# Autonomous Radiation Monitoring of Small Vessels

*K.P. Ziock, A. Cheriyadat, L. Fabris, J. Goddard,*
*D. Hornback, T. Karnowski, R. Kerekes, J. Newby*

Oak Ridge National Laboratory, Oak Ridge, TN

Corresponding author email address: ziockk@ornl.gov

**Abstract**
Small private vessels are one avenue by which nuclear materials may be smuggled across international borders. While one can contemplate using the terrestrial approach of radiation portal monitors on the navigable waterways that lead to many ports, these systems are ill-suited to the problem. They require vehicles to pass at slow speeds between two closely-spaced radiation sensors, relying on the uniformity of vehicle sizes to space the detectors, and on proximity to link an individual vehicle to its radiation signature. In contrast to roadways where lanes segregate vehicles, and motion is well controlled by inspection booths; channels, inlets, and rivers present chaotic traffic patterns populated by vessels of all sizes. We have developed a unique solution to this problem based on our portal-less portal monitor instrument that is designed to handle free-flowing traffic on roadways with up to five-traffic lanes. The instrument uses a combination of visible-light and gamma-ray imaging to acquire and link radiation images to individual vehicles. It was recently tested in a maritime setting. In this paper we present the instrument, how it functions, and the results of the recent tests.

## 1. Introduction

In recent years, the fear of nuclear terrorism has grown and considerable research has been performed on how to control the illicit flow of nuclear materials. One route of entry for such materials is via maritime traffic. Original focus has been on the large amount of cargo entering the US on large container vessels. However, of additional concern is the large amount of small-vessel traffic (<300 tons) in US coastal waters. According to the DHS "Small Vessel Security Strategy" [1] there are 13 million registered and ~ 4 million unregistered, US recreational vessels with another 82,000 fishing vessels, and 100,000 other commercial vessels. These craft roam over 95,000 miles of coastline, 10,000 mile of navigable waterways, and have access to 361 ports and presumably many more small marinas. In 2006 there were ~ 70,000 foreign arrivals of such small vessels although the true number is hard to estimate since Customs and Border Patrol relies on self-reporting for most of these arrivals. In short such "Small vessels are, however, readily vulnerable to potential exploitation by terrorists, smugglers of weapons of mass destruction . . ." [2].

Successful control of illicit nuclear materials requires a layered approach comprising many elements. One well recognized tool in such a regime is the use of passive radiation detection to screen for nuclear materials at points of entry into the country. However, when it comes to small vessels their overwhelming numbers coupled with their relative freedom of motion makes even the application of this time-honored approach appear difficult. At terrestrial ports of entry, vehicular traffic is naturally constrained, making the use of passive radiation sensors relatively simple to implement. Portal monitors that comprise one or more large radiation sensors

associated with each traffic lane can scan slowly moving traffic at the border. While there are navigational chokepoints where one could consider scanning maritime traffic, there remains the problem that individual lanes do not exist and that vessels are free to move at will within such channels. In this paper we briefly explore the morphology of such navigational chokepoints and present results obtained with the Roadside Tracker, a unique portal-less radiation portal monitor that was recently tested in a maritime environment.

## 2. Channel Structure

If one makes the optimistic assumption that one can devise a radiation sensor that can uniquely link the radiation signatures from multiple vessels in a navigational choke point to the correct vessel, one can ask what the channel morphologies are that such an instrument must deal with. To address this issue, we look at some sample results obtained from overhead imagery of the San Diego Bay area available from Google Earth™ [3]. Fig. 1 shows the overall morphology of the region together with a few close up views of typical marinas, that can be readily identified in the imagery. The marinas have a common layout in that they are protected from the open waters of the bay by a breakwater that has gaps for allowing vessels to enter and leave the facility. Such gaps represent the type of navigational chokepoint that might be suitable places to mount radiation sensors since all vessels must pass through this gap to reach the marina. Of particular concern is the span of the opening since this determines the maximum range over which a detection must be made.

Basic physics indicates that ranges of order 100 m are reasonable as this represents the approximate mean free path of gamma-radiation through the atmosphere at the energies of interest. In fact, recent terrestrial results with 1-$m^2$ class instruments have successfully demonstrated detection at ~ 100 m for $10^7$ Bq class sources, in reasonable integration times ( ~ 60 s) [4-6]. This also represents a reasonable integration time for a boat at 100 m since it takes 38 s to transit a 100 m field of view when traveling at typical “no-wake” velocities of 5 knots (2.6 m/s). While results will depend on many factors, these parameters are not unreasonable for discussion of a generic system.

Examining other marinas in the bay, we find that entrances range from 100 to 200 m, indicating that most such facilities can be covered with two sensors, one on either side of a gap. However if we look at the bay as a whole we also find that it has its own choke point indicated by the yellow line near the top of the figure. Hence, instead of placing instruments at each marina, we can also consider a system architecture that scans all traffic in this main channel. The problem is that the channel width is 600 m, meaning that if we use a pair of sensors with one on either side of the channel, they must each be effective out to a range of 300 m. This distance is uncomfortably large for a single sensor due to atmospheric attenuation. On the other hand, if a sensor is fixed in the middle of the channel, then the range that must be covered by a single instrument is only 150 m. This is close enough to the notional 100 m limit that one could imagine using just three extra large detectors to span the channel with one in the middle.

In exploring other regions of the country, one finds that such “tree-like” branching morphologies are quite common. Generally, the further down the tree (i.e. the closer to the ocean) one goes, the greater the span that must be accommodated. There are significant benefits in moving sensors closer to the ocean including the fact that one covers more traffic with less detectors, and that

there will generally be greater time to respond to an alarm. This increases the likelihood that illicit materials would be intercepted further from the heart of a city.

An additional channel configuration worth considering is represented by the San Francisco Bay area. This has a very large entrance channel with a width ~ 1.6 km. Although this is clearly too large to span with just a few detectors on the water, we point out that the Golden Gate Bridge provides a convenient gantry for holding detectors 75 m above the water's surface. In fact high bridges, capable of clearing ocean going vessels, are a common feature of other metropolitan areas and should be considered in designing any comprehensive system architecture.

## 3. Roadside Tracker

While the preceding discussion of navigable chokepoints is by no means exhaustive, it does indicate that if one had a sensor that could effectively scan passing vessels at ranges of 100 m, then one could consider the use of passive radiation sensors. That such a system could be effectively fielded has recently been demonstrated in a series of measurements taken with half of the Roadside Tracker portal-less portal monitor that is under development at ORNL. The instrument is designed to scan freely flowing traffic on roadways up to 5-lanes wide with traffic moving at velocities up to 110 km/h. [7]. The instrument uses visible-light images to automatically track vehicles through its field of view. Based on the tracks, the system collects a series of short-exposure 1D gamma-ray images. The gamma-ray images are offset based on the vehicle motion to compensate for motion blur and summed into a composite image that represents the integrated data for the vehicle's transit through the field of view. When used on a roadway, the instrument is housed in two 4.3 m long trailers, that are the mirror images of each other and are designed for deployment on opposite sides of the roadway. The final result is the summed output from the gamma-ray imagers in the two trailers. For the tests in San Diego, only one of the trailers was used (see Fig. 2).

### 3.1 Gamma-ray imager

The gamma-ray detector used to collect the data (see Fig. 3) comprises 27, 4.3 x 4.3 x 40 $cm^3$ bars of CsI(Na) mounted in a vertical orientation behind a 1D rank-23 modified uniformly redundant array coded-aperture mask [8]. Each bar is instrumented with a single photomultiplier tube (PMT) at the top of the bar. This is read out by a custom electronics system described in detail elsewhere [9]. In brief, there is a preamplifier mounted on a custom tube base that sends a high-level signal from each PMT to a chassis containing a readout card for each PMT. The card includes a complete, low-power, nuclear-quality readout with a fast channel for triggering and a slow channel that concludes with a 14-bit analog to digital converter (ADC). When an event has been processed by the card, it sets a logic flag and is subsequently readout by a custom digital board that houses a field programmable gate array (FPGA). The FPGA reads out the events on a first-come-first-serve basis and aggregates them into blocks that include all events falling within a 1-ms time window. These blocks are transferred to the host computer for processing. In addition to the gamma-ray events, the FPGA also records two additional event types that serve to synchronize the gamma-ray and video events. The first of these is the logic pulse that is used as a shutter release for the video images. The video system is run at a frame rate of 20 Hz. A second 1-Hz event type includes the absolute time from a small computer running a real time version of Linux that serves as a network time protocol server (NTP). The absolute time is referenced to a customized global positioning system receiver [10]. This instrument provides the 20-Hz shutter-

release signal that is synchronized to universal coordinated time (UTC) and an additional 1-Hz signal that is used as explained in the description of the video system below.

### 3.2 Video System

The video system comprises three Ethernet-based megapixel video cameras [11]. In operation the cameras sit at the top of a 15-m-tall deployable mast on a 2-m-long Al pipe (see Fig. 2 inset). The cameras at either end of the pipe form a stereo pair that is used to track the vessels through the field of view. The third camera views optical targets mounted on the top of the trailer housing the system. The images from this camera are used to determine the orientation of the stereo pair with respect to the gamma-ray imager. This information is required since the location of the top of the mast changes from deployment to deployment and can also move with respect to the gamma-ray imager on a frame-by-frame basis due to buffeting by the wind and or motion of personnel within the trailer.

The video cameras were selected primarily for their unique capability to include the value of a counter driven by their local Ethernet clock with each image. This value is used to verify that all of the cameras in the system are synchronized to each other and, in combination with the "time-events" recorded by the gamma-ray system to also absolutely synchronize the time between the two systems. The Ethernet clock counter is reset every second with an additional 1-Hz signal to remove drift of the local clocks.

### 3.3 Data Processing

The data for the system is collected and processed using three computers connected via a local area network. There is a Master computer that runs a graphical user interface used to control the full system. It exchanges commands and data using a local area network. The other two computers control the gamma-ray and video systems.

The video computer is used to collect the images from the three video cameras and spools the images to disk as they come in. The images are also analyzed by a software engine using the following steps. First the time of the images is checked to make sure that they all correspond to the same frame number. The images from one of the two stereo cameras are then analyzed to find "good features to track" [12]; points in the image where there are strong gradients in image features in two dimensions. The points are tracked from frame to frame using the Lucas-Kanade algorithm [13]. The point tracking information is used to assist stereo point matching, where each point is linked to the corresponding point in the other stereo camera and a 3D coordinate of the point is found. A list of these points is generated and sent to the master computer.

In the master computer the video points are searched using a coherent motion region clustering approach [14] to find points that track across multiple frames in a consistent fashion. Points segregated by this algorithm are identified with a vehicle and the vehicle is tracked from frame to frame until it has left the system. The points that survive the transit are then used to determine a bounding box that is used to define the location of the vessel in each frame. The time that the front edge of the bounding box crosses from one gamma-ray image pixel (GRIP) to the next is calculated by interpolation of the data from sequential video frames to generate a list of video events that are sent to the gamma-ray computer. The locations of the GRIPs are known from an off-line spatial calibration of the overall system. This calibration links the gamma-ray imager

field of view to the optical targets on the roof of the trailer that are monitored by the third video camera. Based on this information and a detailed calibration of the visible-light system (also performed off line) a full correspondence between the gamma-ray and visible-light imaging systems is known. This calibration is only conducted infrequently and the relationship is reestablished on deployment based on the optical targets, (i.e. no in-situ calibration is required.)

During a run, the third computer continuously collects data from the gamma-ray imager. The data is saved to disk in list mode and is also buffered in local memory in a time-ordered data structure organized by frame number and millisecond within a frame [15]. Upon receipt of the video events, the gamma-ray system creates a series of images that are synchronized to the motion of the vessel (see Fig. 4). Each image is made from all data in the buffer that falls in the time interval between the video events, e.g. all of the data between the times that the front of the vessel transited from one GRIP to the next. The images are generated using the properties of the MURA coded aperture mask [8] mounted in front of the detector array. The portions of each gamma-ray image that include the vessel are appropriately shifted in location to compensate for the vessel motion and summed into a stationary image that represents the composite image of that vehicle.

The system localizes sources to a single vessel by coherently adding the images based on their angular velocity through the field of view. However, this velocity is proportional to range and the system is sensitive enough that the change in angular velocity over range differences as short as 30 cm results in some motion blur. To compensate for this, the code that generates video events actually subdivides each vessel into a series of slices that are ~ 30 cm wide and generates a full set of video events for each slice. Any source that is present will appear strongest at the correct range so that this approach not only improves the system response, it also allows one to localize a source in range as well as along the vessel.

In addition to the composite image, the gamma-ray system also keeps track of the counting statistics and spectral information associated with each of the image pixels. The former is used to determine the significance of the counts in a given composite image pixel by dividing the summed counts in the pixel by the variance associated with that pixel [LAI]. The latter is available for on line viewing by the operators. It provides a background-subtracted spectrum of the counts associated with each pixel for use in identifying the isotope(s) that are present at that location.

### 3.4 Adaptations for Maritime venue

The system was designed for use in terrestrial applications and a number of issues had to be addressed in the maritime setting. These were all related to the visible-light processing and resulted primarily from the differences between the texture and content of the visual data. The first was a difference in the overall exposure of the images. The water surface was significantly darker then the pavement normally encountered so that when this was “correctly exposed” some of the alignment targets on the trailer were overexposed. This was easily fixed by adjusting the exposure parameters.

The second issue had to do with the handling of background data points. In the terrestrial application, a number of these are routinely found but are stationary. The software engine is

designed to ignore such points based on the fact that they do not move. In the maritime application we would routinely detected a large number of points from ripples on the surface of the water and from objects in the water drifting with the current that swept past our location. Since they were moving, the points were processed through the system which would frequently become compute-bound due to the large numbers of such points detected. This issue was solved by adjusting the parameters used to find points and by selecting only those points that were above the surface of the water as determined from the stereo information. We note that wakes from tracked vessels had to be handled in a similar fashion.

Finally, we had to adapt several other aspects of the visual parameters. These included compensating for the fact that the imager was mounted several meters above the surface of the water, increasing the region of the image that was scanned for vessels since vessel traffic is not as constrained as vehicles, and we had to enlarge the scan window used by the CMR region to account for the larger size of "small" vessels when compare vehicle traffic.

**4. System Performance**

The system was deployed on a wharf in the San Diego Bay area (Fig. 2.) Small vessels with different sources on board were piloted past the pier at different velocities and stand-off distances. Sample results demonstrating the performance of the system are shown in Fig. 5 to 7. In the first of these we show the results from a vessel that had both an Am-Be neutron source and a $^{239}$Pu source on board. Both sources were easily detected by the system as seen in the image. The spectra associated with the two peaks are also shown in the figure. They clearly identify the source at the front of the vessel as the Am-Be source. In Fig. 6 we show the results for a vessel with a sphere of highly enriched uranium [16] and in Fig. 7 for a $^{239}$Pu source > 100 g.

**5. Conclusions**

Once the system was suitably adjusted to work in the new environment, it performed well, detecting sources at velocities up to 11 knots and across the full range of distances covered by the visible light cameras. However, it was immediately obvious after watching the vessels pass, that maritime venues include much greater distances than those encountered in a 5-lane roadway (~ 30 m). The discussions in the first parts of the paper combined with the latest results for standoff detection in terrestrial settings indicate that such ranges can be covered in the gamma-ray band with reasonably sized instruments. In fact if one is willing to accept instruments positioned across a channel with separations every 200 or 300 m, then the gamma-ray detection issue is largely solved. The work presented above indicates that the solution is to correct for motion blur of the target vessels using information on their locations as they transit the gamma-ray imager's field of view. We note that for systems that must be mounted in the middle of a channel there may also be the additional complication that the imager itself is likely to undergo wave- and current-induced motions (e.g. if it is mounted on a floating platform.) However our earlier work on standoff detection [15], clearly indicates that one can compensate for detector motion. Hence, correcting for modest imager motion while simultaneously tracking vessels represents only a modest notional departure from the work that has already been performed.

While correcting for the motion blur in radiation images is straightforward if one has the coordinates of the target in the field of view of the radiation imager, it remains a very challenging task to accurately determine those coordinates. To be effective, the visible light

system must be able to operate autonomously and include not just accurate target tracks but also automatically acquire all of the targets as well. As one transitions to more complex target tracks one must also correct for rotations of the target around the imager adding an additional level of complexity to the problem. Further, the greater spans in the maritime setting increase the difficulty of tracking, especially under inclement conditions and with less than ideal lighting. In fact tracking after dark is an even greater challenge. While this can be approached using other parts of the electromagnetic spectrum such as IR or radar, there is the additional complication that one must also be able to identify any vessel that sets off an alarm. While a visible-light image can provide unambiguous cues for response teams to identify a boat, things become more difficult if one must turn to IR cameras or radar to track a vessel when visible light fails. As such, we are rapidly approaching an interesting future, where the ability to use passive radiation detectors for standoff detection depends not just on the radiation sensors, but on careful analysis of data from sensors in other parts of the electromagnetic spectrum.

Nevertheless, the results clearly indicate that this is a viable approach and that careful design of a system architecture that includes the geography of a site and couples it with a properly designed video and gamma-ray imager can result in useful component of a layered defense against smuggled nuclear materials.

**References**

[1] Small Vessel Security Strategy available from www.dhs.gov/xlibrary/assets/small-vessel-security-strategy.pdf
[2] Ibid, p. i.
[3] Google Earth available from
[4] R. Penny1, R. Polichar, W. Hood, K.P. Ziock, S. Labov, K. Nelson, L. Fabris, Nucl. Inst. Meth., This issue.
[5] Shawn R. Tornga, et al., Nucl. Inst. Meth., This issue.
[6] S. Zelakiewicz, R. Hoctor, A. Ivan, W. Ross, E. Nieters, D. Mcdevitt, M. Wittbrodt, B. Milbrath, Nucl. Inst. Meth., This issue.
[7] K.P. Ziock, E.C. Bradley, A. Cheriyadat, M.F. Cunningham, L. Fabris, J.S. Goddard D.E. Hornback, T.P. Karnowski, R.A. Kerekes, Presented at IEEE Nucl. Sci. Symp., Oct. 2009, Orlando, FL.
[8] S. Gottesman, E. Fenimore, App. Opt. 28 (1989) 4344-4352.
[9] K.P. Ziock, J. Collins, M.F. Cunningham, L. Fabris, F. Habte, T. Karnowski, IEEE Trans. Nucl. Sci. 55 (2008) 3643-3653.
[10] Modified model 3203 from PTF Inc., 50-L Audubon Road, Wakefield, MA 01880.
[11] Gig E camera from GEViCAM Inc., 673 S. Milpitas Blvd, Suite 100, Milpitas, CA 95035.
[12] J. Shi, C. Tomasi, Proc. Computer Vision and Pattern Recognition'94 IEEE Comp. Soc. Conf. (1994), 593–600.
[13] B.D. Lucas, T. Kanade, Proc, of the 7th Int. Joint Conf. on Art. Intel. (1981) 674-679.
[14] A.M. Cheriyadat, B.L. Bhaduri, R.J. Radke, Proc. Sixth IEEE Computer Society Workshop on Perceptual Organization in Computer Vision (POCV), in conjunction with IEEE CVPR (2008)
[15] K.P. Ziock, J. Collins, M.F. Cunningham, L. Fabris, T. Gee, J. Goddard, F. Habte, T. Karnowski, IEEE Trans. Nucl. Sci., 55 (2008) 3654-3664.

[16] C. Blessinger, C. Pickett, J. H. Leckey, A. DeMint, J. Gooch, T. Hawk, 50th INMM Conf., 2009, Tucson, AZ.

Fig. 1. Overhead imagery of the San Diego Bay. The entrance to the bay (yellow line) has a width of 600 m. The two inset marinas have an entrance span of 100 m (left) and 180 m (right).

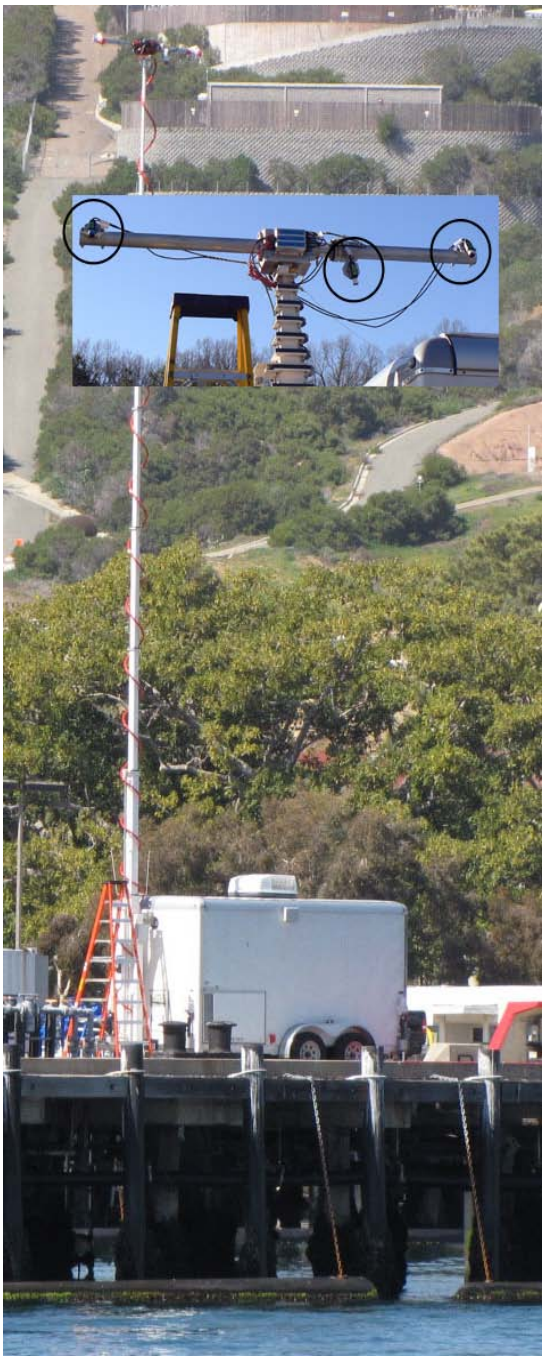

Fig. 2. The instrument in place at a wharf on the side of the San Diego Bay. The trailer houses the gamma-ray imager. The three video cameras are mounted on a 2 m boom (see inset) on the top of a 15 m deployable mast.

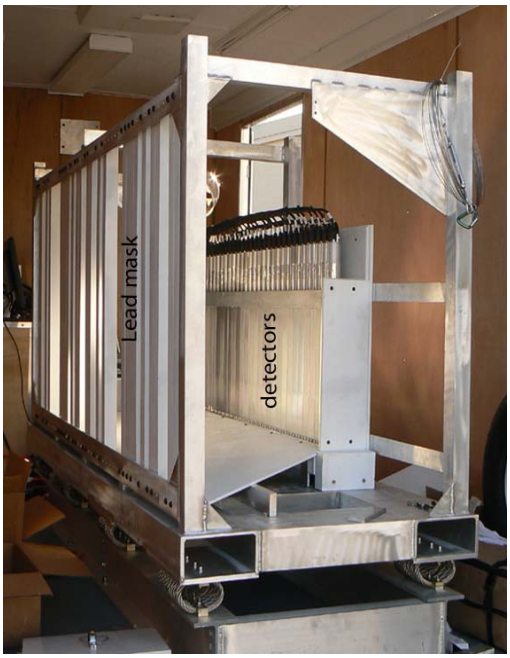

Fig. 3. The coded-aperture gamma-ray imager.

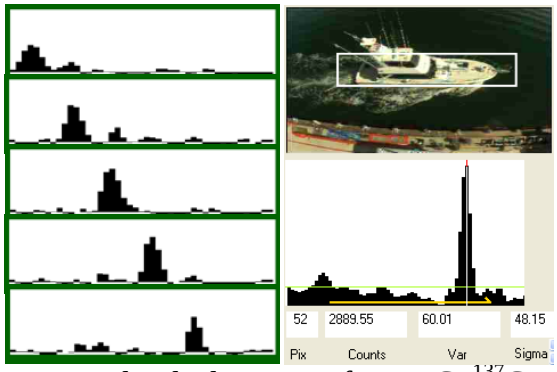

Fig. 4 Individual images of a 5-mCi $^{137}$Cs source passing in front of the gamma-ray imager (left). These are shifted horizontally so that the peaks fall on top of each other and then summed to generate the overall result from the vessel (bottom right). The video image of the targeted vessel is shown with its bounding box shown in white to identify it (top right). Below the composite gamma-ray image is an arrow indicating the length of the vessel. The thin line near the bottom of the image represents the line for zero counts. The pixel at the peak of the image has a statistical significance of 48 sigma.

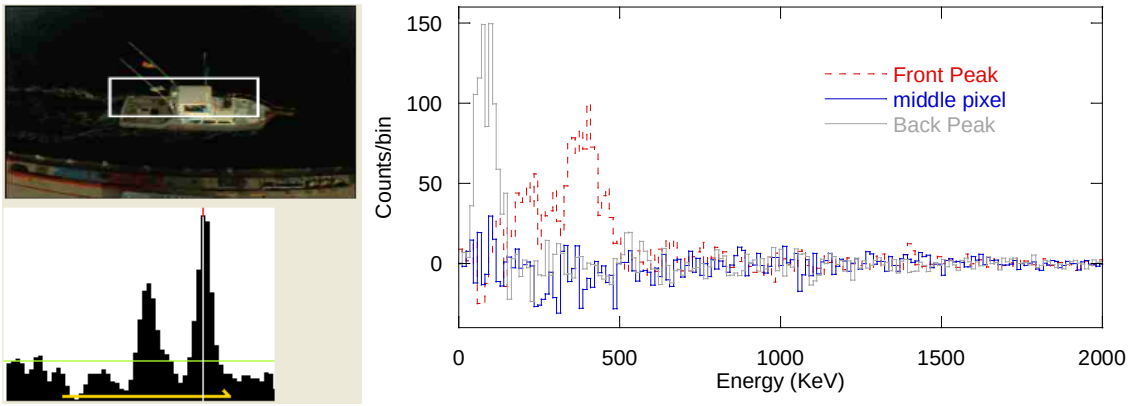


Fig. 5 Results for a vessel with two sources. On the top right is a picture of the vessel, on the bottom right is the 1D gamma-ray image.

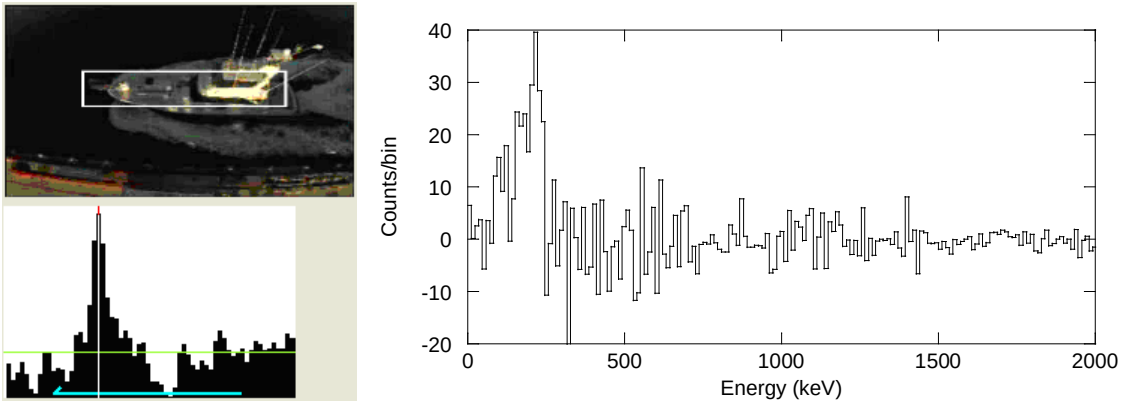


Fig. 6. Results for a vessel with an HEU sphere and the associated spectrum.

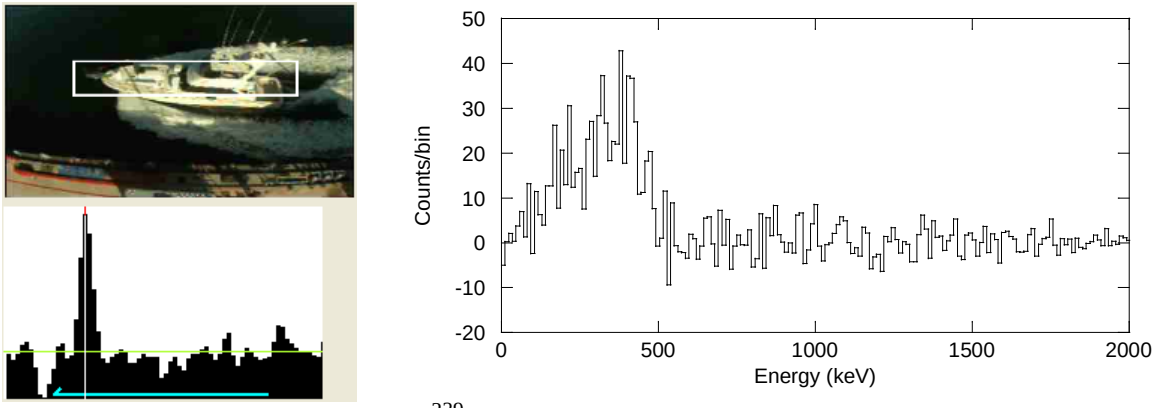


Fig. 7. Results for a vessel with a $^{239}$Pu source and the associated spectrum.