

Mathematical Modeling of Cyber Resilience

Alexander Kott, Michael J. Weisman
U.S. Army Combat Capabilities Development Command
Army Research Laboratory
Adelphi, MD
{alexander.kott1, michael.j.weisman2}.civ@army.mil

Joachim Vandekerckhove
University of California, Irvine
Department of Cognitive Sciences
Irvine, CA
joachim@uci.edu

Abstract— We identify quantitative characteristics of responses to cyber compromises that can be learned from repeatable, systematic experiments. We model a vehicle equipped with an autonomous cyber-defense system and which also has some inherent physical resilience features. When attacked by malware, this ensemble of cyber-physical features (i.e., “bonware”) strives to resist and recover from the performance degradation caused by the malware’s attack. We propose parsimonious continuous models, and develop stochastic models to aid in quantifying systems’ resilience to cyber attacks.

I. INTRODUCTION

Resilience continues to gain attention as a key property of cyber and cyber-physical systems, for the purposes of cyber defense. Although definitions vary, it is generally agreed that cyber resilience refers to the ability of a system to resist and recover from a cyber compromise that degrades the mission-relevant performance of the system [1]. Resilience should not be conflated with risk or security [2].

To make the discussion more concrete, consider the example of a military ground logistics vehicle, possibly unmanned, which performs a mission of delivering heavy supplies along a difficult route. The adversary’s malware successfully gains access to the Controller Area Network (CAN bus) of the vehicle [3]. Then, the malware executes cyber attacks by sending a combination of messages intended to degrade the vehicle’s performance and diminish its ability to complete its delivery mission. We assume that the malware is at least partly successful, and the vehicle indeed begins to experience a degradation of its mission-relevant performance.

At this point, we expect the vehicle’s resilience-relevant elements to resist the degradation and then to

recover its performance to a satisfactory level, within an acceptably short time period. These “resilience-relevant elements” might be of several kinds. First, because the vehicle is a cyber-physical system, certain physical characteristics of the vehicles mechanisms will provide a degree of resilience. For example, the cooling system of the vehicle will exhibit a significant resistance to overheating even if the malware succeeds in misrepresenting the temperature sensors data. Second, appropriate defensive software residing on the vehicle continually monitors and analyzes the information passing through the CAN bus [4]. When the situation appears suspicious, it may take actions such as blocking or correcting potentially malicious messages. Third, it is possible that a remote monitoring center, staffed with experienced human cyber defenders, will detect a cyber compromise and will provide corrective actions remotely [5].

For the purposes of this paper, we assume that the remote monitoring and resilience via external intervention is impossible [6]. This may be the case if the vehicle must maintain radio silence for survivability purposes, or if the malware spoofs or blocks communication channels of the vehicle. Therefore, in this paper we assume that resilience is provided by the first two classes of resilience-relevant elements. Here, by analogy with malware, we call these “bonware” – a combination of physical and cyber features of the vehicle that serve to resist and recover from a cyber compromise.

A key challenge in the field of cyber resilience is quantifying or measuring resilience. Indeed, no engineering discipline achieved significant maturity without being able to measure the properties of phenomena relevant to the discipline [5]. Developers of systems like the notional vehicle in our example must be able to quantify the resilience of the vehicle under development in order to know whether the features they introduce in the vehicle improve its cyber resilience, or make it worse. Similarly, buyers of the vehicle need to know how to quantitatively specify and test resilience in order to determine whether the product meets their specifications.

In this paper, we report some of the results of a project called *Quantitative Measurement of Cyber Resilience*

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein. This work was partially funded by Cyber Technologies, Deputy CTO for Critical Technologies/Applied Technology, Office of the Under Secretary of Defense Research and Engineering. Dr. Vandekerckhove’s research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-21-2-0284.

(QMoCR) in which our research team seeks to identify quantitative characteristics of systems' responses to cyber compromises that can be derived from repeatable, systematic experiments. Briefly, we have constructed a test-bed in which a surrogate vehicle is subjected to controlled cyber attacks produced by malware. The vehicle is equipped with an autonomous cyber-defense system [6, 4] and also has some inherent physical resilience features. This ensemble of cyber-physical features (i.e., "bonware") strives to resist and recover from the performance degradation caused by the malware's attack. The test bed is instrumented in such a way that we can measure observable manifestations of this battle between the malware and bonware, especially the mission-relevant performance parameters of the vehicle.

The details of the test bed and the experiment are given in a companion paper [7]. The focus of this paper is different -- here we concentrate on constructing mathematical models that can be used to describe the dynamics of the malware-bonware battle. We seek models that are parsimonious in the number of empirical parameters and allow us to easily derive parameters of the model from experimental data.

The remainder of the paper is organized as follows. In the next section, we briefly describe prior work related to quantification of cyber resilience. We provide formal definitions of accomplishment and functionality. We propose a class of parsimonious models in which effects of both malware and bonware are approximated as deterministic, continuous differentiable variables, and we explore several variations of such models. In the following section we propose a different class of models -- stochastic models, and we show how this class is related to the previously proposed class of deterministic models. Then we show how these models are used to approximate experimental data obtained with our surrogate vehicle. We show how to determine the parameters of the models from experimental data. We discuss whether these parameters might be considered quantitative characteristics (i.e., measurements) of the bonware's cyber resilience.

II. PRIOR WORK

A growing body of literature explores quantification of resilience in general and cyber resilience in particular. Very approximately, the literature can be divided into two categories: (1) qualitative assessments of a system (actually existing or its design) by subject matter experts (SMEs) [8], [9] and (2) quantitative measurements based on empirical or experimental observations of how a system (or its high-fidelity model; [10]) responds to a cyber compromise [1, 11]. In the first category, a well-cited example is the approach called the cyber-resilience matrix [12]. In this approach, a system is considered as

spanning four domains: (1) physical (i.e., the physical resources of the system, and the design, capabilities, features and characteristics of those resources); (2) informational (i.e., the system's availability, storage, and use of information); (3) cognitive (i.e., the ways in which informational and physical resources are used to comprehend the situation and make pertinent decisions); and (4) social (i.e., structure, relations, and communications of social nature within and around the system). For each of these domains of the system, SMEs are asked to assess, and to express in metrics, the extent to which the system exhibits the ability to (1) plan and prepare for an adverse cyber incident; (2) absorb the impact of the adverse cyber incident; (3) recover from the effects of the adverse cyber incident; and (4) adapt to the ramifications of the adverse cyber incident. In this way, the approach defines a 4-by-4 matrix that serves as a framework for structured assessments by SMEs.

Another example within the same category (i.e., qualitative assessments of a system by SMEs) is a recent, elaborate approach proposed by [13]. The approach is called Framework for Operational Resilience in Engineering and System Test (FOREST), and a key methodology within FOREST is called Testable Resilience Efficacy Elements (TREE). For a given system or subsystem, the methodology requires SMEs to assess, among others, how well the resilience solution is able to (1) sense or discover a successful cyber-attack; (2) identify the part of the system that has been successfully attacked; (3) reconfigure the system in order to mitigate and contain the consequences of the attack. Assessment may include tests of the system, although the methodology does not prescribe the tests.

Undoubtedly, such methodologies can be very valuable in finding opportunities in improvements of cyber-resilience in a system that is either at the design stage or is already constructed. Still, these are essentially qualitative assessments, not quantitative measurements derived from an experiment.

In the second category (i.e., quantitative measurements based on empirical or experimental observations of how a system, or its high-fidelity model, responds to a cyber compromise), most approaches tend to revolve around a common idea we call here the area under the curve (AUC) method [14, 15]. In an experiment/test, a system is engaged into a performance of a representative mission, and then is subjected to an ensemble or sequence of representative cyber attacks. A mission-relevant quantitative functionality of the system is observed and recorded. The resulting average functionality, divided by normal functionality, can be used as a measure of resilience.

However, AUC-based resilience measures are inherently cumulative, aggregate measures, and do not tell us much about the underlying processes. For example, is it

possible to quantify the resilience impact of the bonware of the given system? Similarly, is it possible to quantify the impact of malware? In addition, is it possible to gain insights into how these values of impactfulness vary over time during an incident? We offer steps toward answering such questions.

III. QUANTITATIVE MEASURES OF CYBER RESILIENCE

Every mission has a goal, and we postulate that for a given mission, there exists a function $\mathcal{A}(t)$ that represents accomplishment and is cumulative from the mission start time up until the present time t . We define functionality, $F(t)$, to be the time derivative of mission accomplishment. Thus,

$$F(t) = \frac{d\mathcal{A}}{dt}, \quad \mathcal{A}(t) = \int_{t_0}^t F(\tau) d\tau. \quad (1)$$

The normal functionality, when the system performs normally and does not experience effects of a cyber attack, may, in general, vary with time. For simplicity, throughout this paper, we assume the normal functionality to be constant in time, $F_N(t) = F_N$. Thus, the functionality of our system prior to an attack or other malfunction is normal at the start time: $F(t_0) = F_N$.

In the following sections we develop models for the behavior of a system's functionality over the course of a mission during which it is being attacked by malware and defended by bonware.

In the first set of models, we assume that there is an observable, sufficiently smooth function representing mission accomplishment, and we define functionality to be its time derivative. Then, we motivate a parsimonious model for the differential equation governing functionality, give the general solution, and discuss a few specific cases. We then develop a stochastic model and show its relationship to the continuous model. Finally, we show, both analytically and empirically, that with sufficiently many instantiations, the average of an ensemble of stochastic curves will approximate the solution to the continuous model differential equation.

IV. CONTINUOUS MODEL

For the first set of models, we make the assumption that mission accomplishment is twice continuously differentiable: $\mathcal{A} \in C^2$, and thus $F \in C^1$.

As a first approximation, we let the impact of malware on the derivative of functionality be linear. The impact of bonware is similarly defined and proportional to the level of functionality below normal. Malware degrades the system while bonware aims to increase functionality over time. Malware impact and bonware impact are assumed to be continuous functions of time, $\mathcal{M}, \mathcal{B} \in C^0$. The

impact on functionality is the sum of the impacts of malware and bonware, and

$$\frac{dF}{dt} + \mathcal{Q}(t)F(t) = F_N\mathcal{B}(t), \quad (2)$$

where $\mathcal{Q}(t) = \mathcal{M}(t) + \mathcal{B}(t)$. Since we expect bonware to help (or at least not harm) and malware to not help, we assume $\mathcal{B}(t) \geq 0$ and $\mathcal{M}(t) \geq 0$. We also assume normal functionality is positive, $F_0 > 0$, and functionality is always nonnegative and less than or equal to normal functionality, $0 \leq F(t) \leq F_N$. This first-order linear differential equation has the following solution:

$$F(t) = e^{-\int_0^t \mathcal{Q}(p) dp} \left(F(0) + F_N \int_0^t e^{\int_0^\tau \mathcal{Q}(p) dp} \mathcal{B}(\tau) d\tau \right). \quad (3)$$

To help us understand how the model works, we find explicit solutions for a number of examples.

A. Constant model

Assuming \mathcal{M}, \mathcal{B} , and \mathcal{Q} are constant, we have

$$\frac{dF}{dt} + \mathcal{Q}F(t) = F_N\mathcal{B}. \quad (4)$$

1) *No bonware*: If $\mathcal{B} = 0$, then Equation 4 reduces to $\frac{dF}{dt} + \mathcal{M}F(t) = 0$ and $F(t) = F(0)e^{-\mathcal{M}t}$. If also $\mathcal{M} = 0$ (no bonware and no malware), then $\frac{dF}{dt} = 0$ and $F(t) = F(0)$.

2) *Bonware*: With bonware present, the solution is

$$F(t) = \left[F(0) - \frac{F_N\mathcal{B}}{\mathcal{Q}} \right] e^{-\mathcal{Q}t} + \frac{F_N\mathcal{B}}{\mathcal{Q}}. \quad (5)$$

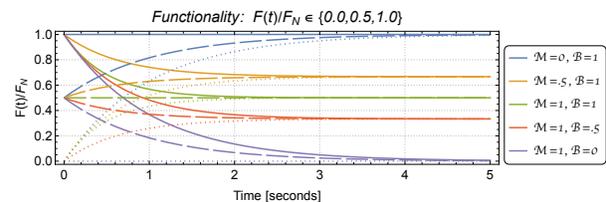


Figure 1. Normalized functionality, $F(t)/F_N$, is shown for various values of \mathcal{M} (malware attacking) and \mathcal{B} (bonware defending) and at initial conditions $F(0)/F_N \in \{0.0, 0.5, 1.0\}$. The functionality over time depends on the relative strengths of bonware and malware and on the initial condition. When the system initially is at normal functionality and when malware overpowers bonware, functionality exhibits exponential decay. When functionality initially is low, and when bonware overpowers malware, the system recovers (via Eq. 5) to $\frac{F_N\mathcal{B}}{\mathcal{M}+\mathcal{B}}$.

If $F(0) > F_N\mathcal{B}/\mathcal{Q}$, then $F(t)$ will initially, at time $t = 0$, be at $F(0)$ and decrease to $F_N\mathcal{B}/\mathcal{Q}$. If $F(0) < F_N\mathcal{B}/\mathcal{Q}$, then the function $F(t) = F(0)$ will be constant. If $F(0) > F_N\mathcal{B}/\mathcal{Q}$, the function will start at $F = F(0)$ and increase to $F_N\mathcal{B}/\mathcal{Q}$. Examples of these situations are shown in Figure 1. The plot of ($\mathcal{M} > 0$) in Figure 1 shows that even in the presence of bonware, malware will still have an impact on the system. The steady-state

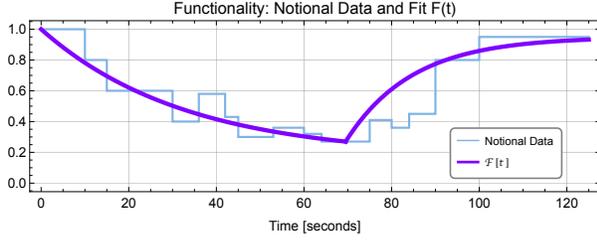


Figure 2. The smooth line is an example functionality curve with piecewise constant malware and bonware impacts. The notional data and piecewise constant model fit are described below in Section VI-A.

of the system is obtained either by setting $\frac{dF}{dt} = 0$ in Equation 4 or letting $t \rightarrow \infty$:

$$F_{\infty} = \lim_{t \rightarrow \infty} F(t) = F_N \frac{\mathcal{B}}{\mathcal{M} + \mathcal{B}} \quad (6)$$

so that the antidote to malware is to overwhelm it with bonware. The exponent, $-Qt = (-\mathcal{M} - \mathcal{B})t$ in the solution given by Equation 5 indicates that increasing the impact of either malware or bonware will cause the system to more quickly approach steady-state. At steady-state,

$$\frac{F_N - F_{\infty}}{F_{\infty}} = \frac{\mathcal{M}}{\mathcal{M} + \mathcal{B}}. \quad (7)$$

Equation 7 gives us further insight into the trade-off between impacts of both malware and bonware. The relative decrease of the function from normal functionality is equal to the ratio of malware impact to the sum of malware and bonware impacts.

B. Piecewise constant model

If either malware's or bonware's impact diminishes at some point in the incident, the model may switch from one set of constants defining malware and bonware to another set of constants. The differential equation (Eq. 2) may now be expressed as

$$\frac{dF}{dt} = \sum_{j=0}^{N-1} (F_N - F(t))\mathcal{B}_j(t) - F(t)\mathcal{M}_j(t), \quad (8)$$

where the vectors $\mathcal{M} = (\mathcal{M}_0, \mathcal{M}_1, \dots, \mathcal{M}_{N-1})$ and $\mathcal{B} = (\mathcal{B}_0, \mathcal{B}_1, \dots, \mathcal{B}_{N-1})$ contain the malware impact and bonware impacts within time windows whose end points are defined by $\{t_0, t_1, \dots, t_N\}$. The solution will be a function which, in each time interval, is the solution found in Equation 5. The purple curve in Figure 2 is a realization of this model.

C. Linear model

The impacts of malware and bonware may also be linear functions of t , so that $\mathcal{M}(t) = \nu - \mu t$, $\mathcal{B}(t) = \alpha - \beta t$, and $Q(t) = \lambda - \omega t$, where $\lambda = \alpha + \nu$ and $\omega = \beta + \mu$. Under this linear model, Equation 2 becomes

$$\frac{dF}{dt} + (\lambda - \omega t)F(t) = F_N(\alpha - \beta t) \quad (9)$$

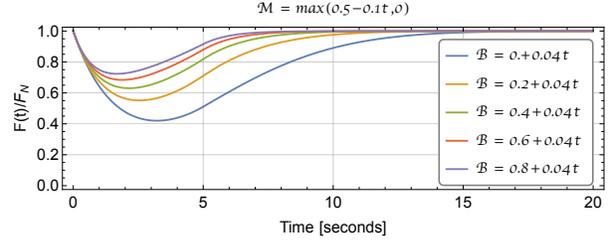


Figure 3. Normalized functionality, $F(t)/F_N$, for piecewise linear models. Impacts of malware and bonware are linear functions of time.

The solution can be expressed in terms of the error function $\text{erf}(z) = \frac{2}{\sqrt{\pi}} \int_0^z e^{-\tau^2} d\tau$:

$$\begin{aligned} \frac{F(t)}{F_N} = \frac{1}{\Omega(t)} & \left\{ \frac{F(0)}{F_N} - \frac{\beta}{\omega} (1 - \Omega(t)) + (\alpha\omega - \beta\lambda) \right. \\ & \left. \times \frac{\sqrt{\frac{\pi}{2}} e^{\Lambda^2}}{\omega^{3/2}} \left[\text{erf}(\Lambda) + \text{erf}\left(\frac{\omega t}{\sqrt{2\omega}} - \Lambda\right) \right] \right\} \end{aligned} \quad (10)$$

where $\Omega(t) = e^{\lambda t - \frac{1}{2}\omega t^2}$, and $\Lambda = \lambda/\sqrt{2\omega}$.

D. Piecewise linear model

Both malware and bonware impacts may initially be linear, but if the situation changes and a different linear model holds after a time, the model should be able to account for it. In particular, if malware impact is decreasing over time, at some point we will reach $\mathcal{M} = 0$ and the model switches to a new linear model. Equation 9 can be written

$$\frac{dF}{dt} = \sum_{j=0}^{N-1} [(\lambda_j - \omega_j t)F(t) - F_N(\alpha_j - \beta_j t)].$$

The solution follows from Equation 10. Example realizations of the piecewise linear models are shown in Figure 3. The shapes of the curves resemble experimental data discussed in [7].

V. STOCHASTIC DIFFERENTIAL EQUATION MODEL

In this section, we extend the previously defined differential equation (DE) model to a *stochastic* differential equation (SDE) model. The extension is motivated by the discontinuous nature of the notional data in Figure 2. Whereas the DE model assumed a smooth functionality curve, our stochastic version allows for a more punctuated attack-and-restoration pattern.

In the SDE model, both malware and bonware may be active at random (or a priori unknown) times with random (or a priori unknown) effectiveness. Let malware activity $A^m(t) \in \{0, 1\}$ indicate whether malware was successful at time t , bonware activity $A^b(t) \in \{0, 1\}$ indicate whether bonware was successful at time t , malware effectiveness $E^m(t) \in [0, 1)$ express the proportion of functionality reduced by the malware's success at time

t , and bonware effectiveness $E^b(t) \in [0, 1]$ express the proportion of damage undone by the bonware's success at time t . The SDE analog to Equation 2 is then

$$\frac{dF}{dt} = (F_N - F(t)) A^b(t) E^b(t) - F(t) A^m(t) E^m(t). \quad (11)$$

Rather than changing deterministically over time, these model parameters are assumed to vary stochastically according to these distributions:

$$A^m(t) \sim \text{Bern}(\theta^m(t)), \quad (12)$$

$$A^b(t) \sim \text{Bern}(\theta^b(t)), \quad (13)$$

$$E^m(t) \sim \text{Unif}(0, \gamma^m(t)), \quad (14)$$

$$E^b(t) \sim \text{Unif}(0, \gamma^b(t)), \quad (15)$$

where $\text{Bern}(\theta)$ indicates the Bernoulli distribution with rate θ and $\text{Unif}(0, \gamma)$ indicates a uniform distribution with lower bound 0 and upper bound γ . Hence, $\theta^m(t) \in [0, 1]$ is the probability that malware is successful at time t , $\theta^b(t) \in [0, 1]$ is the probability that bonware is successful at time t , $\gamma^m(t) \in (0, 1]$ is the maximum fraction of damage inflicted by malware, and $\gamma^b(t) \in (0, 1]$ is the maximum fraction of damage undone by bonware.

Like the ordinary differential equation (ODE) model, the SDE model allows for a number of interesting variants. In the remainder of this section, we introduce some useful simplifications and extensions.

A. Constant parameters

In the simplest version of the SDE model, we assume that the rate of malware attacks and their maximum efficiency are constant in time, and that the rate of bonware restoration and its maximum efficiency are both constant in time. Specifically, under this model it is assumed that $\theta^m(t) = \theta^m$, $\theta^b(t) = \theta^b$, $\gamma^m(t) = \gamma^m$, and $\gamma^b(t) = \gamma^b$ (cf. Eqs. 12–15).

This version of the model is parsimonious, with only four free parameters, each with a useful interpretation.

B. Piecewise constant parameters

In a first extension of the SDE model, we assume that malware and bonware have an activity rate of 0 until they activate at time points t^m and t^b , respectively. After activation, both activity rates are constant. Additionally, both efficiency parameters are assumed to be constant in time. Specifically, $\theta^m(t) = \theta^m u(t - t^m)$, $\theta^b(t) = \theta^b u(t - t^b)$, $\gamma^m(t) = \gamma^m$, and $\gamma^b(t) = \gamma^b$. Here $u(\cdot)$ is the unit step function.

The piecewise constant formulation here is particularly apposite for our experiments, in which the timing of malware attacks is known. The analyst can then choose between considering t^m and t^b as known, and retain the parsimony of a four-parameter model, or considering them unknown and estimate them from data in a six-parameter model. The difference between the known

onset time of an attack and the estimated t^m may then be interpreted as the time it takes for an attack to take effect. Similarly, the difference between that onset time and the estimated t^b may be interpreted as the delay until bonware begins to restore functionality after an attack.

C. Parameter expansion of the SDE model

The SDE model can be conveniently stated as a hidden Markov model and implemented as a directed acyclic graph [16] for efficient parameter estimation with a general-purpose Bayesian inference engine (e.g., JAGS; [17]). Our implementation relied on a sequential definition for the likelihood function: $(F(t+1) | F(t), \dots) \sim \text{Unif}(L(t), U(t))$, with $L(t) = F(t) - A^m(t) E^m(t) F(t)$ and $U(t) = F(t) + A^b(t) E^b(t) (1 - F(t))$.

Since this likelihood function depends on the unknown stochastic parameters $A^m(t)$, $A^b(t)$, $E^m(t)$, and $E^b(t)$, we applied a parameter expansion approach [18, 19] using Equations 12–15.

D. Relationship between continuous and SDE model

With the parameters of the stochastic model selected appropriately, we show that as the number of stochastic realizations increases, the expectation of the solution to the stochastic differential equation model approaches that of the ODE model. We show this for the simple constant parameter case. The general result follows by extension.

Theorem. *Let $y_k^m \sim \text{Bern}(2\mathcal{M})$, $y_k^b \sim \text{Bern}(2\mathcal{B})$, $z_k^m \sim \text{Unif}(0, F_k)$, $z_k^b \sim \text{Unif}(0, F_N - F_k)$, and*

$$F_{k+1} = F_k - y_k^m z_k^m + y_k^b z_k^b, \quad (k = 1, \dots, K). \quad (16)$$

Let $\mathcal{F}_{kn} = \frac{F_{k,j}}{n}$, ($j = 1, \dots, n$), then

$\mathcal{F}_k = \mathbb{E}(F_k) = \lim_{n \rightarrow \infty} \mathcal{F}_{kn}$ and $\mathcal{F}_k \approx F(k)$, for large k , where $F(t)$ is the solution to the initial value problem given by Equation 4 with $F(0) = \mathcal{F}_0$.

Proof. Take the expectation of Equation 16. Then $\mathcal{F}_k - \mathcal{F}_{k-1} + (\mathcal{M} + \mathcal{B})\mathcal{F}_{k-1} = F_N \mathcal{B}$. With $\mathcal{F}_{0n} = \mathcal{F}_0$, the solution is $\mathcal{F}_k = [\mathcal{F}_0 - \frac{F_N \mathcal{B}}{\mathcal{Q}}] (1 - \mathcal{Q})^k + \frac{F_N \mathcal{B}}{\mathcal{Q}}$, which approximates Equation 5 for large k . \square

VI. AN APPLICATION OF THE MODEL

A. Obtaining model parameters

Given notional data that represents a typical curve of functionality over the course of an incident where malware and bonware are active, we develop a fast method to estimate the continuous model parameters for a curve that approximates the data, and use these parameters to generate further realizations based on this model. In Figure 2 an example of such notional data is plotted (in light blue). In this section, we illustrate our method to extract the model parameters from this curve.

The set $P = \{t_0, \dots, t_K\}$ partitions the mission timeline and malware and bonware are constant in each interval (t_{i-1}, t_i) , $i = 1, \dots, K$. In each interval, $Q_i = \mathcal{M}_i + \mathcal{B}_i$ and the differential equation governing Continuous Model I is $\frac{dF(t)}{dt} + Q_i F(t) = F_N(t) \mathcal{B}_i$. Thus, in each interval (t_{i-1}, t_i) , the solution is

$$F(t) = \left[F(t_{i-1}) - \frac{F_N \mathcal{B}_i}{Q_i} \right] e^{-Q_i(t-t_{i-1})} + \frac{F_N \mathcal{B}_i}{Q_i}.$$

We compute the effectiveness, E^m , and activity, A^m , of malware and of bonware (E^b, A^b), in each interval: $\mathcal{M}_i = E_i^m A_i^m$, $\mathcal{B}_i = E_i^b A_i^b$.

We observe that there is a unique switching time t^* where the functionality's trend reverses, and thus we take $K = 2$. Before the switch, the impact of malware is greater than that of bonware. From the time of the switch until the end of the mission, bonware is stronger. To estimate the switching time t^* , we find the minimum of the data to occur over the interval from 64 s to 75 s. There, the minimum value of the data curve is $m = 0.27$. Taking the midpoint, our estimate for t^* is 69.5 s. We estimate the activity of malware before switching to be the number of times the data curve decreases divided by the switching time. Similarly, our estimate of bonware activity is the number of times the data curve increases prior to the switching time. We thus have $A_1^m \approx 7/69.5 \approx 0.101$ and $A_1^b \approx 2/69.5 \approx 0.014$.

To determine the remaining parameters, we numerically solve this system of equations:

$$\begin{aligned} \alpha m &= F_N \frac{\mathcal{B}_1}{Q_1}, \\ m &= F(0) - F_N \frac{\mathcal{B}_1}{Q_1} e^{-Q_1 t^*} + F_N \frac{\mathcal{B}_1}{Q_1}. \end{aligned}$$

The first equation says that where the curve meets the minimum of the data, it has experienced exponential decay of α toward the asymptotic minimum. We take α to be $\alpha = 1 - 1/e$. The second equation says that the minimum occurs at the switching time (the time when the model switches from malware dominating bonware, to bonware dominating malware). Solving this system of equations yields (with $\mathcal{M}_1 = Q_1 - \mathcal{B}_1$), $\mathcal{M}_1 \approx 0.025$ and $\mathcal{B}_1 \approx 0.005$, so that $E_1^m = \mathcal{M}_1/A_1^m \approx 0.503$ and $E_1^b = \mathcal{B}_1/A_1^b \approx 0.362$.

To the right of t^* , we fit an exponentially increasing function. Similar to before the switching time, we compute the activities of the malware and bonware: $A_2^m \approx \frac{1}{100-69.5} \approx 0.033$ and $A_2^b \approx \frac{4}{100-69.5} \approx 0.131$.

To determine the remaining parameters, we numerically solve this system of equations:

$$\begin{aligned} \zeta &= \frac{F(0) \mathcal{B}_2}{Q_2}, \\ \tilde{\alpha} \zeta &= \left(m - \frac{F_N \mathcal{B}_2}{Q_2} \right) \left(e^{-Q_2(125-t^*)} + \frac{F_N \mathcal{B}_2}{Q_2} \right). \end{aligned}$$

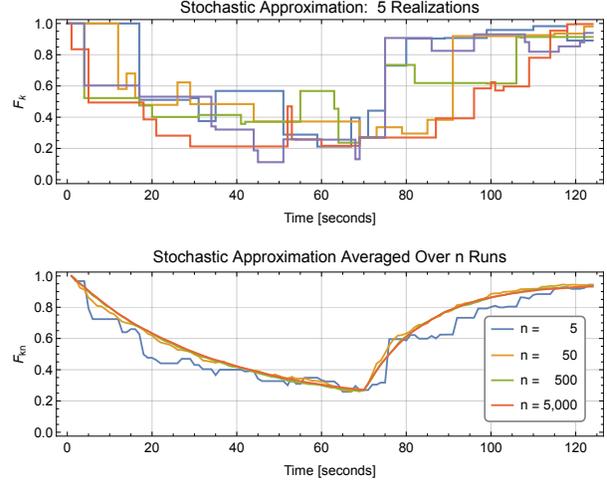


Figure 4. (Top) Five realizations of the stochastic model generated with the parameters obtained by fitting the notional data shown in Figure 2. Each realization is different but each roughly follows an exponential decay when $t < t^*$ and an exponential recovery for $t \geq t^*$. (Bottom) Averages of n stochastic runs for $n \in \{5, 50, 500, 5000\}$. As n increases, the average of the ensemble approaches the fitted curve as predicted in the theorem of Section V-D.

We have found that $\tilde{\alpha} = 1 - e^{-4}$ and $\zeta = 0.95$ are satisfactory values to use for these hyperparameters.

We compute $\mathcal{M}_2 \approx 0.005$ and $\mathcal{B}_2 \approx 0.088$, so that $E_1^m = \mathcal{M}_1/A_1^m \approx 0.201$ and $E_1^b = \mathcal{B}_1/A_1^b \approx 0.957$.

B. Generating stochastic realizations

Using the parameters found in Section VI-A, we can now generate stochastic realizations. In Section V-D, we showed that for large sample sizes, the average of our ensemble will approach the solution to the continuous ODE model. We show empirically that this is indeed the case. To illustrate, we generated five realizations (Figure 4) of the stochastic model with parameters found from the notional data of Section VI-A. By averaging n curves when $n \in \{5, 50, 500, 5000\}$ we see how the ensemble average approaches the solution of the corresponding differential equation as predicted in the Theorem of Section V-D.

VII. DISCUSSION AND CONCLUSION

We have presented a broadly applicable framework for the analysis of the cyber resilience of military artifacts. Our framework relies on the construction of a custom differential equation time series model that shows good qualitative correspondence to the functionality of vehicles performing missions. Seeking to move beyond the use of area-under-the-curve quantifications of cyber resilience, our proposed models have the advantage that their parameters have domain-relevant interpretations such as *the activity of malware* and *the effectiveness*

of *bonware*. Such interpretable parameters can provide a more nuanced interpretation of cyber resilience data being experimentally obtained in our lab.

Our formal models come in two families with complementary advantages. A series of continuous models is parsimonious, mathematically convenient, and easy to fit. A series of discrete, stochastic models shows greater verisimilitude but is slightly less parsimonious and requires more computationally onerous parameter estimation techniques.

Both types of models can be extended to a large variety of custom circumstances, including the case where model parameters change gradually, abruptly, or predictably as a result of experimental manipulation. Future work will include an extension to the cases of multiple simultaneous objectives and to the case of multiple vehicles to be analyzed jointly.

REFERENCES

- [1] A. Kott and I. Linkov. *Cyber resilience of systems and networks*. Springer International Publishing, New York, NY, 2019.
- [2] I. Linkov, B. D. Trump, and J. Keisler. Risk and resilience must be independently managed. *Nature*, 555:7694, 2018.
- [3] M. Bozdal, M. Samie, and I. Jennions. A survey on CAN bus protocol: Attacks, challenges, and potential solutions. In *2018 International Conference on Computing, Electronics & Communications Engineering*, pages 201–205. IEEE, August 2018.
- [4] A. Kott, P. Théron, M. Drašar, E. Dushku, B. LeBlanc, P. Losiewicz, ..., and K. Rządca. Autonomous intelligent cyber-defense agent (aica) reference architecture, 2018. arXiv:1803.10664.
- [5] A. Kott, M. S. Golan, B. D. Trump, and I. Linkov. Cyber resilience: by design or by intervention? *Computer*, 54(8):112–117, 2021.
- [6] A. Kott and P. Théron. Doers, not watchers: Intelligent autonomous agents are a path to cyber resilience. *IEEE Security & Privacy*, 18(3):62–66, 2020.
- [7] J. Ellis, T. Parker, J. Vandekerckhove, B. Murphy, S. Smith, A. Kott, and M. Weisman. Experimental infrastructure for study of measurements of resilience. *Proceedings of IEEE Military Communications Conference*, pages 841–846, Dec. 2022.
- [8] A. Alexeev, D. Henshel, K. Levitt, P. McDaniel, Rivera B., S. Templeton, and M. Weisman. Constructing a science of cyber-resilience for military systems. *NATO IST-153 Workshop on Cyber Resilience*, pages 23–25, 2017.
- [9] D. S. Henshel, K. Levitt, S. Templeton, M. G. Cains, A. Alexeev, B. Blakely, P. McDaniel, G. Wehner, J. Rowell, and M. Weisman. The science of cyber resilience: Characteristics and initial system taxonomy. *Fifth World Conference on Risk*, 2019.
- [10] A. Kott, J. Ludwig, and M. Lange. Assessing mission impact of cyberattacks: Toward a model-driven paradigm. *IEEE Secur Privacy*, 15(5):65–74, January 2017.
- [11] A. K. Ligo, A. Kott, and I. Linkov. How to measure cyber-resilience of a system with autonomous agents: Approaches and challenges. *IEEE Engineering Management Review*, 49(2):89–97, 2021.
- [12] I. Linkov, D. A. Eisenberg, K. Plourde, T. P. Seager, J. Allen, and A. Kott. Resilience metrics for cyber systems. *Environ. Syst. Decis.*, 33(4):471–476, 2013.
- [13] Peter Beling, Barry Horowitz, and Tom McDermott. Developmental test and evaluation (dte&a) and cyberattack resilient systems. *TR SERC-2021-TR-015 (V2)*, September 2021.
- [14] S. Hosseini, K. Barker, and J. E. Ramirez-Marquez. A review of definitions and measures of system resilience. *Reliability Engineering & System Safety*, 145:47–61, 2016.
- [15] A. Kott and I. Linkov. To improve cyber resilience, measure it. *Computer*, 54(2):80–85, Feb. 2021.
- [16] T. L. Griffiths, C. Kemp, and J. B. Tenenbaum. *Bayesian models of cognition*, pages 59–100. Cambridge University Press, Cambridge, MA, 2008.
- [17] M. Plummer. JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In K. Hornik, F. Leisch, and A. Zeileis, editors, *Proceedings of the 3rd International Workshop on Distributed Statistical Computing*. DSC 2003, Vienna, Austria, 2003.
- [18] C. Liu, D. B. Rubin, and Y. N. Wu. Parameter expansion to accelerate em: the px-em algorithm. *Biometrika*, 85(4):755–770, 1998.
- [19] A. Gelman. Parameterization and Bayesian modeling. *Journal of the American Statistical Association*, 99:537–545, 2004.