

Final Report: Non-Visible, Automated Target Acquisition and Tracking



Klaus Ziock
Lorenzo Fabris
James Goddard
Donald Hornback
Tom Karnowski
Jason Newby

October 2016

**Approved for public release.
Distribution is unlimited.**

DOCUMENT AVAILABILITY

Reports produced after January 1, 1996, are generally available free via US Department of Energy (DOE) SciTech Connect.

Website <http://www.osti.gov/scitech/>

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
Website <http://www.ntis.gov/help/ordermethods.aspx>

Reports are available to DOE employees, DOE contractors, Energy Technology Data Exchange representatives, and International Nuclear Information System representatives from the following source:

Office of Scientific and Technical Information
PO Box 62
Oak Ridge, TN 37831
Telephone 865-576-8401
Fax 865-576-5728
E-mail reports@osti.gov
Website <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 & Isotope Technology

**FINAL REPORT: NON-VISIBLE, AUTOMATED TARGET ACQUISITION
AND TRACKING**

Klaus Ziock
Lorenzo Fabris
James Goddard
Donald Hornback
Tom Karnowski
Jason Newby

Date Published: October 2016

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

CONTENTS

LIST OF FIGURES	v
LIST OF TABLES	v
LIST OF ACRONYMS	vi
1. INTRODUCTION	1
2. PROJECT OVERVIEW	2
2.1 PHASE 1	2
2.2 PHASE 2	3
2.3 PHASE 3	4
2.4 PHASE 4	4
3. SYSTEM DEVELOPMENT	5
3.1 COMPUTER UPGRADES	5
3.1.1 Hardware	6
3.1.2 Software	6
3.2 LWIR SYSTEM	12
3.2.1 Hardware	12
4. SYSTEM PERFORMANCE	15
4.1 ORNL TESTS	15
4.1.1 Vehicle detection rates	15
4.1.2 Tracking performance	16
4.1.3 Source tests	17
4.1.4 Stop-and-go Traffic	19
4.2 FREEWAY PERFORMANCE	19
4.3 SMALL VESSEL TRACKING	21
4.4 DEMONSTRATION EXERCISE, PORT OF LOS ANGELES	23
5. CONCLUSIONS	25
6. REFERENCES	27

LIST OF FIGURES

Fig. 1. The original RST deployed on the I-495 Capital Beltway.	1
Fig. 2. RST output.....	2
Fig. 3. Video processing steps.	7
Fig. 4. Gamma-ray image processing steps.	7
Fig. 5. Auto exposure developed from 16-bit images.	9
Fig. 6. The RST deployed at an intersection during testing at the Port of LA.	9
Fig. 7. Street-level camera deployed near the Port of LA.....	10
Fig. 8. AI calibration target.....	11
Fig. 9. LWIR images under varying conditions.....	13
Fig. 10. Old (visible-light) and new (LWIR) stereo vision rigs.....	13
Fig. 11. Trailer roof as seen by its alignment camera.	14
Fig. 12. RST visible-light tracking errors.	17
Fig. 13. Mismatch in local and remote gamma results.	18
Fig. 14. Graph of source detection significance.....	19
Fig. 15. The RST deployed on the eastbound lanes of Interstate 40.....	20
Fig. 16. The RST on either side of the Boynton Beach Channel.	21
Fig. 17. Maritime detection of a weak check source.....	22
Fig. 18. The RST on a 4-lane causeway leading to piers for container ships.	23
Fig. 19. The RST on a four-lane causeway leading to piers for container ships.....	24
Fig. 20. LWIR vs visible-light vessel images.	25

LIST OF TABLES

Table 1. BVR Vehicle Detection Rates	15
Table 2. BVR Winter Vehicle Detection Rates	16
Table 3. Prior RST Visible-Light Vehicle Detection Rates.....	16
Table 4. BVR Test Site Tracking Errors.....	16
Table 5. BVR Source Detection Results.....	18
Table 6. BVR, Stop-and-Go Vehicle Detection Rates.....	19
Table 7. Freeway Vehicle Detection Rates	20
Table 8. Freeway Tracking Errors	20
Table 9. Small-Vessel Detection Rate	22
Table 10. Small-Vessel Tracking Errors.....	22
Table 11. Small-Vessel Detection Rates With Different CMR Parameters.....	23
Table 12. Port of LA Vehicle Detection Rates	24
Table 13. Port of LA Vehicle Detection vs. CMR Parameters	25

LIST OF ACRONYMS

BVR	Bethel Valley Road
CMR	coherent motion region
DHS	Department of Homeland Security
DNDO	Domestic Nuclear Detection Office
GPS	Global positioning system
LED	light-emitting diode
LPR	license plate reader
LTAT	local trailer's target acquisition and tracking
LWIR	long wave infrared
ORNL	Oak Ridge Nation Laboratory
OS	Operating System
RST	Roadside Tracker
RTAT	remote trailer's target acquisition and tracking
RTK	real time kinematic
SRNL	Savannah River National Laboratory
SSD	solid state drive
SWIR	short wave infrared
TAT	target acquisition and tracking

1. INTRODUCTION

The Roadside Tracker (RST) represents a new approach to radiation portal monitors. It uses a combination of gamma-ray and visible-light imaging to localize gamma-ray radiation sources to individual vehicles in free-flowing, multi-lane traffic. Deployed as two trailers that are parked on either side of the roadway (Fig. 1); the RST scans passing traffic with two large gamma-ray imagers, one mounted in each trailer. The system compensates for vehicle motion through the imager's fields of view by using automated target acquisition and tracking (TAT) software applied to a stream of video images. Once a vehicle has left the field of view, the radiation image of that vehicle is analyzed for the presence of a source, and if one is found, an alarm is sounded. The gamma-ray image is presented to the operator together with the video image of the traffic stream when the vehicle was approximately closest to the system (Fig. 2). The offending vehicle is identified with a bounding box to distinguish it from other vehicles that might be present at the same time. The system was developed under a previous grant from the Department of Homeland Security's (DHS's) Domestic Nuclear Detection Office (DNDO).



Fig. 1. The original RST as deployed during final testing on the I-495 Capital Beltway around Washington, DC. The system is in the two white trailers parked on either side off the roadway.

This report documents work performed with follow-on funding from DNDO to further advance the development of the RST. Specifically, the primary thrust was to extend the performance envelope of the system by replacing the visible-light video cameras used by the TAT software with sensors that would allow operation at night and during inclement weather. In particular, it was desired to allow operation after dark without requiring external lighting. As part of this work, the system software was also upgraded to allow the use of 64-bit computers, the current generation operating system (OS), software development environment (Windows 7 vs. Windows XP, and current Visual Studio.Net), and improved software version controls (GIT vs. Source Safe.) With the upgraded performance allowed by new computers, and the additional memory available in a 64-bit OS, the system was able to handle greater traffic densities, and this also allowed addition of the ability to handle stop-and-go traffic.




Fig. 2. RST output showing a video image (top) of the vehicle being scanned (indicated by the white rectangle). The 1D gamma-ray image (bottom) is given below the visible-light image with the vehicle's length and direction of travel given by the arrow at the bottom of image. This vehicle shows a strong detection marked by the vertical cursor line. The green horizontal line represents the baseline of the gamma-ray image and the red horizontal line the 5-sigma threshold. In this case the vehicle has a 24-sigma detection. The spectrum from the peak pixel is shown on the right and indicates this is a ^{137}Cs source.

2. PROJECT OVERVIEW

The project was planned using a phased approach based on DNDO's preferred four-phase development cycle: Proof of Feasibility, Component Development, System Development, and Testing and Evaluation.

2.1 PHASE 1

During Phase 1, we investigated in detail the two baseline non-visible imaging modalities forwarded in the proposal as the modalities of choice; one based on short wave infrared (SWIR) video cameras and the other on low-power, broadband, high-definition radar. The results of Phase 1 have been reported in the End of Phase 1 Project Report submitted to DNDO in September of 2012 and will not be repeated in detail here. In short, the study found little advantage to SWIR over visible-light imagery. There is not sufficient ambient SWIR radiation to allow imaging at night without adding SWIR artificial illumination to the operational area. While this could be done without the distractions and high profile that would have accompanied the use of visible lights; it was considered more problematic for the greater distances that might be encountered in possible long-range applications such as maritime settings. Of more importance, it was found that lighting on the vehicles (headlights and taillights) emitted significant amounts of SWIR radiation, and this caused dynamic range and fogging issues in the images. Further, the lights on the vehicles dominated the images, meaning considerable algorithmic development would have been required for use of this modality. Finally, one of the ancillary goals was to improve visibility in fog, and SWIR showed no significant benefits under such conditions.

Based on these findings, and a further study of cameras that work in other portions of the electromagnetic spectrum, it was decided that long wave infrared (LWIR) was a more promising candidate. In particular, LWIR images are made based on the thermal emissions from objects, meaning external lighting is not required. With a significantly longer wavelength than visible light, LWIR was also expected to show

better performance in fog. Finally, the LWIR cameras were found to be lower in cost than the SWIR cameras originally proposed.

Like the initial choice of SWIR, we also encountered issues with the use of Broadband Radar. However for Radar, these issues turned out to be primarily administrative rather than technical. Specifically, the radar wavelengths used by the commercial, low-cost, radar systems sold for use on small vessels that we originally planned to use were not available for use by governmental organizations. In the end, we were unable to obtain a license to test the systems. We did explore specialty systems based on similar technology used for monitoring traffic at intersections (including for control of red-light cameras). However, initial, fruitful discussions with vendors ultimately failed as our use was viewed outside their business model. In the end, with the improved performance of the LWIR over SWIR and the high costs of further investigation of radar, the system design was changed to include only LWIR cameras as the input to the TAT system.

2.2 PHASE 2

With the findings from Phase 1 that called for a completely different non-visible sensor technology, the work planned for Phase 2 was revised to explore the issues associated with the LWIR cameras. This included a fundamental check on the suitability of the technology reminiscent of work planned for SWIR under Phase 1. In particular, images needed to be collected and the ability to track vehicles tested. This included evaluating performance under different environmental conditions and ambient lighting. In addition, LWIR cameras were not available with the same pixel count as the visible and SWIR cameras, and the impact on the ability to track vehicles due to this, and the inherently lower sharpness of LWIR images, had to be tested. In addition, the arrays of small bolometers used as sensors in the LWIR cameras have a time constant that can be as high as 10 ms, which means some motion blur is present at high vehicle velocities, and the impact of this on tracking needed to be tested. Finally, the SWIR cameras originally selected were expected to perform as a drop-in replacement for the visible-light cameras, with virtually identical infrastructure concerns such as external triggering and image timestamps. This was not the case for any available LWIR cameras, and plans were to develop a translation card that took the LWIR CameraLink output, attached a timestamp to each image in real time, and transmitted the images to the host computer via Ethernet.

One of the unique aspects of the RST visible-light TAT software engine is the use of stereo information to help cluster points found in images into coherent groups that represent vehicles. In particular, the stereo information helps remove uncertainties associated with projection effects that can link points on tall vehicles in a nearer lane with short vehicles in a more distant lane. To continue to utilize the same algorithms meant that a stereo LWIR system was needed, and this is an area that has little reported research in the literature. While conceptually simple, the extension of stereo to LWIR images could be complicated by the inherently softer (less sharp) nature of the LWIR images, and so we planned to build and test an LWIR stereo rig that would eventually be a “drop-in” replacement for one of the current RST visible-light stereo rigs.

While exploring the overall system design in Phase 1, it became clear that the age of the extant RST computers was becoming a reliability issue, and it was decided that the system should be ported to new computers while upgrading the OS to the (at the time) latest version of Windows (Windows 7). With this change, it was decided that the project would benefit if the rest of the development tools were also updated. Finally, with the upgrade, we opted to convert the system to 64-bit (rather than 32-bit) operation. This was expected to ease performance issues uncovered during tests at the end of the original RST program. These were linked to memory issues under conditions of heavy traffic. Finally, with the improved performance during heavy traffic likely for the new computers and software, plans were made to add the ability to handle stop-and-go traffic to the system. Earlier development on this effort had

improved the data handling to deal with stop-and-go, but we were very optimistic about this approach given the improved hardware throughput.

2.3 PHASE 3

Phase 3 was used to integrate the new components (computer software and hardware, and LWIR system) into a single RST trailer and make the full system operational. This included both the upgraded computers with their new software, as well as the LWIR stereo rig. Part of the earlier work had been to determine how best to align the new stereo rig with the trailer. In the past a third optical camera had been used to view and track targets on the top of the trailer. After some experiments and analysis, we opted to continue to use the visible-light camera, and installed lights on top of the trailer that could be seen by the camera if the system was operating after dark. Since these were directed up toward the mast-mounted cameras, they neither attracted significant attention from passing vehicles nor created a driving hazard for those vehicles.

Once integrated, the system was used on the Oak Ridge National Laboratory (ORNL) campus to wring out remaining bugs in the new code and to adjust tracking parameters in the software. One of the issues uncovered during this phase of the project was that the LWIR cameras performed a flat-field calibration at periodic intervals that took them off line for a few seconds. While this could be switched off, it raised concerns that the image quality would suffer. Extensive tests indicated that this was not an issue if a flat field was performed before the start of a run and then those values were used for the run. However, it was not discovered until later that the flat-field data only applied for the type of triggering used to collect it. Originally, we collected the flat-field data with the cameras on internal trigger, while vehicle data were collected with an external trigger. This meant that significant amounts of data were collected without a viable flat field and the system still performed reasonably well. Toward the end of the phase, it was decided to build a second LWIR stereo rig so that the full two-trailer RST could be used during the Test and Evaluation Phase.

2.4 PHASE 4

Phase 4 of the project was organized into a series of measurement campaigns to test the system in different operational settings. In the prior project that developed the RST, we had already demonstrated that if the TAT tracking worked well, then the radiation detections obtained were at levels predicted by simple counting statistics. That meant that we did not need to generate radiation detections to determine the system performance, but could rather determine how well the system tracked vehicles and compare the performance to that of the visible-light system. If the comparison was favorable, then the gamma-ray detection would be good. Overall, the system tests naturally broke into two components: what was the vehicle detection rate, and if detected, how faithful was the reported track to the actual vehicle motion? For the original RST, the vehicle detection rates were slightly better than 99%. To determine the latter for the new cameras, we planned to use real time kinematic (RTK) enabled global positioning system (GPS) receivers in vehicles that passed through the system. These provided the locations of the vehicles to the centimeter level while they traversed the system, and should have allowed us to determine the tracking performance without turning on the radiation sensors. Unfortunately, the majority of the data files collected by the roving receiver were too long for the vendor's software to parse correctly. (While we were aware of this problem and stopped acquisitions after reasonable intervals, the software still only created single large data files.) As a proxy, we relied on the findings from previous work that the GPS tracks are generally linear, barring lane changes, which are extremely rare. With this model, we were able to compute the track error as the deviation from a linear fit.

With the emphasis of the testing on tracking, we focused on running the system under a wide range of environmental conditions, including around the clock, different weather, and at temperature extremes. All

of these were targeted at the goal of determining how well the new camera system would track vehicles across a wide range of operating conditions. To this end, the around-the-clock work emphasized different parts of the diurnal thermal cycle: rapid changes in temperature (occurring after sunrise and sunset), and times with minimal temperature changes (late afternoon, and predawn.) The former could present challenges of rapid contrast changes, while the latter could represent times when thermal uniformity washed out images. Neither of these turned out to be an issue, but we found that the automatic gain system in the cameras sometimes produced images with poor contrast—although the tracking still worked. To improve on this, we developed our own exposure software that took advantage of the 16-bit images available from the LWIR cameras. Our automated contrast system would periodically evaluate the average pixel value obtained from the roadway when no vehicles were present and then set the gain with user-determined values about this point. By selecting this limited range of values to generate the 8-bit images used by the processing stream, we were able to improve the visual contrast of the images. (It did not affect tracking as much as the operator's perception of the quality of an image.) With the timing of Phase 4 (Spring through Summer), we did not have the opportunity to collect cold-weather data, but results obtained late in Phase 3 showed that contrast of images taken on an extremely cold morning (single digits, °F) showed no issues with tracking.

One of the issues already known from the RST tests was the problem of vehicle identification to first responders whose mission was to interdict vehicles that set off the radiation alarm. This was already challenging due to the overhead nature of the visible-light images, and became worse with the LWIR images that provided an apparent grayscale image of the vehicle. While the overall shape of the vehicle could be used to help identify it, color was clearly missing, and the grayscale features did not relate to features visible to the naked eye. To help with this, for the final exercise, we added a visible-light camera at street level. An image from this camera was made available to the operators when an alarm occurred, and this did facilitate vehicle identification.

In addition to tests on the ORNL campus with radiation sources, three off-campus measurement campaigns were run. These included a freeway component on Interstate 40 east of Knoxville, a test with small vessels in Boynton Beach, Florida, and a general test near the port of Los Angeles, California. More details and the results of the testing are described in detail in the System Performance section of this report.

3. SYSTEM DEVELOPMENT

During the course of the project, a number of changes were made to the RST. In the following sections, we describe the final implementations and reasoning behind making them. These do not necessarily follow the chronological order of the project. Details about this can be found in the monthly reports submitted to DNDO throughout the system development cycle. The discussions below assume familiarity with the RST and its various components. Details can be found in the final report submitted to DNDO at the close of that project and in the published literature [1–3].

3.1 COMPUTER UPGRADES

The original computers in the RST were Dell 3400 machines running a 32-bit version of Windows XP. During the course of the current project, support for that OS was being phased out, and it was decided to transition the entire system over to a 64-bit version of Windows 7. At the same time, the machines that were most heavily used were replaced. There had been a number of failures during the late stages of the RST project, and this was partially attributed to the age of the machines and was most problematic for those machines engaged in the heaviest processing. The situation was worse for those machines since there was less computing margin before problems with the machine impacted the performance of the overall system.

3.1.1 Hardware

The computers with the heaviest processing loads (the Master and the TAT machines) were replaced with Dell T7600 machines. These each have 8-core processors with hyper-threading, allowing the software to take better advantage of the multithreaded architecture of the TAT and Master software. One of the biggest issues in the RST system was the write times needed to save the camera images to disk and the time needed to back up the data. This was predicated on the desire to save all of the images from each of the three cameras per trailer to allow software development on the stored data. Although this would not be a problem for a final system where much less data could be saved, it was a real issue that impacted the performance of the prototype. In particular, we found that each time we purchased new hard drives to collect data; we had to worry whether or not the drives could keep up. Even those disks that worked well when new would frequently slow to the point they could not keep up once more than about half of the disk space was utilized. To solve this issue, we changed over to solid state drives (SSDs) on the TAT machines. These had better performance when new, and they did not slow appreciably as they reached their full capacities.

One advantage of the old magnetic drives was that their low cost made them suitable as long-term storage media. Since SSDs are considerably more expensive, we opted to use the SSDs only for data collection and eventually backed the data onto magnetic drives. To ensure data redundancy, and to ease the time required to back up the data during a data collection campaign, we configured each of the TAT machines with a mirrored redundant array of independent disks so that we simultaneously wrote two copies of all of the data. Once a disk was full, the data were copied and the SSDs reused.

3.1.2 Software

3.1.2.1 System Synopsis

Although the TAT software is described elsewhere [4–6], a brief synopsis is useful to understand both the software changes and benefits derived from them. The system collects video images based on an external trigger that is used to synchronize the various video images and the gamma-ray data. The trigger is itself provided by a GPS receiver in each trailer (to avoid issues with distributing the signal across the roadway) and is much more accurate than the required 1 ms accuracy desired. Each of the video images is searched for good-features-to-track points [7]. These are unique locations where there are multidimensional gradients in the contrast (Fig. 3, left). Optical flow routines [8] are then used to link the same points through a series of images (Fig. 3, center). A coherent motion region (CMR) algorithm is used to search for points that retain spatial coherence through the field of view, and based on this a preliminary vehicle detection is generated. Once a vehicle leaves the system’s field of view, its track is refined by filtering the points used and smoothing the motion. A best bounding box is used to determine the extent of the vehicle (Fig. 3, right), and the time when the front of the bounding box crosses from one gamma-ray pixel to the next is determined. These times (in units of shutter releases), together with the location of the front of the vehicle and the distance from the imager to the vehicle,¹ are reported to the gamma-ray imager as a video event.

Meanwhile, the gamma-ray camera has stored its data in a time-ordered list-mode buffer. The time ordering is based on the number of shutter release triggers that have been seen by the system. When it receives a video event, the gamma-ray system finds all of the data for the time that the vehicle was in the gamma-ray pixel that the front of the vehicle just left. The data are used to form a gamma-ray image of the vehicle. A series of such images form the frames of a gamma-ray movie, synchronized with the video

¹ The system actually breaks each vehicle into a set of range slices, each one ~30 cm deep. Video events are generated for each slice, allowing generation of gamma-ray images for each slice as well. This provides information on the location of the source across the vehicle, in addition to the information on the location in length provided by the 1D images.

data stream (Fig. 4, left). To add each frame to a “stationary” image of the vehicle, the gamma-ray image is offset so that the front of the vehicle in the current image lines up with the front of the vehicle in the stationary image (Fig. 4, center). After the vehicle has left the field of view, the stationary image includes all of the data from the vehicle’s transit with the motion blur removed (Fig. 4, right). This image is searched for a radiation source, and if one is found an alarm is generated.



Fig. 3. A single frame from a vehicle is searched for good-features-to-track points (left). The same points are found in subsequent video frames (center) and only those points resulting in smooth motion kept. The filtered points are searched with a CMR and bounding boxes drawn around the vehicles on a frame by frame basis (right). [3]




Fig. 4. The gamma-ray imager generates a series of images based on vehicle events (left--see text). The images are offset so that the front of the vehicle is aligned for each image (center) and a composite image generated (lower right) that is matched to the visible-light image of the targeted vehicle identified by a rectangle (upper right).

3.1.2.2 Environment Upgrades

Several software upgrades were made to the original RST software throughout this project. Some of the upgrades were to the RST software itself, while some were changes in the support software. In the latter category, the most beneficial change was to upgrade to a 64-bit operating system. This provided the primary benefit of making a greater amount of memory available to the TAT processes. The vehicle detection approach used by the system relies on maintaining information for each vehicle for its entire transit through the field of view. With many vehicles in the field of view, particularly if some also had a lot of points, this could require a significant amount of memory. In addition, if a vehicle with many points

was found, the processing sometimes bogged down, further increasing the buffering requirements to the point where the original system occasionally ran out of memory. This was known to happen with heavy multi-lane traffic, and it made the system incapable of handling stop-and-go traffic. With the increased memory available because of the 64-bit OS (and the increased processor throughput and reduced SSD disk write times) the new machines did not suffer this failure mode, and this allowed us to implement routines to handle stop-and-go traffic.

In addition to the change to Windows 7, we also upgraded the remainder of the software suite, including a change to the latest version of the software development environment (Visual Studio), and an upgrade to the latest version of OpenCV [9]. The latter is notable because it includes support for graphical processor units, which could further speed code execution (although those features were not used during the project.) The RST was developed using Microsoft Source Safe version control software, and this product was phased out by the vendor. While revising everything else, we changed to a popular open-source version-control system (GIT). This had a number of desirable features, the most significant being facile reintegration of changes from versions of the code on multiple computers, not all of which needed to be connected to the network. This was particularly useful to integrate changes made to the online code during measurement campaigns as the RST computers themselves were never on an external network.

3.1.2.3 RST Upgrades

System Design

For the RST code base itself, we reorganized the TAT and Master computer code into a single project (these were previously separate) and incorporated many of the TAT development tools as well. Much of this code had been written as needed during development of the RST, usually as separate projects with modules taken from the main code base. Maintaining the separate modules when upgrades were made to the main code base was problematic, and by combining the code into a single project, we simultaneously removed many of these duplicate modules.

Camera Specific

In addition to the reorganization, the code had to be changed to allow use of the new LWIR cameras. These had a lower pixel count but could output either 8- or 16-bit grayscale (vs. 8-bit color for the RST) images. During the course of the project, we also experimented with different alignment cameras to look at the top of the trailer (LWIR vs. visible-light) and contemplated using the two trailers with different stereo rigs (one with LWIR the other with visible-light). To handle all of the different configurations, a modular approach was implemented that allowed the user to set the software to handle the correct cameras when the software was started. This required incorporating flexibility into the code on the image types, but it proved very useful overall. For instance, during the Test and Evaluation phase we were unhappy with the way the contrast in the LWIR images changed over the course of the day when using the cameras' built-in gain control. While in the field, we were able to switch to collection of 16-bit images and implemented our own automated contrast routines that produced images of more value to the operator (Fig. 5). Another change facilitated by the modular design that was implemented in the field was addition of a street-level visible-light camera (Fig. 6 and 7) to aid in vehicle identification to officers responding to alarms.

Overall, the LWIR images were morphologically similar to the visible-light images, and we were able to process the images using the original RST algorithms (and code base). We did have to modify the parameters to optimize performance. In particular, due to less overall detail, the constraints on a successful stereo vision matching pair were relaxed to create sufficient stereo vision points for tracking.

We found that there were more errors in the matches, but the post-processing was robust enough to handle this issue.



Fig. 5. Sample vessel images using the camera auto-exposure (left), and our improved scaling from the 16-bit images (right). Note that the sample image from the camera's auto-exposure mode is an example with particularly good contrast. Conversely the right-hand example from our 16-bit conversion is typical of the imager performance.




Fig. 6. The RST deployed at an intersection during testing at the Port of LA. The picture on the left shows the extra street-level visible-light camera added to the system.

In addition, other improvements were made that would apply to both visible-light imaging and thermal imaging. One such change was an increase in the spacing between acceptable quality points. This resulted in fewer overall candidate points but improved the latency of the system significantly. Without this change, some vehicles with an excessively large number of points taxed the computational resources of the Master computer in the coherent motion region search. This would periodically create significant delays between the times that that vehicle (and subsequent vehicles) left the field of view and the time their final results were provided. We found this condition would persist until the system could catch up with the traffic flow again. Once the change was made, this latency almost completely disappeared.



Fig. 7. Sample image from the street-level visible light camera (right) and the corresponding LWIR image (left).

Stop-and-Go Traffic

To handle stop-and-go traffic, a dynamic approach was selected that allowed for significantly different velocities in different lanes. The main issue with stop and go is the significant amount of data that must be maintained until the vehicle exits the field of view and its track is fully processed. In particular, the stereo reconstruction algorithm becomes a bottleneck. To reduce this, we took advantage of the fact that a slowly moving vehicle does not need to have its location updated as frequently as one moving more quickly. A processing step was added that makes a rapid, monocular camera estimate of the ‘real world’ coordinates of the points from the “good features to track” points [7], and estimates the velocity on the fly. The estimate is made at the step where the relative motion of the camera on the mast is used to correct for false movement caused by the camera motion. Subsequently, the stereo reconstruction is applied less often to points that are moving slower, and a faster interpolation method is applied between frames to assign real-world coordinates to these points in the intermediate frames. Another important but simple change was applied to the retention rate of stationary points. The TAT code assumes vehicles are moving overall, but accommodates a variable period of stationary motion to handle potentially brief instances of stalled motion. This feature also helps to eliminate stationary points on the background (such as on the road surface) that can occasionally get “flagged” as potential points to track. With the improved stereo-vision processing for slower points, we were able to increase the retention time with little or no computational penalty, which allowed us to successfully track periodically stationary traffic. The balance between rejecting true background points and retaining those from stopped vehicles can be tuned by changing this retention parameter, but we found little tuning in real-world data was necessary once we tested the system under stop-and-go conditions.

LWIR Stereo Vision

The LWIR stereo vision system largely was a drop-in replacement for the visible light system with a few issues. First, the mast motion in the visible system was corrected by using fiducial marks on the roof of known position which were tracked by a processing thread using a downward looking “alignment” camera. The positions were then applied to correct for the relative camera motion and so create a stable coordinate system for motion compensation. For the LWIR system, we required a means of tracking the roof coordinates without ambient light. We pursued two options: a visible camera with illuminated roof points and a LWIR camera with some other mechanism such as heat sources or the trailer shape itself. We eventually chose the former because we were able to find very low-cost light-emitting diode (LED) sources (truck brake lights) that worked very well for the roof fiducials. In addition, the LED sources

were very easy to segment from the background as they have virtually no response in the blue plane of the visible-light camera.

The second major issue of using the LWIR cameras was calibration of the stereo vision system and the calibration of the stereo vision to alignment camera. The LWIR stereo system required a large calibration target that would have a strong thermal response; ideally without using active sources. A target was built of aluminum-covered sheeting with black plastic circular landmarks at regular intervals (a standard computer vision target geometry) (Fig. 8). This was found to work well enough, although sometimes environmental conditions caused problems in the detection of the circles (for example, the targets would sometimes get washed out by reflections from other objects in the field of view). For alignment of the visible-light camera to the LWIR stereo system, an “optical” target was needed which could be easily detected in both modes. Again, the aluminum sheet with black plastic targets was a good candidate for this procedure as long as care was taken to avoid environmental conditions that caused problems (bright sunlight could saturate parts of the images obtained with the visible-light camera). A further problem of this calibration was matching the field of view for the visible-light camera to that of the stereo-vision system. We were able to solve this problem by selecting a lens with a wider field of view than used previously for the optical camera. This change was effected without a significant loss in the accuracy of tracking the fiducials on the top of the trailers.

One noteworthy item uncovered in our last several deployments was that few in-field calibrations were needed, in contrast to the visible system. We believe the LWIR system benefited from having smaller field-of-view lenses with less distortion, which made the system more robust to small mechanical movements that forced recalibrations in the visible-light system.




Fig. 8. Visible light (left) and LWIR (right) images of the large Al target used to align the system.

Gamma-Ray System

The code used by the gamma-ray system was both more mature (having been largely kept from an earlier project) and required less computational resources than the TAT system. The software was ported to run under 64-bit Windows 7, but the original machines were retained as they still performed sufficiently well. Since the structure of the video events used to generate the gamma-ray images was unchanged by the switch to the LWIR system, no upgrades to the gamma-ray algorithms were required. However, to deal with the longer dwell times of stop-and-go vehicles in the system, the memory buffer used to store the list-mode data was increased, but this was easily accommodated by the memory available to the software under the 64-bit OS.

However, late in the previous project a conceptual flaw in how the data from both sides of the roadway were combined was uncovered. Specifically, when adding the images, the uncertainties were not properly handled for cases where the vehicle was closer to one trailer than the other. The statistical significance of

a source in the gamma-ray image for each individual trailer is correctly calculated for the peak pixel based on a first-principle propagation of errors using counting statistics [10]. Because of the anti-aliasing approach used by the system, there is covariance in the error for neighboring pixels. (This has only recently been explored for 2D images [11] and that work has not been extended to the 1D images used here.) That means that one cannot just add data from a set of pixels and calculate a total uncertainty based on adding the individual uncertainties in quadrature (the normal approach for uncorrelated errors). When adding the data from the two trailers, the code rebins the data to the lowest resolution image (the one with the vehicle furthest from the trailer), and then adds the two images. In the rebinning process, the errors of neighboring pixels are incorrectly treated as independent, and this is the source of the flaw. It manifests as unreasonably high confidence in the peak pixel.

A procedure to fix the uncertainties was developed that would generate an image from each trailer at the appropriate resolution needed to add the data from the two trailers without rebinning. In that case, simple error propagation would be correct. However, the changes were never fully debugged. At the time that the code was being debugged, data from a stationary source were needed from both trailers to disentangle phase issues in adding the gamma-ray data to the new resolution images, but the system was offline for other development work. By the time that the hardware was up and running, suitable data were collected, but the project had advanced to a point where it was no longer feasible to implement the changes. The recent work on calculating the covariance terms from neighboring pixels correctly [11] should provide both an easier fix, and also allow using the data from several neighboring pixels to enhance the sensitivity from the individual trailers. However, such improvements are left to future development.

3.2 LWIR SYSTEM

The Project used FLIR, Inc. Tau 2 high-resolution LWIR cameras [12]. These are based on a micro-bolometer array with an output from each pixel representing the LWIR flux seen by that pixel. The flux is a function of both the temperature and emissivity of the pixel in the scene being observed. Generally, the emissions are from the object itself, and are not the result of ambient LWIR radiation reflected off of the object (although such reflections can occur). By imaging the thermal emissions, no external lighting is required, and the system works equally well during the day and at night. Because the emission rate is dependent on the absolute temperature of the object, the system also works well over the environmental extremes sampled (257–310 K, Fig. 9). By the end of the project we had constructed LWIR-based stereo rigs for both the local and the remote trailers. These are described in more detail below.

3.2.1 Hardware

Each LWIR stereo rig (Fig. 10) was designed as a plug-in replacement for the visible-light stereo rigs developed for the RST. The design is based on a 1-m-long, 2.54-cm-diameter stainless steel tube that maintains the spacing and pointing between the stereo pair. The cameras attach to aluminum plates that clamp onto the tube and are covered by an open-fronted environmental enclosure. The cameras are recessed in the enclosure and pointed out and down to see the roadway from the top of the 50-ft mast. The enclosure and pointing are sufficient to shield them from rain. A similarly mounted, downward-looking, visible-light camera is used to monitor illuminated targets (Fig. 11) mounted to the top of the trailers and is used by the software to track the motion of the stereo rig with respect to the field of view of the gamma-ray camera in real time. The tube itself mounts with a clamping assembly to an aluminum plate that has a quick release mechanism to either of two mating plates, one in the trailer for transport and another at the top of the mast for use. A central box includes hermetic connectors for the three Ethernet cables (one for each camera) and an additional connector that carries power and the signal used to trigger the cameras' shutters. The power and signals are distributed from the central box to each of the cameras on the rig via permanent connections so that users need only make connections in one place.




Fig. 9. When compared to visible-light pictures (inset), the LWIR images have a remarkably consistent visual morphology. Three of the images shown represent the transition from late afternoon to nightfall. The image on the bottom right was taken on a winter night when it was much colder. The circular vignetting in that image is due to an issue with the flat field correction (see text) and was later fixed. In all instances, many

points (cyan) are found on the vehicles in the field of view.



Fig. 10. Old visible light stereo rig mounted for travel inside one of the trailers (left) and the new LWIR stereo rig mounted on top of the mast (right). The pair of cameras for stereo are in the end housings, and the alignment camera to monitor targets on the roof of the trailer is in the remaining enclosure. This box is larger for the LWIR rig because there was initially space to try both an LWIR and a visible light camera for this task.

For the RST to function effectively, the images from all of the cameras must be taken simultaneously, with a time stamp provided in the images and entrained in the gamma-ray data stream. The RST visible-light cameras accepted an external electrical pulse to trigger the camera shutters and also recorded the time this occurred using a local oscillator with the time included in the exchangeable image file data sent with the image.² The LWIR cameras are designed with a camera-link interface, and while they accept an external trigger, they do not record the time. With the limited length of camera-link cables (without the use of an active repeater) and the lack of synchronization information, this interface was insufficient for our needs. Original plans were to use a local microprocessor board running a real-time operating system to receive the images over camera-link, add a time stamp and send them to the TAT computer via Ethernet. However, a search of available products indicated a commercial solution that performed these functions existed [13], and after acquiring and testing one of the units, it was found to function acceptably. With this extra component, except for the different number of pixels and grayscale vs. color, the LWIR cameras became drop-in replacements for the visible-light cameras.




Fig. 11. Roof of one of the trailers as seen by its alignment camera (left). The black and white optical targets on the roof of the trailer were used by the original RST system. The new LED-based red lights are also visible. At night, the LED lights are all that is visible (right). While the LWIR cameras could effectively “see” in the dark, we still had to determine the alignment of the stereo rig to the trailer. To achieve this, we explored the use of either a third LWIR camera to look at LWIR targets on the top of the trailer, or the use of an extant visible-light camera together with lighted targets on the roof of the trailer. In the end, we found that designing a passive LWIR target that worked for all environmental conditions was difficult and settled on an active (heated) target design. However, once we required an active target, and given the relative costs of the LWIR vs. the visible-light cameras, we decided to use an active (lighted) visible target system. After dark, this made the system more visible to passing vehicles. However, the light was directed upward and did not pose a traffic hazard. Further, the system was never designed for use in low-profile operations, so the light leakage was deemed acceptable.

The final mechanical issue associated with the LWIR system was the overall alignment of the system. This included determining the intrinsic distortions of the LWIR cameras, the extrinsic alignment of the stereo cameras to each other, and the extrinsic alignment of the stereo rig to the gamma-ray imager. To perform the alignments, images are taken of an ‘optical’ target with a known pattern with the target in different orientations, distances, and locations in the field of view. There were two issues in translating this process from visible-light to LWIR. First, we needed to develop a suitable optical target for use in the different alignments. Because of the thermal nature of the LWIR images, making a target with a regular pattern was challenging. After some experimentation, this was solved by placing circles of black tape on an aluminum metal sheet. The different emissivity of the two materials provided sufficient contrast (particularly if the sheet was warm) to allow the calibrations to be performed. The second issue was due to the smaller fields of view of the LWIR cameras (compared to the RST visible-light cameras.)

² The local oscillators were reset every second to provide an absolute time accurate enough to relate the images to the absolute time as recorded by the (non realtime) Windows operating system.

Alignment of the stereo rig to the gamma-ray imager field of view required that the target be visible to both the LWIR cameras and the alignment camera at the same time. With the smaller LWIR field of view, there was less overlap, making the overall calibration difficult. In the end, a large target was used and positioned so that most of it was visible to all of the cameras. Some efforts were made to allow processing of the images even if the full target array was not visible to the cameras, but these efforts were eventually abandoned in favor of using a lens with a larger field of view for the visible-light alignment camera. This provided more overlap between the fields of view of the two imaging systems. The camera was positioned so that the trailer was largely in the bottom area of the camera view with some excess to allow for mast motion.

4. SYSTEM PERFORMANCE

The performance of the upgraded RST was evaluated using data collected late in Phase 3 (particularly for cold weather data), and in several measurement campaigns in Phase 4. As for the original RST, the performance of the system was determined in both the optical (LWIR) and the gamma-ray bands. This follows since failure of the optical tracking system to detect a vehicle means that vehicle is not searched for a gamma-ray signature. If a target is optically detected, then the result from the gamma-ray imager depends on how well the vehicle is tracked since poor tracking causes blurred net gamma-ray images. However, based on earlier work with the RST, we were able to verify that if the tracking was good, then the gamma-ray results were basically determined by the counting statistics associated with the source. This meant that the system performance could be evaluated based simply on the accuracy of the target tracking which was originally planned for the use of RTK-enabled GPS receivers, but as described below, was evaluated using a linear fit model instead. This eased the logistics of dealing with nuisance alarms from real, unplanned sources in the traffic flow, and, consequently, the use of radioactive sources was minimized during offsite testing.

4.1 ORNL TESTS

The primary validation with sources was performed at a standard test site on the main roadway through ORNL. Tests replicating those performed on the original RST with both a source and an RTK-enabled GPS were conducted at this location. However, the RTK-enabled GPS data experienced considerable problems, and therefore as a proxy, the track error was computed as the deviation from a linear fit. This was based on the findings from previous work that the GPS tracks are generally linear, barring lane changes, which are extremely rare.

4.1.1 Vehicle detection rates

The first issue addressed by the tests was determining the TAT vehicle detection rates. These were evaluated by comparing the TAT output with results obtained by manually counting the vehicles observed by replaying the video data. Tables 1 and 2 give the results obtained at the primary ORNL test site for both the Local (LTAT) and Remote (RTAT) trailers, as well as the results for the combined system (where available). Table 2 has the results for cold weather. These were performed before the second trailer had been upgraded to LWIR tracking, so the results are for a single trailer only. Results for temperatures near freezing and a milder winter afternoon are given as well. There appear to be some differences between the two conditions; in particular, sometimes vehicles in cooler weather had fewer points detected. This was likely caused by cooler vehicles that had traveled shorter distances in the morning.

Table 1. Bethel Valley Road (BVR) Vehicle Detection Rates

Season	Vehicles	Combined	LTAT	RTAT	Multiple
--------	----------	----------	------	------	----------

Winter	977	NA	844 (86.4%)	NA	38 (3.8%)
Spring	42	42 (100%)	42 (100%)	42 (100%)	14 (33.3%)

Table 2. BVR Winter Vehicle Detection Rates

Conditions	Vehicles	Combined	LTAT	RTAT	Multiple
Mild	414	NA	411 (99.3%)	NA	21 (5.0%)
Cold	431	NA	410 (95.1%)	NA	14 (3.2%)

Table 3. Prior RST Visible-Light Vehicle Detection Rates

Vehicles	Combined	LTAT	RTAT	Multiple
2030	2015	1867 (92.0%)	1877 (92.5%)	62 (3.0%)

The results for the individual trailers compare favorably with those obtained previously with the optical system (Table 3), although these results are composite from several different locations that do not include Bethel Valley Road (BVR). While the combined results from the spring look excellent, there are really too few vehicles to fully test the system. However, they do give a taste of the improved performance obtained with the final system due to upgraded software, better flat fielding techniques, and improved contrast in the images. More extensive tests were conducted offsite under much more challenging conditions and are reported later in this report.

4.1.2 Tracking performance

For detected vehicles, the gamma-ray sensitivity will be affected by the accuracy with which the vehicle's location is determined. To estimate those we compared the TAT results with results obtained with a simple linear model of the vehicle motion. In short, the errors tell how much the reported vehicle location fluctuated about where one would expect the vehicle to be had it passed through the system at a constant velocity and at a constant range (y value). At the BVR site, the average errors in mm for the different trailer configurations are shown in Table 4.

Table 4. BVR Test Site Tracking Errors

Configuration	Mean X Error (mm)	Mean Y Error (mm)
Both	271.1	225.1
LTAT	294.0	329.2
RTAT	145.2	96.8

For this setup, the error was dominated by the results from the Local trailer. This is presumably due to the setup geometry on the three-lane roadway which had the trailers separated by a gap commensurate with a five-lane roadway. The Local trailer was offset from the nearest actual lane by two lanes of grass so that except for the middle of five lanes all traffic was closer to the Remote trailer. One expects the errors to be larger the further a vehicle is from a trailer since the resolution of the camera decreases with distance. Combining the results for the two trailers yields somewhat better results than the results for the LTAT alone, while the error for the RTAT is significantly lower. The errors are similar in magnitude to those reported in the Independent Study performed on the system by DNDO as seen in Fig. 12. Unfortunately,

those results are only available graphically, but it is clear that tracking errors of the order of a meter are not unrealistic for the system, and with those values the system still generated good gamma-ray detections.

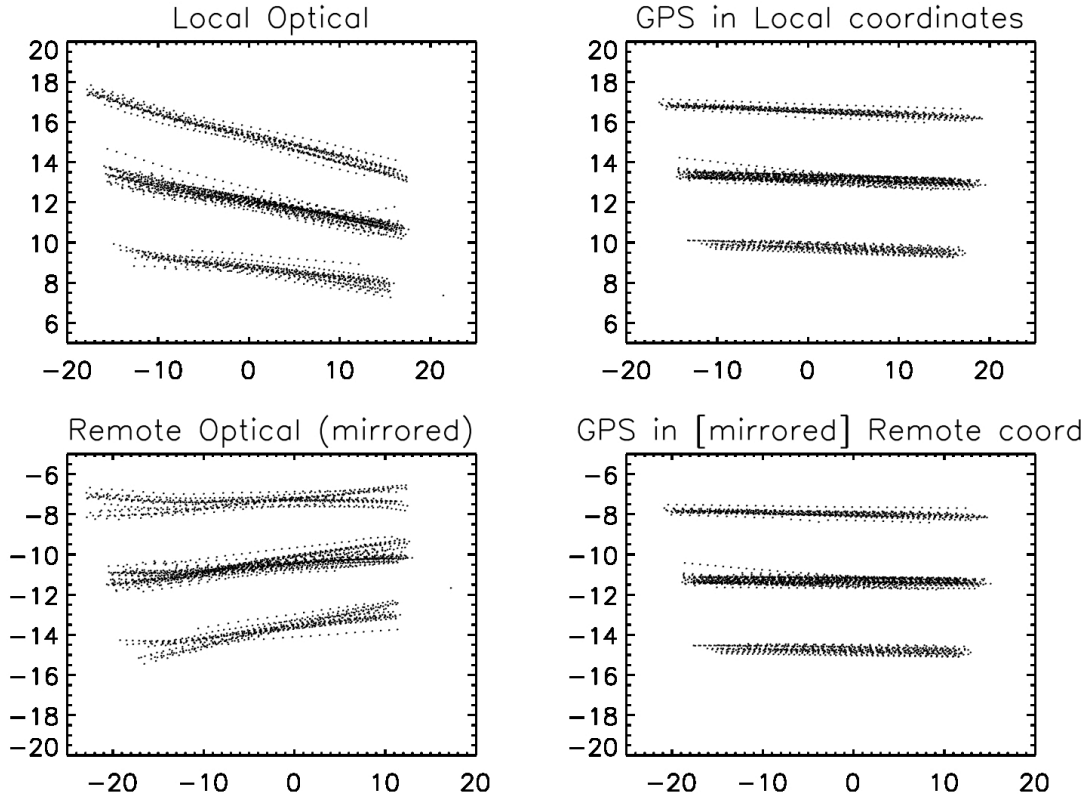


Fig. 12. Tracking (units are in meters) reported for the optical RST from the independent evaluation [14]. Comparing the Optical and GPS results clearly indicates average errors >100 cm. Source tests

The performance of the full system (LWIR plus gamma ray) was also tested using radioactive sources. The results are presented in Table 5. Unfortunately, we note that the results had obvious problems with respect to the local/remote alignment of the gamma-ray images (Fig. 13). Consequently, reporting the combined data as automatically generated by the system does not make sense since the statistical significance of detections will be lower than those obtained with a single trailer. We note that this issue was noted at the time of the measurements, and we assumed that a more careful offline calibration could be used to fix the results. However, such attempts proved fruitless, and in fact a closer look at the results shows an unexplained anomaly in the lengths of the vehicles in the center lane (equidistant from both trailers). These are significantly different even though they should be approximately equal in length. This same error was not seen in later deployments and indicates a fundamental problem with the data that could not be resolved. To compare these results with earlier measurements taken with the visible-light cameras, we have used the theoretical maximum obtained by adding the individual peak values from each side of the roadway in quadrature. While this will overestimate the performance, since we do not expect an exact match in the peak location for the two trailers, it does allow some comparison with the earlier results. The primary goal of the source tests was to validate that the LWIR did not degrade the performance of the system, and the fact that each trailer performed reasonably well on its own proves that point. However, the lack of a properly combined image from the two systems does indicate that some work remains to be done to improve the correspondence between the local and the remote trailers. Although, this was a problem for the visible-light system as well, the different vehicle lengths seen in

these data represent a new failure mode which is presumably due to the use of incorrect parameters in the system setup.

Table 5. BVR Source Detection Results

Velocity	Sigma Vel	# Passes	Local Sigma	L Sigma Var	Remote Sigma	R Sigma Var	RMS Sum	RMS Sigma
52.8	1.1	5	15.4	1.0	9.7	1.1	18.2	1.1
39.1	2.1	4	17.0	0.9	10.1	1.1	19.8	1.2
30.2	0.6	6	19.6	1.4	12.7	1.4	23.4	0.6

To determine the net impact of the use of LWIR cameras, we compare the results obtained at ORNL to those obtained in the tests conducted by our team at Savannah River National Laboratory and subsequently reported in [3]. For the comparison, the most recent results must be scaled for the differences in source strength and background rates of the two evaluations. The recent data were collected with an 863- μCi ^{137}Cs source, comparable to the 944- μCi ^{137}Cs source used at Savannah River National Laboratory (SRNL). To compare the numbers, the statistical significance must be adjusted by the square root of the ratio of the source strengths so that, exclusive of background rates, the ORNL results should be 95% of those obtained at SRNL. However, the backgrounds were different at the two locations as well, with the SRNL rates ~ 1.4 times higher than those at ORNL, meaning the ORNL rates should be higher by the square root of the ratio of the backgrounds or, ~ 1.2 times better. In combination, the two factors indicate that the ORNL values should be 1.15 times those obtained at SRNL. Adding the corrected values to Fig. 17 of [3] we obtain the plot in Fig. 14. As can be seen, the performance with the LWIR system is slightly better than that of the visible-light system, but this can easily be explained by the forced alignment between the results from the two trailers.




Fig. 13. Local and Remote results obtained in testing with sources. The locations of the gamma-ray peaks obtained from the two sides of the roadway are not aligned. Of greater concern is that the lengths of the vehicles are significantly different from the two sides of the roadway, even though the vehicles were approximately equidistant from each system. Overall, the Remote trailer has weaker detections, and this is likely due to the alignment issues discussed above since all of the processing is done in terms of the “local” coordinates. It cannot be traced to different background rates between the two trailers as these were the same to better than 4%, with the local background slightly higher than that of the remote.




Fig. 14. Comparison of latest LWIR results compared to the original visible-light RST data [3]. The new data show a small increase in sensitivity compared to the old data; however, even if statistically significant, which is doubtful, the higher result can be attributed to the optimal combination of the data from the two sides of the roadway (see text). Stop-and-go Traffic

In addition to the standard test location, the system was used during two deployments near a traffic light on the same roadway through ORNL to determine performance of the stop-and-go system. Because of safety concerns by ORNL, the second deployment was made using only a single trailer. Unfortunately, data from the first deployment were lost so that dual results are not available from this location. However, subsequent tests in the Los Angeles area included dual-sided stop-and-go events. The results are summarized in Table 6. As can be seen, the detection rate is similar to that obtained with flowing traffic. The table includes data collected both during the day and at night. While there is some reduction in the performance at night, this is attributed to a greater number of larger vehicles (particularly the rate of multiple detection) in the latter data.

Table 6. BVR, Stop-and-Go Vehicle Detection Rates

Conditions	Vehicles	Combined	LTAT	RTAT	Multiple
–	621	NA	617 (99.4%)	NA	34 (5.5%)
Day	487	NA	485 (99.6%)	NA	18 (3.7%)
Night	49	NA	48 (98.0%)	NA	11 (22.4%)

4.2 FREEWAY PERFORMANCE

To determine the system performance with multi-lane, high-velocity traffic, we deployed on either side of a three-lane section of Interstate 40, approximately 10 miles east of Knoxville, Tennessee (Fig. 15). The location was primarily selected based on a paved U-turn for official vehicles, the flatness of the shoulders, the presence of guard rails on both sides of the traffic lanes (to provide protection for the personnel and the instrument), and the ability to drive behind the guard rails. The speed limit was 65 mph. Data were

collected at all times of day with varying traffic densities. As a major east-west route, a significant fraction of the traffic is large trucks, and this is different than most locations used to test the RST.



Fig. 15. The RST deployed on the eastbound lanes of Interstate 40.

Of the data that were collected, we selected four files to explore the TAT performance. The results are summarized in Table 7. We note that the truck traffic did result in a large number of multiple retirees.

Table 7. Freeway Vehicle Detection Rates

Conditions	Vehicles	Combined	LTAT	RTAT	Multiple
Day	2351	2305 (98.0%)	2220 (94.4%)	2206 (93.8%)	1176 (50.0%)
Night	263	259 (98.4%)	256 (97.3%)	251 (95.4%)	1304 (49.9%)
Total	2614	2564 (98.1%)	2476 (94.7%)	2457 (94.0%)	1304 (49.9%)

Finally, the tracking error was again calculated (Table 8) and found to be significantly lower than that obtained at the BVR location. This is attributed to the smaller road range covered (three lanes plus shoulders). We also surmise that the faster speeds found on the interstate are a contributing factor in the reduced error. This seems counter-intuitive, but in fact the faster, more uniform/constant, speeds are better suited to some aspects of the tracking, which tends to prefer tracks that show stable, continuous motion over those with more erratic behavior. Based on comparison with the ORNL source results, the low tracking error should result in excellent gamma-ray detection performance.

Table 8. Freeway Tracking Errors

Configuration	Mean X Error (mm)	Mean Y Error (mm)
Both	32.9	186.4
LTAT	26.7	289.3
RTAT	94.7	118.9

4.3 SMALL VESSEL TRACKING

To test the performance of the system for use with small vessels at navigable choke points, the instrument was deployed at the Boynton Beach Inlet in Boynton Beach, Florida (Fig. 16). This location was selected because of its unique attributes: a narrow channel (50 m) with parking on either side of the channel. In general, distances associated with vessel detection are larger than those associated with roadways and because of this they presented a challenge to the RST due to its design for use on the latter. The physics of source detection indicates that imaging, square-meter-class radiation sensors are viable at ranges up to ~100 m. With an instrument on either side of a channel, this means a 200-m channel could be effectively handled by a suitably adapted version of the RST, and many channels that meet these criteria can be found [2]. In principle, only the stereo rig would have to be changed so that it points out further from the trailer (the current field of view is fixed to a range of ~30 m), although in practice additional changes would be required (presumably a change in lens focal length, a reduction of the mast height, tracking changes to better handle non-linear motion, etc.).



Fig. 16. The RST on either side of the Boynton Beach Channel.

The tests were scheduled to cover a two-week period at the end of June and beginning of July, with one week set aside to adjust parameters of the TAT algorithm for use with vessels. In practice only a few days were required to make the tracking work effectively. One of the biggest issues was the limited field of view of the LWIR cameras due to their downward orientation. The cameras from each side saw only a small way past the center of the channel. We opted to run in this configuration, using the tracking information from whichever trailer saw a vessel to generate the final solution. This was fairly effective in making the system work.

We did find that the contrast of the vessels on the water obtained from the automatic camera gain control was at times aesthetically unpleasing (Fig. 5). To counter this, we implemented 16-bit image collection and used an automated software loop to select the range of values from within the 16 bits to generate the 8-bit images processed by the system. This capability had been tested earlier, but the camera firmware did not fully support it at that time.

For the data collection, we generally tracked whatever traffic happened to pass by. We hired a boat and collected RTK GPS data one afternoon, but the collection unit was not properly reset between runs, and consequently the resulting file was impossible to process due to its length. In addition, on one day local law enforcement representatives provided a weak ^{137}Cs check source placed on a Sheriff Department Vessel (Fig. 17). Given the source strength, it is not surprising that it was only observed in low-velocity passes near a trailer.

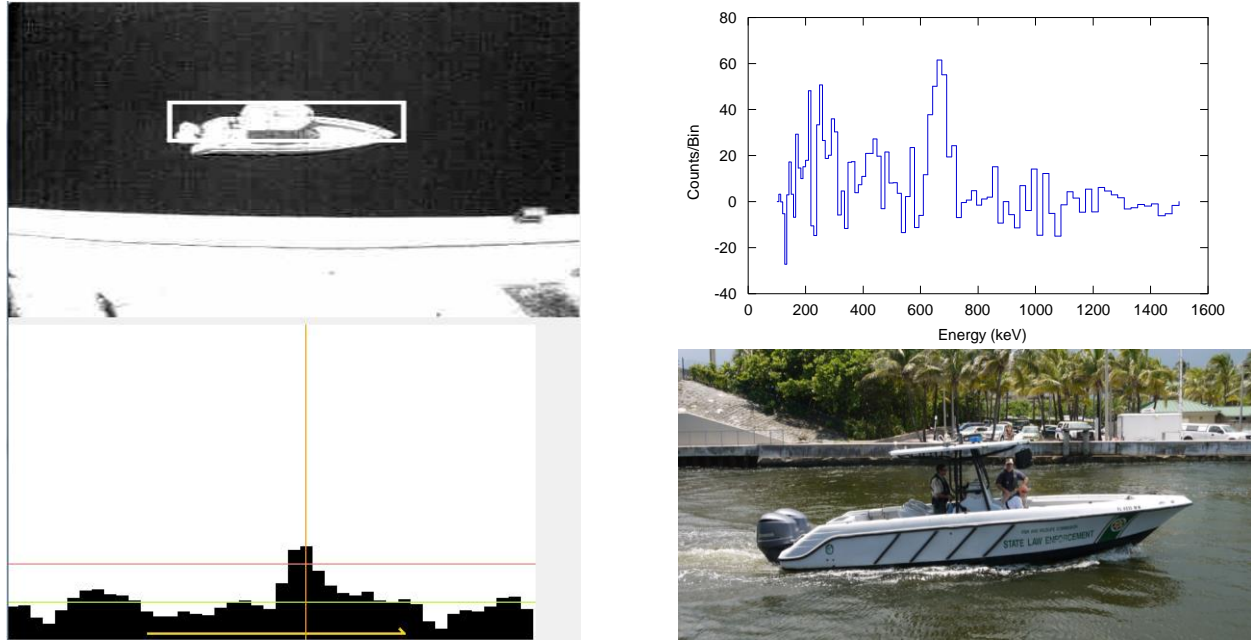


Fig. 17. Picture of the Sheriffs' vessel with a weak ^{137}Cs check source in the black bag hanging from the front of the roof (lower right). The RST output (left) shows the LWIR image (top) of the vessel together with its gamma-ray image (right). The spectrum from the gamma-ray pixel under the orange cursor is shown on the top right. The tests were conducted near the summer solstice with temperatures fairly constant. We noted a change in the water temperature passing through the inlet depending on the direction of the tide. The water coming in from the ocean on the flood tide was cooler than that from the Intracoastal Waterway as the tide receded. Boats were also generally hotter when the sun was near its zenith. A run starting before sunrise and one running from late afternoon through nearly midnight did not have any issues with the tracking (although there was very little traffic at the beginning of the former and near the end of the latter). Based on this, additional runs in the very early morning were not attempted.

Table 9. Small-Vessel Detection Rate

Vessels	Combined	LTAT	RTAT	Multiple
25	25 (100%)	25 (100%)	25 (100%)	15 (60%)

The results of the boat detection by the TAT system for some of the data are presented in Table 9. As can be seen, the detection rate is quite good. When combined with the tracking results presented in Table 10, it is clear the system can be used effectively in such an environment. The higher tracking errors are likely due to the much greater distances involved in the maritime configuration. Since the gamma-ray pixels will be similarly larger, we do not expect a significant degradation in performance, although some optimization might be warranted for a dedicated maritime system.

Table 10. Small-Vessel Tracking Errors

Configuration	Mean X Error (mm)	Mean Y Error (mm)
Both	1,441	709
LTAT	1,066	1,067
RTAT	1,099	446

We took advantage of the extra time afforded by the shorter-than-expected adaptation requirements to explore the issue of multiple detections. In the TAT software, once the good points have been sent to the

Master computer by both TAT machines, the points are combined into vehicles by using “coherent motion regions” (CMRs) [15]. Some of the parameters used by this algorithm are related to general vehicle sizes. We reprocessed a sample maritime run using larger CMR parameters (geared toward larger vessels). These parameter changes resulted in far fewer multiple detections, with no loss in overall detections, as shown in Table 11. Based on this promising result we continued this investigation in the test conducted at the Port of Los Angeles.

Table 11. Small-Vessel Detection Rates With Different CMR Parameters

CMR Region	Vessels	Combined	LTAT	RTAT	Multiple
Large	25	25 (100%)	25 (100%)	25 (100%)	4 (16%)
Default	25	25 (100%)	25 (100%)	25 (100%)	15 (60%)

4.4 DEMONSTRATION EXERCISE, PORT OF LOS ANGELES

The final major deployment of the system was at various locations around the Port of Los Angeles. For most of the settings, we had both emergency response support and different sources available to fully test the system. We also conducted some blind tests where neither the source vehicle nor the source type were known. We deployed at three different locations. The first location was on Navy Way, a four-lane causeway leading to various piers for large container ships (Fig. 18). The traffic at the site (as was much of the traffic in the area) was dominated by trucks hauling cargo containers. Both the second and third sites were on Harry Bridges Boulevard, which has two travel lanes in each direction with alternating left-turn lanes adding a fifth lane. The second site was just east of the N. Marine Avenue intersection, and the third was just west of the Broad Avenue intersection. The West Broad intersection included a traffic light, and we performed tests with stop-and-go traffic at this location. One trailer at the site was parked halfway onto the sidewalk (Fig. 19). The sites were selected to represent an urban roadway, although traffic was never very heavy. The proximity to the port also meant a large number of heavy trucks used the roadway. At the third site, we collected data through the evening hours, but unfortunately traffic fell off precipitously after ~6:00 p.m.



Fig. 18. The RST on a 4-lane causeway leading to piers for container ships.

Due to the large number of trucks at all of these sites, we experimented with the size of the CMR used to select individual vehicles from the set of good-features-to-track points generated by the code. This significantly reduced the number of multiple retires, but also had the unfortunate side effect of sometimes grouping multiple vehicles into single large vehicles. It also caused us to misidentify the source vehicle for one of the source runs. In this case the source vehicle was following another vehicle, and the system never recorded the source vehicle—attributing the detection to the lead vehicle.

During the trials in Los Angeles, we added a street-level, visible-light camera to the system (Fig. 6). The images from the camera were synchronized and recorded with the rest of the system, and images were matched with the time when vehicles were near the center of the system’s field of view. This proved very useful for describing the offending vehicles to prepositioned chase vehicles. One other change to the code made in Los Angeles also facilitated communications with responding officers. That was an increase in the minimum spacing between points that were processed. For some trucks, this greatly reduced the number of points and virtually eliminated long-standing issues with latency developing when such vehicles transited the system. It did not seem to affect the detection rate. The performance of the TAT system for sample runs is given in Table 12



Fig. 19. The RST deployed at an intersection on Harry Bridges Road near the Port of LA.

Table 12. Port of LA Vehicle Detection Rates

Area	Vehicles	Combined	LTAT	RTAT	Multiple
Surface Street	84	83 (98.8%)	83 (98.8%)	53 (63.1%)	17 (20.2%)
Terminal Access Road	109	109 (100%)	109 (100%)	109 (100%)	104 (95%)

A final test site by the side of a channel to work with vessels turned out to be non-productive. The channel was ~300 m across, and only one trailer was set up. However it was ~15 m from the edge of the water, so the system field of view did not extend far enough into the channel to see the few small vessels that passed by. We did detect a marine patrol vessel that was available for a short time and passed near the system.

As in the maritime case, another comparison was made regarding the merging of multiple vehicles into a single detection. We reprocessed two Port of Los Angeles runs with a larger CMR that was geared more toward truck detections. As before, these parameter changes resulted in far fewer multiple detections, with little loss in overall detections, as shown in Table 13.

Table 13. Port of LA Vehicle Detection vs. CMR Parameters

CMR	Vehicles	Combined	LTAT	RTAT	Multiple
Large	375	372 (99.2%)	373 (99.4%)	261 (69.6%)	22 (5.8%)
Default	375	374 (99.7%)	366 (97.6%)	369 (71.3%)	184 (49.1%)

5. CONCLUSIONS

The work developing a non-visible TAT approach that was the primary thrust of this project resulted in a robust tracking system that made effective use of the algorithms previously developed for the visible-light RST. In fact, in some ways, the LWIR images are preferred for tracking even during the daytime since these images are remarkably similar regardless of the ambient conditions (Fig. 9). This was found to be particularly useful in vessel tracking where sun glints, ripples, waves and even the whitewater from vessel wakes largely disappear in the LWIR images (Fig. 20). While use in dense fog was never fully tested, results of images in lighter fog and mist did show some benefit for the use of LWIR. Of more benefit than the ability to see through the fog was the complete elimination of exposure issues associated with low sun angles causing very bright scattered light in the field of view. Of greater importance, the system worked around the clock with no observed impact due to the diurnal temperature cycle.




Fig. 20. Sample images obtained with the LWIR system (right) compared to an image obtained with the original RST (left). While the velocity of the former was measured at twice that of the latter, the wake and associated white water is just barely visible in the LWIR image. This makes the TAT tracking much more reliable. The use of LWIR did exacerbate one of the long-standing issues with the RST found during field trials, the ability to identify a specific vehicle to first responders. Due to the overhead images provided by the system, it is hard to identify a vehicle even in visible-light images. The use of LWIR further complicates identification because the grayscale in the images has no bearing on the colors or even grayscale seen by the eye. To address this, addition of a street-level, color, visible-light camera late in the project proved very useful. The addition of further cameras with multiple views to overcome obscuration by other traffic should go a long way toward identification of a suspect vehicle. Identification was also compounded by the sometimes long detection latencies (10s of seconds) and the subsequent manual reporting via radio to the first responders. The long delays were virtually eliminated by late changes to the point detection software that significantly reduced the system latency, providing alarms within a few seconds of a vehicle exiting the system field of view. Although beyond the scope of the project, the issues associated with verbal vehicle descriptions could be removed by automatically sending street-level, color images to chase vehicles via an internet link. Finally, for help in vehicle identification, the system could be integrated with automated license plate readers (LPRs) triggered or synched to the system. Discussions with LPR vendors indicated that it might be relatively simple to add the RST results to their extant database for rapid distribution to interdiction vehicles. The biggest obstacle to full integration with LPR systems will be obscuration, since the system is parked on the side of a multi-lane roadway.

The project also included significant upgrades to the code base and migration to newer software and hardware platforms. Both of these proved valuable. In particular, migration to a 64-bit Operating System provided the TAT software with more memory, resulting in much more reliable long-term operation. While not as obvious in terms of performance, the software suite underwent significant overhaul, removing redundancy, and improving modularity while making work on the code easier. This allowed some of the infield upgrades (16-bit to 8-bit automated gain handling, extra street-level camera, tests of different CMR vehicle sizes) during the final testing phases of the project to be implemented successfully.

While the overall project achieved its major goals, there are still some issues that need to be addressed to make a fully deployable system. Chief among these is the alignment between the coordinate systems of the Local and Remote trailers. This issue arose repeatedly throughout the life of the project and still requires some work. The alignment is currently accomplished by manually finding common points in the field of view between the two systems. Errors in the associations are computed as part of this process, and discrepancies of order a meter are not uncommon, particularly near the edges of the fields of view. While these are small enough for the system to function, it is clear that this represents an error of many pixels in the images, and one should be able to do better than this. One of the problems is the difficulty in finding points that are common to all of the cameras throughout the field of view, leading to selection of a small set of points, often with many of them in only one region of the image. One possible solution to this problem would be to use vehicles transiting the system as a source of points. While views from opposite sides of the roadway may make point-to-point correspondence difficult, use of aggregate results from many vehicles (vehicle bounding box or mean location) could provide a solution that is optimized in the portion of the image that is needed to track the vehicles.

With a detection miss rate of order 1%, the TAT system exceeds the performance obtained by conventional traffic detection systems. Nevertheless, there are still some issues associated with poor bounding boxes and multiple detections of the same vehicle. The former should be significantly eased if one can improve the correspondence between the Local and Remote coordinate systems. Some improvement of the latter was achieved by tuning the CMR parameters, although sometimes at the expense of missing one of several closely following vehicles, or at the expense of combining them into a single large vehicle. Further optimization of the parameters and/or the use of multiple CMR threads, each tuned for a different vehicle sizes with a conflict resolution to combine the results, could improve the overall performance.

The gamma-ray system also has a few outstanding issues that must be addressed. Prime among these is obtaining the correct combined detection significance for vehicles far from the center of the roadway. While the approach that was developed but not fully implemented as part of this project is conceptually sound, it does not allow for proper addition of more than one pixel of the source. Recent work on covariance in neighboring pixels of coded-aperture images should help address that issue as well.

Beyond the aforementioned issues, there remain several upgrades that are commonly incorporated in deployed systems but beyond the scope of a developmental instrument. This includes such things as automated gain calibration and automated isotope ID for the gamma system, overall improvements in the user interface, and a reduction in the amount of data that is saved.

Because the imaging approach pioneered by the RST ties the radiation signature to a specific location of a specific target, it means that the data from different imagers can be combined to increase overall sensitivity. The potential value of this approach can be applied to system geometries beyond that specifically pursued by the RST. This includes incorporation of an imaged system for fixed multi-lane installations such as border crossings and toll plazas; potentially significantly reducing the problems where a single vehicle causes alarms on multiple vehicles in several lanes. It also means that rapid

response systems based on using several smaller imagers in parallel (at the same location of a roadway) or serially (one after the other on a roadway) could be developed.

6. REFERENCES

1. K.P. Ziock, J. Collins, M.F. Cunningham, L. Fabris, T. Gee, J. Goddard, F. Habte, T. Karnowski, "The Use of Gamma-Ray Imaging to Improve Portal Monitor Performance," *IEEE Trans. Nucl. Sci.*, **55**, 3654-3664, 2008.
2. K.P. Ziock, A. Cheriyyadat, J. Goddard, D. Hornback, T. Karnowski, R. Kerekes, J. Newby, "Autonomous Radiation Monitoring of Small Vessels," SORMA 2010 (invited), *Nucl. Inst. Meth. in Phys. Res. A.*, **652**, 10-15, 2011.
3. K.P. Ziock, E.C. Bradley, 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. Nucl. Sci.*, **60**, 2237-2246, 2013.
4. R.A. Kerekes, J.S. Goddard, A.M. Cheriyyadat, T.P. Karnowski, M.F. Cunningham, D.E. Hornback, L. Fabris, K.P. Ziock, *Proceedings of 2011 Future Instrumentation International Workshop*, Oak Ridge, Tenn., Nov. 7-8, 2011.
5. 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, "Design of Dual-Road Transportable Monitoring System for Visible Light and Gamma-Ray Imaging," *Proc. SPIE*, **7665**, doi:10.1117/12.850191, 2010.
6. 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*, **7538**, doi:10.1117/12.838428, 2010.
7. J. Shi, C. Tomasi, "Good features to track," *Proc. Computer Vision and Pattern Recognition '94 IEEE Comp. Soc. Conf.*, pp., 593-600, 1994.
8. 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.
9. More information at, [Opencv.org](http://opencv.org).
10. E.E. Fenimore, "Coded aperture imaging: predicted performance of uniformly redundant arrays," *Applied Optics*, **17**, 3562-3570, 1978.
11. M.C. Fleenor, M.A. Blackston, K.P. Ziock, "Correlated statistical uncertainties in coded-aperture imaging," *Nucl. Inst. & Meth in Phys. Res. A*, **784**, 370-376, 2015.
12. FLIR Tau 2, information at: <http://www.flir.com/cores/display/?id=54717>.
13. Pleora iPort PT1000-CL, 340 Terry Fox Drive, Suite 300, Kanata, ON K2K 3A2, Canada. Information at: <http://www.pleora.com/our-products/frame-grabbers/iport-pt1000-cl>.

14. “The Roadside Tracker Technology Characterization Report,” Domestic Nuclear Detection Office, Department of Homeland Security, 2013.
15. A.M. Cheriyyadat, B.L. Bhaduri, and R.J. Radke, “Detecting multiple moving objects in crowded environments with coherent motion regions,” in *Proc. IEEE Comp. Soc. Conf. Comp. Vision Pattern Recognit.*, pp. 245–252, 2008.