of the stereo rigs, due to concerns that they would lose calibration over time. Although we have implemented the monitor and introduced a message into the run log regarding the positions and the drift, an in-depth analysis has not been completed.

5. PROJECT EVOLUTION

The project was based on a modified version of the standard DNDO program structure with the following stages: proof of feasibility, proof-of-concept design, POC construction, and test and evaluation. Because of the similarity of the gamma-ray imaging needs to those used in the earlier Large Area Imager (LAI) project, we were able to end the POF stage with a simple working prototype that used the LAI as the gamma-ray imager.

The timeline for the project evolved over its lifetime. The final version, with various highlights and achievements, together with an approximate phasing, is shown in Fig. 16. We note that there is some overlap between the different phases. This was partially in response to the unique starting point with the LAI to facilitate a POF demonstration. That allowed a successful POF review based on initial design efforts using LAI results prior to completion of the experimental integration of the visible and gamma-ray imaging needed for the POF experiments. This was leveraged to start the Phase II POC design work before the POF experiments had been concluded, and allowed some acceleration of the schedule from that originally proposed to DNDO. Unfortunately, some of that advance was later lost, primarily due to getting the complex POC to operate reliably.

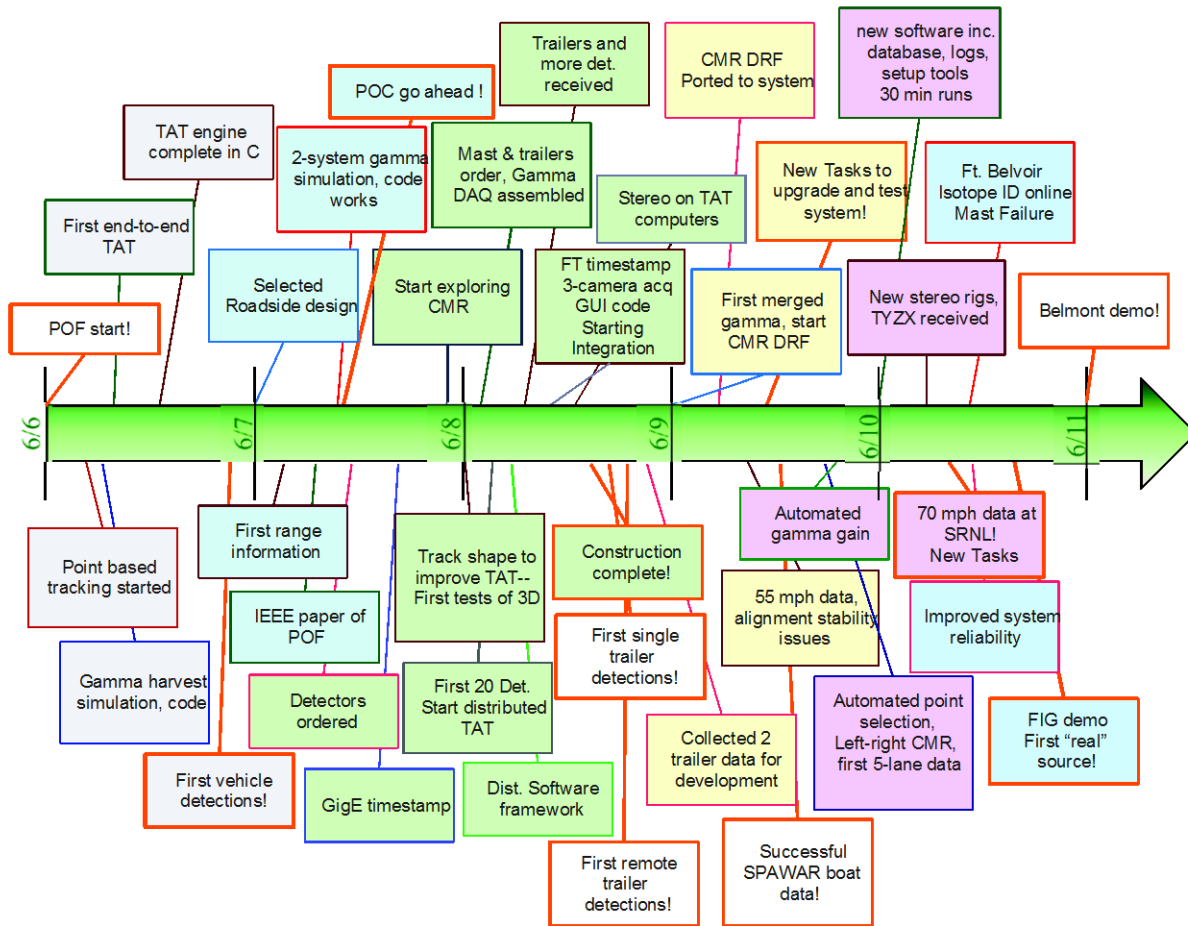


Fig. 16. Project timeline with significant milestones and accomplishments

In the POF system the video images were collected with a single camera mounted on a 25-foot-high mast. The results obtained with this system demonstrated the overall concepts and provided the initial precepts that guided development of the overall system. Noteworthy results included:

- Recognition that the primary challenge of the project was TAT development and from this, a program structure that included a two-part approach to this part of the problem. The two components were a basic research strand that allowed different basic concepts and algorithms to be tested in a rapid-turn-around software environment such as MATLAB [27], and a working prototype strand that developed the basic TAT engine and then incorporated the most promising results from the other strand as needed to improve the tracking results.
- The use of robust point-based segmentation to enable tracking at arbitrary locations without requiring prior images of the roadway that have no vehicles in them for background subtraction.
- Development of simulation codes that generated random traffic flow in different roadway configurations so that the overall system performance could be estimated for different imager design configurations.

Details of the results from the first phase of the project are summarized in a final report and the interested reader is referred to that report (included in Appendix B).

Phase 2: POC design

Phase 2 covered the design of the instrument, leading to an 80-90% complete design for presentation to DNDO in a preliminary design review. The design was then completed (including any feedback from DNDO) to arrive at the final instrument design for use in Phase 3, construction of the POC prototype. During this time we were also collecting and analyzing data from the POF prototype as indicated in the timeline segment shown in Fig. 17.

Important decisions during this period included developing for the roadside (versus chokepoint) scenario, a two-trailer system with one on either side of the roadway, the use of stereo video, as well as compromise to a 50-foot tall mast to overcome some but not all vehicle occlusion issues.

The work during the phase was broken into 4 tasks with two tasks from

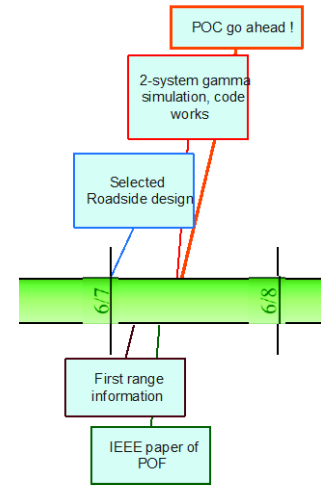


Fig. 17. Major items during the POC design phase.

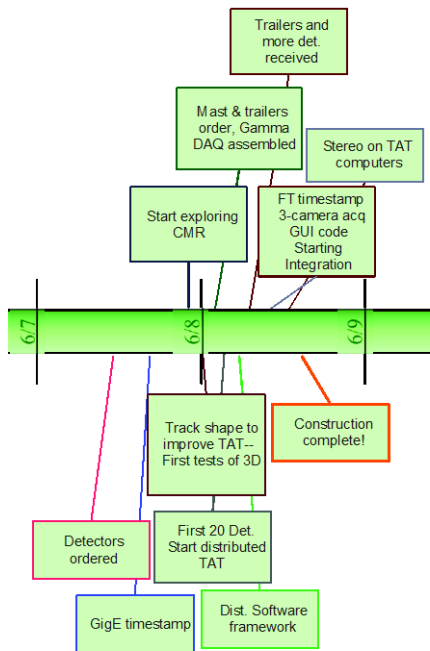


Fig. 18. POC construction phase highlights

Phase 3: POC Construction

Phase 3 covered the construction of the instrument. At the same time, much of the basic TAT software was also assembled. This included developing the initial stereo algorithms; improving and perfecting the data pipeline processing including distributed processing, porting, and improving the track clustering (co-

herent motion region algorithm); and the alignment camera processing. The major highlights are shown in the segment of the project timeline in Fig. 18.

Phase 4: Integration and Test and Evaluation

Phase 4 included integrating the various elements from the construction phase to result in a working instrument. Major highlights of the phase are shown in Fig. 19. The performance of the instrument was then improved in a cyclic process that included collecting data, using this to improve the performance, and then collecting more data with the improved system. It was during this phase that the first major issue surfaced due to the overall complexity of the system. While we had anticipated the difficulties associated with bringing our new work on line, we were not prepared for the continual issues that arose with commercial off-the-shelf (COTS) technology. This included such things as cameras stopping and starting, wireless connection problems, computer problems, generator failures, etc. When folded into the growing pains associated with stability of the TAT software, there was significant time spent debugging as opposed to collecting data whenever the system was run. This situation slowly eased over the course of the development cycle as the software was improved. However, the system was always running near the limits of its capabilities as we saved all of the data it generated, including 3 full frames of video at 20 Hz.

One issue that did impact system development was the prohibition on using the instrument on public roadways without a response capability based on concerns that any detected sources could not be ignored. Without such a response capability—and we were reluctant to use a developmental prototype as a primary detection system for use in public—the system was limited to use on non-public roadways, which primarily meant on the ORNL reservation. Because of this, most of the development was performed on a two-lane side road that had minimum traffic and a posted speed limit of 45 mph (Fig. 20). For heavier traffic flows and to simulate the broader expanse of a 5-lane roadway we found a suitable location on the main ORNL thoroughfare (Fig. 21). This also had a 45 mph speed limit, but during peak times had almost continuous traffic in one direction or the other. The location has two travel lanes and center turn lane which could be used as a travel lane for testing purposes only. To mimic a 5-lane roadway, one of the trailers was parked adjacent to the road and the other was setup in a parking area that allowed establishing the appropriate 5-lane-plus-shoulder separation of a 5-lane freeway. However, onsite usage did limit exposure to different deployment sites and traffic conditions. The importance of a broader exposure was evident during later deployments and demonstrations of the system, particularly during the earliest off-site work since each new site represented unique operational, logistical, and visual challenges. Exposure to each new site broadened our experience with the system leading to both operational and algorithmic improvements to the system.

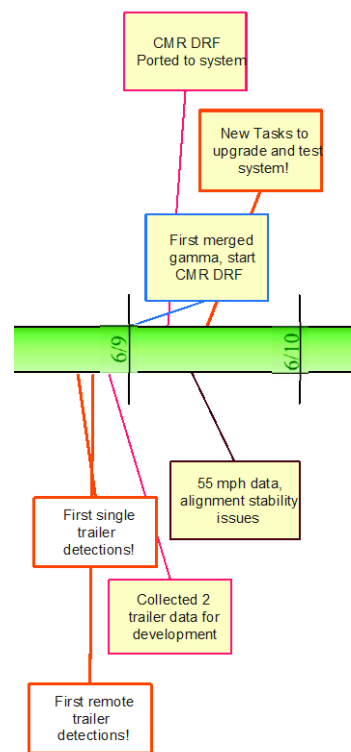


Fig. 19. Phase 4.



Fig. 20. RST setup at the standard 2-lane development site.

To determine TAT performance at higher velocities, we did collect data on a nearby 5-lane roadway (two travel lanes in each direction with a center turn lane) with a posted speed limit of 55 mph (Fig. 22). However the system was run without collecting gamma-ray data to avoid the issue of source detection without the presence of interdiction capabilities. Although these tests did provide data that were useful to develop and test the TAT system, their overall value was diminished because we could not simultaneously run with radiation sources to test overall system performance.



Fig. 21. RST set up on the main ORNL thoroughfare.

Before the end of Phase 4, we were able to pull together test data over a range of velocities and with different sources but we were unable to collect high-speed data on the ORNL reservation. As such, we were unable to complete a full characterization.

Phase 5: Additional Development

With the performance issues uncovered in the system and with the difficulty in getting a full set of calibration data, an additional development phase was added to the project. In the end, this new Phase extended for two years and included upgrades to the system performance, the collection of a set of high speed performance data, demonstrations of the system with vessels, and participation in a Limited Use Exercise (LUE), that provided further opportunities to exercise the system in public. Highlights are presented in Figs. 23 and 24.

Details on the initial tasks, designed to improve the system capabilities, are provided below:

Extended Operations: The goal of this task was to extend the weather conditions during which the system could operate. While the bulk of the instrument is safely ensconced out of the weather inside of the trailers, the cameras, at the top of the mast, are exposed to the elements. A minor element of the work included developing housings for each of the cameras that did not impair the visual performance. A greater effort included adapting and testing the TAT algorithms to handle the different visual morphologies presented by vehicles under inclement conditions. Primarily this includes reflections off of the pavement and the fact that headlights are on.

The task was eventually completed, however an initial housing design that fully enclosed the cameras, relied on conductively coupling the waste heat to the mounting bar through the mounting surface. Unfortunately the heat transfer was insufficient, as discovered when the cameras were deployed during hot weather. We did take advantage of the error to also redesign the camera rig itself. In the new design, the environmental enclosures were also used as a mechanical shield of the cameras, carrying stresses from accidental contact to the mounting bar, independent of the cameras. The goal was to minimize changes in rig alignment over time. This problem was further eased by shortening the length (and also diameter) of the bar. Prior results indicated that the 2 m separation originally selected for stereo resolution was not required. The rig redesign had the additional benefit that it made setting the system up and tearing it down much easier because the new lighter weight rigs were much easier to handle.

Isotope ID: The goal of this task was to take advantage of the spectral capabilities of the system via an automated isotope ID code module. For this we turned to work performed under the SORDS program where the spectral information from the coded-aperture images available from the LAI-based SAIC [28]



Fig. 22. First offsite data collection site.

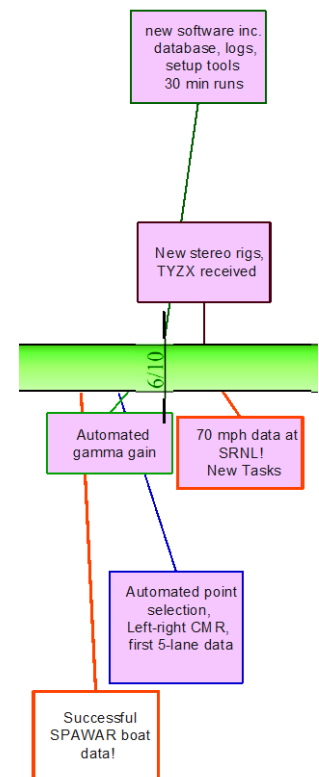


Fig. 23. Improvement and Test and Evaluation

instrument were analyzed using the LLNL-developed RNAK algorithms. The advantage was that we were able to leverage the prior SORDS work that developed RNAK to handle the unique statistical issues associated with coded-aperture data. As it turned out, most of the effort required to make this function correctly was establishing the energy calibration of the different detectors. In particular, there is a subtle problem in how the individual gamma-ray detectors are triggered that means high count rates of low-energy gamma rays can depress the response at higher energies. This made it difficult to find a low energy calibration line that worked for all of the detectors. While the correct fix would have been to change the FPGA code in the gamma data acquisition system, this was not possible within the budget so we ended up using the 511-keV and 1274-keV emissions from a ^{22}Na source. While the isotope ID was successfully implemented and demonstrated during various exercises, it was never completely automated in that the user has to select the source peak in the gamma image from an alarm with the cursor and “push” a software button to request the analysis.

Fixed Portal Study: The goal of this task was to look at how the imaging ap-

proach of the RST could be used to improve the performance of a fixed multilane portal system such as those at border crossings. However, the funds were reprogrammed before the task started to help cover deployment costs such those incurred in the SPAWAR demonstration.

Improve TAT: The purpose of this task was to continue a separate research strand focused on improving the performance of the TAT algorithms. The work was conducted on data collected with the system but in an offline setting where ideas could be easily coded and tried without worrying about the extra complexities involved in adding the code to the working software or impacting the readiness of the system. Several concepts were initially planned for exploration, including: alternate point selection algorithms, developing a True 3D coherent motion region, dense Stereo mapping and detection, and harvesting of occlusion regions. In the end we focused on two concepts. The first was using true 3D coherent motion regions where 3D points measured by the stereo system were projected to the plane of the road and used to perform track clustering. The second concept was to improve calibration methods. This permitted easier ‘in the field’ camera and rig calibrations. Offline dense stereo processing was investigated to estimate the utility of this in conjunction with a dedicated processing unit (see robustness below). The improved calibration methods allowed faster recalibration of the camera rig parameters, particularly the transformations between the stereo / reference camera pair and the alignment / reference camera pair.

Improved Deployability: This was a set of tasks designed to reduce the time required to set the system up at a new location, while at the same time improving the overall accuracy achieved. The four sub tasks included Auto Calibration, which added an automated selection of points from the camera images needed to initialize each trailer’s tracking zone while also providing cross calibration between the two trailers coordinate systems; Continuous Calibration, which continuously monitors the overall system calibration while it is running; and operation with limited or no guy wires, which looked at improving the digital image stabilization of the video images to reduce the impact of mast motion and allow system deployment without guy wires except under very windy conditions. While a variety of improvements were made to




Fig. 24. Limited Use Exercise 1

the system, the project focus prevented full adoption of the different techniques in the final system as tested in the characterization effort and LUE experiments.

Robustness: This task had two sub tasks; the first would add large networked disks to the system, improving data transfer rates and also improving data archival by simply swapping disks for new ones. The solution adopted was a variant of this approach with large removable drives added to the system computers and a networked drive system used for off line backup and analysis. The second subtask was to procure custom stereo cameras from TYZX, Inc. for use with the system. The goal for this was twofold. First, it was hoped that a commercial unit might have less issues associated with calibration stability, and second, the commercial unit had a built-in stereo processor that supplied high density stereo data in real time. This should have simultaneously eased the computational burdens on the TAT computers while providing potential tracking improvement with the dense stereo information. Unfortunately, there were several problems with the commercial units. They were designed to be drop-in replacements for the current stereo rigs but the initial version came with an incompatible triggering interface. While this was eventually sorted out, we then encountered a number of issues in getting the units to record data with sufficient reliability to allow their use in the system.

In addition to these tasks, there was the standard project management task, as well as a task to support additional demonstrations of the instrument.

During the first year of the extended phase a number of major accomplishments were achieved. First, we participated in an exercise conducted by DNDO and hosted by SPAWAR in San Diego. This provided an opportunity to run the instrument on small vessels, in an approximation of a “channel-side” tracker. To make this function we had to modify the TAT algorithms and parameters to track the small vessels. The trial was ultimately successful with the system detecting all of the sources at the maximum speed with which the vessels could run past the pier where the measurements took place (~ 11 knots). While successful, the performance did indicate that more work would be required to develop a truly robust small-vessel tracking capability. Second, between the additional testing and improvements to software we significantly improved the system stability. During testing on ORNL, even during periods of heavy traffic, the system would run for half an hour without running into memory issues. This was deemed sufficient for further work since sifting through runs longer than this would become an issue in any event. As such the system was considered ready for a full evaluation, although as pointed out below, other issues continued to impact reliability.

Further funding was received to perform additional tasks. This included testing the system with a broader range of sources than originally planned (more isotopes and also a look at weak source performance.) The additional tasks also expanded the work on dense stereo in anticipation of getting the commercial stereo rigs operational. The instrument was added to the Limited Use Experiment (LUE) DNDO was planning with the SORDS instruments. This was meant to provide end-users exposure to the instrument and obtain their feedback on the potential value of its capabilities. Finally, the additional demonstrations task was meant to cover participation in some of the LUE exercises and other targets of opportunity that might arise.



Fig. 25. RST at SRNL where 70 mph data were collected.

Shortly after the new funding was received, we tackled the problem of collecting a definitive data set that shows the system met the stated goal of detecting an unshielded 1-mCi source in any of 5 lanes of traffic at velocities of 70 mph. Initially we attempted to perform these measurements on the ORNL reservation (to avoid the need for law enforcement participation required on public roadways). However, after unsuccessful attempts to allow off-hours runs at 70 mph on the ORNL site, we turned to Savannah River National Laboratory as a test venue. They have a stretch of 4-lane divided highway where one of the sides is abandoned. There was also enough space on either side of the roadway to allow establishing an effective 5-lane-plus-shoulder spacing of the trailers. This was the site eventually adopted and a test campaign conducted at that location (Fig. 25). The emphasis of the data collection was on the center of 5 travel lanes as this would present the weakest signal in both detectors. The results (Fig. 26) conclusively show that the system achieves this goal. But even at this late stage some of the data collections were painful. One of the problems was that the system, which had been working fine for periods up to half an hour of data collection, suddenly started having memory issues after only a few minutes of run time. This significantly slowed data collection since we had to carefully stage each pass of the source vehicle with the system start. This was exacerbated by the high temperatures during the campaign which meant that some of the video cameras started to have problems initializing. We finally found that the primary source of the memory problems was due to a new set of hard drives purchased for the campaigns, and how they were formatted. Once resolved, data collection would proceed smoothly for ~ half a day when problems started to recur. This was finally traced to a significant slowing of the write speed to the disks once they were ~ half full. Another example of issues with COTS system was the generators used during the campaign. These were fueled from a truck that came to the site and at some point introduced particulates into the gas tank. This clogged the fuel filters so that the generators would shut down unexpectedly after each refueling (presumably by stirring up the sediments).




Fig. 26. RST performance on different 1-mCi-class sources as a function of velocity.

Overall, these issues arose because of the complexity of the system and are not included as an indictment of building a complex system but rather that development of such systems requires extra care in selecting components with high reliability, the need for spares, and most importantly, the need for extensive trials of the system under various conditions to expose failure modes and weaknesses—many of which can be corrected once identified. For example, during later evolutions we modified the video camera enclosures to allow them to run cooler, while also monitoring their temperatures. We also tested new hard drives for their sustained write speeds and swapped them out before they filled up, and we added extra filtering capabilities to the generator fuel lines.

The plans for the LUE included participation in three activities, an initial demonstration at Ft. Belvoir in the DC area, an extended test at Fort Indiantown Gap, and participation in a real-world event (finally selected to be the Belmont Stakes horse race.) During the initial demonstration, there were a number of issues that surfaced with the system performance such that the live demonstration was erratic. This included a very windy day such that even with guy-wire stabilization the masts moved considerably (normally we would not have deployed during such weather), a recurrence of the clogged fuel filter problem on one of the trailers, and (we found later) a virus on the Master computer that meant attempts to run the system remotely using Remote Access and a long Ethernet line failed. Later in the day we were able to run somewhat successfully with guy wires, albeit with periodic crashes.

Based in part on these poor results, we returned to the issues of system reliability and under the additional demonstration task we performed a number of hardware and software maintenance and upgrades. For the hardware, we had the masts restrung (we were having problems with their deploying), flushed and cleaned the fuel tanks while adding larger fuel filters. For the software, we upgraded the camera code so that connections to the cameras were maintained, rather than started and stopped with each run, set a maximum number of points to track, and improved the throughput by removing some debugging checks.

(This was in response to a problem with the wind induced motion of the masts overloading the software because background points appeared as moving points.) Finally, we found and removed malware from the system. Overall, these fixes were successful and we have had better results with the system since.

With the fixes, the system successfully completed participation in the LUE. The results obtained at Fort Indiantown Gap were excellent and we even saw our first non-planned source in the form of what must have been a medical patient (based on isotope ID). One outcome of this exercise was the clear need to develop a suitable Concept of Operations (CONOPS) to handle source detections. During this and subsequent exercises, we relied on verbal descriptions to alert responding officers. However, the time for operators to recognize an unplanned detection had occurred, get a high resolution image, and look at the Isotope ID was too long in the current system. Also, the overhead view was not always the optimum for finding a particular vehicle and this whole aspect of the system needs to be properly addressed in an advanced prototype. (The goal of this project was to demonstrate that detections could be made and localized to the correct vehicle—it did not include how to interdict that vehicle, which clearly must include timely presentation of positive ID information to responding officers.)

The final LUE exercise was held during the Belmont Stakes horse race. The RST was set up in the vicinity of the race on one direction of a 6-lane divided roadway (it covered 3 lanes—see Fig. 27). The system worked effectively, including detection of a few additional medical patients, until the traffic changed over to stop-and-go conditions. This is not a condition for which the system was designed although straightforward modifications should address this traffic condition.



Fig. 27. RST deployed near the Belmont Raceway.

A separate report on the LUE was written by independent third parties conducting the exercise. It was completed with neither input from, nor review by ORNL. Subsequently, ORNL believes the report has a number of factual errors and that the performance of the RST was faulted under questionable conditions. In particular, the system was found to work very well whenever test sources were used, correctly detecting and identifying the correct vehicle at all times unless the source was so shielded that there was no hope of detecting it. At no time did the RST alarm on a known, isolated, non-source vehicle. However, the system did detect several unexpected medical patients in incidental traffic and for those, the performance in the report was judged considerably worse—in essence because the response teams were unable to validate the medical patient detections. Given that correct performance with known sources generated

neither false alarms nor identified the incorrect vehicle, we consider it highly unlikely that the system would suddenly perform erratically only for those situations where the ground truth was unknown. It should be further noted that medical patients represent very strong detections, further reducing the likelihood of an incorrect detection by the system.

Independent Characterization

Late in the lifecycle of the project, DNDO initiated an independent characterization study of the instrument. This work included two major data collection campaigns, one at ORNL and a second in the DC area. The ORNL data included a series of runs with different sources in different lanes of travel. It also included a set of data where the system was deliberately misaligned to determine the impact of setup precision on the results. One of the prior difficulties in developing the system was a lack of ground truth knowledge of the location of a target vehicle. Prior efforts used manually assisted analysis of the same video data that was run through the TAT software. Since calibration and alignment errors were the same in both instances, this was a bad means of chasing down systematic errors in the tracking information. For the ORNL data collection, we found a low-cost compensated GPS system [29] that could be used on the source vehicle. This provided ground truth data for the vehicle location good to ~ 10 cm (3-sigma) on a real-time basis. Since the data were completely independent of the video system, they helped correct several small issues with the tracking data. One of the biggest issues discussed earlier were errors between the local and remote trailer coordinate systems. With the real-time data we were able to reduce these errors. We discovered that when the two systems started at different frames, we were correctly modifying the video data so that they were time-synchronized, but the corresponding frame offset was not propagating to the gamma event data (so that the gamma data was delayed from one trailer to the other). We also found that the DAD started counting an initial frame prior to TAT triggering, such that there has always been a 50 ms offset between the two systems which has a slight impact on the shift of the image at high speeds (Fig. 28). Finally, we also found some errors and bugs in the video event generation code, which sometimes negatively affected the results.

ORNL conducted some initial analysis of the data acquired in the system alignment characterization study. In particular, we found that a trailer offset of up to 4 meters (the largest offset tested) was not significant in the final gamma detections. The offset was measured both by GPS and by the coordinate transformation created in the REMLOC process. The other misalignments studied were a relative angle between the two trailers. This study was slightly more complicated. In principle we believed the relative angle would be corrected by the coordinate transformation, and indeed the angle could be estimated by the transformation. When the local trailer was oriented with the front of the imager closer to the road at an extreme angle (+11 degrees), we found a significant drop in the detection performance. This was not mirrored at -13 degrees, and was found in both the ground truth data and in the video-based data. Consequently, we believe at the extreme angles, there must be an assumption that fails in the video event formation code.




Fig. 28. Improved performance at 70 mph after removal of gamma/video offset. Without the correction, the source is just beyond the back of the vehicle (center); with it, it is correctly positioned at the rear of the vehicle.

In an indirect fashion, the advantages of imaging were clearly highlighted by one of the side studies of the characterization. One of the external teams looked at the RST data collected during the Belmont Stakes deployment of the LUE as non-imaged data. By combining all of the data, they found a slight elevation in the count rate for intervals when vehicles were in the field of view compared to times that no vehicle was present. This was eventually attributed to a short interval when both trailers saw a significant increase in counts. With this information, we were able to go back to the saved vehicle output results and saw immediately that the elevated rates were due to a medical patient driving through the system. That was logged as a detection by the system at the time of the demonstration and localized to a single vehicle in the imaged data.

Data collected in the DC area included work with vessels on the Potomac, some work with pedestrians at RSL Andrews, and several days' worth of data collection on the outer loop of the DC beltway (Fig. 1). The final analysis is being conducted by independent outside teams and results are not included here. We do note that the system did have some problems during the tests. In particular there were issues with the system tracking for extended periods in all of the situations, but particularly the nonstandard vessel and pedestrian tracking, due to out-of-memory errors. This was eventually traced to a failure of one of the TAT computers only recognizing 2 of the 4 gigabytes of installed memory. Once the machine was replaced, operation became much smoother. However, we did end up with one of the two TAT machines ~10% slower than the other simply by virtue of its age. This machine was barely able to keep up with the level of traffic on I-495, an indication that an upgrade to the latest computer technology (some of the machines are ~ 5 years old) could significantly improve overall system performance.

We also had some trouble with one of the gamma-ray data collection chassis. When they were first assembled early in the project we found that if too many of the system slots were filled, then a voltage watchdog would indicate one of the supply voltages sagged as the system booted so it never got through the initial boot sequence. Since we did not need the number of channels that caused problems, we simply ran with the needed number of boards. Unfortunately, this condition recurred during the trial, but would happen with only a few (4-5) of the slots filled. Since each trailer requires 27 boards, this meant that we could not bring the affected system on line. While the problem was patched during most of the characterization campaign, it did mean that for some of the days we only worked with one trailer (notably when collecting data from vessels, where we could not get on both sides of the channel anyway.) On analysis of the problem after returning to ORNL we found that the original problem was a bad batch of resistors (more than 5-sigma out of specification) and presumably this coupled with aging in the system caused the failure during the campaign. A final positive result from the DC campaign was interest expressed by a

DTRA team when they viewed the system during a demonstration on the last day of the campaign. We are working together with SPAWAR and hope to obtain DTRA funding to build a version of the RST suitable for military operations.

Summary InterAgency Agreement History

The work discussed above was performed at ORNL via three InterAgency Agreements (IAA) between DNDO and DOE. The history of the project is summarized via the IAA's in Table 2

Table 2: Project IAA History

IAA	Amend- ment	Date	Period of Performance (POP)	Tasks/Comments	Amount (\$K)
HSHQDC-06-X-00423		05/31/06	06/01/06-02/01/07	1. Project Management 2. Gamma-ray imager roadside algorithm development 3. Video tracking algorithm development 4. Video tracking hardware acquisition	523.8
	1	02/01/07	06/01/06-04/02/07	POP extension	
	2	02/06/08		Close out	-6.8
HSHQDC-07-X-00181		02/06/07	02/06/07 11/30/07	1. Project management (Phase I & II) 2. Video Gamma Integration (I-5) 3. POF Experiments (I-6) 4. TAT software development (II-1) 5. POC prototype design (II-2) 6. POC Final Design (II-3)	1,373.0
	1	11/30/07	2/06/07-01/30/08	Change in POP	
	2	01/16/08	01/31/08-01/30/09	Exercise Option for Phase III	370.0
	3	02/15/08		No change to RST	
	4	03/12/08	01/31/08-01/30/09	Fully Fund Phase III	1,499.3
	5	12/04/08	12/04/08-01/30/10	Exercise Option for Phase IV	1,467.7
	6	12/08/10		Close Out	-3.7
HSHQDC-10-X-00049		12/03/09	12/03/09-12/02/10	1. Project Management 2. System Testing, Characterization and Demo Support 3. Extend Allowable Operating Conditions 4. Automated Isotope ID 5. Fixed Portal Design Study 6. Improve TAT Performance 7. Improved Deployability 8. Improved System Robustness	1,450.0
	1	07/29/10	12/03/09-09/30/11	9. Custom Commercial Stereo Imager Development 10. Additional Demonstrations 11. Additional Source Tests 12. Fixed Portal Design Study 13. Dense Stereo Implementation 14. Pilot Study Support	220.6
	2	08/20/10	12/03/09-09/30/11	For Tasks 10-13	380
	3	09/01/10	12/03/09-09/30/11	For Task 9	45
	4	11/22/10	12/03/09-09/30/11	For Task 10	99.8
	5	08/16/11	12/03/09-09/30/12	For Task 11 --Expanded to Phase II in next Amendment	350
	6	02/29/12	02/29/12-09/30/12	Phase II Tasks 1-3 1. Project Management 2. Roadside Tracker Maintenance 3. Data Collection & Analysis Support	200
	7	08/17/12	12/03/09-11/30/12	Complete Phase II	50

6. CONCLUSION

Development of the RST has resulted in a new capability for detecting illicit nuclear materials. With such an instrument one can perform targeted, low-impact, scans of vehicular traffic at random locations providing a new deterrent to combat nuclear smuggling. There is the hope that shielding might be at a minimum for in-transit sources, further increasing the chances of detection. The system detects and scans over 99% of the passing traffic and can isolate detections to a single vehicle. Further, the imaging provides natural immunity to systematic effects such as baseline suppression, source confusion, and background variations while providing background-subtracted spectra for use with isotope ID algorithms. In addition to the roadside application, there are a number of other places where this approach could prove of value including, maritime against small vessels at natural navigational chokepoints, pedestrians, railways and classic fixed chokepoint traffic settings. Some tests have been conducted for the first two of these, but for any such application, additional work would be required to upgrade the TAT algorithms for tracking in these different situations. Results with both pedestrians and vessels indicate that the point-based tracking is robust and flexible, meaning the likelihood of successful adaptation to any of these applications is high.

The system has been developed to a technical readiness level (TRL) of 5-6, meaning we have a working fieldable prototype that can be used to demonstrate the capabilities and utility of the approach. With the extensive trials, we have worked out quite a few of the issues encountered during those trials and at this point the primary concerns are those associated with failure of the COTS components, particularly due to their age after years of testing. A further round of developments would be required to develop an instrument ready for use by non-specialist field personnel. Such development should include several aspects: First, a more ruggedized design, particularly including selection of rugged COTS components such as the computers, Second, software upgrades that primarily involve wrapping the proven software elements in overarching control software suitable for non-specialist use. This includes adapting the graphical user interface, as well as implementing various state-of-health monitors and watchdog timers to make sure that the system is truly functioning as expected. Third, work needs to be performed in alerting response personnel and identifying suspect vehicles to them. In addition, a proper CONOPS should be developed. Even with the few tests where we conducted live exercises with response personnel, it is clear that positive identification and rapid interdiction are important. Suspect vehicles can rapidly be lost in the flow of traffic and the radiation sensors in response vehicles have neither the sensitivity nor the specificity to pick the suspect vehicle from among even a few others. Finally, there are a number of easily implemented improvements to the current system such as adding the ability to handle stop-and-go traffic, as well as automating the gain stabilization, and isotope ID.

Finally, we wish to thank DNDO for their continued support throughout the project and particularly their forbearance at those times when the work encountered unexpected setbacks.

7. REFERENCES

1. P.E. Fehlau, C. Garcia, Jr., R.A. Payne, E.R. Shunk. "Vehicle monitors for domestic perimeter safeguards," LA-9633-MS UC-15, Los Alamos National Laboratory, 1983.
2. R.L. York, P.E. Fehlau, "1997 Update for Applications Guide to Vehicle SNM Monitors," LA-13247-MS, Los Alamos National Laboratory, 1997.
3. R.T. Kouzes, E.R. Siciliano, "The response of radiation portal monitors to medical radionuclides at border crossings," *Rad. Meas.*, Vol. 41, pp. 499-512, 2006.
4. C.A. Lo Presti, D.R. Weier, R.T. Kouzes, J.E. Schweppe, "Baseline suppression of vehicle portal monitor gamma count profiles: A characterization study," *Nucl. Inst. Meth. A.*, vol. 562, pp. 281–297, 2006.
5. K.P. Ziock, K.E. Nelson, "Maximum detector sizes required for orphan source detection," *Nucl. Inst. Meth. A.*, vol. 579, pp. 357-362, 2007.
6. K.P. Ziock, L. Fabris, D. Carr, J. Collins, M. Cunningham, F. Habte, T. Karnowski, W. Marchant, "A Fieldable-Prototype, Large-Area, Gamma-Ray Imager for Orphan Source Search", *IEEE Trans. Nuc. Sci.*, vol. 55, pp. 3643 – 3653, 2008.
7. E.E. Fenimore, T.M. Cannon, "Coded aperture imaging with uniformly redundant arrays," *Applied Optics*, vol. 17, no. 3, pp. 337–347, 1978.
8. E.E. Fenimore, "Coded aperture imaging: Predicted performance of uniformly redundant arrays," *Appl. Opt.*, vol. 17, no. 22, pp. 3562-3570.
9. K.P. Ziock, E.C. Bradley, A.M. Cheriyyadat, M. Cunningham, L. Fabris, C.L. Fitzgerald, J.S. Goddard, D.E. Hornback, R.A. Kerekes, T.P. Karnowski, W.T. Marchant, R.J. Newby, "Performance of the Roadside Tracker portal-less portal monitor," *IEEE Trans. Nuc. Sci.*, vol. 60, no. 2, pp. 2237-2246, 2013.
10. K.P. Ziock, J. Collins, M. Cunningham, L. Fabris, T. Gee, J. Goddard, F. Habte, T. Karnowski, "The Use of Gamma-Ray Imaging to Improve Portal Monitor Performance," *IEEE Trans. Nuc. Sci.*, vol. 2, pp. 1198 - 1208, 2007.
11. K.P. Ziock, A. Cheriyyadat, L. Fabris, J. Goddard, D. Hornback, T. Karnowski, R. Kerekes, J. Newby, "Autonomous radiation monitoring of small vessels," *Nucl. Inst. Meth. A.*, vol. 652, pp. 10-15, 2011.
12. T.P. Karnowski, M.F. Cunningham, J.S. Goddard, A.M. Cheriyyadat, D.E. Hornback, L. Fabris, R.A. Kerekes, K.P. Ziock, T.F. Gee, "Motion estimation accuracy for visible-light/gamma ray imaging fusion for portable portal monitoring," *Proc. SPIE*, vol. 7538, DOI 10.1117/12.838428, 2010.
13. T.P. Karnowski, M.F. Cunningham, J.S. Goddard, A.M. Cheriyyadat, D.E. Hornback, L. Fabris, R.A. Kerekes, K.P. Ziock, E.C. Bradley, J. Chesser, W. Marchant, "Desing of Dual-Road Transportable Portal Monitoring System for Visible Light and Gamma-Ray Imaging," *Proc. SPIE*, vol. 7665, DOI 10.1117/12.850191, 2010.
14. R.A. Kerekes, J.S. Goddard, A.M. Cheriyyadat, T.P. Karnowski, M.F. Cunningham, D.E. Hornback, L. Fabris, K.P. Ziock, "Accuracy of a stereo vision system for portable road side portal monitoring," in Proc. 2011 Future of Instrumentation International Workshop, Oak Ridge, TN, Nov. 2011.
15. K.P. Ziock, J.W. Collins, Fabris, S. Gallagher, B.K.P. Horn, R.C. Lanza, N.W. Madden, "Source-Search Sensitivity of a Large-Area, Coded-Aperture, Gamma-Ray Imager," *IEEE Trans. Nuc. Sci.*, vol. 53, pp. 1614-1621, 2006.
16. W.H. Press, S.A. Teukolsky, W. T. Vetterling, B.P. Flannery, Numerical Recipies in C, (Cambridge University Press, New York, New York, 1995) p. 293.

17. Denise Pelowitz (ed.). "MCNPX User's Manual, version 2.5.0," Los Alamos National Laboratory, LA-CP-05-0369, 2005.
18. S. Agostinelli *et al.*, "GEANT4—a simulation toolkit," *Nucl. Inst. Meth in Phys. Res. A*, vol 506, pp. 250-303, 2003.
19. S. Gottesman, E. Fenimore, "New family of binary arrays for coded aperture imaging," *App. Opt.*, vol. 28, No. 20, pp. 4344-4352, Oct. 1989.
20. PCI-6534, National Instruments Corporation, 11500 N Mopac Expwy, Austin, TX 78759, www.ni.com.
21. K.P. Ziock, M.T. Burks, W. Craig, L. Fabris, E.L. Hull, N.W. Madden, "Real time generation of images with pixel-by-pixel spectra for a coded aperture imager with high spectral resolution," *Nucl. Inst. Meth B*, vol. 505, pp. 420-424, 2003.
22. GEViCAM Inc., 673 S. Milpitas Blvd, Suite 100, Milpitas, CA 95035.
23. J. Shi, C. Tomasi, "Good features to track," *Proc. Computer Vision and Pattern Recognition '94 IEEE Comp. Soc. Conf.* pp., 593–600, 1994.
24. B.D. Lucas, T. Kanade, "An iterative image registration technique with an application to stereo vision," *Int. Joint Conf. on Artificial, Intelligence*, vol. 3, pp. 674-679, 1981.
25. F.M. Dickey, L.A. Romero, "Normalized Correlation for Pattern Recognition," *Optics Letters*, vol. 16, Issue 15, pp. 1186-1188, 1991.
26. A.M. Cheriyyadat, B.L. Bhaduri, R.J. Radke, "Detecting multiple moving objects in crowded environments with coherent motion regions", *IEEE Computer Vision and Pattern Recognition Workshops*, pp. 1-8, 2008.
27. Mathworks, 3 Apple Hill Dr., Natick, MA 01760.
28. R.D. Penny, W.E. Hood, R.M. Polichar, F.H. Cardone, L.G. Chavis, S.G. Grubs, B.P. Huntley, R.A. Kuharski, R.T. Shyffer, K.P. Ziock, L. Fabris, S.E. Labov, K. Nelson, "A Dual-Sided Coded-aperture Radiation Detection System," SORMA 2010, *Nucl. Inst. Meth. in Phys. Res. A.*, vol. 652, pp. 578-581, 2011.
29. Spectra Precision, ProMark 100, 10355 Westmoor Dr., Suite #100, Westminster, CO 80021, <http://www.spectraprecision.com/welcome-85.kjsp?RF=PRO-EN>

Appendix A: List of Acronyms

Acronym	Description
ADC	Analog to Digital Converter
BLR	Baseline Restorer
CMR	Coherent Motion Regions
CONOPS	Concept of Operations
COTS	Commercial off the Shelf
DAD	Gamma Data Collection Computer
DAC	Digital to Analog Converter
DAQ	Data Acquisition System
DHS	Department of Homeland Security
DNDO	Domestic Nuclear Detection Office
DOE	Department of Energy
DRF	Duel Road Fusion
FPGA	Field Programmable Gate Array
GEANT	Monte Carlo Radiation Transport Software
GPS	Global Positioning System
GRIP	Gamma-Ray Image Pixel
IAA	Inter-Agency Agreement
LAI	Large Area Imager
LUE	Limited Use Exercise
MCNP	Monte Carlo Radiation Transport Software
NORM	Normally Occurring Radioactive Materials
PMT	Photomultiplier Tube
POC	Proof-of-Concept
POF	Proof-of-Feasibility
POP	Period of Performance
REMLOC	Remote to Local
RNAK	Radio Nuclide Analysis Kit
RPM	Radiation Portal Monitors
RST	Roadside Tracker
SGR	Segmentation Results
SORDS	Stand-Off Radiation Detection System
TAT	Target Acquisition and Tracking
TRL	Technical Readiness Level

(09/07)

Gamma-Ray and Video Roadside Tracker

CFP 06 – TA-02-OR18

ORNL DNDO Program Manager: James D. White

2367-T118-06, 43WT11801 (expired 04/02/07)

2367-T303-07, 43WT30301 (expires 11/30/09)

Report on the Proof of Feasibility Experiments with a Prototype, Roadside Tracker, Gamma-Ray-Imaging Portal Monitor

K.P. Ziock, J. Collins, M. Cunningham, L. Fabris

T. Gee, J. Goddard, F. Habte, T. Karnowski

Prepared by: *Klaus Ziock, Principal Investigator*

ORNL

PO Box 2008, MS-6010

Oak Ridge, TN 37831

ziockk@ornl.gov

Abstract

We report on the results of the Proof of Feasibility experiments conducted with our novel “Roadside Tracker” imaging portal monitor approach. An experimental version of the Roadside Tracker was assembled by using the hardware from the “Large-Area Imager,” wide-area search instrument, combined with a visible-light video camera and new software. The Roadside Tracker concept uses automated target acquisition and tracking (TAT) software to identify and track vehicles in visible-light images. The field of view of the visible camera overlaps with and is calibrated to that of a one-dimensional gamma-ray imager. The TAT code passes information on when vehicles enter and exit the system field of view, and when they cross gamma-ray pixel boundaries. Based on this information, the gamma-ray imager “harvests” the gamma-ray data specific to each vehicle; integrating its radiation signature for the entire time that it is in the field of view. In this fashion we are able to generate vehicle-specific radiation signatures, and avoid source confusion problems that plague non-imaging portal monitors. The results obtained during the experiments conclusively demonstrate the value of this approach. We were able to automatically alarm on vehicles that carried a source, present the spectrum from that source, and even localize the source within the vehicle. If “innocent” (sourceless) vehicles were in the field of view simultaneously, no false alarm on these vehicles was obtained. The system was able to monitor two lanes of traffic simultaneously from one side of the roadway, thereby demonstrating that this approach can be implemented without affecting the flow of traffic.

Report on the Proof of Feasibility Experiments with a Prototype, Roadside Tracker, Gamma-Ray-Imaging Portal Monitor

K.P. Ziock, J. Collins, M. Cunningham, L. Fabris, T. Gee, J. Goddard, F. Habte, T. Karnowski

Introduction

With heightened concerns over the smuggling of fissile material, the use of radiation sensors as portal monitors is increasing. Of particular interest are larger systems that can be rapidly deployed to different locations and that can detect significant quantities of fissile material that might be present in vehicular traffic. An ideal system should have the following properties:

- A high probability of detection.
- A low false alarm rate.
- Minimized disruption of traffic.
- Threat linkage to an individual vehicle.
- Threat (isotope) identification.
- Immediate intuitive feedback.

The classic approach to any vehicular portal monitor is to use one or more detectors for each lane. The detectors integrate the radiation signature from a vehicle stopped, or moving slowly in front of it, with integration intervals based on occupancy sensors [1-5]. It is clear that implementing such a system is inherently difficult, particularly in a rapid deployment scenario. Further, setting up and running multiple sensors in the lanes of traffic will have a considerable impact on traffic flow. A more advanced approach using a single sensor at the side of a roadway is problematic due to source confusion issues (i.e., identifying which vehicle, in which lane, set off an alarm) unless traffic is extremely light. Even if multiple lanes are instrumented, the range of radiation from a strong source in a single “hot” vehicle may cause alarming in multiple lanes and for several vehicles in each lane. This means that traffic must be stopped and each vehicle inspected with secondary means, or that vehicles must be allowed to pass without further challenge. Neither is an acceptable solution.

Radiation Imaging Approach

It has been successfully shown that a vehicle-mounted, one-dimensional, gamma-ray imager can be used to improve on wide-area search applications [6]. Such an instrument keeps track of its location to map instantaneous, 1-D “snap-shot” gamma-ray images to world images that are corrected for different ranges from the imager. In this way a 2-D image of the world is created (see Fig. 1). Similar images could be built if the imager is stationary and the world moves past it. In fact, if a vehicle with a very strong source were to pass in front of an imager, short-exposure images would reveal the source’s location, allowing a “movie” of its passage to be generated (see Fig. 2). Each of the movie frames could then be shifted so that it is centered on the source location, and the data summed to integrate the gamma-ray signature from the vehicle’s entire transit in front of the imager. This “harvesting” of the gamma-ray data increases the system sensitivity to the source and is the concept for the Roadside Tracker System.

In reality, most vehicles are invisible to this approach because they are not significant sources of radiation. In the search application, this is not a problem because knowledge of an imager’s location at all times allows the imaged data to be properly mapped to the correct world pixels of the final image. To map the data from a non-




Fig. 1. Gamma-ray image made by the Large Area Imager tracking past a 1 mCi source at 20 mph at a range of 50 m. This represents a 6-sigma detection. Darker pixels indicate more radiation.




Fig. 2. One-dimensional “images” (histograms) of a 1-mCi source transiting the imager’s field of view. The source location is given by the peak in the histogram and progresses from right to left in the images from top to bottom.




Fig. 3. Final vehicle image created by harvesting (summing and shifting) the gamma-ray data from the same vehicle transit where a few snapshots are shown in Fig. 2.

data to the correct pixels of a vehicle-specific image that is accumulated for the entire time that a vehicle is in the instrument's field of view. After a vehicle has transited the system, this overall image is automatically checked to see if a source is present. If a detection is made, then the visible-light image is immediately available to operators (see Fig. 4) so they can identify the suspect vehicle.

Radioactive source detection is accomplished based on a *statistical* threshold of the counts in each pixel of the final gamma-ray image. As the data is collected, the number of counts used at each step of the image generation process is known. Based on work with the Large Area Imager, we have shown that straight-forward error propagation of such information provides a reasonable first-order estimate of the uncertainty in each pixel of the final image [7]. Further, we have shown that the counts obtained in “no-source” pixels of an image randomly fluctuate about zero solely based on these counting statistics errors, and therefore, the significance of a detection in a single pixel is just the number of counts in the pixel divided by the calculated variance. By setting an alarm when this quotient (or sigma value) exceeds 5-sigma, we obtain a high confidence of detection with a very low false-alarm probability.

Note that this approach is immune to the normal problems associated with systematic, vehicle-induced changes in the background count rate. In a classic portal monitor, the very presence of a vehicle changes the background count rate in an unknown fashion due to two competing effects. First, the vehicle can act as a shield between the detector and the rest of the world, thereby reducing the count rate. Second the vehicle can include materials and distributed cargo that are naturally radioactive, thereby increasing the count rate [1, 8]. These two competing effects mean that the alarming threshold has a large systematic component that makes determining the true sensitivity of such a system uncertain. The imaging process removes at least the first-order effects of such variations. Further, since threat sources are generally considered to be compact, by comparing the average radiation levels from the vehicle to any compact source seen in the image, one can overcome second-order effects as well. As with the large area imager, statistical outliers of the process must be determined by testing any final system on a large sample of vehicles, both with and without sources, and then performing an overall analysis using standard statistical techniques such as receiver-operator-characteristics (ROC) curves [9].

The roadside tracker uses radiation sensors (CsI(Na)) with medium energy resolution that provide the additional capability to check for alarms in different regions of the gamma-ray spectrum. By using pre-selected regions of interest or energy windows, where threat isotopes emit radiation, we can optimize sensitivity to such isotopes. If a detection is made, one has the further benefit that *background-subtracted* spectra from the “hot” pixels of the vehicle are available immediately [10]. Such spectra can be used to help determine the threat-level posed by any source that is detected.

Instrument Design

To obtain proof-of-feasibility (POF) data for the Roadside Tracker, we used the Large Area Imager (LAI) wide-area search instrument [11]. The LAI comprises two front-

radioactive vehicle to the correct pixel of its overall gamma-ray image requires knowing where the vehicle is at all times. With such information, a gamma-ray image could be “harvested” from the continuous data stream from a gamma-ray imager by adding the individual 1-D snap shots (see Fig 3).

For the Roadside Tracker, we obtain the requisite information on a vehicle's location by using automated target acquisition and tracking (TAT) routines to process visible-light (video) images of the traffic flow. The visible-light images are mapped to the gamma-ray imager pixels. This information is used by the TAT algorithms to determine each time a vehicle's front edge crosses the boundary of a gamma-ray pixel. This information is passed to the gamma-ray imager which uses it to map the gamma-ray



Fig. 4. Video image of a vehicle containing the 1-mCi source detected in Figs. 2 and 3.



Fig. 5. Picture of the Large Area Imager. It is mounted in a 16-foot trailer shown in Fig. 6

to-back, 1-D coded aperture imagers. Each imager has 27, $4.2 \times 4.2 \times 40$ cm CsI(Na) bars as a 1-D detector array behind a coded-aperture mask made of bars of linotype metal (a Pb alloy) (see Fig. 5.) The combined active detector area of 0.8 m^2 is sensitive to millicurie-class sources at distances of ~ 50 m passing by at speeds of 30 km/h [11]. For the POF experiments, a video camera was mounted on a 7.5-m high extendible boom fixed to the trailer and angled downward to image the roadway immediately in front of the gamma-ray imager (Fig. 6.) The shutter signal from the camera was used to provide time synchronization between the video and gamma-ray imagers.



Fig. 6. The Large Area Imager with the boom for deploying the video camera (inset).

Data Handling

The data handling is separate for the video and gamma-ray data and occurs in real time. Synchronization of the two is provided by the frame number, which is the count of the number of shutter signals since a data acquisition was started. It is independently maintained by both systems. The video system takes the incoming images and searches them for vehicles using an automated target acquisition and tracking (TAT) engine. The system generates a video event whenever a vehicle enters the system, exits the system or crosses a gamma-ray image pixel (GRIP) boundary. (Lane changes are an additional category of video event that is not yet implemented.) The video events are actually computed only after a vehicle has exited the system to allow the optimized performance available by using information from the entire transit of a vehicle past the instrument. The video analysis requires ~ 250 msec after the vehicle leaves the last GRIP. The list of video events is sent to the gamma-ray imager together with a picture showing the vehicle. The gamma-ray code uses the information to harvest the gamma-ray data unique to that vehicle. A one-dimensional radiation image of the vehicle emissions is the end result of this process. The statistical significance of the radiation levels seen in any one pixel is known from the imaging process



Fig. 7. Points detected on a passing vehicle.



Fig. 8. Points are grouped into vehicles. Here the results are shown as two vehicles pass each other in different lanes.



Fig. 9. Bounding box automatically drawn to locate a passing vehicle.

where the intensity gradient is simultaneously large in more than one direction. Identified points (see Fig. 7) are checked for motion from frame-to-frame and a pair-wise test of all points is used to find those points with similar motion. These are grouped into “blobs” based on their motion and proximity. The blobs are further merged into individual vehicles (see Fig. 8) using area overlap in a simple non-conflicting graph approach [14]. An overall bounding box for the vehicles is drawn based on the extrema of points associated with the blobs. The size of the bounding box is expected to vary as the vehicle enters and exits the field of view. Because the vehicle itself is fixed in size, a region of interest (ROI) is determined from the entire suite of bounding boxes available after the vehicle has fully transited the system (see Fig. 9.) Further, the ROI is assumed to move with constant velocity throughout the initial POF runs. (This approximation improves overall performance when the bounding box changes size and or fails to move due to segmentation errors.) Once a vehicle has fully left the gamma-ray imager field of view, the ROI is determined and the frame number where its leading edge crosses each GRIP boundary is calculated. This process results in a series of fractional frame numbers that are passed as video events to the gamma-ray code. In addition, a special “retire vehicle” event is also passed to the gamma-ray code. This event carries a video image of the vehicle together with its bounding box and a unique identification number.

Gamma-ray data handling

The gamma-ray data is collected using only a minor modification to the firmware of the Large Area Imager system [11]. In short, the data is collected asynchronously from individual electronics cards, one for each of the 54 CsI(Na) detectors. The cards provide full event handling, including a fast and slow channel for triggering and spectroscopy, respectively. After an event is processed, the card sets a flag to a digital board to indicate that digitized data is ready. The digital board collects the events in a time-ordered manner, groups them in millisecond increments together with header information, and sends them to the host computer via a commercial digital PCI interface card [15]. In addition to the gamma-ray data, the system also responds to pulse events. In the original LAI these were used to input data from a fifth-wheel velocity sensor. In the new system, the pulses come from the video camera and correspond to the shutter release. The system counts the number of these pulses and reports this number as another event class that is also entrained in the millisecond structure by the data handler.

At the data acquisition computer, the raw incoming data is first saved to disk, thereby allowing later replay as if the instrument were “live.” As shown in Fig. 10, the gamma-ray events are then saved into a data structure by frame number and millisecond within the frame number. Information includes the detector number and the energy of

the event. At a frame rate of 15/s, there are 68 millisecond time stamps available per frame. The gamma-ray data is retained for a total of 500 frames, comprising over half a minute of data.

The system periodically checks for video events which are sent by the video processing engine via an Ethernet link. The system currently accepts four classes of video event: New Vehicle, Cross Front Imager GRIP Boundary, Cross Back Imager GRIP Boundary, and Retire Vehicle. In addition to event type, each event carries information on the vehicle number (a unique number assigned by the TAT engine), the type of event, and the time of the event. The latter number is based on fractional frame time.

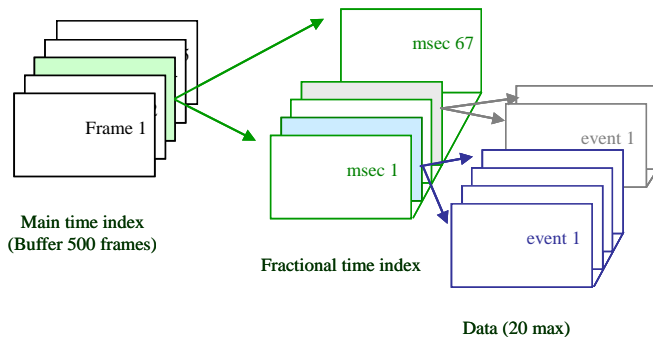


Fig. 10. Data structure for the gamma-ray events. The events are grouped by frame, each frame having 68 possible msec buffers. Each msec buffer can hold up to 20 events.




Fig. 11. Vehicle display window. One can choose from among the list of “retired” vehicles displayed in the center of the form, by selecting them with the mouse. The video image (top left) and the 1-d gamma image (bottom left) of the selected vehicle are shown together with the vehicle attributes in text form (on the right). The “NS” designation in the list indicates that a detection of greater than 5-sigma occurred for that vehicle. The cyan arrow in the 1-d gamma image shows the vehicle length and direction. In this vehicle a source is detected with ~ 12.6 sigma significance in the back seat of the crew-cab pickup truck.

If a New Vehicle event occurs, a new vehicle data structure is added to a list of currently active vehicles. For each video event indicating that a GRIP boundary has been crossed, the system finds the correct fractional frame number and then sums all of the data from the buffer structure between the time of the previous video event and the current event into a detector array. The data in the array is then used to generate a gamma-ray image, and the data from this image is then harvested based on the vehicle’s location. The harvested data is offset to the correct vehicle location and added to the global image for that vehicle. If a Retire Vehicle event is received, the system removes the vehicle data from the active vehicle list and saves it to a list of completed vehicles. At the same time a check is made to determine the pixel with the strongest detection, and this pixel is checked to see if an alarm has occurred.

The vehicles that have exited the system are shown in a separate window (see Fig. 11) that allows one to select any of the vehicles from a list-box. When selected, the data from that vehicle is displayed by the system. The display includes a histogram of the gamma-ray image, together with an arrow indicating vehicle length and direction of travel, and a video image of the vehicle. If “Full Image Processing” is enabled, then a spectrum of the gamma-ray data for the vehicle as a whole, a region of the vehicle, or for individual pixels of the vehicle can be shown (see Fig. 12). The vehicles displayed in the list can be filtered based on alarm conditions, lanes of travel, and frame number(s).

System Setup

The proof of feasibility tests were conducted in two stages. For the first stage, the system was setup in a parking area away from the actual roadway with a single lane of traffic defined using orange cones. For the second location, we moved to a gravel parking area just off of a two-lane roadway that had modest traffic loads and a maximum speed limit of 45 mph. For each deployment, the video image was calibrated to a number of reflective disks placed in the field of view. The distances in between the disks were measured using a tape measure. Cross-calibration between the gamma-ray and the video pixels was achieved by simultaneously imaging a gamma-ray source at three separate locations with the video and gamma-ray instruments. As the video camera was not weather hardened, we removed it from its “Hi-Pod” mount each evening. In principal, its view could be reestablished strictly by realigning on the reflective disks, whose location was marked on the ground. In practice, we recollected simultaneous video and gamma-ray images daily to minimize any residual uncertainties.




Fig. 12. Energy spectrum obtained from a source in a passing vehicle. The gamma-image is shown in the top panel. The spectrum from the pixel at the center of the peak (under the white-line cursor) is given in the middle panel. The bottom spectrum is from pixel 15 (indicated by the arrow).

Preliminary Tests

In the first location the functionality of the system was verified after the final system integration. The tests included checking that slow moving vehicles were tracked in both directions and that source data was correctly harvested based on the video images. Tests were conducted using the 662 keV line from a 100 μCi ^{137}Cs source (580-750 keV). The instrument was set up ~ 4 m from the “roadway” and vehicle velocities were typically of order 5 mph. The results were extremely encouraging as can be seen in Figs. 13 to 15. In Fig. 13 we show one of the first positive results with the system, where a 13-sigma detection is made for a vehicle moving from right to left in front of the system. In the same figure we show results from a no-source run later that day. In this case the maximum significance of any GRIP is only 2.4-sigma. The system display shows the vehicle length reported by the TAT code via the gold or blue arrow superimposed at the bottom of the gamma-ray image (histogram). The vehicle’s direction of motion is given by both the color of the arrow and the location of the arrow head.

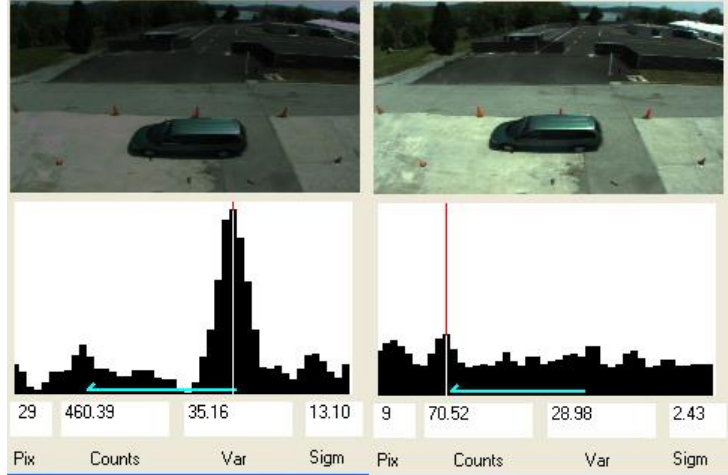


Fig. 13. First results with the system. The left panel shows the results for the vehicle moving at 3.4 mph with a 100 μCi source. A 13.1 sigma detection is made. In the panel on the right, the vehicle is moving at 2.2 mph with no source. The pixel with the maximum significance is only 2.4 sigma.

In Fig. 14 we demonstrate that the instrument correctly locates the source within a vehicle. As the source is progressively moved from the back to the front of the vehicle in successive runs past the instrument (right to left panels in the figure), the source location in the images shifts appropriately. Note that we could not fit the source (which had to be in a large box for radiation regulations) into the front seat of the vehicle so that in the rightmost panel, the vehicle was driven backwards to put the source at the front of the direction of motion.

We did notice that the system performed better on some runs than on others (see Fig. 15.) To isolate the cause of this to either the video or gamma-system, we generated the GRIP events by hand-segmenting the video images. The resulting gamma-ray image is shown on the right side of Fig. 16. The improvement in the height and width of the gamma-ray peak is obvious, and indicates that fluctuating GRIP crossing times provided by the automated TAT engine can degrade the system results. By plotting the frame number versus GRIP boundary crossing times generated automatically and by hand, (Fig. 16) it is obvious that the TAT engine had problems from frame 1540 to frame 1565 in this run. One of the immediate results of this test was to implement a uniform motion approximation in the TAT algorithms.

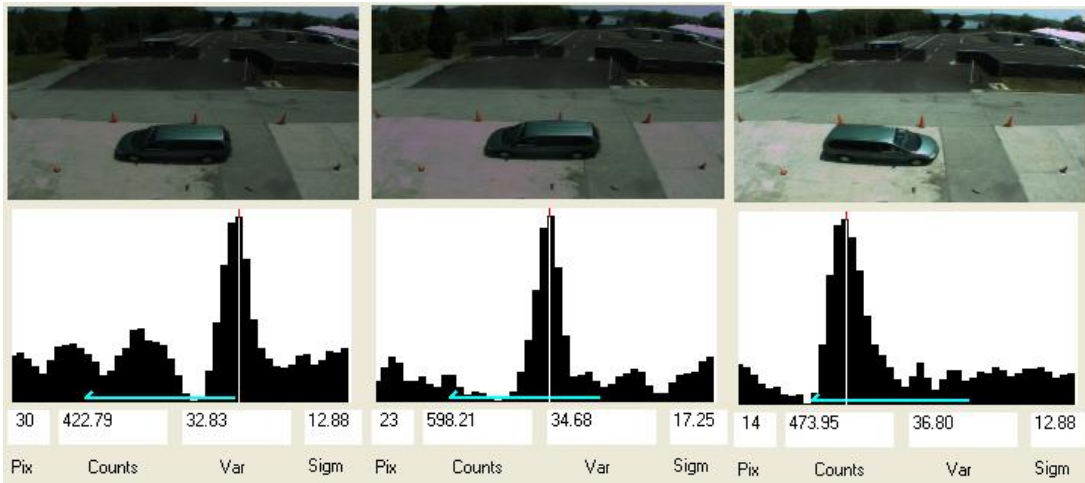


Fig. 14. Source location within the vehicle is also detected. The source is moved from the back (left) to the middle (center) to the front (right) of the vehicle. Note that we could not fit the source and its enclosure in the front seat so we put it in the back of the vehicle and drove backwards for the rightmost image.

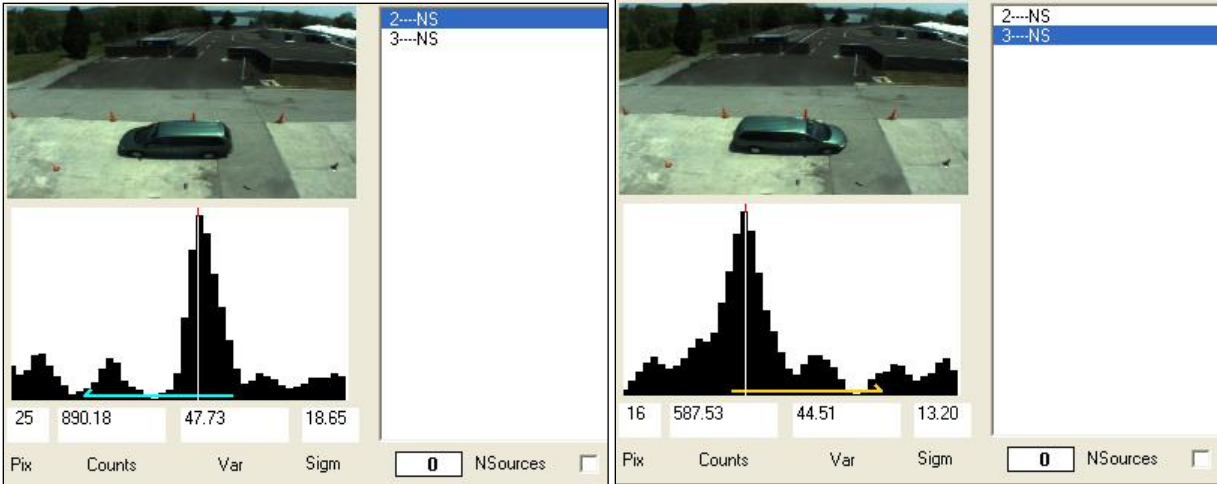


Fig. 15. Selecting different vehicles from the same run. The left image shows the display when the first vehicle is selected, the right image shows the second. The direction of travel is given by the arrow at the bottom of the gamma-ray image. Note also that the right gamma image is “broader” than the one on the left.

Proof of Feasibility Results

The results obtained at the second location constitute the real proof of feasibility tests. We collected several days’ worth of data, all of which was run real time to assess performance of the system. It was also saved to disk for later analysis and tests of improvements to the system. The data shown in the previous and following sections are all from replaying the data. Results that are described as real time or as on the day of data collection, use the original video event files saved to disk and are representative of results obtained the day of the experiment. However, we did correct a coding error in the gamma-ray code that resulted in over estimates of the significance of source detection because some data was added to the results multiple times. The bulk of the sample results shown below use later versions of the code that include improvements made to the TAT system since the data were collected.

Despite some obvious failures that are being addressed by the ongoing TAT development effort, the system performance was remarkable. The results at the first location demonstrated that the system properly identifies the vehicle with a source, as well as providing the source location along the length of the vehicle. Tests at the new loca-




Fig. 16. Plot of the GRIP vs frame number for the second vehicle of Fig. 15. The graphs show the time (frame number) when the front of the vehicle crosses a GRIP boundary. The original tracking (left) failed briefly around frame number 1550. This failure broadened the gamma-ray image (upper left.) Segmenting the data by hand (right) provides good tracking throughout the vehicle’s passage, giving a narrower and stronger gamma-ray detection (upper right).

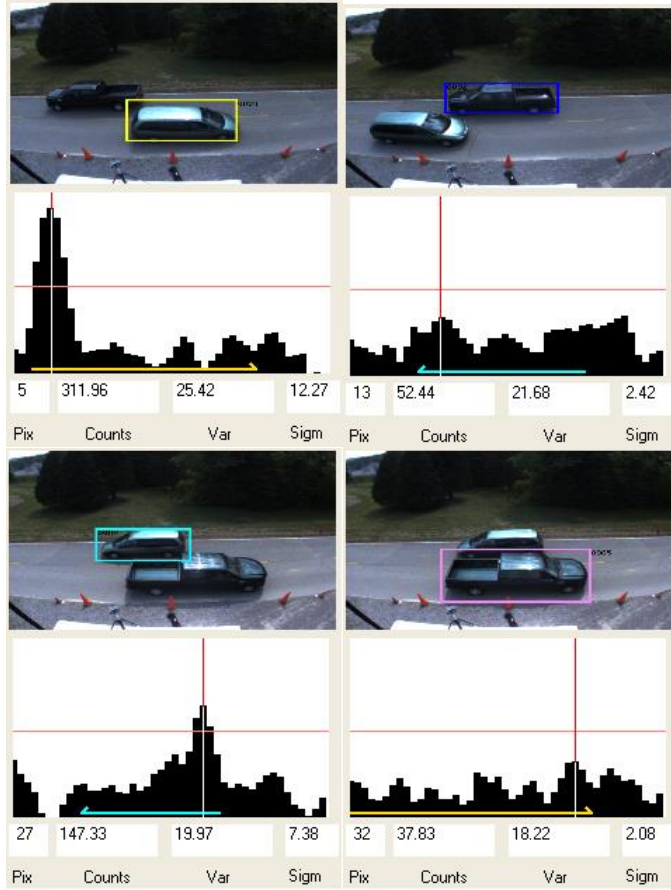


Fig. 17. Vehicles crossing. The Van carries a 100 μCi source and is detected if it is in the near (top) or far (bottom) lane. The pickup truck traveling in the opposite direction at the same time has no source and does not alarm. The horizontal orange line represents the 5-sigma detection threshold. The histogram below the video frames is of the sigma values of the gamma-ray pixel, rather than the standard count histogram. The non-imaged data for this sequence is shown in Fig. 18.




Fig. 19. Strong (in white car) and weak (in green van) sources in the field of view simultaneously. Both sources are detected and attributed to the correct vehicles. The non-imaged data does not indicate the presence of two sources. (The vertical scale in the gamma image is the same in both panels.)




Fig. 18. The non-imaged gamma-ray intensity data associated with Fig. 17. The imaged data is collected between the colored vertical lines with the color of the line matching the vehicle region of interest shown in Fig. 17. It would be impossible to link a given vehicle to the count data without the imaging information.

tion added the ability to have multiple vehicles in the field of view simultaneously, as well as to test the system at higher velocities. We also used a stronger, 1 mCi source in some of the runs. In addition, we ran tests with both sources in the field of view simultaneously. Samples of system behavior in these more demanding situations are described in the following sections, followed by a general discussion of the system performance overall.

Performance with multiple vehicles

The primary goal of the Proof of Feasibility experiments was to demonstrate that the combined gamma/video imaging approach can detect sources *and* that it can correctly attribute them to a single vehicle in realistic, complex, traffic situations. The results already presented on the preliminary tests conclusively show that both detection and attribution to a single vehicle are possible. The tests with multiple vehicles address the remaining issues, and also show that the instrument correctly identifies the lane a vehicle is in. There are two multi-vehicle traffic geometries that were investigated, one with several vehicles traveling in the same direction, the other with vehicles crossing in opposite directions. Both of these tests represent situations where a non-imaging system set at the side of a roadway would be unable to correctly associate a source(s) with a given vehicle(s).

In Fig. 17 we show sample results from the system with two vehicles traveling in opposite directions that cross the center of the field of view at nearly the same time. Non-imaged data from the same vehicle crossing are shown in Fig. 18. This vehicle configuration is perhaps the most difficult for a non-imaging system to

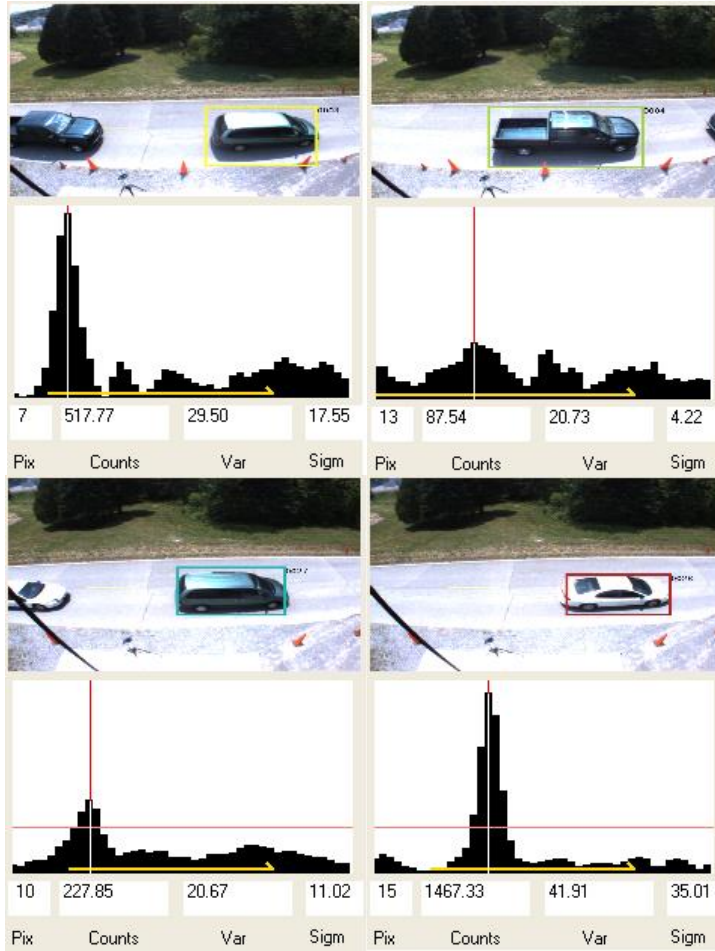


Fig. 20. A vehicle with no source (top right) is correctly passed with no alarm, even if an alarm triggers on the previous vehicle (top left.) Even in the challenging case where the 1 mCi source follows the 100 μ Ci source, the system still correctly sees both sources (bottom.) In the bottom figures the histograms show the statistical significance of the gamma-ray pixels, rather than the counts as shown in the top. This better displays the results for this high dynamic range example.

resolve. In contrast, the roadside tracker not only identifies which vehicle has the source, it also shows the relative location of the source within that vehicle. Of a total of 23 crossing configurations where the vehicle with the source(s) was reported by the latest TAT code, we correctly detect the source(s) 19 out of 23 times. The source was never attributed to the wrong vehicle. The few times that the source was missed were presumably due to shielding of the weaker source in the far lane by the vehicle in the near lane. In Fig. 19 we show the results when both vehicles have a source. In this case, both sources are seen and correctly attributed to the proper vehicle, as is seen by comparing the strengths of the relative detections.

While crossing sources are arguably more difficult for a non-imaging system to resolve, correctly identifying which vehicles within a stream of traffic moving in the same direction can also be very challenging. Sample results with the Roadside Tracker are shown in Fig. 20 (top). Again, the imaging approach correctly identifies which of two vehicles driving past the imager has the source, while this is almost impossible to determine from the non-imaged data. The system still correctly identifies what happens, even if the strong 1 mCi source follows the source that is 10 times weaker (Fig. 20, bottom). For completeness, we present the count-rate histogram for non-imaged data (Fig. 21) for the two-source case. There is no obvious indication that two sources are present in this data.

We worked to deliberately confuse the system by having the vehicles track through the field of view with progressively less separation while maintaining the same speed. It is not a surprise that at separations of order half a




Fig. 21. Non-imaged gamma-ray intensity data associated with the bottom panels of Fig. 20. The imaged data is collected between the colored vertical lines whose color matches the vehicle region of interest shown in Fig. 20. No time structure indicating two sources are present is visible in the plot.




Fig. 22. Currently, the TAT software will merge the two vehicles into one if they follow too closely and travel at the same speed. The source is still detected and correctly positioned with respect to the front of the combined vehicle.




Fig. 23. Peak detected counts versus one over the velocity (top) and sigma versus velocity (bottom). The upper curve is expected to be linear while the bottom should be a power law to the minus one half. The solid curves represent the best fit to a line and a power law, respectively.

results while we were collecting data. Despite our ability to replay the data (with the corrected code) this meant that the data set was not as complete as we might have wished. To span a wider velocity range, we turned to the longest run of the trials and looked at the results with the 1 mCi source in the near lane. This included passes in front of the system over the velocity range from 4.3 mph to 36 mph. The results are shown in Fig. 23. As can be seen, the performance follows the expected behavior. However, there is a large scatter in the data.

There are a number of factors that can contribute to the few points that fall far from the expected lines. These include poor TAT tracking, the presence of the second source, changes in the source location within the vehicle, the location of the source with respect to the gamma-ray pixel boundaries, and changes in the vehicle location within the lane. The first factor has already been discussed, while the second is obvious in that the second source will add to the correct pixel for the source we are looking at during some of its passage through the field of view. The presence of the second source also changes size of the background, which will affect the sigma of the detection. A change in the source location with respect to the vehicle can change the perceived source intensity if it moves behind a structural pillar, thereby partially shielding it from the detector. Since the instrument integrates the signal over a large range of angles during the transit, it is more likely to see a source than an instrument that looks only at a stationary vehicle, but this is a concern for any passive system. The location of the source with respect to the gamma-ray signal is a subtle effect that was already optimized in the Large Area Imager. It occurs because we look at the counts only in the peak pixel of the image. Suppose that the imager were designed so that a source in the center of a pixel falls entirely within that pixel. The counts in the peak pixel would then equal the detected flux—say N counts. However

vehicle length, the TAT system has difficulty separating the vehicles. In the right panel of Fig. 21 we show the results when vehicles moving at the same velocity follow too closely. Currently the system combines the two into a single long vehicle. If the source is in the front vehicle, it is detected. At the moment the gamma-ray code does not allow for vehicles longer than the instantaneous field of view so that sources in the second vehicle (or in long trucks) would be missed.

Performance at different velocities

As shown by the multiple vehicle results, the POF tests were completely successful, allowing us to correctly detect a radiation source, localize it to, and within, a single vehicle from a stream of multilane traffic. Beyond these results, we are interested in the performance of the system as a function of velocity as this is an important parameter to understand the ultimate capabilities of a road-side tracking instrument. Specifically, we are interested if the behavior follows expectations as the velocity increases. Since the integration time on a given vehicle decreases inversely as the velocity increases, we should collect less signal counts at higher velocities. However, as the integration times decrease, we also collect proportionally less background counts so we expect the significance of detection to decrease as the square-root of the signal counts, or that as the velocity increases, sigma decreases as the inverse of the velocity squared. To achieve this behavior requires that both the coded-aperture imaging behaves as expected, and that the video tracking behaves well throughout the velocity regime.

Recognizing the importance of verifying this behavior, we sought to collect data at a range of velocities once we were on a roadway with a reasonable speed limit. Unfortunately, we soon found that the video tracker would not capture vehicles at 25 mph and above during the tests. Although the source of this problem has since been found and corrected, it did mean that we could not validate the

if the source falls exactly between two pixels, then each of these will see half of the original counts ($N/2$) and the apparent detection is half as strong. The LAI design was chosen to minimize this effect by oversampling the image so that the source always shows up in several pixels, however some variation still occurs. In a deployed system, one would not look at the response of only the peak signal but at a range of pixels surrounding it. However, selecting an optimum pixel set is beyond the scope of the POF tests. The final cause for deviations from the expected source intensities has to do with the location of the vehicle within the lane. Depending on whether the vehicle is near the center of the road, the center of the lane, or near the edge of the road, the distance from the imager to the source can vary by a meter. This introduces a 20% uncertainty in the distance to the vehicle. Since we integrate over the track of the vehicle, the change in measured flux falls off as the distance, meaning a 20% variation in the counts can occur between two vehicles from this source alone.

Continuing Development

The performance of the TAT code continues to improve. Since the data were originally collected we have made a number of improvements that include revision of the point detection-threshold, Kalman filtering to allow non-constant velocities, improved region of interest definitions, etc. A full description of this work is beyond the scope of this document. However, to demonstrate the improvements achieved, we present the performance of the system as observed during the experimental campaign and using the latest TAT system at the time this document was written in Tables 1 and 2. The first table gives the overall performance of the TAT system, and the improvement in vehicle tracking is obvious from the increase in the number of vehicles successfully retired. These represent vehicles for which data is passed to the gamma-ray code for harvest of the radiation signature. In May we detected only 58% (144) of the vehicles with 4 of these reported too late for the gamma-ray data buffers. In addition, 5 vehicles were missed because they were merged with the previous vehicle. With the current version of the TAT software, when the same data is replayed, we detect 91% (224) of the vehicles with the latency problems completely fixed. In addition, the separation between closely following vehicles has been improved with a total of only 7 merged vehicles. In fact, the performance has improved in all aspects except that there is now a greater tendency for vehicles to be reported more than once. The overall improvement is not just in terms of the vehicles detected, but carries through to the gamma-ray performance as shown by Table 2. The first improvement is that the regions of interest are better defined, so that less gamma-ray sources are missed because the peak of the signal is outside the vehicle. Further, improved tracking means that the gamma-ray images are sharper, so that the significance of detection is higher, and fewer sources are missed.

Table 1: TAT System Performance

	5/23/07 and 5/24/07 Results	Late August Replay Results
# vehicle passes	247	247
Vehicles successfully retired	144	225
Vehicles missed	98	14
Vehicles merged with previous	5	8
Seen Multiple times	5	24
Gamma Latency Problems	4	0
Good Vehicles for Gamma test	140	225

Table 2: Gamma Results with TAT vehicle detection

	5/23/07 and 5/24/07 Results			Late August Replay Results		
	1 mCi	100 μ Ci	No Source	1 mCi	100 μ Ci	No Source
# Passes with good vehicle	28	67	45	48	93	84
# Detected (Alarmed)	27	39	NA	48	81	NA
# False Alarms	NA	NA	0	NA	NA	0
# Significant but outside vehicle	1	13	NA	0	0	NA
Gamma Missed	0	15	NA	0	12	NA

Summary

The results presented above clearly indicate that the roadside-tracking approach is a viable means to link vehicles with their radiation signatures without using individual detectors for each lane. Sources were correctly attributed to the proper vehicle, even in complicated situations where two sources were present in the field of view simultaneously. This successful result was the case even for the most difficult situation where the source types were identical, i.e. the same energy cuts were used on both vehicles. Of equal importance is the fact that, even when the strong one-milliCurie source was present, no false alarms were generated for other sourceless vehicles simultaneously in the field of view. This performance cannot be matched by non-imaging systems, even if lanes are individually instrumented. The system also has the capability to provide the energy spectrum from a source in real-time, meaning that improved differentiation between threat and non-threat sources can be used to decide if vehicles require further attention.

Beyond testing the basic concept, the Proof of Feasibility experiments provided a number of other benefits. A large set of data was collected which has been used to further refine the TAT algorithms. System performance has been improved, both by replaying the PoF data and comparing the gamma-ray results as tracking parameters were changed. In addition, some of the runs were segmented by hand and the system was compared to “ground-truth” data. We were also able to verify that the source intensity and significance of detection scale linearly and as the square-root of the velocity, respectively. In combination, these results provide a valuable basis for the design of the proof of concept prototype.

References

1. P.E. Fehlau, C. Garcia, Jr., R.A. Payne, E.R. Shunk. Vehicle monitors for domestic perimeter safeguards, LA-9633-MS UC-15, Los Alamos National Laboratory, 1983.
2. R.L. York, P.E. Fehlau, “1997 Update for Applications Guide to Vehicle SNM Monitors,” Los Alamos National Laboratory report LA-13247-MS, 1997.
3. “AT-900 Series Radiation Portal Monitor (RPM),” SAIC Exploranium, Mississauga, ON, Canada L5T 2V7, www.saic.com/products/security/at-900s/
4. “RadSentry Security Portals for SNM and Other Radionuclides,” Canberra Corp, 800 REsearch Parway, Meriden, CT 06450, www.canberra.com/products/1211.asp
5. “Guardian Continuous Radiation Monitoring Systems (Radiation Portal Monitors),” NucSafe Corp., 765 Emory Valley Road, Oak Ridge, TN 37830, www.nucsafe.com/Products/guardian_crms.htm
6. Large Area Imaging Detector for Long-Range, Passive Detection of Fissile Material, K.P. Ziock, W.C. Craig, L. Fabris, R.C. Lanza, S. Gallagher, B.K.P. Horn, N.W. Madden, *IEEE Trans. Nuclear Science* **51**, 2238-2244, 2004.
7. Source-Search Sensitivity of a Large-Area, Coded-Aperture, Gamma-Ray Imager, K. P. Ziock, J. W. Collins, Fabris, S. Gallagher, B. K. P. Horn, R. C. Lanza, N. W. Madden *IEEE Trans. Nucl. Sci.* **53**, 1614-1621, 2006.
8. “Baseline suppression of vehicle portal monitor gamma count profiles: A characterization study,” C.A. Lo Presti, D.R. Weier, R.T. Kouzes, J.E. Schweppe, *Nucl. Inst. Meth.*, **A562**, 281-297, 2006.
9. “Comparing Imaging and Non-Imaging Techniques for Reducing Background Clutter and Resolving Distant Point Sources,” R. Wurtz, K.P. Ziock, L. Fabris, R. Graham, *Proceedings IEEE Nuclear Science Symposium and Medical Imaging Conference*, Las Croabas, Puerto Rico, 10/23-10/29, 2006.
10. “Real time generation of images with pixel-by-pixel spectra for a coded aperture imager with high spectral resolution,” K.P. Ziock, M.T. Burks, W. Craig, L. Fabris, E.L. Hull, N.W. Madden, *Nucl. Inst. Meth.*, **B505**, 420-424, 2003.
11. “A Fieldable-Prototype Large-Area Gamma-ray Imager for Orphan Source Search,” K. P. Ziock, J. Collins, M. Cunningham, L. Fabris, F. Habte, T. Karnowski, *Proceedings IEEE Nuclear Science Symposium and Medical Imaging Conference*, Honolulu, Hawaii, 10/27-11/3, 2007.
12. “Vehicle segmentation and tracking from a low-angle off-axis camera,” N.K. Kanhere; S.J. Pundlik; S.T. Birchfield, *Proc. IEEE Comp. Soc. Conf. on Computer Vision and Pattern Recognition*, 1152-1157, June 2005.
13. “A Real-time Computer Vision System for Measuring Traffic Parameters,” D. Beymer, P. McLauchlan, B. Coifman, J. Malik, *Proc. IEEE Comp. Soc. Conf. on Computer Vision and Pattern Recognition*, 495-501, June, 1997.
14. “Detection and classification of vehicles,” S. Gupte, O. Masoud, R.F.K. Martin, N.P. Papanikolopoulos, *IEEE Trans. Intel. Transp. Syst.*, **3**, 37-47, 2002.
15. National Instruments PCI-6534, www.ni.com