A Conversation with Xiaokui Shu: The pursuit of speed in cybersecurity

INNOVATION LEADERS

by Bushra Anjum

Editor's Introduction

In this interview, Bushra Anjum, senior editor Ubiquity, sits down with Xiaokui Shu, a Research Staff Member at IBM, and they chat about the latest emerging issues in cybersecurity. Is it possible to discover modern cyber threats before attackers accomplish their goals? The discussion then turns to the new era of dynamic cyber defense, which consists of detection strategies and procedures developed on the fly by security analysts based on live observations.

A Conversation with Xiaokui Shu: The pursuit of speed in cybersecurity

INNOVATION LEADERS
by Bushra Anjum

Xiaokui Shu is a Research Staff Member at IBM and a member of the ACM Future of Computing Academy. Starting from designing penetration tests and creating risk assessments in college, Dr. Shu has been researching in cybersecurity to identify, formalize, and develop creative solutions to cutting-edge security problems. His research has been highlighted in ACM News and featured in *Communications of the ACM*, and his paper is among the most downloaded publications recognized by the IEEE Signal Processing Society. At IBM Research, Dr. Shu leads the cyber reasoning initiative and works with researchers, engineers, and faculty members to advance cyber defense. Before joining IBM, Dr. Shu received his Ph.D. degree in computer science at Virginia Tech with the Outstanding Ph.D. Student Award. His homepage is at <https://xshu.net>.

What is your big concern about the future of computing to which you are dedicating yourself?

Speed in cybersecurity. More specifically, can we discover modern cyber threats before attackers accomplish their goals or in advance of the threat impact escalation?

In the traditional cyber cat-and-mouse game between attackers and defenders, the speed used to refer to the *execution* of attacks and defenses, e.g., propagation and detection of a virus. A concrete example: if a virus is detected and quarantined before it infects other devices, the defenders win. Besides execution, the development of strategy, technology, and executables on both attacker and defender sides used to occur offline and the speed of the development phases were not critical. That is to say, attackers were free to use a year creating ransomware, and security vendors took months to years to build anti-virus and firewalls. The landscape shifted when a group of advanced attackers realized development-before-execution blinds them to possible attacks since they have a limited view of the attack target in the development phase. So they started to bring development online and to interleave attack development and execution. For instance, before developing malware to collect sensitive data from an enterprise, attackers may first develop and execute early attack phases to gain access to the target network and collect information about the environment and vulnerabilities, a.k.a., phase reconnaissance. The attackers may then develop and execute procedures moving through the target network to locate valuable assets and clear accessing obstacles, a.k.a., phase lateral movement. After these phases, they finally start malware development, delivery, and execution.

Let's assume a piece of code does not modify itself, then a compiled piece of attack or defense code contains a limited predefined set of knowledge such as attack tactics from the attackers' code or the definition of malicious behaviors from the defenders' code. When attackers extend the limited set of knowledge by developing attack procedures during cyberattack campaigns, the old cyber game loses its balance—no defense program can know new attack tactics created by live human attackers on the fly. For instance, after identifying the anti-virus software in the attack target, attackers develop and test malware to make it undetectable by this anti-virus before launching it at the target. To discover the threats, defenders need to understand security measurement limitations and gaps for each protected environment. Moreover, without detection guarantee of each attack steps by existing detectors, defenders need some new capabilities: first, to obtain hints of attacks from various perspectives, e.g., digging into process activities besides malware signatures to discovery unknown malware; second, to reveal the root cause of security incidents, e.g., backtracking unknown zero-day exploits that allow malware to get in; third, to evaluate covered impacts besides confirmed attack activities, e.g., checking whether the explicit ransom is the ultimate goal or just a deceptive goal to stop investigation; fourth, to verify possible threat hypotheses, e.g., finding clues and proofs of the real attack goal of long-term customer privacy breach besides the explicit data ransom. The dynamics here is: attackers don't know what entire attack campaigns look like at the beginning since they are building campaigns on the fly. Nor can defenders know what such modern threats are and how to detect them ahead of time.

Now we enter a new era of cybersecurity with dynamic cyber defense against modern cyber threats. Similar to what attackers do, the dynamic cyber defense enables defenders to develop detection strategies and procedures on the fly based on what they observe. While traditional static security defense requires a nearly complete understanding of a threat before defense development, dynamic cyber defense creates new knowledge of threat for each protected environment during detection. Before we dive into more details, you may already catch the concern—now the speed in cybersecurity is no longer just the execution of attacks or defenses. It also involves development. For defenders, the development includes not only programming, but also threat model creation, requirement analysis, algorithm design, data normalization, and system assembling even with pre-developed detection modules. This is why it takes an average of 279 days to discover a data breach today [1]. Even worse, this number is increasing steadily over the past years, which indicates a growing speed advantage of attackers in the cat-and-mouse game. If this trend continues, the loss of security will impede the use and development of computing in the future.

How did you first become interested in cybersecurity, especially the traditional cyber cat-and-mouse game between attackers and defenders?

Cybersecurity brought me a lot of fun during my teenage years. After cleaning a virus from my friend's PC in the early 2000s, I started studying mechanisms behind malware and created

countermeasures against remote control and system locks in middle school. Dynamically developed attacks beyond static malware is always an interesting possibility in my mind. I believe its detection requires a broad spectrum of knowledge, which motivates me to study various aspects of computer science and mathematics from operating system to abstract algebra. I also try to understand the feasibility of such attacks by developing deep penetration tests in college and writing risk assessments to summarize discovered security gaps and mitigation suggestions.

In recent years, cyberattacks evolve quickly and advanced persistent threat (APT) becomes a major concern, which are the modern cyber campaigns I described in the first question. During my Ph.D., I studied one well-known modern threat, the Target data breach. And I advanced the detection of modern threats at different stages from both theoretical and practical aspects. From the scientific perspective, I proposed a formal method to explain program anomaly detection and shed light on future detection model development. From the engineering perspective, I designed and implemented multiple detection approaches seeking program anomalies, data leaks, suspicious user behaviors, etc.

A few years ago, DARPA started a series of cybersecurity programs to understand both attackers and defenders in the cyber world. After graduating and joining IBM, I worked on one of the DARPA programs named Transparent Computing to recognize the limits of cyber defense and develop new methodologies to break the limits. Building a multi-angle detection platform and orchestrating the MARPLE team discovering a series of red team APT campaigns in four years [2], I gained a clear understanding of the modern cyber game, especially the dynamics that attackers bring to the table, which largely weakens existing defenses. Dynamic cyber defense is a path to the future. However, we need to solve the speed issue pointed out in the first question—it is challenging to accelerate dynamic cyber defense, especially knowledge expression, composition, reuse, and application. And it is challenging to foster creative artificial intelligence to play the game. In summary, the cybersecurity industry has a long way to go towards efficient dynamic cyber defense using existing human knowledge and new knowledge learned and organized by machines.

Please share about some of the initiatives you are a part of that aim to accelerate the dynamic cyber defense.

Why is it difficult to speed up dynamic cyber defense? Because dynamic cyber defense is as challenging as scientific discovery. Many scientific discoveries start from interesting observations, then hypothesis development, experiment design, predication verification, and finally theory establishment. In dynamic cyber defense, security analysts start from suspicious observations such as repeated larger-than-expected network traffic. Then they make threat hypotheses such as a data exfiltration. Next they verify detailed adversary tactics against more observations, and revise the hypotheses along with new observations. The aha moment comes

when a threat hypothesis is verified. To complete the defenses, specific cyber responses are then proposed, evaluated, and executed against each attack.

New to the security industry, dynamic cyber defense has neither systematic methodology nor proper tooling. Existing best practices rely on ad-hoc threat hunting with incomplete and incompatible pieces of threat intelligence and knowledge. The practices are working but slow, and they are far from the full strength of dynamic cyber defense.

At IBM Research, I am currently leading the cyber reasoning initiative that focuses on the methodology and tooling development for dynamic cyber defense. I propose a new methodology named threat intelligence computing [3] to formalize the reasoning aspect of dynamic cyber defense as a graph computation problem. The new methodology provides both human and machine with one unified way to quickly encode knowledge and verify threat hypotheses against system and network data. It enables easy cyber knowledge composition of large attack plots from small steps and tactics, techniques, and procedures (TTP). In addition, proprietary knowledge from existing detection and reasoning systems can be trivially embedded via labels in our approach. We build a proof of concept system in the DARPA Transparent Computing program for threat hunting, for automatic TTP detection, and for comprehensive policy reasoning, which turns out to be a great success [4]. The next challenge is to apply the methodology and technology outside the DARPA environment, and we are working on new platforms and new industry standards to bring security vendors and threat hunters together towards practical dynamic cyber defense. Once accomplished, it will rebalance the game with orders of magnitude shorter time to discover modern threats. It will also unlock further opportunities for defense design automation and security artificial intelligence which will take the speed of cyber defense to the next level.

Dynamic cyber defense is an emerging approach. It connects multiple sub-disciplines in cybersecurity and beyond such as system security, computer language, formal methods, machine learning, etc. And it opens lots of opportunities with big data and AI. Feel free to reach me if you are interested, want to join our effort, think about future industry standards, or plan for research collaborations. Thanks ACM for the great questions.

References

- [1] IBM Security and Ponemon Institute. [2019 Cost of a Data Breach Report](#). 2019. IBM.
- [2] Rao, J. R. and Shu, X. [Unleashing cyber reasoning potential in the era of AI security](#). IBM Research Blog. May 1, 2020 IBM, Armonk, NY, USA.
- [3] Shu, X., Araujo, F., Schales, D. L., Stoecklin, M., Ph., Jang, J., Huang, H., and Rao, J. R. Threat intelligence computing. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM, New York, 2018, 1883–1898.

[4] Shu, X. Unleashing Cyber Reasoning: DARPA Transparent Computing Threat Hunting Retrospective. Sponsored talk at the 2020 Annual Computer Security Applications Conference (ACSAC). 2020.

Biography

Bushra Anjum is a software technical lead at Amazon in San Luis Obispo, CA. She has expertise in Agile Software Development for large scale distributed services with special emphasis on scalability and fault tolerance. Originally a Fulbright scholar from Pakistan, Dr. Anjum has international teaching and mentoring experience and has served in academia for over five years before joining the industry. In 2016, she has been selected as an inaugural member of the [ACM Future of Computing Academy](#), a new initiative created by ACM to support and foster the next generation of computing professionals. Dr. Anjum is a keen enthusiast of promoting diversity in the STEM fields and is a mentor and a regular speaker for such. She received her Ph.D. in computer science at the North Carolina State University (NCSU) in 2012 for her doctoral thesis on Bandwidth Allocation under End-to-End Percentile Delay Bounds. She can be found on Twitter [@DrBushraAnjum](#).

DOI: 10.1145/3457416