An Information Fusion Framework for Threat Assessment

Justin M. Beaver
Applied Software Engineering Research Group
Oak Ridge National Laboratory
Oak Ridge, TN
beaverjm@ornl.gov

Ryan A. Kerekes
Image Science and Machine Vision Group
Oak Ridge National Laboratory
Oak Ridge, TN
kerekesra@ornl.gov

Jim N. Treadwell
Applied Software Engineering Research Group
Oak Ridge National Laboratory
Oak Ridge, TN
treadwelljn@ornl.gov

Abstract - *Modern enterprises are becoming increasingly sensitive to the potential destructive power of small groups or individuals with malicious intent. In response, significant investments are being made in developing a means to assess the likelihood of certain threats to their enterprises. Threat assessment needs are typically focused in very specific application areas where current processes rely heavily on human analysis to both combine any available data and draw conclusions about the probability of a threat. A generic approach to threat assessment is proposed, including a threat taxonomy and decision-level information fusion framework, that provides a computational means for merging multi-modal data for the purpose of assessing the presence of a threat. The framework is designed for flexibility, and intentionally accounts for the accuracy of each data source, given the environmental conditions, in order to manage the uncertainty associated with any acquired data. The taxonomy and information fusion framework is described, and discussed in the context of real-world applications such as shipping container security and cyber security.*

Keywords: Information fusion, threat assessment, Bayesian belief networks, data analysis, threat signatures.

1 Introduction

As enterprises are forced to account for the threats posed by the tactics of small, subversive groups or individuals, they are becoming more aware of the vulnerabilities in their infrastructures. In the freight industry, for example, the high demand for speed in moving goods internationally inhibits an organization's ability to take the time to thoroughly inspect shipping container contents [1]. As a result, shipping containers present an opportunity for a small group or individual to pose a significant threat. Similarly, the volume of enterprise-level Internet traffic, coupled with the lack of scalable data analysis tools, prevents an organization from thoroughly analyzing its network transactions. The high business demand for Internet availability often forces organizations to expose their critical and sensitive information resources to the threat of unauthorized access.

Threat assessment is a means to quantify the risk associated with process or system vulnerabilities. Often, threat assessment is applied in cases where the vulnerability cannot be adequately addressed due to limitations in resources or the intractability of a thorough analysis. The goal of threat assessment is to provide decision support so that a human operator can reliably identify the presence of a threat and take corrective actions. Threat assessment provides a more thorough understanding of the likelihood of a threat, given the current conditions. It is different from threat prediction where the future threat potential is forecasted. Rather it provides a threat quantification given the current conditions or state of the environment, and makes no claims about the future state of the environment.

In the freight example, it is desirable to quantify the level of confidence that a shipped container's contents pose no threat to various personnel, resources, or customers. By understanding the risk, a human operator may elect to route a smaller, more manageable number of shipping containers to a manual inspection station based on the threat likelihood. In the cyber security example, it is helpful to know the likelihood that your enterprise is under a cyber attack, given the current conditions. A human operator can take the necessary actions to respond to the potential threat, once it has been identified.

This research addresses the need for a technology that can combine information from disparate data sources in order to provide an automated and reliable threat assessment. We propose an information fusion framework that is generic in its structure, but specific in its data. The goal of this framework is to provide a means for combining data for threat assessment in a variety of domains. Information is fused at the decision level, after the data has been acquired and analyzed in order to extract the relevant features. As such, the proposed framework can support many modes of raw data, including textual data, numeric data, and image data;

provided that analysis technologies exist that can highlight those features that are of interest. The goal of this research is to merge the relevant features, revealed during data analysis, to provide an accurate threat assessment.

Section 2 provides a summary of the related work in this field. Section 3 describes in detail the methodology used in creating and applying the information fusion framework. Section 4 outlines the application of this framework to the domains of shipping container security and cyber security. Section 5 concludes the paper.

2 Related Work

Information fusion (IF) is defined as the combination of data from disparate sources to produce an outcome that is superior to any provided by an individual source. A superior outcome typically includes an improvement in accuracy, higher confidence through complementary information, or improved performance in the presence of countermeasures [7].

IF can occur on multiple levels [12]. Sensor-level fusion is the level at which relevant data is extracted from the source signal. Feature-level fusion is the combination of data to produce a composite feature vector that characterizes the object under test. Decision-level fusion is the layer that provides a projection of a future state of the object based on the feature vector provided, and is the information presented to an operator to facilitate a human decision. Related to these different levels, Dasarathy [13] characterized IF in terms of the input/output characteristics of a given fusion function: Data in-Data out, Data in-Feature out, Feature in-Feature out, Feature in-Decision out, and Decision in-Decision out. Thus, an IF architecture is simply the combination of these different types of fusion functions to produce a holistic decision support IF system.

The Joint Directors of Labs (JDL) have developed the most prominent model of information fusion. The JDL fusion model and its revisions [6][10][11] focus on maximizing the automation of fusion. It breaks data fusion into five levels, each of which further refines the data from the acquired state to a form that both adequately represents the entities and their environment and is actionable. Much of the literature surrounding IF focuses on the various levels of the JDL model to create and optimize algorithms that merge sensor data in a complex and dynamic space. Automated target location, identification, and tracking are central themes in this type of fusion.

Situational awareness (SA) is an extension of IF which focuses on incorporating human decision-making in the IF process. Endsley's model of SA [9] defines three levels that include Perception of the various relevant elements in the environment, Comprehension of the patterns that are recognized through analysis or evaluation, and Projection of the likely future states based on the understanding of the current state. The levels proposed in the SA model are analogous to the sensor, feature, and decision levels described in [12]. SA systems are by design semi-automated and allow for a Human in the Loop (HIL) to make decisions.

The application of information fusion and situational awareness to threat assessment is prevalent in the literature. The identification of targets in a battle space [14][16], threat assessment for cyber attacks [17][18], and automobile collision avoidance [19] are a subset of domains where IF and SA have been applied to assess threats. The information fusion approaches applied to these problems are varied, and include technologies such as neural networks, fuzzy logic, Bayesian belief networks, and more traditional statistics. These are the various means by which data can be combined.

There have also been several approaches to generic information fusion frameworks for threat assessment. Generic frameworks are typically both an organization of the threat assessment data and a complementary means to combine the data. In addition, generic frameworks typically address multiple levels in the JDL model. Steinberg [15] proposed a generic model for identifying and predicting threat situations that combined the application of hypotheses, threat prediction, threat assessment and consequence assessment. The model focused primarily on the ontological relationships between entities associated with a perpetrator's capability, opportunity, and intent in order to quantify the threat potential. Benavoli, et al, [20] defined an approach to threat assessment that merges/marginalizes available data using evidential networks. Their approach centered on the use of valuation-based systems as a generic framework for uncertainty management. Although the focus of the work was its application to battlespace situation assessment, the approach could be easily adapted to other domains.

This research builds on these prior approaches to information fusion for threat assessment. We focus on the decision-level (JDL Level 3) of information fusion for threat assessment. The proposed framework centers on a taxonomy of threat data that captures the signatures of a given threat and the observables that reveal the presence of the signatures. The application-specific structure of the taxonomy is represented in the structure of the probabilistic model that assesses the likelihood of a threat. The model accounts for the uncertainty associated with the acquisition of data and the detection of threat observables. To our knowledge, no work has been done to

provide a decision-level fusion model that characterizes threat data or assesses threat likelihood in the same way.

3 Methodology

This section describes the technical approach to applying a decision-level information fusion framework to the problem of threat assessment. In Section 3.1, we first describe our assumptions about the underlying data being fused in the framework. These may also be viewed as the initial conditions that must be satisfied in order to apply the framework. Section 3.2 describes the threat taxonomy that we propose to dissect each threat into smaller, more manageable questions to analyze and answer. Section 3.3 describes the architecture of the information fusion framework and Section 3.4 details the methods used to form the probability network that implements the threat assessment taxonomy.

3.1 Assumptions/Initial Conditions

As the proposed information fusion framework operates at a level above data acquisition and feature extraction, assumptions must be made about the data acquired that is to be fused in terms of both its nature and structure. The following assumptions must hold in order to apply this decision-level fusion framework:

- Raw data has been acquired:
  This approach assumes that the raw data has been acquired and has been translated into an analyzable form (e.g., images, measurements).

- Feature extraction has been performed:
  It is assumed that data-specific processing and analysis has been performed on the raw data to extract the relevant features. This may include shape detection for images, or document similarity for raw text.

- Features can be expressed as discrete values:
  This framework assumes that all relevant features can be sufficiently expressed in either the nominal or ordinal scales.

All of these conditions/assumptions must hold in order to effectively use the proposed approach.

3.2 Threat Taxonomy

In order to assess the likelihood that a given threat exists, an entity-relationship framework must be in place that can adequately characterize the threat. Typically, the approach is to decompose the threat into smaller more easily quantified elements, which, when combined, form a characterization of that threat. The challenge in establishing such a representation is to discriminate those factors that reliably identify the presence of a threat.

There is a temptation to incorporate a variety of factors in threat characterization that are intuitive, but have little relevance in a threat assessment. For example, the intent of a perpetrator seems logical to incorporate in to a threat assessment model. However, intent is a very abstract concept and a difficult variable to reliably assess. We chose to ignore the more abstract attributes of threat assessment. Instead, we focus on characterizing threats in terms of their more tangible properties. That is, we seek concrete observables that are mapped to more tangible signatures, attributes, or properties associated with the threat. By focusing solely on observables, we eliminate the error associated with inferring the states of abstract and subjective variables.

Figure 1 depicts the different entities associated with threat assessment and the relationships between those entities. The proposed taxonomy consists of a high-level threat that is composed of one or more *threat signatures*. A *threat signature* is an attribute or property of the threat that is detectable. A *threat observable* is an embodiment of the threat signature that reveals the presence of the signature. While the terms *signature* and *observable* are more typically associated with the physical characteristics of an object, we extend this language to include non-physical observables as well. In threat assessment, the detection of an observable such as the frequency of a term/phrase in a collection of documents may be as significant as detecting the density of a material. Both the physical and non-physical signatures of threats are treated equally in the threat taxonomy. Its structure does not attempt to convey significance of each threat signature, as that is handled through the Threat Assessment Engine (see Section 3.4).

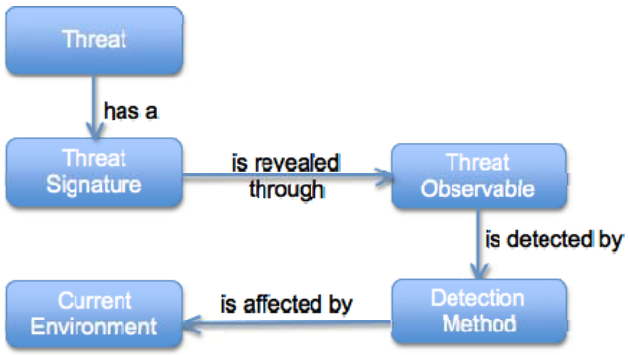

Figure 1. Threat taxonomy entities and relationships.

To illustrate the concept of signatures and observables in threat assessment, consider an example. In cargo container processing, the presence of unauthorized explosive materials is a driving concern. Scanning technologies are commonly applied to identify explosive materials that may be a threat. The “Explosives” threat in

this application can be represented as the collection of properties, or *signatures*, associated with explosives, such as the shape or density of common explosive materials. Each of these *signatures* is further described in terms of the *observable* that distinguishes whether the signature has been detected. For example, in the case of the "Explosives Shape" signature, the *observable* would be whether the image analysis software identified an object within the container that has a shape consistent with that of a commercially available explosives device.

Since this research is focused on the concrete observables associated with a threat, it must consider the reliability of the *detection method* associated with an observable, and how the method's ability to detect the observable is affected by the threat *environment*. Continuing with the example of detecting explosive materials in shipping containers, consider the approaches to detecting explosives material in shipped containers. One detection method for the explosive shape might be a Computed Tomography (CT) scan. However, the accuracy and resolution of CT scans are highly dependent on the size and density of the container. Thus, it is valuable to capture these dependencies in the taxonomy in order to manage the uncertainly associated with detecting observables.

Note the flexibility of this taxonomy. Each threat is comprised of one or more signatures, which are distinct characteristics of the threat that may be observed, with some degree of uncertainty, by a detection method. This structure allows for the easy addition of signatures, observables, or detection methods to strengthen the body of evidence associated with the threat. The characterization of a threat is limited only by the detection methods available for the given signature. It should be noted that the detection of observables is accomplished through an appropriate data analysis or feature extraction method performed during the lower levels of data fusion, and is not addressed in this document.

3.3 Fusion Architecture

This research assumes the availability of multi-modal, multi-source data for reliable threat assessment through the proposed decision-level fusion framework. Any system that incorporates multiple approaches to acquiring and analyzing data must also have an approach for merging the results of those independent analyses in order to present a consistent assessment to an operator. The information fusion system designed for this research assumes the existence of a support framework that includes acquisition of data from documents, measurement and/or scanning devices (Sensor-level fusion), and the analysis of the acquired data item to produce a set of relevant features (Feature-level fusion). Information is fused at a high level and incorporates the various features uniquely produced for each data source as depicted in Figure 2.

This architecture is based on Endsley's model of SA [9] with stages for Perception, Comprehension, and Projection. The focus of this work is the Projection stage – using extracted feature sets to predict the presence of a threat through decision-level information fusion. As data is acquired from the various data sources, knowledge becomes available through feature extraction and is added to the body of evidence for assessing any potential threats.

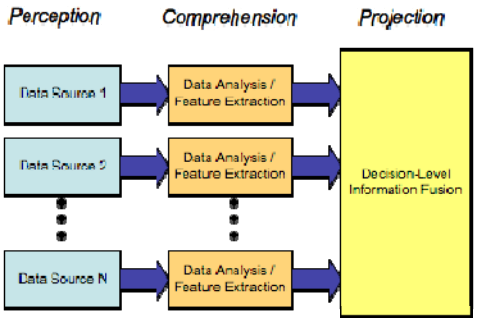

Figure 2. Information fusion architecture.

3.4 Threat Assessment Engine

Determining the probability of a given threat requires representing the threat signature and observable taxonomy described in Section 3.2 in a mathematical framework such that the detected observables affect the likelihood that the threat exists. We call this framework the threat assessment engine. In addition, any probability model for threat assessment must tolerate data that is uncertain or unavailable. For example, a practical scenario in cyber security threat assessment is one where the server housing one of the system's intrusion detection tools has ceased operations. In such a situation, it is undesirable for the threat assessment system to simply cease until the server is back online; rather, it should carry out analysis with minimal loss in accuracy despite the now unavailable data source.

Bayesian Belief Networks were selected as the mechanism for the threat assessment engine and implementing the signature/observable taxonomy. A Bayesian Belief Network (BBN) is a network of nodes connected by directed arcs. Each node in the network represents a random variable in the model, and each arc signifies a cause-effect relationship between the variables. Thus, there may be several arcs leading to or from any given node, but there can be no cyclic relationships. The probability function associated with each node is the joint probability distribution of inputs to outputs. BBN node values are represented as discrete variables, and so can accommodate both subjective and objective data. They

can adapt to an environment as data is processed, can infer unknown model elements based on known model elements, and perform well in the presence of uncertain or unavailable data [5].

The design of the BBN for threat decision-level information fusion, shown in Figure 3 and Figure 4, reflects the hierarchy of the threat taxonomy. Figure 3 depicts the design for determining the presence of each individual threat component. For each threat signature, one or many detection methods may be available to perceive the observables associated with the signature. Each detection method performs more reliably under some conditions and less reliably under other conditions. The BBN captures this uncertainty by accounting for the characteristics of the environment that affect the accuracy of a given detection method. The detection method's accuracy is modeled as an effect of the Accuracy Affecters in the environment, based on what is described in the taxonomy.

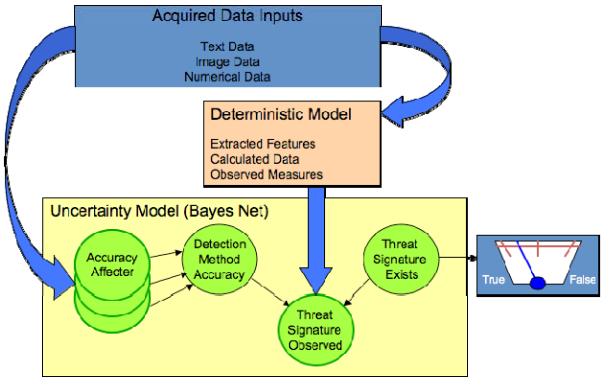

Figure 3. Threat signature Bayesian network design.

Acquired data, in the form of raw text, images, and numerical measurements, are processed using state-of-the-practice techniques to extract features and entities, and to make the determination of whether a threat signature was observed. There is no uncertainty associated with the state of this node – the lower level fusion algorithms in feature extraction either observed a signature or not. Thus, the "Threat Signature Observed" node has two states (*True*, *False*) and is the means for incorporating the results of the deterministic, lower-level fusion models into the uncertainty model.

The uncertainty model for each threat signature accounts for what was observed, and the accuracy of the detection method to compute the probability that the threat signature actually exists. Once values for the "Detection Method Accuracy" and "Threat Signature Observed" nodes are established, the structure of the network allows for the propagation of belief to the "Threat Signature Exists" node, which is also a two-state node (*True*, *False*). The output of the threat signature model is a level of belief, or probability, that the given signature actually exists based on detection result and the environmental factors. This framework is applied to each threat signature in the taxonomy described in Section 3.2.

Once the presence of each threat signature is probabilistically determined, the second tier of the BBN (see Figure 4) combines each signature to determine the likelihood that the overall threat exists. As with each threat signature, the "Threat Exists" node is a two-state node (*True*, *False*) that reflects whether the system believes the threat exists based on the fusion of all of the information. In addition to a *True*/*False* value, the BBN provides a probability, or degree of confidence, that the assessment is accurate. The probability is useful in conveying to an operator the level of confidence that a particular threat exists.

Prior probabilities are essential in the operation of a BBN. That is, in order for the probability network to produce a reliable assessment, it must have historical data that is representative. This system is designed to rely on both expert input and the shadowing of real-world operations to accumulate a significant historical data set.

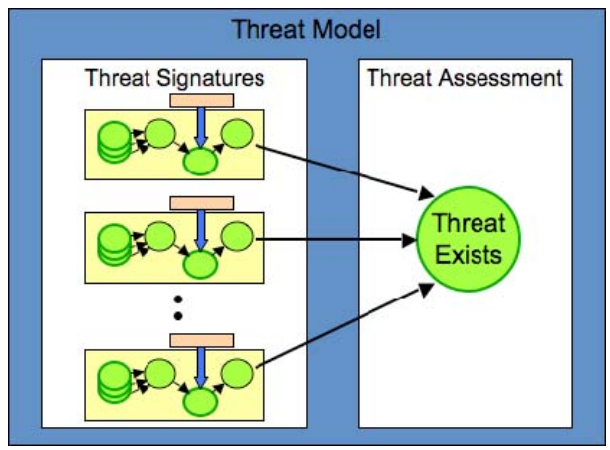

Figure 4 Threat assessment Bayesian network design.

The use of threat signatures in our design provides flexibility in combining data elements. Consistent with the threat taxonomy, threat signatures are easily added to or removed from the BBN depending on the types of data sources that are available and the types of discrete tests performed on each extracted data item. Similarly, dependencies between threat signatures with respect to data source accuracy are adequately captured in both the taxonomy and the model. However, any temporal dependencies between elements are not reflected in the framework. That is, this approach makes no claims to predict a future threat, but is focused on assessing the likelihood that a threat exists given the current state of all known data.

4 Applications

This section outlines our work in applying the framework described in this paper to real-world information fusion problems. Section 4.1 explains our work in the domain of threat assessment for shipped cargo container security. This fusion framework is applied to automate the fusion of text, image, and numerical measurement data in order to provide reliable decision support in assessing whether a threat is concealed in a cargo container. Section 4.2 discusses the application of this taxonomy and framework to the domain of cyber security, in order to minimize the number of intrusion detection false alarms associated with identified alerts and attacks. We analyze the ways in which our information fusion approach can be applied to reliably assess the threat an attack or alert poses to a cyber defense.

4.1 Shipped Container Threat Assessment

We applied the information fusion framework to the problem of threat assessment for shipped containers. We have built a working prototype to demonstrate how data can be fused and presented to an operator in a notional container processing system. The notional container processing system is modeled after real-world systems, and is comprised of multiple testing stations. Each testing station provides some level of data acquisition, whether it be an automated download of a shipping manifest, or equipment to produce an image of the container's contents. The prototype includes an operator interface that provides a view of the fused results, and a back-end simulation that represents the environment and systems used for container processing. The software is written in the Java programming language and leverages the Netica libraries [2] for BBN implementation and the Visualization Toolkit (VTK) for image and volume visualization [3].

The development of the threat taxonomy was dependent on the data available for each shipping container. Table 1 describes a subset of the data used in this research, and the source from which the data was acquired. These data are representative of the types of data available for shipping container information fusion, as specified in [8]. Table 2 contains examples of how these data were translated into signatures and observables. The prototype implements the information fusion framework that evaluates the threat probability, via the described BBN, given these signatures and observables.

Table 1. Example shipping container data.

Data Item Description	Data Source
Shipper Information: The name and address of the shipper.	Shipping Manifest
Destination Information: The name and address of the consignor.	Shipping Manifest
Commodity: A classification of the nature of the goods being shipped.	Shipping Manifest
Shipped Weight: The recorded weight of the container.	Shipping Manifest
Measured Weight: The weight as measured during processing.	Scales
Measured Radiation: Radiation levels detected during processing.	Radiation Portal
2-D Scanned Images: Two-dimensional images using X-ray or similar technology.	Scanning Station
3-D Scanned Image: Three-dimensional image using computed tomography (CT) or similar technology.	Scanning Station

In the shipping container threat assessment application, there are three major threats being assessed: Explosive, Nuclear, and Radiological. For each of these threats, we defined a set of signatures and observables that are unique to that threat. In keeping with the information fusion framework taxonomy, we also identified detection methods and accuracy affecters for each of these observables. The BBN derived from this taxonomy was instrumented with prior probabilities based on expert analysis, and was implemented with an operator interface that gives a user insight into the computed threat assessment.

Table 2. Example signatures and observables for shipped container threats.

Signature	Observable
Object Shape	Shape of an object in the container is identified through image processing as consistent with the shape of a threat (e.g., explosive material).
Object Density	Density of an object in the container is identified through image processing as consistent with the density of a threat.
Container Weight	Measured weight is consistent with recorded weight.
	Measured weight is within the expected weight distribution for that commodity.
Known Shipper	Shipper is part of the Known Shipper program.
Radiation Level	Detected radiation counts are consistent with radioactive material.

Figure 5 shows the developed operator interface. The focus of the interface is the visualization that combines multiple imaging modalities in order to present all relevant imaging data simultaneously. Anomalies and threat assessments are elevated to the operator's attention through graphics.

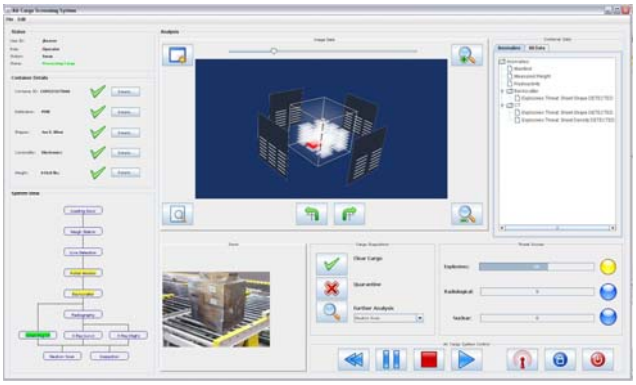

Figure 5. Container threat assessment operator interface.

The threat score is the probability of a given threat being present in the container, and is communicated through both a meter and a colored icon located in the southeast area of Figure 5. The threat score icon's color maps to a 0.2 interval in the threat probability range, and is based on the Department of Homeland Security's Color-coded Threat Level System [4]. This provides the operator with an immediate visual cue of the potential for the threat to exist. In addition, the results of detecting observables associated with container signatures are visually presented to the operator in the form of green checkmark or red 'X' icons in the Container Details panel. The icons reflect whether a signature, or collection of signatures was found. We envision this framework to ultimately be implemented as part of a decision support tool for use in the freight industry.

4.2 Cyber Security Threat Assessment

Enterprise cyber defense systems are typically overwhelmed with the volumes of data that must be processed and analyzed to attempt to detect unauthorized intrusions. Several off-the-shelf intrusion detection tools alert an operator to suspicious situations, but are very unreliable in discriminating actual attacks from suspicious activities [21]. Analysts are overwhelmed with the number of false alarms and are incapacitated to effectively respond to an attack.

The intrusion detection domain provides another opportunity to apply the decision-level information fusion framework. Similar to container processing, it is a domain where data is acquired from multiple sources and in multiple modes, and must be merged to provide decision support to an operator. The goal in this application is to provide a reliable assessment of the threat of an attack based on the current conditions.

The taxonomy described in Section 3.2 is easily applied to the intrusion detection domain. The data sources are the raw data feeds being analyzed including network packets, operating system events, and other logged events in the enterprise network architecture. Table 3 contains a small subset of examples of individual intrusion detection data items and their data sources, and Table 4 provides examples of signatures and observables that may be indicators of an intrusion threat. The detection methods are the niche intrusion detection tools. Each tool provides insight into a dimension of intrusion detection space, e.g., on a host, on the network. However, the accuracy of these tools in identifying an actual attack is unreliable based on the environmental conditions (data rates, rule set version, etc).

Table 3. Example intrusion detection data.

Data Item Description	Data Source
Detected host root access: A user has gained root access to the system.	Host event log
Detected host access: A user has gained access to the system.	Host event log
Detected Malware Probe: An external entity is probing for planted malware.	Intrusion detection log
Detected Service Scan: An external entity is probing all host's ports.	Intrusion detection log
Vulnerability Profile: An itemization of the services available on a host.	IT analyst

Table 4. Example signatures and observables for an intrusion detection threat.

Signature	Observable
Focused Probing	Event logs reveal several probing events focused on a single target
Distributed Probing	Event logs reveal several probing events sent from a single source.
Focused Access Attempts	Event logs reveal several attempts to access a host machine.
Distributed Access Attempts	Event logs reveal several attempts to access host machines from a single source.

The threats associated with intrusion detection are represented in terms of different attack categories. Each of these categories has signatures that are indicators of that specific attack type. The observables are those alerts that are detected through the niche intrusion detection tools, and reveal the presence of the signature. When the information fusion framework is applied, the probability of each threat signature actually existing is propagated to determine the likelihood that the overall threat, or cyber attack, actually exists.

5 Conclusion

This research details an approach to the fusion of disparate information to produce a probabilistic assessment of the presence of a threat. We have proposed a taxonomy for organizing threats in terms of signatures and observables. We have also proposed an associated information fusion framework that reflects the taxonomy, and provides a generic structure that easily accommodates

the flexibility required in real-word applications. The fusion structure also manages the uncertainty associated with acquired data by accounting for the environmental factors that affect observable detection method accuracy.

Information fusion was achieved by leveraging Bayesian Belief Networks for probabilistic threat assessment. The design of the Bayesian network allows for flexibility in terms of the taxonomy of threats. That is, the approach easily accommodates multiple instances of threats and their signatures, and can be applied to several domains. Specifically, we described the application of the framework to both shipping container and cyber security threat assessment.

Our future work for this research includes the validation of the proposed decision-level information fusion model and architecture for the applied domains. In addition, we expect to extend the current approach to incorporate temporal relationships between entities and in the resultant BBN implementation.

6 Acknowledgements

Prepared by Oak Ridge National Laboratory, P.O. Box 2008, Oak Ridge, Tennessee 37831-6285; managed by UT-Battelle, LLC, for the U.S. Department of Energy under contract DE-AC05-00OR2225. This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 for the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.

References

[1] G. Schneider. "Terror Risk Cited for Cargo Carried on Passenger Jets; 2 Reports List Security Gaps." *The Washington Post*, June 10, 2002.

[2] Norsys Software Corporation. "Netica (online)." In http://www.norysys.com, Vancouver, Canada, 2008. Norsys Software Corporation.

[3] Kitware, Inc. *The VTK User's Guide.* Kitware Inc., USA.

[4] Department of Homeland Security. "Homeland Security Advisory System (online)." http://www.dhs.gov/xinfoshare/programs/Copy_of_press_release_0046.shtm, Washington, D.C., 2008. United States of America Department of Homeland Security.

[5] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach.* Prentice Hall, Upper Saddle Ridge, NJ, 2nd edition, 2003.

[6] Data Fusion Subpanel of the Joint Directors of Laboratories. "Data fusion lexicon." In Technical Panel for C3, 1991. United States of America Department of Defense.

[7] E.P. Blasch and S. Plano. "JDL Level 5 fusion model: user refinement issues and applications in group tracking." SPIE Vol. 4729, *Aerosense*, 2002, pp. 270-279.

[8] Stanford University Study Group, Center for International Security and Cooperation. "Detecting Nuclear Material in International Container Shipping: Criteria for Secure Systems." *Journal of Physical Security*, Vol. 1, No. 1, 2004.

[9] M. Endsley. "Toward a Theory of Situational Awareness in Dynamic Systems." *Human Factors Journal*, Vol. 37, pp. 32-64, 1996.

[10] A.N. Steinberg, C. Bowman, and F. White. "Revisions to the JDL Data Fusion Model." *NATO/IRIS Conference*, October 1998.

[11] J. Llinas, C. Bowman, G. Revora, A. Steinberg, E. Waltz, and F. White. "Revisions and extensions the JDL Data Fusion Model II." In *Proceedings of The 7th International Conference on Information Fusion*, pp. 1218-1230, 2004.

[12] L.A. Klein. *Sensor and Data Fusion: A Tool for Information Assessment and Decision Making.* SPIE Press, Bellingham, Washington, USA, 2004.

[13] B.V. Dasarathy, *Decision Fusion*, IEEE Computer Society Press, 1994.

[14] P. Valin, E. Bosse, and A. Jouan. "Airborne application of information fusion algorithms to classification.", Technical Report TR 2004-282, Defense Research and Development Canada – Valcartier, May 2006.

[15] A.N. Steinberg. "An Approach to Threat Assessment." In Proceedings of the 7th International Conference on Information Fusion. Philadelphia, PA. July 25-29, 2005.

[16] Y. Liang. "An Approximate Reasoning Model for Situation and Threat Assessment." In Proceedings of the 4th International Conference on Fuzzy Systems and Knowledge Discovery. 2007.

[17] S.J. Yang, J. Holsopple, and M. Sudit. "Evaluating Threat Assessment for Multi-Stage Cyber Attacks." In Proceedings of the 2006 Military Communications Conference. Washington, DC. Oct. 23-25, 2006.

[18] R. Chinchani, A. Iyer, H.Q. Ngo, and S. Upadhyaya. "Towards a theory of insider threat assessment." In Proceedings of the 2005 International Conference on Dependable Systems and Networks, 2005.

[19] L. Yang, J.H. Yang, E. Feron, and V. Kulkarni. "Development of a performance-based approach for a rear-end collision warning and avoidance system for automobiles." In Proceedings of the IEEE Intelligent Vehicles Symposium, 2003.

[20] A. Benavoli, B. Ristic, A. Farina, M. Oxenham, and L. Chisci. "An Approach to Threat Assessment Based on Evidential Networks." In Proceedings of the 10th International Conference on Information Fusion. July 9-12, 2007.

[21] S.T. Brugger and J. Chow. "An Assessment of the DARPA IDS Evaluation Dataset Using Snort." UC Davis Technical Report CSE-2007-1. Davis, CA. January, 2007.