A Survey on Security Metrics[1]

MARCUS PENDLETON, The University of Texas at San Antonio
RICHARD GARCIA-LEBRON, The University of Texas at San Antonio
SHOUHUAI XU, The University of Texas at San Antonio

The importance of security metrics can hardly be overstated. Despite the attention that has been paid by the academia, government and industry in the past decades, this important problem stubbornly remains open. In this survey, we present a survey of knowledge on security metrics. The survey is centered on a novel taxonomy, which classifies security metrics into four categories: metrics for measuring the system *vulnerabilities*, metrics for measuring the *defenses*, metrics for measuring the *threats*, and metrics for measuring the *situations*. The insight underlying the taxonomy is that situations (or outcomes of cyber attack-defense interactions) are caused by certain threats (or attacks) against systems that have certain vulnerabilities (including human factors) and employ certain defenses. In addition to systematically reviewing the security metrics that have been proposed in the literature, we discuss the gaps between the state of the art and the ultimate goals.

1. INTRODUCTION

Security metrics is one of the most important open problems in security research. It has been recognized on the Hard Problem List of the United States INFOSEC Research Council (both 1999 and 2005 editions) [Council 2007], has been reiterated in 2011 by the United States National Science and Technology Council [Science and Council 2011], and most recently has been listed as one of the five hard problems in Science of Security (August 2015) [Nicol et al.].

The security metrics problem certainly has received a lot of attention, including government and industry bodies [Chew et al. ; (IATAC) 2009; Institute ; for Internet Security 2010]. For example, the United States National Institute of Standards and Technology proposed three categories of security metrics—*implementation*, *effectiveness*, and *impact* [Chew et al.]; the Center for Internet Security defined 28 security metrics in another three categories—*management*, *operational*, and *technical* [for Internet Security 2010]. However, these efforts are almost exclusively geared towards cyber defense administrations and operations. They neither discuss how the security metrics may be used as parameters in security modeling (i.e., theoretical use of security metrics), nor discuss what the gaps are between the state-of-the-art and the ultimate goals and how these gaps may be bridged. This motivates us to survey the knowledge in the field, while hoping to shed some light on the difficulties of the problem and the directions for future research. To the best of our knowledge, this is the first survey of security metrics, despite that there have been some efforts with a much narrower focus (e.g., [Landwehr et al. 1994; Chandola et al. 2009; Milenkoski et al. 2015; Roundy and Miller 2013; Ugarte-Pedrero et al. 2015]).

The paper is organized as follows. Section 2 discusses the scope and methodology of the survey. Section 3 describes security metrics for measuring system vulnerabilities. Section 4 reviews security metrics for measuring defenses. Section 5 presents security metrics for measuring threats. Section 6 describes security metrics for measuring situations. Section 7 discusses the gaps between the state-of-the-art and the security metrics that are desirable. Section 8 concludes the paper.

[1]Author's addresses: M. Pendleton and R. Lebron-Garcia and S. Xu, Department of Computer Science, The University of Texas at San Antonio. Correspondence: Shouhuai Xu (shxu@cs.utsa.edu)

2. SCOPE AND METHODOLOGY

2.1. Terminology

The term *security metrics* has a range of meanings, with no widely accepted definition [Jansen 2009]. It is however intuitive that security metrics reflect some security attributes quantitatively.

Throughout the paper, the term *systems* is used in a broad sense, and is used in contrast to the term *building-blocks*, used to describe concepts such as cryptographic primitives. The discussion in the present paper applies to two kinds of systems: (i) *enterprise systems*, which include networked systems of multiple computers/devices (e.g., company networks), clouds, and even the entire cyberspace, and (ii) *computer systems*, which represent individual computers/devices. This distinction is important because an enterprise system consists of many computers/devices, and measuring security of an enterprise system naturally requires to measure security of the individual computers.

The term *attacking computer* represents a computer or IP address from which cyber attacks are launched against others, while noting that the attacking computer itself may be a compromised one (i.e., not owned by a human attacker). The term *incident* represents a successful attack (e.g., malware infection or data breach).

For applications of security metrics, we will focus on two uses. The *theoretical use* is to incorporate security metrics as parameters into some security models that may be built to understand security from a more holistic perspective. There have been some initial studies in pursuing such models, such as [LeMay et al. ; Xu 2014a], which often aim to characterize the evolution of the global security state. The *practical use* is to guide daily security practice, such as comparing the security of two systems and comparing the security of one system during two different periods of time (e.g., last year vs. present year).

2.2. Scope

We have to limit the scope of the literature that is surveyed in the present paper. This is because *every* security paper that improves upon a previous result—be it a better defense or more powerful attack—would be considered relevant in terms of security metrics. However, most security publications did *not* address the security metrics perspective, perhaps because it is sufficient to show, for example, a newly proposed defense can defeat an attack that could not be defeated by previous defenses. This suggests us to survey the literature that made a reasonable effort at defining security metrics. This selection criterion is certainly subjective, but we hope the readers find the resulting survey and discussion informative. It is worth mentioning that our focus is on security metrics, rather than the specific approaches for analyzing them. We treat the analysis approaches as an orthogonal issue because a security metric may be analyzed via multiple approaches.

Even within the scope discussed above, we still need to narrow down our focus. This is because security, and security metrics thereof, can be discussed at multiple levels of abstractions, including systems and building-blocks as mentioned above. For building-blocks, great success has been achieved in measuring the *concrete security* of cryptographic primitives [Bellare et al.], while other notable results include metrics for measuring privacy [Dwork ; Shokri et al.], information flow [Mardziel et al. 2014], side-channel leakage [Schneider and Moradi 2015], and hardware security [Rostami et al. 2014]. On the other hand, our understanding of security metrics for measuring security of *systems* lags far behind, as the present paper shows. One thing that is worth clarifying is that the exposure of cryptographic keys, due to the use of weak randomness in the key generation algorithm or Heartbleed-like at-

tacks, is treated as a systems security problem. This is plausible because the formal framework for analyzing cryptographic security assumes that the cryptographic keys are not exposed.

The aforementioned discrepancy between the metrics for systems security and the metrics for building-blocks security is unacceptable because the former is often needed and used in the process of business decision-making. This suggests us to focus on systemizing the underdeveloped field of systems security metrics. The importance of this underdeveloped field can be seen by efforts that have been made by government and industrial bodies [Chew et al. ; (IATAC) 2009; Institute ; for Internet Security 2010]. This prompts us to consider both these metrics and those that appeared in academic venues.

2.3. Survey methodology

Our survey methodology is centered on the perspective of cyber attack-defense interactions, which applies to both *enterprise systems* and *computer systems* mentioned above.

Figure 1 illustrates a snapshot of an *enterprise system*. At time t, the enterprise system consists of n computers (or devices, virtual machines in the case of cloud), denoted by the vector $C(t) = \{c_1(t), \ldots, c_n(t)\}$, where n could vary with time t (i.e., n could be a function of time t). Each computer, $c_i(t)$, may have a vector $v_i(t)$ of vulnerabilities, including the computer user's vulnerability to social-engineering attacks, the vulnerability caused by the use of weak passwords, and the software vulnerabilities that may include some zero-day and/or some unpatched ones. Attacks are represented by red arrows, and defenses are represented by blue bars. Defenses accommodate both the defenses that are installed on the individual computers (e.g., anti-malware tools) and the defenses that are employed at the perimeter of the enterprise system (e.g., firewalls). The thickness of red arrows and blue bars reflect the attack and defense power, respectively. Some attacks penetrate through the defenses (e.g., attacks against computer $c_n(t)$), while others fail (e.g., attacks against computer $c_1(t)$). The outcome of the attack-defense interaction at time t is reflected by a security state vector $S(t) = \{s_1(t), \ldots, s_n(t)\}$, where $s_i(t) = 0$ means computer $c_i(t)$ is secure at time t and $s_i(t) = 1$ means computer $c_i(t)$ is compromised at time t. However, the defender's observation of the security state vector $S(t)$, denoted by $O(t) = \{o_1(t), \ldots, o_n(t)\}$, may not be perfect because of false-positive, false-negative, or noise.

Figure 2 illustrates a snapshot of a *computer system* $c_i(t)$ at time t, where we also use blue bars to represent defenses, use red arrows to represent attacks, and use their thickness to reflect their defense/attack power. At a high level, $c_i(t)$ may have a range of vulnerabilities, which correspond to $v_i(t)$ in Figure 1. The vulnerabilities include the computer user's vulnerability (or susceptibility) to social-engineering attacks, the vulnerability caused by the use of weak passwords, and software vulnerabilities. The defense may include (i) the use of some *filtering* mechanisms that are deployed at the enterprise system perimeter to block (for example) traffic from malicious or blacklisted IP addresses, (ii) the use of some *attack detection* mechanisms to detect and block attacks before they reach computer $c_i(t)$, and (iii) the use of some proactive defense mechanisms (e.g., address space randomization) to try to prevent the exploitation of some vulnerabilities. Suppose the attacker has a vector of 12 attacks. Attack 1 successfully compromises $c_i(t)$ because the user is lured into (e.g.) clicking a malicious URL. Attack 4 successfully compromises $c_i(t)$ because the password in question is correctly guessed. Attacks 6 and 7 successfully compromise $c_i(t)$ because they exploit a zero-day vulnerability, despite the possible employment of proactive defense mechanisms on $c_i(t)$. Attack 9 successfully compromises $c_i(t)$ because the vulnerability is unpatched and the attack is neither filtered nor detected. Attack 12 successfully compromises $c_i(t)$ because the cryptographic key in question is exposed by (e.g.,) Heartbleed-like attacks
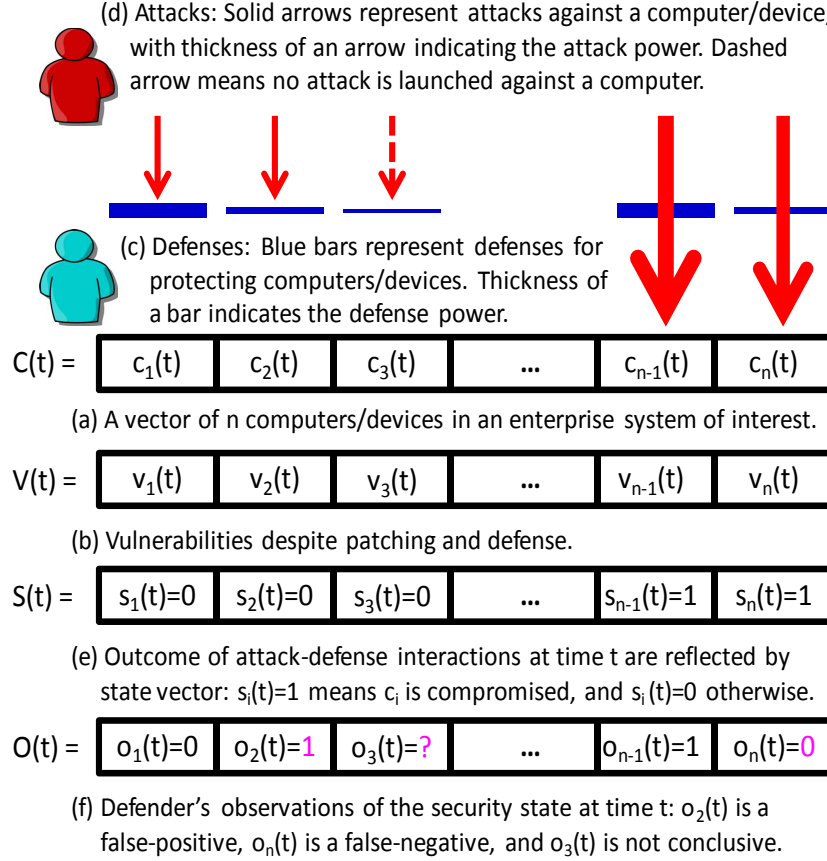
(d) Attacks: Solid arrows represent attacks against a computer/device, with thickness of an arrow indicating the attack power. Dashed arrow means no attack is launched against a computer.

(c) Defenses: Blue bars represent defenses for protecting computers/devices. Thickness of a bar indicates the defense power.

$C(t) = $ | $c_1(t)$ | $c_2(t)$ | $c_3(t)$ | ... | $c_{n-1}(t)$ | $c_n(t)$

(a) A vector of n computers/devices in an enterprise system of interest.

$V(t) = $ | $v_1(t)$ | $v_2(t)$ | $v_3(t)$ | ... | $v_{n-1}(t)$ | $v_n(t)$

(b) Vulnerabilities despite patching and defense.

$S(t) = $ | $s_1(t)=0$ | $s_2(t)=0$ | $s_3(t)=0$ | ... | $s_{n-1}(t)=1$ | $s_n(t)=1$

(e) Outcome of attack-defense interactions at time t are reflected by state vector: $s_i(t)=1$ means c_i is compromised, and $s_i(t)=0$ otherwise.

$O(t) = $ | $o_1(t)=0$ | $o_2(t)=1$ | $o_3(t)=?$ | ... | $o_{n-1}(t)=1$ | $o_n(t)=0$

(f) Defender's observations of the security state at time t: $o_2(t)$ is a false-positive, $o_n(t)$ is a false-negative, and $o_3(t)$ is not conclusive.

Fig. 1. An abstract representation of a snapshot of an enterprise system at time t, where the defenses (i.e., blue bars) accommodate both the defenses that are installed on the individual computers (e.g., anti-malware tools) and the defenses that are employed at the perimeter of the enterprise system (e.g., firewalls). Some attacks penetrate through the defenses (e.g., attacks against computer $c_n(t)$), while others fail (e.g., attacks against computer $c_1(t)$). Computer $c_i(t)$ may have a vector $v_i(t)$ of vulnerabilities, some of which may not be known to the defender (i.e., zero-day). The outcome of the attack-defense interaction at time t is reflected by a security state vector $S(t) = \{s_1(t), \ldots, s_n(t)\}$, where $s_i(t) = 0$ means computer $c_i(t)$ is secure at time t and $s_i(t) = 1$ means computer $c_i(t)$ is compromised at time t. However, the defender's observation of the security state vector $S(t)$, denoted by $O(t) = \{o_1(t), \ldots, o_n(t)\}$, may not be perfect because of false-positives, false-negatives, or non-decisions.

against which no defense is employed (i.e., the lack of blue bars). All of the other attacks are blocked by some of the defense mechanisms or the patch of the vulnerability in question.

Our methodology leads to 4 categories of security metrics with respect to vulnerabilities, defenses, threats, and situations. As we review each category of security metrics, we also discuss their theoretical and practical uses mentioned above as well as what the *ideal metrics* may be, which hints at the gap between the state-of-the-art and the ideal metrics we need to close. The insight behind the taxonomy is that, in principle, *situations* (or outcomes of cyber attack-defense interactions) are caused by certain *threats* (or attacks) against systems that have certain *vulnerabilities* (including human factors) and employ certain *defenses*. We here give a brief overview of the categories, which will be respectively elaborated in Sections 3-6.
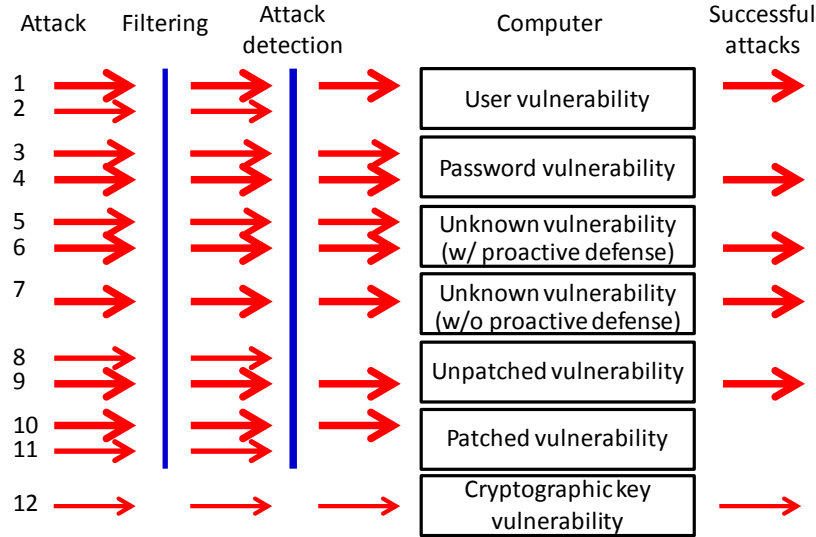
Fig. 2. A snapshot of attacks against a computer (or device), say $c_i(t)$, in the enterprise system at time t, where blue bars also represent defenses and red arrows represent attacks (their thickness reflect their defense/attack power). If $c_i(t)$ was compromised at time $t_1 < t$ and is not cleaned up at time t, or if $c_i(t)$ is compromised at time t, then $s_i(t) = 1$. If the defender correctly observes the state of $c_i(t)$ at time t, the observation is $o_i(t) = 1$. For cryptographic key vulnerabilities (e.g., Heartbleed-like vulnerabilities that are not patched), there is essentially no defense that can block the attack.

2.3.1. Metrics for measuring vulnerabilities. This category of metrics aim to measure the vulnerabilities of systems. As illustrated in Figure 1(a), an enterprise system consists of a vector $C(t) = (c_1(t), \ldots, c_n(t))$ of computers at time t. As illustrated in Figure 1(b), the enterprise system may have a vector $V(t) = (v_1(t), \ldots, v_n(t))$ of set of vulnerabilities at time t, where $v_i(t)$ is, as illustrated in Figure 2, the vector of vulnerabilities with respect to $c_i(t)$. Vulnerabilities include user vulnerabilities, password guessability, and software vulnerabilities. Software vulnerabilities can be known or unknown (i.e., zero-day) to the defender.

2.3.2. Metrics for measuring defenses. This category of metrics aim to measure the power or effectiveness of the defense mechanisms that are employed to protect enterprise and computer systems. As Figure 1(c) and Figure 2 illustrate, we use blue bars to represent defenses and their thickness to indicate their power. In practice, some computers may be well defended (illustrated by thick blue bars), some computers may be poorly defended (illustrated by thin blue bars), and some computers or zero-day vulnerabilities may not be defended at all (illustrated by the absence of blue bars).

2.3.3. Metrics for measuring threats. This category of metrics measure the threat landscape as well as the power or effectiveness of attacks. The threat landscape describes aspects of the attacking computers. As illustrated in Figure 1(d) and Figure 2, we use red arrows to represent attacks and their *thickness* to indicate their attack power. Some computers may be attacked by powerful attacks (illustrated by thick arrows), some computers may be attacked by less powerful attacks (illustrated by thin arrows), and some computers may not be attacked at all (illustrated by dash arrows).

2.3.4. Metrics for measuring situations. This category of metrics measure outcomes of attack-defense interactions, especially the evolution of the global security state $S(t)$ over time t [LeMay et al. ; Xu 2014a]. As illustrated in Figure 1(e) and Figure 2, secu-

rity state $s_i(t) = 1$ means computer $c_i(t)$ is compromised at time t, and $s_i(t) = 0$ otherwise. However, the defender may not know the true state vector $S(t) = (s_1(t), \ldots, s_n(t))$ because of measurement or observation errors such as false-positives, false-negatives, and non-decisions, as illustrated in Figure 1(f). In other words, it is possible that the observation vector $O(t) = (o_1(t), \ldots, o_n(t))$ observed by the defender is not equal to $S(t)$.

3. METRICS: MEASURING SYSTEM VULNERABILITIES

These metrics aim to measure the vulnerabilities of enterprise and computer systems via their users, the passwords of their users, their interfaces, their software vulnerabilities, and the vulnerabilities of the cryptographic keys they use.

3.1. Measuring system users' vulnerabilities

One metric is *user's susceptibility to phishing attacks* [Sheng et al. 2010]. This online study of 1,001 users shows that phishing education can reduce the user's susceptibility to phishing attacks and that young people (18 to 25 years old) are more susceptible to phishing attacks. This metric is measured via the false-positive rate that a user treats legitimate email or website as a phish, and the false-negative rate that a user treats a phishing email or website as legitimate and subsequently clicks the link in the email or submits information to the website.

Another metric is *user's susceptibility to malware infection* [Lalonde Levesque et al.]. This clinical study of interactions between human users, anti-malware software, and malware involves 50 users, who monitor their laptops for possible infections during a period of 4 months. During this period of time, 38% of users are found to be exposed to malware, which indicates the value of the anti-malware tool (because these laptops would have been infected if anti-malware software was not used). The study also shows that user demographics (e.g., gender, age) are *not* significant factors in determining a user's susceptibility to malware infection, which contradicts the aforementioned finding in regards to users' susceptibility to phishing attacks [Sheng et al. 2010]. Nevertheless, it is interesting to note that (i) users installing many applications are more susceptible to malware infections, because the chance of installing malicious applications is higher, and (ii) users visiting many websites are more susceptible to malware infections, because some websites are malicious [Lalonde Levesque et al.].

It is important to understand and measure the degrees of users' susceptibilities to each individual class of attacks and to multiple classes of attacks collectively (e.g., multiple forms of social-engineering attacks). For this purpose, research needs to be conducted to quantify how the susceptibilities are dependent upon factors that affect users' security decisions (e.g., personality such as high vs. low attention control [Neupane et al. 2015]). This area is little understood [Howe et al. 2012; Sheng et al. 2010; Lalonde Levesque et al.], but the reward is high. For the theoretical use of security metrics, these metrics can be incorporated into security models as parameters to model (e.g.) the time or effort that is needed in order for an attacker to exploit user vulnerabilities to compromise a computer or to penetrate into an enterprise system. For the practical use of security metrics, these metrics can be used to tailor defenses for individual users (e.g., a careless employee may have to go through some security proxy in order to access Internet websites). It would be appropriate to say that being able to measure these security metrics is as important as being able to measure individual users' susceptibility to cancers because of (e.g.) her genes. As the ability to quantify an individual's predisposition to diseases can lead to proactive treatment, the ability to quantify security can lead to tailored and more effective defenses.

3.2. Measuring password vulnerabilities

The *parameterized password guessability* metric measures the number of guesses an attacker with a particular cracking algorithm (i.e., a particular threat model) needs to make before recovering a password [Weir et al. 2010; Bonneau 2012a; Kelley et al. ; Ur et al.]. This metric is easier to use than earlier metrics such as *password entropy* [Burr et al. 2006], which cannot tell which passwords are easier to crack than others, and *statistical password guessability* [Bonneau 2012b; Bonneau 2012a; Kelley et al.], which is more appropriate for evaluating passwords as a whole (rather than for evaluating them individually).

The parameterized password guessability metric should be used with caution if a single password cracking algorithm is used, because different cracking algorithms can have very different strategies with varying results [Ur et al.]. When the defender is uncertain about the threat model, multiple cracking strategies need to be considered. For both theoretical and practical uses of password vulnerability metrics, we might need to consider the worst-case and/or the average-case parameterized password guessabilities. This is one of the few sub-categories of security metrics that are better understood.

3.3. Measuring interface-induced vulnerabilities

The interface to access an enterprise or computer system from the outside world (e.g., service access points) offers potential opportunities for launching cyber attacks against the system. The *attack surface* metric measures the number and severity of attack vectors that can be launched against a system through its service access points such as sockets and RPC endpoints [Manadhata and Wing 2011]. It is worth mentioning that the attack surface is not necessarily dependent upon software vulnerabilities. The attack surface should be used with caution because reducing the attack surface (e.g., uninstalling a security software) does not necessarily improve security [Nayak et al.]. It has been suggested to define a variant of attack surface as the portion of the attack surface that has been exercised [Nayak et al.]. This variant, while useful in analyzing historical data (i.e., incidents that have occurred), may or may not be appropriate for measuring security in the future because an attack surface not exercised in the past may be exercised in the future.

Suppose we are to model security from higher levels of abstractions by treating system interface. We would need to measure *interface-induced system susceptibility*, which measures how the exercise of attack surface is dependent upon the features of attack surfaces. For practical purposes, it is ideal to be able to predict interface-induced system susceptibilities, namely the interfaces that will be exploited to launch attacks in the near future. Knowing which interfaces are more likely to be abused to launch attacks would allow the defender to employ tailored defenses that pay particular attention to these interfaces.

3.4. Measuring software vulnerabilities

Software vulnerabilities are the main venue for launching cyber attacks. We classify the metrics for measuring software vulnerabilities into three sub-categories: spatial characteristics, temporal characteristics, and severity.

3.4.1. Measuring software vulnerability spatial characteristics. These metrics reflect how spatially vulnerable an enterprise or computer system is. The number of *unpatched vulnerabilities* at time t can be determined by using vulnerability scanners [Chew et al. ; for Internet Security 2010]. The *vulnerability prevalence* metric measures the popularity of a vulnerability in a system [Zhang et al. 2014b]. This metric is important because a single vulnerability may exist in many computers of an enterprise or cloud system,

and because an attacker can launch a single attack against all the computers that possess a prevalent vulnerability. Another variant metric is the number of *exploited vulnerabilities* that have been exploited in the past [Nayak et al. ; Allodi]. This metric is important because some vulnerabilities may never get exploited in the real world because, for example, many vulnerabilities are difficult to exploit or their exploitation does not reward the attacker with many compromised computers. For example, one study [Nayak et al.] shows that at most 35% of the known vulnerabilities have been exploited, with a small number of vulnerabilities (e.g., CVE-2008-4250 and CVE-2009-4324) being responsible for many attacks. Another study [Allodi] shows that 10% of the vulnerabilities are responsible for 90% attacks.

When using these metrics as parameters in security modeling, we would need to estimate the susceptibility of a computer to attacks that exploit software vulnerabilities at time t. When using these metrics to compare the security of two systems or the security of a system during two periods of time, one must be cautious about (i) some vulnerabilities never being exploited, (ii) the varying capabilities of scanners in terms of their scanning depth and completeness, and (iii) the threats may be different (e.g., two systems may be targeted by different attackers and there may be zero-day attacks that are not detected yet). In other words, the theoretical and practical uses of these security metrics require us to estimate, or even predict, the *vulnerability situation awareness* metric. This metric measures the number of vulnerabilities of a system at time t and the likelihood of each of these vulnerabilities being exploited at time $t' \geq t$.

3.4.2. Measuring software vulnerability temporal characteristics. Temporal characteristics of software vulnerabilities include their evolution and lifetime.

Measuring evolution of software vulnerabilities. The *historical vulnerability* metric measures the degree that a system is vulnerable, or the number of vulnerabilities, in the past [Al-Shaer et al. 2008; Ahmed et al. 2008]. The *future vulnerability* metric measures the number of vulnerabilities that will be discovered during a future period of time [Al-Shaer et al. 2008; Ahmed et al. 2008]. Interesting variants of these metrics include historical exploited vulnerabilities, namely the number of vulnerabilities that were exploited in the past, and future exploited vulnerabilities, namely the number of vulnerabilities that will be exploited during a future period of time. The *tendency-to-be-exploited* metric measures the tendency that a vulnerability may be exploited, where the "tendency" may be computed from (e.g.) the information that was posted on Twitter before vulnerability disclosures [Sabottke et al. 2015]. This metric may be used to prioritize vulnerabilities for patching.

Measuring software vulnerability lifetime. It is ideal that each vulnerability is immediately patched upon its disclosure. Despite the enforcement of patching policies, some vulnerabilities may never get patched. The *vulnerability lifetime* metric measures how long it takes to patch a vulnerability since its disclosure. Different vulnerability lifetimes may be exhibited at the client-end, the server-end, and the cloud-end.

Client-end vulnerabilities are often exploited to launch targeted attacks (e.g., spear-fishing) [Hardy et al. 2014; Marczak et al. 2014]. These vulnerabilities are hard to patch completely because of their prevalence (i.e., a vulnerability may appear in multiple programs) [Nappa et al. 2015]. A study conducted in year 2010 [Frei and Kristensen 2010] shows that 50% of 2 million Windows users in question are exposed to 297 vulnerabilities over a period of 12 months. A more recent study [Nappa et al. 2015] shows that despite the presence of 13 automated patching mechanisms (other than the Windows update), the median fraction of computers that are patched when exploits are available is no greater than 14%, the median time for patching 50% of vulnerable computers is 45 days after disclosure.

One would think that server-end vulnerabilities are more rapidly patched than client-end ones. Let us consider the disclosure of two severe vulnerabilities in OpenSSL. First, for the pseudorandom-number-generation vulnerability in Debian Linux's OpenSSL, a study [Yilek et al.] shows that 30% of the computers that were vulnerable 4 days after disclosure remain vulnerable almost 180 days later (i.e., 184 days after disclosure). This is somewhat surprising because the private keys generated by the vulnerable computers might have been exposed to the attacker. Second, for the Heartbleed vulnerability in OpenSSL that can be remotely exploited to read a vulnerable server's sensitive memory that may contain cryptographic keys and passwords, a study [Durumeric et al. 2014] estimates that 24%-55% of the HTTPS servers in Alexa's Top 1 Million websites were initially vulnerable. Moreover, 11% of the HTTPS servers in Alexa's Top 1 Million remain vulnerable 2 days after disclosure, and 3% of the HTTPS servers in Alexa's Top 1 Million were still vulnerable 60 days after disclosure.

One may think that vulnerabilities in the cloud are well managed, perhaps because cloud users can run public virtual machine images (in addition to their own images). A study [Zhang et al. 2014b] shows that many of the 6,000 public Amazon Machine Images (AMIs) offered by Amazon Web Services (AWS) Elastic Compute Cloud (EC2), contain a considerable number of vulnerabilities, and that Amazon typically notifies cloud users about vulnerabilities 14 days after their disclosure.

Summarizing the temporal metrics discussed above, we observe that defenders need to do a substantially better job at reducing the lifetime of software vulnerabilities after disclosure. Because vulnerability lifetime may never be reduced to 0, it is important to know the vulnerability vector $V(t)$ or $v_i(t)$ at any time t. For using vulnerability lifetime in security modeling, we need to know its statistical distribution and how the distribution is dependent upon various factors.

3.4.3. Measuring software vulnerability severity. This metric measures the degree of damage that can be caused by the exploitation of a vulnerability. A popular example is the *CVSS score*, which considers the following three factors [of Incident Response and (FIRST)]. The base score reflects the vulnerability's time- and environment-invariant characteristics, such as its access condition, the complexity to exploiting it, and the impact once exploited. The temporal and environmental scores reflect its time- and environment-dependent characteristics. Another example is the *availability of exploits* in black markets [Bilge and Dumitras], which is interesting because the public release of vulnerabilities is often followed by the increase of exploits.

However, many vulnerabilities have the same CVSS scores [Jansen 2009; Allodi and Massacci 2014]. The practice of using CVSS scores (or base scores) to prioritize the patching of vulnerabilities has been considered both harmful, because information about low-severity bugs can lead to the development of high-severity attacks [Allodi and Massacci 2014; Arnold et al. ; Brumley et al.], and ineffective, because patching a vulnerability solely because of its high CVSS score makes no difference than patching vulnerabilities randomly [Allodi and Massacci 2014]. For practical use, it would be ideal if we can precisely define the intuitive metric of *patching priority*. For theoretical use, it would be ideal if we can quantify the *global damage* of a vulnerability to an enterprise system upon its exploitation, which may in turn help measure the *patching priority*.

3.5. Measuring cryptographic key vulnerabilities

Cryptographic keys are vulnerable when the underlying random number generators are weak, as witnessed by the pseudorandom-number-generation vulnerability in Debian Linux's OpenSSL [Yilek et al.]. Here we highlight the weak cryptographic keys

caused by using the fast `/dev/urandom` as a replacement of the slow `/dev/random` in Linux [Heninger et al. 2012]. The difference between them is that the former returns the requested number of bytes immediately even though not enough entropy has been collected, while the latter returns the requested number of bytes only after the entropy pool contains the required fresh entropy. As a consequence of using `/dev/urandom`, the same key materials (e.g., prime numbers) can be generated and used by multiple devices. A study shows that RSA private keys for 0.50% of the TLS hosts examined and 0.03% of SSH hosts examined can be exposed because their RSA moduli shared non-trivial factors [Heninger et al. 2012]. The study also shows that the DSA private keys for 1.03% of the SSH hosts examined can be extracted due to the insufficient randomness in their digital signatures. These problems mainly exist in embedded devices, including routers and firewalls, because they generate cryptographic keys on their first boot.

This kind of vulnerability should have been prevented by prudential engineering in the use of randomness, which requires the programmer to understand, for example, the difference between `/dev/random` and `/dev/urandom`. Nevertheless, it would be ideal to know whether a newly generated cryptographic key is weak or not.

4. METRICS: MEASURING DEFENSES

These metrics measure the defenses employed to protect enterprise and computer systems via the effectiveness of blacklisting, the power of attack detection, the effectiveness of software diversification, the effectiveness of memory randomization, and the overall effectiveness of these defenses.

4.1. Measuring the effectiveness of blacklisting

Blacklisting is a useful, lightweight defense mechanism. Suppose a malicious entity (e.g., attacking computer, IP address, malicious URL, botnet command-and-control server, and dropzone server) is observed at time t. Then, the traffic flowing to or from the malicious entity can be blocked starting at some time $t' \geq t$. The *reaction time* is the delay $t' - t$ between the observation of the malicious entity at time t and the blacklisting of the malicious entity at time t' [Kührer et al.]. The *coverage* metric measures the portion of malicious entities that are blacklisted. For example, a study shows that the union of 15 malware blacklists covers only 20% of the malicious domains that are compromised by some major malware families [Kührer et al.].

These metrics are with respect to the observers and blacklists in question. For practical use, these metrics can be used to compare the effectiveness of different blacklists and can guide the design of better blacklisting solutions (e.g., achieving a certain reaction time and a certain degree of coverage). For theoretical use in security modeling, we might need to accommodate them into a unified metric, which may be called *blacklisting probability* and may be measured by the conditional probability that a malicious entity at time t (e.g., URL or IP address) is blacklisted at time t. This would require us to understand the various factors that can impact malicious entities to be blacklisted.

4.2. Measuring the power of attack detection

Attack detection tools, such as cyber instruments (e.g., honeypots and blackholes that monitor unused IP addresses for attacks), intrusion detection systems and anti-malware programs, aim to detect attacks. The effectiveness of attack detection can be measured by their individual effectiveness, relative effectiveness, and collective effectiveness.

4.2.1. Measuring the individual detection power. For instrument-based attack detection, the *detection time* metric measures the delay between the time t_0 at which a compro-

mised computer sends its first scan packet and the time t that a scan packet is observed by the instrument [Rajab et al. 2005]. This metric depends on several factors, including malware spreading model, the distribution of vulnerable computers, the size of the monitored IP address space, and the locations of the instrument.

For intrusion detection systems (including anomaly-based, host-based, and network-based), the initial set of metrics for measuring their detection power are: The *true-positive* rate, denoted by $\Pr(A|I)$, is defined as the probability that an intrusion (I) is detected as an alert that indicates attack (A). The *false-negative* rate, denoted by $\Pr(\neg A|I)$, is defined as the probability that an intrusion is not detected as an attack. The *true-negative* rate, denoted by $\Pr(\neg A|\neg I)$, is defined as the probability that a non-intrusion is not detected as an attack. The *false-positive* rate, also called *false alarm rate*, denoted by $\Pr(A|\neg I)$, is defined as the probability that a non-intrusion is detected as an attack. Note that $\Pr(A|I) + \Pr(\neg A|I) = \Pr(\neg A|\neg I) + \Pr(A|\neg I) = 1$. The *receiver operating characteristic* (ROC) curve reflects the dependence of the true-positive rate $\Pr(A|I)$ on the false-positive rate $\Pr(A|\neg I)$, and therefore can help determine the trade-off between the true-positive rate and the false-positive rate.

When using the preceding security metrics to compare the effectiveness of intrusion detection systems, care must be taken. One issue is the event unit, such as packet vs. flow in the context of network-based intrusion detection [Gu et al.]. Another issue is the base rate of intrusions, which can lead to misleading results if not adequately treated — a phenomenon known as the *base-rate fallacy* [Axelsson]. In order to deal with the base-rate fallacy, one can treat the input to an intrusion detection system as a stream $\mathcal{I}$ of 0/1 random variables (0 indicates benign or normal, 1 indicates malicious or abnormal), and treat the output of the intrusion detection system as a stream $\mathcal{O}$ of 0/1 random variables (0 indicates no alert or normal, 1 indicates alert or abnormal). Let $H(\mathcal{I})$ and $H(\mathcal{O})$ denote the entropy of $\mathcal{I}$ and $\mathcal{O}$, respectively. The mutual information $I(\mathcal{I}, \mathcal{O})$ between $\mathcal{I}$ and $\mathcal{O}$, namely $I(\mathcal{I}, \mathcal{O}) = H(\mathcal{I}) - H(\mathcal{I}|\mathcal{O})$, indicates the amount of uncertainty of $\mathcal{I}$ reduced after knowing $\mathcal{O}$. The *intrusion detection capability* metric is defined as the normalization of $I(\mathcal{I}, \mathcal{O})$ with respect to $H(\mathcal{I})$, which reflects the base rate [Gu et al.].

Intrusion detection may also be measured via the *cost* metric in the decision-theoretic framework [Gaffney Jr and Ulvila 2001]. The cost includes both the operational cost of intrusion detection and the damage caused by false negatives. Cardenas et al. [Cardenas et al. 2006] unified these metrics into a single framework of multi-criteria optimization, which allows fair comparisons between intrusion detection systems in different operational environments. We refer to a recent survey [Milenkoski et al. 2015] for more details.

The metrics mentioned above are mainly geared towards the practical use of measuring the detection power of each individual detection system and comparing the detection power of two detection systems. When modeling intrusion detection systems as a component in a broader or holistic security model, we may need to define and measure the *detection probability* metric as the conditional probability that a compromised computer at time t is also detected as compromised at time t, namely $\Pr(o_i(t) = 1|s_i(t) = 1)$. This would require us to study how this probability is dependent upon other factors.

4.2.2. Measuring the relative detection power. This metric [Boggs and Stolfo 2011; Boggs et al.] reflects the effectiveness of a defense tool when employed in addition to other defense tools. A defense tool does not offer any extra power if it cannot detect any attack that cannot be detected by the already deployed defense tools. Let $\mathcal{A}$ denote a set of attacks, $\mathcal{D} = \{d_1, \ldots, d_n\}$ denote the set of n defense tools, and X_d denote the set of attacks that are detected by defense tool $d \in \mathcal{D}$. The relative power of de-

fense tool $d' \in \mathcal{D}$ with respect to the set $D \subset \mathcal{D}$ of deployed defense tools is defined as $\frac{|X_{d'} - \cup_{d \in D} X_d|}{|A|}$.

This kind of metric can be used to help decide whether to purchase/install a new defense tool or not, depending on its relative detection power. This kind of metric could be used as parameters in security models that aim to characterize the global effectiveness of employing a defense tool. It is worth mentioning that these metrics are measured based on attacks that have been seen in the past. We might need to estimate these metrics with respect to future attacks that may contain some unknown attacks, which we may call *relative effectiveness against known and unknown attacks*. This may require us to estimate the base rate of unknown attacks.

4.2.3. Measuring the collective detection power. This metric has been proposed for measuring the collective effectiveness of intrusion detection systems and anti-malware programs [Boggs and Stolfo 2011; Morales et al. 2012; Boggs et al. ; Mohaisen and Alrawi 2014; Yardon 2014]. Let $\mathcal{A}$ denote a set of attacks, $\mathcal{D} = \{d_1, \ldots, d_n\}$ denote the set of n defense tools, X_d denote the set of attacks that are detected by defense tool $d \in \mathcal{D}$. The collective detection power of defense tools $D \subseteq \{d_1, \ldots, d_n\}$ is defined as $\frac{|\cup_{d \in D} X_d|}{|A|}$ [Boggs and Stolfo 2011; Boggs et al.]. For malware detection, experiments [Morales et al. 2012; Mohaisen and Alrawi 2014; Yardon 2014] show that the collective use of multiple anti-malware programs still cannot detect all malware infections. For example, one recent estimation [Yardon 2014] shows that anti-malware tools are only able to detect 45% of attacks.

The practical use of these metrics include the comparison of the collective effectiveness between two combinations of detection tools and the evaluation of the effectiveness of defense-in-depth. The theoretical use of these metrics include the incorporation of them as parameters into security models that aim to characterize the global collective effectiveness of employing a combination of defense tools. Like in the case of relative effectiveness mentioned above, these metrics may need to be measured or estimated with respect to known and unknown attacks, which we may call *collectiveness effectiveness against known and unknown attacks*. This may also require us to estimate the base rate of unknown attacks.

4.3. Measuring the effectiveness of Address Space Layout Randomization (ASLR)

Code injection was a popular attack that aims to inject some malicious code into a running program and direct the processor to execute it. The attack requires the presence of a memory region that is both executable and writable, which was possible because operating systems used to not distinguish programs and data. The attack can be defeated by deploying Data Execution Prevention (DEP, also known as W⊕X), which ensures that a memory page can be writable or executable at any point in time, but not both. The deployment of DEP made attackers move away from launching code injection attacks to launching code reuse attacks, which craft attack payloads from pieces or "gadgets" of executable code that is already running in the system. In order to launch a code-reuse attack, the attacker needs to know where to look for gadgets. This was possible because the base addresses of code and data (including stack and heap) in the virtual memory used to be fixed.

One approach to defending against code reuse attacks is to use ASLR to "blind" the attacker, by randomizing the base addresses (i.e., shuffling the code layout in the memory) such that the attacker cannot find useful gadgets. Coarse-grained ASLR has the vulnerability that the leak or exposure of a single address gives the attacker adequate information to extract all code addresses. Fine-grained ASLR do not suffer from this problem (e.g., page-level randomization [Backes and Nürnberger 2014;

Larsen et al.]), but are still susceptible to attacks that craft attack payloads from Just-In-Time (JIT) code [Snow et al. 2013]. This attack can be defeated by *destructive code read*, namely that the code in executable memory pages is garbled once it is read [Tang et al. 2015]. ASLR can also be enhanced by preventing the leak of code pointers, while rendering leaks of other information (e.g., data pointers) useless for deriving code pointers [Lu et al. 2015].

There are two metrics for measuring the effectiveness of ASLR. One metric is the *entropy* of a memory section, because a greater entropy would mean a greater effort in order for an attacker to compromise the system. For example, a brute-force attack can feasibly defeat a low-entropy ASLR on 32-bit platforms [Shacham et al. 2004]. The related *effective entropy* metric measures the entropy in a memory section that the attacker cannot circumvent by exploiting the interactions between memory sections [Herlands et al. 2014].

The two metrics mentioned above *indirectly* reflect the effectiveness of ASLR. For practical use, we would need to measure the *direct* security gain offered by the deployment of ASLR and/or the extra effort that is imposed on the attacker in order to circumvent ASLR. Being able to measure the effectiveness of ASLR on individual computers, the resulting metrics could be be incorporated into theoretical cyber security models to characterize their global effectiveness.

4.4. Measuring the effectiveness of enforcing Control-Flow Integrity (CFI)

Despite the employment of defenses such as the aforementioned DEP and ASLR, control-flow hijacking remains to be a big threat [Szekeres et al. 2013; Larsen et al.]. Enforcing CFI has a great potential in assuring security. The basic idea underlying CFI is to extract a program's Control-Flow Graph (CFG), typically from its source code via static analysis, and then instrument the corresponding binary code to abide by the CFG at runtime. This can be implemented by runtime checking of the tags that were assigned to the indirect branches in the CFG, such as indirect calls, indirect jumps, and returns. Since it is expensive to enforce CFI according to the entire CFG [Abadi et al.], practical solutions enforce weaker restrictions via a limited number of tags. It was known that coarse-grained enforcement of CFI uses a small number of tags and can be compromised by code reuse attacks [Göktas et al. ; Davi et al. 2014; Carlini and Wagner 2014]. This inadequacy led to fine-grained CFI, such as the enforcement of forward-edge control integrity (i.e. indirect calls but not returns) [Tice et al. 2014], and the use of message authentication code to prevents unintended control transfers in the CFG [Mashtizadeh et al. 2015]. Even a fully accurate, fine-grained CFG can be compromised by the *control flow bending* attack [Carlini et al. 2015], which however can be mitigated by per-input CFI [Niu and Tan 2015].

How should we measure the power of CFI? First, the power of CFI is fundamentally limited by the *accuracy* of the CFGs [Evans et al. 2015]. Because CFGs are computed via static analysis, their accuracy depends on *sound* and *complete* pointer analysis, which is undecidable in general [Ramalingam 1994]. The trade-off of using an unsound pointer analysis is that all of the due connections may not be reported and therefore can cause false-positives. The trade-off of using an incomplete pointer analysis is that excessive connections may be reported (i.e., over-approximation), which can be exploited to run arbitrary code despite the enforced fine-grained CFI [Evans et al. 2015]. Second, we need to measure the *resilience* of a CFI scheme against control flow bending attacks, which ideally reflects the effort (or premises) that an attacker must make (or satisfy) in order to evade the CFI scheme. This metric would allow us to compare the resilience of two CFI schemes. Third, it would be ideal if we can measure the power

of CFI via the classes of attacks that it can defeat. The key issue here is to have a formalism by which we can *precisely* classify attacks. The challenge is that there could be infinitely many attacks and it is not clear what would be the right formalism.

4.5. Measuring overall defense power

In the above we discussed the defense power of individual defense mechanisms. In practice, different kinds of defense mechanisms are used together and therefore we need to consider the overall defense power. This is a field that is much less understood. The well known approach is to use penetration test to evaluate an enterprise or computer system's *penetration resistance*. This metric can be measured as the effort (e.g., person-day or cost) it requires for the red team to penetrate into the system [Levin 2003]. This metric can be used to compare the effectiveness of two systems against the same red team. On the other hand, there have been some initial studies at measuring the power of Moving Target Defense (MTD), which may deploy a set of MTD mechanisms together. The analytic approach aims to indirectly measure the degree that an enterprise system can tolerate some undesirable security configurations, under which the global security state $S(t)$ may evolve towards many computers are compromised [Han et al.]. Complementary to the analysis approach, there have been proposals for evaluating the effectiveness of MTD via experimentation [Zaffarano et al.], emulation [Eskridge et al.], and simulation [Prakash and Wellman].

When using these metrics, one must be cautious about what the penetration resistance is with respect to the specific red team, which may reassemble real-world attackers to a certain extent. Moreover, one should bear in mind that the identification of security holes does not offer any quantitative security [Sanders 2014]. For theoretical and practical uses, it is ideal that the *penetration resistance* should at least consider some hypothetical new attacks, which may exploit some known or zero-day vulnerabilities (i.e., "what if" analysis). Moreover, systematic metrics need to be devised to directly measure the effectiveness of MTD, which may be in terms of some *direct effectiveness* metrics, such as the security gained by launching MTD and/or the extra effort imposed on the attacker in order to achieve attack goals.

5. METRICS: MEASURING THREATS

These metrics measure the threats against an enterprise or computer system via the threat landscape, the threat of zero-day attacks, the power of individual attacks, the sophistication of obfuscation, and the evasion capability.

5.1. Measuring the threat landscape

The threat landscape can be characterized via multiple attributes. One attribute is the attack vector. The number of *exploit kits* metric describes the number of automated attack tools that are available in the black market [Ablon et al. 2014]. This metric can be extended to accommodate, for example, the vulnerabilities that the exploits are geared for. This is a good indicator of cyber threats because most attacks would be launched from these exploit kits [Nayak et al. ; Allodi].

The *network maliciousness* metric [Zhang et al. 2014a] measures the fraction of blacklisted IP addresses in a network. The study [Zhang et al. 2014a] shows that there were 350 autonomous systems which had at least 50% of their IP addresses blacklisted. Moreover, there was a correlation between mismanaged networks and malicious networks, where "mismanaged networks" are those networks that do not follow accepted policies/guidelines. The related *rogue network* metric measures the population of networks that were abused to launch drive-by download or phishing attacks [Stone-Gross et al. 2009]. The *ISP badness* metric quantifies the effect of spam from one ISP or Autonomous System (AS) on the rest of the Internet [Johnson et al. 2012].

The *control-plane reputation* metric quantifies the maliciousness of attacker-owned (i.e., rather than legitimate but mismanaged/abused) ASs based on their control plane information (e.g., routing behavior), which can achieve an *early-detection time* of 50-60 days (before these malicious ASs are noticed by other defense means) [Konte et al.]. Malicious, rogue, and bad networks, once detected, can be filtered by enterprise systems via blacklisting.

The *cybersecurity posture* metric measures the dynamic threat imposed by the attacking computers [Zhan et al. 2014]. It may include the attacks observed at honeypots, network telescopes, and/or production enterprise systems. One related metric, the *sweep-time*, measures the time it takes for each computer or IP address in a target enterprise system to be scanned or attacked at least once [Zhan et al. 2014]. Another related *attack rate* metric measures the number of attacks that arrive at a system of interest per unit time [Zhan et al. 2013; Zhan et al. 2015]. These metrics reflect the aggressiveness of cyber attacks.

Although the security metrics mentioned above can reflect some aspects of the threat landscape, we might need to define what may be called *comprehensive cyber threat posture*, which reflects the holistic threat landscape. This metric is useful because the threat landscape could be used as the "base rate" (in the language of intrusion detection systems) that can help fairly compare the overall defense effectiveness of two enterprise systems. It is interesting to investigate how these security metrics should be incorporated into security models as parameters for analyzing, for example, the evolution of the global security state $S(t)$ over time t.

5.2. Measuring zero-day attacks

The number of *zero-day attacks* measures how many zero-day attacks were launched during a past period of time. For example, Symantec estimated that there were 8-15 zero-day vulnerabilities between 2006 and 2011, among which 9 were exploited to launch zero-day attacks in 2008, 12 were exploited to launch zero-day attacks in 2009, 14 were exploited to launch zero-day attacks in 2010, and 8 were exploited to launch zero-day attacks in 2011 [Corporation 2012]. The *lifetime of zero-day attacks* measures the period of time between when the attack was launched and when the corresponding vulnerability is disclosed to the public. One study shows that the lifetime of zero-day attacks can last 19-900 days, with a median of 240 days and an average of 312 days [Bilge and Dumitras]. The number of *zero-day attack victims* measures the number of computers that were compromised by zero-day attacks. It has been reported that most zero-day attacks affected a small number of computers, with a few exceptions (e.g., Stuxnet) [Bilge and Dumitras]. This suggests that zero-day attacks are mainly used for launching targeted attacks.

The preceding metrics are typically measured after the fact. For both theoretical and practical uses of security metrics measuring zero-day attacks, one important metric that has yet to be studied is the *susceptibility of a computer to zero-day attacks*. Metrics of this kind, once understood, not only could help allocate defense resources to carefully monitor the computers that are more susceptible to zero-day attacks, but also could be incorporated into security models to analyze the evolution of the global security state $S(t)$ over time t.

5.3. Measuring the attack power

Now, we discuss the metrics for measuring the power of targeted attacks, the power of botnets, the power of malware spreading, and the power of attacks that exploit multiple vulnerabilities in enterprise systems.

5.3.1. Measuring the power of targeted attacks. The success of targeted attacks or Advanced Persistent Threats (APT) often depends on the delivery of malware and the tactics that are used to lure the target to open malicious email attachments. Let α denote the social engineering tactics, ranging from the least sophisticated to the most sophisticated (e.g., $\alpha \in \{0, \ldots, 10\}$). Let β denote the technical sophistication of the malware that are used in the attacks, also ranging from the least sophisticated to the most sophisticated (e.g., $\beta \in [0, 1]$). The *targeted threat index* metric, which measures the level of targeted malware attacks, can be defined as $\alpha \times \beta$ [Hardy et al. 2014].

The preceding metric represents a first step in measuring the power of targeted attacks, and much research has yet to be done. It would be ideal if we can measure the *susceptibility of a computer to targeted attacks*, for which the user's susceptibility to social-engineering attacks metric mentioned above would be an integral factor. Being able to measure metrics of this kind allows the defender to tailor defenses for the computers that are more vulnerable to targeted attacks. These metrics could also be immediately incorporated into security models that analyze (e.g.) the evolution of $S(t)$ over time t.

5.3.2. Measuring the attack power of botnets. The threat of botnets can be characterized by several metrics. The first metric is *botnet size*. It is natural to count the number x of bots belonging to a botnet. It is important to count the number of bots that can be instructed to launch attacks (e.g., distributed denial-of-service attacks) at a point in time t, denoted by $y(t)$. Due to factors such as the diurnal effect, which explains why some bot computers are powered off during night hours at local time zones, $y(t)$ is often much smaller than x [Dagon et al. 2006]. A related metric is the network bandwidth that a botnet can use to launch denial-of-service attacks [Dagon et al. 2007]. The second metric is *botnet efficiency*, which can be defined as the network diameter of the botnet network topology [Dagon et al. 2007]. This metric measures a botnet's capability in communicating command-and-control messages and updating bot programs. The third metric is *botnet robustness*, which measures the robustness of botnets under random or intelligent disruptions [Dagon et al. 2007]. There has been a body of literature [Albert and Barabasi 2002] on measuring complex network robustness that can be adopted for characterizing botnets.

Although the above metrics measure botnets from some intuitive aspects, it remains elusive to define the intuitive metric of *botnet attack power*, which is important because it can prioritize the countermeasures against botnets. Moreover, the intuitive *botnet resilience* metric would need to take into consideration the counter-countermeasures that may be employed by the attacker during the process that the defender launches countermeasures against botnets.

5.3.3. Measuring malware spreading power. The *infection rate* metric, denoted by γ, measures the average number of vulnerable computers that are infected by a compromised computer (per time unit) at the early stage of spreading [Chen and Ji]. Intuitively, γ depends on the scanning strategy. Here we only mention the random scanning strategy. With random scanning, the malware scans over the vulnerable computers uniformly at random. Denote by z the number of scans and infections that are made by an infectious computer (per time unit). Denote by w the number of vulnerable computers. For the IPv4 address space, the infection rate is $\gamma = zw/2^{32}$ [Chen and Ji].

The infection rate metric can be used to compare the infection rate of different malware and has been used in various models (e.g., [Chen and Ji]) to derive interesting results. It should be noted that the infection rate is defined with respect to the early stage of infection. We would need other metrics to measure, such as the *sweep-time* that all or a fraction of the vulnerable computers that will be infected. This metric is hard to compute for arbitrary scanning strategies. Another interesting metric is the

number of *wasted scans*, namely those which arrive at some infected computers. This metric reflects the unstealthly nature of malware spreading.

5.3.4. Measuring the power of attacks that exploit multiple vulnerabilities. Vulnerabilities can be exploited in a chaining fashion. There is a large body of literature in this field, including *attack graphs* (see, e.g., [Phillips and Swiler 1998; Ritchey and Ammann 2000; Sheyner et al. 2002;   Jha et al. 2002;   Ammann et al. ;   Albanese et al. 2012; Homer et al. 2013; Cheng et al. 2014]), *attack trees* [Schneier 2000], and *privilege trees* [Dacier et al. 1996; Ortalo et al. 1999]. At a high level, these models accommodate system vulnerabilities, vulnerability dependencies (i.e., prerequisites), firewall rules, etc. In these models, the attacker is initially in some security state and attempts to move from the initial state to some goal state, which often corresponds to the compromise of computers. These studies have led to a rich set of metrics, such as the following.

The *necessary defense* metric measures the minimal set of defense countermeasures that must be employed in order to thwart a certain attack [Sheyner et al. 2002]. The greater the necessary defense, the more powerful the attack.

The *weakest adversary* metric measures the minimum adversary capabilities that are needed in order to achieve an attack goal [Pamula et al.]. This metric can be used to compare the power of two attacks with respect to some attack goal(s). For example, one attack has the required weakest adversary capabilities, but the other does not.

The *existence, number, and lengths of attack paths* metrics measure the these attributes of attack paths from an initial state to the goal state [Ritchey and Ammann 2000; Sheyner et al. 2002; Jha et al. 2002; Cheng et al. 2014]. These metrics can be used to compare two attacks. For example, the attack that has a set X of attack paths is more powerful than another attack that has a set Y of attack paths, where $Y \subset X$.

The *k-zero-day-safety* metric measures the number of zero-day vulnerabilities that are needed in order for an attacker to compromise a target [Wang et al.]. This metric can be used to compare the power of two attacks as follows: An attack that requires k_1 zero-day vulnerabilities in order to compromise a target is more powerful that an attack that requires k_2 zero-day vulnerabilities, where $k_1 < k_2$.

The *effort-to-security-failure* metric measures what the attacker needs to do in order to move from an initial set of privileges to the goal set of escalated privileges [Dacier et al. 1996; Ortalo et al. 1999]. An attack that incurs a smaller effort-to-security-failure is more powerful than an attack that requires a greater effort, assuming the efforts are comparable.

Although the metrics mentioned are useful, it would be ideal if we can measure what we call *multi-stage attack power*, which may be able to incorporate all the metrics mentioned above into a single one. This metric could also be incorporated into security models to analyze (e.g.) the evolution of security state $S(t)$ over t. One barrier is to systematically treat unknown vulnerabilities.

5.3.5. Measuring the power of evasion against learning-based detection. Sophisticated attacks can evade the defense system by, for example, manipulating some features that are used in the detection models (e.g., classifiers). This problem is generally known as *adversarial machine learning* [Dalvi et al. 2004; Lowd and Meek 2005; Huang et al. ; Šrndic and Laskov 2014]. There is a spectrum of evasion scenarios, which vary in terms of the information the attacker knows about the detection models, such as (i) knowing only the feature set used by the defender, (ii) knowing both the feature set and the training samples used by the defender, and (iii) knowing the feature set, the training samples, and the attack detection model (e.g., classifiers) used by the defender

[Šrndic and Laskov 2014; Xu et al. 2014b]. The evaluation of effectiveness is typically based on metrics such as false-positive and false-negative rates as a consequence of applying a certain evasion method.

It is ideal if we can measure the *evasion capability* of attacks. This not only allows us to compare the evasion power of two attacks, but also can possibly be used to compute the damage that can be caused by evasion attacks. Despite the many efforts (cf. [Šrndic and Laskov 2014] for extensive references), this aspect of security is far from understood.

5.4. Measuring obfuscation sophistication

Obfuscation based on tools such as run-time packers have been widely used by malware-writers to defeat static analysis. Despite the numerous studies that have been surveyed elsewhere [Roundy and Miller 2013; Ugarte-Pedrero et al. 2015], we understand very little about how to quantify the obfuscation capability of malware. Nevertheless, there have been some notable initial efforts. The *obfuscation prevalence* metric measures the occurrence of obfuscation in malware samples [Roundy and Miller 2013]. The *structural complexity* metric measures the runtime complexity of packers in terms of their layers, granularity etc. [Ugarte-Pedrero et al. 2015].

It is ideal if we can measure the *obfuscation sophistication* of a malware, perhaps in terms of the amount of effort that is necessary for unpacking the packed malware. One practical use is to automatically differentiate the malware samples that must be manually unpacked from those that can be automatically unpacked. One possible theoretical use is to incorporate it as a parameter in a model for analyzing the evolution of security state $S(t)$ over time t.

6. METRICS: MEASURING SITUATIONS

Situation is the comprehensive manifestation of attack-defense interactions with respect to an enterprise or computer system. These metrics measure situations via security states, security incidents, and security investments.

6.1. Measuring the evolution of security state

As illustrated in Figures 1-2, the security state vector of an enterprise system $S(t) = (s_1(t), \ldots, s_n(t))$ and the security state $s_i(t)$ of computer $c_i(t)$ both dynamically evolve as a outcome of attack-defense interactions. These metrics aim to measure the dynamic security states. However, the measurement process often incurs errors, such as false-positives and false-negatives. As a consequence, the observed state $O(t) = (o_1(t), \ldots, o_n(t))$ is often different from the true state $S(t)$. The *fraction of compromised computers* is $|\{i : i \in \{1, \ldots, n\} \wedge s_i(t) = 1\}|/n$. It has been shown that under certain circumstances, there can be some fundamental connection between the global security state and a very small number of nodes that can be monitored carefully [Xu et al. 2012a]. An alternative metric is the *probability a computer is compromised at time* t, namely $\Pr[s_i(t) = 1]$ as illustrated in Figure 1. This metric has been proposed in some recent studies that aim to quantify the security in enterprise systems (e.g., [LeMay et al. ; Xu and Xu 2012; Da et al. ; Xu 2014a; Zheng et al. 2015; Xu et al. 2015b; Xu et al. 2015a; Xu 2014b; Han et al. 2014; Xu et al. 2014a; Lu et al. 2013; Xu et al. 2012b; Xu et al. 2012a; Li et al. 2011]). These studies represent some early-stage investigations towards modeling security from a holistic perspective.

Knowing the dynamic security state can help the defender make the right decision. For example, knowing the probabilities that computers are compromised at time t,

namely $\Pr[s_i(t) = 1]$ for every i, allows the defender to use an appropriate threshold cryptographic mechanism [Desmedt and Frankel] to tolerate the compromises. However, faithful security models may require to accommodate many, if not all, of the aforementioned metrics as parameters. Moreover, it is important to know $S(t)$ and $s_i(t)$ for any t, rather than for $t \to \infty$. These impose many open problems that remain to be investigated [Xu 2014a].

6.2. Measuring security incidents

6.2.1. Measuring spatial characteristics of security incidents. Spatial characteristics of incidents can be described by the *incident rate* metric [for Internet Security 2010], which measures the fraction of computers that are successfully infected or attacked at least once during a period of time. The incident rate is often smaller than the *encounter rate*, which measures the fraction of computers that encountered some malware or attack during a period of time [Yen et al. ; Mezzour et al.]. This is because some malware encounters and attacks are defeated by the deployed defense, which can be measured through the *blocking rate* metric, namely the difference between the *encounter rate* and the *incident rate*. For example, Microsoft reports 21.2% and 19.2% of the reporting computers in question encountered malware respectively in 2013 and 2014, and malware infection rates are much less than malware encounter rates [Microsoft 2014]. Another study based on the anti-virus software reports from an enterprise of 62,884 computers shows a 15.31% malware encounter rate during a period of 4 months [Yen et al.].

These metrics may be used as alternative to the global security state $S(t)$, especially when $S(t)$ is difficult to obtain for arbitrary t (rather than for $t \to \infty$). It would be ideal if we can measure the *incident occurrence frequency* as an approximation of the number of compromised computers, namely $|\{i : i \in \{1, \ldots, n\} \wedge s_i(t) = 1\}|$, which may be represented as some mathematical functions of system features. A recent study shows the possibility of predicting data breaches from symptoms of the network in question (e.g., network mismanagement and blacklisted IP addresses) [Liu et al. 2015]. One should be cautions when using these metrics to compare the security of two systems, because they may be attacked with different attack vectors.

6.2.2. Measuring temporal characteristics of security incidents. The temporal characteristic of incidents can be described by the *delay in incident detection* [for Internet Security 2010], which measures the time between when an incident occurred and when the incident is discovered. Another metric is the *time between incidents* [for Internet Security 2010; Holm 2014], which measures the period of time between two incidents. Yet another metric is the *time-to-first-compromise* metric [Jonsson and Olovsson 1997; Madan et al. 2002; Holm 2014], which measures the duration of time between a computer starts to run and the first malware alarm is triggered on the computer, (alarm indicating detection rather than infection). A study based on a dataset of 5,602,097 malware alarms, which correspond to 203,025 malware attacks against 261,757 computers between 10/15/2009 and 8/10/2012, shows that the time-to-first-compromise follows the Pareto distribution [Holm 2014].

These metrics may be used as alternative to the global security state $S(t)$, especially when $S(t)$ is difficult to predict for arbitrary t (rather than for $t \to \infty$). It would be ideal if we can predict the *incident occurrence frequency* as an approximation of the number of compromised computers at a future time t, namely $|\{i : i \in \{1, \ldots, n\} \wedge s_i(t) = 1\}|$. One should be cautions when using these metrics to compare the security of a system during two different periods of time, because the threats would be different.

6.2.3. Measuring the damage of security incidents. The damage caused by incidents can be measured by the *cost of incidents* metric [for Internet Security 2010], which measures the count or cost of detected security incidents (i.e., successful infections or attacks)

that have occurred by in a system during a period of time. The cost of incidents may include both the direct cost (e.g., the amount of lost money) and the indirect cost (e.g., negative publicity and/or the recovery cost). A very recent survey shows that the re-mediation cost per insider attack is $500,000 [Infosecbuddy 2015].

These metrics are useful, but we understand these metrics very little. It would be ideal if we can predict the cost of incidents that occur in a future period of time, which may depend on factors such as the *delay in detection* metric mentioned above.

6.3. Measuring security investment

These metrics include *security spending* as percentage of IT budget [Chew et al. ; for Internet Security 2010] and *security budget* allocations [for Internet Security 2010]. The former is important because enterprises want to know whether their security expenditure is justified by the security performance and is comparable to other organizations' security investment. The latter indicates how the security budget is allocated for various security activities and resources.

It is important to understand the *payoff of security investment*, which requires to investigate whether or not there is some inherent relationship between the cost of security incidents and factors such as the delay in the detection of security incidents, which may depend on security investment.

7. DISCUSSION

Table 7 summarizes the taxonomy and security metrics systemized above. The column of *desirable security metrics* shows that there are big gaps between the state-of-the-art and the ultimate goals. Now we discuss some fundamental questions that must be addressed in order to bridge these gaps.

7.1. What should we measure?

First, Table 7 shows that there are big gaps between the state-of-the-art metrics (i.e., second column) and the desirable metrics (i.e., third column). For example, we used the thickness of blue bars and red arrows in Figures 1 and 2 to reflect the defense power and attack power, respectively. However, the existing security metrics do not measure this intuitive thickness metric. This resonates Pfleeger's observation [Pfleeger 2009] that metrics in the literature often correspond to "what *can be* easily measured," rather than "what *need to be* measured"—a fundamental problem that is largely open. In what follows we discuss the 4 categories systemized above and highlight what kinds of research are needed.

For measuring system vulnerabilities, we considered metrics for measuring users' vulnerabilities, password vulnerabilities, system interface-induced vulnerabilities, software vulnerabilities, and cryptographic key vulnerabilities. These classes of met-rics appear to be complete. For example, the problem of insider threats could be treated by some users' vulnerability metrics, such as *user's susceptibility to insider threats*. (The survey did not include security metrics for measuring insider threats, simply because there are no well-defined metrics of this kind despite the efforts [Azaria et al. 2014; Martinez-Moyano et al. 2008].) However, it is not clear what kind of formalism would be sufficient for reasoning about the completeness of metrics. A related open problem is: How can we define a metric that may be called *overall vulner-ability* of an enterprise or computer system, which reflects the systems' overall suscep-tibility to attacks?

For measuring defense power, we considered metrics for measuring the effectiveness of blacklisting, the attack detection power, the effectiveness of ASLR, the effectiveness of assuring CFI, and the overall defense power. It is ideal that the overall defense power metric can accommodate the other kinds of metrics. An important open problem

Table I. Summary of the taxonomy, the representative examples of security metrics systemized in the present paper, and the desirable metrics that are discussed in the text († indicates that the desirable metric is little understood)."(Avoidable via prudential engineering)" means the attacks in question can be avoided by prudential engineering.

Measuring what?	representatives of metrics systemized in the paper	desirable security metrics
Measuring system vulnerabilities		
users' vulnerabilities	user's susceptibility to phishing attacks [Sheng et al. 2010], user's susceptibility to malware infection [Lalonde Levesque et al.]	user's susceptibility to class(es) of attacks (e.g., social-engineering)†
password vulnerabilities	parameterized/statistical password guessability [Weir et al. 2010; Bonneau 2012a; Kelley et al. ; Ur et al. ; Bonneau 2012b], password entropy [Burr et al. 2006]	worst-case and average-case parameterized password guessability
system interface	attack surface [Manadhata and Wing 2011], exercised attack-surface [Nayak et al.]	interface-induced susceptibility†
software vulnerabilities		
vulnerability spatial characteristics	unpatched vulnerabilities [Chew et al. ; for Internet Security 2010], exploited vulnerabilities [Nayak et al. ; Allodi], vulnerability prevalence [Zhang et al. 2014b]	vulnerability situation awareness†
vulnerability temporal characteristics	historical (exploited) vulnerability [Al-Shaer et al. 2008; Ahmed et al. 2008], future (exploited) vulnerability [Al-Shaer et al. 2008; Ahmed et al. 2008], tendency-to-be-exploited [Sabottke et al. 2015], vulnerability lifetime [Frei and Kristensen 2010; Nappa et al. 2015; Yilek et al. ; Durumeric et al. 2014; Zhang et al. 2014b],	vulnerability vector at any time†, distribution of vulnerability lifetime†
vulnerability severity	CVSS score [of Incident Response and (FIRST)], availability of exploit [Bilge and Dumitras]	patching priority†, global damage†
cryptographic key vulnerabilities	vulnerable cryptographic keys [Yilek et al. ; Heninger et al. 2012; Durumeric et al. 2014]	(avoidable via prudential engineering)
Measuring defenses		
effectiveness of blacklisting	reaction time [Kührer et al.], coverage [Kührer et al.]	blacklisting probability†
attack detection power		
individual detection power	detection time [Rajab et al. 2005], false-positive, false-negative, true-positive, true-negative, ROC, intrusion detection capability [Gu et al.], cost [Gaffney Jr and Ulvila 2001]	detection probability†
relative detection power	relative effectiveness [Boggs and Stolfo 2011; Boggs et al.]	relative effectiveness against unknown attacks†
collective detection power	collective effectiveness [Boggs and Stolfo 2011; Morales et al. 2012; Boggs et al. ; Mohaisen and Alrawi 2014; Yardon 2014]	collective effectiveness against unknown attacks†
ASLR effectiveness	entropy [Shacham et al. 2004], effective entropy [Herlands et al. 2014]	security gain†, extra attack effort†
CFI effectiveness	CFG accuracy [Evans et al. 2015],	CFI resilience†, CFI power†
overall defense power	penetration resistance [Levin 2003], indirect MTD effectiveness [Han et al.]	resistance against unknown attacks†, direct MTD effectiveness†
Measuring threats		
threat landscape	exploit kits [Ablon et al. 2014], malicious network [Zhang et al. 2014a], rogue network [Stone-Gross et al. 2009], ISP badness [Johnson et al. 2012], control-plan reputation [Konte et al.], early-detection time [Konte et al.], cybersecurity posture [Zhan et al. 2014], sweep-time [Zhan et al. 2014], attack rate [Zhan et al. 2013; Zhan et al. 2015]	comprehensive cyber threat posture†
zero-day attacks	number of zero-day attacks [Corporation 2012], lifetime of zero-day attacks [Bilge and Dumitras], number of zero-day attack victims [Bilge and Dumitras]	susceptibility of a computer to zero-day attacks†
attack power		
power of targeted attacks	targeted threat index [Hardy et al. 2014]	susceptibility to targeted attacks†
power of botnet	botnet size [Dagon et al. 2006], botnet efficiency [Dagon et al. 2007], botnet robustness [Dagon et al. 2007]	botnet attack power†, botnet resilience with counter-countermeasures†
power of malware spreading	infection rate [Chen and Ji]	attack power†, wasted scans†
power of multi-stage attacks	necessary defense [Sheyner et al. 2002], weakest adversary [Pamula et al.], attack paths [Ritchey and Ammann 2000; Sheyner et al. 2002; Jha et al. 2002; Cheng et al. 2014], k-zero-day-safety [Wang et al.], effort-to-security-failure [Dacier et al. 1996; Ortalo et al. 1999]	multi-stage attack power†
power of evading detection	(no nontrivial metrics defined)	evasion capability†
obfuscation sophistication	obfuscation prevalence [Roundy and Miller 2013], packer structural complexity [Ugarte-Pedrero et al. 2015]	obfuscation sophistication†
Measuring situations		
security state	fractions of compromised computers at time t, probability a computer is compromised at time t [LeMay et al. ; Da et al. ; Xu 2014a]	$S(t)$ and $s_i(t)$ for any t †
security incidents		
incident spatial characteristics	incident rate [for Internet Security 2010; Microsoft 2014; Maier et al. ; Yen et al. ; Lalonde Levesque et al.], encounter rate [Yen et al. ; Mezzour et al. ; Microsoft 2014; Maier et al. ; Lalonde Levesque et al.]	incident occurrence frequency†
incident temporal characteristics	delay in incident detection [for Internet Security 2010], time between incidents [for Internet Security 2010; Jonsson and Olovsson 1997; Madan et al. 2002; Holm 2014], time-to-first-compromise [Jonsson and Olovsson 1997; Madan et al. 2002; Holm 2014]	predictive incident occurrence frequency†
incident damage	cost of incidents [for Internet Security 2010]	predictive incident damage†
security investments	security spending [Chew et al. ; for Internet Security 2010], security budget [for Internet Security 2010]	payoff of security investment†

is: Can the study of the overall defense power metric help determine the completeness of the other classes of defense power metrics? For example, is there a formalism by which we can rigorously show that the overall defense power metric can or cannot be derived from the other kinds of metrics?

For measuring threats, we considered metrics for measuring threat landscape, zero-day attacks, attack power, and obfuscation sophistication. The question is: How can we rigorously show that these metrics collectively reflect the intuitive metric that may be called *overall attack power*?

For measuring situations, we considered metrics for measuring the global security state $S(t) = (s_1(t), \ldots, s_n(t))$, security incidents, and security investments. These metrics appear to be complete because security state $S(t)$ itself does not reflect the damage of security incidents, which may or may not be positively correlated to security investments. Nevertheless, it may be helpful to integrate these metrics and the metrics for measuring system vulnerabilities and defense power into a single category, which may more comprehensively reflect the situations. This is because, for example, a user's susceptibility to attacks may vary with time t. An important open problem is: How can the defense power metrics and the attack power metrics be unified into a single framework such that, for example, a single algorithm would allow us to compute the outcome (or consequence) of the interaction between an attacker of a certain attack power and a defender of a certain defense power? This would formalize the intuitive representation of attack power and defense power as highlighted in Table 7 (third column), namely the thickness of blue bars and red arrows in Figures 1 and 2. Resolving this problem would immediately lead to a formal treatment of the arms race between the attacker and the defender, as exemplified by the discussion on the effectiveness of ASLR (Section 4.3) and CFI (Section).

Second, what would be the *complete* set of security metrics from which any useful security metric can be derived? The concept of *completeness* not only applies to the categories of security metrics, but also applies to the security metrics within each category. In order to shed light on this fundamental problem, let us look at the example of healthcare. In order to determine a person's health state, various kinds of blood tests are conducted to measure things such as glucose. The tests are subsequential to the medical research that discovered (for example) that glucose is reflective of a certain aspect of human body's health state, which answers the *what to measure* question. This example highlights that more research is needed to understand *what security metric is reflective of which security attribute or property*; otherwise, our understanding of security metrics will remain heuristic.

7.2. How to measure what we should measure in practice: random variables or numbers?

Security metrics are difficult to measure in practice. For example, vulnerabilities are dynamically discovered and the attacker may identify some zero-day vulnerabilities that are not known to the defender; the defender does not know for certain what exploits the attack possess; there may be some attack incidents that are never detected by the defender. These indicate that uncertainty is inherent to the threat model the defender is confronted with. As a consequence, security metrics should often be treated as random variables, rather than numbers. This means that we should strive to characterize the distributions of the random variables representing security metrics, rather than the means of random variables only. Another uncertainty is caused by the measurement error, such as $S(t) \neq O(t)$ as illustrated in Figure 1. This further highlights the importance of treating security metrics via random variables.

7.3. To aggregate, or not to aggregate?

The need to consider security at multiple levels/aspects of abstractions can be at least traced back to Lampson [Lampson 2006] in stopping control-flow hijack attacks [Carlini et al. 2015]. It is ideal to aggregate multiple security metrics into a single one. But, the problem is *how* [Pfleeger and Cunningham 2010; Pfleeger 2009] and *to what extent*. For the *how* part of the problem, the difficulty lies in the need to tackle the *dependence* between the security metrics, which often exists because the same system with the same vulnerabilities and the same defense deployment would get compromised by the same attack. For example, there may exist some dependence between the *security investment*, *security coverage*, and *damage of incidents*. When the mean of random variables (representing metrics) is the only concern, we can indeed aggregate multiple measures into a single one via some linear combination of them. However, the mean of random variables is just one decision-making factor, because we often need to consider the variance of the random variable in question. This requires us to tackle the dependence between the measures. For the *to what extent* part of the problem, the difficulty lies in the determination of degree of aggregation. At the highest level of aggregation, one may suggest to have metrics such as *degree of confidentiality* and *degree of integrity*. Even if it is feasible to do so, the operational utilities of such metrics would be limited because they typically measure the outcome of attack-defense interactions and do not necessarily give guidance for adjusting the defense during the course of attack-defense interactions.

7.4. What are the desirable properties of security metrics?

There have been proposals for characterizing "good" metrics, such as the following. From a conceptual perspective, a good metric should be *easy to understand* not only to researchers, but also to defense operators [Lippmann et al. 2012]. From a measurement perspective, a good metric should be *relatively easy to measure*, consistently and repeatably [Jaquith 2007]. From a utility perspective, a good metric should allow both *horizontal comparison* between enterprise systems, and *temporal comparison* (e.g., an enterprise system in the present year vs the same enterprise system in the last year) [of State 2010; Lippmann et al. 2012].

However, we also need to understand the mathematical properties security metrics should possess. These properties not only can help us differentiate the good metrics from the bad metrics, because for example we can conduct transformations between metrics. These properties may also ease the measurement process. To see this, let us look at the particular property of *additivity*. The property of *additivity* can be understood from the following example. When we talk about the measurement of mass, we are actually seeking a mapping $mass$ from $\mathcal{A} = (Objects, \text{heavier-than}, \mathbf{o})$ to $\mathcal{B} = (\mathbb{R}^+ \cup \{0\}, >, +)$, where $\mathbf{o}$ can be the "putting together" operation and $\mathbb{R}^+$ is the set of positive reals. Then, $mass$ should satisfy: (i) For two objects a and b, if a heavier-than b, then $mass(a) > mass(b)$. (ii) For any objects a and b, $mass(a\mathbf{o}b) = mass(a) + mass(b)$. Although the above (i) is relatively easy to achieve when measuring security, the above (ii), namely the additivity property, rarely holds in this domain. However, the additivity is useful because it substantially eases the measurement operations. A partial explanation for the lack of additivity is that security exhibits emergent behavior [Pfleeger and Cunningham 2010; Xu 2014b]. This suggests to investigate whether there are some additivity-like properties that can help ease the measurement of security.

7.5. What partnership is needed to tackle the problem?

The problem of security metrics may not be solved without a good partnership between the government, industry, and academia. We strongly suggest that whenever possible, any security publication in the future should include an **explicit** definition of the security metric that is relevant to the study. The reality is that most security publications did not define the security metrics they use. For example, a defense mechanism is often published with a security analysis that often makes a qualitative statement like "the mechanism can defeat a specific attack in a specific threat model." Because of the diversity in terms of the kinds of security metrics, both *bottom-up* and *top-down* approaches are important. For the bottom-up approach, each publication should strive to define, as precise as possible, the security metric and its attributes. In terms of attributes of security metrics, one can consider its *temporal* characteristics, its *spatial* characteristics, and its connection to high-level security concepts such as confidentiality, integrity and availability. Moreover, we strongly suggest that the security curriculum should have substantial materials for educating and training future generations of security researchers and practitioners with a systematic body of knowledge in security metrics. This is largely hindered by the lack of systematic treatment. We hope the present survey can catalyst the development of such materials, which however might not be possible until after security publications include serious effort at explicitly defining the security metrics in question.

While academic researchers are perhaps obligated to propose *what to measure*, they often do not have access to real data. The industry has the data, but is often prohibited from sharing data with academic researchers because of legitimate concerns (e.g., privacy). Although the government has already incentivized data sharing through projects such as PREDICT (www.predict.org), our research experience hints that semantically richer data is imperative for tackling the problem of security metrics.

8. CONCLUSION

We have presented a survey of security metrics. The survey is centered on the insight that "situations (or outcomes of cyber attack-defense interactions) are caused by certain threats (or attacks) against systems that have certain vulnerabilities (including human factors) and employ certain defenses." The resulting taxonomy contains 4 categories of security metrics: those for measuring the system *vulnerabilities*, those for measuring the *defenses*, those for measuring the *threats*, and those for measuring the *situations*. In addition to systematically reviewing the security metrics that have been proposed in the literature, we propose what we believe to be desirable security metrics. We discuss the gaps between the state-of-the-art metrics and the desirable metrics. We also discussed some fundamental problems that must be adequately addressed in order to bridge these gaps, including what academia, industry and government should do in order to catalyze the research in security metrics.

REFERENCES

Martín Abadi, Mihai Budiu, Úlfar Erlingsson, and Jay Ligatti. Control-flow Integrity. In *Proc. ACM CCS'05*. 340–353.

Lillian Ablon, Martin C. Libicki, and Andrea A. Golay. 2014. *Markets for Cybercrime Tools and Stolen Data: Hackers' Bazaar*. RAND Corporation.

M.S. Ahmed, E. Al-Shaer, and L. Khan. 2008. A Novel Quantitative Approach For Measuring Network Security. In *IEEE INFOCOM'2008*.

Ehab Al-Shaer, Latifur Khan, and Mohammad Salim Ahmed. 2008. A Comprehensive Objective Network Security Metric Framework for Proactive Security Configuration. In *Proc. CSIIRW'08*. 42:1–42:3.

Massimiliano Albanese, Sushil Jajodia, and Steven Noel. 2012. Time-efficient and Cost-effective Network Hardening Using Attack Graphs. In *Proc. IEEE DSN'12*. 1–12.

R. Albert and A. Barabasi. 2002. Statistical mechanics of complex networks. *Reviews of Modern Physics* 74 (2002), 47–97.

Luca Allodi. The Heavy Tails of Vulnerability Exploitation. In *Proc. International Symposium on Engineering Secure Software and Systems (ESSoS'15)*. 133–148.

Luca Allodi and Fabio Massacci. 2014. Comparing Vulnerability Severity and Exploits Using Case-Control Studies. *ACM Trans. Inf. Syst. Secur.* 17, 1 (2014), 1:1–1:20.

Paul Ammann, Duminda Wijesekera, and Saket Kaushik. Scalable, Graph-based Network Vulnerability Analysis. In *Proc. ACM CCS'02*. 217–224.

Jeff Arnold, Tim Abbott, Waseem Daher, Gregory Price, Nelson Elhage, Geoffrey Thomas, and Anders Kaseorg. Security Impact Ratings Considered Harmful. In *Proc. Hot Topics in Operating Systems (HotOS'09)*. 16–16.

Stefan Axelsson. The Base-rate Fallacy and Its Implications for the Difficulty of Intrusion Detection. In *Proc. ACM Conference on Computer and Communications Security (CCS'09)*. 1–7.

A. Azaria, A. Richardson, S. Kraus, and V.S. Subrahmanian. 2014. Behavioral Analysis of Insider Threat: A Survey and Bootstrapped Prediction in Imbalanced Data. *Computational Social Systems, IEEE Transactions on* 1, 2 (2014), 135–155.

Michael Backes and Stefan Nürnberger. 2014. Oxymoron: Making Fine-Grained Memory Randomization Practical by Allowing Code Sharing. In *Proc. USENIX Security Symposium*. 433–447.

Mihir Bellare, Anand Desai, E. Jokipii, and Phillip Rogaway. A Concrete Security Treatment of Symmetric Encryption. In *Proc. FOCS'97*. 394–403.

Leyla Bilge and Tudor Dumitras. Before We Knew It: An Empirical Study of Zero-day Attacks in the Real World. In *Proc. ACM CCS'12*. 833–844.

Nathaniel Boggs, Senyao Du, and SalvatoreJ. Stolfo. Measuring Drive-by Download Defense in Depth. In *Proc. RAID'14*. 172–191.

Nathaniel Gordon Boggs and Salvatore Stolfo. 2011. ALDR: A New Metric for Measuring Effective Layering of Defenses. *Fifth Layered Assurance Workshop (LAW'11)* (2011).

Joseph Bonneau. 2012a. The Science of Guessing: Analyzing an Anonymized Corpus of 70 Million Passwords. In *Proc. IEEE Symposium on Security and Privacy*. 538–552.

Joseph Bonneau. 2012b. Statistical Metrics for Individual Password Strength. In *Proc. International Conference on Security Protocols*. 76–86.

David Brumley, Pongsin Poosankam, Dawn Song, and Jiang Zheng. Automatic Patch-Based Exploit Generation is Possible: Techniques and Implications. In *Proc. 2008 IEEE Symposium on Security and Privacy*. 143–157.

William E. Burr, Donna F. Dodson, and W. Timothy Polk. 2006. *Electronic Authentication Guideline*. Technical Report Special Publication 800-63 Version 1.0.2. NIST. http://csrc.nist.gov/publications/nistpubs/800-63/SP800-63V1_0_2.pdf

A.A. Cardenas, J.S. Baras, and K. Seamon. 2006. A framework for the evaluation of intrusion detection systems. In *Proc. IEEE 2006 Symposium on Security and Privacy*.

Nicholas Carlini, Antonio Barresi, Mathias Payer, David Wagner, and Thomas R. Gross. 2015. Control-Flow Bending: On the Effectiveness of Control-Flow Integrity. In *24th USENIX Security Symposium*. 161–176.

Nicholas Carlini and David Wagner. 2014. ROP is Still Dangerous: Breaking Modern Defenses. In *Proc. USENIX Security Symposium*. 385–399.

Varun Chandola, Arindam Banerjee, and Vipin Kumar. 2009. Anomaly Detection: A Survey. *ACM Comput. Surv.* 41, 3 (2009), 15:1–15:58.

Zesheng Chen and Chuanyi Ji. Measuring Network-Aware Worm Spreading Ability. In *INFOCOM'2007*. 116–124.

Yi Cheng, Julia Deng, Jason Li, ScottA. DeLoach, Anoop Singhal, and Xinming Ou. 2014. Metrics of Security. In *Cyber Defense and Situational Awareness*. Vol. 62.

Elizabeth Chew, Marianne Swanson, Kevin Stine, Nadya Bartol, Anthony Brown, and Will Robinson. *NIST Special Publication 800-55 Revision 1: Performance Measurement Guide for Information Security*.

Symantec Corporation. April 2012. Symantec Internet security threat report, volume 17. http://www.symantec.com/threatreport/. (April 2012).

INFOSEC Research Council. 2007. Hard Problem List. http://www.infosec-research.org/docs_public/20051130-IRC-HPL-FINAL.pdf. (2007).

Gaofeng Da, Maochao Xu, and Shouhuai Xu. A new approach to modeling and analyzing security of networked systems. In *Proc. 2014 Symposium and Bootcamp on the Science of Security (HotSoS'14)*. 6.

M. Dacier, Y. Deswarte, and M. Kaâniche. 1996. IFIP SEC Information Systems Security Conference. Chapter Models and Tools for Quantitative Assessment of Operational Security, 177–186.

David Dagon, Guofei Gu, Christopher P. Lee, and Wenke Lee. 2007. A Taxonomy of Botnet Structures. In *23rd Annual Computer Security Applications Conference (ACSAC'07)*. 325–339.

D. Dagon, C. Zou, and W. Lee. 2006. Modeling Botnet Propagation Using Time Zones.. In *Proc. NDSS'06*.

Nilesh Dalvi, Pedro Domingos, Mausam, Sumit Sanghai, and Deepak Verma. 2004. Adversarial classification. In *KDD'04*. 99–108.

Lucas Davi, Ahmad-Reza Sadeghi, Daniel Lehmann, and Fabian Monrose. 2014. Stitching the Gadgets: On the Ineffectiveness of Coarse-grained Control-flow Integrity Protection. In *Proceedings of the 23rd USENIX Conference on Security Symposium*. 401–416.

Y. Desmedt and Y. Frankel. Threshold cryptosystems. In *Proc. CRYPTO 89*. 307–315.

Zakir Durumeric, James Kasten, David Adrian, J. Alex Halderman, Michael Bailey, Frank Li, Nicolas Weaver, Johanna Amann, Jethro Beekman, Mathias Payer, and Vern Paxson. 2014. The Matter of Heartbleed. In *Proc. ACM IMC'14*. 475–488.

Cynthia Dwork. Differential Privacy. In *Proc. Automata, Languages and Programming, 33rd International Colloquium (ICALP'06)*. 1–12.

Thomas C. Eskridge, Marco M. Carvalho, Evan Stoner, Troy Toggweiler, and Adrian Granados. VINE: A Cyber Emulation Environment for MTD Experimentation. In *Proc. ACM MTD'15*. 43–47.

Isaac Evans, Fan Long, Ulziibayar Otgonbaatar, Howard Shrobe, Martin Rinard, Hamed Okhravi, and Stelios Sidiroglou-Douskos. 2015. Control Jujutsu: On the Weaknesses of Fine-Grained Control Flow Integrity. In *Proc. ACM CCS'15*. 901–913.

The Center for Internet Security. 2010. The CIS Security Metrics (ver 1.1.0). http://benchmarks.cisecurity.org/downloads/metrics/. (2010).

S. Frei and T. Kristensen. Feb. 2010. THE SECURITY EXPOSURE OF SOFTWARE PORTFOLIOS. https://secunia.com/gfx/pdf/Secunia_RSA_Software_Portfolio_Security_Exposure.pdf. (Feb. 2010).

John E. Gaffney Jr and Jacob W. Ulvila. 2001. Evaluation of Intrusion Detectors: A Decision Theory Approach. In *Proceedings of the 2001 IEEE Symposium on Security and Privacy*. 50–61.

Enes Göktas, Elias Athanasopoulos, Herbert Bos, and Georgios Portokalidis. Out of Control: Overcoming Control-Flow Integrity. In *Proc. IEEE 2014 Symposium on Security and Privacy*. 575–589.

Guofei Gu, Prahlad Fogla, David Dagon, Wenke Lee, and Boris Skorić. Measuring Intrusion Detection Capability: An Information-theoretic Approach. In *Proc. AsiaCCS'06*. 90–101.

Yujuan Han, Wenlian Lu, and Shouhuai Xu. Characterizing the power of moving target defense via cyber epidemic dynamics. In *Proc. 2014 Symposium and Bootcamp on the Science of Security (HotSoS'14)*. 10.

Yujuan Han, Wnelian Lu, and Shouhuai Xu. 2014. Characterizing the Power of Moving Target Defense via Cyber Epidemic Dynamics. In *Proc. 2014 Symposium and Bootcamp on the Science of Security (HotSoS'14)*. 10:1–10:12.

Seth Hardy, Masashi Crete-Nishihata, Katharine Kleemola, Adam Senft, Byron Sonne, Greg Wiseman, Phillipa Gill, and Ronald J Deibert. 2014. Targeted threat index: Characterizing and quantifying politically-motivated targeted malware. In *Proc. USENIX Security Symposium*.

Nadia Heninger, Zakir Durumeric, Eric Wustrow, and J. Alex Halderman. 2012. Mining Your Ps and Qs: Detection of Widespread Weak Keys in Network Devices. In *Proc. USENIX Security Symposium*.

William Herlands, Thomas Hobson, and P Donovan. 2014. Effective entropy: Security-centric metric for memory randomization techniques. In *Workshop on Cyber Security Experimentation and Test*.

Hannes Holm. 2014. A Large-Scale Study of the Time Required to Compromise a Computer System. *IEEE Trans. Dependable Sec. Comput.* 11, 1 (2014), 2–15.

John Homer, Su Zhang, Xinming Ou, David Schmidt, Yanhui Du, S. Raj Rajagopalan, and Anoop Singhal. 2013. Aggregating Vulnerability Metrics in Enterprise Networks Using Attack Graphs. *J. Comput. Secur.* 21, 4 (2013), 561–597.

Adele E. Howe, Indrajit Ray, Mark Roberts, Malgorzata Urbanska, and Zinta Byrne. 2012. The Psychology of Security for the Home Computer User. In *IEEE Symposium on Security and Privacy*. 209–223.

Ling Huang, Anthony D. Joseph, Blaine Nelson, Benjamin I.P. Rubinstein, and J. D. Tygar. Adversarial Machine Learning. In *Proc. ACM AISec'11*. 43–58.

Information Assurance 09 SOAR Technology Analysis Center (IATAC). 2009. Measuring Cyber Security and Information Assurance. https://www.csiac.org/sites/default/files/cybersecurity.pdf. (2009).

Infosecbuddy. Retrieved July 6, 2015. Insider Threat: Spotlight Report (2015). http://www.infosecbuddy.com/wp-content/uploads/2015/06/Insider-Threat-Report-2015.pdf. (Retrieved July 6, 2015).

The SANS Institute. Critical Security Controls. (????).

Wayne Jansen. 2009. Directions in Security Metrics Research. http://csrc.nist.gov/publications/nistir/ir7564/nistir-7564_metrics-research.p (2009).

Andrew Jaquith. 2007. *Security Metrics: Replacing Fear, Uncertainty, and Doubt*. Addison-Wesley Professional.

S. Jha, O. Sheyner, and J. Wing. 2002. Two Formal Analys s of Attack Graphs. In *Proc. IEEE Workshop on Computer Security Foundations*. 49–59.

Benjamin Johnson, John Chuang, Jens Grossklags, and Nicolas Christin. 2012. Metrics for Measuring ISP Badness: The Case of Spam. In *Proc. Financial Cryptography and Data Security (FC'12)*. 89–97.

Erland Jonsson and Tomas Olovsson. 1997. A Quantitative Model of the Security Intrusion Process Based on Attacker Behavior. *IEEE Trans. Softw. Eng.* 23, 4 (1997), 235–245.

Patrick Gage Kelley, Saranga Komanduri, Michelle L. Mazurek, Richard Shay, Timothy Vidas, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor, and Julio Lopez. Guess Again (and Again and Again): Measuring Password Strength by Simulating Password-Cracking Algorithms. In *Proc. IEEE Symposium on Security and Privacy*. 523–537.

Maria Konte, Roberto Perdisci, and Nick Feamster. ASwatch: An AS Reputation System to Expose Bulletproof Hosting ASes. In *Proc. ACM SIGCOMM'15*. 625–638.

Marc Kührer, Christian Rossow, and Thorsten Holz. Paint It Black: Evaluating the Effectiveness of Malware Blacklists. In *Proc. Research in Attacks, Intrusions and Defenses (RAID'14)*. 1–21.

Fanny Lalonde Levesque, Jude Nsiempba, José M. Fernandez, Sonia Chiasson, and Anil Somayaji. A Clinical Study of Risk Factors Related to Malware Infections. In *Proc. ACM Conference on Computer and Communications Security (CCS'13)*. 97–108.

Butler Lampson. 2006. Practical Principles for Computer Security. (2006).

Carl E. Landwehr, Alan R. Bull, John P. McDermott, and William S. Choi. 1994. A Taxonomy of Computer Program Security Flaws. *ACM Comput. Surv.* 26, 3 (1994), 211–254.

Per Larsen, Andrei Homescu, Stefan Brunthaler, and Michael Franz. SoK: Automated Software Diversity. In *Proc. 2014 IEEE Symposium on Security and Privacy*. 276–291.

Elizabeth LeMay, Michael D. Ford, Ken Keefe, William H. Sanders, and Carol Muehrcke. Model-based Security Metrics Using ADversary VIew Security Evaluation (ADVISE). In *International Conference on Quantitative Evaluation of Systems (QEST'11)*. 191–200.

David Levin. 2003. Lessons Learned in Using Live Red Teams in IA Experiments. In *3rd DARPA Information Survivability Conference and Exposition (DISCEX-III)*. 110–119.

X. Li, P. Parker, and S. Xu. 2011. A Stochastic Model for Quantitative Security Analysis of Networked Systems. *IEEE Transactions on Dependable and Secure Computing* 8, 1 (2011), 28–43.

R.P. Lippmann, J.F. Riordan, T.H. Yu, and K.K. Watson. 2012. *Continuous Security Metrics for Prevalent Network Threats: Introduction and First Four Metrics*. Technical Report IA-3. MIT Lincoln Laboratory. https://www.ll.mit.edu/mission/cybersec/publications/publication-files/full_papers/2012_05_22_Lippmann_TechReport_FP.pdf

Yang Liu, Armin Sarabi, Jing Zhang, Parinaz Naghizadeh, Manish Karir, Michael Bailey, and Mingyan Liu. 2015. Cloudy with a Chance of Breach: Forecasting Cyber Security Incidents. In *USENIX Security Symposium*. 1009–1024.

Daniel Lowd and Christopher Meek. 2005. Adversarial learning. In *KDD'05*. 641–647.

Kangjie Lu, Chengyu Song, Byoungyoung Lee, Simon P. Chung, Taesoo Kim, and Wenke Lee. 2015. ASLR-Guard: Stopping Address Space Leakage for Code Reuse Attacks. In *Proc. ACM CCS'15*. 280–291.

Wenlian Lu, Shouhuai Xu, and Xinlei Yi. 2013. Optimizing Active Cyber Defense Dynamics. In *Proceedings of the 4th International Conference on Decision and Game Theory for Security (GameSec'13)*. 206–225.

B.B. Madan, K. Gogeva-Popstojanova, K. Vaidyanathan, and K.S. Trivedi. 2002. Modeling and quantification of security attributes of software systems. In *Proc. Dependable Systems and Networks (DSN'02)*. 505–514.

Gregor Maier, Anja Feldmann, Vern Paxson, Robin Sommer, and Matthias Vallentin. An Assessment of Overt Malicious Activity Manifest in Residential Networks. In *Proc. DIMVA'11*. 144–163.

Pratyusa K. Manadhata and Jeannette M. Wing. 2011. An Attack Surface Metric. *IEEE Trans. Software Eng.* 37, 3 (2011), 371–386.

William R. Marczak, John Scott-Railton, Morgan Marquis-Boire, and Vern Paxson. 2014. When Governments Hack Opponents: A Look at Actors and Technology. In *Proc. USENIX Security Symposium*. 511–525.

Piotr Mardziel, Mario S. Alvim, Michael Hicks, and Michael R. Clarkson. 2014. Quantifying Information Flow for Dynamic Secrets. In *Proc. 2014 IEEE Symposium on Security and Privacy*. 540–555.

Ignacio J. Martinez-Moyano, Eliot Rich, Stephen Conrad, David F. Andersen, and Thomas R. Stewart. 2008. A Behavioral Theory of Insider-threat Risks: A System Dynamics Approach. *ACM Trans. Model. Comput. Simul.* 18, 2 (2008), 7:1–7:27.

Ali Jose Mashtizadeh, Andrea Bittau, Dan Boneh, and David Mazières. 2015. CCFI: Cryptographically Enforced Control Flow Integrity. In *Proc. ACM CCS'15*. 941–951.

Ghita Mezzour, Kathleen M. Carley, and L. Richard Carley. An Empirical Study of Global Malware Encounters. In *Proc. Symposium and Bootcamp on the Science of Security (HotSoS'15)*. 8:1–8:11.

Microsoft. 2013-2014. Security Intelligence Report. http://www.microsoft.com/security/sir/default.aspx. (2013-2014).

Aleksandar Milenkoski, Marco Vieira, Samuel Kounev, Alberto Avritzer, and Bryan D. Payne. 2015. Evaluating Computer Intrusion Detection Systems: A Survey of Common Practices. *ACM Comput. Surv.* 48, 1 (2015).

Aziz Mohaisen and Omar Alrawi. 2014. Av-meter: An evaluation of antivirus scans and labels. In *Proc. DIMVA'2014*. 112–131.

Jose Andre Morales, Shouhuai Xu, and Ravi Sandhu. 2012. Analyzing Malware Detection Eciency with Multiple Anti-Malware Programs. *ASE Science Journal* 1, 2 (2012), 56–66.

Antonio Nappa, Richard Johnson, Leyla Bilge, Juan Caballero, and Tudor Dimitras. 2015. The Attack of the Clones: A Study of the Impact of Shared Code on Vulnerability Patching. In *Proc. IEEE Symposium on Security and Privacy*.

Kartik Nayak, Daniel Marino, Petros Efstathopoulos, and Tudor Dumitras. Some Vulnerabilities Are Different Than Others. In *Proc. RAID'14*.

Ajaya Neupane, Md. Lutfor Rahman, Nitesh Saxena, and Leanne Hirshfield. 2015. A Multi-Modal Neuro-Physiological Study of Phishing Detection and Malware Warnings. In *Proc. ACM CCS'15*. 479–491.

David Nicol, Bill Sanders, Jonathan Katz, Bill Scherlis, Tudor Dumitra, Laurie Williams, and Munindar P. Singh. The Science of Security 5 Hard Problems (August 2015). http://cps-vo.org/node/21590. (????).

Ben Niu and Gang Tan. 2015. Per-Input Control-Flow Integrity. In *Proc. ACM Conference on Computer and Communications Security (CCS'15)*. 914–926.

Forum of Incident Response and Security Teams (FIRST). Common Vulnerability Scoring System (CVSS). (????).

United States Department of State. 2010. iPost: Implementing Continuous Risk Monitoring at the Department of State. http://www.state.gov/documents/organization/156865.pdf. (2010).

Rodolphe Ortalo, Yves Deswarte, and Mohamed Kaâniche. 1999. Experimenting with Quantitative Evaluation Tools for Monitoring Operational Security. *IEEE Trans. Softw. Eng.* 25, 5 (Sept. 1999), 633–650.

Joseph Pamula, Sushil Jajodia, Paul Ammann, and Vipin Swarup. A Weakest-adversary Security Metric for Network Configuration Security Analysis. In *Proc. ACM QoP'06*. 31–38.

S.L. Pfleeger and R.K. Cunningham. 2010. Why Measuring Security Is Hard. *Security Privacy, IEEE* 8, 4 (July 2010), 46–54.

Shari Lawrence Pfleeger. 2009. Useful Cybersecurity Metrics. *IT Professional* 11, 3 (2009), 38–45.

Cynthia Phillips and Laura Painton Swiler. 1998. A Graph-based System for Network-vulnerability Analysis. In *Proc. 1998 Workshop on New Security Paradigms (NSPW '98)*. 71–79.

Achintya Prakash and Michael P. Wellman. Empirical Game-Theoretic Analysis for Moving Target Defense. In *Proc. ACM MTD'15*. 57–65.

Moheeb Abu Rajab, Fabian Monrose, and Andreas Terzis. 2005. On the Effectiveness of Distributed Worm Monitoring. In *Proceedings of the 14th Conference on USENIX Security Symposium*.

G. Ramalingam. 1994. The Undecidability of Aliasing. *ACM Trans. Program. Lang. Syst.* 16, 5 (1994), 1467–1471.

Ronald W. Ritchey and Paul Ammann. 2000. Using Model Checking to Analyze Network Vulnerabilities. In *Proc. IEEE Symposium on Security and Privacy*. 156–165.

M. Rostami, F. Koushanfar, and R. Karri. 2014. A Primer on Hardware Security: Models, Methods, and Metrics. *Proc. IEEE* 102, 8 (Aug 2014), 1283–1295.

Kevin A. Roundy and Barton P. Miller. 2013. Binary-code Obfuscations in Prevalent Packer Tools. *ACM Comput. Surv.* 46, 1 (July 2013), 4:1–4:32.

Carl Sabottke, Octavian Suciu, and Tudor Dumitras. 2015. Vulnerability Disclosure in the Age of Social Media: Exploiting Twitter for Predicting Real-World Exploits. In *24th USENIX Security Symposium*. 1041–1056.

William H. Sanders. 2014. Quantitative Security Metrics: Unattainable Holy Grail or a Vital Breakthrough within Our Reach? *IEEE Security & Privacy* 12, 2 (2014), 67–69.

Tobias Schneider and Amir Moradi. 2015. Leakage Assessment Methodology - a clear roadmap for side-channel evaluations. *IACR Cryptology ePrint Archive* (2015).

Bruce Schneier. 2000. *Secrets & Lies: Digital Security in a Networked World*. John Wiley & Sons, Inc.

National Science and Technology Council. 2011. Trustworthy Cyberspace: Strategic Plan for the Federal Cybersecurity Research and Development Program. https://www.nitrd.gov/SUBCOMMITTEE/csia/Fed_Cybersecurity_RD_Strategic_Plan_2011.pdf. (2011).

Hovav Shacham, Matthew Page, Ben Pfaff, Eu-Jin Goh, Nagendra Modadugu, and Dan Boneh. 2004. On the Effectiveness of Address-space Randomization. In *Proc. ACM Conference on Computer and Communications Security (CCS'04)*.

Steve Sheng, Mandy Holbrook, Ponnurangam Kumaraguru, Lorrie Faith Cranor, and Julie Downs. 2010. Who Falls for Phish?: A Demographic Analysis of Phishing Susceptibility and Effectiveness of Interventions. In *Proc. CHI'10*. 373–382.

Oleg Sheyner, Joshua Haines, Somesh Jha, Richard Lippmann, and Jeannette M. Wing. 2002. Automated Generation and Analysis of Attack Graphs. In *Proc. 2002 IEEE Symposium on Security and Privacy*. 273–284.

Reza Shokri, George Theodorakopoulos, Jean-Yves Le Boudec, and Jean-Pierre Hubaux. Quantifying Location Privacy. In *Proc. IEEE Symposium on Security and Privacy*. 247–262.

Kevin Z. Snow, Fabian Monrose, Lucas Davi, Alexandra Dmitrienko, Christopher Liebchen, and Ahmad-Reza Sadeghi. 2013. Just-In-Time Code Reuse: On the Effectiveness of Fine-Grained Address Space Layout Randomization. In *Proc. IEEE Symposium on Security and Privacy*. 574–588.

B. Stone-Gross, C. Kruegel, K. Almeroth, A. Moser, and E. Kirda. 2009. FIRE: FInding Rogue nEtworks. In *Computer Security Applications Conference (ACSAC'09)*. 231–240.

Laszlo Szekeres, Mathias Payer, Tao Wei, and Dawn Song. 2013. SoK: Eternal War in Memory. In *Proc. IEEE Symposium on Security and Privacy*. 48–62.

Adrian Tang, Simha Sethumadhavan, and Salvatore Stolfo. 2015. Heisenbyte: Thwarting Memory Disclosure Attacks Using Destructive Code Reads. In *Proc. ACM CCS'15*. 256–267.

Caroline Tice, Tom Roeder, Peter Collingbourne, Stephen Checkoway, Úlfar Erlingsson, Luis Lozano, and Geoff Pike. 2014. Enforcing Forward-edge Control-flow Integrity in GCC & LLVM. In *Proc. USENIX Security Symposium*. 941–955.

Xabier Ugarte-Pedrero, Davide Balzarotti, Igor Santos, and Pablo Garcia Bringas. 2015. SoK: Deep Packer Inspection: A Longitudinal Study of the Complexity of Run-Time Packers. In *2015 IEEE Symposium on Security and Privacy*. 659–673.

Blase Ur, Sean M. Segreti, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor, Saranga Komanduri, Darya Kurilova, Michelle L. Mazurek, William Melicher, and Richard Shay. Measuring Real-World Accuracies and Biases in Modeling Password Guessability. In *Proc. 24th USENIX Security Symposium*.

Nedim Šrndic and Pavel Laskov. 2014. Practical Evasion of a Learning-Based Classifier: A Case Study. In *Proceedings of the 2014 IEEE Symposium on Security and Privacy*. 197–211.

Lingyu Wang, Sushil Jajodia, Anoop Singhal, and Steven Noel. k-Zero Day Safety: Measuring the Security Risk of Networks against Unknown Attacks.

Matt Weir, Sudhir Aggarwal, Michael Collins, and Henry Stern. 2010. Testing Metrics for Password Creation Policies by Attacking Large Sets of Revealed Passwords. In *Proc. CCS'10*. 162–175.

Li Xu, Zhenxin Zhan, Shouhuai Xu, and Keying Ye. 2014b. An evasion and counter-evasion study in malicious websites detection. In *IEEE Conference on Communications and Network Security (CNS'14)*. 265–273.

Maochao Xu, Gaofeng Da, and Shouhuai Xu. 2015a. Cyber Epidemic Models with Dependences. *Internet Mathematics* 11, 1 (2015), 62–92.

Maochao Xu and Shouhuai Xu. 2012. An Extended Stochastic Model for Quantitative Security Analysis of Networked Systems. *Internet Mathematics* 8, 3 (2012), 288–320.

Shouhuai Xu. 2014a. Cybersecurity Dynamics. In *Proc. Symposium and Bootcamp on the Science of Security (HotSoS'14)*. 14:1–14:2.

Shouhuai Xu. 2014b. Emergent Behavior in Cybersecurity. In *Proceedings of the 2014 Symposium and Bootcamp on the Science of Security (HotSoS'14)*. 13:1–13:2.

Shouhuai Xu, Wenlian Lu, and Hualun Li. 2015b. A Stochastic Model of Active Cyber Defense Dynamics. *Internet Mathematics* 11, 1 (2015), 23–61.

Shouhuai Xu, Wenlian Lu, and Li Xu. 2012a. Push- and Pull-based Epidemic Spreading in Arbitrary Networks: Thresholds and Deeper Insights. *ACM Transactions on Autonomous and Adaptive Systems (ACM TAAS)* 7, 3 (2012), 32:1–32:26.

Shouhuai Xu, Wenlian Lu, Li Xu, and Zhenxin Zhan. 2014a. Adaptive Epidemic Dynamics in Networks: Thresholds and Control. *ACM Transactions on Autonomous and Adaptive Systems (ACM TAAS)* 8, 4 (2014), 19.

Shouhuai Xu, Wenlian Lu, and Zhenxin Zhan. 2012b. A Stochastic Model of Multivirus Dynamics. *IEEE Trans. Dependable Sec. Comput.* 9, 1 (2012), 30–45.

D. Yardon. May 4, 2014. Symantec Develops New Attack on Cyberhacking. http://www.wsj.com/articles/SB10001424052702303417104579542140235850578. (May 4, 2014).

Ting-Fang Yen, Victor Heorhiadi, Alina Oprea, Michael K. Reiter, and Ari Juels. An Epidemiological Study of Malware Encounters in a Large Enterprise. In *Proceedings of ACM Conference on Computer and Communications Security (CCS'14)*. 1117–1130.

Scott Yilek, Eric Rescorla, Hovav Shacham, Brandon Enright, and Stefan Savage. When Private Keys Are Public: Results from the 2008 Debian OpenSSL Vulnerability. In *Proc. ACM IMC'09*. 15–27.

Kara Zaffarano, Joshua Taylor, and Samuel Hamilton. A Quantitative Framework for Moving Target Defense Effectiveness Evaluation. In *Proc. ACM MTD'15*. 3–10.

Zhenxin Zhan, Maochao Xu, and Shouhuai Xu. 2013. Characterizing Honeypot-Captured Cyber Attacks: Statistical Framework and Case Study. *IEEE Transactions on Information Forensics and Security* 8, 11 (2013), 1775–1789.

Zhenxin Zhan, Maochao Xu, and Shouhuai Xu. 2014. A Characterization of Cybersecurity Posture from Network Telescope Data. In *Proceedings of The 6th International Conference on Trustworthy Systems (InTrust'14)*.

Zhenxin Zhan, Maochao Xu, and Shouhuai Xu. 2015. Predicting Cyber Attack Rates With Extreme Values. *IEEE Transactions on Information Forensics and Security* 10, 8 (2015), 1666–1677.

Jing Zhang, Zakir Durumeric, Michael Bailey, Mingyan Liu, and Manish Karir. 2014a. On the Mismanagement and Maliciousness of Networks. In *Proc. NDSS'14*.

Su Zhang, Xinwen Zhang, and Xinming Ou. 2014b. After We Knew It: Empirical Study and Modeling of Cost-effectiveness of Exploiting Prevalent Known Vulnerabilities Across IaaS Cloud. In *Proceedings of the 9th ACM Symposium on Information, Computer and Communications Security (ASIA CCS '14)*. 317–328.

Ren Zheng, Wenlian Lu, and Shouhuai Xu. 2015. Active Cyber Defense Dynamics Exhibiting Rich Phenomena. In *Proc. 2015 Symposium and Bootcamp on the Science of Security (HotSoS'15)*. 2:1–2:12.