Research Statement

David Mohaisen (mohaisen@ucf.edu)

My research interests are in the area of computer systems security and online privacy. My current work attempts to understand and improve the security and privacy in various domains, including Internet of Things, blockchain systems, virtual reality and wearable systems, and next generation networks. Over the years, I developed interests in expertise in evidence-based security for domains such as software, mobile, Internet. and web applications. A common theme in my work is the development of learning techniques and tools for security: to understand codes, traffic, and infrastructure usage in real systems.

To date, I have published more than 150 peer-reviewed research papers, many of which are in top conferences (CCS [1], NDSS [2–4], PETS [5], IMC [6, 7], DSN [8–10], ICDCS [11–19], and INFOCOM [20–24], to name a few) and premier journals (IEEE Transactions [25–33]). In addition to scholarly publications, I contributed 10 patents, 1 standard (IETF), and several systems. During my academic tenure, I built a research program by attracting several research grants, awards, and gifts as a PI. I mentored two postdocs, who are now tenure-track assistant professors, and graduated seven doctoral students, six of whom went to academia; four as tenure-track assistant professors–e.g., Wayne State University and Layola University Chicago, and two as postdoctoral researchers at Georgia Tech and TAMU, respectively. Finally, I expect to graduate three more PhD students by the end of this year.

1 Recent Research Results

Broadly, my recent work has focused on networked systems security and privacy and malicious software (malware) analysis, including detection and attribution. My work has been motivated by the ever-increasing size of the attack surface of critical software and networked systems, which make automation for security very important. The basic questions that motivated my work are:

Q1. what are the features that can be used to fingerprint and attribute adversaries in different contexts?
Q2. how to retrieve, represent, and use those features to efficiently detect adversaries in context?
Q3. how robust are those features to temporal or adaptive changes in the adversaries' behavior?
Q4. how to use a refined understanding of those features to defend against those adversaries?

In pursuing those questions, I have built and used expertise in the domain of program analysis, protocol analysis, measurements, machine learning, vulnerability analysis, and secure systems design. Key results from my work included the development of systems for binary analysis and security [17,34–36], web systems security with threat attribution [37,38], mobile systems security [39,40], and security against denial of service attacks [32, 41, 42]. In the following, I summarize some of those results.

Software Security My earliest work on software security has been a dynamic analysis-based system, called AMAL, which classifies and clusters malware using autonomous feature extraction and expert labeled training data. AMAL sandboxes malicious binaries to collect fine-granularity behavioral artifacts that characterize malware's usage of the file system, memory, network, and registry. Expert labeling and unsupervised clustering produced models that accurately determine malware status and family, providing operationally acceptable precision and recall, and reasonable overhead through benchmarks and cost estimates.

While AMAL provides a high accuracy, it is computationally expensive. Thus, my next system, Chatter, leveraged an elegant representation of select events in the binary execution for fingerprinting. Whereas calculating and exposing low-level features might have ill scalability or gamesmanship effects, this system tersely and efficiently captured execution patterns. By creating an alphabet and a language to represent runtime behavior, techniques from n-gram processing are used to train a classifier capable of distinguishing different binaries types with high accuracy. We also generalize the approach to mobile malware [27].

Researchers rely on outputs of antivirus scanners in establishing ground-truth for their methods, with a lack of research validating scanners performance. We developed AV-Meter, a system for evaluating the performance of antivirus scans and labels. Utilizing malware samples that have been manually labeled by expert analysts, AV-Meter revealed the gap in correctness, coverage, and consistency of popular scanners.

In this line, my other results included techniques for authorship attribution using deep learning networks [1], extending those capabilities to multi-author identification [5] and web contents fingerprinting [23].

Mobile Security Mobile security threats have emerged because of the fast growth in mobile technologies. To address threats associated with mobile software, various techniques are developed in the literature, including ones that utilize static, dynamic, on-device, off-device, and hybrid approaches for identifying, classifying, and defend against mobile threats. Those techniques fail at times, while creating a trade-off of performance and operation. My contribution to bridge the gap in the literature includes the following: 1) *Andro-AutoPsy.* The key insight in the design of this system is that focusing on malware-centric features may not be sufficient, thus *Andro-AutoPsy* employs a new set of malware creator-centric information readily available from the developer profile. Using *Andro-AutoPsy*, we detect and classify malware samples into

similar subgroups by exploiting the behavior profiles extracted from integrated footprints, which are implicitly equivalent to distinct behavior characteristics. *Andro-AutoPsy* provide an accuracy greater than 98%, and was shown to be capable of identifying zero-day mobile malware. 2) *AndroTracker.* A large number of malicious mobile applications are created by a small number of professional underground actors, although previous studies overlooked such information as a detection feature. We exploited this feature in *AndroTracker*, a system that enables fast detection of malware by using creator information such as serial number of certificate. *AndroTracker* further analyzes behaviors associated with permissions use to increase detection accuracy. *AndroTracker* classified malicious mobile apps based on similarity scoring, and achieved detection and classification accuracy of 99% and 90%, respectively.

My work in this space also addressed the security of automated mobile fare collection [22], smartphone authentication [4, 43], inference attacks on GPS-enabled mobile devices [12], among other results.

Internet Security Instructed by adversaries, malware infect infrastructure to launch attacks, requiring addressing. As such, I focused on endpoints, name servers, and DNS entities for their security as infrastructure.

Networked machines serving as binary distribution points, C&C channels, or drop sites are a ubiquitous aspect of malware infrastructure. By sandboxing malware one can extract the network *endpoints* contacted during execution. Some endpoints are benign, while exclusively malicious destinations can serve as signatures, a behavioral distinction drawn by expert analysts, resulting in considerable cost. This line of work leveraged 28,000 expert-labeled endpoints from ≈100k malware binaries for malicious endpoints detection, using endpoints' static metadata and not network payloads or routing dynamics. Performance validates this approach, achieving 99.4% accuracy at binary threat classification and 93% accuracy on the more granular task of severity prediction. This performance is driven by features capturing a domain's behavioral history and registration properties, and highlight the prominent role that dynamic DNS providers and "shared-use" public services play as perpetrators seek agile and cost-effective hosting infrastructure.

Malicious webpages on the Internet are another prevalent threat, and their detection is important. To address this threat, we proposed systems exploiting machine learning over network metadata derived from the sandboxed execution of webpage content. The goal of this research is to detect malicious webpages and identify the type of vulnerability using dynamic network artifacts—and their representations. These features are collected during webpages rendering. Using a real-world dataset that includes different type of malicious behaviors, our results show that dynamic network artifacts can be used effectively to detect most types of vulnerabilities achieving an accuracy reaching 96%. The system was also able to identify vulnerability types with high accuracy achieving an exact match in 91% of the cases.

Another piece of infrastructure is DNS resolvers, where open resolvers, for example, are exploited for malicious activities such as amplification. In this line, we studied the DNS resolution system and its security. In one study, we measured DNS open resolvers longitudinally, highlighting more than 3M resolvers in the IPv4 space (as of 2019), with many deviating from standard behavior and serving malicious contents [**?**]. In another study, we explored intrinsic behaviors of malicious domains through their interaction with name servers, and found a significant number of domains that have frequently switched their name servers. As a result, we developed an identifier called "NS-Switching Footprint" (NSSF) and used it to cluster domains, enabling us to detect those with suspicious behavior. We also designed a model that represents a time series, which could be used to predict the number of name servers that a domain will interact with. Experiments show the efficacy of our model using data capturing all `.com` and `.net` zone change transactions.

These are only examples. Other related work includes infrastructure abuse with typosquatting [44–46], open resolvers [33, 47], and sybils in online gaming platforms [3, 48]

Security Against Denial of Service DDoS attacks are a constant threat and are difficult to defeat because: 1) it is hard to know in advance when an attack is launched, 2) where the attacking machines are from, 3) how many attacking machines are involved, and 4) how long an attack will last. Most Internet DDoS attacks, however, are attributed to larger interconnected and overly complex entities that belong to various botnets, and understanding those botnets can help address the above challenges. To pursue this work, we relied on 50,704 Internet DDoS attacks across the globe, of which data is collected for a seven-month periods from 674 botnet generations, from 23 different botnet families, against 9026 victim IPs in 1074 organizations located in 186 countries by addressing the modeling of source, target, and predictions.

In the first line [8], we analyze overall characteristics of DDoS attacks. Some highlights of this work include: (1) geolocation analysis shows that the geospatial distribution of the attacking sources follows certain patterns, enabling very accurate source prediction of future attacks for most active botnets; (2) from the target perspective, multiple attacks to the same target exhibit strong patterns of inter-attack time interval, allowing accurate start time prediction of the next anticipated attacks; (3) different botnets launch DDoS attacks targeting the same victim, simultaneously or in turn. These findings provided better understanding of today's Internet DDoS attacks, offering new insights for designing new defense schemes at different levels.

As the source of the attacks is various botnets, understanding the sources and their adaptation and evolution in light of detection efforts is important. In this line [42,49], we conducted a measurement study on some

of the most active botnets on the Internet based on the aforementioned data. We first examine and compare the attack capabilities of different families of today's active botnets, highlighting their magnitude, collaborative patterns, spatial-, temporal, and spatio-temporal characteristics. We supplement this line by translating our analytical findings concerning the shifts of the attackers into concrete and generative models [41, 50]. Based on our previous findings that most attacks are not widely distributed, i.e., highly regionalized, we explore the dynamics behind the scenes and find that there are certain shift patterns of each botnet family, indicating strategic attack resources deployment. We generalize those findings into spatial, temporal, and spatiotemporal models, and propose various defenses employing those insights and predictions.

PhD Dissertation and Earlier Work My doctoral work focused on measuring, analyzing, and improving properties for building trustworthy computing on social graphs; e.g., sybil defenses, mixing networks, etc. Using large-scale measurements [6, 51], I proposed improving these properties to build better systems [20]. Moreover, I worked on various topics in systems security and privacy, including results on studying privacy in various domains, including spectrum sharing [19], domain name system [33] named data networks [52], DNSSEC privacy [29], and data mining that includes shortest path calculation [53]. In the broad area of network security, my earlier results included primitives for new and efficient authentication [4, 43] and set operations [54], fair scheduling scheduling [21], secure routing [2], and domain name redirection [7].

2 Future Research Plan

My mid-term research plan for the next five years will be focused on exploring various uncharted research questions in the broad area of systems security and motivated by addressing the security and privacy of various application domains, including Internet of Things (system), blockchain systems (protocol), online communities (data; artifacts), virtual reality and wearable systems (system), and next generation networked systems (system). Employing the various toolsets and techniques developed so far, my research will be enabled by several strategies. 1) The development of new tools to analyze, characterize, and understand protocols, algorithms, and systems, by examining them and artifacts resulting from their usage. 2) Identifying features and variations in the design, deployment, and/or usage of those systems and protocols that will answer questions concerning their vulnerability with respect to various security and privacy requirements. 3) The development of techniques that will ensure those systems' security and privacy in light of the identified vulnerabilities and root cause analysis associated with them. I anticipate my research will enable a better comprehension of the security and privacy of those systems by benefiting from advances in the areas of program/protocol analysis, representation, and learning.

Internet of things security. Motivated by the ad-hoc approaches for improving IoT security, we plan to pursue a systematic and cross-layer approach. By breaking a typical IoT software into application and Internet layers, we plan to build a set of capabilities that will run in parallel with the IoT applications while a) performing various security analyses within a given layer, b) using cross-layer artifacts towards understanding high-level security events for detection, attribution, and defense. Initiated in various recent studies, my approach of understanding IoT through a management layer, rather than rewriting the IoT stack, is for two reasons. First, a new design realizing the principle would take long to implement, assuming its acceptance by the many stakeholders in the IoT ecosystem. Second, such a new design would not address security of legacy IoT devices in deployment today. Moreover, having a deep understanding of how a IoT malware functions is essential before some efforts are made to re-write the stack to achieve security. Our approach can serve both purposes of (1) providing in-depth understanding of IoT malware at different layers that future application development can leverage, and (2) achieving protection capabilities for legacy IoT applications. Encouraged by some preliminary results in this space, our analysis will be comprehensive, employing representation of artifacts obtained from the disassembled IoT binaries across strings, graphs, and endpoints. Our strings representation analysis will trace residual strings, including functions, special words, flags, etc. appearing in the statically analyzed IoT applications and use them for fingerprinting malicious software and understanding the adversary's intent. Graph representations that address obfuscation will further incorporate graphical structures from the code analysis (e.g., control flow and data flow) to attribute code and detect maliciousness behaviors through graph mining. Endpoints and shell codes appearing in the residual strings of the binaries will shed a deeper understanding of the lifecycle of IoT malware, and help devise effective strategies for their containment. All in all, we anticipate the various capabilities to be employed for detecting malicious IoT applications using machine learning systems that we will develop and study for their robustness to changes in the analyze applications and adversaries' evasive behaviors.

Blockchain Systems Security. Current blockchain designs, such as Bitcoin and Ethereum, cannot meet future blockchain applications requirements for they suffer from various security and scalability issues [55]. As blockchain networks scale up, we observed that network synchronization decreases which increases the risk for partitioning attacks [13]. Given several research gaps in this space, we plan to pursue various directions for improving the security of blockchain systems through analysis and design.

This research direction will include developing and advancing the following. 1) Methods for improving the network resilience to asynchrony and partitioning by a systematic root-cause analysis that will look into understanding bottlenecks of latency in today's blockchain systems and appropriate designs to reduce it for a consistent view of blockchain nodes. 2) New communication models that will improve the design assumptions and requirements for blockchain systems, providing a compromise between synchrony and smaller attack surface by addressing temporal variations in blockchain views, and the appropriate evaluations to understand the security of those communication models. 3) Methods for scaling up blockchain systems, especially customizable payment channel networks (PCNs), which alleviate some of the overhead on the main chain by aggregating transactions offline, and their appropriate security analysis. 4) Refined secure consensus algorithms with desirable features, such as low latency, robustness to failure and fairness. 5) Designs that employ advances in secure hardware, e.g., Trusted Execution Environment (TEE), to harden the security of smartphone-based blockchain applications (e.g., blockchain-backed health information systems). This direction is nontrivial, and requires answering the following questions: a) how to overcome the constraints of migrating sensitive applications from rich execution environment to TEE with limited support, and (2) how much security improvement can be achieved through smartphone-based TEE?

Virtual Reality and Wearable Systems Security and Privacy The extensive use of smart and wearable devices has facilitated many useful applications. Many applications can gather, process, and share rich data/metadata, such as geolocation, trajectories, elevation, voice commands, hand movement indicators, time, etc. Fitness applications utilize such information for activity tracking, and have recently witnessed a boom in popularity. Virtual reality (VR) headsets, such as Oculus Rift and Magic Leap 1, allow for an immersive experience, and are facilitated by novel input mechanisms and mediums, such as tapping keyboards projected in the virtual environments. Smartwatches and smartphones are equipped with artificial intelligence (AI) assistants that employ state-of-the-art voice recognition and analysis capabilities using high-quality built-in microphones. While the technologies behind those new devices have enabled many useful applications, our understanding of those technologies' security is quite limited.

Motivated by the popularity of those technologies and a gap in the security and privacy literature, this line of work will look into systematically examining the security and privacy of those wearable technologies, and how this security is impacted by the convergence with other technologies. The key objective of the research is to examine the ample opportunities an every-day adversary can have by employing various advanced learning techniques over low-level observations to launch powerful attacks with significant implications in the application space. Motivated by some of our recent preliminary and promising results that show that an adversary can identify location trajectory on a map by only observing the elevation profile of an activity obtained from a smartwatch [12], we envision a set of side-channel based attacks on those technologies that combine similar observations. Examples of the research questions that we will address in this research line include: 1) Would an adversary engaging with a target in a phone conversation, while the target is typing on her keyboard, be able to recognize what is being typed by exploiting the background sound in the recorded conversation? 2) Would the same adversary be able to detect what is being typed by exploiting recordings by a compromised smartwatch of the target? 3) Would an adversary observing a target typing sensitive login information in the air on a virtual keyboard be able to recognize what is being typed? 4) How can the target defend her input from the adversary without sacrificing the usability of those technologies? The answer to those questions is nontrivial, and we will build on our expertise in signal processing, machine learning, pattern recognition, and defense design to answer those questions.

Next Generation Networked Systems By 2023, there will be around 100 zettabytes of data. The increasing data volumes accelerated the development of processing, storage, I/O devices, and network infrastructure, with 100 Gbps per-port speed and high-end switches capable of processing 25.6 Tbps. For security, traffic measurement is of paramount importance, and the current state-of-the-art employing measurement information in the router's TCAM are impractical. Our recent work highlighted the role of sampling in realizing efficient traffic measurement systems for security by employing advances in memory systems (in memory computing using SRAM) [15, 24, 56]. In this direction, we will pursue research to enable in-data-plane scalable traffic measurement and timely detection of security events (e.g., anomaly, DoS, and DDoS).

The technical aspects we will advance in this work include the following. 1) Sketch-based flow sampling/counting techniques: building on our preliminary results, we will explore scalable fine-grained per-flow measurements enabled by new sampling/counting techniques that employ sketches. 2) Per-flow spectral density distribution measurement systems: To mitigate the per-flow memory usage of the sketch, we will use a computational distance to indicate the volume of flows, and to minimize the memory space assigned for each flow to scale up the flow counting and retention capacities. 3) Per-flow spread distribution measurement systems: spreader detection is essential for anomaly detection, although cannot be computed in the data-plane because of the computational overhead (extensive memory read). With most security applications, the measurement and detection need to be lightweight and online. In this direction, we mainly focus on designing an in-data-plane decodable spreader detector that can be deployed in switches (data-plane

programmable switch) to detect spreader-related anomalies in real-time. 4) Online anomaly detection and mitigation: employing counting using the advances we will make, as described above, we anticipate to also develop online detection techniques of anomalies. For achieving more intelligent mitigation and control of networks, we expect using machine learning techniques to dynamically predict and define action sets.

References

[1] M. Abuhamad, T. AbuHmed, A. Mohaisen, and D. Nyang. Large-scale and language-oblivious code authorship identification. In *ACM Conference on Computer and Communications Security, CCS*, 2018.

[2] M. Schuchard, E. Vasserman, D. Mohaisen, N. Hopper, and Y. Kim. Losing control of the internet: using the data plane to attack the control plane. In *Network and Distributed System Security Symp. (NDSS)*, 2011.

[3] E. Lee, J. Woo, H. Kim, D. Mohaisen, and H. K. Kim. You are a game bot!: Uncovering game bots in mmorpgs via self-similarity in the wild. In *Proceedings of NDSS*, 2016.

[4] Z. Ba, S. Piao, X. Fu, D. Koutsonikolas, D. Mohaisen, and K. Ren. Abc: Enabling smartphone authentication with built-in camera. In *ISOC NDSS*, 2018.

[5] M. Abuhamad, T. AbuHmed, D. Nyang, and D. A. Mohaisen. Multi-χ: Identifying multiple authors from source code files. *Proc. Priv. Enhancing Technol.*, 2020(3):25–41, 2020.

[6] D. Mohaisen, A. Yun, and Y. Kim. Measuring the mixing time of social graphs. In *Proceedings of ACM SIGCOMM Conference on Internet Measurement (IMC)*, pages 383–389. ACM, 2010.

[7] G. Namata, A. West, and D. Mohaisen. Web of redirection: Measuring domain forwarding and applications at internet-scale. In *ACM IMC*, 2014.

[8] A. Wang, D. Mohaisen, W. Chang, and S. Chen. Delving into internet ddos attacks by botnets: Characterization and analysis. In *IEEE International Conf. on Dependable Systems and Networks (DSN)*, 2015.

[9] J. Park, A. Khormali, M. Mohaisen, and D. Mohaisen. Where are you taking me? behavioral analysis of open DNS resolvers. In *IEEE International Conference on Dependable Systems and Networks, DSN*, 2019.

[10] J. Jeon, J. Kim, J. Kim, K. Kim, D. Mohaisen, and J. Kim. Privacy-preserving deep learning for geo-distributed medical big-data platforms. In *IEEE Conf. on Dependable Systems and Networks, DSN*, 2019.

[11] H. Alasmary, D. Mohaisen, and et al. Soteria: Detecting adversarial examples in control flow graph-based malware classifiers. In *IEEE International Conf. on Distributed Computing Systems, ICDCS*, 2020.

[12] U. Meteriz, N. F. Yildiran, J. Kim, and D. Mohaisen. Understanding the risks of sharing elevation information on fitness applications. In *IEEE Int'l Conf. on Distributed Computing Systems, ICDCS*, 2020.

[13] M. Saad, V. Cook, L. Nguyen, M. T. Thai, and D. Mohaisen. Partitioning attacks on bitcoin: Colliding space, time, and logic. In *IEEE International Conf. on Distributed Computing Systems, ICDCS*, Jul 2019.

[14] A. Abusnaina, D. Mohaisen, and others. Adversarial learning attacks on graph-based iot malware detection systems. In *IEEE International Conference on Distributed Computing Systems, ICDCS*, 2019.

[15] R. Jang, S. Moon, Y. Noh, D. Mohaisen, and D. Nyang. Instameasure: Instant per-flow detection using large in-dram working sets. In *IEEE International Conf. on Distributed Computing Systems, ICDCS*, 2019.

[16] S. Ahn, J. Kim, W. Choi, D. Mohaisen, and S. Kang. Shmcaffe: A distributed deep learning platform with shared memory. In *IEEE International Conf. on Distributed Computing Systems, ICDCS*, 2018.

[17] S. Baek, Y. Jung, D. Mohaisen, S. Lee, and D. Nyang. SSD-Insider: Internal defense of SSD against ransomware. In *IEEE Int'l Conf. on Distributed Computing Systems, ICDCS*, 2018.

[18] S. Chen, K. Ren, S. Piao, C. Wang, Q. Wang, L. Su, and D. Mohaisen. Defending against voice impersonation attacks on smartphones. In *IEEE Int'l Conf. on Distributed Computing Systems, ICDCS*, 2017.

[19] C. Guan, D. Mohaisen, L. Su, K. Ren, and Y. Yang. When smart TV meets CRN: privacy-preserving fine-grained spectrum access. In *37th IEEE Int'l Conf. on Distributed Computing Systems, ICDCS*, 2017.

[20] D. Mohaisen, N. Hopper, and Y. Kim. Keep your friends close: Incorporating trust into social network-based sybil defenses. In *30th IEEE International Conf. on Computer Communications (INFOCOM)*, 2011.

[21] T. Zhu, D. Mohaisen, Y. Ping, and D. Towsley. DEOS: dynamic energy-oriented scheduling for sustainable wireless sensor networks. In *IEEE INFOCOM*, 2012.

[22] F. Dang, P. Zhou, Z. Li, E. Zhai, D. Mohaisen, Q. Wen, and M. Li. Large-scale invisible attack on AFC systems with nfc-equipped smartphones. In *IEEE Conf. on Computer Communications, INFOCOM*, 2017.

[23] A. Abusnaina, R. Jang, D. Nyang, and D. A. Mohaisen. DFD: adversarial learning-based approach to defend against website fingerprinting. In *IEEE Conf. on Computer Communications, INFOCOM*, 2020.

[24] R. Jang, D. Min, S. Moon, D. A. Mohaisen, and D. Nyang. Sketchflow: Per-flow systematic sampling using sketch saturation event. In *39th IEEE Conference on Computer Communications, INFOCOM 2020, Toronto, ON, Canada, July 6-9, 2020*, pages 1339–1348, 2020.

[25] R. Jang, J. Kang, D. Mohaisen, and D. Nyang. Catch me if you can: Rogue access point detection using intentional channel interference. *IEEE Trans. Mob. Comput.*, 19(5):1056–1071, 2020.

[26] F. Dang, E. Zhai, Z. Li, D. Mohaisen, K. Bian, Q. Wen, and M. Li. Pricing data tampering in automated fare collection with nfc-equipped smartphones. *IEEE Trans. Mob. Comput.*, 18(5):1159–1173, 2019.

[27] F. Shen, J. D. Vecchio, D. Mohaisen, S. Y. Ko, and L. Ziarek. Android malware detection using complexflows. *IEEE Trans. Mob. Comput.*, 18(6):1231–1245, 2019.
[28] D. Nyang, D. Mohaisen, and J. Kang. Keylogging-resistant visual authentication protocols. *IEEE Trans. Mob. Comput.*, 13(11):2566–2579, 2014.
[29] D. Mohaisen, Z. Gu, K. Ren, Z. Li, and D. Nyang. Look-aside at your own risk: Privacy implications of DNSSEC look-aside validation. *IEEE Trans. Dependable Secur. Comput.*, 17(4):745–759, 2020.
[30] A. Wang, W. Chang, S. Chen, and D. Mohaisen. A data-driven study of ddos attacks and their dynamics. *IEEE Trans. Dependable Secur. Comput.*, 17(3):648–661, 2020.
[31] D. Mohaisen, D. F. Kune, E. Y. Vasserman, M. Kim, and Y. Kim. Secure encounter-based mobile social networks. *IEEE Trans. Dependable Sec. Comput.*, 10(6):380–393, 2013.
[32] A. Wang, W. Chang, S. Chen, and D. Mohaisen. Delving into internet ddos attacks by botnets: Characterization and analysis. *IEEE/ACM Trans. Netw.*, 26(6):2843–2855, 2018.
[33] D. Mohaisen and K. Ren. Leakage of .onion at the DNS root: Measurements, causes, and countermeasures. *IEEE/ACM Trans. Netw.*, 25(5):3059–3072, 2017.
[34] D. Mohaisen and O. Alrawi. Behavior-based automated malware analysis and classification. *Elsevier Computers & Security*, 2015.
[35] D. Mohaisen, A. G. West, A. Mankin, and O. Alrawi. Chatter: Classifying malware families using system event ordering. In *IEEE Conference on Communications and Network Security (CNS)*, 2014.
[36] D. Mohaisen and O. Alrawi. Av-meter: An evaluation of antivirus scans and labels. In *International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment (DIMVA)*, 2014.
[37] A. G. West and D. Mohaisen. Metadata-driven threat classification of network endpoints. In *International Conf. on Detection of Intrusions and Malware, and Vulnerability Assessment (DIMVA)*, 2014.
[38] D. Mohaisen, M. Bhuiyan, and Y. Labrou. Name server switching: Anomaly signatures, usage, clustering, and prediction. In *International Conference on Information Security Applications (WISA)*, 2014.
[39] J. wook Jang, J. Woo, D. Mohaisen, and H.-K. Kim. Andro-dumpsys: Anti-malware system based on malware creator and malware centric information. *Elsevier Computers & Security*, 58:125–138, 2016.
[40] J. Jang, J. Woo, D. Mohaisen, and H. Kim. Andro-autopsy: Anti-malware system based on similarity matching of malware creator-centric information. *Elsevier Digital Investigation Journal*, 2015.
[41] A. Wang, D. Mohaisen, and S. Chen. An adversary-centric behavior modeling of ddos attacks. In *37th IEEE International Conference on Distributed Computing Systems, ICDCS*, 2017.
[42] A. Wang, W. Chang, S. Chen, and D. Mohaisen. A data-driven study of ddos attacks and their dynamics. *IEEE Transactions on Dependable and Security Computing*, 2018.
[43] M. Abuhamad, T. AbuHmed, D. A. Mohaisen, and D. Nyang. Autosen: Deep-learning-based implicit continuous authentication using smartphone sensors. *IEEE Internet Things J.*, 7(6):5008–5020, 2020.
[44] J. Spaulding, S. Upadhyaya, and D. Mohaisen. The landscape of domain name typosquatting: Techniques and countermeasures. In *International Conf. on Availability, Reliability and Security (ARES)*, 2016.
[45] J. Spaulding, S. J. Upadhyaya, and D. Mohaisen. You've been tricked! A user study of the effectiveness of typosquatting techniques. In *IEEE International Conf. on Distributed Computing Systems, ICDCS*, 2017.
[46] J. Spaulding, D. Nyang, and D. Mohaisen. Understanding the effectiveness of typosquatting techniques. In *Proceedings of the fifth ACM/IEEE Workshop on Hot Topics in Web Systems and Technologies, HotWeb*, 2017.
[47] J. Park, J. Choi, D. Nyang, and D. Mohaisen. Transparency in the new gtld era: Evaluating the DNS centralized zone data service. *IEEE Trans. Netw. Serv. Manag.*, 16(4):1782–1796, 2019.
[48] H. Kwon, D. Mohaisen, J. Woo, Y. Kim, E. Lee, and H. K. Kim. Crime scene reconstruction: Online gold farming network analysis. *IEEE Trans. Inf. Forensics Secur.*, 12(3):544–556, 2017.
[49] W. Chang, D. Mohaisen, A. Wang, and S. Chen. Measuring botnets in the wild: Some new trends. In *ACM Symposium on Information, Computer and Communications Security (ASIACCS)*, 2015.
[50] A. Wang, D. Mohaisen, W. Chang, and S. Chen. Revealing ddos attack dynamics behind the scenes. In *International Conf. on Detection of Intrusions and Malware, and Vulnerability Assessment (DIMVA)*, 2015.
[51] D. Mohaisen, H. Tran, N. Hopper, and Y. Kim. On the mixing time of directed social graphs and security implications. In *ACM Symp. on Information, Computer and Communications Security (ASIACCS)*, 2012.
[52] D. Mohaisen, X. Zhang, H. Xie, and Y. Kim. Protecting access privacy of cached contents in information centric networks. In *ACM Symposium on Info., Computer and Communications Security (ASIACCS)*, 2013.
[53] Q. Wang, K. Ren, M. Du, Q. Li, and D. Mohaisen. Secgdb: Graph encryption for exact shortest distance queries with efficient updates. In *Financial Cryptography and Data Security, FC*, 2017.
[54] M. Kim, D. Mohaisen, J. H. Cheon, and Y. Kim. Private over-threshold aggregation protocols over distributed databases. In *Proceedings of IEEE TKDE*, 2016.
[55] M. Saad, J. Spaulding, C. Kamhoua, S. Shetty, D. Nyang, and D. Mohaisen. Exploring the attack surface of blockchain: A comprehensive survey. *IEEE Commun. Surv. Tutorials*, 22(3):1977–2008, 2020.
[56] S. Kim, R. Jang, D. Mohaisen, and D. Nyang. Count-less: A better count-min sketch counting less counters generate bibtex entry for paper. In *IEEE Conf. on Computer Communications, INFOCOM*, 2021.