Research Statement

Aziz Mohaisen

April 28, 2015

My research interests are in the areas of networks security and online privacy. My work has broadly combined principles of the *design*, *analysis*, and *development* of security and privacy primitives and tools for real-world systems. Over the past ten years, my interests have evolved to include topics in *security analytics*, *social networks security and privacy*, *Internet security*, *networks security*, and *privacy*. My approach in conducting research in these areas considered exploratory, constructive, and empirical methods. A common theme in my most recent research work is the use of advanced machine learning techniques for security analytics: to understand codes, traffic, and infrastructure usage in real-world deployments. My earlier work focused on understanding various security issues in multiple networking contexts, by design and analysis. My doctoral thesis work draw from both directions: measurements of social graphs properties used for security, and improved security by better properties in those graphs. In the following, I provide a summary, mostly focused on recent work, and open directions.

1 Security Analytics

Software is part of everything electronic around us, including software running on personal computers, refrigerators, security cameras, security gates, and smartphones. There is the good software, and the "good software made to do bad things" by malicious authors; also known as malware. Understanding families of malware is an essential step for disinfection, risk assessment, and mitigation. With the ever-increasing volumes of infections reported everyday, automation of malware analysis and classification became essential. Malware classification and family identification are not new problems. However, the rapid evolution of the malware attack and defense ecosystem have enabled much fruitful research of analysis systems by capitalizing on a finer understanding of the attack posture of today's adversaries and malware authors. Systems I developed in this line of research include malware analysis and classification for binaries (AMAL [9], Chatter [29], AV-Meter [10]), web infrastructure (ADAM [6], TI [39], and NSSF [11]), and mobile platforms (AndroTracker [4] and Andro-AutoPsy [3]). The following is a summary of those systems and their operational relevance.

1.1 Analytics for Rogue Codes

AMAL: AMAL is a behavior-based tool to classify and cluster malware that utilizes autonomous feature extraction and expert labeled training data. AMAL sandboxes malicious binaries to collect fine-granularity behavioral artifacts that characterize malwares usage of the file system, memory, network, and registry. Expert labeling by analysts and unsupervised clustering enable the production of models that can accurately determine malware status and family (e.g., Zeus). Precision and recall metrics for the technique have been excellent with a great operational relevance, including benchmarks, cost estimates, and other metrics endorsing AMALs approach.

Chatter: While AMAL provides a high accuracy, it comes also at some significant cost. Chatter, on the other hand, is a system for representing and leveraging the sequence of events in a malware execution. Whereas calculating and exposing low-level feature values as in AMAL might have ill scalability or gamesmanship effects, Chatter tersely and efficiently captures execution patterns. By creating an alphabet and a language to represent runtime behavior, techniques from n-gram processing are used to train a binary classifier that is capable of distinguishing different malware samples with high accuracy.

AV-Meter: Both AMAL and Chatter rely on a fine ground truth for their performance evaluation, as well as a large-body of literature. However, researchers relied in the past heavily on outputs of antivirus scanners in establishing ground-truth for their methods and companies use then to guide mitigation and disinfection efforts. However, there is a lack of research that validates the performance of these antivirus vendors. To this end, AV-Meter is a system for evaluating the performance of antivirus scans and labels. Utilizing malware samples that have been manually labeled by expert analysts we reveal dramatic errors in the correctness, coverage, and consistency of current antivirus offerings. We release a call to the community to challenge relying on antivirus scans and labels as a ground truth for malware analysis and classification, by highlighting various risks.

1.2 Analytics for Rogue Infrastructure

My work on building analytics for rogue infrastructure includes ADAM; for automation of malicious webpages classification, TI; for network endpoints appearing in malware execution, and NSSF; for characterizing anomalous domains based on their name server switching patterns. The following provides a description of such work and their operational relevance.

TI: Networked machines serving as binary distribution points, C&C channels, or drop sites are a ubiquitous aspect of malware infrastructure. By sandboxing malcode one can extract the network endpoints (i.e., domains and URL paths) contacted during execution. Some endpoints are benign, e.g., connectivity tests. Exclusively malicious destinations, however, can serve as signatures enabling network alarms. Often these behavioral distinctions are drawn by expert analysts, resulting in considerable cost and labeling latency. Leveraging 28,000 expert-labeled endpoints derived from ≈100k malware binaries this paper characterizes those domains towards prioritizing manual efforts and automatic signature generation. Our analysis focuses on endpoints' static metadata properties and not network payloads or routing dynamics. Performance validates this straightforward 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. More qualitatively we discover the prominent role that dynamic DNS providers and "shared-use" public services play as perpetrators seek agile and cost-effective hosting infrastructure.

ADAM: Malicious webpages are a prevalent threat in the Internet security landscape. Building on this existing literature, we introduced ADAM, a system that uses machine learning over network metadata derived from the sandboxed execution of webpage content. ADAM aims at detecting malicious webpages and identifying the type of vulnerability using simple set of features as well. Machine-trained models are not novel in this problem space. Instead, it is the dynamic network artifacts (and their subsequent feature representations) collected during rendering that are the greatest contribution of this work. Using a real-world operational dataset that includes different type of malice behavior, our results show that dynamic cheap network artifacts can be used effectively to detect most types of vulnerabilities achieving an accuracy reaching 96%. The system was also able to identify the type of a detected vulnerability with high accuracy achieving an exact match in 91% of the cases. We identify the main vulnerabilities that require improvement, and suggest directions to extend this work to practical contexts.

NSSF: There exists a significant number of domain names that have frequently switched their name servers for several reasons. In this work, we delved into the analysis of name-server switching behavior and presented a novel identifier called "NS-Switching Footprint" (NSSF) that can be used to cluster domains, enabling us to detect domains 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. We performed the experiments with the dataset that captured all `.com` and `.net` zone changing transactions (i.e., adding or deleting name servers for domains) from March 28 to June 27, 2013.

1.3 Analytics for Rogue Apps

Mobile security threats have recently emerged because of the fast growth in mobile technologies and the essential role that mobile devices play in our daily lives. For that, and to particularly address threats associated with malware, 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, and succeed at other times, while creating a trade-off of performance and operation. To this end, we contribute two systems: Andro-AutoPsy and AndroTracker.

Andro-AutoPsy: Our first contribution to the mobile security defense posture is Andro-AutoPsy, an anti-malware system based on similarity matching of malware-centric and malware creator-centric information. 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. The experimental results demonstrate that Andro-AutoPsy is scalable, performs well in detecting and classifying malware with an accuracy greater than 98%, and is capable of identifying zero-day mobile malware.

AndroTracker: A large number of malicious mobile applications are created by a small number of professional underground actors, however previous studies overlooked such information as a feature in detecting and classifying malware, and in attributing malware to creators. Guided by this insight, we propose a method to improve on the performance of Android malware detection by incorporating the creators information as a feature and classify malicious applications into similar groups. We developed a system called AndroTracker that implements this method in practice. AndroTracker enables fast detection of malware by using creator information such as serial number of certificate. Additionally, it analyzes malicious behaviors and permissions to increase detection accuracy. AndroTracker also can classify malware based on similarity scoring. Finally, AndroTracker shows detection and classification performance with 99% and 90% accuracy respectively.

1.4 Analytics for Denial of Service Attacks

My final thrust of work on security analytics has been focused on analyzing and understanding distributed denial of service (DDoS) attacks. Enormous efforts have been continuously made from both academia and industry to understand the DDoS attacks and defend against them. With an ever-improving defense posture, the attack strategies are constantly changing as well; making DDoS attacks some of the most severe threats on the Internet. DDoS attacks, by nature, are difficult to defend against 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 (among others). This is particularly true for the DDoS attacks launched for demonstration of capabilities, which are driven by reasons other than the commercial ones.

However, most Internet DDoS attacks are today attributed to larger interconnected and overly complex entities that belong to various botnets. For such botnet-based (commercialized) DDoS attacks, understanding the underlying relationships between various attacks and attackers is fundamental in defending against the attacks. Particularly, are those relationships and efforts totally random? How do the attackers manage their resources? Can we estimate attack origins, sizes, duration, start time, and magnitude based on historical data? If there are some patterns in these attacks, can we learn and utilize them to improve the existing defenses? Apparently, understanding the latest attacking strategies and postures is key to the success of any defense.

To pursue this work, we relied on 50,704 different Internet DDoS attacks across the globe, of which data is collected for a seven-month periods operationally. These attacks were launched by 674 botnet generations from 23 different botnet families with a total of 9026 victim IPs belonging to 1074 organizations that are collectively located in 186 countries.

Target Analytics: In this first direction [37], 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, which enables very accurate source prediction of future attacks for most active botnet families; (2) from the target perspective, multiple attacks to the same target also exhibit strong patterns of inter-attack time interval, allowing accurate start time prediction of the next anticipated attacks from certain botnet families; (3) there is a trend for different botnets to launch DDoS attacks targeting the same victim, simultaneously or in turn. These findings add to the existing literature on the understanding of today's Internet DDoS attacks, and offer new insights for designing new defense schemes at different levels.

Source Analytics: As many of the reported attacks are launched by Botnets, they remain the most powerful attack platform by constantly and continuously adopting new techniques and strategies in the arms race against various detection schemes. Thus, it is essential to understand the latest of the botnets in a timely manner so that the insights can be utilized in developing more efficient defenses. In this direction [1], we conduct a measurement study on some of the most active botnets on the Internet based on the aforementioned data. We first examine and compare the attacking capabilities of different families of today's active botnets, highlighting their magnitude. Most interestingly, our analysis clearly shows that different botnets start to collaborate when launching DDoS attacks.

Modeling: in this direction [38], we present analytical findings concerning the geolocagtion shifts of the attackers, of today's DDoS attacks. Based on our previous findings that most of these DDoS attacks are not widely distributed as the attackers are mostly from the same region, 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 force deployment. Furthermore, the shift patterns could be shared among families indicating potential collaborations

2 Security with Social Networks

There have been many attempts in the past decade to build applications for distributed and peer-to-peer systems that exploit social networks properties. These properties are believed to be hard to alter and to be intrinsic to social networks. For example, social networks are used for building censorship-resistant Internet storage, content sharing and publishing, routing protocols, and Sybil defenses. In each of these applications, social networks are assumed to be well-connected and trusted. For instance, in social networks-based Sybil defenses a high quality of trust reduces the attackers' ability to produce multiple identities, thus, defending against the Sybil attack—caused by nodes with multiple identities in distributed systems. For these Sybil defenses to work, social networks are assumed to be trust possessing and fast-mixing, a formal quality of high-connectivity of these networks. Other applications require other properties, such as betweenness, expansion, well-balance, etc. Despite their importance to these applications, there has been not much efforts spent understanding the quality of these properties in social graphs and how such quality affects the performance of these designs.

To enable trustworthy computing on social networks, my doctoral thesis has been focused on measuring, analyzing, and improving properties used for building applications using these networks. I proceeded to this goal by large-scale *measurements* and *analyses* to understand these properties. Mindful of lessons learned from the measurement, I then proceeded to improving properties and to discovering others that can be easily achieved in variety of social networks to build trustworthy systems.

2.1 Understanding Social Networks for Security Applications

It is widely believed that social networks are "fast-mixing", and many Sybil defenses make crucial use of this property. An experimental verification of this property, and its quality, has not been done before. To address this problem, we explored mathematical tools and used them to measure the mixing time—the time that it takes a random walk on a graph to approach the stationary distribution—of several social graphs, using two techniques [30]. First, we used the second largest eigenvalue modulus, a well-known technique in random walk theory, to bound the mixing time. Second, we sampled initial distributions to compute the random walk length required to achieve probability distributions close to stationary. Our findings show that the mixing time of social graphs is much larger than used in literature, which leads to several striking results. First, designs based on the fast-mixing property utilize a weaker property concerning the average mixing in the social graph (as opposed to the worst mixing [30]), making most theoretical provable guarantees of these systems inaccurate. Second, current security systems based on fast-mixing properties have weaker guarantees and have to be less efficient in order to compensate for the slower mixing graphs. The work presented a breakthrough, and is highly cited.

While many social graphs are directed by nature, many applications are often evaluated on undirected versions of them. We measured the mixing time of several directed graphs and their undirected counterparts and found that directed graphs are slower mixing than undirected ones [28]. Furthermore, we found that evaluation of applications on the undirected graphs always overestimates the security provided by these applications. To understand why some graphs are fast-mixing and why some others are not, we related the mixing time to degeneracy, which captures cohesiveness of the graph. We show that fast-mixing graphs have a larger single core, whereas slow mixing graphs have smaller multiple cores. We build on those observations by designing techniques to improve the mixing time of slow mixing graphs using auxiliary links [12].

2.2 Better Security Applications with Social Networks

Mindful of properties in social networks we explored earlier, we propose several applications and systems on top of social networks utilizing properties that are easy to achieve. Below I describe some of these systems.

SocialCloud. We introduce SocialCloud [26, 27], a computing system in which computing nodes are governed by social ties driven from a trusted social graph. We show that incentives to adopt this paradigm are intuitive and natural, and security guarantees provided by it are solid. We propose metrics for measuring the utility of this computing paradigm, and consider several design trade-offs for its operation. Using real-world social traces, we run an event-driven simulator of SocialCloud, and demonstrate the potential of this paradigm for ordinary users. Interestingly, we find graphs known to perform poorly for Sybil defenses [30] are good candidates for our SocialCloud for their "self load-balancing" features.

MeetUp. MeetUp [19] is an encounter-based social network (ESN). MeetUp addresses several availability, security, and privacy challenges that are unaddressed in prior ESN designs. Mindful of the possible pitfalls of prior work, we construct a flexible framework for a secure ESN, which can be used to construct networks that offer better security, privacy, and availability guarantees. We describe two example constructions derived from this framework, and consider each in terms of the ideal requirements. Some of our new designs meet more system security, reliability, and privacy requirements than previous work. We also evaluate real-world performance of one of our designs by implementing a proof-of-concept iPhone application called MeetUp. Experiments highlight the potential of our system and hint at the deployability of our designs on a large scale.

DynaMix. Existing solutions for anonymous communication on social structures undermine the impact of networks dynamics on anonymity guarantees. We build DynaMix, an anonymous communication system that exploits dynamic structures in social networks [18]. We formally show an intuitive connection between anonymity on dynamic graphs and random walks on weighted graphs in which weights summarize the history of edges and allow for future dynamics to weight adjustment. We showed several measurements of our proposed model on dynamic graphs extracted from real-world social networks and compared it to static structures driven from the same graphs, highlighting potential of our proposed system enriching graph structure and improving quantitative anonymity as both entropy and anonymity sets.

Keep your friends close. Social network-based Sybil defenses do not consider the different amounts of trust in different graphs, though trust is a crucial requirement in these systems. To address this problem, we introduced several theoretical and data-driven designs to tune the performance of Sybil defenses by accounting for differential trust [16]. Surprisingly, we find that the cost of operating Sybil defenses is greater in graphs with high trust than in graphs with low trust values. We discovered that this behavior is due to the community structure in high-trust graphs, requiring higher costs to traverse multiple communities. Furthermore, we showed that our proposed designs to account for trust—while they increase the cost of operating Sybil defenses on graphs with low trust value—greatly decrease the advantage of attacker.

3 Other Topics and Early Work

While the two research thrusts above represent two core areas of research competencies, I also developed interests in various aspects related to security, privacy, and measurements over the past ten years as detailed in the following.

Privacy enhancing technologies. My work in this direction consisted of studying privacy in various domains, including location-based services [15], domain name system [21] privacy in named data networks [31, 22], .onion leakage [36] and data mining [13, 2, 17]. For example, my work on data mining focused on the design of privacy-preserving data clustering [13, 2] and association rule mining algorithms [17] that are both efficient and robust against known attacks. I designed an algorithms that defend against known attacks, and maintain the same computational advantages of the previously known algorithms at small loss of data utility [13, 2]. I developed algorithms for discovering relations between data variables while maintaining data privacy [14, 17]. At the theoretical end, we formalized an average-case privacy notion of individual data records. Our algorithms improved existing works and struck a balance between privacy and resources, by utilizing randomization techniques already known in the literature.

Networks security. Another of interest area has been networks security in general. In this vein of research, I worked on building primitives for efficient cryptographic key distribution [23], revocation [25], management [20], authentication [33], and set operations [5] in sensor networks. Furthermore, I researched security analysis of authentication and key distribution designs [24]. I also designed efficient scheduling algorithms in sensor [40] and other opportunistic networks [7]. I finally contributed to the design of usable authentication protocols using augmented reality [34] and visual cryptography [8], as well as studying routing instability in BGP [35] and the use of domain name redirection and applications [32], among other topics.

4 Future Directions

A substantial short to mid term research goal is seen in continuing this work. In particular, I intended to continue working on extending the two core areas of competencies: security analytics and social networks analysis for security applications. First, I believe that there is a rewarding research direction in investigating scalability, better accuracy, and alternatives with possible operational benefits to the proposed techniques for security analytics. Second, extending the studies of the DDoS analytics to possibly realize defenses utilizing the insight into the attack posture would also be a great goal. Finally, I would like to utilize insights into social network properties for new security primitives, including limiting abuse and misuse in distributed systems when coupled with social networks.

In the longer term, I would like to build a research agenda around utilizing analytics outcomes for efficient information sharing: labels provided on malware samples, threat indicators, mobile apps, etc., can be best utilized as actionable indicators for other usage. Realizing platforms for representation of such outcomes, and methods for their transportation while guarding the context of the original indicators, including privacy, is a clearly rewarding long-term goal.

References

[1] W. Chang, A. Mohaisen, A. Wang, and S. Chen. Measuring botnets in the wild: Some new trends. In *Proceedings of the 10th ACM Symposium on Information, Computer and Communications Security*, ASIA CCS '15, pages 645–650. ACM, 2015.

[2] D. Hong and A. Mohaisen. Augmented rotation-based transformation for privacy-preserving data clustering. *ETRI Journal*, 31(3):351–361, June 2010.

[3] J.-w. Jang, H. Kang, J. Woo, A. Mohaisen, and H. K. Kim. Andro-autopsy: Anti-malware system based on similarity matching of malware and malware creator-centric information. *Elsevier Digital Investigation Journal*, 2015.

[4] H. J. Kang, J.-w. Jang, A. Mohaisen, and H. K. Kim. Androtracker: Creator information based android malware classification system. In *Information Security Applications - 15th International Workshop, WISA*, volume 8909, 2014.

[5] M. Kim, A. Mohaisen, J. H. Cheon, and Y. Kim. Private over-threshold aggregation protocols. In *Information Security and Cryptology - ICISC 2012 - 15th International Conference*, LNCS, pages 472–486. Springer, 2012.

[6] A. E. Kosba, A. Mohaisen, A. G. West, and T. Tonn. ADAM: automated detection and attribution of malicious webpages. In *IEEE Conference on Communications and Network Security, CNS 2013, National Harbor, MD, USA, October 14-16, 2013*, pages 399–400, 2013.

[7] Y. Li, A. Mohaisen, and Z. Zhang. Trading optimality for scalability in large-scale opportunistic routing. *IEEE T. Vehicular Technology*, 62(5):2253–2263, 2013.

[8] Y. Maeng, A. Mohaisen, M. Lee, and D. Nyang. Transaction authentication using complementary colors. *Computers & Security*, 48:167–181, 2015.

[9] A. Mohaisen and O. Alrawi. AMAL: high-fidelity, behavior-based automated malware analysis and classification. In *Information Security Applications - 15th International Workshop, WISA 2014. Revised Selected Papers*, volume 8909 of *LNCS*, pages 107–121. Springer, 2014.

[10] A. Mohaisen and O. Alrawi. Av-meter: An evaluation of antivirus scans and labels. In *Detection of Intrusions and Malware, and Vulnerability Assessment - 11th International Conference, DIMVA 2014, Egham, UK, July 10-11, 2014. Proceedings*, volume 8550 of *LNCS*, pages 112–131. Springer, 2014.

[11] A. Mohaisen, M. Bhuiyan, and Y. Labrou. Name server switching: Anomaly signatures, usage, clustering, and prediction. In *Information Security Applications - 15th International Workshop, WISA 2014. Revised Selected Papers*, volume 8909 of *LNCS*, pages 202–215. Springer, 2014.

[12] A. Mohaisen and S. Hollenbeck. Improving social network-based sybil defenses by rewiring and augmenting social graphs. In *Information Security Applications - 14th International Workshop, WISA 2013, Revised Selected Papers*, volume 8267, pages 65–80. Springer, 2013.

[13] A. Mohaisen and D. Hong. Mitigating the ica attack against rotation based transformation for privacy preserving clustering. *ETRI Journal*, 30(6):868–870, December 2008.

[14] A. Mohaisen and D. Hong. Privacy preserving association rule mining revisited. In *The 9th International Workshop on Information Security Applications (WISA 2008)*, pages 1–16, September 2008.

[15] A. Mohaisen, D. Hong, and D. Nyang. Privacy in location based services: Primitives toward the solution. In *NCM 2008, The Fourth International Conference on Networked Computing and Advanced Information Management*, pages 572–579. IEEE Computer Society, 2008.

[16] A. Mohaisen, N. Hopper, and Y. Kim. Keep your friends close: Incorporating trust into social network-based sybil defenses. In *INFOCOM 2011. 30th IEEE International Conference on Computer Communications*, pages 1943–1951. IEEE Computer Society, 2011.

[17] A. Mohaisen, N.-S. Jho, D. Hong, and D. Nyang. Privacy preserving association rule mining revisited: Privacy enhancement and resources efficiency. *IEICE Transactions*, 93-D(2):315–325, February 2010.

[18] A. Mohaisen and Y. Kim. Dynamix: anonymity on dynamic social structures. In *8th ACM Symposium on Information, Computer and Communications Security, ASIA CCS '13*, pages 167–172, 2013.

[19] A. Mohaisen, D. F. Kune, E. Y. Vasserman, M. Kim, and Y. Kim. Secure encounter-based mobile social networks: Requirements, designs, and tradeoffs. *IEEE Trans. Dependable Sec. Comput.*, 10(6):380–393, 2013.

[20] A. Mohaisen, Y. Maeng, J. Kang, D. Nyang, K. Lee, D. Hong, and J. Han. Protection techniques of secret information in non-tamper proof devices of smart home network. In *International Conference on Ubiquitous Intelligence and Computing*, LNCS, pages 548–562. Springer, 2008.

[21] A. Mohaisen and A. Mankin. Evaluation of privacy for dns private exchange. IETF Internet Draft, `https://tools.ietf.org/html/draft-am-dprive-eval-00`, March 2015.

[22] A. Mohaisen, H. Mekky, X. Zhang, H. Xie, and Y. Kim. Timing attacks on access privacy in information centric networks and countermeasures. *IEEE Trans. Dependable Sec. Comput.*, 2015.

[23] A. Mohaisen and D. Nyang. Hierarchical grid-based pairwise key predistribution scheme for wireless sensor networks. In *European Conference on Wireless Sensor Networks (EWSN)*, volume 3868 of *LNCS*, pages 83–98. Springer, 2006.

[24] A. Mohaisen and D. Nyang. On the inefficiency of the resources optimal key pre-distribution scheme for wireless sensor network. *Journal of Communications*, 5(2):164–168, March 2010.

[25] A. Mohaisen, D. Nyang, Y. Maeng, and K. Lee. Structures for communication-efficient public key revocation in ubiquitous sensor network. In *MSN*, volume 4864 of *LNCS*, pages 822–833. Springer, 2007.

[26] A. Mohaisen, H. Tran, A. Chandra, and Y. Kim. Trustworthy distributed computing on social networks. In *8th ACM Symposium on Information, Computer and Communications Security, ASIA CCS '13*, pages 155–160, 2013.

[27] A. Mohaisen, H. Tran, A. Chandra, and Y. Kim. Trustworthy distributed computing on social networks. *IEEE T. Services Computing*, 7(3):333–345, 2014.

[28] A. Mohaisen, H. Tran, N. Hopper, and Y. Kim. On the mixing time of directed social graphs and security implications. In *7th ACM Symposium on Information, Compuer and Communications Security, ASIACCS '12*, pages 36–37, 2012.

[29] A. 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*, pages 283–291, 2014.

[30] A. 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.

[31] A. Mohaisen, X. Zhang, M. Schuchard, H. Xie, and Y. Kim. Protecting access privacy of cached contents in information centric networks. In *8th ACM Symposium on Information, Computer and Communications Security, ASIA CCS '13*, pages 173–178, 2013.

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

[33] D. Nyang and A. Mohaisen. Cooperative public key authentication protocol in wireless sensor network. In *International Conference on Ubiquitous Intelligence and Computing*, volume 4159 of *LNCS*, pages 864–873. Springer, 2006.

[34] D. Nyang, A. Mohaisen, and J. Kang. Keylogging-resistant visual authentication protocols. *IEEE Trans. Mob. Comput.*, 13(11):2566–2579, 2014.

[35] M. Schuchard, E. Y. Vasserman, A. Mohaisen, D. F. Kune, N. Hopper, and Y. Kim. Losing control of the internet: using the data plane to attack the control plane. In *In NDSS'11: the 18th Annual Network and Distributed System Security Symposium*, pages 1–18. ISOC, February 2011.

[36] M. Thomas and A. Mohaisen. Measuring the leakage of onion at the root: A measurement of tor's .onion pseudo-tld in the global domain name system. In *Proceedings of the 13th Workshop on Privacy in the Electronic Society, WPES 2014, Scottsdale, AZ, USA, November 3, 2014*, pages 173–180, 2014.

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

[38] A. Wang, A. Mohaisen, W. Chang, and S. Chen. Revealing ddos attack dynamics behind the scenes. In *DIMVA*, 2015.

[39] A. G. West and A. Mohaisen. Metadata-driven threat classification of network endpoints appearing in malware. In *Detection of Intrusions and Malware, and Vulnerability Assessment - 11th International Conference, DIMVA 2014, Egham, UK, July 10-11, 2014. Proceedings*, volume 8550, pages 152–171, 2014.

[40] T. Zhu, A. Mohaisen, Y. Ping, and D. Towsley. DEOS: dynamic energy-oriented scheduling for sustainable wireless sensor networks. In *Proceedings of the IEEE INFOCOM 2012*, pages 2363–2371, 2012.