The Geneva Papers, **2018**

www.genevaassociation.org

Reducing Informational Disadvantages to Improve Cyber Risk Management[†]

Sachin Shetty[a], Michael McShane[a], Linfeng Zhang[b], Jay P. Kesan[b], Charles A. Kamhoua[c], Kevin Kwiat[c] and Laurent L. Njilla[c]
[a]Strome College of Business, Old Dominion University, 2026 Constant Hall, Norfolk, VA 23529, USA.
E-mail: mmcshane@odu.edu
[b]University of Illinois at Urbana-Champaign, Champaign, IL, USA.
[c]Air Force Research Lab, Rome, NY, USA.

Effective cyber risk management should include the use of insurance not only to transfer cyber risk but also to provide incentives for insured enterprises to invest in cyber self-protection. Research indicates that asymmetric information, correlated loss, and interdependent security issues make this difficult if insurers cannot monitor the cybersecurity efforts of the insured enterprises. To address this problem, this paper proposes the Cyber Risk Scoring and Mitigation (CRISM) tool, which estimates cyberattack probabilities by directly monitoring and scoring cyber risk based on assets at risk and continuously updated software vulnerabilities. CRISM also produces risk scores that allow organisations to optimally choose mitigation policies that can potentially reduce insurance premiums.
The Geneva Papers (2018). https://doi.org/10.1057/s41288-018-0078-3

Keywords: cyber risk management; cyber insurance; vulnerability assessment; security risk scores; Bayesian belief networks; attack graphs

Article submitted 12 May 2017; accepted 20 December 2017; published online NaN NaN

Introduction

Losses from malicious cyber events are the result of deliberate actions by intelligent attackers who can adapt and change tactics as defences are implemented to inflict damage on organisations. The market for cyber insurance[1] is growing but has been limited by multiple issues, such as asymmetric information, correlated losses, and interdependent security. These issues also hinder carriers from using lower insurance premiums to reward appropriate cybersecurity investments by policyholders.[2]

Various researchers have argued that the lack of historical data plagues the development of the cyber insurance market. However, the dynamic nature of cyberattacks and rapid adaptation of the attackers poses another problem: the half-life of historical cyber data in effectively quantifying cyber risk and determining premiums is short.[3] A dynamic method for directly monitoring and measuring cyber risk beyond historical data is required. This

[†] Approved for public release: distribution unlimited. Case Number: 88ABW-2017-2684, dated 26 May 2017.
[1] We use the terms "cyber risk insurance" and "cyber insurance" interchangeably.
[2] See Gordon *et al*. (2003); Kesan *et al*. (2004); Böhme (2005); Ogut *et al*. (2005); Majuca *et al*. (2006).
[3] See Eling and Schnell (2016).

paper proposes the Cyber Risk Scoring and Mitigation (CRISM) tool, which allows insurers to reduce informational disadvantages. CRISM maps an organisation's systems and networks, ranks asset importance, identifies vulnerabilities using a continuously updated external source, and then computes Bayesian probabilities to generate attack graphs to produce risk scores that incorporate vulnerability scores and asset importance related to the vulnerability.

CRISM can be compared to devices installed in vehicles that monitor exactly what a specific driver is doing instead of basing the driver's risk on historical data and other rough underwriting measures. In effect, this direct monitoring can substantially reduce insurer informational disadvantages and allow premiums to be based on the individual insured rather than on risk pool averages. This paper deems such direct monitoring methods even more essential to market development of cyber insurance compared to auto insurance, where historical data and indirect risk proxies are effective underwriting measures. CRISM reduces the information advantage that both insureds and cybercriminals have relative to the insurer. With CRISM, an organisation would work with the insurer as a trusted partner in managing cyber risk. Insured organisations can also use CRISM risk scores to determine optimal cyber risk mitigation.

The following section documents the difficulty insurers face in using lower premiums as incentives for insureds to invest in effective cyber risk mitigation. The subsequent section offers an overview of the cyber insurance market followed by a terminology section that defines important terms, including those applied in the security area but not commonly used for risk management and insurance research. Then the CRISM tool is described in detail, while the final section concludes and suggests potential research directions that build on this paper.

Cyber insurance and incentives for investing in cyber risk mitigation

As for other types of insurance, cyber insurers face adverse selection and moral hazard. These are asymmetric information problems where the insured has an information advantage over the insurer concerning the insured's risk. Adverse selection is a hidden information problem where the organisation knows more about its own risk than the insurer, making it difficult for the insurer to charge a risk-based premium. Moral hazard is a hidden action problem where the insurer cannot observe the insured's actions after the policy is issued. The incentive is for the insured not to implement or maintain sufficient cyber risk mitigation actions because the costs are being shared with the insurer.[2]

Insurers can mitigate these information disadvantages by expending effort on understanding the insured risk. Underwriting to alleviate the adverse selection problem by understanding an organisation's cyber risk profile can include a detailed cybersecurity questionnaire and on-site risk assessments that include physical and technical analysis of the organisation's networks. To overcome the moral hazard problem, a cyber insurer needs to invest in methods that allow near continuous observation of the organisation's network while premiums are updated periodically based on the organisation maintaining cybersecurity measures. Some insurers partner with third-party firms to perform this ongoing monitoring of organisations.[2] Majuca *et al.*[2] argue that IT departments have devoted too much effort to the impossible task of avoiding risks instead of applying the full

complement of risk management techniques that include the optimal balance of risk retention, mitigation, and transfer.

An important cyber risk management question is the relation between buying insurance and the amount of self-protection provided by the insured. Ehrlich and Becker[4] show that a risk-based pricing signal for insurance encourages firms to increase self-protection. In other words, an incentive of lower insurance costs results in insureds investing more in mitigating the insured risk. Additional difficulties facing cyber insurers are correlated cyber losses and interdependent security, which can create externalities providing an incentive for organisations to freeride on other organisations that do invest in cybersecurity. Cyber losses can be correlated because organisations are connected by networks running similar software platforms. Interdependent security means that an organisation's cybersecurity depends not only on its own cybersecurity posture but also on effective cybersecurity being implemented by other organisations. Böhme[2] argues that correlated loss issues prevent cyber insurance from evolving into a mature market. One main problem is the difficulty in observing whether an organisation maintains effective security measures, which prevents insurers from offering discounts to encourage those measures. Ogut *et al.*[2] find that interdependent security can cause firms that buy cyber insurance to invest less in cyber risk mitigation than socially optimal if the insurer cannot observe the organisation's self-protection efforts. They find that providing lower premiums induces greater self-protection only if the insurer can observe the self-protection level. Schwartz *et al.*[5] investigate interdependent security in which an organisation's security can be dependent on a malicious actor in the network. The findings suggest that if malicious actors buy insurance and their security practices cannot be monitored, then cyber insurers have no incentive to base the insurance premium on the security practices of any users. They argue that a way to overcome this issue is the monitoring of insureds' cybersecurity practices by a third party who is trusted by insurer and insured.

In summary, the research covered in this section implies that a main impediment to the sustainable development of a sustainable cyber insurance market is the inability of insurers to monitor the insured's cybersecurity measures. This paper proposes CRISM as a tool for potentially overcoming this obstacle.

Cyber insurance issues

Research on the insurability of cyber risk has mainly developed in the computer science and IT disciplines. Gordon *et al.*[2] address the asymmetric information issues of adverse selection and moral hazard that hinder the development of the cyber insurance market. Böhme[2] describes how the near monoculture of software platforms results in correlation of cyber losses that slows the development of a cyber insurance market. Ogut *et al.*[2] investigate the difficulties caused by interdependent security, meaning that cybersecurity for an organisation depends on the actions/inactions of other organisations. Marotta *et al.*[6] investigate the difficulties of cyber insurance growing into a mature insurance market.

[4] See Ehrlich and Becker (1972).
[5] See Schwartz *et al.* (2010).
[6] See Marotta *et al.* (2017).

In one of the few papers from the risk management and insurance literature, Biener *et al.*[7] analyse the insurability of cyber risk using the Berliner[8] criteria and conclude that most problems of the insurability criteria are related to randomness of loss occurrence, information asymmetry, and cover limits. Risks that are problematic to insure are those that are correlated, suffer from lack of data, and are dynamic. Cyber risks by their nature are prone to moral hazard and adverse selection. Insurers mitigate these uncertainties with a higher safety loading added to premium, lower cover limits, and increased exclusions, resulting in polices that cover a small portion of cyber losses incurred by organisations.

Cyber insurance was first introduced to the market in 1997.[9] In terms of mechanism and functionality, cyber insurance has become much like other traditional insurance products. Typical cyber insurance policies include first-party coverages that indemnify insureds for remediation, notification, and other incurred expenses and third-party coverage for losses caused to third parties and related legal expenses. However, cyber insurance providers are struggling to assess cyber risk. A survey conducted by Risk Management Solutions (RMS)[10] shows that there are a wide variety of indicators used by different insurers, and underwriting often consists of questionnaires and interviews. Other types of insurance, such as auto and life, do not require complicated underwriting procedures because they are mature products backed up by historical data that are an effective predictor of future losses, which greatly mitigates the uncertainty in risk exposure faced by insurers.

A survey carried out by the SANS Institute[11] indicates that most cyber insurance underwriting practices today primarily rely on underwriters' experience and judgement. In other words, since cyber insurers do not have sufficient useful data regarding cyber risk or possess sophisticated cyber risk assessment techniques, information regarding the insured is opaque, and the link between cyber incidents and financial losses is not well established. This leads to high premium rates or limited coverage because the potential losses caused by moral hazard and adverse selection need to be taken into consideration during pricing to maintain insurers' profitability. In addition, coverages with high limits are rarely available in the cyber insurance market as insurers want to avoid high exposure. As a result, the insureds who need high limits must stack up multiple policies (to form so-called "tower policies"), which is usually a suboptimal solution because the coverages provided by different policies are misaligned most of the time.

Even with the previously described difficulties, the cyber insurance market is gaining an increasing number of insurance companies that are attracted by a potentially large market. Aon[12] used 2016 accounting statements for cyber insurers doing business in the U.S. and found that on average the cyber insurance line was profitable. Total U.S. cyber insurance premiums increased 30 per cent to USD 1.34 billion with 29 new entrants writing cyber insurance in 2016 for a total of 138 insurers. With the complications of determining risk-based premiums, this substantial increase of insurers entering the cybermarket is

[7] See Biener *et al.* (2015).
[8] See Berliner (1982).
[9] AIG offered the first cyber insurance policy to cover policyholders for third-party losses caused by breaches that originated from outside the company: http://www.aei.org/publication/cyber-insurance-why-is-the-market-still-largely-untapped/.
[10] See RMS (2016).
[11] See SANS Institute (2016).
[12] See Aon (2017).

worrisome. The equivalent of a "cyber hurricane" could have dire consequences for insurers and insureds. Lloyds[13] poses a plausible cyber catastrophe scenario resulting in a USD 243 billion to USD 1 trillion economic impact on the U.S. economy and USD 21.4 to USD 71.4 billion dollars in insured losses. Even a smaller loss could reverse the number of insurers jumping into the market and likely make more organisations willing to implement a direct monitoring system that allows full and effective cyber risk management. The terrorist attacks on the World Trade Center in 2001 collapsed the market for terrorism insurance until the U.S. government got involved as insurer of last resort to cap the maximum amount of insurer claims paid in future terrorist attacks. A similar scenario could occur as a result of catastrophic cyber losses.

Until insurers can better handle asymmetric information, correlated loss, interdependent security, and historical cyber data issues, organisations will not be fully satisfied with cyber insurance policies. In addition, without the ability to monitor cyber defence efforts, insurers will not lower insurance premiums to reward enterprises for improved cybersecurity, which is an impediment to effective cyber risk management. Developing a protocol to maintain business continuity and to recover expeditiously from a cyber loss requires a comprehensive *ex ante* assessment of the cyber risks posed by an enterprise's cyber infrastructure. The CRISM tool described later can potentially alleviate these issues.

Cyber insurance can be an effective way to improve resilience because it speeds up the process of recovery from financial losses, helps insured entities resume daily operations swiftly after cyber incidents, and can promote improved cybersecurity if informational issues can be overcome.

Terminology

After a review of the cyber risk literature, this paper adopts the "cyber risk" definition of Cebula and Young[14] and used by cyber insurance researchers in the risk management and insurance literature, such as Biener *et al.*[7]: "operational risks to information and technology assets that have consequences affecting the confidentiality, availability, or integrity of information or information systems".

To be effective, cyber risk management and insurance require cooperation across disciplines. Such necessary collaboration entails terminology confusion. The terms "peril" and "hazard" are used in the insurance literature, with "peril" meaning the immediate "cause of loss". Insurance policies name perils that are covered, such as fire, or those that are excluded, such as flood. A "hazard" is defined as "a condition that increases the frequency or severity of a loss",[15] for example, oily rags increase the likelihood of loss by fire. This paper uses terms from the security field that do not exactly map to insurance terminology.[16] "Vulnerability" is a flaw or weakness in software or hardware design that results in loss of confidentiality, integrity, and availability (CIA triad). "Threat" is potential for a specific vulnerability to be exploited.

[13] See Lloyds (2015).
[14] See Cebula and Young (2010).
[15] See Filtner (2010).
[16] See NIST (2012).

Information security risks are evolving with classification schemes most mature in the computer science and engineering literature. Jouini *et al.*[17] aggregate various cyber risk taxonomies into the two main categories of "external and internal" with further subcategorisation levels: "human, environmental, and technological; malicious and non-malicious; and accidental and intentional". The CRISM tool is designed for directly monitoring and scoring cyber risks based on continuously updated software vulnerabilities related to these seven categories under the "human, malicious, intentional" classification by Jouini *et al.*[17]: "destruction of information, corruption of information, theft, illegal usage, disclosure of information, denial of use, and elevation of privilege".

Advantages of CRISM over other tools

The goal of a risk assessment tool for cyber insurance is to perform quantitative assessment of exploitability and impact of attack surfaces within the cyber infrastructure. The following list gives five key requirements for an effective cyber risk assessment tool gathered by the authors from the National Institute of Standards and Technology (NIST) publications.[18,16]

Automatic discovery of vulnerabilities: due to constantly evolving attack surfaces, there is a need to discover the presence of exploitable vulnerabilities on a daily basis.
Lateral propagation analysis: adversaries exploit a chain of vulnerabilities to reach the attack goal, requiring analysis of the propagation of attackers through the chain of exploits. The analysis provides information on stepping stones, pivot points, attack paths, and vulnerable nodes, which in turn provides insights into potential strategies employed by adversaries.
Security metrics: the quantification of attack surfaces based on exploitability and impact analysis is required to develop baseline security risk scores. These metrics aid the decision maker in developing effective risk management policies.
Prioritised mitigation plan: the mitigation plan should provide an ordered list of vulnerabilities to patch or apply security controls to achieve a desired security score.
Compliance with a cybersecurity framework: NIST provides a policy framework for organisations to assess and enhance their ability to prevent, detect, and respond to attacks. Cyber risk assessment tools should take the recommendations of the framework into consideration for assessment of cyber risks.

There are several other cybersecurity risk assessment tools that provide a subset of the above requirements, but none satisfy all requirements. Tools such as Nessus, SAINT, open vulnerability assessment system (OpenVAS), and Nikto only perform key requirement (1) very well. Core Impact, Nexpose, Metasploit, and Qualys only perform (1) and (4) very well, while Bitsight and SecurityScorecard rate high on (1) and (3).[19] In the following sections, we describe how CRISM meets each of the requirements.[20]

[17] See Jouini *et al.* (2014).
[18] See NIST (2011).
[19] This list was compiled by the authors using a website for each tool.
[20] Traditional risk scoring tools, for example used by underwriters for determining policyholder risk for auto insurance, depend on proxy categories, such as driver age, gender, vehicle type, address, and policyholder-specific information such as credit score and claims history. CRISM can be considered analogous to the use of

Modelling description

CRISM provides an end-to-end automated capability to provide security scores and prioritised mitigation plans for cyber infrastructure. Based on the definitions provided previously, CRISM models threats that have potential to exploit vulnerabilities. The same vulnerability can be exploited by multiple threats. First, vulnerabilities are identified, and then the potential of exploiting these vulnerabilities, which are termed threats, is determined. Next, CRISM models the likelihood of the presence of threats, meaning the likelihood of exploiting the vulnerabilities, with different threats representing differing likelihoods of exploiting the vulnerabilities. The vulnerabilities in cyber systems and networks need to be identified because they increase the potential of threat exploitation. The ability to estimate and rank the impact of the attacks based on the importance of the assets being targeted is becoming increasingly crucial.

CRISM models specific threats by developing attack graphs.[21] The analysis of the attack graphs provides insight into attack surfaces and contributes to better understanding of cyber risk. Attack graph analysis looks beyond threat detection to focus on minimising the impact of threats. The data generated from vulnerability assessments are fed into a risk model to quantify the impact of the vulnerability. The identification of threats to the cyber infrastructure will necessitate developing security metrics to better ascertain the impact of the threats and the resultant response.

Cybersecurity risk modelling is an emerging research topic with no well-founded model to develop quantifiable metrics. This paper proposes security risk modelling to produce metrics such as security risk scores that allow the impact of threats to be ranked. With the availability of security metrics, enterprises can develop mitigation policies to ensure operational resiliency in the presence of cyber threats.

Cyber Risk Scoring and Mitigation (CRISM) tool

Figure 1 illustrates the five phases of the CRISM tool, which uses Bayesian attack graphs to measure security risk. CRISM is built over a platform optimised for vulnerability detection, attack graph analysis, and risk assessment that produces cyber risk scores. The platform can handle diverse network configurations and dynamic scaling. The tool also provides options to choose among several risk assessment models for generating, analysing, and evaluating attack paths based on security requirements for both cloud and non-cloud configurations. CRISM provides quantitative risk assessment and categorises attack paths based on the impact of vulnerabilities being exploited, and can illustrate the security risk scores via different visual metaphors, which allows practitioners to process information at several levels of granularity.

CRISM is implemented in five phases. In the first phase, CRISM maps the physical/virtual systems and networks. This phase is implemented using Network Mapper (Nmap),

Footnote 20 continued
telematics in determining auto insurance premiums where the continuously monitored actual behaviour of the specific policyholder is added as a factor in determining the premium (Baecke and Bocca, 2017).

[21] An attack graph models system security vulnerabilities and all potential sequences through which the vulnerabilities can be exploited.

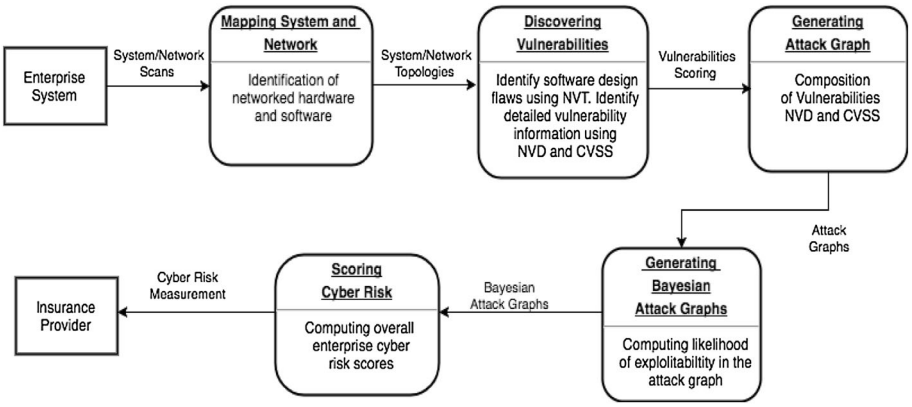

Figure 1. Five phases of the Cyber Risk Scoring and Mitigation (CRISM) tool.

which uses network probes to discover network topology and identify hosts, ports, and services.

The second phase of CRISM identifies vulnerabilities in the target environment. CRISM uses the OpenVAS to scan the network and detect vulnerabilities in every host. The OpenVAS project maintains a public feed of network vulnerability tests (NVTs). Based on the services it discovers through the scanning process, OpenVAS performs real attack tests using the NVTs to gather the vulnerability information. The common vulnerability scoring system (CVSS) provides score information for the vulnerabilities.[22] Finally, CRISM combines the network, hosts, port, service, and vulnerability information in a report.

In the third phase, an attack graph is automatically generated based on the vulnerabilities identified by the second phase.[23] Other state-of-the-art attack graph models make several assumptions that render the calculation of security risk ineffective. For instance, some attack graph models assume known starting points to launch an attack, which is usually not a reasonable assumption in the light of insider attacks when it is not clear how and where the attack originated. An incorrect assumption of where the attack is launched can easily change the risk rank in an attack graph. Other attack graph models ignore the importance of the asset to the overall domain while evaluating risk. In such models, risks are not normalised, so they cannot be compared and ranked. CRISM considers where the attack originated and normalises risks by incorporating the importance of assets in determining vulnerability scores.

The attack graphs are generated based on results of actual exploits conducted on the cyber infrastructure, not just by looking up the discovered software in vulnerability databases as is done by most vulnerability scanners. CRISM conducts actual exploits on the target machines. A successful exploit indicates that a known vulnerability can be truly exploited even if the information is not present in the vulnerability database.

[22] CVSS is an open framework for estimating and quantifying software vulnerabilities of various vendors.

[23] The CRISM approach to using attack graphs and converting the CVSS base scores into probabilities leverages work from other researchers, such as Wang *et al.* (2008); Poolsappasit *et al.* (2012); Homer *et al.* (2013).

CRISM uses the open source multi-host multi-stage vulnerability (MulVal) analysis language to generate the attack graph. First, the scanning result of OpenVAS is analysed and the topology and vulnerability information is extracted, then data from the National Vulnerability Database (NVD) and CVSS mentioned earlier are leveraged to collect the scores of the risks.[24]

The attack graph model incorporates asset importance and vulnerability scores. The authors developed an algorithm based on PageRank to compute asset values. For example, the impact of exploiting vulnerability X on Node A vs on Node B depends on the importance of the nodes. Node A may represent a system used across the enterprise, whereas Node B use may be limited to a single department. To determine vulnerability scores, CRISM again draws on the CVSS. The Common Vulnerabilities Exposures (CVEs)[25] dictionary is used to acquire the unique identifiers for known vulnerabilities. The authors developed a technique to combine the vulnerability scores and asset values into an attack graph model for assets before they are subject to known vulnerabilities.

In the fourth phase, CRISM computes Bayesian probabilities to generate Bayesian attack graphs. In particular, the authors adapt the notion of Bayesian belief networks to encode the contribution of different security conditions during system compromise.

CVSS can be used to produce a score that includes the likelihood of a vulnerability being exploited. The three components of the score are determined by base, temporal, and environmental metrics.[26] The base metric reflects vulnerability characteristics that do not change across user environments and through time. The temporal metric represents scores that can vary over time but are constant across user environments. The environmental metric reflects vulnerability characteristics that apply to a specific user environment. A simplified base metric example follows to illustrate how the likelihood of a vulnerability being exploited by an attacker is estimated:

$$\text{Exploitability} = 2 * \text{AccessVector} * \text{AccessComplexity} * \text{Authentication}, \quad (1)$$

where the 2 in the equation normalises values between 0 and 1; AccessVector represents the way the vulnerability can be exploited, such as requiring local, adjacent network, or network accessibility; AccessComplexity represents the complexity required for the

[24] NVD is a repository that provides CVSS scores for all known vulnerabilities for software and operating systems. The NVD was created by the Department of Homeland Security to inform the public about common computer vulnerabilities (http://nvd.nist.gov/). It is maintained by the National Institute of Standards and Technology (NIST). Before CVSS, there was no common platform to identify the vulnerabilities, and therefore vendors used their own methods for scoring the vulnerabilities. The National Infrastructure Assurance Council (NIAC) launched CVSS in 2005 (http://www.dhs.gov/national-infrastructure-advisory-council). Several major organisations such as CERT, IBM, and Cisco were involved in the development of CVSS. These organisations also use these metrics to prioritise the response to the vulnerabilities they encounter in their day-to-day activities. CVSS is currently maintained by the Forum of Incident Response and Security Teams (FIRST).

[25] CVE is a dictionary that assigns unique identifiers for all the security vulnerabilities that are publicly known (http://cve.mitre.org/about/index.html). CVE is used as the industry standard for vulnerability and exposures names. Once a vulnerability is discovered, it is assigned a unique CVE Identifier (e.g. CVE-2012-0015), which includes a brief description and references, such as advisories or vulnerability reports. CVE was quickly adopted by organisations, and its use is so widespread that organisations are producing "CVE Compatible" products and services.

[26] For CVSS details related to these components, refer to https://www.first.org/cvss/v2/guide.

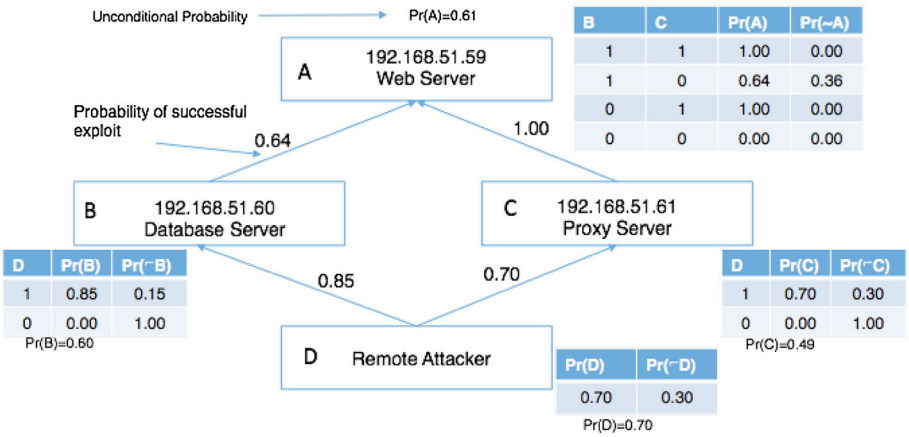

B	C	Pr(A)	Pr(~A)
1	1	1.00	0.00
1	0	0.64	0.36
0	1	1.00	0.00
0	0	0.00	0.00

D	Pr(B)	Pr(¬B)
1	0.85	0.15
0	0.00	1.00

D	Pr(C)	Pr(¬C)
1	0.70	0.30
0	0.00	1.00

Pr(D)	Pr(¬D)
0.70	0.30

Figure 2. Example of Bayesian attack graph modelling in CRISM.

vulnerability to be exploited; and Authentication indicates how many instances of authentication are required to exploit the vulnerability.

Figure 2 illustrates an example of a Bayesian attack graph to model the exploitability of vulnerabilities in three connected servers (web, database, and proxy servers). The nodes (A–C) represent software applications running on a computing device (such as a server) with the remote attacker represented by D. The edges (indicated by arrows connecting nodes) represent the presence of exploitable vulnerabilities. Each of the three servers can host several software applications, and each application can have multiple exploitable vulnerabilities reported in the CVSS. The number on each edge represents the probability of exploitability of the vulnerability. For the sake of simplicity and without loss of generality, each server is assumed to be running one software application with each application having one vulnerability. We assume that the goal of attacker D is to compromise the web server (A). The attacker can reach A by either compromising the vulnerability in B or the vulnerability in C.

Pr(D) is the probability of D attacking the network and is estimated to be 0.7 in this example, which is derived based on past attacks. A base exploitability score for the exploitability of node A from node B is determined from CVSS where in this example AccessVector = 1 (indicating that just network accessibility is required), AccessComplexity = 0.71 (indicating low complexity), and Authentication = 0.45 (indicating that multiple instances of authentication are required). Using these values in Equation (1), the probability of exploitability of Node A from Node B is $2 * 1 * 0.71 * 0.45 \simeq 0.64$. The probabilities on the other edges are calculated in a similar manner. Pr(B) is computed as $\Pr(B|D = 1) * \Pr(D = 1) = 0.85 * 0.7 = 0.60$. Pr(C) is computed as $\Pr(C|D = 1) * \Pr(D = 1) = 0.7 * 0.7 = 0.49$. Pr(A) is the joint unconditional probability of exploiting Node A from either Node B or C: $\Pr(A) = 0.64 * \Pr(B) * \Pr(D) + 1 * \Pr(C) * \Pr(D) = 0.64 * 0.6 * 0.7 + 1 * 0.49 * 0.7 = 0.61$.

In the fifth phase, CRISM analyses the Bayesian attack graphs to compute the overall risk scores with this output illustrated in Figure 3, which shows the risk scores for two network segments. The top half of Figure 3 indicates the overall risk scores for each

Figure 3. Comparative risk scores output by the fifth phase of CRISM.

network segment. A score of '0' indicates no cyber risk and a score of '10' indicates maximum cyber risk. The bottom half of the figure indicates the risk scores for each of the servers and software within the network segments. Most vulnerability scanners and penetration testing software provide scores. However, they do not provide a prioritised mitigation plan that provides recommendations on improving the scores. CRISM provides an ordered list of vulnerabilities that need to be patched to improve the score. This is also a process to gain insight into how the effectiveness of patching vulnerabilities affects the security score. Organisations can prioritise patching vulnerabilities based on the impact on the overall security score. CRISM can benchmark across various enclaves in an organisation to observe the variability in risk scores. In addition, the interplay between patching the vulnerabilities and changes to the risk score provides insight into factors contributing to the overall risk score.

What has just been described addresses the problem of developing cyber risk assessment to compute security risk scores for known vulnerabilities in enterprise systems. The availability of reliable and trustworthy risk metrics facilitates optimal choice of mitigation policies and reallocation of assets to minimise damage from cyberattacks. The development of this risk assessment capability involves understanding attack surfaces impacting systems and networks, which is achieved by identifying and continuously monitoring attack surfaces present in the hardware, software, virtualised systems, networks, and other areas for service providers and consumers. CRISM models specific cyber threats by analysing various attack paths to provide insight into the attack surfaces and contribute to better understanding of cyber risk. CRISM looks beyond detection of specific known vulnerabilities and focuses on the mitigation of the vulnerabilities.

Attack graph modelling techniques are employed to represent possible external or internal attacks. This type of modelling provides insights into the cost of fixing vulnerabilities, the vulnerabilities that need to be patched first, and the ranking of attack surfaces that are valuable to attackers. For instance, the model provides an optimal set of vulnerabilities that the administrator can patch to protect an asset with the least effort. It is critical for the security administrators to use risk scores to prioritise patches based on the importance of the assets. For

providing estimates of attack probabilities, CRISM draws on a database that is continuously updated as threats evolve to bypass the insurers' problem of historical data being made irrelevant by the dynamic nature of cyber threats. CRISM provides the ability to monitor enterprise networks and provide information that allows insurers to more reliably use risk-based pricing to incentivise investment in effective cybersecurity measures.

Conclusion

The demand for cyber insurance is growing, but insurers are wary of expanding coverage due to lack of credible data, interdependent security, asymmetric information issues such as adverse selection and moral hazard, and the potential for catastrophic aggregate losses in the face of correlated exposures among policyholders. The result is cyber insurance policies with gaps in coverage and lower limits that do not indemnify insureds for many cyber losses. Cyber risk is difficult to measure and price, with underwriting left to rely on experience and judgement instead of credible quantitative methods possible in other lines of insurance. The CRISM tool draws on a continuously updated vulnerability database that accounts for the dynamic nature of cyber threats to estimate attack probabilities.

The development of a sustainable cyber insurance market is also hampered by the lack of historical data as is the case for any relatively new risk. Even more worrisome is the dynamic nature of cyber risk, which may render historical data less useful than for other types of insurable risks. This paper proposes the CRISM tool to alleviate these issues and enable insurers to offer policies that more fully cover cyber losses. Previous work finds that a key issue for handling these problems is the inability of insurers to know the extent to which organisations are implementing cybersecurity defences.

CRISM is a direct monitoring tool that surveys organisational networks with continuously updated vulnerability information integrated with potential effects on related corporate asset values. CRISM produces risk scores that can be used by insurance underwriters and also by enterprise risk managers and information security officers to prioritise cyber risk mitigation. While CRISM can be a useful tool for both insurers and the insured, organisations are naturally hesitant to allow such outsider access to corporate networks. Enterprises are realising the existential challenge posed by cyber risks and have started to allow managed security service providers (MSSPs) access to their networks.[27] MSSPs trusted by both insurer and insured could be intermediaries to provide CRISM cyber scores to insurers. Until being allowed more access to organisational networks, cyber insurance markets will remain lacking as insurers will not offer the full cyber insurance coverage that policyholders desire. Cyber risk necessitates a full risk management process in which organisations mitigate cyber risks to the point where the costs of mitigation equal the benefits, and then decide on the amount to retain and to transfer. Direct monitoring by a tool such as CRISM will be important in implementing effective and holistic cyber risk management.

Future research based on this paper can follow various paths. The Bayesian probabilities used in CRISM to generate risk scores can be improved by incorporating more data. For example, by analysing databases containing large amounts of cyber event information, other

[27] See Tøndel *et al*. (2016).

factors could be determined to update probabilities and increase accuracy. Another potentially fertile research track is using CRISM output in other models to further improve risk-based pricing decisions for cyber insurance. One possibility is usage of CRISM risk scores in an option theoretic approach for modelling the occurrence of cyber intrusions. To determine the appropriate amount of capital reserving and the premium to charge for coverages, an insurer needs to predict the claim count in a portfolio during the policy period, which in the case of cyber insurance is the number of cyber incidents that are experienced by policyholders. However, the number of occurrences and types of cyberattacks changes dramatically from time to time, making predictions difficult using historical data. Option theoretic approaches have been used for credit risk assessment and have been applied to many other fields. This approach is based on incentives that motivate hackers and treat hacking activity not as a random action, but triggered when the benefit to the hacker is greater than the cost. With CRISM providing a quantitative measure of the security level, option theoretic models can translate CRISM scores into more tangible risk assessment measures to help insurers make better cyber risk measurement and insurance pricing decisions.

The attack graphs generated by CRISM do not have the capability of capturing zero-day attacks, meaning attacks not based on known vulnerabilities. Research modelling human behaviour to capture attacker intent and the effect on exploitability and impact scores could be a productive avenue for understanding zero-day attacks. The Bayesian attack graph framework allows the introduction of new causal relationships. Once developed, an adversarial behaviour model can be integrated into the framework.

Acknowledgements

This work was supported by the Office of the Assistant Secretary of Defense for Research and Engineering [OASD (R&E)] Agreement FA8750-15-2-0120 and Department of Homeland Security Grant 2015-ST-061-CIRC01.

References

Aon (2017) *Cyber Update: 2016 Cyber Insurance Profits and Performance*, from http://thoughtleadership.aonbenfield.com/Documents/20170504-ab-cyber-naic-supplemental-study.pdf, accessed 19 August 2017.

Baecke, P. and Bocca, L. (2017) 'The value of vehicle telematics data in insurance risk selection processes', *Decision Support Systems* 98: 69–79.

Berliner, B. (1982) Limits of Insurability of Risks, Englewood Cliffs, NJ: Prentice-Hall.

Biener, C., Eling, M. and Wirfs, J.H. (2015) 'Insurability of cyber risk: An empirical analysis', *The Geneva Papers on Risk and Insurance—Issues and Practice* 40(1): 131–158.

Böhme, R. (2005) 'Cyber-insurance revisited', in *Proceedings of the Fourth Workshop on the Economics of Information Security (WEIS 2005)*, Cambridge, MA: Kennedy School of Government, Harvard University.

Cebula, J.J. and Young, L.R. (2010) *A Taxonomy of Operational Cyber Security Risks*, Technical Note CMU/SEI-2010-TN-028, Software Engineering Institute, Carnegie Mellon University, from http://www.dtic.mil/get-tr-doc/pdf?AD=ADA537111, accessed 18 August 2017.

Ehrlich, I. and Becker, G.S. (1972) 'Market insurance, self-insurance, and self-protection', *Journal of Political Economy* 80(4): 623–648.

Eling, M. and Schnell, W. (2016) *Ten Key Questions on Cyber Risk and Cyber Risk Insurance*, The Geneva Association, from https://www.genevaassociation.org/media/954708/cyber-risk-10-key-questions.pdf, accessed 11 May 2017.

Filtner, A. (2010) *Foundations of Risk Management and Insurance*, CPCU Series, Malvern, PA: American Institute for Chartered Property Casualty Underwriters, p. 1.16.

Gordon, L.A., Loeb, M.P. and Sohail, T. (2003) 'A framework for using insurance for cyber-risk management', *Communications of the ACM* 46(3): 81–85.
Homer, J., Zhang, S., Ou, X., Schmidt, D., Du, Y., Rajagopalan, S. R. and Singhal, A. (2013) 'Aggregating vulnerability metrics in enterprise networks using attack graphs', *Journal of Computer Security* 21(4): 561–597.
Jouini, M., Rabai, L.B.A. and Aissa, A.B. (2014) 'Classification of security threats in information systems', *Procedia Computer Science* 32: 489–496.
Kesan, J.P., Majuca, R.P. and Yurcik, W.J. (2004) *The Economic Case for Cyberinsurance*, University of Illinois Law and Economics working papers.
Lloyds (2015) *Business Blackout*, Lloyds of London and University of Cambridge Centre for Risk Studies, from https://www.lloyds.com/news-and-insight/risk-insight/library/society-and-security/business-blackout, accessed 19 August 2017.
Majuca, R.P., Yurcik, W. and Kesan, J.P. (2006) *The Evolution of Cyberinsurance*, working paper, from https://arxiv.org/ftp/cs/papers/0601/0601020.pdf, accessed 18 August 2017.
Marotta, A., Martinelli, F., Nanni, S., Orlando, A. and Yautsiukhin, A. (2017) 'Cyber-insurance survey', *Computer Science Review* 24: 35–61.
NIST (2011) *Managing Information Security Risk: Organization, Mission, and Information System View*, NIST Special Publication 800-39, from https://dl.acm.org/citation.cfm?id=2206253, accessed 3 December 2017.
NIST (2012) *Guide for Conducting Risk Assessments*, NIST Special Publication 800-30, from http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-30r1.pdf, accessed 3 December 2017.
Ogut, H., Menon, N. and Raghunathan, S. (2005) 'Cyber insurance and IT security investment: Impact of interdependent risk', in *Workshop on the Economics of Information Security (WEIS)*, Harvard University.
Poolsappasit, N., Dewri, R. and Ray, I. (2012) 'Dynamic security risk management using Bayesian attack graphs', *IEEE Transactions on Dependable and Secure Computing* 9 (1): 61–74.
RMS (2016) *Managing Cyber Insurance Accumulation Risk*, Risk Management Solutions, Inc., and the Cambridge Centre for Risk Studies, from http://cambridgeriskframework.com/getdocument/39, accessed 11 May 2017.
SANS (2016) *Bridging the Insurance/InfoSec Gap: The SANS 2016 Cyber Insurance Survey*, from https://www.sans.org/reading-room/whitepapers/legal/bridging-insurance-infosec-gap-2016-cyber-insurance-survey-37062, accessed 11 May 2017.
Schwartz, G., Shetty, N. and Walrand, J. (2010) Cyber-insurance: Missing market driven by user heterogeneity, from https://pdfs.semanticscholar.org/d1db/6af4b7c93315e48c8ab407f1f75187a88687.pdf, accessed 18 August 2017.
Tøndel, I.A., Seehusen, F., Gjære, E.A. and Moe, M.E.G. (2016) 'Differentiating cyber risk of insurance customers: The insurance company perspective' in Buccafurri, F., Holzinger, A., Kieseberg, P., Tjoa, A.M. and Weippl, E. (eds.) *International Conference on Availability, Reliability, and Security*, Springer, pp. 175–190.
Wang, L., Islam, T., Long, T., Singhal, A. and Jajodia, S. (2008) 'An attack graph-based probabilistic security metric', in *Proceedings of The 22nd Annual IFIP WG 11.3 Working Conference on Data and Applications Security*, London, UK: pp. 283–296.

About the Authors

Sachin Shetty is an Associate Professor of Virginia Modeling, Analysis and Simulation Center at Old Dominion University. His research interests lie at the intersection of computer networking, network security, and machine learning. His laboratory conducts cloud and mobile security research and has received over USD 10 million in funding from federal agencies. He has authored and co-authored over 125 research articles in journals and conference proceedings and two books.

Michael McShane is an Associate Professor of Risk Management and Insurance at Old Dominion University. Prior to joining Old Dominion University, he received his Doctoral Degree in 2007. His main fields of research include enterprise risk management and flood insurance.

Linfeng Zhang is a Research Assistant at the Information Trust Institute at University of Illinois at Urbana Champaign. He has an Undergraduate Degree in Actuarial Science and Statistics, and a Master's Degree in Finance. His research interest is in predictive models for forecasting the claim number in a portfolio of cyber insurance policies.

Jay P. Kesan is appointed in the College of Law and the Department of Electrical and Computer Engineering. His academic interests are in the areas of technology, law, and business. Specifically, his recent work focuses on patent law and policy, cybersecurity and privacy, and biofuel regulation. He is best known for employing empirical, computational, and analytical methods in his research.

Charles A. Kamhoua is a Research Electronics Engineer at the Cyber Assurance Branch of the U.S. Air Force Research Laboratory, Rome, New York. His current research interests include the application of game theory to cyber security, survivability, cloud computing, hardware Trojan, online social network, wireless communication, and cyber threat information sharing. He has more than 60 technical publications in prestigious journals and for international conferences along with a Best Paper Award at the 2013 IEEE FOSINTSI.

Kevin Kwiat has been with the U.S. Air Force Research Laboratory (AFRL)in Rome, New York for over 32 years. Currently he is assigned to the Cyber Assurance Branch. His main research interest is dependable computer design. He is an advisor for the National Research Council. He has been by recognised by the AFRL Information Directorate with awards for best paper, excellence in technology teaming, and for outstanding individual basic research.

Laurent L. Njilla is a Research Electronics Engineer at the Cyber Assurance Branch of the U.S. Air Force Research Laboratory (AFRL), Rome, New York. He is responsible for conducting basic research in the areas of hardware design, game theory applied to cybersecurity and cyber survivability, hardware Trojan, online social network, cyber threat information sharing, and category theory. He is a member of the National Society of Black Engineers (NSBE).