September 2021

Robot Hacking Games

China's Competitions to Automate the Software Vulnerability Lifecycle

CSET Issue Brief

AUTHOR
Dakota Cary

Executive Summary

Robot Hacking Games (机器人网络安全大赛, RHG) are government-backed competitions that China uses to advance automatic software vulnerability discovery, patching, and exploitation technologies.[1] These tools offer both offensive and defensive capabilities that promise to increase the scale and pace of vulnerability discovery. If successful, countries could use these tools to find software vulnerabilities quicker than their adversaries. A fully developed capability would allow defenders to patch vulnerabilities as quickly as they are found; attackers could build new exploits equally fast. The Defense Advanced Research Project Agency's Cyber Grand Challenge in 2016 spurred China's interest in this area. The DARPA effort resulted in the creation of state-of-the-art tools in each of these areas, which have since been siloed into separate programs. China, by contrast, has hosted at least seven competitions since 2017.

China's competition structure embodies its military-civil fusion strategy, attracting a collection of academic, military, and private-sector teams. Just two years after the People's Liberation Army's National University of Defense Technology won the first competition in 2017, the military started managing competitions of its own.[2] By 2021, a laboratory run by the PLA Equipment Development Department hosted its first RHG competition.[3] These management and oversight roles situate the PLA in an ideal position to evaluate and attract the best tools and talent. Other state hacking teams, like those of the Ministry of State Security (MSS), will benefit from the technology's development, too.

Leading Chinese cybersecurity experts and government strategy documents tie automated software vulnerability discovery, patching, and exploitation tools to Chinese President Xi Jinping's goal for China to become a "cyber powerhouse" (网络强国).[4] These policy documents create a de facto political mandate for China's cybersecurity community to develop the desired tools. Although they will not make China a "cyber powerhouse" on their own, their development illustrates one important capability that China has chosen in pursuit of its goal.

Introduction

A collection of seven server racks, each with their team's color and logo splashed on the covers of the bulky boxes, stood on stage in a Las Vegas conference room in August 2016.[5] Professional commentators narrated as each team's Cyber Reasoning System (CRS) hacked away on their own code, and their competitors' servers, trying to find vulnerabilities. Over the course of the competition, these machines earned points by patching their own vulnerabilities while maintaining system performance and submitting successful attacks against opposing teams' servers.[6] But the event lacked the chatter of fingers furiously striking keys that normally accompanies hacking competitions. A few feet away from the flowing bits and bytes, a collection of PhDs, researchers, and private-sector innovators who created the CRSs watched the scoreboard update after every five-minute round. Like coaches at a swim meet, all they could do was sit back and watch.

Figure 1. DARPA's CGC in Las Vegas[7]

Source: DARPA

DARPA hoped to show that software vulnerability discovery, patching, and exploitation could be automated. Together, these three phases constitute the "vulnerability lifecycle." Once a software vulnerability is found, what happens next depends on who found it. Attackers exploit those vulnerabilities, allowing them to access protected systems. Defenders patch those same vulnerabilities to prevent compromise. Both offense and defense

want to automate software vulnerability discovery, a well-developed field of research consisting of corporate developers and cybersecurity experts using tools to find software flaws. Automated patching and exploitation are relatively less-developed and not as widely used. DARPA's CGC and China's RHGs lump these three distinct phases together because they rely on similar technical processes and techniques. This paper refers to these capabilities as "tools to automate the vulnerabilities lifecycle," or AVL tools.

Currently, software vulnerability discovery, exploitation, and patching can be labor-intensive.[8] Software developers, often with years of experience, must pore over code looking for ways it can break. Even with existing tools and techniques, such as symbolic execution, it is impossible to consider all possible avenues of failure. Mistakes leave behind vulnerabilities that attackers may exploit. Open source fuzzing tools, such as American Fuzzy Lop, help researchers locate cracks in their code by generating inputs to cause software crashes. But the time dedicated to this process during software development is constrained by economics. Corporate requirements and shareholder value dictate the amount of time spent securing products, often resulting in insufficient attention. High labor costs and slow product development dent profits. AVL tools would pay huge dividends to companies and governments able to deploy the technology.

The Cyber Grand Challenge provided a glimpse of the future by automatically identifying vulnerabilities, building and applying patches, and exploiting vulnerable programs. Although the event targeted relatively simple software compared to more widely-used programs, it demonstrated that AVL tools are viable. The day after CGC's machine-only event concluded, another important event unfolded. DEF CON, a conference for hackers that was co-located with DARPA's event, invited CGC's winning team to enter their system in a capture-the-flag game against DEF CON's human finalists.[9] ForAllSecure, the CGC's winning team from Carnegie Mellon University, agreed to submit their CRS, Mayhem. In the end, Mayhem lost to all of the competition's 14 human teams.[10] But it was not a resounding defeat. In the first 10 hours of the CTF, Mayhem led some of the human teams. In the hours that followed,

the teams overtook the machine. By the end of the CTF, humans remained undefeated in hacking competitions.

The events in Las Vegas set off a firestorm of articles touting the impact automation would have on cybersecurity.[11] The articles were optimistic and half right. ForAllSecure eventually received a contract to deploy their CRS on DOD systems.[12] Other competitors sold their systems to cybersecurity firms.[13] But AVL tools are still only deployed piecemeal—as specialized vulnerability discovery tools, not as fully developed vulnerability lifecycle products. On this front, the article's predictions were wrong. The technology is not trustworthy enough to automatically patch software, and most exploit generation requires a hands-on approach. Still, the competition and its results were so consequential that the Smithsonian National Museum of American History displayed Mayhem in an exhibition on innovation in defense technology.[14]

CGC changed some fundamental assumptions about software and security that underpin the cyber domain. All the hard work spent engineering and fine-tuning the CRSs to do a human's job was sure to go somewhere impressive.

But as far as is publicly known, the technology has not. DARPA never planned a second CGC. The agency pushed its research on automated exploitation and automated patching into siloed DARPA programs, reducing the public incentive to assemble such systems while simultaneously removing the technology's focal point for the cybersecurity community. The grand challenge model used for CGC isn't intended to support annual competitions, but rather tries to spur innovation and evaluate the best technology for a particular field at a single point in time. For Chinese Communist Party (CCP) policymakers, CGC did just that.

China's Robot Hacking Games (机器人网络安全大赛)

China's cybersecurity policymakers began monitoring the development of DARPA's CGC when it was first announced in 2014. Chinese policy publications and industry magazines hyped up the importance of the competition for cybersecurity.[15] That same year, Xi signaled that the party wanted China to become a

"cyber powerhouse" (网络强国), an intentionally vague term meant to inspire.[16] Following Xi's announcement, CCP policymakers began releasing strategy documents to define what the political objective meant in technical terms.

Xu Guibao (徐贵宝) was one of those policymakers. Xu served as a senior manager at a government think tank under China's Ministry of Industry and Information Technology.[17] Throughout his time in government service, he authored numerous policy documents related to China's 13th Five-Year Plan and received 10 patents for his technical innovations. When policymakers needed someone to serve as lead author for China's 2015 *"Internet + Artificial Intelligence Three-Year Action Plan,"* to support China's AI-related technology development, he was a perfect fit.

In the span of a single year, Xu witnessed three events that shaped his perspective on the technologies China needed to achieve its "cyber powerhouse" ambitions. The first event drew global attention. In early 2016, DeepMind's AlphaGo beat Lee Sedol, a world-champion Go player, in four of five games.[18] The event concentrated minds around the world on the potential impact of machine learning. A few months later in late 2016, DARPA's CGC concluded with a human versus machine competition, where ForAllSecure's Mayhem led two of the fourteen human teams from DEF CON before ultimately losing.[19] For Xu, Mayhem's short-lived lead over humans reminded him of Lee Sedol's loss. The influential academic referred to Mayhem's performance against those teams as the "AlphaGo incident in the field of cybersecurity"—a bit hyperbolic, but indicative of his thinking at the time.[20] As winter turned to spring, Xu watched as the WannaCry ransomware tore its way across networks and pillaged computers around the world, including networks in China. WannaCry shook Chinese policymakers. For a regime that prizes stability and the government's ability to solve society's problems, the WannaCry incident concentrated minds on the powerful and uncontrollable effects malware could have. The trend in cybersecurity and AI was clear to Xu—automation promised scale and capabilities that could best humans.

Xu published an influential article in 2017 titled “U.S. Intelligent Cyberattack and Cyber Defense: Inspiration for China's Cyber Powerhouse Strategy.”[21] The title alone was a clear indication Xu thought AVL tools should be one part of China’s journey to becoming a “cyber powerhouse.” Xu argued that China must “accelerate the development of a networked system of autonomous repair and offensive and defensive robots” to achieve its cyber powerhouse strategy.[22] (These so-called robots are the cyber reasoning systems tested at DARPA’s CGC and Xu’s inspiration). The CCP channeled Xu’s recommendations in its *New Generation Artificial Intelligence Development Plan*, released around the same time in 2017, generically stating China must “strengthen AI cybersecurity technology research and development.”[23] Chinese policymakers would later expressly echo Xu’s recommendation. But by the time China’s first Robot Hacking Game competitors filed into a Wuhan conference room in late 2017, it was already becoming clear that China viewed AVL tools as important to becoming a “cyber powerhouse.”[24]

Figure 2: China’s First RHG in 2017[25]

Source: Zhejiang University

The International Robot Hacking Game (国际机器人网络安全大赛, RHG) attempted to recreate DARPA’s CGC.[26] Teams even decorated their server racks with the same combination of lights, colored trim, and logos. Despite differences in scoring and

structure, China's first RHG tested the same types of technologies as the CGC: automated vulnerability discovery, patching, and exploitation.[27] In the same way that DARPA oversaw the CGC, the Ministry of State Security's 13th Bureau, the Central Cyberspace Administration of China, and the Ministry of Education supervised the International RHG competition.[28] Despite including "international" in its name, China's first RHG attracted only three of the competition's 22 teams from abroad—one was a CGC finalist.[29] In a humorous twist, a source familiar with the competition claimed one Chinese team just copied open-source code published by a CGC participant and hoped for the best.[30] They didn't win. In the end, China's National University of Defense Technology, a PLA military academy, beat the CGC finalist and other entrants to win the competition.[31] The "international" component of RHG competitions has since been dropped.

Chinese policymakers saw what they needed to see. In the months following its first RHG, China doubled down on Xu's recommendations in its "*Internet +" Artificial Intelligence Three-Year Action Plan* covering the 2018 to 2020 time period.[32] The plan stated that "in order to solve the security technology problems such as vulnerability discovery, security testing, threat warning, attack detection, and emergency response, enterprises should promote the advanced application of advanced AI technology in the field of cybersecurity."[33] By 2018, AVL tools had solidified their place in China's technology development strategies.

The subsequent promulgation and standardization of RHG competitions was swift. Including preliminary rounds, China has hosted at least a dozen competitions for AVL technology since DARPA's CGC in 2016.[34] As an indication of its now prominent role, a 2019 article published by *Civil-Military Integration in Cyberspace* promoted the RHG model as a new standard for cybersecurity competitions in China, joining classic cyber games like capture the flag and jeopardy.[35]

Figure 3: A Timeline of RHG Finals.[36]

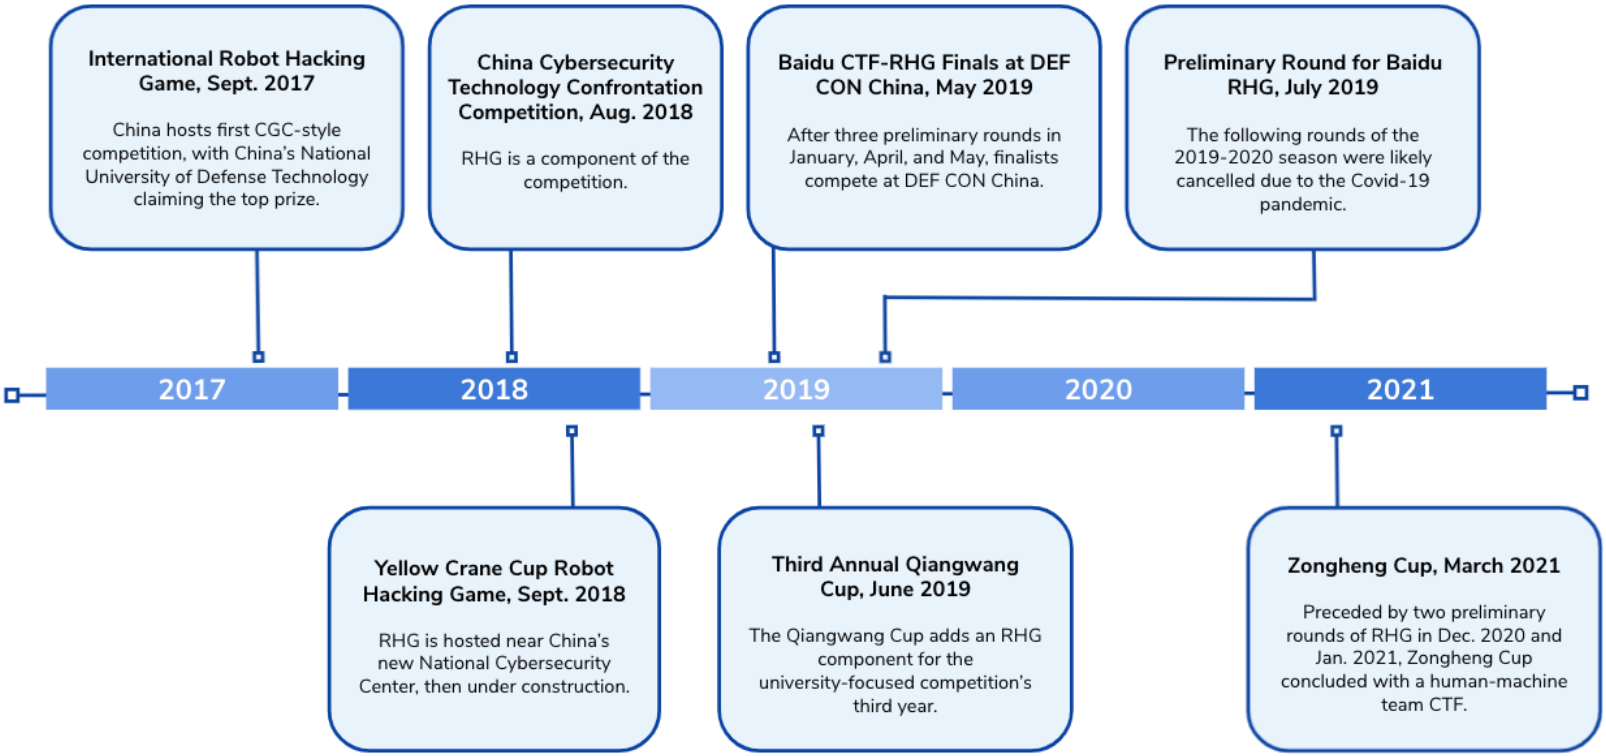

Source: Center for Security and Emerging Technology

Implications of China's RHGs

China's Pursuit Will Endure

Xi wants China to become a "cyber powerhouse." Strategic policy documents signal that AVL tools are key to achieving Xi's ambitions.[37] As a result, the Party expects organizations able to research the technology to do so.[38] Efforts to develop AVL tools will persist until new strategic documents redefine what it means to be a "cyber powerhouse" or the technology meets the needs of the government. The widespread adoption of the RHG competition model provides strong incentives for Chinese academics, firms, and PLA laboratories to develop the technology. Although the prize money for winners of RHG competitions is paltry compared to private-sector competitions ($50K vs. $250K), party committees at universities and companies are able to encourage their organization's participation. In the United States, such small awards would fall short of the costs for just one researcher to work on AVL tools. In China, the CCP's political mandate to pursue the technology ensures that the competitions and technology remain a focal point for the cybersecurity community, regardless of the rewards offered. Organizations that are able to support the technology's development but choose not to would be out of step with the party—a politically untenable position.[39] China's crackdown on tech firms will concentrate minds on the need to be on the same team as party policymakers. The strong political signal by the CCP mobilizes resources across China to focus on the technologies' development.

Increasing PLA Involvement

The Ministry of State Security 13th Bureau and Ministry of Education served as government "steering organizations" (指导单位) responsible for managing the first three RHG competitions.[40] Some regional offices of the MSS 13th Bureau run cyber operations in partnership with regional State Security Bureaus.[41] But the 13th Bureau is also responsible for general cybersecurity issues within government agencies. The motivation behind the bureau's involvement in the first three RHGs is unclear—the

Ministry of Education's involvement may suggest benign intentions. Although the MSS 13th Bureau has not hosted an RHG competition since late 2018, research on AVL tools may have been moved in-house. The technology's offensive and defensive uses, combined with the bureau's dual-purpose missions, obfuscate the nature of its interest. Few questions remain about the interest of the PLA, however.

The Third Annual Qiangwang Cup (强网杯), which is self-described as having "a natural tendency towards military-civil fusion (军民融合)," marked the shift towards PLA involvement. Qiangwang Cup was the first competition overseen by the Central Cyberspace Administration of China and PLA Information Engineering University.[42] The shift from MOE and MSS 13th Bureau oversight suggests increased military interest in the technology. PLA Information Engineering University is part of the PLA Strategic Support Force's Network Systems Department, which is responsible for military hacking operations.[43] The university's oversight of the RHG may reflect an interest in recruiting students with knowledge of AVL tools, since the Qiangwang Cup is a competition for college students.

In 2021, military oversight of RHGs expanded further. The Key State Laboratory for Information System Security Technology (信息系统安全技术重点实验室), a lab administered by the PLA's Equipment Development Department, managed the 2021 Zongheng Cup (纵横杯).[44] According to the U.S.-China Security and Economic Review Commission, the Equipment Development Department "plays a central role in military modernization by overseeing weapons development across the entirety of the PLA."[45] The lab's oversight of the competition indicates an uptick in the PLA's responsibility for developing, and possibly deploying, the technology.

RHGs Are Evolving

As long as AVL tools are central to the competition, hosts can change game structures and experiment with operational concepts. The Zongheng Cup introduced human-machine team competitions, where an automated AVL system supports two people in a 3-vs.-3

capture-the-flag style competition.[46] This human-in-the-loop concept is behind one of DARPA's follow-on programs to the CGC—Computers and Humans Exploring Software Security (CHESS).[47] Overseen by a lab affiliated with the PLA Equipment Development Department, the Zongheng Cup demonstrates converging operational concepts between the United States and China. RHGs are no longer changing their structures to match those of the Ministry of Education, but instead those of the PLA.

Experience and Collaboration

China's system of competitions attracts new participants, facilitates hands-on experience, and fosters relationships between institutions and competing teams. "Promot[ing] the training and selection of talents in the field of AI-based cybersecurity" was a key objective for China's first RHG and remains a goal of each subsequent competition.[48] Although automated software vulnerability discovery, patching, and exploitation promise to be more efficient than human professionals alone, these systems still require specialized knowledge to deploy. Operators with experience using the technology can more easily diagnose and solve errors as they arise during deployment.

Competitions also encourage relationships between participants. These relationships can be formal, such as teams representing multiple institutions, or informal—social gatherings after the competition. Having a cohort of researchers familiar with the technology is crucial to its successful deployment. Close professional connections could provide networks for troubleshooting technical issues or helping the PLA deploy the technology.

Conclusion

China's state hacking teams, which involve the PLA and Ministry of State Security, stand ready to adopt AVL tools. A report from *MIT Technology Review* detailed how China's government monitored cybersecurity competitions for new tools and techniques, then rapidly acquired and deployed them against domestic surveillance targets in Xinjiang.[49] RHGs are likely no different. But a full life-cycle AVL tool has not been compiled yet. Instead, individual parts of the tools—like fuzzers, symbolic execution, or automatic exploit generation—may progress in a piecemeal fashion. Automated vulnerability tools are already widely deployed in software development, so improvements in the technology are building on past success. Still, the CEO of Qihoo360, the cybersecurity firm responsible for China's Cybersecurity Military-Civil Fusion Innovation Center—among other state ties—called automated vulnerability discovery tools an "Assassin's Mace" for China.[50] The arcane term references the military strategy of creating an asymmetric advantage over a more powerful enemy—in DOD jargon, it is the Chinese Offset Strategy. For China's military, attacking an adversary's command and control system to disrupt "system-of-systems" communication would fit the bill.[51] AVL tools could help the PLA foment such an attack.

U.S. policymakers should consider whether current support for developing AVL technologies is enough. China's largest tech firms and universities are now competing at events hosted by the PLA's labs. Those competitions, in turn, spur innovation, connect researchers, and create a platform for iteratively testing and improving the technology. The United States, by contrast, supports three DARPA programs: Assured Micropatching, CHESS, and Harnessing Autonomy for Countering Cyberadversary Systems.[52] Combined with any classified programs or allocations, these three programs represent the USG's best efforts to develop AVL tools.

To get the most out of the technology and maintain any lead over China in this technology, the United States may need to invest more in developing AVL tools. Public competitions with cash prizes large enough to turn winners into businesses could be a good first step. DARPA's CGC in 2016 helped launch a few new companies.

But increasing investment and public interest in the technology by the cybersecurity community could yield even greater dividends. With some luck and more public investment, new businesses and a more secure U.S. cyber domain could be in the offing.

Author

Dakota Cary is a research analyst at CSET, where he works on the CyberAI project.

Acknowledgments

Thanks to Perri Adams, Sean Heelan, Conrad Stosz, John Bansemer, Chris Rohlf, Ben Murphy, Kady Arthur, Rael Baird, Kayla Goode, Ngor Luong and Andrew Lohn.

Document Identifier: doi: 10.51593/2021CA005

Endnotes

[1] 系统管理员, "浙江大学 phrack 战队在首届国际机器人网络安全大赛中荣获三等奖." 浙江大学电气工程学院, September 26, 2017, https://perma.cc/W5VH-J7F5. The translation to "Robot Hacking Game" from Mandarin is both a direct translation, and the translation used in China's own translations. Figure 2 shows each server rack embossed with "RHG" in large white letters at the top to drive home the competition's branding. Although the name evokes thoughts of animated machinery moving about, the more appropriate English-language idea might be a "bot"— used to denote automated bits of software from virtual assistants to automated web scrapers.

[2] "国防科技大学电子科学学院'Halfbit'代表队夺得首届国际机器人网络安全大赛冠军," 国防科技大学, November 8, 2017, https://perma.cc/ESL5-8YNL; 中共中央网络安全和信息化委员会办公室, "第三届‘强网杯’全国网络安全挑战赛正式启动," April 23, 2019, https://perma.cc/9E4N-CGY4; 安全 419, "RHG 赛事平台落地‘纵横杯’ 人工智能自动化攻防演练或成行业常态比赛," Sohu, March 30, 2021, https://perma.cc/7E53-2FBZ.

[3] 奇安信, "冠军！IQ 战队夺魁 RHG 国际机器人网络安全对抗赛," 奇安信, March 30, 2021, https://perma.cc/93CH-QXNN.

[4] Translator's note: For a more in-depth discussion in English of the Chinese term 网络强国, which can be rendered as "cyber powerhouse" or "cyber superpower," see Rogier Creemers et al., "Lexicon: 网络强国 Wǎngluo Qiángguó," *New America*, May 31, 2018, https://www.newamerica.org/cybersecurity-initiative/digichina/blog/lexicon-wangluo-qiangguo/.

[5] Dustin Fraze, "Cyber Grand Challenge," Defense Advanced Research Projects Agency, accessed August 27, 2021, https://perma.cc/65W8-XEEK.

[6] Defense Advance Research Projects Agency, "Cyber Grand Challenge Rules, Version 3," Massachusetts Institute of Technology. November 18, 2014, 12-13, https://archive.ll.mit.edu/cybergrandchallenge/docs/CGC_Rules_18_Nov_14_Version_3.pdf. The "attacks" were, in fact, proofs of concepts that exploited other teams' vulnerabilities. An automated referee system evaluated whether the attacks would work as intended, and if so, awarded points. The victims were docked points, but no malware was installed on the targeted system. This structure prevented teams from permanently impairing their opponents and focused the game on vulnerability discovery, exploitation, and patching.

[7] Fraze, "Cyber Grand Challenge."

[8] Tamulyn Takakura, "Top 5 Takeaways from the 'ForAllSecure Makes Software Security Autonomous' Livestream," ForAll Secure, April 17, 2019, https://perma.cc/A7KV-Q4E3.

[9] Defense Advanced Research Projects Agency, "DARPA Celebrates Cyber Grand Challenge Winners," U.S. Department of Defense, August 5, 2017, https://www.darpa.mil/news-events/2016-08-05a.

[10] Vito Genovese, "2016 DEF CON CTF Final Scores," Legitimate Business Syndicate, accessed August 27, 2021, https://blog.legitbs.net/2016/09/2016-def-con-ctf-final-scores.html.

[11] Devin Coldewey, "Carnegie Mellon's Mayhem AI Takes Home $2 Million from DARPA's Cyber Grand Challenge," *TechCrunch*, August 5, 2016, https://perma.cc/NMN4-ZF8L; Stephanie Kanowitz, "DARPA's All-Machine Cyber Challenge," GCN, July 21, 2016, https://perma.cc/M92F-5W6Q; "How ForAllSecure's 'Mayhem' Won DARPA's Cyber Grand Challenge." *Insider*, August 23, 2016, https://perma.cc/2T4X-4VLS; Sean Gallagher, "The World Series of Hacking—without Humans," *ArsTechnica*, August 15, 2016, https://perma.cc/HA4K-W26T; Mark Ward, "Can Machines Keep Us Safe from Cyber-Attack?," *BBC News*, August 2, 2016, https://perma.cc/FF84-X594; Tom Simonite, "Pentagon Bot Battle Shows How Computers Can Fix Their Own Flaws," *MIT Technology Review*, August 4, 2016, https://perma.cc/4CAY-73DQ; Cade Metz, "Security Bots Will Battle in Vegas for DARPA's Hacking Crown," *WIRED*, July 28, 2016, https://perma.cc/SY84-W6PA; David Brumley, "Mayhem Wins DARPA CGC," ForAllSecure, August 6, 2016, https://forallsecure.com/blog/mayhem-wins-darpa-cgc.

[12] Tom Simonite, "This Bot Hunts Software Bugs for the Pentagon," *WIRED*, June 1, 2020, https://www.wired.com/story/bot-hunts-software-bugs-pentagon/.

[13] GrammaTech, "Five Points Capital Completes Acquisition of GrammaTech," PR Newswire, November 12, 2019, https://www.prnewswire.com/news-releases/five-points-capital-completes-acquisition-of-grammatech-300955576.html.

[14] Behring Center, "Innovations in Defense: Artificial Intelligence and the Challenge of Cybersecurity," Smithsonian National Museum of American History, 2016, https://perma.cc/US2D-5XHD.

[15] 孙宝云, "全球网络部队建设、网络安全人才培养与网络安全教育：2014 年新动向," *北京电子科技学院学报* (March 2015): 66–74, 87; 编辑部, "国际动态（2014 年 5 月）," *中国信息安全* (June 2014): 22–23.

[16] Creemers et al., "Lexicon: 网络强国 Wǎngluo Qiángguó."

[17] 徐贵宝, "美国智能网络攻防对我国网络强国的启示," 世界电信, no. 3: 57–60.

[18] "The Google DeepMind Challenge Match," DeepMind, March 2016, https://deepmind.com/alphago-korea.

[19] Defense Advanced Research Projects Agency, "DARPA Celebrates Cyber Grand Challenge Winners," August 5, 2017, https://www.darpa.mil/news-events/2016-08-05a.

[20] 徐贵宝, "美国智能网络攻防对我国网络强国的启示," 世界电信, no. 3: 57–60.

[21] 徐贵宝, "美国智能网络攻防对我国网络强国的启示."

[22] 徐贵宝, "美国智能网络攻防对我国网络强国的启示."

[23] Creemers et al., "Lexicon: 网络强国 Wǎngluo Qiángguó." Author's note: Although both policy documents referenced here discuss "AI" for cybersecurity, the RHG and CGC did not use what most experts would call AI. Instead, these automated systems relied on prescribed reasoning systems to make decisions. None used machine learning, deep learning, reinforcement learning, or any other framework typically associated with AI.

[24] 系统管理员, "浙江大学 phrack 战队在首届国际机器人网络安全大赛中荣获三等奖," 浙江大学电气工程学院, September 26, 2017, https://perma.cc/W5VH-J7F5.

[25] 系系统管理员, "浙江大学 phrack 战队在首届国际机器人网络安全大赛中荣获三等奖."

[26] 系系统管理员, "浙江大学 phrack 战队在首届国际机器人网络安全大赛中荣获三等奖"; 雷锋网, "22 支机器人战队比赛‘收割’漏洞，国家队这样赢了," 百科 TA 说, September 22, 2017, https://perma.cc/ZAD6-6JZC; 360 网络攻防实验室, I 春秋, "首届国际机器人网络安全大赛（武汉）– 安全客，安全资讯平台," 安全客, September 21, 2017, https://perma.cc/PH7T-66PL.

[27] 雷锋网, "22 支机器人战队比赛‘收割’漏洞，国家队这样赢了," 百科 TA 说, September 22, 2017, https://perma.cc/ZAD6-6JZC.

[28] 360 网络攻防实验室, I 春秋, "首届国际机器人网络安全大赛（武汉）- 安全客，安全资讯平台." That cybersecurity firm, IntegrityTech, has hosted all of China's competitions and owns the RHG competition model as its intellectual property. "首届国际机器人网络安全大赛 - CTF 大本营 - 网络安全竞赛平台," i 春秋, September 2019, https://perma.cc/G7BF-2YST; Peter Mattis and Matthew Brazil, *Chinese Communist Espionage: An Intelligence Primer* (Annapolis, MD: Naval Institute Press, 2019); "RHG 赛事平台落地‘纵横杯’人工智能自动化攻防演练或成行业常态比赛," Sohu, March 30, 2021, https://perma.cc/7E53-2FBZ;

IntegrityTech, "'纵横杯' 网络安全竞赛," 纵横杯, accessed August 27, 2021, https://perma.cc/5JRP-6W3F.

[29] 360 网络攻防实验室, l 春秋, "首届国际机器人网络安全大赛（武汉）- 安全客，安全资讯平台." The U.S. team, Highlander, came from University of California, Riverside. It did not compete under the same name as any DARPA CGC participants, however. Chengyu Song's CV identifies himself as an employee of University of California, Riverside faculty and a participant in DARPA CGC. https://perma.cc/8XN6-VUBM Dr. Song likely led the Highlander team from UC Riverside at China's first RHG, but there is no sourcing that directly supports this claim. The two other international teams were from Ukraine and Russia.

[30] Interview held under Chatham House Rules. Chris Bing, "Huawei Tried to Acquire Technology from the Winners of the Cyber Grand Challenge," *CyberScoop*, November 9, 2017, https://perma.cc/7KUE-NNEF.

[31] 360 网络攻防实验室, l 春秋, "首届国际机器人网络安全大赛（武汉）- 安全客，安全资讯平台"; 雷锋网, "22 支机器人战队比赛'收割'漏洞，国家队这样赢了"; 网业创新, "机器人网络安全大赛成功举办," 中国信息安全, 96–98.

[32] Paul Triolo, Elsa Kania, and Graham Webster, "Translation: Chinese Government Outlines AI Ambitions through 2020," *New America*, January 26, 2018, https://www.newamerica.org/cybersecurity-initiative/digichina/blog/translation-chinese-government-outlines-ai-ambitions-through-2020/.

[33] Triolo, Kania, and Webster, "Translation: Chinese Government Outlines AI Ambitions through 2020."

[34] "RHG 赛事平台落地 '纵横杯' 人工智能自动化攻防演练或成行业常态比赛," Sohu, March 30, 2021, https://perma.cc/7E53-2FBZ.

[35] 邬江兴, "网络安全学科竞赛的创新与发展," 网信军民融合, June 11–13.

[36] "国防科技大学电子科学学院'Halfbit'代表队夺得首届国际机器人网络安全大赛冠军,"国防科技大学, November 8, 2017, https://perma.cc/ESL5-8YNL; 中国网, "2018 中国网络安全技术对抗赛在京召开 人工智能安全夺旗赛成焦点," *新华网*, August 16, 2018, https://perma.cc/8D5A-M9ZX; 国海洋大学计算机系, "中国海洋大学战队荣获'黄鹤杯'RHG 机器人网络安全大赛季军," 中国海洋大学, July 16, 2019, https://perma.cc/BG6R-FKMJ; 企鹅号–丁牛科技, "丁牛科技荣获'黄鹤杯'机器人网络安全大赛企业第一名，大赛第七名!," 腾讯云, September 21, 2018, https://perma.cc/3TSZ-W8VQ; 小度房产, "BCTF-RHG 场景攻防热身赛闭幕，人工智能网络安全竞赛的新战场开启," 百科, April 16, 2019, https://perma.cc/SU6W-K49W; 科技正能量, "BCTF 的"黄埔军校"，将培养一支网络安全的'AI 国防军,'" 360Doc, June 25, 2020, https://perma.cc/526Y-U2JB; 百度安全, "百度国际网络安全技术对抗赛," 百度, January 2019,

https://perma.cc/2YRV-CPM6; 360 网络攻防实验室, I 春秋, "第三届强网杯全国网络安全挑战赛人工智能挑战赛 - CTF 大本营," 网络安全竞赛平台-I 春秋, 2018, https://perma.cc/98QH-UQ57; 赛宁网安 Cyberpeace, "第三届'强网杯'全国网络安全挑战赛圆满落幕-腾讯 eee 成功卫冕," 百度, June 19, 2019, https://perma.cc/4DMS-N7GN; 中共中央网络安全和信息化委员会办公室, "第三届'强网杯'全国网络安全挑战赛正式启动," April 23, 2019, https://perma.cc/9E4N-CGY4; 恒安嘉新, "恒安嘉新 EversecLab 战队获 BCTF AI 对抗赛漏洞挖掘方向第一名," 恒安嘉新（北京）科技股份公司, July 10, 2019, https://perma.cc/CA5R-VU6V; 新华网, "2020'黄鹤杯'网络安全人才与创新峰会在汉举办," 百度, September 13, 2020, https://perma.cc/NBF8-LBJD; IntegrityTech, "'纵横杯'网络安全竞赛," 纵横杯, accessed August 27, 2021, https://perma.cc/5JRP-6W3F.

[37] China's "*Internet +" Artificial Intelligence Three-Year Action Plan* released for 2018 to 2020, the sequel to the 2015 edition authored by Xu, doubled-down on his recommendations. That document argued that AI technology could be applied to such problems as "vulnerability discovery, security testing, threat warning, attack detection, and emergency response." Triolo, Kania, and Webster, "Translation: Chinese Government Outlines AI Ambitions through 2020."

[38] Alex Stone and Peter Wood, "China's Military-Civil Fusion Strategy: A View From Chinese Strategists" (China Aerospace Studies Institute, accessed August 2021), 8, https://static1.squarespace.com/static/5e356cfae72e4563b10cd310/t/5ee37fc2fcb96f58706a52e1/1591967685829/CASI+China%27s+Military+Civil+Fusion+Strategy-+Full+final.pdf.

[39] The small cash-sums at RHGs are the result of explicit government policy to decrease payouts when a government organization is involved in a cybersecurity competition. Office of the Chinese Communist Party Central Cyberspace Affairs Commission (中央网络安全和信息化委员会办公室, and 网信办) and the PRC Ministry of Public Security (公安部), "关于规范促进网络安全竞赛活动的通知," 国家互联网办公室, September 17, 2018, https://perma.cc/5423-H72X.

[40] 360 网络攻防实验室, I 春秋, "首届国际机器人网络安全大赛（武汉）- 安全客，安全资讯平台"; i 春秋, "首届国际机器人网络安全大赛（武汉）."

[41] Insikt Group, "Recorded Future Research Concludes Chinese Ministry of State Security Behind APT3," Recorded Future, May 17, 2017, https://www.recordedfuture.com/chinese-mss-behind-apt3/. Mattis identifies CNITSEC as the 13th Bureau of the MSS. Mattis and Brazil, *Chinese Communist Espionage.*

[42] 邬江兴, "网络安全学科竞赛的创新与发展," 网信军民融合, June 11–13; 中共中央网络安全和信息化委员会办公室, "第三届'强网杯'全国网络安全挑战赛正式启动."

[43] Rachael Burton, "The People's Liberation Army Strategic Support Force: Leadership and Structure," Project2049, September 25, 2018, https://project2049.net/2018/09/25/the-peoples-liberation-army-strategic-support-force-leadership-and-structure/.

[44] "RHG 赛事平台落地‘纵横杯’人工智能自动化攻防演练或成行业常态比赛"; IntegrityTech, "'纵横杯'网络安全竞赛"; 奇安信, "冠军！IQ 战队夺魁 RHG 国际机器人网络安全对抗赛." Lab Attribution: The Key State Laboratory for Information System Security Technology (信息系统安全技术重点实验室) is not well publicized. A 2015 research paper on cyber attack and defense technologies, show three authors from the Information System Security Technology Laboratory. https://perma.cc/9T26-WUYR Each of the three authors also reports and affiliation with the Beijing Institute of System Engineering (北京系统工程研究所) in their by-line. The authors are either all concurrently employed by the Beijing Institute of System Engineering or the Information System Security Laboratory is an organization within the Institute. A media organization owned by the Shanghai Municipal Government shows that the Beijing Institute of System Engineering was under the PLA General Armament Department as of 2016. See "一种基于概率转移的 Cyber 攻击场景感知推理技术," Baidu, June 29, 2020, https://perma.cc/99RB-9GAU?type=image. Following the PLA reorganization, the General Armament Department became the Equipment Development Department. See U.S.-China Economic and Security Review Commission, "Section 2: China's Military Modernization in 2017," 2017, 201, https://perma.cc/ZP8X-X6JG. Barring the splitting of pieces of the General Armaments Department during the PLA reforms, the Beijing Institute of System Engineering should now reside under the Equipment Development Department of the CMC.

[45] U.S.-China Economic and Security Review Commission, "Section 2: China's Military Modernization in 2017," 201.

[46] 奇安信, "冠军！IQ 战队夺魁 RHG 国际机器人网络安全对抗赛."

[47] The Computers and Humans Exploring Software Security (CHESS) program aims to improve automated vulnerability discovery with a human-machine team approach to analysis. Researchers hope to secure mission critical systems from zero-day exploits. Dustin Fraze, "Computers and Humans Exploring Software Security (CHESS)," Defense Advanced Research Projects Agency, accessed August 27, 2021, https://www.darpa.mil/program/computers-and-humans-exploring-software-security.

[48] i 春秋, "首届国际机器人网络安全大赛（武汉）."

[49] Patrick Howell O'Neill, "How China Turned a Prize-Winning iPhone Hack against the Uyghurs," *MIT Technology Review*, May 6, 2021, https://perma.cc/Y34Q-H82P.

[50] Jiang Jie, "China Unveils Its First Civil-Military Cybersecurity Innovation Center," *People's Daily Online*, December 28, 2017, https://perma.cc/R8QB-VK4J; 网络传播杂志, "360: 自觉担当责任维护网络安全," 中共中央网络安全和信息化委员会办公室, November 6, 2018, https://perma.cc/ENA2-WZ3F.

[51] "China's Cyber Power in a New Era" in "Asia Pacific Regional Security Assessment 2019" (International Institute for Strategic Studies, May 2019), 77–90, https://www.iiss.org/publications/strategic-dossiers/asiapacific-regional-security-assessment-2019/rsa19-07-chapter-5.

[52] Sergey Bratus, "Assured Micropatching (AMP)," Defense Advanced Research Projects Agency, accessed August 27, 2021, https://www.darpa.mil/program/assured-micropatching; Dustin Fraze, "Harnessing Autonomy for Countering Cyberadversary Systems (HACCS)," Defense Advanced Research Projects Agency, accessed August 27, 2021, https://www.darpa.mil/program/harnessing-autonomy-for-countering-cyberadversary-systems; Fraze, "Computers and Humans Exploring Software Security (CHESS)."