SAND2020-6006C

Information Theoretic Metric of Verification System Performance

Jason Reinhardt, Michael Hamel, Ben Bonin, Eva Uribe

Sandia National Laboratories[1]

Abstract

There remain several challenges to verifying potential future nuclear arms control agreements. These include, among others, protecting sensitive information during verification activities, confirming the authenticity of an accountable item, and confirming the disposition of materials. Technology-specific approaches offer potential solutions to these challenges but understanding and comparing the performance of disparate technologies in the context of a verification regime requires metrics and assessment methods that are technology and regime agnostic.

This work proposes an objective metric that can quantify confidence and be used as a comparison tool for different information barriers. For a hypothetical arms-control scenario, we used varying information barriers that only allow certain information taken from a gamma-ray spectrum to be "seen" by the monitor. A Bayesian belief network was created that included the monitor's prior and updates their belief as more information became available through measurement. We use the Kullback-Leibler Divergence as a measure of the information gained by the monitor about a protected fact, given the observation of the verification measurement. The result of this analysis is an objective metric of confidence for the monitoring party, and for information protection for the host. This paper will describe the theory and framework used for this analysis.

Introduction

Nuclear arms control agreements often place limits on activities, materials, or arsenals in an effort to enhance strategic stability for all participants. Many agreements employ a verification regime to enforce the limitations and obligations spelled out in the treaty language. These regimes can often involve cooperative verification activities that build a corpus of evidence that a monitoring nation uses in its decision to verify that the host nation is in compliance with the treaty stipulations. For example, the Strategic Arms Reduction Treaty (START) and its follow-on, New START, specified verification regimes that included on-site inspections to ensure that treaty limits were being observed (See Office of the Secretary of Defense 2020, and Department of State 2016). While these regimes utilized radiation detection equipment to verify non-nuclear status of objects, the verification protocols did not include measurements of actual nuclear warheads.

However, it is believed that future treaties may require that measurements be performed to verify that declared treaty accountable items are indeed warheads (See, for example Fuller 2010). If a future treaty implements a verification regime that requires measurements to be performed by a monitoring nation on a host nation's nuclear weapon or weapons usable nuclear material, protecting the sensitive information about weapon designs becomes a central issue. One approach to solving this problem is to create mechanisms that rely on protocols and technologies to protect against the transmission of sensitive information during the measurement, while also providing the monitoring party confidence that the host is meeting their agreement obligations. Such protocols or technologies designed for this purpose can include an *information barrier*, which is meant to only provide the monitoring party with agreed upon information while preventing the release of any sensitive information that may be collected during the verification process. Examples of proposed system that use an information barrier include the Trusted

[1] Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.

Radiation Identification System (Seager et al. 2001) and the Trusted Radiation Attribute Demonstration System (Mitchell and Tolk 2000).

In general, there has been little examination of the fundamental theoretical principles that underpin the concepts of information barriers. An information barrier must be designed such that the monitoring party can gain confidence that something is true, without revealing other information that may be deemed sensitive. A complication in this paradigm, is that information collected to prove a fact deemed non-sensitive, may be related to another fact that is deemed sensitive. This paper proposes an approach that may provide a pathway towards quantifying the performance of information barriers that is based in the traditions of information theory and probability theory. We propose using these concepts because at a fundamental level, information barriers and detection systems describe communication channels that manage how information is transmitted from one party to another.

This paper proposes a technology agnostic metric for characterizing information barrier performance within a verification regime. An accompanying paper lays out a broader vision and justification for cooperative international development of such metrics and how doing so can promote arms control verification goals (Bonin et al. 2020). The next section establishes some background on the information barriers concept and lays out some fundamental definitions and concepts from information theory. From there, key concepts and insights are developed through a series of examples. A metric for understanding the effectiveness of information barriers is proposed and developed. A method is then proposed for using the metric in two key tasks: (1) to understand what can reasonably be expected from an information barrier technology, and (2) to understand the necessary performance parameters for detection systems to achieve a specified information barrier performance level. Finally, conclusions are drawn and briefly discussed.

Information Barriers

While the concept of information barriers dates back to the 1980s, the term *information barriers* is reported to have been born out of the safeguards community in the late 1990s (Close, MacArthur, and Nicholas 2001, National Academy of Science 2005). Information barriers themselves have been defined as “a suite of hardware and software components and procedures that separate a classified data collection system from an unclassified display and user interface” (MacArthur et al. 1998). The requirements for an information barrier are two-fold:

1. Protect the host, or inspected party, by guaranteeing that sensitive data cannot be transmitted to any other party through a measurement.

2. Provide the monitoring party that the measurement validates a claim by the host, using authentic and accurate data.

The approach to information barrier systems that meet these goals has been design-centric, focused on developing architectures, protocols, and technologies that adhere to a set of heuristic guidelines. Design heuristics include defense-in-depth, minimizing classified side components, physical protection of systems, prevention of unnecessary emissions (such as electromagnetic signatures), and simplicity in design and function (MacArthur et al. 1998). Several proposed or demonstrated examples of such systems exist (Zuboski et al. 1999, Glaser et al. 2014). However, there is no known treatment of the fundamental principles and theory involved in information barriers to guide development past heuristic design approaches.

Probabilistic Models and Information Theory

Previous work has sought to use probabilistic and quantitative assessments of uncertainty in the analysis of arms control approaches and concepts. These have included the use Bayesian networks, game theoretic

concepts, and dynamic systems models to examine information barriers and verification protocols (Beumont et al. 2016). Others have examined how probabilistic approaches and analytic concepts can be used to estimate the confidence of host's cheating through aggregating monitoring and verification results in a Bayesian frame (Snowden 2019). Stepping back and examining a core probabilistic concept as a starting point for this work illustrates both a requirement for an ideal information barrier and the central problem of creating one. Mathematically, a single condition has to hold for a measurement to be an information barrier:

$$p(X|Y) = p(X) \qquad \text{(Eqn. 1)}$$

That is, the probability distribution over the quantity X does not change with the observation of Y. In the case of an information barrier, X would be the information the host is trying to protect, and Y is the signal the host is revealing. This property would be ideal for an information barrier because it means that the monitor learns nothing about the quantity X from observing Y. However, this property is also the definition of probabilistic independence, meaning that X must be informationally unrelated to Y for it to hold perfectly. This is the central dilemma of information barrier design.

Realistically, observing Y may provide some amount of information about X and the relationship in Equation 1 might not hold absolutely. A metric that can assess quantitatively the amount of information that is potentially communicated about X when Y is observed. It should provide a basis for comparison between information barrier designs and protocols, while being agnostic to the technology or approach used.

Fortunately, the field of information theory provides a rich set of concepts for addressing this problem.[2] Information theory is built on the concept of information entropy (reference to Shannon) which measures the average uncertainty in an unknown quantity, and is mathematically defined as:

$$H(X) = -\sum_{x \in X} p(x) \log_2 p(x) \qquad \text{(Eqn. 2)}$$

Information entropy is measured in bits when the base of the logarithm is two and indicates the expected amount of information that one would have to transmit to communicate the outcome, X. Mutual information measures the reduction in uncertainty (or information entropy) of a quantity X due to learning the outcome of a related quantity, Y, and is mathematically defined as:

$$I(X;Y) = \sum_{x \in X, y \in Y} p(x,y) \log_2 \frac{p(x,y)}{p(x)p(y)} \qquad \text{(Eqn. 3)}$$

where $p(x,y)$ is the joint distribution over both quantities, and $p(x)$ and $p(y)$ are the marginal distributions. Substituting Bayes theorem, we can derive the following:

$$I(X;Y) = \sum_{y \in Y} p(y) \sum_{x \in X} p(x|y) \log_2 \frac{p(x|y)}{p(x)} \qquad \text{(Eqn. 4)}$$

Equation 4 can also be expressed in terms of the expected Kullback-Liebler Divergence, D_{KL}, with respect to the observed quantity, Y (Kullback and Liebler 1951). The Kullback-Liebler Divergence, D_{KL}, is expressed as:

---

[2] While a full review of information theory is beyond the scope of this paper, see Cover and Thomas 2012 for an excellent reference text on the subject.

$$D_{KL}\big(p(X|y)||\,p(X)\big) = \sum_{x\in X} p(x|y)\log_2\frac{p(x|y)}{p(x)} \qquad \text{(Eqn. 5)}$$

The Kullback-Leibler divergence, $D_{KL}\big(p(X|y)||\,p(X)\big)$, can be interpreted as the amount of information that the monitor gains about X through the specific observation of $Y = y$. Taking the expectation of the D_{KL} gives the expression of mutual information in a different form:

$$I(X;Y) = \sum_{y\in Y} p(y)D_{KL}\big(p(X|y)||\,p(X)\big) \qquad \text{(Eqn. 6)}$$

Mutual information as expressed in Equation 6 can be interpreted as the amount of information about X that one can expect to learn when one observes Y in general. In this case, it is expressed in the units of bits given the base of the logarithm being two.

Calculating the mutual information as expressed in Equation 6 requires certain probability distributions be specified. It utilizes a prior distribution on the part of the monitor, $p(X)$, that describes the monitor's state of information about the quantity to be protected. It also relies on testable performance metrics of the measurement system, $p(Y|X)$, which describe how likely the system is to give a particular output, $Y = y$, given a specific input, $X = x$. Finally, it depends on the monitor's pre-posterior distribution over possible observed values, $p(y)$, which can be derived from $p(X)$ and $p(Y|X)$. This concept provides the basis for a proposed metric for information barrier performance.

Proposed Metric for Information Barrier Performance

We argue that an objective metric be used to understand information barrier performance, in order to aid in design and the development of mutual trust in information barrier systems. We have also argued mutual information, or the expected Kullback-Liebler divergence is the basis for such a measure. The problem is that such a quantity depends on a prior distribution, $p(X)$. Unfortunately, the host cannot know the monitor's state of information before the measurement. In order to overcome this problem, we may search the monitor's possible probability distributions for the distribution with the worst-case mutual information value for the verification protocol. Simply put, we propose that the maximum mutual information over all possible monitor probability distributions, $I^\star(X;Y)$, be used as a metric of information barrier quality. This can be expressed mathematically as:

$$I^\star(X;Y) = \max_{f\in\Delta X}\sum_{y\in Y} p(y)D_{KL}(p(X|y)||\,f) \qquad \text{(Eqn. 7)}$$

where f is a possible distribution over X in the simplex that describes all possible distributions of X, or $f \in \Delta X$. While this description treats a single protected quantity, X, and a single observed quantity, Y, the proposed approach is extensible to ensembles of both variable types.

Evaluating the metric $I^\star(X;Y)$ is subject to a few requirements. First, the information to be protected, X, must be identified, along with the information to be shared, Y. This would be spelled out in the verification protocol. Second, a model that describes how those two quantities are related must be developed. This model explicitly lays out the verification protocol mathematically and creates a mutual understanding of the process for both sides. Third, the performance of the measurement system $p(Y|X)$ must also be specified. This can be determined using test objects that are not sensitive, through modeling and simulation, much (if not all) of which could be done in partnership between the monitor and the host. Developing this data through cooperative experiments and research would also help to promote mutual trust and transparency.

An Example

As a proof of concept, we defined a hypothetical host/monitor verification regime that uses radiation measurements as a verification tool. In this hypothetical regime we have defined a treaty accountable item (TAI) to consist of two, and only two radioisotopes. In this example, a legitimate TAI consists of 7.263 μCi of Cs-137 and as second isotope of either 9.3 μCi of Co-60 or 10.0 μCi of Y-88. The goal of the agreement is to verify that a single TAI is present in the container presented by the host to the monitor. It is known that a legitimate TAI contains 7.263 μCi of Cs-137, but the host wants to protect the identity of the second isotope from the monitor. For the verification measurement, the monitor will measure gamma radiation between energies 579 keV and 735 keV emitting from the presented container with a NaI(Tl) detector.

For sake of simplicity, assume that the container can only be in one of six possible states as shown in Table 1. The container can be empty, registering a background spectrum at the detector (State 0). It can contain one of the two legitimate TAI (States 1 and 2). Or it can contain one of the three spoof items (States 3 through 5). Spectra for the items and environmental background were simulated using the Gamma Detector Response and Analysis Software (GADRAS) and the mean CPS within the energy window of interest were calculated.

Table 1. Composition of possible items in scenario

Source State	Legitimate or Spoof Item?	Activity of Cs-137 (μCi)	Activity of Co-60 (μCi)	Activity of Y-88 (μCi)	Mean CPS in Window
0	Empty (Background)	0	0	0	4.97
1	Legitimate	7.263	9.3	0	1.51×10^3
2	Legitimate	7.263	0	10.0	1.45×10^3
3	Spoof	7.263	0	0	1.02×10^3
4	Spoof	7.263	14.7	0	1.75×10^3
5	Spoof	7.263	0	12.8	1.70×10^3

Two information barriers designs are proposed to protect the identity of the second isotope. The first, information barrier A, will indicate "Valid" if the count rate between the energies 579 keV to 735 keV are within the range of $C \pm n\sqrt{C}$, of the two legitimate CPS mean values in Table 1. Otherwise, information barrier A will indicate "Invalid" to the monitor.

The second information barrier design proposed, information barrier B, utilizes two separate ranges, $C_1 \pm n\sqrt{C_1}$ and $C_2 \pm n\sqrt{C_2}$, where C_1 is the mean CPS for State 1 in Table 1, and C_2 is the mean CPS for State 2. If the measured CPS in the energy window of concern is within either range, information barrier B will indicate "Valid." Otherwise, it will indicate "Invalid" to the monitor.

Model

We can form a model of the monitor's reasoning in the example above, and then use that model to evaluate the proposed metric for information barrier effectiveness, $I^\star$, in the example scenario. Figure 1 illustrates a Bayesian network version such a model with four variables. The first is the Source State, X, and is described in Table 1. Source State has three child variables Second Isotope, $Q \in \{Co60, Y88\}$, which describes whether Co-60 or Y-88 is present; Legitimate, $W \in \{No, Yes\}$, which describes if the item presented is a legitimate TAI as described in Table 1; and Info Barrier, $Y \in \{Invalid, Valid\}$, which describes the output of the information barriers described above.

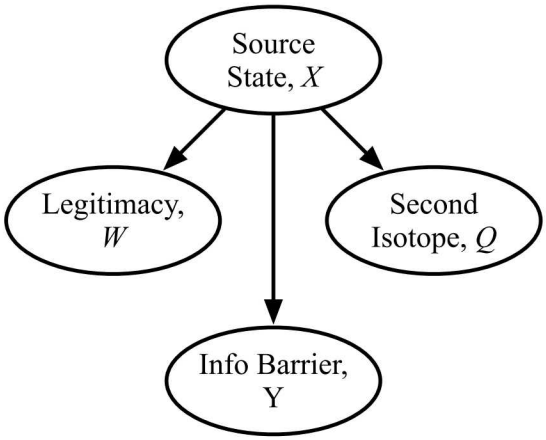

Figure 1: Monitor Reasoning Model

Table 2: Conditional Probability Tables

Source State	$P(Q = q\|X)$		P(W=w|X)	
	$q = Co60$	$q = Co60$	$w = No$	$q = Yes$
0	0.5	0.5	1.0	0.0
1	1.0	0.0	0.0	1.0
2	0.0	1.0	0.0	1.0
3	0.5	0.5	1.0	0.0
4	1.0	0.0	1.0	0.0
5	0.0	1.0	1.0	0.0

The prior distribution over the source state is the space over which the optimization in the proposed metric described in Equation 5 will be performed. The conditional probability tables describing the Second Isotope, $P(Q|X)$, and Legitimacy, $P(W|X)$, are shown in Table 2.

The conditional probability table for the Info Barrier, $P(Y|X)$, is calculated using a Poisson distribution function, with the appropriate mean CPS parameter as described in Table 1, and the ranges for each information barrier described above. An example of such a conditional probability distribution is given in Table 3 for a specific range parameter value of $n = 2.0$ and using the single range information barrier design (information barrier A). Note that the zeros in the table are rounded from small values ($< 10^{-3}$).

Table 3: Conditional Probability Tables

Source State	$P(Y = y\|X)$	
	$y = Invalid$	$y = Valid$
0	1.0	0.0
1	1.0	0.0
2	0.09	0.91
3	0.08	0.92
4	1.0	0.0
5	1.0	0.0

In this case, the host is trying to protect the true value of second isotope, or Q, while transmitting the most information about the legitimacy, or W. Thus, the proposed metric of information barrier performance is:

$$I^{\star}(Q|Y;Q) = \max_{f \in \Delta Q} \sum_{y \in Y} p(y) D_{KL}(p(Q|y) || f) \quad \text{(Eqn. 8)}$$

A second metric of interest can be defined as the amount of information that is transmitted about W at the distribution, $f^{\star}$, that maximizes Equation 6, or:

$$I(W|Y, f^{\star}; W|f^{\star}) = \sum_{y \in Y} p(y) D_{KL}(p(W|y, f^{\star}) || \, p(W|f^{\star})) \quad \text{(Eqn. 9)}$$

Arguably, the monitor should choose an information barrier (either A or B) and a value for the range size, n, that minimizes the value of $I^{\star}(Q|Y;Q)$ while providing for the largest value of $I(W|Y, f^{\star}; W|f^{\star})$ possible. That is, they wish to minimize the protected information transmitted and maximize the information that gives the monitor confidence that the presented object is legitimate.

Results

The model described above was implemented in Python and the metrics described in Equations 8 and 9 were assessed for information barrier designs A and B over various values of the range parameter n. The results of the model are shown in Figure 2. The top plot shows the performance for information barrier design A according to the proposed metric of effectiveness as a solid line. As the range parameter increases, and the valid range for CPS in the energy window 579 keV and 735 keV grows, more and more information about the protected quantity, Q, is potentially transmitted in the measurement, despite the monitor only getting a "Valid" or "Invalid" signal from the information barrier. At $n \approx 3.75$, nearly 0.8 bits of information are transmitted — a significant fraction given there is at most 1 bit of entropy available in Q. The information barrier does transmit a reasonable amount of information about the legitimacy, W, for many values of n, as shown by the dashed line. However, there is a range below $n = 1.0$ where very little information about the protected quantity, Q, is transmitted, and a reasonable amount of information is transmitted about the legitimacy, W.

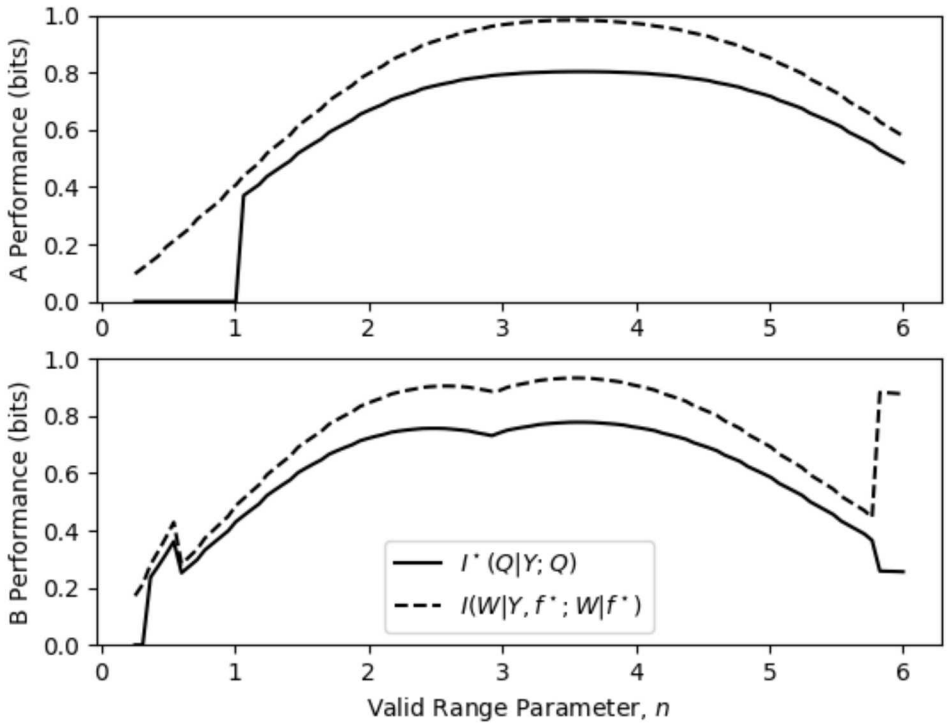

Figure 2: Information Barrier Performance Results

The bottom plot in Figure 2 shows the performance of information barrier B. In this case, a significant of information about the legitimacy, W, is transmitted for all values of n, as shown by the dashed line. In fact, at some values of n, more than 0.6 bits of information are transmitted through the information barrier. This indicates that the signal to the monitor is strong evidence about the validity of the object presented. However, there is also a large amount of information about the protected identity of the second isotope, Q, is also transmitted for all values of n. Information barrier design B does not offer a reasonable trade-off in between protecting sensitive information well and enabling verification.

The discontinuities in Figure 2 are worth noting. For information barrier A (the top plot), the discontinuities occur where the range is wide enough to prevent discrimination between the legitimate TAIs (Source States 1 and 2) but while minimizing confusion with the spoof items (Source States 3, 4, and 5). For information barrier B (bottom plot), as the two ranges grow, they begin to overlap and interact changing the detector's performance, and creating the discontinuities shown.

The results above would suggest that in this proof-of-concept, an information barrier with a single range and a parameter value of $n < 1.0$ might be a reasonable design. However, that result is not generalizable. The result depends heavily on the performance of the detector, the nature of the TAIs, and potential spoof objects.

Discussion

Recall that the proposed metrics of information barrier performance are dependent on the monitor's state of information about the public and protected quantities captured in the model. The host cannot know the monitor's state of information exactly. As a result, the proposed metric seeks the worst-case value by maximizing the information transmission over all possible information states of the monitor. One might propose an alternative approach where one assumes that if a host is uninformed about monitor's state of information, then a uniform distribution would be a reasonable representation of that monitor's state of information. Or, if the host may feel they have a reasonable estimate of the monitor's state of information, they could craft a prior distribution over Source State, X, that they feel accurately describes that information state. However, both of these approaches miss the central point that the results of the model and metric are to be arrived at in collaboration between the monitor and host. If both agree that the information barrier should protect certain information, then both should agree that minimizing the worst-case transmission is a reasonable goal.

The results shown in Figure 2 highlight the central challenge in designing effective information barriers: it is difficult to both convey information that builds confidence in an assertion while protecting related information. Now that a proposed objective metric for information barrier performance, $I^{\star}(Q|Y;Q)$, has been suggested, searches of an information barrier design's parameter space can yield optimal configurations. That is, given a model for a proposed information barrier, ideal parameter values could be determined through optimization that minimize $I^{\star}(Q|Y;Q)$. The result would be a lower bound for protected information transmission, providing a benchmark for implementations of that information barrier concept to be measured against in laboratory testing and joint experiments. These experiments can help to validate such implementations and foster trust on both sides in the verification system.

Finally, an important concept for metrics such as that proposed is that those metrics should be technology agnostic. In this case, the metric should be able to be evaluated against any verification regime proposed or proposed measurement and monitoring technology. This drives metrics development towards metrics that focus on performance against the problem, and less on specific technical aspects. For instance, one might make the argument that scintillator performance (resolution, efficiency, etc.) are important metrics of the measurement system, and they are related to overall performance. However, they may not provide a direct and comparable understanding of performance against the objectives of verification regime, which

is to provide confidence to the monitor and the host that the treaty is being upheld and that information is protected.

Next Steps

There are several possible next steps in the development of the concepts and metrics presented in this paper. First, a study exploring the application of this approach to common information barrier and warhead verification measurement schemas could be instructive and help further develop out the approach. This study would also need to include each of the competing objectives of instilling confidence in the monitor of the warhead authentication through the measurement and protecting the information of the host.

Second, the concepts presented in this paper can be expanded upon to create a design approach for information barriers, and verification regimes more broadly. Formally capturing verification regimes as probabilistic models and principles using these metrics, and then mathematically exploring the design space for verification protocols and technology may yield important insights about protocol and verification equipment design. Doing so would allow for a more explicit examination of the implicit tradeoffs and augment best-practices approaches. Such an approach could also provide useful insights when designing a verification technology with an information barrier by considering what information is transmitted to the monitoring party in case of some degree of failure in the verification technology.

Third, experimental data could be collected from existing system with relevant information barriers that were designed for hypothetical verification scenarios. This data would be used to estimate some of the probabilities required from empirical data. The results would allow for the further development and proof-of-concept of the design approach advocated in the example. An incremental step in this direction would be to recreate the example problem from this paper in the laboratory and experimentally verify the results.

Finally, a set of test problems could be created to evaluate the performance of information barrier systems designed for hypothetical varication scenarios. Their theoretical lower-bound information barrier performance metric could be estimated using the approaches above. Lab tests of these technologies could then be performed using the test problems in order to estimate the information barrier performance. The physical system performance could then be compared to theoretical limits. Each of these steps can help to further the understanding of how systems with information barriers would perform in hypothetical verification scenarios. The work presented in this paper also provides a path towards possibly creating an objective metric for comparing differing systems, designed for the same task, in the hypothetical scenarios.

Conclusions

We have proposed an information theoretic metric for assessing information barrier effectiveness and performance when used as a verification tool in a hypothetical arms-control scenario. We have also provided a proof-of-concept example that illustrates how such a metric could be used to evaluate and compare alternate information barrier designs. This metric is the maximum mutual information between a quantity protected by the host and the quantity observed by the monitor, or $I^{\star}(Q|Y;Q)$. The only time an information barrier can truly transmit no information about one uncertainty (such as weapon design information), while transmitting information about a different uncertainty (such as a measurement of the weapon), is when those uncertainties are probabilistically independent. Information transmission through information barriers can depend not only on the mechanism and protocols for protecting information, but also on the prior information held by the monitoring party. For these reasons, it may impossible or impractical to eliminate certain kinds of information transmission through verification systems, and objective metrics of performance are necessary in order to understand if proposed solutions are effective, and to what degree sensitive information on both sides is protected.

References

Beaumont, Paul, Neil Evans, Michael Huth, and Tom Plant. "Nuclear Arms Control: Optimising Verification Procedures Through Formal Modelling." 2016.

Bonin, Ben, Eva Uribe, Jason Reinhardt, Michael Hamel. “Facilitating International Cooperation through a Metrics Framework for Comparative Assessment of Nuclear Arms Control Verification Methods.” Institute for Nuclear Material Management 2020 Annual Meeting Proceedings. July 2020.

Close, D. A., D. W. MacAuthur, and N. J. Nicholas. “Information Barriers – A Historical Perpsective.” Los Almos National Laboratories. LA-UR-01-2180. 2001

Cover, Thomas M. and Joy A. Thomas. *Elements of Information Theory.* John Wiley & Sons. 2012.

Fuller, James. “Verification on the Road to Zero: Issues for Nuclear Warhead Dismantlement.” Arms Control Today. Vol. 40. No. 10. pp 19-27. 2010.

Glaser, Alexander, Boaz Barak, Robert Goldstein. “A Zero-Knowledge Protocol for Nuclear Warhead Verification.” Nature. Vol. 510. pp. 497-502. June 2014.

Kullback, Solomon and Richard Liebler. “On Information and Sufficiency.” Annals of Mathematical Statistics. Vol. 22 (1). pp. 79-86. 1951.

MacArthur, Duncan, M. William Johnson, Nancy Jo Nicholas, Rena Whiteson. “Use of Information Barriers to Protect Classified Information.” Institute of Nuclear Materials Management Conference. Naples, Florida. July, 1998.

Mitchell, Dean J. and Keith M. Tolk. “Trusted Radiation Attribute Demonstration System.” Sandia National Laboratories. SAND2000-1481C. 2000.

National Academy of Science. *Monitoring Nuclear Weapons and Nuclear-Explosive Materials: An Assessment of Methods and Capabilities.* Committee on International Security and Arms Control, National Research Council. pp. 107-108. 2005.

Office of Secretary of Defense. “Treaty Between the United States of America and The Union of Soviet Socialist Republics on Strategic Offensive Reductions (START I).” Website. Available at: https://www.acq.osd.mil/tc/start1/START1text.htm. Accessed: June 10, 2020.

Seager, Kevin D., Dean. J. Mitchell, T. W. Laub, K. M. Tolk, R. L. Lucero, and K. W. Insch. “Trusted Radiation Identification System.” Proceedings of the 42nd Annual INNM Meeting. Indian Wells, CA. 2001.

Shannon, Claude E. “A Mathematical Theory of Communication.” The Bell System Technical Journal, Vol. 27, pp. 379–423, 623–656, July, October, 1948.

Snowden, Mareena Robinson. “Probabilistic Verification: A New Concept for Verifying the Denuclearization of North Korea.” Arms Control Today. September, 2019. Available at: https://www.armscontrol.org/act/2019-09/features/probabilistic-verification-new-concept-verifying-denuclearization-north-korea. Accessed: June 10, 2020.

U.S. Department of State. “New START.” Website. Available at: http://go.usa.gov/cfXkz. Accessed: March 11, 2016

Zuboski, Peter B., Joseph P. Indusi, Peter E. Vanier. “Building a Dedicated Information Barrier System for Warhead and Sensitive Item Verification.” Brookhaven National Laboratory. BNL-66214. 1999.