Cyber Science 2024 Accepted Papers with Abstracts

A Case Study of a Ransomware and Social Engineering Competition
Aunshul Rege and Rachel Bleiman

Abstract: Abstract. The average cost of a ransomware attack in 2022 was $4.49 million and the recovery time rose from 7.8 days in 2021 to 14.9 days in 2022. Organizations often bring in professional negotiators to represent them during these attacks to lower ransom demands, buy time for incident response, get evidence of decryptors, and determine who is behind the attack; these research and negotiation skills fall under the area of social engineering. This paper shares a case study of a real-time virtual Ransomware and Social Engineering Competition (RSEC), which occurred in summer 2022 and hosted 8 high school and 6 undergraduate teams. The paper demonstrates how students learned about using psychological persuasion in ransomware negotiations, leveraging OSINT and the ATT&CK framework for attack attribution, and interfacing with the clients to best represent their needs. The paper discusses how the skills and knowledge stemming from RSEC are mapped to the NICE framework.

Malware Detection using Machine Learning Algorithms
Nureni Ayofe Azeez, Ogechukwu Juliet Nzeribe, Charles Van der Vyver and Ademola Philip Abidoye

Abstract: Recent years have seen a surge in malware attacks, which have caused considerable financial and operational damage to individuals and businesses alike. As malware techniques continue to evolve, it is essential to be proactive and adaptive in order to detect and reduce the risk of such attacks. To address this issue, this research is focused on creating a sophisticated malware detection system that utilizes machine learning algorithms to detect Malware attacks. With this technique, a comparative assessment of the algorithms used was carried out. The models were trained using four datasets. The machine learning algorithms used are Decision Trees, Xgboost, Logistic Regression, Naïve Bayes and Random Forest, are tested for their ability to detect Malware behaviors. A comparison of the results was carried using four evaluation metrics namely: Accuracy, Precision, Recall and F1-score. At the end of the experiments - For Dataset 1, XGBoost had the highest accuracy of 0.982456. For Dataset 2, Random Forest had the highest accuracy of 0.990432. For Dataset 3, XGBoost, Random Forest and Decision Tree had the highest accuracy of 1.0 meaning they achieved a perfect accuracy of 100%. For Dataset 4, both Random Forest and XGBoost had the highest accuracy with the values respectively 0.628474 and 0.628135. A bar chart was used to graphically represent our results for each evaluation metrics. Boxplot used to represent results showing the minimum, maximum, and quartile ranges.

An Exemplar Incident Response Plan for Security Operations Center Analysts
Enoch Agyepong and Cyril Onwubiko

Abstract:
The significance of a cyber security incident response plan cannot be overemphasised, especially at a time when cyberattacks are becoming increasingly prevalent and sophisticated. The average cost of a singular cyber incident response in 2023 was ~USD 4.45 million. This figure represents a notable increase over the past few years, showing a rising trend in the costs associated with cyber breaches. Therefore, it is imperative that organisations must have a focused and coordinated approach when responding to cyber incidents. Security Operations Centres (SOCs) need a well-defined incident response plan to facilitate the detection and response activities in the event of a cyberattack. While there are several well-known incident response plans, these plans typically have a broad scope and are designed to assist teams and organisations develop their own incident handling processes. Previous studies propose a tailored incident response plan for the Computer Security Incident Response Team (CSIRT). This paper presents an example incident response plan for SOC analysts. The proposed approach leverages the United States National Institute of Standards and Technology (NIST) incident response framework. The proposed approach will be beneficial to analysts with varying levels of experience/expertise, especially junior and early-career analysts, since research indicates that these categories of analysts are burdened by the complexity of security incident analysis.

RingGuard: A privacy protected Peer-to-Peer Federated Learning Framework
Narasimhan Kovalai, Ramsundar Tanikella, Rojalini Tripathy and Padmalochan Bera

Abstract: Recently, collaborative learning-based training is being used to develop accurate ML models. However, due to privacy and regulatory compliance, data owners are constrained to share their data. To address this, Federated Learning (FL) has emerged as a solution, enabling collaborative model training without the need to share raw data. Instead, only model parameters are exchanged among participating parties. However, FL faces many challenges, particularly its reliance on a central server, managing data heterogeneity and security challenges of parameter sharing. In this paper, we propose a peer-to-peer federated learning framework that leverages secure multiparty computation (SMC) to overcome these challenges. Our framework eliminates the dependency on a central server, ensuring enhanced security in parameter sharing. We implement our proposed framework for both Independent and Identically Distributed (IID) and non-Independent and Identically Distributed (non-IID) data distributions, using FedProx for non-IID data and FedAvg for IID data. We evaluated the performance of our framework using MNIST and FEMNIST datasets, incorporating varying numbers of clients, and conducted a comprehensive analysis. Experimental results demonstrate that our framework performs better in terms of test accuracy and smooth learning than existing centralized FL frameworks.

TIPS: Threat Sharing Information Platform for Enhanced Security
Lakshmi Rama Kiran Pasumarthy, Hisham Ali, William Buchanan, Jawad Ahmad, Audun Josang, Vasileios Mavroeidis and Mouad Lemoudden

Abstract: There is an increasing need to share threat information for the prevention of widespread cyber-attacks. While threat-related information sharing can be conducted through traditional information exchange methods, such as email communications etc., these methods are often weak in terms of their trustworthiness and privacy. Additionally, the absence of a trust infrastructure between different information-sharing domains also poses significant challenges. These challenges include redaction of information, the Right-to-be-forgotten, and access control to the information-sharing elements. These access issues could be related to time bounds, the trusted deletion of data, and the location of accesses. This paper presents an abstraction of a trusted information-sharing process which integrates Attribute-Based Encryption (ABE), Homomorphic Encryption (HE) and Zero Knowledge Proof (ZKP) integrated into a permissioned ledger, specifically Hyperledger Fabric (HLF). It then provides a protocol exchange between two threat-sharing agents that share encrypted messages through a trusted channel. This trusted channel can only be accessed by those trusted in the sharing and could be enabled for each data-sharing element or set up for long-term sharing.

Workshop Insights: Navigating Cybersecurity Regulations for Device Manufacturers and Healthcare Operators
Andrea Skytterholm, Lars Halvdan Flå and Martin Gilje Jaatun

Abstract: Both the manufacture and use of medical devices are heavily regulated, but stakeholders have a varying level of maturity, and often struggle to comply with rules and regulations. This paper reports on an empirical elicitation activity that sought to enumerate the challenges faced by (particularly smaller) device manufacturers and device operators (typically: hospitals, with a goal to informing the creation of tools that these stakeholders can use to address the challenges.

TRIST: Towards a container-based ICS testbed for cyber threat simulation and anomaly detection
Carol Lo, Jack Christie, Thu Yein Win, Zeinab Rezaeifar, Zaheer Khan and Phil Legg

Abstract: Cyber-attacks on Industrial Control Systems (ICS), as exemplified by the incidents at the Maroochy water treatment plant and the Ukraine's electric power grid, have demonstrated that cyber threats can inflict significant physical impacts. These incidents caused widespread service disruptions and substantial economic losses, underscoring the urgent need for an in-depth understanding of cyber threats in industrial environments. Industrial security research is usually conducted on physical testbeds to avoid safety issues, production interruptions and other operational

constraints in industrial processes. Nevertheless, security defenders often encounter obstacles in developing or accessing physical testbeds due to associated costs and complexities. These factors hinder research progress to devise early detection mechanisms for cyber threats – essential for effective incident response. To overcome these obstacles, this paper presents a container-based virtual testbed. Its lightweight architecture enables replicable and efficient deployment of testbeds at low cost for simulating cyber threats on Cyber-Physical Systems (CPS) – the cornerstone of industrial automation and control systems. Also, the container-based virtual testbed provides a cost-effective option for producing datasets for training, testing and optimization of unsupervised anomaly detection models. Besides, an evaluation on resource consumption is conducted. The paper also discusses the benefits and limitations of proposed container-based ICS testbeds and suggests future research areas.

Exploring DTrace as an Incident Response Tool for Unix Systems
Joe Duin, Sean McKeown and Mwrwan Abubakar

Abstract: Critical National Infrastructure (CNI) is often the target of sophisticated and sustained cyber-attacks perpetrated by advanced threat actors with considerable resources. These attacks can lead to interruptions in core services such as energy and water supplies, transportation, healthcare, and telecommunications. The effective and swift remediation of such attacks is contingent on the respective Digital Forensics and Incident Response (DFIR) professionals possessing the appropriate tooling and resources for the target environments. However, the Unix systems which often run critical infrastructure are poorly accommodated in comparison to their Windows and Linux counterparts. This paper seeks to expand the options available to DFIR analysts on Unix systems by exploring the potential for DTrace as a potential Incident Response utility. DTrace is included in many Unix operating systems by default, while also having support for Linux, Windows and macOS, making it a useful pre-packaged solution. We explore the utility of DTrace, and the visibility it provides into the OS and kernel, through a variety of proof-of-concept case studies based on tactics and techniques in the MITRE ATT\&CK framework. We find that DTrace's functionality lends itself well to a real-time monitoring and probing solution for Unix systems, which could potentially form the basis of an EDR solution to revolutionise Incident Response on such platforms.

A Survey on Firmware Security Fuzzers
Silje Marie Sørlien, Åse Marie Solnør, Karin Bernsmed and Martin Gilje Jaatun

Abstract: Security testing is challenging with low-level connected devices such as Internet of Things (IoT) devices, with limitations in memory, communication bandwidth, and processing power. As a result, traditional testing processes become difficult and time-consuming in such cases. In this paper, we survey recent fuzz testing platforms

suitable for security testing of embedded firmware, and assess how suitable they are for a specific use case.

Cybersecurity Upskilling for Criminal Justice Professionals
Yan Bai and Juan Li

Abstract: A diverse and interdisciplinary cybersecurity workforce plays a crucial role in enhancing cybersecurity worldwide. Criminal justice (CJ) professionals are at the forefront of investigating cybercrimes and bringing cyber-criminals to justice. Their expertise and efforts in cybersecurity are essential to maintaining the integrity of our digital ecosystem and ensuring the safety and security of individuals, organizations, and society. However, many CJ professionals have not received comprehensive cybersecurity education and training. This knowledge gap poses significant challenges and limitations to effectively combating cybercrime. With the sponsorship of National Science Foundation, The University of Washington Tacoma and North Dakota State University are developing scenario-based offensive security and Web-based showcase labs with interactive simulations and case studies in three progressive courses, revolutionizing cybersecurity education for future CJ professionals. This project integrates artificial intelligence into the curriculum to enhance CJ professionals' capabilities. Our ultimate goal is to develop a skilled workforce of CJ professionals with cybersecurity and privacy knowledge, addressing the critical need for such expertise in the field. Our poster will explain the motivations behind our project, the course framework used, and preliminary implementation results.

vSPACE: Voting in a Scalable, Privacy-Aware and Confidential Election
Se Elnour, William Buchanan, Paul Keating, Mwrwan Abubaka and Sirag Elnour

Abstract: The vSPACE experimental proof-of-concept (PoC) on the [True{\bf\textit{Elect}}][{\bf\textit{Anon}}Creds] protocol presents a novel approach to secure, private, and scalable elections, extending the TrueElect and ElectAnon protocols with the integration of AnonCreds SSI (Self-Sovereign Identity). Such a protocol PoC is situated within a Zero-Trust Architecture (ZTA) and leverages confidential computing, continuous authentication, multi-party computation (MPC), and well-architected framework (WAF) principles to address the challenges of cybersecurity, privacy, and trust over IP (ToIP) protection. Employing a Kubernetes confidential cluster within an Enterprise-Scale Landing Zone (ESLZ), vSPACE integrates Distributed Ledger Technology (DLT) for immutable and certifiable audit trails. The Infrastructure as Code (IaC) model ensures rapid deployment, consistent management, and adherence to security standards, making vSPACE a future-proof solution for digital voting systems.

Perception of cyber risk for assisted living technologies in the Norwegian healthcare sector – key stakeholders perceptions

Alvhild Skjelvik and Bian Yang

Abstract: Facing a demographic shift and healthcare personnel shortage, Norway and other European nations anticipate increased healthcare demands. Assisted living technologies are seen as potential solutions to manage demands, yet they carry inherent cybersecurity risks. While the majority of prior research within cybersecurity for assisted living technology has focused on tech-nical risks, this study examines cybersecurity risk perception among key stakeholders in the Norwegian healthcare sector. Utilizing the BCISQ-questionnaire and incorporating a behavioral simulation, we analyzed risk perception patterns among 185 survey respondents and 33 interviewees. Despite observable differences in risk perception among stakeholders, at-tributed to factors like human error, organizational affiliation, and technol-ogy familiarity, the sample size limits generalizability. Key findings highlight the risk perception related to assisted living technology amongst five stakeholder groups, where their perception of risk is formed by likelihood and consequence in light of confidentiality, integrity and availability. This study not only capture stakeholders' cybersecurity perception but also identify the differences between the groups, and factors influencing their perceptions.

Attention based Image Steganography Using CNNs in Integration with Quantization

Sravan Kumar Gatram, Kamalakanta Sethi, Piyush Joshi and Rakesh Kumar Sanodiya

Abstract: Steganography is a security mechanism that is primarily used to hide secret information in various forms of digital media. This is useful for transmitting sensitive information without disclosing its existence. In this paper, we present a novel deep learning based steganographic approach that integrates Convolutional Neural Networks (CNNs) with Quantization and Convolutional Block Attention Module (CBAM). Our steganography approach overcomes various shortcomings of existing methods, including low embedding capacity and visual distortions in images. The integration of CNNs in our proposed approach enables efficient feature extraction, while Quantization refines and enhances the extracted features which helps to preserve the image quality after embedding. In addition, the CBAM attention mechanism captures useful features while filtering out noise and irrelevant information. This helps us to reduce the feature distortions and improve the feature representations. We have implemented and tested the effectiveness of our steganographic approach with various evaluation metrics. The experimental results show that our model outperforms the state-of-the-art approaches.

C-MRiC
CENTRE FOR MULTIDISCIPLINARY RESEARCH, INNOVATION AND COLLABORATION ®

Improving Credit Card Fraud Detection with Combined Feature Extraction and Class Balance Optimisation Techniques

Boluwatife Ajibade-Ajibosin, Uchenna Daniel Ani, Theocharis Kyriacou and Mark Turner

Abstract: Credit card fraud is a persistent and evolving challenge that poses significant financial harm to cardholders, financial institutions, national, and global economies. The use of Machine Learning (ML) methods has heavily enhanced the detection of credit card fraud, offering improvements to other traditional credit card fraud detection approaches such as manual checks and rule-based methods. However, there are limitations that impact the performance and efficiency of the credit card detection using ML. This work addresses some of the current challenges associated with employing ML for credit card fraud by developing a robust credit card fraud detection model. The proposed model employs Syn-thetic Minority Over-sampling Technique (SMOTE) to address the class im-balance issue typically akin to credit card fraud datasets. A Recursive Feature Elimination with Cross-Validation (RFECV) scheme was utilized as a feature selection technique to select the optimal subset of features that can improve the accuracy of the credit card fraud detection model. Following rigorous evaluation, the proposed credit card fraud detection system demonstrates exceptional performance above other known existing systems across critical metrics including accuracy, recall, precision, F1 score, and ROC-AUC. Thus, is considered a better solution for more accurate detection of credit card fraud in-stances.

Digital Evidence from the Legal Practitioners' Perspective

Ibeabuchi Egbu, Jacques Ophoff and Annelize McKay

Abstract: Cybersecurity, much like traditional security, requires the collaborative effort of various stakeholders, including government, regulators, tech professionals, and legal practitioners. Each party must play their part effectively to achieve the objective of a cyber environment free from danger or threat. Ironically, legal practitioners who are responsible for prosecuting cybercrimes and administering justice face significant challenges in interpreting and applying the technical aspects of digital investigations. Existing research indicates that this difficulty largely stems from their predominantly non-technical background, which hinders their ability to fully grasp the implications of digital evidence. This research study aimed to investigate how employing analogies to traditional forms of evidence could aid legal practitioners in better understanding and applying the technical terms encountered in digital investigations. Building on existing literature on legal practitioners' digital awareness, this research study adopted a five-step methodological approach. First, a digital forensics report was analysed to identify technical terms. These technical terms were simplified and analogies to traditional evidence were constructed. These analogies were published via a website, and made accessible through a custom-built Microsoft

Word add-in. A survey involving 16 participants was carried out to assess the value of the simplification provided by the analogies, and the impact was evaluated against Bloom's taxonomy of cognitive learning. Results show that the analogies improved understanding of the digital forensics report and the way practitioners interpreted and applied technical terms found in the report.

Attack Resilient Federated Learning Framework
Sushant Kumar, Kasturi Routray and Padmalochan Bera

Abstract: Federated learning stands out as a promising paradigm for collaborative training of machine learning models where a server supervises the learning process while keeping sensitive data on the user devices. Here, training is decentralized and conducted on edge devices beyond the control of a server. This increases the potential for malicious clients to tamper with the learning process and compromise the global model, resulting in a significant security risk. The majority of existing solutions are tailored to scenarios where data exhibits independent and identically distributed (IID) characteristics across devices. A notable performance degradation is observed when the data distribution deviates from the independent and identically distributed (non-IID) scenario. In this paper, we first evaluate the performance of existing byzantine robust aggregation schemes in non-IID settings within an adversarial scenario. Then, we introduce a novel attack resilient aggregation scheme named FedResil with the objective of enhancing performance in the same adversarial environment. It leverages non-private data, which is collectively agreed upon by the participating clients before the training process begins, to delineate clusters of clients. Subsequently, the server applies existing byzantine robust aggregation rules to each cluster independently, generating model updates within each cluster. The model update from each cluster is then aggregated to construct the final global model. Through extensive experimentation, we demonstrate that FedResil in malicious settings achieves performance similar to those in scenarios where there is no malicious client.

Design of Pairing Free Attribute-based Encryption for Smart Grid Applications
Anjali Kumari, Sreevallabh Karanam, Swejan Annabathini and Kamalakanta Sethi

Abstract: Smart grids are light-weight, intelligent and reliable devices used daily for power supply. Smart grids continuously measure the incoming and outgoing flow of electricity and other connected smart grids. This data is shared with many stakeholders like the users (the actual user of electricity), phasor units etc. The data security here plays a crucial role. Initially pairing CP-ABE was used in order to provide fine grain access control and data integrity and confidentiality. But the computationally complexity of the cryptographic technique add over-head to the system reducing its performance. In order to improve the performance matrix of smart grids we propose the Pairing Free CP-ABE (PF CP-ABE) which uses replaces the complex bilinear

pairing with simpler scaler multiplication. At last, we have analyzed the security, access policy and collusion resistance properties along with the performance analysis of our system.

Critical Infrastructures in the Cloud
Martin Gilje Jaatun and Geir Kjetil Hanssen

Abstract: Cloud computing is increasingly being used not only to support critical infrastructure applications, but actually forms a vital part of them. This paper outlines some security requirements that are relevant to apply to critical infrastructure cloud applications.

Radio Frequency (RF) Cyber Threats and a novel RF Cyber Kill Chain
Carolyn Swinney and Woods John

Abstract: Cyber threats are defined as those that target a system's availability, integrity, or confidentiality. As the threat landscape increases with 5G and IoT connectivity, so does the potential to directly affect or access connected ‘system of systems’ using Radio Frequency (RF) connections. The open architecture and low-cost nature of today’s Software Defined Radio (SDR) technology present opportunities to easily exploit RF access points with cyber threats in a growing connected cyber terrain. SDRs can provide low-cost, simple discrete attacks on a system’s confidentiality, integrity, and availability, without a requirement for physical access. The Cyber Kill Chain, which focuses on external complex and persistent threats, is not fit for purpose for an RF-based threat. While traditional cyber threats may be increasing in complexity and time to render, it is suggested in this paper that RF-enabled cyber-attacks are becoming less complex and easier to implement due to the availability of low-cost SDR technology and associated high-power amplifiers and antennas. In this paper, a new RF Cyber Kill Chain is proposed to aid penetration testers and improve the cyber-security of connected systems. The RF Cyber Kill Chain is shown through a case study to accurately map three stages of attack: Reconnaissance, Preparation, and Execution. SDR technology is also shown in this paper to have enormous potential when combined with ma-chine learning (ML) to provide early warning of cyber-attacks utilising the RF spectrum.

C-MRiC
CENTRE FOR MULTIDISCIPLINARY RESEARCH, INNOVATION AND COLLABORATION ®

Implementation of Zones and Conduits in Industrial Control and Automation Systems

Lars Halvdan Flå, Mary Ann Lundteigen, Fredrik Gratte and Martin Gilje Jaatun

Abstract: Despite being established concepts in standards we argue that zones and particularly conduits can benefit from more detailed discussions of their architecture and implementation. In this paper we make 3 contributions towards this: describe detailed principles for developing conduits, and principles for connecting zones with potentially different security levels. Both of these processes are expressed in the firm of flow charts. Lastly, we discuss a few highlighted challenges related to the application of zones and conduits in practice.

Towards a Secure Manufacturing Framework for Single Malt Whisky in Industry 4.0

Jacob Connell, Dimitrios Kasimatis, Pavlos Papadopoulos, William J. Buchanan, Panagiotis Sarigiannidis and Nikolaos Pitropakis

Abstract: Recent instances have brought to light the inherent vulnerability of the manufacturing sector to cyber security threats, particularly within the food and beverage industries. With the rapid adoption of Industry 4.0, supply chains are exposed to an expanding surface area of potential attacks. This paper delves into the paramount issue of supply chain security and explores the utilisation of distributed ledger technology to enhance the reliability and security of the whisky manufacturing process. We introduce an experimental framework that is focused on performance enhancements and security optimisations. Furthermore, a review of the implemented testbed evaluates the robustness of the business logic, emphasising additional benefits derived from implemented features. This analysis serves to showcase the advancements made in response to recent disruptions within similar supply chains.

Utilizing YARA: An Effective Method for Phishing Attack Response

Ferenc Leitold

Abstract: Phishing attacks continue to be a major threat to cybersecurity, leveraging deception to extract sensitive information from unsuspecting victims. This paper explores the use of YARA (Yet Another Recursive Acronym), a powerful tool for malware identification and classification, in enhancing the detection and response to phishing attacks. We delve into several phishing techniques such as spear phishing, whaling, and clone phishing, outlining their evolution and the increasing sophistication of these threats. The paper discusses the creation and application of YARA rules to identify phishing indicators within emails, including suspicious language, generic greetings, and malicious links or attachments. By integrating YARA into cybersecurity operations, organizations can create customizable and efficient detection frameworks that adapt to emerging threats. Through practical examples and case studies, we

demonstrate the effectiveness of YARA in identifying both traditional and obfuscated phishing attacks, thereby improving incident response and mitigating the impact of these pervasive threats to individuals and businesses alike. This research highlights YARA's versatility and its significant role in bolstering cybersecurity defenses against phishing.