

Quantifying Cybersecurity Effectiveness of Dynamic Network Diversity

Huashan Chen, Hasan Cam, and Shouhuai Xu

Abstract—The deployment of monoculture software stacks can have devastating consequences because a single attack can compromise all of the vulnerable computers in cyberspace. This one-vulnerability-affects-all phenomenon will continue until after software stacks are diversified, which is well recognized by the research community. However, existing studies mainly focused on investigating the effectiveness of software diversity at the building-block level (e.g., whether two independent implementations indeed exhibit independent vulnerabilities); the effectiveness of enforcing network-wide software diversity is little understood, despite its importance in possibly helping justify investment in software diversification. As a first step towards ultimately tackling this problem, we propose a systematic framework for modeling and quantifying the cybersecurity effectiveness of network diversity, including a suite of cybersecurity metrics. We also present an agent-based simulation to empirically demonstrate the usefulness of the framework. We draw a number of insights, including the surprising result that proactive diversity is effective under very special circumstances, but reactive-adaptive diversity is much more effective in most cases.

Index Terms—Software diversity, network diversity, security quantification, metrics, agent-based simulation, cybersecurity dynamics

1 INTRODUCTION

SOFTWARE monoculture enables the automatic amplification of cyber attack damages because vulnerabilities are replicated network-wide or even cyberspace-wide [1], [2]. As a consequence, a single exploit may allow an attacker to compromise many programs. In order to cope with the problem, researchers have proposed diversifying program implementations [3], [4], [5], leading to the notion of *software diversity*. There are various flavors of software diversity, such as: *N-version programming* (i.e., a program specification has multiple independent implementations [6], [7]); *natural diversity* (e.g., multiple browsers incurred by market competition [8]); *compiler-based diversification* (i.e. random executables generated from a given source code [9], [10], [11], [12], [13]); and *software runtime environment diversification* (e.g., address space layout randomization [14], [15], [16], [17], instruction set randomization [18], [19], system calls randomization [20], and replicated execution [21]).

It is intuitive that employing software diversity could increase security. For example, the U.S. Navy developed the RHIMES system [22] to enhance security of shipboard systems, by introducing diversity to each programmable logic controller. However, real-world software diversity is often employed in an ad hoc fashion, which can be justified by how different OSes (e.g., Windows vs. various kinds of Unix) and browsers (e.g., Safari vs. Firefox vs. Chrome) are employed in practice. One exception is the investigation of employing software diversity to enhance Byzantine Fault-Tolerance (BFT), namely how to employ software diversity

in the replica implementations so that they do not contain common vulnerabilities [23], [24], [25], [26], [27], [28]. This is important because the theoretical fault-tolerance guarantee can be ruined otherwise. Another exception is the investigation of employing software diversity in detecting cyber attacks by leveraging the behavioral discrepancy between diversified replicas [29], [30], [31]. Despite these studies, some fundamental questions remain open, such as: *How should software diversity be employed in practice to amplify, if not maximize, security?*

The preceding question leads to the notion of *network diversity*, which deals with the employment of diversified program implementations in network-wide software stacks [1], [2], [3], [32], [33]. A simpler version of the notion may be called *static diversity*, where diversified implementations are employed once and for all (i.e., unchanged after initial employment). There are studies on optimizing static diversity via some flavor of *graph coloring* algorithm; by treating colors as diversified implementations, the research problem is to minimize defective edges (i.e., adjacent nodes have the same color or run the same implementation) [32], [33], [34], [35], [36], [37]. There are also studies on quantifying the effectiveness of *static diversity* [37], [38], [39], [40], [41].

Despite these studies, there is no systematic understanding on the network-wide effectiveness of employing software diversity, for multiple reasons. First, most studies use coarse-grained models (i.e., treating each computer as a unit) and do not consider attack-defense interactions. Second, it is not clear how to quantify the network-wide cybersecurity effectiveness of employing network diversity. Third, it is not clear how network diversity should be dynamically employed, leading to the notion of *dynamic diversity*. Addressing these problems can deepen our understanding of software diversity, help decision-makers determine whether or not to invest in software diversity, and guide practitioners in intelligently employing diversified

H. Chen is with the Department of Computer Science, University of Texas at San Antonio, San Antonio, TX, USA. H. Cam is with Best Buy, Richfield, MN, USA. S. Xu is with the Department of Computer Science, University of Colorado Colorado Springs, Colorado Springs, Colorado, USA; this work was partly done when he was affiliated with University of Texas at San Antonio. Correspondence: sxu@uccs.edu

implementations in real-world cyber defense operations.

Our contributions. We make three contributions. First, we propose a framework for modeling and quantifying the network-wide cybersecurity of enforcing network diversity. The framework considers *fine-granularity* by distinguishing applications and operating systems. The framework considers the time dimension, which allows us to investigate *dynamic* diversity; that is, the employment of diversified implementations in the network evolves over time.

Second, we propose a suite of metrics to quantify the network-wide effectiveness of employing diversity, including: (i) *time-to-succeed*, which measures how long it takes an attacker to break a defender's goal (if possible); (ii) *attacker slow-down*, which measures the extent an attacker is slowed down by network diversity; (iii) *attack worst damage*, which measures the damage an attacker can cause in the worst case; (iv) *attack extra cost*, which measures the extra investment the defense imposes on an attacker in order to break the defender's goal; (v) *vulnerability tolerance*, which measures the upper bound of vulnerabilities that can be tolerated when achieving the defender's goal; (vi) *average operational cost*, which measures the average fraction of programs that re-deploy dynamic diversity. These metrics may be of independent value.

Third, we demonstrate the usefulness of the framework and metrics by presenting a multi-agent simulation study with multiple network diversity strategies: monoculture (the baseline with no diversity), static diversity (for comparison purposes), proactive (periodically re-diversifying network-wide software stacks), reactive-adaptive (re-diversifying the network stacks in response to detected attacks), and hybrid (of the last two). According to the simulation study, we draw a number of insights, including: (i) In terms of attacker slow-down, reactive-adaptive diversity is the most effective strategy and the initial diversity configuration matters. (ii) In order to reduce the attack worst damage, different diversity strategies should be used in different parameter regimes. (iii) Reactive-adaptive diversity leads to a higher vulnerability-tolerance than proactive diversity does. (iv) Proactive diversity improves security only when dynamic diversity is widely re-employed at a high frequency, which however incurs a high operational cost. (v) The more the diversified implementations, the higher the attacker slow-down, the higher the attack extra cost, and the higher the vulnerability tolerance. This is especially true for reactive-adaptive diversity. Note that these findings may not be universally true because they are derived from the parameter settings used in the simulation study.

Paper outline. Section 2 presents the framework. Section 3 describes the simulation study. Section 4 reviews related prior work. Section 5 discusses the limitations of the present study. Section 6 concludes the paper. Table 1 summarizes the notations used throughout the paper.

2 THE FRAMEWORK

Consider a network, where each node represents a computer. Each computer has a software stack, running an operating system (OS) in the kernel space and some application(s) in the user space. Each *program*, OS and application alike, may have diversified implementations (e.g.,

\mathcal{A}, \mathcal{D}	attacker \mathcal{A} and defender \mathcal{D}
n, i, j	n is the number of computers in a network, $\text{app}_{i,j}$ is the j -th application running in computer $i \in [1, n]$
ℓ, z	ℓ is the number of phases of \mathcal{A} 's strategy, $z \in [1, \ell]$
h, k	h is the number of different programs, $k \in [1, h]$
X_k	X_k is the number of diversified implementations of program k
Q_k	software quality of diversified implementations of the k -th program (functionality), $Q_k \in [0, 1]$ interpreted as probability of containing vulnerability
G	$G = (V, E)$ is the communication graph of a network (rather than the network's physical topology), where $v \in V$ represents a program (application or operating system)
C_t	$C_t : V \rightarrow \text{SW}$ is the network diversity configuration at time t , where SW is the set of diversified implementations of all programs
ϕ_t	$\phi_t : C_t(V) \rightarrow 2^{\text{VUL}}$ is the mapping from diversified programs to the vulnerabilities present in them at time t . $\Phi_t = (\phi_t(C_t(v)))_{v \in V}$ is the vulnerabilities in the network.
$s_{v,t}$	$s_{i,t} \in \{0, 1, 2\}$ is the cybersecurity state of node (i.e., program running at) $v \in V$ at time t (0: <i>vulnerable</i> ; 1: <i>compromised</i> ; 2: <i>invulnerable</i>). Vector $S_t = (s_{v,t})_{v \in V}$
$s_{i,t}$	$s_{i,t} \in \{0, 1, 2\}$ is the cybersecurity state of computer i at time t , similarly defined as $s_{v,t}$. Vector $S'_t = (s_{i,t})_{i \in [1, n]}$
Σ_t	$\Sigma_t = (G, C_t, \Phi_t, S_t)$ is the cybersecurity situation of a network at time t
$\Sigma_{\mathcal{A},t}$	$\Sigma_{\mathcal{A},t} = (G_{\mathcal{A},t}, C_{\mathcal{A},t}, \Phi_{\mathcal{A},t}, S_{\mathcal{A},t})$ is \mathcal{A} 's perception of the target network Σ_t at time t
$\Omega_{\mathcal{A},t}$	$\Omega_{\mathcal{A},t} = (\omega_{v,\mathcal{A},t})_{v \in V_{\mathcal{A},t}}$ is \mathcal{A} 's goal at time t at the <i>program</i> level, where $\omega_{v,\mathcal{A},t} \in \{\perp, 1\}$; $\Omega'_{\mathcal{A},t} = (\omega_{i,\mathcal{A},t})_{i \in [1, n]}$ is \mathcal{A} 's goal at time t at the <i>computer</i> level, where $\omega'_{i,\mathcal{A},t} \in \{\perp, 1\}$
$\Gamma_{\mathcal{A}}$	$\Gamma_{\mathcal{A}} = \{\gamma_{\mathcal{A},1}, \dots, \gamma_{\mathcal{A},\ell}\}$ is attacker's strategy of ℓ phases
$\Delta_{\mathcal{A}}$	$\Delta_{\mathcal{A}} = (\Delta_{\mathcal{A},z})_{z \in [1, \ell]}$ is an attacker's capability where $\Delta_{\mathcal{A},z} = \{\psi_{\mathcal{A},z,1}, \dots, \psi_{\mathcal{A},z,m_z}\}$ is the exploits that are applicable and available to the attacker at phase z
$\mathcal{F}_{\mathcal{A},t}$	$\mathcal{F}_{\mathcal{A},t}$ is the attacker's decision-making algorithm to make an attack plan $\Lambda_{\mathcal{A}} = (\lambda_{\mathcal{A},1}, \dots, \lambda_{\mathcal{A},\ell})$ at time t , where $\lambda_{\mathcal{A},z} \in \Delta_{\mathcal{A},z}$ for $1 \leq z \leq \ell$
$\Sigma_{\mathcal{D},t}$	$\Sigma_{\mathcal{D},t} = (G_{\mathcal{D},t}, C_{\mathcal{D},t}, \Phi_{\mathcal{D},t}, S_{\mathcal{D},t})$ is \mathcal{D} 's perception of Σ_t at time t
$\Omega_{\mathcal{D},t}$	$\Omega_{\mathcal{D},t} = (\omega_{v,\mathcal{D},t})_{v \in V}$ is \mathcal{D} 's goal at time t at the <i>program</i> level, where $\omega_{v,\mathcal{D},t} \in [0, 1]$; $\Omega'_{\mathcal{D},t} = (\omega_{i,\mathcal{D},t})_{i \in [1, n]}$ is \mathcal{D} 's goal at time t at the <i>computer</i> level, where $\omega_{i,\mathcal{D},t} \in [0, 1]$
$\Gamma_{\mathcal{D}}$	$\Gamma_{\mathcal{D}} = \{\gamma_{\mathcal{D},1}, \dots, \gamma_{\mathcal{D},j}\}$ is defender's strategy
$\eta_{\mathcal{D},1}$	$\eta_{\mathcal{D},1}$ represents the proportion of nodes in V employing diversified implementations
$\eta_{\mathcal{D},2}$	$\eta_{\mathcal{D},2}$ represents the <i>frequency</i> at which the diversified implementations will be dynamically re-employed
$\eta_{\mathcal{D},3}$	$\eta_{\mathcal{D},3}$ represents the <i>condition</i> under which diversified implementations are re-employed
$\Delta_{\mathcal{D}}$	$\Delta_{\mathcal{D}} = (\Delta_{\mathcal{D},k})_{k \in [1, h]}$ is defender's capability, where $\Delta_{\mathcal{D},k} = \{\delta_{\mathcal{D},k,1}, \dots, \delta_{\mathcal{D},k,X_k}\}$ is a set of X_k diversified implementations of program k ;
$\mathcal{F}_{\mathcal{D},t}$	the decision-making algorithm to make a defense plan $\Lambda_{\mathcal{D},t} = (\delta_{\mathcal{D},t}(v))_{v \in V}$ at time t
τ	the compromise probability of computer $i \in [1, n]$ at time $t \in [0, T]$ that can be tolerated by the defender
TTS	$\text{TTS}(\mathcal{A}, \mathcal{D}) = \min\{t : \text{cc}(t) > \tau\}$ measures how long it takes for attacker \mathcal{A} to break defender \mathcal{D} 's mission goal τ
ASD	$\text{ASD}_{\mathcal{D},q} = \text{TTS}(\mathcal{A}, \mathcal{D}_q) - \text{TTS}(\mathcal{A}, \mathcal{D}_1)$ measures the extent at which \mathcal{A} is slowed down by defense strategy $\gamma_{\mathcal{D},q}$
AWD	$\text{AWD}_{G,h,X,Q}(\mathcal{A}, \mathcal{D}, T) = \max\{\text{cc}(t) : t \leq T\}$ measures attack worst damage during mission lifetime $t \in [0, T]$
AEC	$\text{AEC}_{\mathcal{D},q} = \text{AI}(\gamma_{\mathcal{D},q}) - \text{AI}(\gamma_{\mathcal{D},1})$ measures the number of extra exploits \mathcal{A} needs to obtain to make $\text{AWD}_{G,h,X,Q}(\mathcal{A}, \mathcal{D}_q, T) > \tau$ against defense strategy $\gamma_{\mathcal{D},q}$
VT	$\text{VT}_{\mathcal{D},q} = \max\{Q : \text{AWD}_{G,h,X,Q}(\mathcal{A}, \mathcal{D}_q, T) \leq \tau\}$ captures the upper bound of tolerable vulnerabilities such that \mathcal{D} can still achieve its mission goal τ in lifetime $[0, T]$
AOC	$\text{AOC}_{\mathcal{D},q}(T) = \sum_{t=1}^T \text{OC}_{\mathcal{D},q}(t)/T$ is the average operational cost to achieve the defender's goal; $\text{AOC}_{\mathcal{D},q}^{\max} = \max\{\text{AOC}_{\mathcal{D},q}(T) : \text{AWD}_{G,h,X,Q}(\mathcal{A}, \mathcal{D}_q, T) \leq \tau\}$ and $\text{AOC}_{\mathcal{D},q}^{\min} = \min\{\text{AOC}_{\mathcal{D},q}(T) : \text{AWD}_{G,h,X,Q}(\mathcal{A}, \mathcal{D}_q, T) \leq \tau\}$ are maximum and minimum $\text{AOC}_{\mathcal{D},q}(T)$, respectively.

TABLE 1: Summary of main notations.

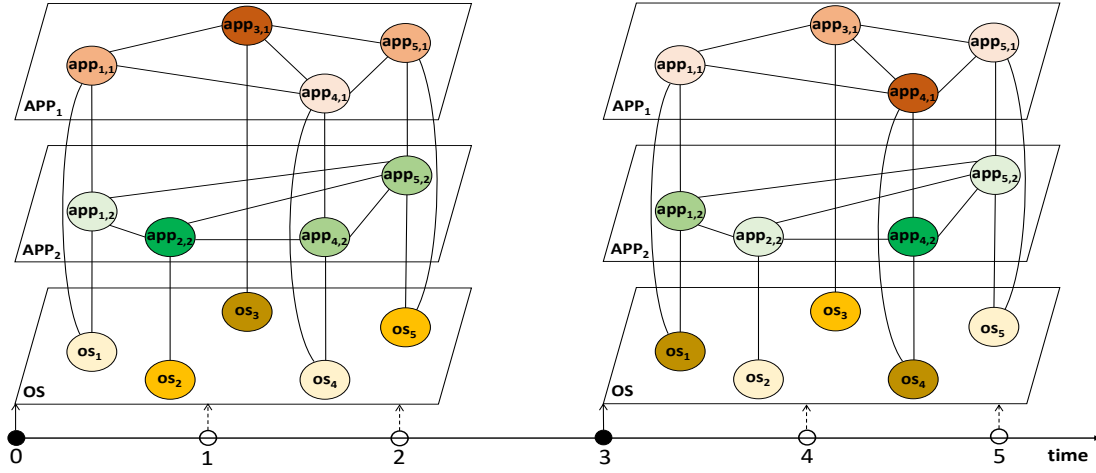


Fig. 1: Illustration of dynamic diversity in a network of 5 computers ($1 \leq i \leq 5$). The time horizon shown is $t = 0, 1, \dots, 5$.

Safari vs. Firefox vs. Chrome for browser). These diversified implementations may be obtained by using some of the diversification methods mentioned above (e.g., N -version programming or natural diversity via market competition). Each program may or may not be vulnerable.

2.1 Intuition of Dynamic Network Diversity

Given diversified implementations of programs, the aforementioned notion of *static* network diversity is to employ diversified programs in the computers' software stacks, where the employment (or configuration) will not change during the time horizon of interest. We initiate the aforementioned *dynamic* network diversity, which aims to dynamically employ diversified programs at the computers' software stacks, where the employment (or configuration) does change during the time horizon of interest.

Figure 1 illustrates the idea via a network of 5 computers, denoted by $1 \leq i \leq 5$. The time horizon is $t = 0, \dots, 5$. Each computer runs an OS, and we use os_i to represent the OS running in computer i . There are two applications, denoted by APP_j for $1 \leq j \leq 2$ (e.g., APP_1 is browser and APP_2 is email). Computers 1, 4 and 5 run both applications; computer 2 runs APP_2 ; computer 3 runs APP_1 . We use $app_{i,j}$ to represent the application APP_j running in computer i . (i) There are three implementations of OS (e.g., Windows vs. Linux vs. macOS), which are respectively indicated by three colors. For example, at time $t = 0, 1, 2$, computers 1 and 4 run the same OS, causing os_1 and os_4 to have the same color; computers 2 and 5 run another OS, causing os_2 and os_5 to have another color; computer 3 runs yet another OS, causing os_3 to have a different color. (ii) There are three implementations of APP_1 (e.g., Safari vs. Firefox vs. Chrome), which are respectively indicated by three colors. For example, at time $t = 0, 1, 2$, computers 1 and 5 run the same implementation of APP_1 , causing $app_{1,1}$ and $app_{5,1}$ to have the same color; computers 3 and 4 run two other implementations of APP_1 , causing $app_{3,1}$ and $app_{4,1}$ to have different colors; computer 2 does not run APP_1 . (iii) There are three implementations of APP_2 (e.g., Outlook vs. Thunderbird vs. eM Client), which are respectively indicated by three colors. For example, at time

$t = 0, 1, 2$, computers 4 and 5 run the same implementation of APP_2 , causing $app_{4,2}$ and $app_{5,2}$ to have the same color; computers 1 and 2 run two other implementations of APP_2 , causing $app_{1,2}$ and $app_{2,2}$ to have different colors; computer 3 does not run APP_2 .

Figure 1 illustrates dynamic diversity as follows. At time $t = 0$, the network-wide diversity is configured to run a certain combination of specific implementations of OS and applications as indicated by colors. The solid arrow at $t = 0$ indicates that a new (i.e., initial) diversity configuration is employed. This configuration remains unchanged for $t = 1, 2$, as indicated by dashed arrows at $t = 1, 2$. At time $t = 3$, the network-wide diversity is re-configured to run another combination of the diversified implementations of OS and applications. The solid arrow at $t = 3$ indicates this new employment. The configuration remains unchanged for $t = 4, 5$, as indicated by dashed arrows at $t = 4, 5$.

2.2 Framework Overview

Figure 2 highlights the Cyber Observation-Orientation-Decision-Action Loop (COODAL) framework for describing attack-defense interactions in a network. COODAL adapts the military operation concept of OODA Loop [42] to cyberspace. At a high level, the network is abstracted as a communication graph. Both attacker and defender have their own Observation (collecting data), Orientation (analyzing the collected data while leveraging relevant intelligence or information), Decision (determining what to do), and Action (executing the decision). We consider a discrete-time model with a finite time horizon $t = 0, 1, \dots, T$, where T is the defender's mission lifetime (i.e., certain security requirements must be satisfied in order to achieve mission assurance in time interval $[0, T]$). For simplicity, we assume an attack or defense action takes effect instantly (if effective). When desired, this assumption can be eliminated by explicitly modeling the delay for an action to take effect (e.g., the framework can be extended to accommodate that an attack or defense action occurs at time t and takes effect at $t + 1$).

We stress that the framework does not make restrictive assumptions. For example, we do not make any restriction on how vulnerabilities may be distributed among the diversified implementations by equally accommodating the

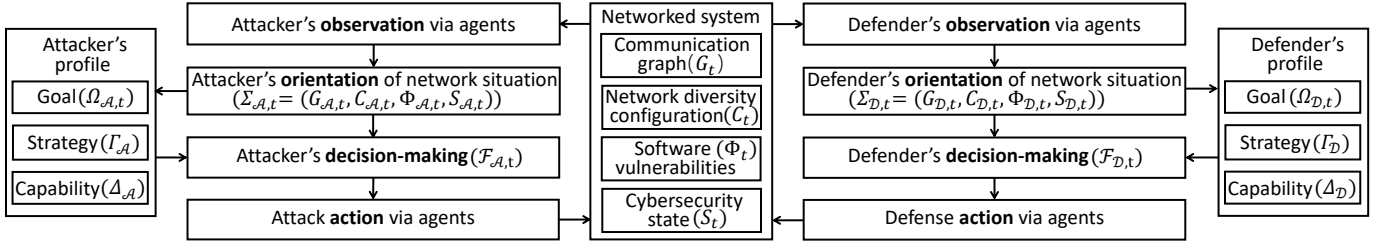


Fig. 2: The COODAL framework for characterizing the effectiveness of dynamic network diversity.

scenarios where diversified implementations may or may not have common vulnerabilities. When the attacker has an exploit against a vulnerability that is common to multiple programs, the attacker can compromise all of these vulnerable programs. A successful attack against an OS implies a successful attack against the applications running on top of the OS; a successful attack against an application paves a way for the attacker to attack the OS beneath the application (e.g., privilege escalation).

2.3 Modeling Network

Modeling computers' software stacks. A network (e.g., enterprise or mission network) consists of n interconnected computers or devices. Each computer runs a *software stack*, which has two layers: application and operating system (OS). We treat any software running in the *user space* as application program. By contrast, an OS runs in the *kernel space*. Consistent with the notations used above, we use APP to denote the universe of application programs. As mentioned above, computer i runs one or multiple applications, denoted by $\text{app}_{i,j}$, where $1 \leq i \leq n$, j is the index of the application, and $\text{app}_{i,j} \in \text{APP}$. We use OS to denote the universe of OSes, and use os_i to denote the specific OS running in computer i , where $\text{os}_i \in \text{OS}$. Let h denote the number of different programs running in the network.

Modeling communications. We explicitly model the communications between applications because they can be leveraged by the attacker to wage attacks. There are two types of communications: *intra-computer* and *inter-computer* [43]. Intra-computer communications are conducted by the programs running in a computer and can be represented as *edges* in the terminology of Graph Theory. For example, Figure 1 shows that the two applications running in computer 1 (i.e., $\text{app}_{1,1}$ and $\text{app}_{1,2}$) are designed to communicate with each other, and that both $\text{app}_{1,1}$ and $\text{app}_{1,2}$ can communicate with os_1 (e.g., for making system calls). Inter-computer communications are conducted by applications running in different computers and can be represented as *edges*. Note that OSes may not communicate with each other. We distinguish intra-computer and inter-computer communications for two reasons: (i) they are often leveraged by the attacker for different purposes — the former for privilege escalation and the latter for lateral movement between computers; and (ii) they are defended by different security mechanisms — the former is defended by host-based mechanisms (e.g., intrusion prevention) and the latter is defended by network-based mechanisms (e.g., firewall).

Formalizing the preceding discussion, we naturally obtain the notion of *communication graph* $G = (V, E)$, where

each vertex or node $v \in V$ represents (and runs) a program and each edge $(u, v) \in E$ represents that a pair of nodes are permitted to communicate with each other. Note that a computer is represented by a set of nodes in G because it runs a set of programs. Since *programs* run at *nodes* in G , we use *programs* and *nodes* interchangeably to make succinct statements. We use the term *edges* in the standard way to indicate *undirected* graphs; a communication graph can be *directed* in principle, which is however rare in practice. In the example illustrated in Figure 1, where have $|V| = 13$ (i.e., 4 nodes running APP₁, 4 nodes running APP₂, and 5 nodes running OS), and the edge set E is as illustrated.

We stress that communication graph $G = (V, E)$ is different from a networking-induced graph, for two reasons. First, a vertex or node in a communication graph represents a program running in a computer and each edge represents the communication between two programs. In contrast, a node in a networking-induced graph often represents a computer, which can run multiple programs. Second, a communication graph can encode access control policies, which may regulate which programs are (not) allowed to communicate with which other programs running in the same computer or different computers. This means that a communication graph may not be a complete graph because some programs may only be allowed to communicate with some of the others. Whereas, a networking-induced graph cannot encode access control policies and would be a complete graph because any computer can communicate with any other computer as long as they are routable.

In this paper we assume a communication graph $G = (V, E)$ is *time-independent*, meaning that the applications running in a computer are fixed. This means that the applications running in a computer do not change, but their specific implementations may change over time. This is plausible because we focus on quantifying the effectiveness of dynamic network diversity. Nevertheless, $G = (V, E)$ can be extended to time-dependent $G_t = (V_t, E_t)$ when desired.

Modeling network diversity configuration. Let SW denote the set of diversified implementations of the h programs running in the network, including applications and OSes. The *dynamic* network diversity configuration at time t refers to the mapping from the set of programs to their specific implementations. This mapping can be described by a function $C_t : V \rightarrow \text{SW}$ such that $C_t(v)$ is a specific implementation of the program running at node $v \in V$ at time t . For example, Figure 1 shows that there are 3 different programs (i.e., APP₁, APP₂, OS); each program has 3 diversified implementations (indicated by 3 different colors) and thus $|\text{SW}| = 9$, where computer 1 runs a

specific implementation of each of the three programs at $t \in [0, 2]$. Note that the preceding representation is general enough to accommodate the *static* network diversity, which corresponds to $C_0 = C_1 = \dots = C_T$, and the *monolithic* software stack, which corresponds to each program having exactly one implementation (i.e., $|SW|$ equals the number of different programs and $C_0 = C_1 = \dots = C_T$).

Modeling software vulnerabilities. The software running in a computer, application and OS alike, may contain vulnerabilities. We use VUL to denote the universe of software vulnerabilities that may be present in the software stacks running in the computers of the network in question. Since vulnerabilities are associated with the nodes in V , we use function $\phi_t : C_t(V) \rightarrow 2^{VUL}$ to describe the set of vulnerabilities that are present in the implementations of application and OS programs running in the computers, where $\phi_t(C_t(v)) = \emptyset$ means the implementation running at node v is not vulnerable. We use vector $\Phi_t = (\phi_t(C_t(v)))_{v \in V}$ to denote the *ground-truth* vulnerabilities that are present in the software running in computers at time t . Note that this ground-truth vulnerability set may or may not be known to the attacker or defender, and that defining $\phi_t : C_t(V) \rightarrow 2^{VUL}$ is equivalent to defining it as a function from the software set SW to 2^{VUL} . We choose the former because it simplifies subsequent discussion (in regards to dynamic network diversity). Let $Q_k \in [0, 1]$, $1 \leq k \leq h$, denote the quality of program, which is measured by the ratio of the number of vulnerable implementations to the total number of diversified implementations of program k . The parameter Q_k , which is interpreted as the probability that program k is vulnerable, would depend on the technologies that are employed to reduce vulnerabilities in the course of software development (e.g., vulnerability detection [44], [45], [46], [47], [48]). The smaller the Q_k , the higher the diversity quality of program k .

Modeling network-wide cybersecurity state and situation. Cybersecurity state and situation can be defined at the *program* of $v \in V$ and at the computer level. The former is suitable for modeling and simulation purposes, and the latter is more suitable for cyber defense operation and management purposes. So, we consider both.

At the program level, we consider a communication graph $G = (V, E)$, which abstracts a network as described above. We use $s_{v,t} \in \{0, 1, 2\}$ to denote the state of node $v \in V$ at time t : $s_{v,t} = 0$ means v (i.e., the program running at v) is *vulnerable* but not compromised, namely that v contains a vulnerability but the vulnerability is not exploited by the attacker and the underlying OS is not compromised; $s_{v,t} = 1$ means v is *compromised* either because its vulnerability is exploited or because the underlying OS is compromised; and $s_{v,t} = 2$ means v is *invulnerable* (i.e., containing no vulnerabilities) and the underlying OS is not compromised. The *program-level* network-wide cybersecurity state at time t is represented by vector $S_t = (s_{v,t})_{v \in V}$.

At the computer level, we say computer i is *vulnerable* (or $s_{i,t} = 0$) if any program running in the computer is vulnerable, *compromised* (or $s_{i,t} = 1$) if any program running in the computer is compromised, and *invulnerable* (or $s_{i,t} = 2$) if all programs running in the computer are invulnerable. Similarly, we can define *computer-level* network-wide cyber-

security state at time t as $S'_t = (s_{i,t})_{i \in [1,n]}$, where S'_t can be derived from G and S_t . We further define the vector $((vc(t), cc(t), ic(t)))$ to succinctly describe the computer-level network-wide cybersecurity effect at time t , where $vc(t) = |\{i : s_{i,t} = 0\}|/n$ is the fraction of *vulnerable* computers, $cc(t) = |\{i : s_{i,t} = 1\}|/n$ is the fraction of *compromised* computers, and $ic(t) = |\{i : s_{i,t} = 2\}|/n$ is the fraction of *invulnerable* computers. Note that $cc(t) + vc(t) + ic(t) = 1$.

The network-wide *cybersecurity situation* can be described by $\Sigma_t = (G, C_t, \Phi_t, S_t)$, or a tuple of the communication graph, the network diversity configuration, the set of vulnerabilities associated with each implementation, and the network-wide cybersecurity state. Note that we do not mention S'_t because it can be derived from S_t .

2.4 Modeling Attacker

We model an attacker \mathcal{A} (i.e., threat model) with five attributes: *knowledge* (what \mathcal{A} knows about a network), *goal* (what \mathcal{A} attempts to achieve), *strategy* (what strategy \mathcal{A} uses), *capability* (what exploits \mathcal{A} possesses), and *decision-making* (i.e., what decision-making algorithms \mathcal{A} uses). Intuitively, \mathcal{A} with a certain knowledge attempts to achieve a goal by leveraging some capabilities to compromise some nodes or computers according to some strategies.

Attacker's knowledge. We define attacker \mathcal{A} 's knowledge as vector $\Sigma_{\mathcal{A}} = (\Sigma_{\mathcal{A},t})_{t \in [0,T]}$ such that $\Sigma_{\mathcal{A},t} = (G_{\mathcal{A},t}, C_{\mathcal{A},t}, \Phi_{\mathcal{A},t}, S_{\mathcal{A},t})$ is \mathcal{A} 's perception of the target network $\Sigma_t = (G, C_t, \Phi_t, S_t)$ at time t , where $G_{\mathcal{A},t} = (V_{\mathcal{A},t}, E_{\mathcal{A},t}) \subseteq G$ is the attacker's perception of G , $C_{\mathcal{A},t}$ is the attacker's perception of C_t , $\Phi_{\mathcal{A},t}$ is the attacker's perception of Φ_t , and $S_{\mathcal{A},t}$ is the attacker's perception of S_t . Note that $S_{\mathcal{A},t} = S_t$ because the attacker knows which programs are compromised by the attacker itself. Note also that the notion of *initial compromise* can be modeled as part of the attacker's knowledge at time $t = 0$, because it describes which nodes $v \in V$ are compromised at $t = 0$. We use $\text{IniComp} = \{v \in V : s_{v,0} = 1\}$ to denote the set of programs that are compromised at $t = 0$.

Attacker's goal. We define attacker \mathcal{A} 's goal as vector $\Omega_{\mathcal{A}} = (\Omega_{\mathcal{A},t})_{t \in [0,T]}$ where vector $\Omega_{\mathcal{A},t} = (\omega_{v,\mathcal{A},t})_{v \in V_{\mathcal{A},t}}$ is \mathcal{A} 's goal at time t , where $\omega_{v,\mathcal{A},t} \in \{\perp, 1\}$, $\omega_{v,\mathcal{A},t} = \perp$ means that \mathcal{A} does not care about the state of $v \in V_{\mathcal{A},t}$ at time t , $\omega_{v,\mathcal{A},t} = 1$ means \mathcal{A} attempts to make $v \in V$ compromised at time t . Since $\Omega'_{\mathcal{A},t} = (\omega_{i,\mathcal{A},t})_{i \in [1,n]}$ can be derived from $\Omega_{\mathcal{A},t}$, in what follows we do not have to mention $\Omega'_{\mathcal{A},t}$ explicitly. This representation is flexible because it can accommodate intuitive attack goals, such as: (i) attempting to compromise a fixed set of nodes or computers at time t ; (ii) attempting to compromise as many nodes or computers as possible at time t ; and (iii) attempting to cumulatively compromise as many nodes or computers as possible at time t .

Attacker's strategy. Inspired by the state-of-the-art industrial characterization of sophisticated cyber attacks, such as the Cyber Kill Chain [49] and Mitre's ATT&CK [50], we propose abstracting them into attacker's strategy. Since strategy is often fixed for $t \in [0, T]$, we specify attacker \mathcal{A} 's strategy via $\ell \geq 1$ phases, denoted by $\Gamma_{\mathcal{A}} = \{\gamma_{\mathcal{A},1}, \dots, \gamma_{\mathcal{A},\ell}\}$. For example, we have $\ell = 7$ for the Cyber Kill Chain and $\ell = 12$ for Mitre's ATT&CK version 7. Since there are

different kinds of strategies, we will demonstrate how to use a specific strategy in our case study.

Attacker's capability. Given an ℓ -phase strategy $\Gamma_{\mathcal{A}}$, attacker \mathcal{A} 's capability at phase z , where $z \in [1, \ell]$, is defined as a set of m_z exploits applicable at phase z and available to \mathcal{A} , denoted by $\Delta_{\mathcal{A},z} = \{\psi_{\mathcal{A},z,1}, \dots, \psi_{\mathcal{A},z,m_z}\}$. We define \mathcal{A} 's capability as vector $\Delta_{\mathcal{A}} = (\Delta_{\mathcal{A},z})_{z \in [1, \ell]}$. Since \mathcal{A} 's capability depends on exploits, we define \mathcal{A} 's investment as $AI = \sum_{z=1}^{\ell} \sum_{q=1}^{m_z} \text{cost}_{\mathcal{A},z,q}$, where $\text{cost}_{\mathcal{A},z,q}$ is the cost to obtain the q -th exploit that is applicable at phase z .

Attacker's decision-making. At time t , attacker \mathcal{A} uses a decision-making algorithm $\mathcal{F}_{\mathcal{A},t}$ to make an attack plan, which specifies the utilization of some of \mathcal{A} 's capabilities, denoted by $\Lambda_{\mathcal{A},t} = (\lambda_{\mathcal{A},1}, \dots, \lambda_{\mathcal{A},\ell})$ where $\lambda_{\mathcal{A},z} \in \Delta_{\mathcal{A},z}$ for $1 \leq z \leq \ell$. This can be denoted by

$$\lambda_{\mathcal{A},z} \leftarrow \mathcal{F}_{\mathcal{A},t}(\Sigma_{\mathcal{A},t}, \Omega_{\mathcal{A},t}, \Gamma_{\mathcal{A}}, \Delta_{\mathcal{A},z}), \quad (1)$$

where $\Sigma_{\mathcal{A},t} = (G_{\mathcal{A},t}, C_{\mathcal{A},t}, \Phi_{\mathcal{A},t}, S_{\mathcal{A},t})$ is \mathcal{A} 's knowledge at time t as described above, $\Omega_{\mathcal{A},t}$ is \mathcal{A} 's goal at time t , $\Gamma_{\mathcal{A}}$ is \mathcal{A} 's strategy, and $\Delta_{\mathcal{A},z}$ is \mathcal{A} 's capability at phase z of the attack strategy. Note the Eq. (1) can be extended to consider multiple exploits (rather than a single exploit $\lambda_{\mathcal{A},z}$) that may be used in an appropriate manner (e.g., sequential).

Putting the description together, we denote attacker $\mathcal{A} = (\mathcal{A}_t)_{t \in [0, T]}$ with $\mathcal{A}_t = (\Sigma_{\mathcal{A},t}, \Omega_{\mathcal{A},t}, \Gamma_{\mathcal{A}}, \Delta_{\mathcal{A}}, \mathcal{F}_{\mathcal{A},t})$.

2.5 Modeling Defender

Similar to the attacker (or threat) model, we describe a defender \mathcal{D} via five attributes: *knowledge* (what \mathcal{D} knows), *goal* (what \mathcal{D} aims to achieve), *strategy*, *capability* (what tools \mathcal{D} possesses), and *decision-making* (what algorithms \mathcal{D} uses).

Defender's knowledge. We define defender \mathcal{D} 's knowledge as vector $\Sigma_{\mathcal{D}} = (\Sigma_{\mathcal{D},t})_{t \in [0, T]}$ such that $\Sigma_{\mathcal{D},t} = (G_{\mathcal{D},t}, C_{\mathcal{D},t}, \Phi_{\mathcal{D},t}, S_{\mathcal{D},t})$ is \mathcal{D} 's perception of the ground-truth situation $\Sigma_t = (G, C_t, \Phi_t, S_t)$ at time t . In the case of *full* knowledge, we have $\Sigma_{\mathcal{D},t} = \Sigma_t$, meaning the defender knows everything about the ground-truth situation. In the more realistic case of *partial* knowledge, the defender only knows: (i) the communication graph G , namely $G_{\mathcal{D}} = G$ because the network is managed by the defender; (ii) the network configuration C_t , namely $C_{\mathcal{D},t} = C_t$ because the defender decides which nodes run which specific implementations; (iii) some information about the ground-truth vulnerabilities associated with the programs, namely $\Phi_{\mathcal{D},t} \subseteq \Phi_t$ because the defender may not know the 0-day ones that are known to the attacker; and (iv) some noisy information about the network's ground-truth cybersecurity state S_t because of the false-positives and/or false-negatives in measuring or inferring cybersecurity states.

Defender's goal. Corresponding to the program-level vs. computer-level distinction, \mathcal{D} 's goal can be defined at two levels. At the *program* level of $v \in V$, we define \mathcal{D} 's goal as vector $\Omega_{\mathcal{D}} = (\Omega_{\mathcal{D},t})_{t \in [0, T]}$ such that $\Omega_{\mathcal{D},t} = (\omega_{v,\mathcal{D},t})_{v \in V}$ is \mathcal{D} 's goal at time t , where $\omega_{v,\mathcal{D},t} \in [0, 1]$ is the tolerable probability that program running at $v \in V$ is compromised at time t . For example, $\omega_{v,\mathcal{D},t} = 0$ means that a successful attack against v cannot be tolerated; $\omega_{v,\mathcal{D},t} = 0.5$ can be interpreted as that compromise of v for at most 50% of the time can be tolerated. At the *computer* level, we define \mathcal{D} 's goal

as vector $\Omega'_{\mathcal{D}} = (\Omega'_{\mathcal{D},t})_{t \in [0, T]}$ such that $\Omega'_{\mathcal{D},t} = (\omega_{i,\mathcal{D},t})_{i \in [1, n]}$ is \mathcal{D} 's goal at time t , where $\omega_{i,\mathcal{D},t} \in [0, 1]$ is the tolerable probability that computer i is compromised at time t . One computer-level goal of particular interest is: $\omega_{i,\mathcal{D},t} \leq 1/3$ for $i \in [1, n]$ and $t \in [0, T]$; it describes cyber defense using Byzantine fault-tolerance techniques to tolerate the compromise of a certain threshold of computers [51]. Since $\Omega'_{\mathcal{D}}$ can be derived from $\Omega_{\mathcal{D}}$, we do not have to mention $\Omega'_{\mathcal{D}}$ except when we discuss computer-level effectiveness.

Defender's strategy. We specify defender \mathcal{D} 's strategies in employing network diversity as a set $\Gamma_{\mathcal{D}} = \{\gamma_{\mathcal{D},1}, \dots, \gamma_{\mathcal{D},5}\}$. For example, $\gamma_{\mathcal{D},1}$ represents monoculture software stacks (i.e., the baseline strategy), $\gamma_{\mathcal{D},2}$ represents static diversity, $\gamma_{\mathcal{D},3}$ represents proactive diversity with fixed intervals, $\gamma_{\mathcal{D},4}$ represents reactive-adaptive diversity where the employment is triggered by security alerts, and $\gamma_{\mathcal{D},5}$ represents hybrid diversity (i.e., a combination of proactive diversity and reactive-adaptive diversity). A strategy can be accompanied by *some* of the following parameters. (i) the *proportion* of nodes in V (re-)employing diversified implementations (e.g., all or some nodes), denoted by $\eta_{\mathcal{D},1}$; (ii) the *frequency* at which the diversified implementations will be dynamically re-employed at the nodes, denoted by $\eta_{\mathcal{D},2}$; (iii) the *condition* under which diversified implementations are re-employed, denoted by $\eta_{\mathcal{D},3}$. As an example showing that not every parameter is relevant to every strategy, we note that the preceding (ii) and (iii) are not relevant to the static diversity strategy; this can be indicated by setting $\eta_{\mathcal{D},2} = \text{NULL}$ and $\eta_{\mathcal{D},3} = \text{NULL}$.

Defender's capability. We define defender \mathcal{D} 's capability as the diversified implementations that are available to \mathcal{D} . Recall that h different programs running in the network (including both applications and operation systems) and each program may have diversified implementations. \mathcal{D} 's capability with respect to program k , where $1 \leq k \leq h$, is a set of X_k diversified implementations, denoted by $\Delta_{\mathcal{D},k} = \{\delta_{\mathcal{D},k,1}, \dots, \delta_{\mathcal{D},k,X_k}\}$. We define \mathcal{D} 's capability as the vector of diversified implementations that are available to \mathcal{D} , denoted by $\Delta_{\mathcal{D}} = (\Delta_{\mathcal{D},k})_{k \in [1, h]}$. Since \mathcal{D} 's capability depends on the diversified implementations, we define the defender's investment as $DI = \sum_{k=1}^h \sum_{w=1}^{X_k} \text{cost}_{\mathcal{D},k,w}$, where $\text{cost}_{\mathcal{D},k,w}$ is the cost for obtaining the w -th diversified implementation of program k .

Defender's decision-making. At time t , \mathcal{D} uses decision-making algorithm $\mathcal{F}_{\mathcal{D},t}$ to make a defense plan $\Lambda_{\mathcal{D},t} = (\delta_{\mathcal{D},t}(v))_{v \in V}$, which specifies how to employ the diversified implementations of the h programs at nodes $v \in V$ at time t . Formally, $\Lambda_{\mathcal{D},t}$ is the output of the defender's decision-making algorithm $\mathcal{F}_{\mathcal{D},t}$ on a number of inputs, denoted by

$$\Lambda_{\mathcal{D},t} \leftarrow \mathcal{F}_{\mathcal{D},t}(\Sigma_{\mathcal{D},t}, \Omega_{\mathcal{D},t}, \Gamma_{\mathcal{D}}, \Delta_{\mathcal{D}}), \quad (2)$$

where $\Sigma_{\mathcal{D},t}$ is the defender's perception of the ground-truth situation $\Sigma_t = (G, C_t, \Phi_t, S_t)$ at time t , $\Omega_{\mathcal{D},t}$ is the defender's goal, $\Gamma_{\mathcal{D}}$ is defender's strategy as described above, and $\Delta_{\mathcal{D}}$ is the defender's capability.

Putting the description together, we denote defender $\mathcal{D} = (\mathcal{D}_t)_{t \in [0, T]}$ with $\mathcal{D}_t = (\Sigma_{\mathcal{D},t}, \Omega_{\mathcal{D},t}, \Gamma_{\mathcal{D}}, \Delta_{\mathcal{D}}, \mathcal{F}_{\mathcal{D},t})$.

2.6 Modeling Effect of Dynamic Diversity

Modeling local effect of dynamic diversity at the node level. At any point in time, a program (i.e., node $v \in V$) is in one of the following three states: *vulnerable*, meaning that the program contains a vulnerability but the vulnerability has not been exploited by the attacker; *invulnerable*, meaning that the program contains no vulnerability; and *compromised*, either because the program contains a vulnerability that has been exploited, or because the underlying OS is compromised (causing any application program running on top of it to be compromised, no matter whether the program contains vulnerability or not).

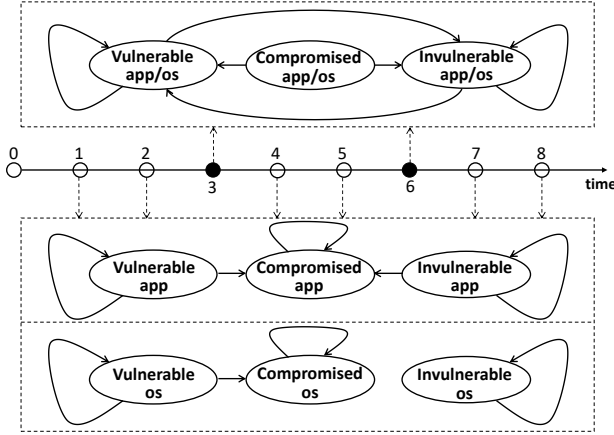


Fig. 3: Modeling the local effect of dynamic network diversity at the program level (i.e., program running at $v \in V$).

Figure 3 highlights the effect of employing dynamic diversity at the program level. Suppose at time $t = 3, 6$ dynamic diversity is employed such that a different implementation of a program (i.e., application or OS) is employed at time t to replace the implementation employed at time $t - 1$. The employment of dynamic diversity at time t leads to one of the following state transitions (see upper half of Figure 3): (i) a vulnerable program is replaced with another vulnerable or an invulnerable program of the same functionality; (ii) an invulnerable program is replaced with a vulnerable or another invulnerable program; (iii) a compromised program is replaced with a vulnerable or invulnerable program. However, a vulnerable or invulnerable program is never replaced with a compromised program because a successful attack only occurs to a vulnerable program that has been employed and exploited (as diversified implementations are stored in a secure environment).

At any other point in time than the employment of dynamic diversity (e.g., $t \in \{1, 2, 4, 5, 7, 8\}$ as shown in Figure 3), the security state transition of application programs is different from that of OSes. For applications, the state transitions are: (i) a vulnerable application program stays vulnerable or become compromised, either because its vulnerability is exploited or the underlying OS is compromised; (ii) a compromised application program stays in the compromised state because we focus on the defense based on dynamic diversity, without considering reactive defense that may detect and clean up the compromised application programs; (iii) an invulnerable application program stays invulnerable or becomes compromised because

the underlying OS is compromised. For OSes, the state transitions are: (i) a vulnerable OS stays vulnerable or become compromised because its vulnerability is exploited; (ii) a compromised OS stays in the compromised state because we do not consider reactive defense; (iii) an invulnerable OS stays invulnerable.

Quantifying the global effect of dynamic network diversity at the computer level. We quantify the global effect at the computer level via the following metrics.

Definition 1 (time-to-succeed or TTS). This metric measures how long it takes attacker \mathcal{A} to break defender \mathcal{D} 's goal specified by a program-level vector $\Omega_{\mathcal{D},t} = (\omega_{v,\mathcal{D},t})_{v \in V}$ or computer-level vector $\Omega'_{\mathcal{D},t} = (\omega_{i,\mathcal{D},t})_{i \in [1,n]}$, where $t \in [0, T]$, and $\omega_{v,\mathcal{D},t} \in [0, 1]$ ($\omega_{i,\mathcal{D},t} \in [0, 1]$) is the compromise probability of program v (computer i) at time t that can be tolerated. As mentioned above, a defense goal of particular interest is $\omega_{i,\mathcal{D},t} \leq 1/3$ for computers $i \in [1, n]$ at any time $t \in [0, T]$ because such compromises can be tolerated by Byzantine fault-tolerant techniques [51]. Formally, we define $\text{TTS}(\mathcal{A}, \mathcal{D}) = \min\{t : \text{cc}(t) > \tau\}$, where $\text{cc}(t)$ is the fraction of compromised computers (i.e., attack damage) at time t .

The *time-to-succeed* metric is reminiscent of the well-known *mean-time-to-compromise* metric. However, the former is defined as a random variable with respect to a specific goal (e.g., breaking defender's goal), where randomness is rooted in different attack strategies, capabilities and decision-making algorithms. In contrast, the latter is defined as a number (or the mean value of a random variable).

Definition 2 (attacker slow-down or ASD). Consider attacker \mathcal{A} (i.e., a fixed threat model) that attempts to break defender \mathcal{D} 's goal $\Omega'_{\mathcal{D}} = (\Omega'_{\mathcal{D},t})_{t \in [0,T]}$, where $\Omega'_{\mathcal{D},t} = (\omega_{i,\mathcal{D},t})_{i \in [1,n]}$ and, for concreteness, $\omega_{i,\mathcal{D},t} \leq \tau$ is the tolerable compromise probability for every computer $i \in [1, n]$ and $t \in [0, T]$, while recalling that $\tau = 1/3$ corresponds to assuring the assumption that is needed by Byzantine fault-tolerant techniques [51]. We define *attacker slow-down* as $\text{ASD}_{\mathcal{D},q} = \text{TTS}(\mathcal{A}, \mathcal{D}_q) - \text{TTS}(\mathcal{A}, \mathcal{D}_1)$, to measure the extent at which \mathcal{A} is slowed down by the employment of network diversity defense strategy $\gamma_{\mathcal{D},q}$ (denoted by \mathcal{D}_q), where $2 \leq q \leq 5$ in this study, when compared with the baseline strategy $\gamma_{\mathcal{D},1}$ of monoculture software stacks (denoted by \mathcal{D}_1).

Definition 3 (attack worst damage or AWD). This metric measures how much damage an attacker \mathcal{A} can cause in the worst case during $t \in [0, T]$, namely the maximum fraction of compromised computers at any time $t \in [0, T]$. Formally, we define *attack worst damage* as $\text{AWD}_{G,h,X,Q}(\mathcal{A}, \mathcal{D}, T) = \max\{\text{cc}(t) : t \leq T\}$ where G, h, X, Q are defined above.

Definition 4 (attack extra cost or AEC). Consider attacker \mathcal{A} (i.e., a fixed threat model) attempting to break defender \mathcal{D} 's goal $\Omega'_{\mathcal{D}} = (\Omega'_{\mathcal{D},t})_{t \in [0,T]}$, where $\Omega'_{\mathcal{D},t} = (\omega_{i,\mathcal{D},t})_{i \in [1,n]}$ and, for concreteness, $\omega_{i,\mathcal{D},t} \leq \tau$ is the tolerable compromise probability for computer $i \in [1, n]$ and $t \in [0, T]$. We define *attack extra cost* (AEC) metric, $\text{AEC}_{\mathcal{D},q} = \text{Al}(\gamma_{\mathcal{D},q}) - \text{Al}(\gamma_{\mathcal{D},1})$, to measure the number of extra exploits \mathcal{A} needs to obtain (i.e., extra investment)

such that $AWD_{G,h,X,Q}(\mathcal{A}, \mathcal{D}_q, T) > \tau$ when \mathcal{D} employs strategy $\gamma_{\mathcal{D},q}$ (i.e., $AI(\gamma_{\mathcal{D},q})$) than the baseline strategy employing monoculture software stacks (i.e., $AI(\gamma_{\mathcal{D},1})$)

Inspired by the notion of *fault-tolerance*, we define *vulnerability-tolerance* metric to capture the upper bound vulnerabilities in the network that can be tolerated by defender \mathcal{D} in achieving its goal $\Omega_{\mathcal{D}}$. The importance of this metric can be seen as follows. (i) When each computer vulnerable with probability at least $1/3 + \epsilon$ for some ϵ , the attacker able to compromise all of them can render Byzantine fault-tolerance techniques useless. (ii) When dynamic network diversity is employed, the attacker may only be able to compromise $1/3$ of the computers because the diversity configuration may have changed before the attacker compromises all of the vulnerable computers. Intuitively, the larger the ϵ , the higher the vulnerability tolerance.

Definition 5 (vulnerability-tolerance or VT). Consider defender goal $\omega_{i,\mathcal{D},t} \leq \tau$ for every computer $i \in [1, n]$ and each time $t \in [0, T]$ and a fixed diversification quality $Q \in [0, 1]$ for diversified implementation (when applying the same quality-enhancement techniques), we define $VT_{\mathcal{D},q} = \max\{Q : AWD_{G,h,X,Q}(\mathcal{A}, \mathcal{D}_q, T) \leq \tau\}$ with $AWD_{G,h,X,Q}(\mathcal{A}, \mathcal{D}_q, T)$ specified in Definition 3.

Definition 6 (average operational cost or AOC). We define the *operational cost* incurred by defender's strategy $\gamma_{\mathcal{D},q}$ at time t , denoted by $OC_{\mathcal{D},q}(t) \in [0, 1]$, as the fraction of programs that are replaced at time t , where $2 \leq q \leq 5$. We define the *average operational cost* up to time T as $AOC_{\mathcal{D},q}(T) = \sum_{t=1}^T OC_{\mathcal{D},q}(t)/T$, which refers to the average fraction of programs that are re-installed at each time $t \in [1, T]$. Intuitively, the larger the $AOC_{\mathcal{D},q}(T)$, the higher the operational cost. Given a defender's goal $\omega_{i,\mathcal{D},t} \leq \tau$ for computer $i \in [1, n]$ and $t \in [0, T]$, we define $AOC_{\mathcal{D},q}^{\max} = \max\{AOC_{\mathcal{D},q}(T) : AWD_{G,h,X,Q}(\mathcal{A}, \mathcal{D}_q, T) \leq \tau\}$ and $AOC_{\mathcal{D},q}^{\min} = \min\{AOC_{\mathcal{D},q}(T) : AWD_{G,h,X,Q}(\mathcal{A}, \mathcal{D}_q, T) \leq \tau\}$ as the maximum and minimum average operational cost to meet the defender's goal, respectively.

Figure 4 illustrates the relationship between the metrics. The effectiveness metrics TTS, ASD, AWD, AEC and VT depend on the attacker's goal in τ , mission lifetime T , and attack damage $cc(t)$. The $cc(t)$ depends on the attack investment AI , the defense investment DI , and the defender's operational cost AOC . The attack investment AI depends on the number of attack phases (ℓ), the number of exploits applicable at each phase (m_z), and the cost to obtain exploits ($cost_{\mathcal{A},z,q}$). The defense investment depends on the number of programs (h), the number of diversified implementations of each program (X_k), and the cost to obtain each implementation ($cost_{\mathcal{D},k,w}$). The defender's average operational cost AOC depends on the proportion $\eta_{\mathcal{D},1}$ and frequency $\eta_{\mathcal{D},2}$ of re-employment.

2.7 Research Questions (RQs)

In order to characterize the effectiveness of dynamic network diversity, we propose investigating the following RQs, where RQ1-RQ3 correspond to defender's gains, RQ4-RQ5 corresponds to defender's cost.

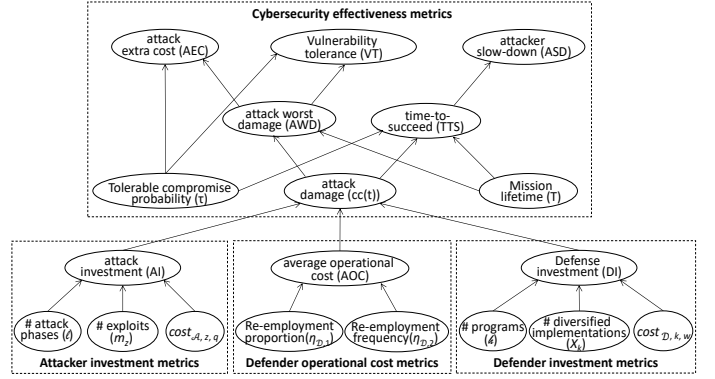


Fig. 4: Illustration of the relationships between the metrics, where “ $X \rightarrow Y$ ” means X is a factor in determining Y .

- RQ1: To what extent can dynamic network diversity slow down the attacker?
- RQ2: How much extra cost can dynamic network diversity impose on the attacker?
- RQ3: To what extent can dynamic network diversity increase the defender's vulnerability tolerance?
- RQ4: To what extent can dynamic network diversity increase the defender's average operational cost?
- RQ5: Is it true that the more diversified implementations the better?

3 SIMULATION EXPERIMENTS

The preceding framework is meant to be as realistic as we can be. This means that it does not make strong assumptions that would warrant analytical treatment. This explains why we pursue simulation-based empirical study to answer the RQs. We adopt agent-based simulation because agents can conduct activities concurrently, which can mimic real-world attack-defense interaction better than sequential simulation. We implement the agent-based simulation via multithreading, where each active local agent, attack and defense alike, is instantiated as one thread such that multiple events can take place concurrently. The simulation experiment is conducted on a computer with 32-CPU and 128GB RAM in the Python environment. In order to accommodate the randomness in the simulation experiment, we conduct 500 simulation runs for each experiment and take their average as the result. In the discrete-time models, all events occur at discrete points in time. In our experiment, we make every attack or defense activity take effect instantly, which can be extended to accommodate any delays if desired.

3.1 Agent-based Simulation of Network System

Figure 5 illustrates our agent-based simulation of attack-defense experiments. We use a master-agent architecture, where two master agents are responsible for scheduling attacks and defenses, respectively. Each simulated host runs a local defense agent by default, which receives instruction from the defense master server. Once a computer is compromised, a local attack agent is instantiated on the compromised computer to receive instructions from the attack master server. When the employment of dynamic network diversity is conducted *manually*, the defense agent can be

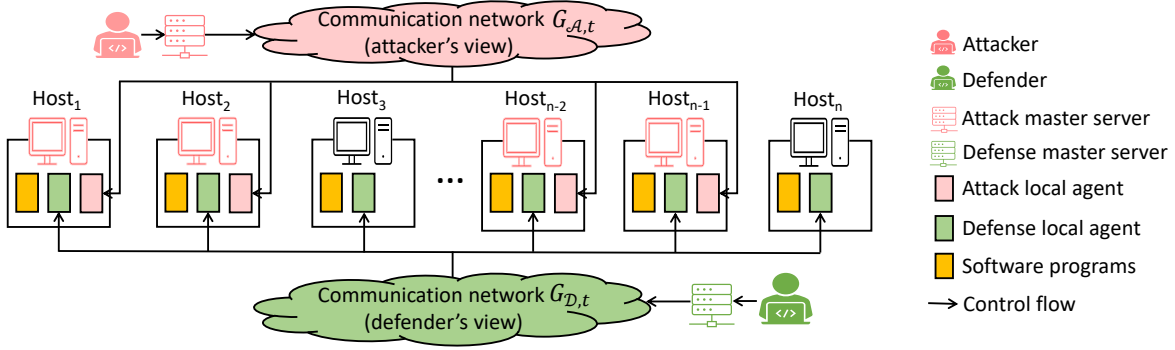


Fig. 5: Illustration of our agent-based simulation of attack-defense experiments.

compromised and re-installed from a clean version when the computer is compromised; when the re-employment process is *automatic*, the defense agent must not be compromised (even if the computer is compromised) so as to assure the re-employment of diversified program implementations. The master attack server updates its knowledge $\Sigma_{A,t}$ based on the information received from the attack agents and makes decisions correspondingly. The master defense server updates its knowledge $\Sigma_{D,t}$ in a similar fashion.

Simulating computers' software stack and communications. In order to make the communication graph $G = (V, E)$ as realistic as possible, we adopt two real-world social networks in Twitter and Friendfeed (which is a social media aggregator) [52], where the former has 5,702 users and the latter has 5,540 users. Together, there are 6,325 users because many users use both Twitter and Friendfeed. We construct the communication graph as follows: a Twitter user corresponds to a Twitter client program and a Friendfeed user corresponds to a Friendfeed client program. A user of Twitter and Friendfeed runs both client programs (i.e., the user's computer is represented as three nodes in the communication graph: one OS and two applications). An edge between two users in the social network means that they communicate with each other using the social network client program. Therefore, the communication graph accommodates the relationships in both social networks.

Simulating network diversity configuration. Since there are many ways to configure network diversity and the notion of an *optimal* algorithm is elusive (e.g., the algorithms for generating configuration C_t for $t > 0$ can be different from the one for generating configuration C_0), we will consider multiple algorithms and empirically contrast them. In contrast, monoculture software stack is trivial to configure because each program has exactly one implementation.

Simulating software vulnerabilities. In order to simulate vulnerabilities contained in the diversified program implementations, namely Φ_t , we assume that each diversified program has the same diversity quality, namely $Q_1 = Q_2 = \dots = Q_h = Q$, and each implementation of a program is equally vulnerable with a certain probability; this is a somewhat simplifying assumption but is reasonable in the sense that the same vulnerability prevention and detection techniques may be equally applicable to all implementations. Since there are studies showing that different implementations often do not have the same vulnerability [53], [54], we assume that the vulnerabilities are distinct in the

sense that each requiring a different exploit.

3.2 Simulating Attacker

Simulating attacker's knowledge. For describing attacker's knowledge at time $t = 0$, we assume that the attacker already compromised some vulnerable programs running at some nodes $v \in V$, namely the initial compromise denoted by IniComp . This is reasonable because initial compromise typically follows reconnaissance or is waged by an insider threat, which is orthogonal to the main purpose of the present study. As attack proceeds, the attacker can increase its knowledge by learning more information about the communication graph $G = (V, E)$, and the programs running at the other nodes $v \in V$ that may not be known to the attacker at time $t = 0$. We assume the attacker knows the vulnerabilities associated with the diversified programs.

Simulating attacker's goal. For describing attacker's goal $\Omega_{A,t}$, we assume the attacker wants to cumulatively compromise as many programs as possible till time $t = T$. That means $\omega_{v,A,T} = 1$ for $v \in V_{A,t}$, where $t \in [0, T]$.

Simulating attacker's strategy. The framework aims to accommodate many attack strategies. In order to make our simulation experiment concrete, we adopt MITRE's ATT&CK [50] because it is widely used. The attack *tactics* in ATT&CK can be naturally mapped to the attack phases in our framework. In terms of attacker's strategy Γ_A , we focus on the following phases: (i) *installation*, which corresponds to $\gamma_{A,1}$ in the framework and installing attack agents in compromised computers (this phase is called *command-and-control* in ATT&CK); (ii) *discovery*, which corresponds to $\gamma_{A,2}$ in the framework and allows the attacker to concurrently explore the other programs running at the nodes in the other computers of the network; (iii) *privilege escalation*, which corresponds to $\gamma_{A,3}$ in the framework and occurs when the attacker gains the root privilege in the compromised computer; (iv) *lateral movement*, which corresponds to $\gamma_{A,4}$ in the framework and allows the attacker or its malware to exploit the vulnerabilities in the a remote software stack; and (v) *causing damages*, which corresponds to $\gamma_{A,5}$ and allows the attacker to causes damages to a compromised computer (this phase accommodates the tactics of *collection*, *exfiltration*, and *impact* in ATT&CK).

Simulating attacker's capability. In order to describe attacker's capability Δ_A , we adopt ATT&CK's attack *procedures* (i.e., attacks) as exploits. We assume the attacker possesses the following exploits: (i) an exploit for achieving *re-*

remote access, denoted by $\psi_{A,1,1}$; (ii) an exploit for achieving remote system discovery or obtaining information about the software running in other computers, denoted by $\psi_{A,2,1}$; (iii) an exploit for discovering system information or getting detailed information about a compromised computer, denoted by $\psi_{A,2,2}$; (iv) a set of exploits for escalating privilege or compromising a vulnerable OS from a compromised application running on top of it, denoted by $\{\psi_{A,3,1}, \dots, \psi_{A,3,m_3}\}$; (iv) a set of exploits for remote exploitation capability or compromising a remote, vulnerable computer for lateral movement purposes, denoted by $\{\psi_{A,4,1}, \dots, \psi_{A,4,m_4}\}$; and (v) an exploit causes some damages (e.g., collecting sensitive data from the compromised software), denoted by $\psi_{A,5,1}$. For our purposes, it suffices to assume $cost_{A,z,q} = 1$ for $1 \leq z \leq \ell$, $1 \leq q \leq m_z$, meaning that each exploit incurs the same cost to the attacker (e.g., purchasing or developing

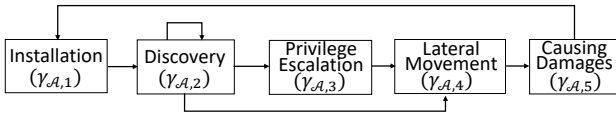


Fig. 6: Attacker \mathcal{A} 's 5-phase strategy (i.e., $\gamma_{A,1}, \dots, \gamma_{A,5}$) and possible decisions (i.e., the arrows), which are adopted from ATT&CK's attack simulator called CALDERA [55].

Simulating attacker's decision-making. In order to simulate the attacker's decision-making function \mathcal{F}_A , we adopt the decision-making component of ATT&CK that is used by ATT&CK's sub-system known as CALDERA [55]. Figure 6 shows how attacks proceed according to the 5-phase strategy mentioned above: install remote access tools at compromised computers; discover local and remote targets; compromise vulnerable OS via privilege escalation (if applicable); compromise remote computers for lateral movement (if applicable); conduct malicious activities; and repeat these processes. Basically, the decision-making algorithm outputs the next exploit that is to be executed, as follows: (i) if the next phase is $\gamma_{A,1}$ or $\gamma_{A,5}$, the attacker will use exploit $\psi_{A,1,1}$ or $\psi_{A,5,1}$ because it only possesses one exploit at each phase; (ii) if the next phase is $\gamma_{A,2}$, the attacker will use exploits $\psi_{A,2,1}$ and $\psi_{A,2,2}$ simultaneously, where $\psi_{A,2,1}$ targets one or more remote computers and $\psi_{A,2,2}$ targets the compromised local computer; (iii) if the next phase is $\gamma_{A,3}$, the attacker will select one exploit from the set $\{\psi_{A,3,1}, \dots, \psi_{A,3,m_3}\}$ according to the vulnerability information discovered in phase $\gamma_{A,2}$; (iv) if the next phase is $\gamma_{A,4}$, the attacker will select one exploit from the set $\{\psi_{A,4,1}, \dots, \psi_{A,4,m_4}\}$ according to the vulnerability information discovered in phase $\gamma_{A,2}$.

3.3 Simulating Defender

Simulating defender's knowledge. As described in the framework, the defender naturally knows, as a part of its knowledge $\Sigma_{D,t}$, the communication graph G (i.e., $G_D = G$) and the network configuration C_t (i.e., $C_{D,t} = C_t$) for any past and present time t . At time $t = 0$, we set that the defender's perception of the cybersecurity state S_0 as $S_{D,0} = (0, 0, \dots, 0)_{|V|}$ because the attack-detection tool (if

applicable) may start to run at $t = 0$. The defender does not know the information about the vulnerabilities because we allow zero-day vulnerabilities, meaning that $\hat{G}_{D,0} = \emptyset$. Putting these together, the defender's initial knowledge is $\Sigma_{D,0} = (G_D, C_{D,0}, S_{D,0}, \hat{G}_{D,0})$.

Simulating defender's goal. For describing a defender's goal $\Omega_{D,t}$, we assume the defender aims to keep the compromise rate at any time $t \leq T$ under a certain threshold τ_D . That is, we have $\omega_{i,D,t} = \tau_D$ for $i \in [1, n]$ and $t \in [0, T]$ on average, where the average is over the 500 simulation runs of each experiment.

Simulating defender's strategy. We consider the following 5 defense strategies Γ_D , (i) *monoculture* software stacks, denoted by $\gamma_{D,1}$, which corresponds to $\eta_{D,1} = \text{NULL}$, $\eta_{D,2} = \text{NULL}$, $\eta_{D,3} = \text{NULL}$. (ii) *Static diversity*, denoted by $\gamma_{D,2}$, which corresponds to $\eta_{D,1} = V$, $\eta_{D,2} = \text{NULL}$, $\eta_{D,3} = \text{NULL}$. (iii) *Proactive diversity*, denoted by $\gamma_{D,3}$, which corresponds to $\eta_{D,1} \neq \text{NULL}$, $\eta_{D,2} \neq \text{NULL}$, $\eta_{D,3} = \text{NULL}$. (iv) *Reactive-adaptive diversity*, denoted by $\gamma_{D,4}$, which corresponds to $\eta_{D,2} = \text{NULL}$, $\eta_{D,3} \neq \text{NULL}$, because dynamic network diversification is triggered by some security events, which may come from an intrusion detection system or the observed network-wide cybersecurity state $S_{D,t}$. Since such reactive intelligence is often noisy in practice, we incorporate false-negative rate (FNR) and false-positive rate (FPR) into such intelligence. When such intelligence is provided by an employed attack detection system, we assume that the attack detection system cannot be compromised in the present study. (v) *Hybrid diversity*, denoted by $\gamma_{D,5}$, which combines the aforementioned proactive diversity and reactive-adaptive diversity and corresponds to $\eta_{D,2} \neq \text{NULL}$, $\eta_{D,3} \neq \text{NULL}$. In this case, dynamic network diversity is triggered periodically or by security events.

Simulating defender's capability. For simplicity, we assume that defender's capability Δ_D includes: (i) the same number of diversified implementations for each program (for simplicity), namely $X_1 = X_2 = \dots = X_h = X$; and (ii) $cost_{D,k,w} = 1$ for $1 \leq k \leq h$ and $1 \leq w \leq X_k$, meaning that each diversified implementation incurs the same cost to the defender. These simplifying assumptions, while arguably reasonable when diversification is an automated process, need to be extended to consider more general cases. It is worth mentioning that different versions of a program should not be counted as diversified implementations because they would have many vulnerabilities in common.

Simulating defender's decision-making. For describing defender's decision-making function $\mathcal{F}_{D,t}$, we first consider $\mathcal{F}_{D,0}$ at time $t = 0$, which outputs the initial network diversity configuration C_0 . We consider three decision-making algorithms for $\mathcal{F}_{D,0}$ at time $t = 0$: (i) the baseline *random coloring* algorithm, which assigns a random implementation of the program in question to run at node v ; (ii) the *color flipping* algorithm [32], which uses the random coloring algorithm mentioned above as a starting point and then iteratively let nodes change their colors to reduce the number of defective edges (i.e., the edges with two end nodes having the same color or running the same implementation of a software); (iii) a new algorithm we propose, which leverages the degrees of the nodes to assign software implementations to nodes $v \in V$, by giving the large-degree nodes a high

priority in running diversified programs. Our algorithm is different from the *color flipping* algorithm mentioned above because we prioritize large-degree nodes in the initial assignment. Due to space limit, we defer the pseudo-code of the algorithm to the Supplementary Material.

For describing defender's decision-making function $\mathcal{F}_{\mathcal{D},t}$ for $t > 0$, we consider $\mathcal{F}_{\mathcal{D},1} = \dots = \mathcal{F}_{\mathcal{D},T} = \text{NULL}$ for defense strategy $\gamma_{\mathcal{D},2}$ because the software deployment stays unchanged over time in the case of static network diversity. For dynamic diversity $\gamma_{\mathcal{D},q}$ where $3 \leq q \leq 5$, we consider a simple random decision-making function $\mathcal{F}_{\mathcal{D},t}$ for $t \in [1, T]$, namely that the defender randomly selects diversified implementations to replace the currently-employed implementations at some or all of the nodes $v \in V$ according to defender's strategy. More sophisticated decision-making functions are left for future studies.

3.4 Simulating Local Effect and Quantifying Global Effect of Dynamic Diversity

We simulate the local effect of dynamic network diversity at each node $v \in V$ as described in Figure 3 and the global effect according to the COODAL based attack-defense interactions described in Figure 2. We collect the network-wide cybersecurity state and situation over time t to quantitatively answer the RQs (Section 2.7).

3.5 Answering RQs

Our simulation study is centered at measuring the metrics to quantify the security effectiveness of employing network diversity, and then leveraging these quantitative effectiveness to draw insights. While the model is general, the simulation study can only consider some specific parameter settings because it is not feasible to consider all parameter settings. That is, the simulation study only corresponds to some scenarios of the general model, and the findings drawn from the simulation study may not be generalized to other parameter settings. Researchers and practitioners can apply our model to their specific parameter settings.

3.5.1 RQ1: To what extent can dynamic network diversity slow down the attacker?

In order to answer RQ1, we investigate how the attacker slower-down metric $\text{ASD}_{\mathcal{D},q}$ depends on defender's goal τ with $2 \leq q \leq 5$. The experimental parameters are: $h = 3$ (3 different programs running in the network: Twitter, Friendfeed, OS), $X = 10$ (each program has 10 diversified implementations), $Q = 1$ (every diversified implementation is vulnerable), $|\text{IniComp}| = 10$ (10 programs/nodes are initially compromised), $m_3 = 5$ (the attacker has 5 exploits against OS), $m_4 = 10$ (the attacker has 5 exploits against Twitter and 5 exploits against Friendfeed), and $T = 500$ (the simulation stops at $t = 500$).

Figure 7(a) plots attacker slow-down $\text{ASD}_{\mathcal{D},2}$ with respect to defender's goal $\tau \in [0, 0.5]$, where the defender uses static diversity and different decision-making algorithm $\mathcal{F}_{\mathcal{D},0}$ at time $t = 0$. We observe: (i) $\text{ASD}_{\mathcal{D},2} = 0$ when $\tau \leq 0.1$, meaning that static diversity cannot slow down the attacker when the defender can only tolerate no more than 10% of the nodes being compromised. (ii) $\text{ASD}_{\mathcal{D},2} > 0$ when $\tau \geq 0.15$, meaning that static diversity can slow down the attacker at

an extent that increases with the degree of tolerable compromise. This is reasonable because the attacker has to do more lateral movements in order to disrupt the defender's goal. (iii) For a fixed tolerable compromise threshold τ , the (initial diversity) decision-making algorithm $\mathcal{F}_{\mathcal{D},0}$ matters and our algorithm slows down the attacker most, with an average slow-down that is almost 2X of that of the random coloring algorithm, where the average is over the τ 's.

Figure 7(b) plots attacker slow-down $\text{ASD}_{\mathcal{D},3}$ with respect to defender's goal $\tau \in [0, 0.5]$, where the defender uses different decision-making algorithms $\mathcal{F}_{\mathcal{D},0}$ for initial diversity and proactively uses $\mathcal{F}_{\mathcal{D},t}$ for $t > 0$ with parameters $\eta_{\mathcal{D},1} = 0.5$ and $\eta_{\mathcal{D},2} = 0.2$ (diversified implementations are re-deployed at 50% of the nodes every 5 time steps). We make the same observations as that of $\text{ASD}_{\mathcal{D},2}$. This is reasonable because proactive diversity can replace compromised programs with secure programs (i.e., benefiting the defender), but can also replace secure programs with vulnerable ones (i.e., benefiting the attacker).

Figure 7(c) plots attacker slow-down $\text{ASD}_{\mathcal{D},4}$ with respect to defender's goal $\tau \in [0, 0.5]$, where the defender uses some decision-making algorithm $\mathcal{F}_{\mathcal{D},0}$ for initial diversity and reactive-adaptively uses $\mathcal{F}_{\mathcal{D},t}$ for $t > 0$ while assuming parameters $\text{FPR} = 0.1$ and $\text{FNR} = 0.1$ (i.e., the attack-detection or threat intelligence has a 10% false-positive rate and a 10% false-negative rate). We make the following observations: (i) $\text{ASD}_{\mathcal{D},4} > 0$ for $\tau \in (0, 0.5]$, meaning that reactive-adaptive diversity can always slow down the attacker at an extent that concavely increases with the tolerable compromise threshold τ . This is reasonable because the defender can replace a likely-compromised software with another diversified implementation to benefit the defender. (ii) For a fixed tolerable compromise threshold τ , the initial diversity algorithm $\mathcal{F}_{\mathcal{D},0}$ matters because our algorithm slows down the attacker most, with an average of almost 2X slow-down than that of the the random coloring algorithm.

Figure 7(d) plots attacker slow-down $\text{ASD}_{\mathcal{D},5}$ with respect to defender's goal $\tau \in [0, 0.5]$, where the defender uses some decision-making algorithm $\mathcal{F}_{\mathcal{D},0}$ for initial diversity and hybrid (of proactive and reactive-adaptive) decision-making algorithm $\mathcal{F}_{\mathcal{D},t}$ for $t > 0$ with parameters $\eta_{\mathcal{D},2} = 0.2$, $\text{FPR} = 0.1$ and $\text{FNR} = 0.1$. We observe the same phenomena as in the case of reactive-adaptive diversity, except that the degree of slow-down incurred by hybrid diversity increases with the tolerable compromise threshold in a *convex* (rather than *concave*) fashion.

By comparing Figures 7(a)-7(d), we observe that for a fixed (initial diversity) decision-making algorithm $\mathcal{F}_{\mathcal{D},0}$, we have $\text{ASD}_{\mathcal{D},4} > \text{ASD}_{\mathcal{D},5} \gg \text{ASD}_{\mathcal{D},3} \approx \text{ASD}_{\mathcal{D},2}$ for $\tau \in (0, 0.45]$, indicating that reactive-adaptive diversity outperforms hybrid diversity, which significantly outperforms proactive diversity and static diversity. Consider $\tau = 1/3$ as an example, we observe $\text{ASD}_{\mathcal{D},4} = 412$, $\text{ASD}_{\mathcal{D},5} = 269$, $\text{ASD}_{\mathcal{D},3} = 10$, and $\text{ASD}_{\mathcal{D},2} = 10$ when using our algorithm as (initial diversity) decision-making algorithm $\mathcal{F}_{\mathcal{D},0}$. Note that the curves in Fig. 7(a) are not smooth because, under the static diversity strategy, compromised programs are not cleaned up and can attack others. This also explains why the curves in Fig. 7(b) are not smooth, namely that proactive diversity always periodically selects random programs for re-employing dynamic diversity, causing some compromised

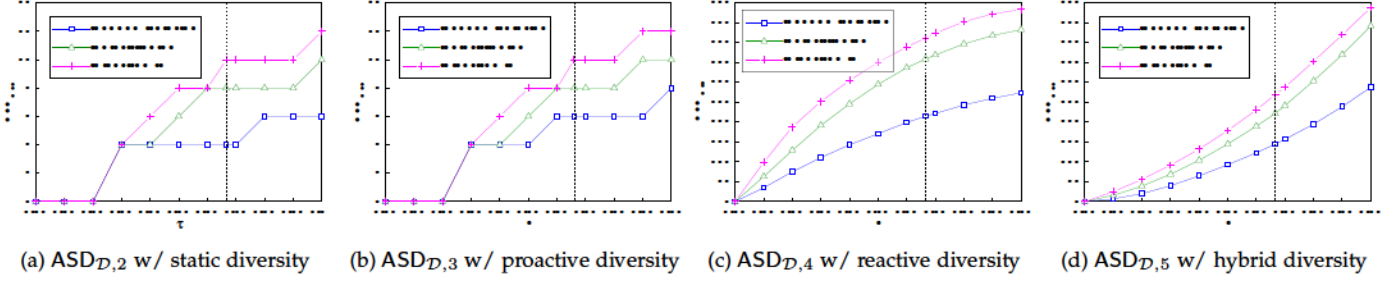


Fig. 7: Plots of attacker slow-down $ASD_{D,q}$ under different diversity strategies (where reactive is short for reactive-adaptive) with varying tolerable compromise threshold τ , where $2 \leq q \leq 5$ and dotted vertical lines indicate $\tau = 1/3$.

computers to remain compromised for extended durations and allowing them to conduct further attacks. In contrast, the curves in Fig. 7(c) and Fig. 7(d) are smooth because these two strategies leverage reactive defense systems to identify and replace the compromised programs timely, which can limit the abrupt spreading of attacks and slow down the attacker. One reason for the reactive-adaptive strategy to perform better than the hybrid strategy is that the former can immediately clean up the compromised programs, but the latter waits until a period of time, which allows the attacker to compromise other vulnerable programs. That is, the difference is caused by whether there is a gap between when compromised programs are detected and when compromised programs are cleaned up.

Insight 1. In terms of the attacker slow-down metric, reactive-adaptive diversity is the most effective strategy and the initial diversity configuration matters.

3.5.2 RQ2: How much extra cost can dynamic network diversity impose on the attacker?

In order to answer RQ2, we investigate how the *attack extra cost* metric $AEC_{D,q}$ increases with defender's goal τ , where $2 \leq q \leq 5$. For this purpose, we need to see how attack cost affects the network-wide cybersecurity state, especially attack worst damage $AWD_{G,h,X,Q}(\mathcal{A}, \mathcal{D}_q, T)$, where $2 \leq q \leq 5$. The simulation experiment parameters are: $h = 3$ (i.e., 3 different programs running in the network: Twitter, Friendfeed, and OS), $X = 10$ (i.e., each program has 10 diversified implementations), $Q = 1$ (i.e., every diversified implementation is vulnerable), $|IniComp| = 10$ (i.e., 10 programs or nodes are initially compromised), $\mathcal{F}_{D,0}$ is our algorithm for employing initial diversity, $\eta_{D,1} = 0.5$ (i.e., diversified implementations are dynamically re-employed at 50% of all nodes), $\eta_{D,2} = 0.2$ (i.e., diversified implementations are re-employed every 5 time steps), $FPR = 0.1$ (i.e., 10% false-positive rate in detecting attacks), $FNR = 0.1$ (i.e., 10% false-negative rate in attack detection), $T = 500$. We assume that the attacker uses the available exploits together, for the sake of reducing the uncertainty in the outcomes that may be incurred by the orders of exploits usage.

Figure 8(a) plots $AWD_{G,h,X,Q}(\mathcal{A}, \mathcal{D}_q, 500)$ for $2 \leq q \leq 5$, with respect to the total number $m_3 + m_4$ of exploits possessed by the attacker, where m_3 is the number of exploits that provide the attacker with privilege escalation capability and m_4 is the number of exploits that provide the attacker with lateral movement capability. We observe the following *phase-transition* phenomenon. When $0 < m_3 + m_4 \leq 12$ (i.e.,

the attacker possessing no more than 40% of the total 30 exploits, which correspond to all of the 30 vulnerabilities in the diversified implementations), the reactive-adaptive diversity strategy $\gamma_{D,4}$ leads to the lowest attack worst damage (i.e., the $\gamma_{D,4}$ curve); when $12 < m_3 + m_4 \leq 21$, the hybrid diversity strategy $\gamma_{D,5}$ leads to the lowest attack worst damage (i.e., the $\gamma_{D,5}$ curve); when $21 < m_3 + m_4 \leq 30$, the proactive diversity strategy $\gamma_{D,3}$ leads to the lowest attack worst damage (i.e., the $\gamma_{D,3}$ curve). This phenomenon can be explained as follows. In terms of the *attack worst damage* metric, reactive-adaptive diversity is the most effective strategy against a less capable attacker because the defender can detect and replace the small number of compromised programs; proactive diversity is the most effective strategy against a more capable attacker, which compromises a large number of computers and demands periodic enforcement of dynamic diversity at most, if not all, of the computers.

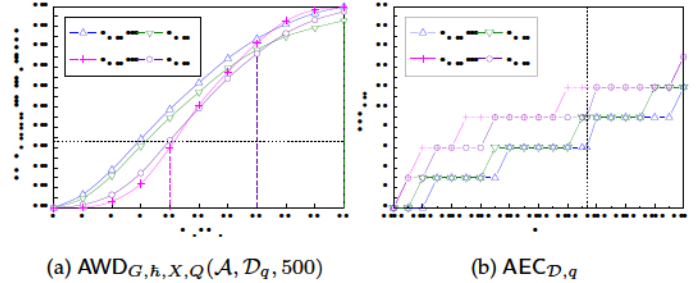


Fig. 8: Plots of $AWD_{G,h,X,Q}(\mathcal{A}, \mathcal{D}_q, 500)$ and $AEC_{D,q}$ with $2 \leq q \leq 5$.

From Figure 8(a), we can derive the attack extra cost $AEC_{D,q}$ with respect to the defender's goal $\tau \in [0, 0.5]$ (i.e., no more than a τ fraction of the nodes are compromised at any time $t \in [0, T]$) where $2 \leq q \leq 5$, which is plotted in Figure 8(b). We make the following observations: (i) $AEC_{D,4} > 0$ and $AEC_{D,5} > 0$ when $\tau \geq 0.025$, $AEC_{D,3} > 0$ when $\tau \geq 0.05$, and $AEC_{D,2} > 0$ when $\tau \geq 0.075$, meaning that employing dynamic network diversity defense strategy can impose attack extra cost on the attacker at an extent that increases with the degree of tolerable compromises. (ii) $AEC_{D,4} \geq AEC_{D,5} \geq AEC_{D,3} \geq AEC_{D,2}$ always holds for any $\tau \in (0, 0.5]$. Consider $\tau = 1/3$ as an example, we observe $AEC_{D,4} = 40\% \geq AEC_{D,5} = 30\% \geq AEC_{D,3} = 30\% \geq AEC_{D,2} = 20\%$, meaning that reactive-adaptive diversity imposes extra cost on the attacker more than hybrid diversity, which incurs more than proactive diversity and even more than static diversity.

Insight 2. In order to reduce the attack worst damage, different diversity strategies should be used in different parameter regimes.

3.5.3 RQ3: To what extent can dynamic network diversity increase the defender's vulnerability tolerance?

In order to answer RQ3, we investigate how the *vulnerability tolerance* metric $VT_{D,q}$ depends on defender's goal τ , where $2 \leq q \leq 5$. Since this dependence would rely on attack worst damage $AWD_{G,h,X,Q}(\mathcal{A}, \mathcal{D}_q, T)$ and the parameters G and h are largely determined by the applications and X is largely determined by what is available, we will investigate the impact of diversification quality Q . The simulation experiment parameters are: $h = 3$ (3 different programs running in the network: Twitter, Friendfeed, and OS), $X = 20$ (each program has 20 diversified implementations), $|IniComp| = 10$ (10 programs or nodes are initially compromised), $\mathcal{F}_{D,0}$ is our algorithm for employing initial diversity, $\bar{m}_3 = 0.5 \times X \times Q$ (the attacker has 50% of the total $X \times Q$ exploits against the OSes on average), $\bar{m}_4 = 0.5 \times 2X \times Q$ (the attacker has 50% of the total $X \times Q$ exploits against Twitter and 50% of the total $X \times Q$ exploits against Friendfeed on average), $\eta_{D,1} = 0.5$ (diversified implementations are dynamically re-employed at 50% of all nodes), $\eta_{D,2} = 0.2$ (diversified implementations are re-employed every 5 time steps), $FPR = 0.1$ (10% false-positive rate in detecting attacks), $FNR = 0.1$ (10% false-negative rate in attack detection), and mission lifetime $T = 500$.

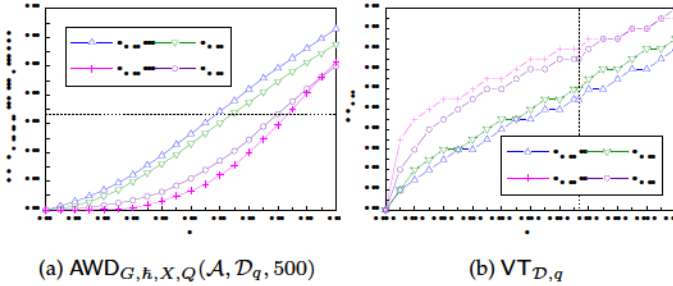


Fig. 9: Plots of $AWD_{G,h,X,Q}(\mathcal{A}, \mathcal{D}_q, 500)$ and $VT_{D,q}$ with $2 \leq q \leq 5$.

Figure 9(a) plots $AWD_{G,h,X,Q}(\mathcal{A}, \mathcal{D}_q, 500)$ with respect to diversity quality Q when the defender employs network diversity defense strategy $\gamma_{D,q}$, where $2 \leq q \leq 5$. We observe that $cc_{D,4} < cc_{D,5} < cc_{D,3} < cc_{D,2}$ when $0 < Q \leq 0.95$, and $cc_{D,5} < cc_{D,4} < cc_{D,3} < cc_{D,2}$ when $0.95 \leq Q \leq 1$. This means that employing dynamic network diversity always leads to higher security than static diversity regardless of the diversification quality Q . We observe that the attack worst damage $AWD_{G,h,X,Q}(\mathcal{A}, \mathcal{D}_q, 500)$ increases when the quality of diversified implementations drops, where $2 \leq q \leq 5$. This means that dynamic network diversity leads to an even higher security when the diversity quality is high (i.e., low Q 's).

From Figure 9(a), we can derive the vulnerability tolerance $VT_{D,q}$ with respect to the defender's goal $\tau \in [0, 0.5]$ where $2 \leq q \leq 5$, which is plotted in Figure 9(b). We observe that $VT_{D,4} \geq VT_{D,5} > VT_{D,3} \geq VT_{D,2}$ when $0 < \tau \leq 0.475$. Consider $\tau = 1/3$ as an example, we observe $VT_{D,4} = 0.8$, $VT_{D,5} = 0.75$, $VT_{D,3} = 0.6$, and $VT_{D,2} = 0.55$.

This means that given an attacker that can exploit 50% of the vulnerabilities in the network on average, reactive-adaptive diversity strategy $\gamma_{D,4}$ makes the defender tolerate a $0.8 - 2/3 = 0.14$ or 14% extra vulnerabilities, and hybrid diversity strategy $\gamma_{D,5}$ make the defender tolerate a $0.75 - 2/3 = 0.09$ or 9% extra vulnerabilities; however, proactive and static diversity strategies cannot achieve this effectiveness.

Insight 3. Reactive-adaptive diversity leads to a higher vulnerability-tolerance than proactive diversity does.

3.5.4 RQ4: To what extent can dynamic network diversity increase the defender's average operational cost?

In order to answer RQ4, we investigate how the *average operational cost* metric $AOC_{D,q}$ increases with defender's goal τ , where $2 \leq q \leq 5$. For this purpose, we need to know how attack worst damage $AWD_{G,h,X,Q}(\mathcal{A}, \mathcal{D}_q, T)$ depends on defense diversity strategies $\gamma_{D,q}$, where $2 \leq q \leq 5$. The simulation experiment parameters are: $h = 3$ (3 different programs running in the network: Twitter, Friendfeed, OS), $X = 10$ (each program has 10 diversified implementations), $Q = 1$ (every diversified implementation is vulnerable), $|IniComp| = 10$ (10 programs/nodes are initially compromised), $\mathcal{F}_{D,0}$ is our algorithm for employing initial diversity, $m_3 = 5$ (the attacker has 5 exploits against OS), $m_4 = 10$ (the attacker has 5 exploits against Twitter and 5 exploits against Friendfeed), and mission lifetime $T = 500$.

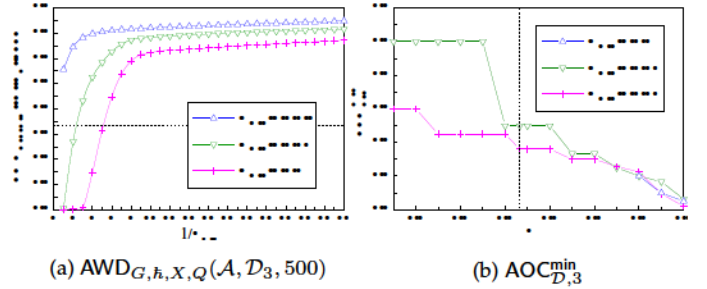


Fig. 10: Plots of $AWD_{G,h,X,Q}(\mathcal{A}, \mathcal{D}_3, 500)$ and $AOC_{D,3}^{min}$.

Figure 10(a) plots $AWD_{G,h,X,Q}(\mathcal{A}, \mathcal{D}_3, 500)$ with respect to $\eta_{D,1}$ and $1/\eta_{D,2}$, where $\eta_{D,1}$ is the the proportion of nodes where diversified programs are dynamically employed and $1/\eta_{D,2}$ is the time interval between two consecutive diversity employments. We observe the following: (i) When $\eta_{D,1} = 0.1$, proactive diversity always leads to high attack worst damage, even if $1/\eta_{D,2}$ is small (i.e., high-frequency in employment). This means that proactive diversity is useless when dynamic diversity is employed at a few nodes. (ii) When $\eta_{D,1} = 0.5$ or 0.9 , proactive diversity can lead to low attack worst damage only when $1/\eta_{D,2}$ is small. Consider defense goal $\tau = 1/3$ as an example, $1/\eta_{D,2}$ must be no more than 2 when $\eta_{D,1} = 0.5$ and no more than 5 when $\eta_{D,1} = 0.9$. This means that proactive diversity is effective only when employed broadly and frequently. (iii) When $\eta_{D,1} = 0.1$ and $1/\eta_{D,2} \geq 2$, or when $\eta_{D,1} = 0.5$ and $1/\eta_{D,2} \geq 6$, or when $\eta_{D,1} = 0.9$ and $1/\eta_{D,2} \geq 14$, proactive diversity leads to higher attack worst damage than static diversity. This means that proactive diversity can do more harm than good by making making more nodes exploitable over time.

Insight 4. Proactive diversity improves security only when employed at most nodes at high frequency.

From Figure 10(a), we can derive the minimum average operational cost $AOC_{D,3}^{\min}$ with respect to the defender's goal $\tau \in [0, 0.7]$, which is shown in Figure 10(b). Consider $\tau = 1/3$ as an example, the minimum average operational cost is $\eta_{D,1} \times \eta_{D,2} = 0.5 \times 1/2 = 0.25$ for $\eta_{D,1} = 0.5$, and $\eta_{D,1} \times \eta_{D,2} = 0.9/5 = 0.18$ for $\eta_{D,1} = 0.9$. Note that in the case $\eta_{D,1} = 0.1$, $AOC_{D,3}^{\min}$ is undefined when $\tau \in [0, 0.55]$ because the strategy can never prevent the attacker from breaking the defender's goal. We further observe that $\eta_{D,1} = 0.9$ leads to a lower $AOC_{D,3}^{\min}$ than $\eta_{D,1} = 0.5$ when $0 \leq \tau \leq 0.5$, meaning that a higher proportion of dynamic diversity re-employment leads to a lower average operation cost.

Insight 5. When proactive diversity is effective, a higher proportion of dynamic re-employment leads to a lower operational cost than what is incurred by a higher re-employment frequency.

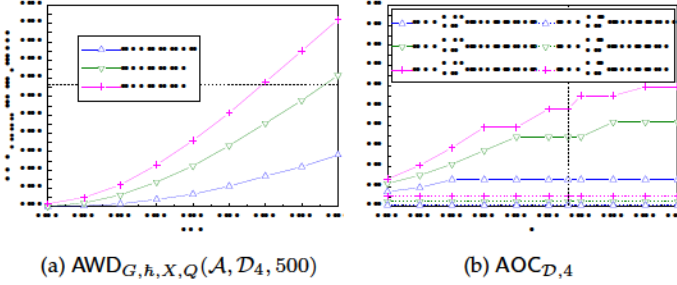


Fig. 11: Plots of $AWD_{G,h,X,Q}(\mathcal{A}, \mathcal{D}_4, 500)$ and $AOC_{D,4}$.

Figure 11(a) plots $AWD_{G,h,X,Q}(\mathcal{A}, \mathcal{D}_4, 500)$ with respect to FPR (false-positive rate) and FNR (false-negative rate), where $FPR, FNR \in [0.02, 0.1]$. We observe that a lower FNR and FPR (i.e., higher attack-detection capability) leads to a lower attack worst damage, meaning the effectiveness of reactive-adaptive diversity largely depends on the attack-detection capability. Consider defender's mission goal of $\tau = 1/3$ as an example, we observe that reactive-adaptive diversity can assure the mission when $FPR = 0.02$ and $0.02 \leq FNR \leq 0.1$, when $FPR = 0.06$ and $0.02 \leq FNR \leq 0.09$, and when $FPR = 0.1$ and $0.02 \leq FNR \leq 0.07$. From Figure 11(a), we can derive the $AOC_{D,4}^{\min}$ and $AOC_{D,4}^{\max}$ for a given FPR with respect to $\tau \in [0, 0.5]$, which is plotted in Figure 11(b). We make the following observations. (i) For a fixed FPR, $AOC_{D,4}^{\min}$ can be very low and remains stable as τ increases, because $AOC_{D,4}^{\min}$ is always achieved when $FNR = 0.02$. The operational cost is low because a high attack-detection accuracy can detect compromised computers before they attack the others. (ii) For a fixed FPR, a higher τ often leads to a higher $AOC_{D,4}^{\max}$, because a higher τ (i.e., higher compromise-tolerance) can be achieved at a lower operational cost from a diversity-based defense standpoint. (iii) The defender's operational cost falls into a wide range as τ increases, meaning that the defender's operational cost largely depends on the attack-detection effectiveness.

Figure 12(a) plots $AWD_{G,h,X,Q}(\mathcal{A}, \mathcal{D}_5, 500)$ with respect to FPR and FNR, where $FPR, FNR \in [0.02, 0.1]$ and re-deployment frequency $\eta_{D,2} = 0.2$. From Figure 12(a), we derive $AOC_{D,5}^{\max}$ and $AOC_{D,5}^{\min}$, which is plotted in Figure

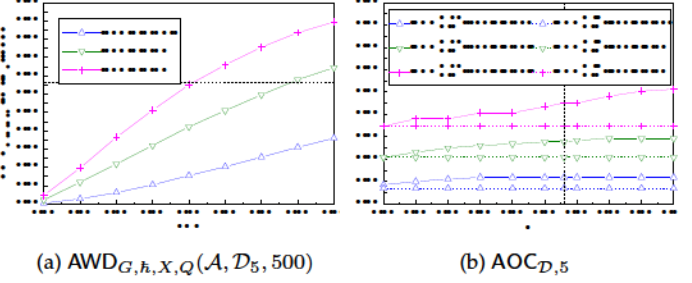


Fig. 12: Plots of $AWD_{G,h,X,Q}(\mathcal{A}, \mathcal{D}_5, 500)$ and $AOC_{D,5}$.

12(b). We observe that the defender's average operational cost is small. Consider $\tau = 1/3$ as an example, we observe $AOC_{D,5}^{\min} = 0.0349$ when $FPR = 0.02$ and $AOC_{D,5}^{\max} = 0.2254$ when $FPR = 0.1$, meaning $AOC_{D,5} \in [0.0349, 0.2254]$ when $FPR \in [0.02, 0.1]$ and $FNR \in [0.02, 0.1]$.

By comparing Figures 10(b), 11(b), and 12(b), we observe that for $\tau \in [0.05, 0.5]$, we have $AOC_{D,3} \in [0.15, 0.5]$ when $\eta_1 \in [0.5, 0.9]$; we have $AOC_{D,4} \in [0.0071, 0.5948]$ and $AOC_{D,5} \in [0.0349, 0.2566]$ when $FPR \in [0.02, 0.1]$ and $FNR \in [0.02, 0.1]$, meaning that reactive-adaptive diversity incurs an average operational cost falling into a wider range than proactive diversity and an even wider range than hybrid diversity. We also observe that proactive diversity incurs a higher operational cost than reactive-adaptive diversity when the defender's tolerable compromise threshold τ is small, and the opposite is true when τ is large.

3.5.5 RQ5: Is it true that the more diversified implementations the better?

In order to answer RQ5, we investigate how the *attacker slow-down* metric $ASD_{D,q}$, the *attack extra cost* metric $AEC_{D,q}$, and the *vulnerability tolerance* metric $VT_{D,q}$ depend on the number of diversified implementations X , where $2 \leq q \leq 5$. The experimental parameters are: $h = 3$ (3 different programs running in the network: Twitter, Friendfeed, and OS), $Q = 1$ (every diversified implementation is vulnerable), $|IniComp| = 10$ (10 programs or nodes are initially compromised), $m_3 = 1$ (the attacker has 1 exploit against OS), $m_4 = 2$ (the attacker has 1 exploit against Twitter and 1 exploit against Friendfeed), $\mathcal{F}_{D,0}$ is our algorithm for employing initial diversity, $\eta_{D,1} = 0.5$ (diversified implementations are dynamically re-employed at 50% of all nodes), $\eta_{D,2} = 0.2$ (diversified implementations are re-employed every 5 time steps), $FPR = 0.1$ (10% false-positive rate in detecting attacks), $FNR = 0.1$ (10% false-negative rate in attack detection), and mission lifetime $T = 500$.

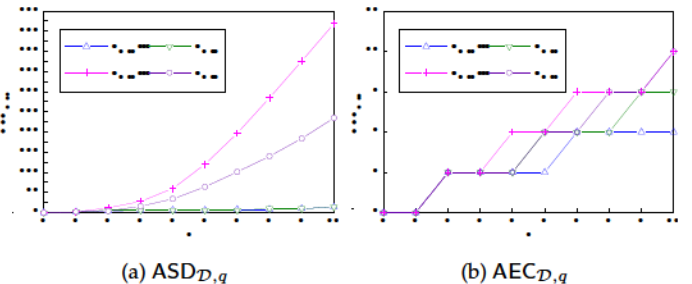


Fig. 13: Plots of $ASD_{D,q}$ and $AEC_{D,q}$, where $2 \leq q \leq 5$.

Figure 13(a) plots $ASD_{D,q}$ with respect to X and $\tau = 0.01$, where $2 \leq q \leq 5$. We take $\tau = 0.01$ as an example because the attacker possessing one exploit for each program may only compromise about 1% computers during the mission lifetime $T = 500$ with $X = 10$. We observe that $ASD_{D,4}$ increases rapidly with X , followed by $ASD_{D,5}$, while $ASD_{D,3}$ and $ASD_{D,2}$ increases slowly with X (e.g., $ASD_{D,3} = ASD_{D,2} = 2$ when $X = 2$ and $ASD_{D,3} = ASD_{D,2} = 14$ when $X = 10$). When $X = 10$, reactive-adaptive diversity slows down the attacker most.

Figure 13(b) plots $AEC_{D,q}$ with respect to X and $\tau = 1/3$, where $2 \leq q \leq 5$. We observe that $AEC_{D,q}$ shows an upward trend as X increases for $q = 2, 3, 4, 5$, meaning that the more diversified implementations, the higher the attack extra cost for disrupting the defender's mission goal (i.e., no more than a τ fraction of computers are compromised at any time $t \in [0, T]$). In addition, we observe $AEC_{D,4} \geq AEC_{D,5} \geq AEC_{D,3} \geq AEC_{D,2}$ for any $X \in [1, 10]$, meaning that reactive-adaptive diversity benefits most from more diversified implementations, followed by hybrid diversity, proactive diversity and static diversity.

In order to characterize the impact of the number of diversified implementations on the vulnerability tolerance $VT_{D,q}$ where $2 \leq q \leq 5$, we observe $VT_{D,q} = 0$ when $X = 1, 2$ and $\tau = 1/3$. This is reasonable because the attacker having one exploit for each program can easily break the defender's goal when the total number of diversified implementations is no more than 2. In contrast, given $\tau = 1/3$, we have $VT_{D,q} = 1$ when $X \geq 3$, where $2 \leq q \leq 5$, because Figure 13(b) shows that an attacker having one exploit for each program will never compromise $1/3$ computers when $X \geq 3$; this is true even if all of the diversified implementations are vulnerable. This means that when the attacker has a limited capability, a substantially higher vulnerability rate (than τ) can effectively be tolerated until the number of diversified implementations reaches a certain amount.

Insight 6. The more diversified implementations with a similar quality, the higher the attacker slow-down, the attack extra cost, and the vulnerability tolerance.

4 RELATED WORK

Prior studies in software diversity. Prior studies mainly aim at obtaining diversified, ideally *independent*, implementations of a program specification [4]. There is a body of literature on software diversity (e.g., [6], [10], [11], [12], [13], [14], [15], [17], [18], [19]). However, the effectiveness of these building-block techniques is not well understood. For example, one study shows that independent software implementation does not lead to independent vulnerabilities because programmers tend to make the same mistakes [56]. Other studies show positive results, such as: the same vulnerabilities in different software may demand different exploits [53]; few vulnerabilities simultaneously appear in different OSes [57]. In contrast, the effectiveness of run-time diversity is better understood: (i) address space layout with 32-bit randomization can be compromised by brute-forcing [58]; (ii) address space layout with higher entropy can be compromised by side-channel attacks [59], [60], [61], [62]; (iii) address space layout randomization is vulnerable because of modern cache architectures [63]; (iv) fine-grained

address space layout randomization is subject to just-in-time code reuse attacks [64]; and (v) instruction set randomization is subject to brute-forcing attacks [65], [66].

The preceding studies consider standalone software diversification techniques. In contrast, we investigate the effectiveness of software diversification techniques from a network standpoint. Since we consider diversifying network-wide software stacks, multiple diversification techniques can be used together. This is reminiscent of the notion of N -variant systems [67], which however do not quantify the network-wide effectiveness dynamic diversity.

Prior studies in network diversity. Network diversity has been investigated in some contexts [1], [2], [3], [32], [33]. There are proposals on measuring static network diversity via (i) the entropy of the distribution of software vulnerability in a network [38] and (ii) the diversity index of the shared vulnerabilities between different software implementations in a network [41]. There are attempts at optimizing static diversity by using some flavor of *graph coloring* algorithm (i.e., treating a diversified implementation as a color and a network as a graph) and minimizing defective edges (i.e., adjacent nodes have the same color or run the same software implementation) [32], [33], [34], [35], [36], [37].

The closely related prior studies are [39], [40], which measure network diversity via *resource richness* and *attack effort*. Our study is different as follows: (i) they investigate how to quantify diversity, whereas we study how to employ *dynamic* diversity; (ii) they do not consider attack-defense interactions, whereas we explicitly model attack-defense interactions; (iii) we quantify the network-wide effectiveness of dynamic diversity, which is not studied by them.

Prior studies related to Moving-Target Defense (MTD). Proactive diversity is one form of MTD, which includes other kinds of proactive defense techniques (e.g., proactively changing IP addresses or port numbers) [68]. Quantifying the security effectiveness of MTD in the broader context is beyond the scope of the present paper (see, e.g., [69]). While our finding that reactive-adaptive is more effective than proactive diversity (i.e., MTD when applied to diversity) may sound counter-intuitive at a first glance, it can be understood as follows: MTD (i.e., proactive diversity in this case) can be employed when the defender does not know the situation information of the network (e.g., which and how many nodes are compromised); in contrast, adaptive diversity can leverage the situational information to adaptively employ diversity, leading to potentially higher effectiveness. In addition, our simulation study focuses on dynamic diversity against malware-like attacks after the attacker establishes footholds at some compromised computers in a network. This means that the simulation study does not consider earlier stages of cyber attacks, such as reconnaissance. Our findings do not contradict the usefulness of MTD in defending against such earlier-stage attack activities (e.g., MTD can effectively disrupt attacker's reconnaissance processes [70], [71], [72], [73]).

Prior studies related to whole-network security analysis. We analyze the security effectiveness of dynamic diversity from a whole-network perspective. There are studies in this perspective, but tackling different problems and using different approaches. The *attack graph* approach studies how

an attacker may exploit multiple vulnerabilities to achieve a certain goal and how to harden a network (see, e.g., [74], [75], [76], [77], [78], [79], [80], [81], [82]). This approach is *combinatorial* in nature and does not consider the *time* dimension [83], [84], [85]. Another approach is the *cybersecurity dynamics* framework [83], [84], [85], which explicitly models attack-defense interactions over time and includes a rich family of models and results (e.g., [69], [86], [87], [88], [89], [90], [91], [92], [93], [94], [95], [96]). These studies aim at analytical results while making simplifying assumptions, such as the *independence* assumption between attacks. In order to eliminate such assumptions, initial efforts have been made in both theoretical studies [89], [97], [98] and empirical studies [43], [99]. Our framework does not make the independence assumption, while characterizing the *transient* behaviors (i.e., the dynamics before converging to an equilibrium); whereas, analytical models so far can only offer asymptotic results (i.e., $t \rightarrow \infty$ or when the dynamics converges to the equilibrium). Last but not the least, our framework goes much beyond [99] by characterizing (e.g.) dynamic attack-defense interactions and decision-making.

5 LIMITATIONS

The present study has several limitations that need to be addressed in future research. First, the framework has two limitations. (i) We assume G is fixed. This implicitly assumes the network defense tools are not compromised because a successful attack against a defense tool can effectively change G . While this is reasonable for missions with a short lifetime T , it is interesting to accommodate dynamic G_t for missions of long lifetime T and the case that the network defense tools can be compromised, as outlined in [83], [84], [85]. (ii) We assume that the attacker selects one tool to use at each phase of an attack strategy. This can be extended to using multiple tools in a sequential manner.

Second, the simulation study has some limitations. (i) We only consider simple decision-making algorithms, which are sufficient for demonstrating the usefulness of the framework but need to accommodate more sophisticated decision-making algorithms. (ii) We use some “synthetic” and simplifying scenarios owing to the lack of real data, meaning that the findings may not be generalized to other scenarios. Specifically, the assumption that each attack phase takes place at one time step may limit the validity of Insight 1; the assumption that each exploit incurs the same cost to the attacker may limit the validity of Insight 2; the assumption that each implementation of a program is equally vulnerable may limit the validity of Insight 3; the independence assumption that different implementations do not have common vulnerabilities holds in some settings [53], [54] but may not hold in general. In order to see the potential impact of these assumptions, we conduct additional experiments where different implementations can contain common vulnerabilities. While omitting the experimental details owing to space limit, they do show that the independence assumption does cause overestimates of the effectiveness of employing network diversity. This resonates the results of earlier theoretical studies [89], [97], [98].

6 CONCLUSION

We proposed a framework for quantifying the cybersecurity effectiveness of enforcing (dynamic) network diversity, including a suite of security metrics for measuring attacker’s cost (incurred by obtaining exploits) and defender’s operational cost (incurred by re-employing network diversity). We conducted simulation experiments to measure these metrics with respect to a number of dynamic diversity strategies and drew insights from the experimental results.

There are many open problems. In addition to those mentioned in Section 5, we highlight two: How can we obtain analytic results without making strong assumptions? How can we quantify attack power and defense power without making strong assumptions?

Acknowledgement. We thank the reviewers for their constructive comments that guided us in revising the paper. We thank Eric Ficke for proofreading the paper. We thank Dr. Alexander Kott for illuminating discussions. This work was supported in part by ARO Grant #W911NF-17-1-0566, NSF Grants #2115134 and #2122631 (#1814825), and by a Grant from the State of Colorado.

REFERENCES

- [1] D. Geer, R. Bace, P. Gutmann, P. Metzger, C. Pfleeger, J. Quarterman, and B. Schneier, “Cyberinsecurity: The cost of monopoly,” <http://cryptome.org/cyberinsecurity.htm>, 27 September 2003.
- [2] M. Stamp, “Risks of monoculture,” *Commun. ACM*, vol. 47, pp. 120–, Mar. 2004.
- [3] Y. Zhang, H. Vin, L. Alvisi, W. Lee, and S. Dao, “Heterogeneous networking: a new survivability paradigm,” in *Proceedings of the 2001 workshop on New security paradigms*, pp. 33–39, 2001.
- [4] B. Baudry and M. Monperrus, “The multiple facets of software diversity: Recent developments in year 2000 and beyond,” *ACM Computing Surveys (CSUR)*, vol. 48, no. 1, pp. 1–26, 2015.
- [5] P. Larsen, A. Homescu, S. Brunthaler, and M. Franz, “Sok: Automated software diversity,” in *IEEE S&P*, pp. 276–291, 2014.
- [6] L. Chen, “N-version programming: A fault-tolerant approach to reliability of software operation,” in *Proc. FTCS*, pp. 3–9, 1978.
- [7] A. Avizienis, “The n-version approach to fault-tolerant software,” *IEEE Transactions on software engineering*, no. 12, pp. 1491–1501, 1985.
- [8] M. Hiltunen, R. Schlichting, C. Ugarte, and G. Wong, “Survivability through customization and adaptability: The cactus approach,” in *Proc. DISCEX*, vol. 1, pp. 294–307, IEEE, 2000.
- [9] T. Jackson, B. Salamat, A. Homescu, K. Manivannan, G. Wagner, A. Gal, S. Brunthaler, C. Wimmer, and M. Franz, “Compiler-generated software diversity,” in *Moving Target Defense*, pp. 77–98, Springer, 2011.
- [10] A. Homescu, T. Jackson, S. Crane, S. Brunthaler, P. Larsen, and M. Franz, “Large-scale automated software diversity - program evolution redux,” *IEEE TDSC*, vol. 14, no. 2, pp. 158–171, 2017.
- [11] M. Franz, “From fine grained code diversity to JIT-ROP to execute-only memory: The cat and mouse game between attackers and defenders continues,” in *Proc. MTD’15*, p. 1, 2015.
- [12] S. Crane, A. Homescu, S. Brunthaler, P. Larsen, and M. Franz, “Thwarting cache side-channel attacks through dynamic software diversity,” in *Proc. NDSS*, 2015.
- [13] P. Larsen, S. Brunthaler, and M. Franz, “Security through diversity: Are we there yet?,” *IEEE Security & Privacy*, vol. 12, no. 2, pp. 28–35, 2014.
- [14] S. Bhatkar, D. DuVarney, and R. Sekar, “Address obfuscation: An efficient approach to combat a broad range of memory error exploits,” in *Proc. USENIXSecurity*, no. 2, pp. 291–301, 2003.
- [15] S. Forrest, A. Somayaji, and D. Ackley, “Building diverse computer systems,” in *Operating Systems*, 1997, pp. 67–72, IEEE, 1997.
- [16] H. Etoh, “Gcc extensions for protecting applications from stack-smashing attacks,” <http://www.research.ibm.com/trl/projects/security/ssp/>, 2000.
- [17] J. Xu, Z. Kalbarczyk, and R. Iyer, “Transparent runtime randomization for security,” in *Proc. IEEE SRDS’2003*, pp. 260–269, 2003.

- [18] G. Kc, A. Keromytis, and V. Prevelakis, "Countering code-injection attacks with instruction-set randomization," in *Proc. ACM CCS'2003*, pp. 272–280, 2003.
- [19] E. Barrantes, D. Ackley, T. Palmer, D. Stefanovic, and D. Zovi, "Randomized instruction set emulation to disrupt binary code injection attacks," in *Proc. CCS*, pp. 281–289, ACM, 2003.
- [20] M. Chew and D. Song, "Mitigating buffer overflows by operating system randomization," 2002.
- [21] B. Salamat, T. Jackson, G. Wagner, C. Wimmer, and M. Franz, "Runtime defense against code injection attacks using replicated execution," *IEEE TDSC*, vol. 8, no. 4, pp. 588–601, 2011.
- [22] "A new defense for navy ships: Protection from cyber attacks." <https://www.onr.navy.mil/en/Media-Center/Press-Releases/2015/RHIMES-Cyber-Attack-Protection.aspx>.
- [23] R. Rodrigues, M. Castro, and B. Liskov, "Base: Using abstraction to improve fault tolerance," *ACM SIGOPS Operating Systems Review*, vol. 35, no. 5, pp. 15–28, 2001.
- [24] T. Roeder and F. B. Schneider, "Proactive obfuscation," *ACM Transactions on Computer Systems (TOCS)*, vol. 28, no. 2, pp. 1–54, 2010.
- [25] T. Wood, R. Singh, A. Venkataramani, P. Shenoy, and E. Cecchet, "Zz and the art of practical bft execution," in *Proceedings of the sixth conference on Computer systems*, pp. 123–138, 2011.
- [26] M. Platania, D. Obenshain, T. Tantillo, R. Sharma, and Y. Amir, "Towards a practical survivable intrusion tolerant replication system," in *Proc. IEEE SRDS*, pp. 242–252, 2014.
- [27] M. Garcia, A. Bessani, and N. Neves, "Lazarus: Automatic management of diversity in bft systems," in *Proc. Middleware*, pp. 241–254, 2019.
- [28] H. Moniz, N. Neves, M. Correia, and P. Verissimo, "Ritas: Services for randomized intrusion tolerance," *IEEE TDSC*, vol. 8, no. 1, pp. 122–136, 2008.
- [29] J. Just, J. Reynolds, L. Clough, M. Danforth, K. Levitt, R. Maglich, and J. Rowe, "Learning unknown attacks—a start," in *Proc. RAID*, pp. 158–176, Springer, 2002.
- [30] E. Totel, F. Majorczyk, and L. Me, "Cots diversity based intrusion detection and application to web servers," in *Proc. RAID*, pp. 43–62, Springer, 2005.
- [31] M. Xu and T. Kim, "Platpal: Detecting malicious documents with platform diversity," in *Proc. USENIXSecurity*, pp. 271–287, 2017.
- [32] A. J. O'Donnell and H. Sethu, "On achieving software diversity for improved network security using distributed coloring algorithms," in *Proceedings of the 11th ACM Conference on Computer and Communications Security (CCS'04)*, pp. 121–131, 2004.
- [33] R. Garcia-Lebron, D. J. Myers, S. Xu, and J. Sun, "Node diversification in complex networks by decentralized colouring," *Journal of Complex Networks*, vol. 7, pp. 554–563, 5 2019.
- [34] C. Huang, S. Zhu, and R. Erbacher, "Toward software diversity in heterogeneous networked systems," in *Proc. IFIP DBSec*, pp. 114–129, Springer, 2014.
- [35] Y. Yang, S. Zhu, and G. Cao, "Improving sensor network immunity under worm attacks: a software diversity approach," in *Proc. Mobihoc'2008*, pp. 149–158, 2008.
- [36] M. Garcia, A. Bessani, I. Gashi, N. Neves, and R. Obelheiro, "Analysis of operating system diversity for intrusion tolerance," *Softw. Pract. Exper.*, vol. 44, pp. 735–770, June 2014.
- [37] D. Borbor, L. Wang, S. Jajodia, and A. Singhal, "Diversifying network services under cost constraints for better resilience against unknown attacks," in *Proc. IFIP DBSec*, pp. 295–312, 2016.
- [38] S. Neti, A. Somayaji, and M. E. Locasto, "Software diversity: Security, entropy and game theory," in *HotSec*, 2012.
- [39] L. Wang, M. Zhang, S. Jajodia, A. Singhal, and M. Albanese, "Modeling network diversity for evaluating the robustness of networks against zero-day attacks," in *European Symposium on Research in Computer Security*, pp. 494–511, Springer, 2014.
- [40] M. Zhang, L. Wang, S. Jajodia, A. Singhal, and M. Albanese, "Network diversity: a security metric for evaluating the resilience of networks against zero-day attacks," *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 5, pp. 1071–1086, 2016.
- [41] O. Temizkan, S. Park, and C. Saydam, "Software diversity for improved network security: optimal distribution of software-based shared vulnerabilities," *Information Systems Research*, vol. 28, no. 4, pp. 828–849, 2017.
- [42] J. Boyd, "The essence of winning and losing," 28 June 1995.
- [43] H. Chen, J. Cho, and S. Xu, "Quantifying the security effectiveness of firewalls and dmzs," in *Proc. HoTSoS'2018*, pp. 9:1–9:11, 2018.
- [44] Z. Li, D. Zou, S. Xu, H. Jin, H. Qi, and J. Hu, "Vulpecker: an automated vulnerability detection system based on code similarity analysis," in *Proc. ACSAC'16*, pp. 201–213, 2016.
- [45] Z. Li, D. Zou, S. Xu, X. Ou, H. Jin, S. Wang, Z. Deng, and Y. Zhong, "Vuldeepecker: A deep learning-based system for vulnerability detection," in *Proc. NDSS'18*, 2018.
- [46] D. Zou, S. Wang, S. Xu, Z. Li, and H. Jin, "μvuldeepecker: A deep learning-based system for multiclass vulnerability detection," *IEEE TDSC*, 2020.
- [47] Z. Li, D. Zou, S. Xu, H. Jin, Y. Zhu, Z. Chen, S. Wang, and J. Wang, "Sysevr: A framework for using deep learning to detect software vulnerabilities," *IEEE Transactions on Dependable and Secure Computing (accepted for publication)*, 2021.
- [48] Z. Li, D. Zou, S. Xu, Z. Chen, Y. Zhu, and H. Jin, "Vuldeelocator: A deep learning-based fine-grained vulnerability detector," *IEEE TDSC*, 2021.
- [49] E. Hutchins, M. Cloppert, and R. Amin, "Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains," *Leading Issues in Information Warfare & Security Research*, vol. 1, no. 1, p. 80, 2011.
- [50] "Mitre att&ck®." <https://attack.mitre.org/versions/v7/>.
- [51] N. Lynch, *Distributed Algorithms*. Morgan Kaufmann, 1996.
- [52] M. Magnani and L. Rossi, "The ML-Model for Multi-layer Social Networks," in *IEEE ASONAM*, pp. 5–12, 2011.
- [53] J. Han, D. Gao, and R. Deng, "On the effectiveness of software diversity: A systematic study on real-world vulnerabilities," in *Proc. DIMVA*, pp. 127–146, Springer, 2009.
- [54] M. Garcia, A. Bessani, I. Gashi, N. Neves, and R. Obelheiro, "Os diversity for intrusion tolerance: Myth or reality?," in *Proc. IEEE/IFIP DSN*, pp. 383–394, 2011.
- [55] A. Applebaum, D. Miller, B. Strom, C. Korban, and R. Wolf, "Intelligent, automated red team emulation," in *Proc. ACSAC*, pp. 363–373, 2016.
- [56] J. Knight and N. Leveson, "An experimental evaluation of the assumption of independence in multiversion programming," *IEEE TSE*, no. 1, pp. 96–109, 1986.
- [57] M. Henriques, "Diverse intrusion-tolerant systems," 2019.
- [58] H. Shacham, M. Page, B. Pfaff, E. Goh, N. Modadugu, and D. Boneh, "On the effectiveness of address-space randomization," in *Proc. CCS'2004*, pp. 298–307, 2004.
- [59] L. Davi, C. Liebchen, A. Sadeghi, K. Snow, and F. Monrose, "Isomeron: Code randomization resilient to (just-in-time) return-oriented programming," in *Proc. NDSS*, 2015.
- [60] I. Evans, S. Fingeret, J. Gonzalez, U. Otgonbaatar, T. Tang, H. Shrobe, S. Douskos, M. Rinard, and H. Okhravi, "Missing the point (er): On the effectiveness of code pointer integrity," in *IEEE Symposium on Security and Privacy*, pp. 781–796, 2015.
- [61] R. Strackx, Y. Younan, P. Philippaerts, F. Piessens, S. Lachmund, and T. Walter, "Breaking the memory secrecy assumption," in *Proc. European Workshop on System Security'2009*, pp. 1–8.
- [62] R. Rudd, R. Skowrya, D. Bigelow, V. Dedhia, T. Hobson, S. Crane, C. Liebchen, P. Larsen, L. Davi, M. Franz, A. Sadeghi, and H. Okhravi, "Address oblivious code reuse: On the effectiveness of leakage resilient diversity," in *Proc. NDSS'2017*.
- [63] B. Gras, K. Razavi, E. Bosman, H. Bos, and C. Giuffrida, "Aslr on the line: Practical cache attacks on the mmu," in *Proc. NDSS*, vol. 17, p. 13, 2017.
- [64] K. Snow, F. Monrose, L. Davi, A. Dmitrienko, C. Christopher, and A. Sadeghi, "Just-in-time code reuse: On the effectiveness of fine-grained address space layout randomization," in *Proc. IEEE S&P'2013*, pp. 574–588.
- [65] A. Sovarel, D. Evans, and N. Paul, "Where's the feeb? the effectiveness of instruction set randomization," in *Proc. USENIX Security'2005*.
- [66] Y. Weiss and E. Barrantes, "Known/chosen key attacks against software instruction set randomization," in *Proc. ACSAC'2006*, pp. 349–360, IEEE, 2006.
- [67] B. Cox, D. Evans, A. Filipi, J. Rowanhill, W. Hu, J. Davidson, J. Knight, A. Tuong, and J. Hiser, "N-variant systems: A secretless framework for security through diversity," in *Proc. USENIXSecurity*, pp. 105–120, 2006.
- [68] J. Cho, D. S. P. H. Alavizadeh, S. Yoon, N. Asher, T. Moore, D. Kim, H. Lim, and F. Nelson, "Toward proactive, adaptive defense: A survey on moving target defense," *IEEE Commun. Surv. Tutor*, vol. 22, no. 1, pp. 709–745, 2020.

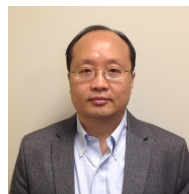
- [69] Y. Han, W. Lu, and S. Xu, "Characterizing the power of moving target defense via cyber epidemic dynamics," in *HotSoS*, pp. 1–12, 2014.
- [70] Y. Huang and A. Ghosh, "Introducing diversity and uncertainty to create moving attack surfaces for web services," in *Moving target defense*, pp. 131–151, Springer, 2011.
- [71] Y. Luo, B. Wang, and G. Cai, "Effectiveness of port hopping as a moving target defense," in *Proc. ICCST*, pp. 7–10, IEEE, 2014.
- [72] T. Carroll, M. Crouse, E. Fulp, and K. Berenhaut, "Analysis of network address shuffling as a moving target defense," in *Proc. ICC'14*, pp. 701–706, IEEE, 2014.
- [73] J. Jafarian, E. Al-Shaer, and Q. Duan, "An effective address mutation approach for disrupting reconnaissance attacks," *IEEE TIFS*, vol. 10, no. 12, pp. 2562–2577, 2015.
- [74] C. Phillips and L. Swiler, "A graph-based system for network-vulnerability analysis," in *Proc. NSPW'1998*, pp. 71–79.
- [75] O. Sheyner, J. Haines, S. Jha, R. Lippmann, and J. Wing, "Automated generation and analysis of attack graphs," in *Proc. 2002 IEEE Symposium on Security and Privacy*, pp. 273–284.
- [76] R. Ritchey and P. Ammann, "Using model checking to analyze network vulnerabilities," in *Proc. IEEE S&P*, pp. 156–165, 2000.
- [77] M. Albanese, S. Jajodia, and S. Noel, "Time-efficient and cost-effective network hardening using attack graphs," in *Proc. DSN'2012*, pp. 1–12.
- [78] P. Ammann, D. Wijesekera, and S. Kaushik, "Scalable, graph-based network vulnerability analysis," in *Proc. CCS*, pp. 217–224, 2002.
- [79] Y. Cheng, J. Deng, J. Li, S. DeLoach, A. Singhal, and X. Ou, "Metrics of security," in *Cyber defense and situational awareness*, pp. 263–295, Springer, 2014.
- [80] J. Homer, S. Zhang, X. Ou, D. Schmidt, Y. Du, S. Rajagopalan, and A. Singhal, "Aggregating vulnerability metrics in enterprise networks using attack graphs," *Journal of Computer Security*, vol. 21, no. 4, pp. 561–597, 2013.
- [81] M. Zhang, L. Wang, S. Jajodia, and A. Singhal, "Network attack surface: Lifting the concept of attack surface to the network level for evaluating networks' resilience against zero-day attacks," *IEEE TDSC*, 2018.
- [82] L. Wang, M. Zhang, and A. Singhal, "Network security metrics: from known vulnerabilities to zero day attacks," in *From Database to Cyber Security*, pp. 450–469, Springer, 2018.
- [83] S. Xu, "The cybersecurity dynamics way of thinking and landscape," in *ACM Workshop on Moving Target Defense*, 2020.
- [84] S. Xu, "Cybersecurity dynamics: A foundation for the science of cybersecurity," in *Proactive and Dynamic Network Defense*, pp. 1–31, Springer, 2019.
- [85] S. Xu, "Cybersecurity dynamics," in *Proc. HotSoS'14*, pp. 14:1–14:2, 2014.
- [86] X. Li, P. Parker, and S. Xu, "A stochastic model for quantitative security analyses of networked systems," *IEEE Transactions on Dependable and Secure Computing*, vol. 8, no. 1, pp. 28–43, 2011.
- [87] S. Xu, W. Lu, and L. Xu, "Push- and pull-based epidemic spreading in networks: Thresholds and deeper insights," *ACM TAAS*, vol. 7, no. 3, 2012.
- [88] S. Xu, W. Lu, and Z. Zhan, "A stochastic model of multivirus dynamics," *IEEE Transactions on Dependable and Secure Computing*, vol. 9, no. 1, pp. 30–45, 2012.
- [89] M. Xu and S. Xu, "An extended stochastic model for quantitative security analysis of networked systems," *Internet Mathematics*, vol. 8, no. 3, pp. 288–320, 2012.
- [90] W. Lu, S. Xu, and X. Yi, "Optimizing active cyber defense dynamics," in *Proc. GameSec'13*, pp. 206–225, 2013.
- [91] H. Cam, *Controllability and Observability of Risk and Resilience in Cyber-Physical Cloud Systems*, pp. 325–343. New York, NY: Springer New York, 2014.
- [92] S. Xu, W. Lu, L. Xu, and Z. Zhan, "Adaptive epidemic dynamics in networks: Thresholds and control," *ACM TAAS*, vol. 8, no. 4, 2014.
- [93] R. Zheng, W. Lu, and S. Xu, "Active cyber defense dynamics exhibiting rich phenomena," in *Proc. HotSoS*, 2015.
- [94] R. Zheng, W. Lu, and S. Xu, "Preventive and reactive cyber defense dynamics is globally stable," *IEEE TNSE*, vol. 5, no. 2, pp. 156–170, 2018.
- [95] L. Yang, X. Yang, and Y. Tang, "A bi-virus competing spreading model with generic infection rates," *IEEE Trans. Netw. Sci. Eng.*, vol. 5, no. 1, pp. 2–13, 2018.
- [96] Z. Lin, W. Lu, and S. Xu, "Unified preventive and reactive cyber defense dynamics is still globally convergent," *IEEE/ACM Trans. Netw.*, vol. 27, no. 3, pp. 1098–1111, 2019.
- [97] G. Da, M. Xu, and S. Xu, "A new approach to modeling and analyzing security of networked systems," in *Proc. HotSoS'14*, pp. 6:1–6:12, 2014.
- [98] M. Xu, G. Da, and S. Xu, "Cyber epidemic models with dependencies," *Internet Mathematics*, vol. 11, no. 1, pp. 62–92, 2015.
- [99] H. Chen, J. Cho, and S. Xu, "Quantifying the security effectiveness of network diversity," in *Proc. HoTSoS'2018*, p. 24:1, 2018.



Huashan Chen received the M.S. degree from the Institute of Information Engineering, Chinese Academy of Sciences, in 2016. He is currently pursuing the Ph.D. degree with the Department of Computer Science, University of Texas at San Antonio. His primary research interests are in cybersecurity, especially moving target defense and security metrics.



Hasan Cam is a Principal Machine Learning Scientist at Best Buy in USA. His current research interests include machine learning, cybersecurity, data analytics, autonomous agents, swarm intelligence, and networks. He currently works on the projects involved with applying machine learning and graph techniques to anomaly and fraud detection. Prior to joining Best Buy, he was a Computer Scientist at the Army Research Laboratory, where he served as the government lead for the Risk area in Cyber Collaborative Research Alliance. Previously, he also worked as a faculty member in the academia and a senior research scientist in the industry. He has published more than 35 journal and 70 conference papers, while having two patents and several patent applications. He has served as an editorial member of two journals, a guest editor of two special issues of journals, an organizer of symposiums and workshops, and a Technical Program Committee Member in numerous conferences. He received the Ph.D. degree in electrical and computer engineering from Purdue University, and the M.S. degree in computer science from Polytechnic University, New York. He is a Senior Member of IEEE.



Shouhuai Xu (M'14–SM'20) is the Gallogly Chair Professor in the Department of Computer Science, University of Colorado Colorado Springs (UCCS). Prior to joining UCCS, he was with University of Texas at San Antonio. He pioneered the Cybersecurity Dynamics approach as foundation of the emerging science of cybersecurity, with three pillars: first-principle cybersecurity modeling and analysis (the x -axis); cybersecurity data analytics (the y -axis); and cybersecurity metrics (the z -axis, to which the present study belongs). He co-initiated the International Conference on Science of Cyber Security and is serving as its Steering Committee Chair. He is/was an Associate Editor of *IEEE Transactions on Dependable and Secure Computing* (IEEE TDSC), *IEEE Transactions on Information Forensics and Security* (IEEE T-IFS), and *IEEE Transactions on Network Science and Engineering* (IEEE TNSE). He received a PhD degree in Computer Science from Fudan University.