

SANDIA REPORT

SAND2003-3385

Unlimited Release

Printed September 2003

High Resolution 3D Insider Detection and Tracking

Cynthia L. Nelson

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,
a Lockheed Martin Company, for the United States Department of Energy's
National Nuclear Security Administration under Contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.



Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: 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, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent 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, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865)576-8401
Facsimile: (865)576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: <http://www.doe.gov/bridge>

Available to the public from

U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800)553-6847
Facsimile: (703)605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: <http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online>



High Resolution 3D Insider Detection and Tracking

Cynthia L. Nelson
Security Technology
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-0780

Abstract

Vulnerability analysis studies show that one of the worst threats against a facility is that of an active insider during an emergency evacuation. When a criticality or other emergency alarm occurs, employees immediately proceed along evacuation routes to designated areas. Procedures are then implemented to account for all material, classified parts, etc. The 3-Dimensional Video Motion Detection (3DVMD) technology could be used to detect and track possible insider activities during alarm situations, as just described, as well as during normal operating conditions. The 3DVMD technology uses multiple cameras to create 3-dimensional detection volumes or zones. Movement throughout detection zones is tracked and high-level information, such as the number of people and their direction of motion, is extracted. In the described alarm scenario, deviances of evacuation procedures taken by an individual could be immediately detected and relayed to a central alarm station. The insider could be tracked and any protected items removed from the area could be flagged. The 3DVMD technology could also be used to monitor such items as machines that are used to build classified parts. During an alarm, detections could be made if items were removed from the machine. Overall, the use of 3DVMD technology during emergency evacuations would help to prevent the loss of classified items and would speed recovery from emergency situations. Further security could also be added by analyzing tracked behavior (motion) as it corresponds to predicted behavior, e.g., behavior corresponding with the execution of required procedures. This information would be valuable for detecting a possible insider not only during emergency situations, but also during times of normal operation.

High Resolution 3D Insider Detection and Tracking

1 Background

The insider is an individual who has authorized access or the authority to authorize access to information or material. Whether the insider has been co-opted by another person or the malevolence is internally derived, the insider has the potential to create maximum damage for a given effort. This potential is based upon the insider's knowledge, access, and authority. The insider may know operational plans, target locations, and routines associated with day-to-day operations. The insider also has ongoing access and, as a result, has the luxury of choosing the optimum time and place for an attack for maximum effectiveness. The insider has time to probe the system and find vulnerabilities. The insider's authorized access to controlled areas and possibly targets means that the insider can bypass some or even many of the technological countermeasures. As a result, many insider protection strategies focus heavily on human- and procedure-based protection measures. However, weaknesses in human- and procedure-based measures can pose a significant insider risk unless they can be verified or enforced by technological measures. This is particularly true in high-security environments where multiple procedures must be followed.

The project discussed in this report provides a technological approach for ensuring that certain required procedures in a facility are followed. It is a passive, video-based system. However, rather than rely on humans to monitor the video, which would introduce another potential weakness into the system, the human operator is notified when an anomaly in procedures occurs. At that time, the operator can analyze video data and alarm information to make decisions about insider activity.

The insider detection and tracking project is discussed in Section 2; the resultant stand-alone prototype system is described in Section 3.

2 Insider Detection and Tracking Project

A major goal of this project is to enhance 3D sensing technologies to enable an automated system to track and record the actions of many people in a monitored area and then alarm on suspected insider activities. The usual mitigation against the risk of insider activity is having procedures in place. There are instances, however, in which the detection of insider activity by following proper procedures is too late to prevent information or material loss. In fact, vulnerability analysis studies show that one of the worst threats against a facility is that of an active insider during an emergency evacuation. When a criticality or other emergency alarm occurs, employees immediately proceed along evacuation routes to designated areas.

Sandia National Laboratories's 3-Dimensional Video Motion Detection (3DVMD) system provides the capability to continually and automatically monitor all activity within a facility. High-level processing can be performed on the 3-dimensional information provided by the system to detect deviations from the required procedures. This technology is described in the following section. Information extracted and processed from the 3DVMD system is described in subsequent sections. Alarm conditions that are currently implemented using this information are also presented.

2.1 3-Dimensional Video Motion Detection

The 3DVMD system is a stand-alone sensor system that operates on a single PC platform. The sensor uses multiple cameras to monitor activity in predefined, three-dimensional (3D) **zones** or **voxels**. Each voxel is a partition of the total volume of video detection, called the **detection volume**. The detection volume is defined by the intersection of the vision fields from combinations of two or more cameras in the system and it is *only* within this volume that activity will be detected. If any part of the volume under observation is not seen by at least two cameras, then it is not part of the detection volume¹. A 3D voxel is defined by masking the pixels from each camera's field-of-view that intersect a selected sub-volume of the total detection volume. The detection volume may be completely partitioned into voxels or may contain just a few defined voxels. The image shown in Figure 1 is an example of individual voxels defined over a region from a single camera's point-of-view. In this image, each voxel is the same size, and the colors represent different levels of voxels.

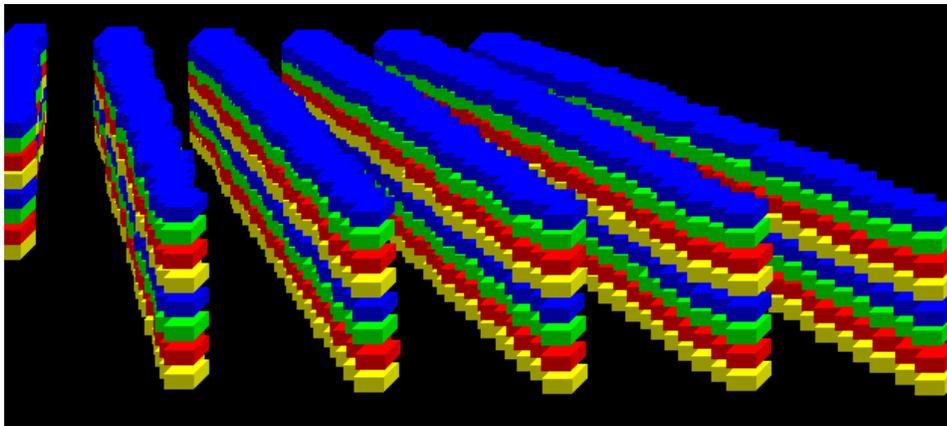


Figure 1. Multiple zones defined within a detection volume

3D motion in the detection volume will only be characterized where voxels are defined. An intruder's movement *between* the voxels would not be detected with 3DVMD. For example, using the voxel mapping shown in Figure 1 and assuming each voxel is 2 ft × 2 ft × 1 ft, a person could avoid detection by creating motion (walking) only between the rows of voxels. Typically, this would not be an issue because the voxel mapping would have to be known by the intruder. However, to avoid this type of situation, a voxel mapping is generally created with a small overlap between rows and columns of voxels such that very little, if any, space in the detection volume remains without being part of at least one voxel. 3D motion for a voxel is declared only if multiple cameras see significant activity in that voxel, causing a **3D alarm**. The size of a voxel is determined by the resolution of the video (number of pixels/foot in the far field of the video detection region) and the size of intruders to be detected. In recent applications, it was found that current PC technology (single 2 GHz processor, 1 Gbyte of RAM) for a 3DVMD system could easily handle large detection volumes (100 meters × 15 meters × 4 meters) of 6-inch voxels or smaller detection volumes (5 meters × 5 meters × 4 meters) of 1-inch voxels. The image shown in Figure 2 provides a 2D view of the 3D motion detected within a detection volume defined by 2-inch voxels.

¹ Viewing by only one camera does not provide the information required to define a "volume."



Figure 2. Voxels from 3D motion detection displayed in 2D image

The 3DVMD technology has been shown to greatly reduce the occurrence of nuisance alarms when compared with traditional 2D video motion detection techniques. As an example, neither the shadow across a reflection on the floor in the image (Figure 1) nor the light reflections resulted in activated voxels. This is particularly important in exterior environments in which moving shadows and reflections can be prevalent.

2.2 3-Dimensional Clustering and Tracking

Information describing the exact 3D location of each activated zone results from the 3DVMD algorithm. If the resolution of the 3D system is high enough, the set of (x,y,z) locations of all activated zones can completely describe the size, exact location, direction of motion, and even certain gestures of a person moving about the monitored room.

A list of all activated voxels is generated by the 3DVMD system. This is read into a 3D clustering algorithm that analyzes the voxels to isolate objects (e.g., people, machines) that may be creating the activity. The 3D clustering filters nuisance voxels that are activated, but are not associated with a valid object. Once objects are clustered, their identity is maintained from frame to frame in the video. As a result, if the objects physically join, such as two people shaking hands, their identities will remain separate. The image displayed in Figure 3 illustrates this condition. This image displays the clustered voxels in 3D space. Two individuals physically touch by handing off an object, as shown by the continuity of the activated voxels across both clusters, but the system keeps the identities separate. When the individuals separate, their identities are still intact. This is extremely important when tracking activity, which is the next step following clustering.

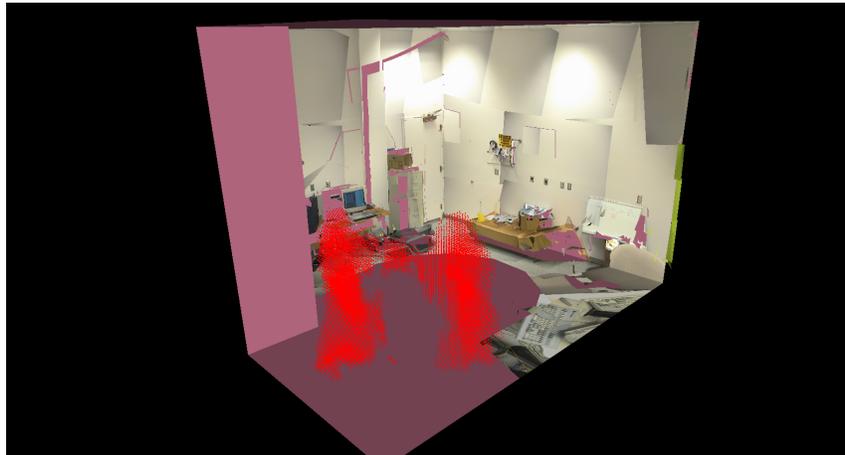


Figure 3. Voxels from two people handing off an object in 3D space

Once an object is clustered, its 3D centroid is located and becomes part of a tracking log for that object. 3D tracking was selected because valid activity could occur at all heights of a facility. An example of this is tracking the movement of barrels of nuclear material by cranes or along ceiling rails. In this case, the object being tracked is actually moving at different levels above the ground.

In order for a track to be initiated, a valid object must be moving for a minimum number of frames. As long as the object is moving and is detected by the 3DVMD algorithms, a track is maintained. If the object stops moving, but is still in the room, the track is maintained. Once the object begins moving again, the track continues. If the track of the object lead to an egress location and the track is then lost, it is assumed that the object left the room. This could, in the future, be tied to another sensor to validate the object's exit (e.g., the breaking of an infrared beam as the object crosses a door threshold). Both a 3D and/or a 2D map of the tracks are available for display and additional processing.

2.3 Detection of Insider Activity Using 3D Information

Many insider-type activities are possible. This project focused on validation of five procedures to help prevent insider activity: (1) enforcement of the two-person rule, (2) process control, (3) protected areas, (4) individual item protection, and (5) evacuation procedures. An alarm condition is generated whenever each of these scenarios is detected. When an alarm condition is generated, information relating to the given alarm is logged for the operator for immediate or future retrieval.

Two-person Rule

Most facilities have implemented several procedure-based protection measures in their highly sensitive areas. One such measure is the implementation of the two-person rule. This can have different levels of enforcement whereby the individuals must be within a certain distance of each other. This procedure can be validated using high-level 3D information. The tracking algorithm indicates the number of people (i.e., objects with specific characteristics) in the area at all times and provides the exact location within the room of those people. Although it is difficult to determine if the people are facing each other, their distance apart and their direction of motion is always known.

Generally, the implementation of the two-person rule procedure requires a second person to enter the room within a specified period of time from when the first person entered. Also, if all but one person has exited the room, that person must exit within a given time period from when the last person exited. These rules are implemented by using track information and system timers. Alarm information is generated if the two-person rule is violated.

Process Control

Process control refers to a procedure in which certain activities must be performed in a specific order. This may also include the requirement that only these activities are to be performed while the room is occupied, and all other activity in the room is to be restricted. The 3D tracking information is used to validate this procedure. A process is selected prior to the room being entered. Once the room is entered, a path that follows each step of the process must be followed. The high-resolution 3D information can be used to determine some details of a process (e.g., if a dial is touched), but cannot determine specific low-level details, e.g., if a dial is turned. If the room is exited before the process is complete or if additional, but unauthorized, activities are performed while the room is occupied, alarm information is generated.

Protected Areas

Although activity is allowed in a monitored room, it may be restricted to certain areas of the room. Restricted areas may exist even though there may be no physical barriers separating them from the rest of the room. Any invasion of the restricted space, either by a person or by throwing an object into the space, can be immediately detected using the high-level 3D and tracking information. This will cause a procedure violation, and alarm information is generated.

Individual Item Protection

Individual items can be secured within the room. An example of a beneficial application is during off-hours or during an emergency evacuation of a machine shop. At these times, sensitive parts may be left on a machine or table. A predefined mask that identifies the location of these items can immediately be loaded into the system. If an item is removed from these masked areas, a violation occurs, and alarm information is generated.

Evacuation Procedures

Evacuation scenarios are typically specific to each facility. The scenario chosen for this project required all individuals to exit a room through a single door within a specified period of time. During this time, no items were to be removed, no protected areas were to be violated, and all motion had to be toward the door. An evacuation alert is provided to the system, either by manual input or a switch closure provided by the facility alarm system. The system immediately begins to monitor for verification of the listed evacuation procedures. If a violation occurs, alarm information is generated.

An additional scenario that can be detected using the high-level 3D information is a situation in which an item is brought into a room and left, or an item is moved from one area of the room to another without authorization. Although these conditions can be detected, alarm information is currently not generated.

High-level 3D information is ideal to monitor the conditions above. Additionally, 3D information may be used to enhance numerous other procedures already implemented to avoid insider activity. Organization 4100 is actively investigating these established procedures for potential 3D application. The current work on this project is structured to allow upgrades for monitoring new situations.

3 Insider Tracking System

A simple user interface exists to run the current insider monitoring system. The main system window is shown in Figure 4. On the left is a viewing window that allows a real-time 3D view of the monitored area or a 2D view from any or all of the cameras. The desired view is selected by pushing the appropriate button located above the window. **CAM-1** through **CAM-4** refer to each of the four camera views, **WORLD** refers to the world or 3D view, and **ALL CAMERAS** refers to all four camera views at one time, each at one-quarter resolution. The view in the figures provides the real-time 2D view from each of the four cameras, one located in each corner of the room.

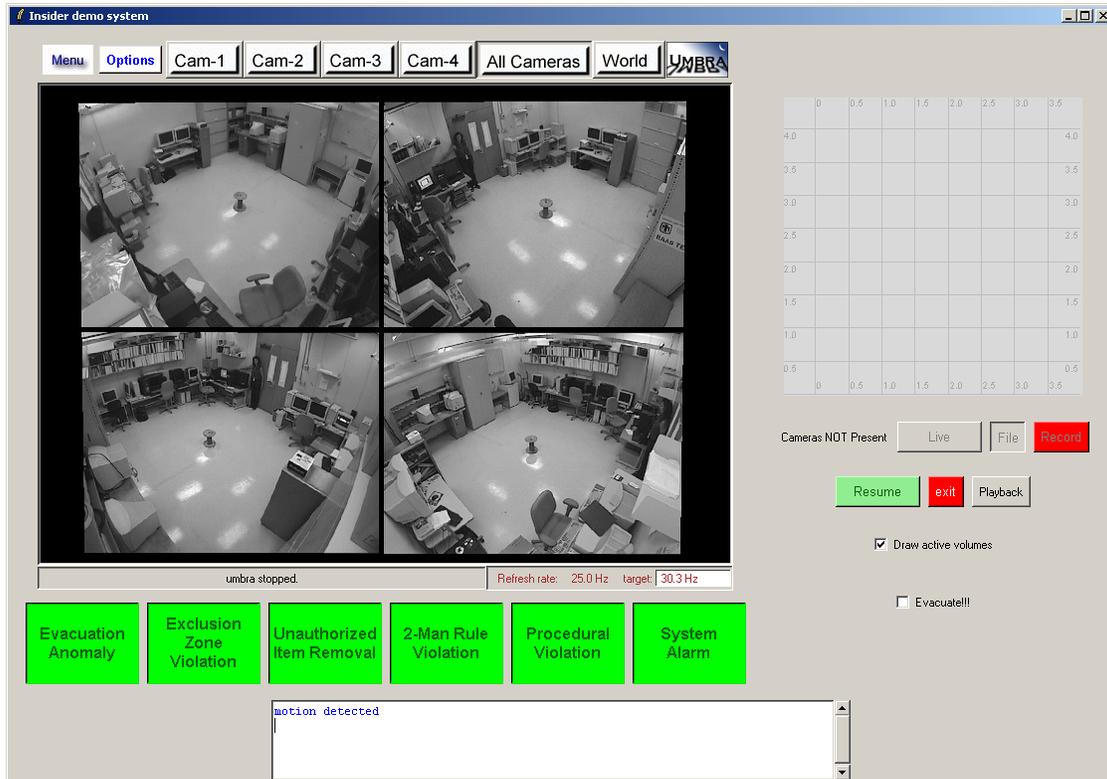


Figure 4. Main window of the current insider detection and tracking system

The 3D activity in the monitored room can be displayed if the **DISPLAY ACTIVE VOXELS** checkbox is checked. In the 2D views, each active voxel is displayed as a cube. In this implementation, all voxel information is shown in white (Figure 5). The active voxels are displayed as points in the **WORLD** or 3D view. In this view, once clusters are created, a box is drawn around each object or person. This can be seen in Figure 6 in which two people are detected. An extension of this display is to show human figures rather than cubes. Efforts in this have begun on a related project and could be integrated with this system.



Figure 5. Active voxels displayed in 2D image views

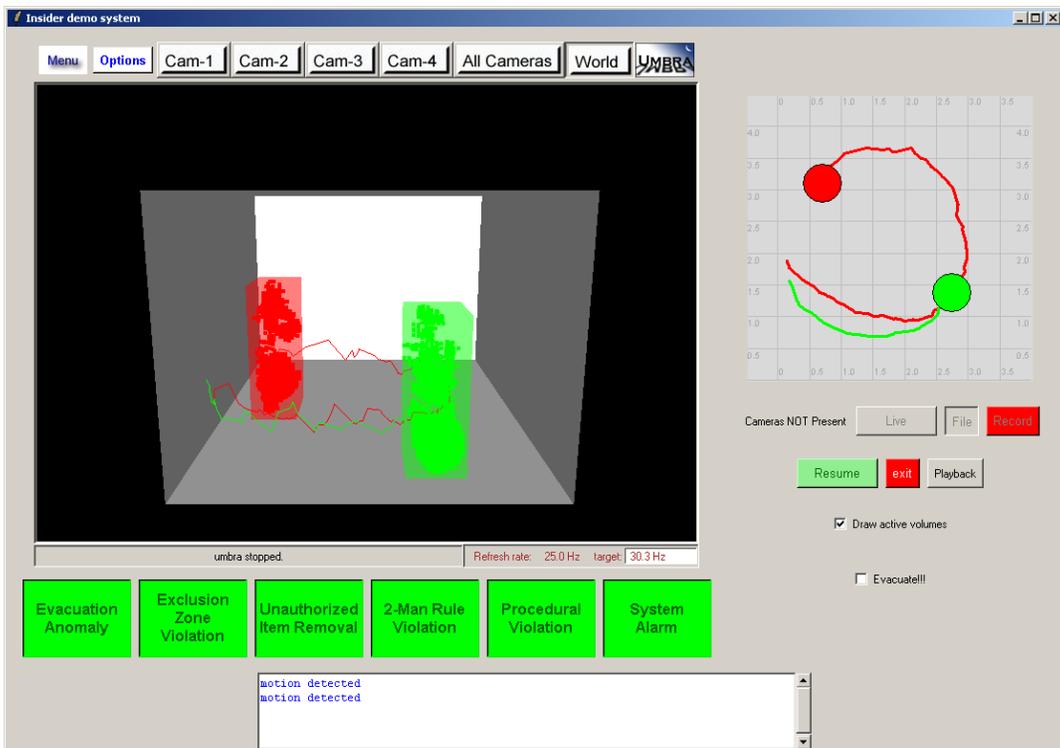


Figure 6. Active voxels display in 3D view with resultant cluster information (no texture mapping in background)

In the real-time 3D view of the active voxels, the 3D tracks can also be observed. Each object and its associated track are associated with a different color. Even when the objects are clustered, they remain separate at all times. Therefore, in the view shown in Figure 6, two objects will always be displayed until one or both are verified as having left the room. The current track information is also displayed in a 2D view. The grid on the right of the user interface associates each 2D object track with a different color and a large circle indicating the current location of the object. Although this information is available at all times, it would be the most useful in a playback scenario in which an alarm condition is being assessed.

An alarm box is associated with each type of alarm violation discussed in Section 2.4. When an alarm condition occurs, the initial green color of the box turns to red. An alarm situation is shown in Figure 7. Here, both a two-person rule violation and an unauthorized item removal alarm are indicated.



Figure 7. Display of alarm situations using 3D information

The capability to assess alarms is available by using a right mouse click on the red alarm box. An assessment is made by entering text that explains the cause and action taken for the given alarm situation. Once an alarm is assessed, the associated box returns to its original green color. A playback option is also provided to allow review of video from the previous few seconds prior to the alarm until the alarm condition occurs.

The user interface can be modified to fit specific facility needs. Currently, an **EVACUATE** button exists for manual indication of an evacuation condition. Eventually, this should also be provided as an input from an external alarm source. Other buttons operate the system using live data, or they can record data and operate the system from the recorded files. An **OPTIONS** button is provided to set up the system (e.g., create masks, define a process order) and to select the

procedures that are to be monitored. The framework for the insider detection and tracking system allows for insertion of additional capabilities. The software is written in C++ with a Tcl/Tk front end and an UMBRA backbone.

4 Summary

Technology was developed to assist in detecting and tracking insider activities within a Department of Energy facility. The basis for the technology is 3-dimensional video motion detection that provides significant high-level information about activity within a monitored room. This high-level information is used to identify certain alarm conditions, such as (1) the violation of a two-person rule, (2) violation of process control procedures, (3) unauthorized access to protected areas, (4) unauthorized item removal, and (5) improper evacuation procedures. The system is designed to accommodate numerous other scenarios that may be specific to a given facility. A data-logging feature records and plays back video in which alarm conditions occurred.

Although the developed technology is video-based, it is **not** a system that requires continual viewing by a human operator. One of the difficulties with detecting insider activity is the reluctance of facilities to use cameras, which could be perceived as a lack of trust in people to do their jobs, i.e., “big brother” syndrome. With the 3DMVD system, however, high-level processing performs all the monitoring activity that otherwise would be accomplished by an operator. It is only when alarm conditions arise that an operator would need to either observe the live video or play back the recorded alarm scenario. This is particularly beneficial in the event of facility alarm conditions that require an evacuation. Although the video would not be observed immediately, it is available for play back when the evacuation condition is terminated. If an alarm condition were raised during the evacuation, a procedure with an increased security level could be implemented immediately, e.g., external fences would be monitored in case an unauthorized item was removed and thrown over the fence for later retrieval.

This technology could be tailored to specific procedures as required by a facility. It is developed to allow the addition of new features and alarm conditions. It could be integrated with existing systems to receive external sensor input and to provide alarm information for external analysis.

Distribution

1	MS	0232	H. R. Westrich, 1011
1		0780	Stephen Ortiz, 4138
14			Cynthia L. Nelson, 4138
1	MS	9018	Central Technical Files, 8945-1
2		0899	Technical Library, 9616