Addressing the Need for Independence in the CSE Model

Robert K. Abercrombie, Erik M. Ferragut, Frederick T. Sheldon
Computational Sciences and Engineering Division
Oak Ridge National Laboratory
Oak Ridge, TN 37831-6418 USA
abercrombier@ornl.gov, ferragutem@ornl.gov, sheldonft@ornl.gov

Michael R. Grimaila
Department of Systems Engineering and Management
Center for Cyberspace Research
Air Force Institute of Technology (AFIT)
Wright-Patterson Air Force Base, Ohio 45433-7765 USA
Michael.Grimaila@afit.edu

***Abstract*—Information system security risk, defined as the product of the monetary losses associated with security incidents and the probability that they occur, is a suitable decision criterion when considering different information system architectures. Risk assessment is the widely accepted process used to understand, quantify, and document the effects of undesirable events on organizational objectives so that risk management, continuity of operations planning, and contingency planning can be performed. One technique, the Cyberspace Security Econometrics System (CSES), is a methodology for estimating security costs to stakeholders as a function of possible risk postures. In earlier works, we presented a computational infrastructure that allows an analyst to estimate the security of a system in terms of the loss that each stakeholder stands to sustain, as a result of security breakdowns. Additional work has applied CSES to specific business cases. The current state-of-the-art of CSES addresses independent events. In typical usage, analysts create matrices that capture their expert opinion, and then use those matrices to quantify costs to stakeholders. This expansion generalizes CSES to the common real-world case where events may be dependent.**

Keywords-Information Security, Cybersecurity Metrics, Risk Analysis, Risk Management, Stakeholder Value, Information Assurance Controls, Migitation Costs, Algorithms

I. INTRODUCTION

Title 44 of the U.S. Code [1] defines Information security as a means of protecting information and information systems from unauthorized access, use, disclosure, disruption, modification, or destruction in order to provide:

- integrity, which means guarding against improper information modification or destruction, and includes ensuring information nonrepudiation and authenticity;
- confidentiality, which means preserving authorized restrictions on access and disclosure, including means for protecting personal privacy and proprietary information; and
- availability, which means ensuring timely and reliable access to and use of information.

Risk is formally defined as "the effect of uncertainty on objectives" (ISO 73 2009 [2]). Risk is often characterized by a probabilistic analysis involving both the likelihood of potential events and their consequences on the organizational objectives. Risk management involves "coordinated activities to direct and control an organization with regard to risk" (ISO 73 2009 [2]). Historically, risk management is a highly introspective knowledge acquisition process that yields benefits not only in the ability to reduce uncertainty in the outcomes of objectives, but also provides the prerequisite knowledge necessary to proactively develop contingency plans which can be placed into action should a risk materialize.

Organizations typically use a risk management process to identify and mitigate risks to assure their organizational mission [3]. Risk management provides a documented, structured, and transparent process to identify critical resources, estimate threats, and vulnerabilities. The intersected set of threats and vulnerabilities cause harm (risks) to those identified resources. Moreover, the process estimates the likelihood of risk occurrence and evaluates tradeoffs among control measures used to mitigate the risks, and periodically revisits the analyses as needed. However, the value of the analysis is a strong function of the accuracy of the inputs to the process.

In this paper, we first review, in Section 2, related works applicable to this investigation, of generalizing and thus expanding Cyberspace Security Econometrics Risk Management tools, techniques, and methodologies to include dependent events. In Section 3, the motivation for the current investigation is discussed as it relates to cyberspace security econometrics risk management. In this section, the concept of dependence is introduced. In Section 4, the foundations of the Cyberspace Security Econometrics System (CSES) are articulated. In Section 6, the expansion of the CSES foundations, explained in Section 5, is developed with an additional explanation of dependency and introducing two

The submitted manuscript has been authored by a contractor of the U.S. Government under contract DE-AC05-00OR22725. Accordingly, the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

assumptions (Assumption One – the probabilities of different matrices are independent; and Assumption Two – Entities in each class (other than Stakeholders) must correspond to disjoint events). In Section 7, the expansion is further generalized by addressing Assumption Two, resulting in accommodating the fact that dependencies between events naturally exist or will occur. Finally, in Section 8, we summarize the paper and discuss where future research is warranted and should proceed.

II. Related Work Applicable to Investigation

Organizations typically implement an IT focused risk management process to identify and mitigate IT related risks in order to assure their organizational mission [4-6]. These enterprises span industries and infrastructures listed in the next several paragraphs.

The European Network and Information Security Agency (ENISA) has generated an inventory of risk management and risk assessment methods [7]. A total of 13 methods were considered. Each method in the inventory has been described through a template. The template used consists of 21 attributes that describe characteristics of a method. The inventory website also provides for the comparison of the risk management methods and also the risk management tools [8].

In the context of applying these techniques to an application space, a *critical infrastructure* is, "an infrastructure so vital that it incapacity or destruction would have a debilitating impact on our defense and national security" [9, 10]. With respect to various subject domains, the associated critical infrastructure consisting of 11 sectors (agriculture and food, water, public health, emergency services, defense industrial base, telecommunications, energy, transportation, banking and finance, chemical and hazardous material, and postal and shipping), and 5 key assets (national monuments and icons, nuclear power plants, dams, government facilities, and commercial key assets) have been identified. Vulnerability analysis of these critical infrastructures is well documented, using traditional techniques [11]. Recently this list is expanding and the electric power industry is applying similar techniques to SCADA (supervisory control and data acquisition) equipment [11], cyber threats [12], cyber security [12], the Internet, and the Smart Grid. The oil and gas industry in particular and the chemical process industry in general store and transport hazardous and energetic chemicals, and operate processes under extreme conditions of temperature and pressure. Terrorists or disgruntled employees may exploit these conditions, which may then lead to toxic release, fire and explosion resulting in mass casualties, property damage, and economic and environmental impacts [13].

Newer approaches are emerging that apply information security techniques of easing the complexity of creating tree and graph structures and deriving probabilistic defense graphs from network architectural models [14]. Further refinements have recently been documented [15] and additional methodologies have been applied dealing with enhanced dynamic decision making [16].

III. Cyber Security Econometrics System (CSES) Motivation for Current Investigation

The Roadmap for Cybersecurity Research articulates that information technology has become pervasive in every way — from our phones and other small devices to our enterprise networks to the infrastructure that supports our economy [12]. Improvements to the security of this information technology are essential for our future. As the critical infrastructures of the United States have become more and more dependent on public and private networks, the potential for widespread national impact resulting from disruption or failure of these networks has also increased. Securing the nation's critical infrastructures requires protecting not only their physical systems but, just as important, the cyber portions of the systems on which they rely.

The motivation for this work, is highlighted by existing and emerging technologies that complement the Roadmap [12] in the context of the survivability of time-critical systems. The President's Comprehensive National Cybersecurity Initiative [17] also emphasizes the need for leap-ahead improvements in security of cyber physical systems. A failure is inclusive of random events, design flaws, and instabilities caused by cyber (and/or physical) attack. One such domain, optimizing investments in critical infrastructure protection, is applicable to the use of the Cyber Security Econometrics System (CSES). We discuss the workings of such a system in this context of the need for optimizing investments, the CSES mathematical foundations, the linear cascading of linear models, and the implications of this computing architecture, in particular, with respect to our nation's critical cyber infrastructure and key resources.

Combining the subject domains of information security, and risk management, CSES is a methodology for estimating security costs to stakeholders as a function of possible risk postures. In earlier works, we presented a computational infrastructure that allows an analyst to estimate the security of a system in terms of the loss that each stakeholder stands to sustain as a result of security breakdowns [18-20]. More recently, we presented how this infrastructure can be used in the subject domain of mission assurance [21]. Additional work has applied CSES to specific business cases [20, 22].

The current state-of-the-art of CSES addresses independent events. In typical usage, analysts create matrices that capture their expert opinion, and then use those matrices to quantify costs to stakeholders. In situations where the underlying events exhibit significant dependencies, the current approach is not appropriate. Unfortunately, significant dependencies are likely to arise with any detailed modeling of a complex system of components, such as an enterprise network.

A. The Introduction of Dependence

There are, in fact, two similar concepts that need to be distinguished. The first is dependence. Two events are dependent if the probability of one event is altered by knowing whether the second event occurred. This is the issue that needs to be dealt with directly to make CSES as generally applicable as possible. The second concept is disjointness. Two events are disjoint if they cannot both occur simultaneously. Inevitably, disjoint events are dependent and independent events cannot be

disjoint (see Figure 1). However, while dependence needs to be addressed, the mathematics of CSES exploits the disjointness of the underlying events. In particular, matrix products are used to compute probabilities; this is a computational application of the Law of Total Probability [23] provided that the events are disjoint. Our new improvement to CSES continues to use matrix products as an efficient way to compute outcome probabilities, but directly addresses dependence among events by enumerating all possible combinations. Although this yields a potentially intractable number of event combinations to consider, we argue that the dependency structure among these events will generally greatly simplify the manual work required.

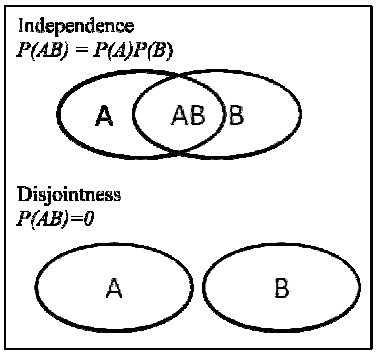

Figure 1. Independent versus disjoint events.

IV. Cyberspace Security Econometrics System (CSES) Foundations

CSES provides many advantages over other known measurement or analysis systems or methodologies such as: (1) it reflects variances existing between different users or stakeholders of the system [18]. Different stakeholders may attach different stakes to the same requirement or service (e.g., a service may be provided by an information technology system, cyber, enterprise or process control system, etc.). (2) For a given stakeholder, CSES can highlight variances that may exist among the stakes attached to satisfying each requirement. For example, a stakeholder may attach or identify different stakes to satisfying different requirements within the overall system. (3) For a given compound specification (e.g., combination(s) of commercial off the shelf software and/or hardware), CSES can identify variances that may exist amongst the levels of verification and validation that are performed on components of the system (or specification). The verification activity may produce higher levels of assurance in satisfying some components of the specification more than others.

The CSES follows a defined process [18, 19, 24, 25]. The initial inputs (1) organizational mission (and components thereof), (2) value of its objectives and assets if uninterrupted, and (3) the components of the enterprise system that support each mission component, are determined by stakeholders.

The stakeholders, with assistance from Subject Matter Experts (SMEs), define the criteria of a quantitative value of an asset. For example, the criteria may include:

- Financial basis (e.g., operational cost of downtime per unit of time defined by hardware/software costs, facilities and staffing versus profit); which is the quantitative measurement to be used within the CSES.
- Federal Information Security Management Act (FISMA) of 2002 [26], stakeholder derived value of assets per NIST 800-60 [27], and/or FIPS 199/200 [28, 29] dictated requirements.
- Stakeholder defined requirements; acceptable and unacceptable impact levels against the value related to IA tenets of confidentiality, availability and integrity may also be examined.

The CSES process proceeds in three steps (generation of stakes matrix, dependency matrix, and threat matrix) to derive the mitigation costs matrix [18]. CSES accounts for failure costs and verification (i.e., mitigation costs). CSES provides a:

- Framework for measuring the appropriate attributes that support the decisions necessary to (1) design security countermeasures, (2) choose between alternative security architectures, (3) respond to events such as intrusions or attacks, and (4) improve security (including reliability and safety) during both design and operational phases.
- Comprehensive basis for choosing courses of action that have the highest risk reduction return on investment, i.e., reduce the most risks for the lowest cost.

The basis of CSES stems from and is consistent with the spirit of Value Based Software Engineering [30, 31].

CSES comprehends the different organizational mission needs for all stakeholders, including reliability and safety [20]. CSES identifies information assurance controls and mitigation costs as an investment toward assuring mission success [21].

V. CSES - A Cascade of Linear Models

In this section, we present the composition of the CSES model and motivate its application.

A. The Stakes Matrix

We consider a system S and we let $H_1, H_2, H_3, \ldots, H_k$ be stakeholders of the system, i.e., parties that have a stake in its operation. We let $R_1, R_2, R_3, \ldots, R_n$, be security requirements that we wish to impose on the system, and we let $ST_{i,j}$, for $1 \le i \le k$ and $1 \le j \le n$ be the stake that stakeholder H_i has in meeting requirement R_j. We let PR_j, for $1 \le j \le n$, be the probability that the system fails to meet security requirement R_j, and we let MFC_i, (Mean Failure Cost), for $1 \le i \le k$, be the random variable that represents the cost to stakeholder H_i that may result from a security failure.

We quantify this random variable in terms of financial loss per unit of operation time (e.g. $/hour); it represents the loss of service that the stakeholder may experience as a result of a security failure. Under some assumptions of statistical independence, we find that the computed expected value of the Mean Failure Cost for stakeholder T_i can be written as:

$$MFC_i = \sum_{1 \le j \le n} ST_{i,j} PR_j .$$

If we let MFC be the column-vector of length k that represents mean failure costs, let ST be the $k \times n$ matrix that represents stakes, and let PR be the column-vector of length n

that represents probabilities of failing security requirements, then this can be written using the matrix product (∘):

$$MFC = ST \circ PR.$$

The Stakes matrix is filled, row-by-row, by the corresponding stakeholders. As for *PR*, we discuss below how to generate it.

B. The Dependency Matrix

We consider the architecture of system *S*, and let C_1, C_2, C_3, …, C_h, be the components of system *S*. Whether a particular security requirement is met or not may conceivably depend on which component of the system architecture is operational. Lets we assume that no more than one component of the architecture may fail at any time. We define the following events:

- E_i, $1 \le i \le h$, is the event: the operation of component C_i is affected due to a security breakdown.
- E_{h+1}: No component is affected.

Given a set of complementary events E_1, E_2, E_3, …, E_h, E_{h+1}, we know that the probability of an event F can be written in terms of conditional probabilities as:

$$P(F) = \sum_{k=1}^{h+1} P(F \mid E_k) P(E_k).$$

We instantiate this formula with *F* being the event: the system fails with respect to some security requirement. Expanding on this effect, we let F_j denote the event that the system fails with respect to requirement R_j and we write (given that the probability of failure with respect to R_j is denoted by PR_j):

$$PR_j = \sum_{k=1}^{h+1} P(F_j \mid E_k) P(E_k).$$

If

- we introduce the *DP* (Dependency) matrix, which has *n* rows and *h*+1 columns, and where the entry at row *j* and column *k* is the probability that the system fails with respect to security requirement *j* given that component *k* has failed (or, for *k*=*h*+1, that no component has failed),
- we introduce vector *PE* of size *h*+1, such that PE_k is the probability of event E_k, then we can write

$$PR = DP \circ PE.$$

Matrix *DP* can be derived by the system's architect, in light of the role that each component of the architecture plays to achieve each security goal. As for deriving vector *PE*, we discuss this matter in the next section.

C. The Impact Matrix

Components of the architecture may fail to operate properly as a result of security breakdowns brought about by malicious activity. In order to continue the analysis, we must specify the catalog of threats that we are dealing with, in the same way that analysts of a system's reliability define a fault model. To this effect, we catalog the set of security threats that we are facing, and we let T_1, T_2, T_3, …, T_p, represent the event that a cataloged threat has materialized, and we let T_{p+1}, be the event that no threat has materialized. Also, we let PT be the vector of size *p*+1 such that

- PT_q, for $1 \le q \le p$, is the probability that threat T_q has materialized during a unitary period of operation (say, 1 hour).
- PT_{p+1} is the probability that no threat has materialized during a unitary period of operation time.

Then, by virtue of the probabilistic identity cited above, we can write:

$$PE_k = \sum_{q=1}^{p+1} P(E_k \mid T_q) PT_q.$$

If

- we introduce the *IM* (Impact) matrix, which has *h*+1 rows and *p*+1 columns, and where the entry at row k and column *q* is the probability that component C_k fails given that threat *q* has materialized (or, for *q*=*p*+1, that no threat has materialized),
- we introduce vector *PT* of size *p*+1, such that PT_q is the probability of event T_q, then we can write

$$PE = IM \circ PT.$$

Matrix *IM* can be derived by analyzing which threats affect which components, and assessing the likelihood of success of each threat, in light of perpetrator behavior and possible countermeasures. Vector *PT* can be derived from known perpetrator behavior, perpetrator models, known system vulnerabilities, etc. We refer to this vector as the *Threat Configuration Vector* or simply as the *Threat Vector*.

VI. Expansion of Method

The current state-of-the-art of CSES addresses independent events. This involves analysts creating matrices that capture their expert opinion, and then using those matrices to quantify costs to stakeholders. This section describes how to expand this to include dependent events.

A. Expert Matrices

Analysts create several matrices that relate various types of entities:

1. A stakeholder-requirement matrix gives the cost (in a common unit, such as USD) of having that requirement fail (for that stakeholder).

2. A requirement-component matrix gives the conditional probability of the requirement failing given that the component fails.
3. A component-threat matrix gives the conditional probability of the component failing given that a particular threat exists.
4. A threat-mitigation matrix gives the probability of a threat existing given that a particular mitigation (or set of mitigations, a.k.a. a risk posture) is in place.

B. Matrix Products

The resulting product of matrices will give a threat-mitigation matrix with rows indexed by stakeholders and columns indexed by mitigations. The entries are the expected costs to each stakeholder for that mitigation. This follows from the nature of the conditional probabilities, but refers to the next subsection for underlying assumptions. This allows decision-makers to determine which mitigation to use.

Other matrix products may also be relevant. In particular, suppose an analyst is interested in knowing the dependence of the results on a particular entry of the component-threat matrix. The effect of that entry on the stakeholder mitigation product matrix can be computed by replacing the component-threat matrix with a matrix of all zeros except for a single one at the entry of interest. The resulting product gives the matrix of derivatives of each matrix product entry with respect to the component-threat entry of interest. This linearity can be used:

1. for sensitivity analysis (to the expert opinions),
2. to search for areas of weakness (e.g., a component that could be hardened against a threat and would yield considerable cost reductions), and
3. to determine non-critical entities (e.g., a component whose security against threats has little or no implications for the cost to stakeholders).

C. Underlying Assumptions

The matrix product makes sense as a product of conditional probabilities. In general,

$$P(A_1 \mid A_2)P(A_2 \mid A_3)\cdots P(A_{n-1} \mid A_n) = P(A_1 A_2 \cdots A_{n-1} \mid A_n)$$

provided that the events are conditionally independent. That is, A_n depends on A_1 only through the influence of A_{n-1}. This gives the first CSES assumption:

Assumption One: The probabilities of different matrices are independent.

For example, the probability of a component failing given the presence of a threat does not depend on which requirement that component is serving, or on the mitigation that is in place (since the threat is already assumed to be present).

We believe this assumption to be fundamentally sound, provided that the components and threats are defined reasonably.

A second assumption involves the summation of terms within a matrix product.

Generally,

$$P(A_1 \mid A_n) = \sum_{a_2, a_3, \ldots a_{n-1}} P(A_1, A_2 = a_2, A_3 = a_3, \ldots, A_{n-1} = a_{n-1} \mid A_n)$$

as long as the joint events, each one described by the set of values a_2, a_3, . . ., a_{n-1}, are disjoint. (This is one version of the Law of Total Probability.) Otherwise, the equation requires correction factors for intersections of events. (See Stanley's standard combinatorics texts for more details on principles of inclusion-exclusion [32].) This leads to the second CSES assumption:

Assumption Two: Entities in each class (other than stakeholders) must correspond to disjoint events.

For example, if two components can fail independently (i.e., neither, either one, or both can fail) then the two component failures are not disjoint (i.e., both can fail because they are not mutually exclusive). Refer to Figure 1.

As a first iteration of CSES, we can suppose that the component failures are rare and relatively independent events. In that case, the probability of multiple component failures is considered negligible. In such applications, Assumption Two is appropriate. However, a general solution for the issue of dependence among events is needed to broaden the field of legitimate applications of CSES. This is a standing issue with the state-of-the-art. The next section describes our general solution that directly addresses dependence. Our method covers the generic case, but specializes to the current CSES when independence and rarity hold.

VII. Enumerating Disjoint Events

We propose to address the lack of generality of Assumption Two by enumerating a finite set of disjoint events. This will make CSES applicable even without Assumption Two. In this section, we first describe the enumeration of all combinations of interest in Subsection A. Next, Subsection B shows that this can quickly lead to manually intractable requirements. However, it is expected that structural relationships and quantitative comparisons between events can be used to greatly simplify the demands on domain experts with respect to interpreting (Subsection C) and populating (Subsection D) the model parameters.

A. The Enumeration

If a system contains components A, B, and C, then there are eight possible component failure states (one of which represents no failures) based on the eight subsets of (*A, B, C*) that fail. (In general, we have 2^k states for k entities.) We propose that the original sets of requirements, components, threats, and mitigations be replaced by all of their possible combinations. The resulting matrices, while being exponentially larger, will (by definition) satisfy Assumption Two, thereby making the calculations relevant.

These larger expert matrices potentially introduce new challenges to the CSES methodology with respect to

1. calculating the product matrix,
2. interpreting the product matrix, and
3. populating the expert matrices.

We discuss these in the next three subsections, respectively.

B. Calcuation

The $n_1 \times n_3$ product of an $n_1 \times n_2$ matrix with an $n_2 \times n_3$ matrix can be computed in work $O(n_1 n_2 n_3)$. For a product of four matrices, the work becomes $O(n_1 n_2 n_3 + n_1 n_3 n_4)$ or $O(n_2 n_3 n_4 + n_1 n_2 n_4)$, depending on the order of operations. (A subfield of combinatorial optimization has addressed how to best order the computations for efficiency [33].) Since we only have a small number of matrices, we can reasonably handle 1000×1000 matrices; significantly larger matrices may pose a problem. In the near term, we propose limiting the analysis to these relatively small matrices. (This allows for up to 10 independent factors prior to the combinatorial enumeration.) With additional improvements in efficiency that exploit structural properties and sparsity, we expect to be able to extend the practical matrix size limits. In particular, since we will be using the dependency structure among the events to populate the matrices (see Subsection D), we should be able to use the same structure to computationally simplify the product calculation. We leave this as future work.

C. Interpretation

The size of the resulting matrix will still only have as many rows as stakeholders originally considered. The number of mitigations may now have increased exponentially. Comparing the resulting mitigations may be a considerable and human-intensive challenge.

We propose a triage process for reducing the number of mitigations that need to be considered. First, if the mitigation costs can be estimated, we can eliminate all mitigations that are too costly. These mitigations are the enumerated combinations of the original set of mitigations. While no original mitigation would be included if it was cost-prohibitive, it is likely (since CSES is being used to support security decisions) that there are combinations of these that are too expensive.

Second, if there are unacceptably high costs to certain stakeholders, we can eliminate the mitigations resulting in those high costs. The use of stakeholder utilities allows for a highly expressive means of accounting for arbitrary stakeholder concerns.

Third, if two mitigations are such that one is both cheaper to implement and costs less to the stakeholders (this is many simultaneous conditions), then the other mitigation can be dropped. These notions should help to reduce even a large set of mitigations to a small set of "best" candidates for analysis by a human.

D. Populating

If each of the four expert matrices is about 1000 × 1000 in size, then there are 4 million entries that need to be assigned. In practice, assigning them individually is inadvisable since

1. it would take far too long,
2. the resulting entries are bound to be inconsistently derived, and
3. any structural relations between the entries must be enforced manually.

We propose explicitly outlining the relationships between the components. For this, we suggest using a Bayesian belief network to capture structural (conditional) dependencies. A Bayesian belief network is a probability model that efficiently captures the potentially highly complex relationships between variables by "factoring" them as conditional dependencies. A collection of variables $X_1, X_2, \ldots, X_k$ are given together with dependencies of the form $X_i \rightarrow X_j$. This is interpreted as saying that the variable X_j depends directly on all X_i such that $X_i \rightarrow X_j$. These dependencies define a directed acyclic graph.

The model must also specify the distribution of X_j given all the X_i on which it directly depends. Since, for our purposes, it suffices to consider only categorical variables, these distributions can be thought of as look-up tables. The rows of the table are indexed by all combinations of values of the input variables (i.e., the X_i), and the value in the table is the probability

$$P(X_j = x_j \mid X_{i1} = x_{i1}, \ldots, X_{ik} = x_{ik}).$$

Additional dependencies do exist in a Bayesian belief network in cases where, say, $X_i \rightarrow X_j$ and $X_j \rightarrow X_k$. However, these can be accounted for with the tables already described by factoring through other variables as needed. The fundamental underlying hypothesis is conditional independence, a generalization of the Markovity assumption for Markov chains. Conditional independence means that the probability of X_j is independent of all X_i that it is *not* directly dependent on, *given* the values of the X_i that it *is* dependent on. This one elegant assumption means that a relatively small number of parameters can be used to fully describe the joint probability distribution of all of the variables. This efficiency is especially pronounced when the dependencies $X_i \rightarrow X_j$ are a sparse set among the set of all variables. Given the joint distribution, it is then possible to perform various types of inference, to learn the tables from data (called training), to compute the probability of various types of events, to synthesize (i.e., simulate) data, and to compare multiple hypotheses for the same data.

Our general scheme for efficiently populating the matrices is to devise a Bayesian belief network for the events and to have experts create the conditional probabilities. In these networks, the original events (prior to the combinatorial enumeration) are the variables. These will include the security requirements, the component failures, and threats. Some dependencies between the events will naturally exist (e.g., tightly coupled components, sub-component to super-component relationships, threats on particular components, mutually exclusive threats), but personal experience with CSES suggests that most pairs of events will be fundamentally independent. The users of CSES are required to establish the dependencies in the network, as well as to populate its conditional probability distributions (i.e., the look-up tables). There is, as there always is in modeling real world situations, a fundamental trade-off between the fidelity of the model and its

simplicity. This corresponds to including more dependencies in the former and fewer in the latter.

Given the model, it is routine to compute the probability for any given combination of events, thereby populating the matrices. This method has a significant advantage over directly populating a large matrix with probabilities because it enforces a high degree of consistency. Where multiple entries in the matrix may take into account a single conditional probability (since a single simple event will occur in half of the enumerated combinations), a manual approach would demand that the user account for this probability in a consistent way in all cases. The Bayesian belief network makes this accounting explicit and, therefore, consistent. Furthermore, it is far more efficient, especially for sparse Bayesian belief networks, which have relatively few dependencies.

We claimed above that our method generalizes the current CSES approach when the events are independent and rare. In this case there are no dependencies $X_i \rightarrow X_j$ in the network. Here, the computation of a combination of events is the product of the component events:

$$P(X) = \prod P(X_i = x_i).$$

When the events are assumed to be rare, the probability of two simultaneous failures becomes negligible. Hence the only effectively non-zero probabilities are the original set of events. The non-zero component of the matrix product is then the same as the original matrix product outlined in Section III. This shows that the original need for 2^k probabilities to specify the matrix entries (for column corresponding to that threat) is reduced to just k probabilities.

As a second example, we consider a slightly more complex situation. Suppose we have two components, a gate and a pump (stored behind the gate). Consider a physical threat against the pump that must first defeat the gate. The probability of defeating the gate can be estimated by experts. Then, experts can estimate the conditional probability distribution of sabotaging the pump given that

1. the gate was defeated, or
2. the gate was not defeated.

The efficiency here is that the four cases can be computed given only three numbers. (These are the probability of success against the gate and the two conditional probabilities against the pump given the result of the gate action.) In larger, more complex cases the reduction in expert values needed is far more dramatic. For example, if there were a sequence of k gates, we would need one probability for the first gate, and two probabilities for each subsequent gate and the pump. The original requirement for $2^{(k+1)}$ probabilities is therefore reduced to $2k+1$ probabilities.

VIII. Summary and Future Work

We have described the state-of-the-art of CSES and the fundamental shortcoming, which we referred to as Assumption Two – Entities in each class (other than Stakeholders) must correspond to disjoint events. We then proposed a method for addressing that assumption in a general setting via the use of combinatorial enumeration. The resulting issues from the combinatorial explosion can be addressed by an application of Bayesian belief networks (to incorporate expert opinion) and basic exploratory heuristics (to analyze the results).

Future work should include (1) additional implementations of this method, and (2) exploration of computational enhancements for particular structured matrices, which would lead to (3) validation via real-world examples. Further investigation is warranted in merging our method with a hybrid analytic dynamic forecasting (HADF) methodology that combines the techniques of analytic hierarchy processes, factor analysis, and spanning trees to the problem of selecting among a set of contingency measures following events which place the organizational mission at risk [16]. The HADF methodology makes use of qualitative subjective assessments by subject matter experts at multiple levels of the organization and uses historical event occurrences (when available) to provide the decision maker with an ongoing recommendation of the best contingency measures to employ to assure the organizational mission objectives. The method is novel because it augments the decision maker's experiential knowledge with a probabilistic forecast of the best contingency measure to take in response to events based upon SME knowledge, historical evidence, and the real-time status critical resources. The methodology provides a structured approach to mitigate operational risk in complex environments and decreases the time required to make decisions under conditions of uncertainty.

In addition to better supporting off-line cyber security decisions, handling of dependencies in CSES improves its support of on-line cyber security analysis. In particular, it could be used to construct utility functions for multi-party cyber conflict simulations, improving the fidelity of the game theoretic analyses in related work [34].

Disclaimer and Acknowledgment

The views expressed in this paper are those of the authors and do not reflect the official policy or position of the United States Air Force, the Department of Defense, the Department of Energy, or the U.S. Government.

References

[1] "Public printing and documents," in *44 USC 3502*, ed. USA, 2009, p. 3542.

[2] ISO, "Iso guide 73:2009 risk management vocabulary," ed. Geneva, Switzerland: International Organization for Standardization, 2009.

[3] D. L. Pipkin, *Information security protecting the global enterprise*: Hewlett-Packard Company, 2000.

[4] G. Stoneburner, A. Goguen, and A. Feringa, "Risk management guide for information technology systems, recommendations of the national institute of standards and technology," National Institute of Standards and Technology, U.S. Department of Commerce, Gaitherburg, MD NIST Special Publication 800-30, 2002.

[5] C. Woody, J. Coleman, M. Fancher, C. Myers, and L. Young, "Applying octave: Practitioners report," Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA CMU/SEI-2006-TN-010, 2006.

[6] COBIT, "Control objectives for information and related technology," Governance, Control and Audit for Information and Related Technology. IT Governance Institute / ISACA / ISACF, 2011.

[7] (2005, Last Accessed October 27, 2010). *Inventory of risk management/risk assessment methods*. Available: http://rm-inv.enisa.europa.eu/rm_ra_methods.html

[8] (2005, Last Accessed October 27, 2010). *Comparison of risk management methods and tools*. Available: http://rm-inv.enisa.europa.eu/comparison.html

[9] "Executive order 13010," The White House, Washington, DC 1996.

[10] R. T. Marsh, "Critical foundations: Protecting america's infrastructures - the report of the president's commission on critical infrastructure protection," The White House, Washington, DC 1997.

[11] T. G. Lewis, *Critical infrastructure protection in homeland security: Defending a networked nation*. Hoboken, NJ John Wiley & Sons, Inc., 2006.

[12] "A roadmap for cybersecurity research," Department of Homeland Security, Washington, DC November 2009.

[13] A. Srivastava and J. P. Gupta, "New methodologies for security risk assessment of oil and gas industry," *Process Safety and Environmental Protection,* vol. 152 in press, p. 6, 2010.

[14] T. Sommestad, M. Ekstedt, and P. Johnson, "Cyber security risks assessment with bayesian defense graphs and architectural models," in *42nd Hawaii International Conference on System Sciences*, Waikoloa, Big Island, Hawaii, 2009, pp. 1-10.

[15] T. Sommestad, M. Ekstedt, and P. Johnson, "A probabilistic relational model for security risk analysis," *Computers & Security,* vol. 29, pp. 659-679, Sep 2010.

[16] M. R. Grimaila and A. Badiru, "A hybrid dynamic decision making methodology for defensive information technology contingency measure selection in the presence of cyber threats," *Operations Research: An International Journal,* 2011.

[17] "The comprehensive national cybersecurity initiative," The White House, Washington, DC 2010.

[18] F. T. Sheldon, R. K. Abercrombie, and A. Mili, "Methodology for evaluating security controls based on key performance indicators and stakeholder mission," in *Proceedings of 42nd Annual Hawaii International Conference on System Sciences (HICSS-42)*, Waikoloa, HI, 2009, pp. 1-10.

[19] R. K. Abercrombie, F. T. Sheldon, and A. Mili, "Synopsis of evaluating security controls based on key performance indicators and stakeholder mission value," in *11th IEEE High Assurance Systems Engineering Symposium (HASE '08)*, Nanjing, China, 2008, pp. 479-482.

[20] A. B. Aissa, R. K. Abercrombie, F. T. Sheldon, and A. Mili, "Quantifying security threats and their potential impacts: A case study," *Innovations in Systems and Software Engineering,* pp. 1-13, 2010.

[21] R. K. Abercrombie, F. T. Sheldon, and M. R. Grimaila, "A systematic comprehensive computational model for stake estimation in mission assurance," in *IEEE International Conference on Social Computing / IEEE International Conference on Privacy, Security, Risk and Trust*, Minneapolis, MN, 2010, pp. 1153-1158.

[22] A. B. Aissa, R. K. Abercrombie, F. T. Sheldon, and A. Mili, "Quantifying security threats and their impact," in *Proceedings of the 5th Annual Workshop on Cyber Security and Information Intelligence Research: Cyber Security and Information Intelligence Challenges and Strategies*, Oak Ridge, Tennessee, 2009, pp. 1-6.

[23] (2010, Last Accessed October 27, 2010). *Law of total probability*. Available: http://en.wikipedia.org/wiki/Law_of_total_probability

[24] A. Mili and F. T. Sheldon, "Challenging the mean time to failure: Measuring dependability as a mean failure cost," in *Proceedings of 42nd Hawaii International Conference on System Sciences (HICSS-42)*, Waikoloa, HI, 2009, p. 10.

[25] R. K. Abercrombie, F. T. Sheldon, and A. Mili, "Managing complex it security processes with valued based measures," in *2009 IEEE Symposium on Computational Intelligence in Cyber Security (CICS 2009)*, Nashville, TN, 2009, p. 7.

[26] FISMA, "Federal information security management act of 2002," in *44 U.S.C.*, ed. US, 2002, p. 3541.

[27] K. Stine, R. Kissel, W. C. Barker, J. Fahlsing, and J. Gulick, "Nist special publication 800-60 volume 1: Guide for mapping types of information and information systems to security categories " NIST Gaithersburg, MD 2008.

[28] FIPS, "Standards for security categorization of federal information and information systems," vol. FIPS PUB 199, ed. Gaithersburg, MD NIST, 2004, p. 13.

[29] FIPS, "Minimum security requirements for federal information and information systems," vol. FIPS PUB 200, ed. Gaithersburg, MD 2006, p. 17.

[30] S. Biffl, A. Aurum, B. W. Boehm, H. Erdogmus, and P. Gruenbacher, *Value based software engineering*. Berlin: Springer-Verlag, 2006.

[31] B. W. Boehm, "A view of 20th and 21st century software engineering," in *Proceedings of the 28th International Conference on Software Engineering*, Shanghai, China, 2006.

[32] R. P. Stanley, *Enumerative combinatorics* vol. 1 and 2: Cambridge University Press, 1999.

[33] T. C. Hu and M. T. Shing, "Computation of matrix chain products. Part ii," *SIAM J. Comput.,* vol. 13, pp. 228-251, 1984.

[34] A. N. Zakrzewska and E. M. Ferragut, "Modeling cyber conflicts using an extended petri net formalism," *unpublished,* p. 8, 2011.