CONF- 9603135--2
SAND96-0686C

Further dissemination only as directed
by HQ DNA/DPSSP, Alexandria, VA
22310-3398, 14 February, 1996, or
higher DoD authority.

SENSOR FUSION FOR INTELLIGENT ALARM ANALYSIS*

March 1995

Cynthia L. Nelson and Deborah S. Fitzgerald
Sandia National Laboratories
Security Technology Department
Albuquerque, NM 87185-0780

RECEIVED
MAR 15 1996
OSTI

Abstract

The purpose of an intelligent alarm analysis system is to provide complete and manageable information to a central alarm station operator by applying alarm processing and fusion techniques to sensor information. This paper discusses the sensor fusion approach taken to perform intelligent alarm analysis for the Advanced Exterior Sensor (AES). The AES is an intrusion detection and assessment system designed for wide-area coverage, quick deployment, low false/nuisance alarm operation, and immediate visual assessment. It combines three sensor technologies (visible, infrared, and millimeter wave radar) collocated on a compact and portable remote sensor module. The remote sensor module rotates at a rate of 1 revolution per second to detect and track motion and provide assessment in a continuous 360° field-of-regard. Sensor fusion techniques are used to correlate and integrate the track data from these three sensors into a single track for operator observation. Additional inputs to the fusion process include environmental data, knowledge of sensor performance under certain weather conditions, sensor priority, and recent operator feedback. A confidence value is assigned to the track as a result of the fusion process. This helps to reduce nuisance alarms and to increase operator confidence in the system while reducing the workload of the operator.

1.0 Introduction

A typical central alarm station for a security environment is an integrated system of people, procedures, and equipment. An alarm communication and display system receives alarm signals from intrusion detection sensors and displays the information to a security operator for action. Although annunciator displays observed by the operator are easy to understand and maintain, they typically contain only a limited amount of information. Therefore, the operator must manually assess each alarm occurrence, which could be difficult, time consuming, and error-prone. The idea behind intelligent alarm analysis (IAA) is to preprocess data from the security sensors and present alarm information to the security operator in a manner that would provide concise and meaningful information and increase confidence in true alarm events while filtering out the necessity to assess nuisance and false alarms. One of the goals of the IAA Project at Sandia National Laboratories is to devise this preprocessing methodology using alarm processing and fusion techniques on available sensor data.

Initial intelligent alarm analysis concepts were developed for the Advanced Exterior Sensor (AES), an intrusion detection and assessment system currently being developed at Sandia. The AES requires the fusion of data

* This work was supported by the Defense Nuclear Agency and the Department of Energy under Contract DE-AC04-94AL85000. Review of this material does not imply Department of Defense endorsement of factual accuracy or opinion.

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

MASTER

from three co-aligned sensors - an infrared, a visible, and a radar sensor. In addition, intelligent alarm analysis utilizes the input from environmental sensors (temperature, ambient light, etc.), knowledge of sensor performance under certain weather conditions, sensor priority, and operator feedback to create a final confidence of current alarm situations to be forwarded to a display panel. The combination (or fusion) of this information could greatly reduce the workload of the operator and provide for a more accurate assessment of alarm situations.

This paper discusses the intelligent alarm analysis process that was initially developed for the AES. Since a major portion of the analysis involves multisensor data fusion, a background on architecture and implementation issues of fusion is given in Section 2.0. A description of the AES fusion problem is then presented in Section 3.0, followed by a solution approach to the fusion problem in Section 4.0. Section 5.0 presents the results obtained from using simulated AES data and discusses future plans for the IAA Project.

2.0 Multisensor Data Fusion

The concept of sensor fusion has become popular with the advent of a number of different types of sensors providing data for simultaneous evaluation. The need for a technique which can transform incomplete, inconsistent, or imprecise data provided by one sensor to more useful information by fusing it with data provided by other sensors is a crucial element to achieve autonomy and efficiency through machine intelligence. **Multisensor data fusion** has been used to provide solutions to problems that are characterized by intensive and diverse sensor information. It can be defined as the process of integrating raw and processed data into some form of meaningful inference that can be used intelligently to improve the performance of the system beyond the level that any one of the components of the system separately could achieve. The product of fusion represents a synthesis of input data appropriate to an individual decision maker's need for meaningful information. Therefore, the implementation of fusion takes many forms and tends to be very problem specific. The output of each fusion process is generally an estimate of an object's (or target's) position, identity, and/or other attributes. For this paper, the fusion process seeks to combine data from sensors to establish the identity of objects, henceforth referred to as **identity fusion**. For our application, identity is simply classified as target or non-target. In this section, architecture and implementation issues for a multisensor data fusion system are discussed. This presents only a synopsis of the fundamental issues that must be considered before designing a data fusion system. Further information on data fusion problems, approaches, and methodologies can be found in [1].

2.1 Architectural Models for Multisensor Data Fusion

A fundamental issue specific to the data fusion process is the choice of when in the processing flow to fuse data. This involves architectural selection of which there are basically three approaches: (a) centralized fusion, (b) decentralized fusion, and (c) hybrid fusion. The three architectures are discussed below.

There are two variations of a **centralized** fusion architecture, one that fuses raw data and one that fuses derived data from multiple sensors. Raw data can be fused if the sensors are either identical or their observational data can be closely compared or merged (e.g., co-aligned visible and infrared images). In such a case, the raw data from each sensor is associated and subsequently fused to create a single data set that can then be processed to estimate the identity of observed entities.

In a second approach to a centralized fusion architecture, preprocessing is applied to each sensor to extract a feature vector. The feature vectors are then associated and subsequently fused. In this approach the features from multiple sensors are concatenated into a single feature vector, which can then be used for identity estimation.

A **decentralized** fusion architecture allows each sensor to perform a maximum amount of preprocessing to generate feature vectors and declarations of identity. That is, the output from each sensor is a decision (i.e., declaration of identity). These decisions are then associated and fused to determine a joint declaration of identity.

A hybrid architecture combines the centralized and decentralized approaches. Raw data, feature level data, and decision level data are all input to a fusion process. This information is then combined to result in fused declarations of identity.

2.2 Implementation Models for Multisensor Data Fusion

As mentioned previously, identity fusion is the focus of this paper. This poses the problem of combining data from multiple sensors to obtain a joint estimate of identity. Ideally, the combined declaration is both more specific and more accurate than declarations from any individual sensor. Identity fusion can occur at the raw data level (prior to feature extraction), at the feature vector level (prior to identity declaration), or at the decision level (after each sensor has made an independent declaration of identity). The choice of when to perform the fusion depends on the types of sensor data available and the types of preprocessing performed by the sensors. These three approaches to fusing identity data are illustrated in Figures 1-3 and are further discussed below.

In **decision level fusion** each sensor performs a feature extraction to obtain an independent declaration of identity. Association is then performed to partition the identity declarations into groups representing observations belonging to the same observed entity. The associated declarations of identity from each sensor are subsequently fused. The decision level fusion architecture is considered one of *decentralized* fusion. (See Figure 1).

In a **feature level fusion** approach each sensor observes an object and feature extraction is performed. The result is a separate feature vector representing the object from each sensor. An association process must then be used to sort feature vectors into meaningful groups. These feature vectors are fused and an identity declaration is made based on the joint feature vector. Feature level fusion uses a *centralized* fusion architecture in which preprocessing is applied to each sensor to extract feature vectors that are subsequently associated and fused. (See Figure 2).

In **data level** (or pixel level, in the case of imagery) fusion data from commensurate sensors are fused directly, with subsequent feature extraction and identity declaration from the fused data. To perform such data level fusion, the sensors must either be identical or commensurate. Association is performed on the raw data to ensure that data being fused relate to the same object. The identification process proceeds identically to the process for a single sensor. This is also a *centralized* fusion architecture, but the raw data is associated and fused rather than feature vectors derived from the raw data. (See Figure 3).

Figure 1. Decision level fusion with a centralized architecture.

Figure 2. Feature level fusion with a centralized architecture.

Figure 3. Data (pixel) level fusion with a centralized architecture.

3.0 Problem Description

One of the goals of intelligent alarm analysis is to help reduce operator workload in safeguards and security applications by processing a variety of inputs such as those from intrusion and environmental sensors, *a priori* knowledge about sensor performance in certain environmental conditions, the type of objects that may pose a threat, and operator feedback. The input is to be combined and possibly condensed to make situation assessment and threat assessment easier for the operator. This paper describes the intelligent alarm analysis concepts applied to the Advanced Exterior Sensor (AES) Project. The next section provides an overview of the AES. Section 3.2 describes the role of multisensor data fusion in the AES Project.

3.1 Advanced Exterior Sensor

The Advanced Exterior Sensor (AES) is a 360-degree scanning, multi-spectral intrusion detection and assessment system that is currently being developed at Sandia National Laboratories. It is a moderate-resolution, true panoramic imaging sensor intended for exterior use for nominal detection of humans out to 500 meters and of vehicles out to 1000 meters. The AES simultaneously uses three sensing technologies (infrared, visible, and millimeter wave radar) along with advanced data processing methods to provide low false-alarm intrusion detection, tracking, and immediate visual assessment. The three sensors are co-located on a portable remote sensor module (shown in Figure 4) that rotates at a rate of one revolution per second to provide detection and assessment in a continuous 360° field-of-regard. Therefore, the images from the infrared and visible detector sets and the radar range data are updated once each second. Figure 5 portrays a flattened view of a 360° image from one revolution of the sensor module. Only a section of the image is displayed at any one time on the console screen. This sensor has been designed for easy use and rapid deployment to cover wide areas beyond or in place of typical perimeters, and tactical applications around fixed or temporary high-value assets. An overview of the AES is provided in [2] and a technical description of the detection process is given in [3].

Figure 4. Advanced Exterior Sensor System Remote Sensor Module

Figure 5. Advanced Exterior Sensor System Image Representation

A robust detection and tracking algorithm was developed to address several design goals of the AES. These goals include

- reducing false alarms caused by background and sensor noise,
- reliably detecting small, slow moving, or erratic targets,
- automatically adapting to various qualities of imagery and environment conditions, and
- reliably detecting targets with low signal-to-clutter ratios and high electrical, spatial, and temporal noise conditions.

There are three major processing pipelines in this algorithm: one for the visible sensor, one for the infrared sensor, and one for the radar sensor. Each pipeline is responsible for processing the image from its associated sensor, extracting necessary features, and creating a set of tracks for that image. A *track* is a set of information or parameters about an object of interest (i.e., target) in the sensor data. These parameters consist of

- location in the image,
- size (total number of pixels),
- height (pixels in the vertical direction),
- width (pixels in the horizontal direction),
- energy (brightness in terms of pixel illumination),
- age (number of image frames from when the target was first detected to its most recent appearance),
- distance traveled (from the image location that the target was first detected to the image location that it has most recently appeared),
- coast time (number of image frames that the target was not detected after its initial detection), and
- error (difference between the predicted location of a target in the next frame and its actual location).

Detection, tracking, and the creation of tracks are performed by a separate processing module for each sensor. Therefore, three sets of tracks (one from each sensor) are created every second. A different processing module is then responsible for obtaining the track data from the three sensors and combining (or fusing) them to create a joint set output for the system operator to observe. A discussion of the sensor fusion module follows.

3.2 Advanced Exterior Sensor Fusion Module

The AES poses a tracker-correlator fusion problem in which tracks from different sensors are obtained and associated, or correlated, into a reduced set of tracks for the operator to observe. As described in the previous section, sequences of images (one image per second) from each of the three AES sensors are individually processed to find purposeful motion that may indicate a target of interest. Parameters of each target are then calculated and stored as a separate track. For each additional image frame, the track information is updated, more tracks may be added, and/or old tracks may be dropped if they've disappeared (coasted) for several frames. Hence, for each image frame there is a list of tracks from each sensor. Each track is essentially an estimated identity declaration for an object seen by the associated sensor (i.e., the object *is* a target as opposed to being a non-target). Identity fusion is then performed to provide a single set of tracks which, ideally, provides a reliable degree of confidence that each track object is indeed a target. The fused track set is referred to as a set of events, where an event is a *condensed* set of information from one or more tracks. Tracks from different sensors may be combined into an event if they are believed to represent the same entity, but tracks from an individual sensor are kept distinct since it is assumed that the detection and tracking processing module will do an appropriate job of separating distinct targets into different tracks.

The resulting list of events is provided for the system operator. Each event in the final set has a confidence associated with it such that the operator can focus on the targets that are above a certain degree of importance. Figure 6 illustrates a simple example of this process. An image for a given time frame is produced by each of the

three sensors. Preprocessing occurs on these images to create a set of tracks for each. The tracks are then fused, a confidence level is assigned to each remaining track, and a final list of events is provided for the operator. Note that in the final set of events, one event (E_1) is created from a visible, infrared, and radar track, one event (E_3) is created from a visible and infrared track, and one event (E_2) is created from only a visible track.

Figure 6. Fusing of Track Information from AES Sensors.

Information in addition to track data is integral to the AES fusion process. This includes environmental data and operator feedback. Both are considered important, particularly when assigning confidence to potential targets. As an example, on an extremely windy day the radar data will produce numerous tracks. However, due to the wind it is inferred that many of the radar tracks are actually false alarms. As a result, the events are assigned low levels of confidence, particularly if they cannot be associated with tracks from the other sensors. As another example, assume the operator is continually presented with events representing targets that tend to move in circles. The operator realizes that these are actually cows grazing within the sensor field of view and, as a result, sends feedback to the fusion system such that the priorities of these erroneous targets (i.e., nuisance alarms) are lowered.

4.0 Problem Approach

As stated in Section 2.1, the choice of architecture is a fundamental issue in developing a data fusion system. The issue revolves around the question of where to combine or fuse the data in the processing flow of two or more sensors. The choice depends on the nature of the sensors involved as well as the nature of the inferences sought. The sensor data for the AES project is preprocessed into tracks whereby each track contains feature information from the sensor data. The feature information is available to determine the confidences of the tracks. However, a decision as to each track's identity has already been made, i.e., it is a target of interest (at this stage, the AES is concerned with targets versus non-targets rather than an exact classification of the targets), and a joint identity with the other sensors is desired in order to determine a final event confidence. Since feature information (at

the track level), decision information (at the event level), and raw environmental data are to be fused, a **hybrid architecture** using **feature/decision level fusion** (Figures 1 and 2) was used for performing intelligent alarm analysis for the AES.

A data fusion system is also dependent on what features are available to fuse. In the AES, the choice is limited since the available features are determined by the information in the tracks (location, size, height, width, energy, age, distance, and coast). These features are reduced to a set of three, *Predictability*, *Reliability*, and *Saliency*, which are believed to provide an accurate representation of the track for the purpose of calculating a confidence in that track. *Predictability* is a function of the distance and direction that the track has traveled and is predicted to travel. *Reliability* is a function of age and how often the track has coasted. *Saliency* is a function of the energy in the track and its size (height, width, and number of pixels). Therefore, each track is assigned a three element feature vector $[P,R,S]$. These features are fused together to produce a *Track Confidence*.

Association between tracks of different sensors is performed using the location and size features. These features are derived from the azimuth and elevation information available from the visible and infrared sensors and from the azimuth and range information available from the radar sensor. The visible and infrared sensors of the AES are adequately co-aligned with the same number of rows (representing elevation) and columns (representing azimuth) per image. Therefore, tracks between the two sensors are combined into a single track if they are within a fixed distance away from each other. For example, in Figure 6 T_{V1} and T_{I1} are combined since they are in close proximity to each other in the field of view, but seen by different sensors. As mentioned earlier, tracks from the same sensor are not associated since the image preprocessing in the AES is assumed to adequately group objects into appropriate tracks.

Although the radar sensor is co-aligned with the other two sensors, the resulting images are very different. Columns still represent the azimuth direction, but the rows in a radar image represent a third dimension, range. Since there is no terrain information available from the visible and infrared sensors, track range cannot be easily determined from their images. Therefore, the combination of radar tracks with other sensor tracks relies mainly on the azimuth parameter, which has a very coarse resolution in the radar images. At the current stage of the AES, two radar pixels in the azimuth direction cover about the same number of columns in one visible or infrared image¹. Therefore, to associate a radar track with a visible or infrared track, the radar track generally has to be in the same half (left or right side) of the image as the visible or infrared track. An example of this association is shown in Figure 7. The radar track is in the left azimuth bin and, hence, is associated with visible tracks in the left half of the visible image. Since terrain information is not available, one radar track could actually be associated with several visible and/or infrared tracks. This results if there are several visible and/or infrared tracks in one half of an image and it is not possible to determine the distance between them. By associating a radar track to one of the other sensor tracks, a range dimension is added to the location information for the associated track. In the figure, radar track T_{R1} is associated with tracks T_{V1} , T_{V2} , and T_{V3} . Therefore, T_{V1} , T_{V2} , and T_{V3} are assigned the range dimension from T_{R1} .

In the final step of the fusion process, a set of events, each with an assigned *Event Confidence*, is created by fusing the *Track Confidence* of associated tracks, the environmental information, and operator feedback. The resulting number of events is the same as the final number of associated tracks. The entire fusion process using the tracks from Figure 6 is illustrated in Figure 8. As shown, the process uses two fusion stages. The first stage fuses the $[P,R,S]$ feature vector from each track into a single parameter called *Track Confidence*. In this stage, the information in a track is condensed and each track remains its own separate entity. In the second stage, the row, column, and size information for each track are used to associate the tracks into events and the *Track Confidences* of the associated tracks are fused to create a final *Event Confidence*.

¹ Dimensions of preliminary AES images are as follows: Visible and infrared images - 480 rows X 512 columns (approximately 3° in azimuth), radar images - 2048 range bins X 6 azimuth bins (approximately 9° in azimuth).

Figure 7. Association of Radar Tracks and Visible Tracks.
It is assumed that each image covers the same 3 degree region.

Figure 8. Fusion Processing in the IAA, based on AES information.

Once the architecture is selected and the features are determined, extracted, and associated, the technique to combine, or fuse, the associated data must be determined. The selection process for techniques requires an overall system perspective. Since a goal of intelligent alarm analysis for the AES is to use environmental information and operator feedback in addition to the intrusion sensor information, a technique is required that would allow for fusing all of this information. It is also desired to instill a degree of uncertainty management in the process since the values of individual features are often unclear. For example, a track with a feature vector of $[0.6, 0.4, 0.7]$ is difficult to relate to a track with a feature vector of $[0.7, 0.3, 0.6]$ in stating that one should have a higher *Track Confidence* than the other. It was decided to incorporate the features into a fuzzy inference system since fuzzy logic is flexible, easy to understand, and is tolerant of imprecise data. An introduction to fuzzy logic can be found in [4]. Fuzzy inference is the actual process of mapping from a given input to an output using fuzzy logic. Figure 9 shows three example rules using the AES Stage 1 fusion inputs and output.

Figure 9. Information Flow in a Fuzzy Inference System.

As shown in the figure, all rules are evaluated in parallel using fuzzy reasoning. Fuzzy reasoning involves determining the degree to which inputs belong to each of the appropriate fuzzy sets via membership functions. The input is always a numerical value and the output is a fuzzy degree of membership (also a numerical value). This process, called fuzzification, amounts mainly to table lookup or function evaluation. As an example, the *Saliency* feature of the AES tracks has a degree of membership in three sets - low, medium, and high - and each set is represented by a Gaussian membership function. The resulting membership functions are shown in Figure 10. Here, a *Saliency* value of 0.5 has a membership value of 0 in the low and high sets and a value of 1 in the medium set.

Figure 10. Membership functions for *Saliency* used during fuzzification.

Each rule is evaluated using the appropriate membership function to fuzzify the feature inputs. A fuzzy operator, such as OR (maximum) or AND (minimum), is applied to the fuzzified inputs for each rule. This gives a value that defines the output membership set for that rule. The output membership set is represented by a portion of the specified output membership function. This process using three example rules is illustrated in Figure 11.

Figure 11. Creation of the Output Fuzzy Set.

In the first rule, S , P , and R are high and, therefore, the high membership function for each parameter (which is Gaussian with a mean of 1, in this example) is used. The parameter values (0.8, 0.5, and 0.6) define a y-axis intersection point on each of the respective membership functions. The OR operator is then applied to these y-axis intersection points. This results in a y-axis intersection of the C output membership function (also a Gaussian with a mean of 1) that is the same as the maximum of the y-axis intersection points of S , P , and R . In the case of the first rule, the maximum is defined by the S parameter. The output value (or fuzzy set) for the rule is the area of the C membership function curve under this intersection point (i.e., the shaded region). This process is followed for each rule. Finally, the union of the output fuzzy sets for each rule is created to produce a joint output set which is defuzzified to produce a numerical value representing the *Track Confidence* of that track. The defuzzification operator used in this case is the centroid of the combined fuzzy sets. A similar process is used to create final *Event Confidence* values.

Fuzzy logic is also used to incorporate information from environmental sensors. For example, two environmental conditions that can impact the AES detection capability are ambient light, which affects the visible sensor, and ambient temperature, which affects the infrared sensor. Following are some examples of rules used to describe the relationships between the AES sensors and these two environmental conditions.

1. IF Visible Track Confidence IS HIGH AND Infrared Track Confidence IS LOW AND Temperature is HOT THEN Confidence IS HIGH
2. IF Visible Track Confidence IS HIGH AND Infrared Track Confidence IS LOW AND Temperature is COLD THEN Confidence IS MEDIUM
3. IF Visible Track Confidence IS LOW AND Infrared Track Confidence IS HIGH AND Light IS POOR THEN Confidence IS HIGH
4. IF Visible Track Confidence IS LOW AND Infrared Track Confidence IS HIGH AND Light IS GOOD THEN Confidence IS MEDIUM

Note that in rules 1 and 2, the visible sensor detects an object with high confidence, however, the infrared imager either does not detect it or detects it with low confidence (low can include zero detection.) In rule 1, the fact that the hot temperature can degrade the accuracy of the infrared imager is taken into account, and the output confidence is still high. In rule 2, the temperature should be ideal conditions for the infrared imager, so the resultant confidence is reduced. Rules 3 and 4 show similar examples of how the ambient light input is used for affecting the confidence of the visible sensor.

5.0 Conclusions

A multisensor data fusion system was developed for intelligent alarm analysis based on the problem description outlined by the Advanced Exterior Sensor project. It is designed as a two stage fusion problem. The first stage is a fuzzy inference system that takes track information (*Predictability*, *Reliability*, and *Saliency*) from each individual sensor as input and produces *Track Confidences* as output. The *Track Confidences* are used as the input to the second stage fuzzy inference system. A data association module is also part of the second stage, using size and location information of the individual tracks as input. The output of the second stage is a list of events with corresponding *Event Confidences*, which are then observed and acted upon by the operator.

Actual data was not available for testing the fusion system. Therefore, a set of simulated tracks was used. Several forms of membership functions were used for both the input and the output of the fuzzy inference systems. However, most testing was performed using a Gaussian distribution for all membership functions.

Simulated data consisted of several tracks from all three sensors. In all cases, appropriate associations were made between tracks from different sensors. This resulted in the creation of the correct number of events for each set of tracks. The success of assigning appropriate *Track Confidences* and *Event Confidences* is a bit more difficult to measure since their values tend to be very subjective. However, comparing them against each other by subjectively choosing a ranking order of the tracks or events, all *Track Confidences* and *Event Confidences* appeared to be appropriately assigned.

When incorporating the environmental data, the resultant event confidences were as expected. Although only simulated data was used, it appears that the addition of the environmental information could be beneficial in actual operations.

There is still a considerable amount of work that could be done in determining the ideal membership functions for both the input parameters and for the output confidences of the stage 1 and stage 2 fuzzy inference systems. However, real track data, rather than simulated data, is required in order to get realistic combinations of the

input parameters. Also, operator feedback will be incorporated into the fusion system. Confidence values could be greatly affected by information from the operator since he/she may be able to make additional inferences about various tracks that cannot be made by the implemented algorithms. This work is expected to proceed once the development of the AES reaches the point at which real data from the co-aligned sensors can be obtained.

REFERENCES

- [1] Hall, David L., Mathematical Techniques in Multisensor Data Fusion, Artech House, 1992.
- [2] Pritchard, Daniel A., "System Overview and Applications of a Panoramic Imaging Perimeter Sensor," Proceedings for the 11th Annual Joint Government-Industry Security Technology Symposium and Exhibition, June 19-22, 1995, Virginia Beach, VA.
- [3] Nichols, Scott A. and R. Brian Naylor, "Reliable Motion Detection of Small Targets in Video with Low Signal-to-Clutter Ratios," IEEE Proceedings of the 1995 International Carnahan Conference on Security Technology, October 18-20, 1995, Sanderstead, Surrey, England.
- [4] Kosko, Bart, Neural Networks and Fuzzy Systems - A Dynamic Systems Approach to Machine Intelligence, Prentice Hall, 1992.

DISCLAIMER

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.