Valdai Papers

Valdai Discussion Club

#95

Why We Must Prohibit Cyberattacks on Nuclear Systems: the Case for Pre-Emptive US–Russia Arms Control

Andrew Futter

valdaiclub.com
#valdaiclub

November, 2018

About the Author

Andrew Futter
Associate Professor of International Politics at the University of Leicester

This publication and other papers are available on
http://valdaiclub.com/a/valdai-papers/

42 Bolshaya Tatarskaya st., Moscow, 115184, Russia

Introduction: When Cyber Meets Nuclear

Almost 35 years ago, US President Ronald Reagan settled down in the White House to watch the latest Hollywood blockbuster *WarGames* as part of his regular Sunday film night. The film, starring a young Matthew Broderick, depicted a teenage computer hacker accidentally breaking into top-secret Pentagon supercomputers that controlled US nuclear weapons. The result was very nearly (a fictional) nuclear World War Three with the Soviet Union. Reagan was so taken by the film that he ordered secret review to be conducted into whether US nuclear weapons could be vulnerable to Computer Network Attacks[1], and whether hackers could somehow launch a US nuclear weapon without authorisation by interfering with computers. Officials reported back to the President that the threat was real and possibly far worse than they expected.[2] What started with a 1983 movie would result in the first proper recognition that nuclear systems were vulnerable to cyberattacks.

A generation later this threat has multiplied considerably. Far more aspects of nuclear operations – from the weapons and delivery vehicles to the command and control apparatus and targeting software – rely on increasingly complex computer code, making them potential targets for malicious attackers. All nuclear-armed states also have plans to modernise their nuclear systems and to incorporate more rather than less computer technology, and to exploit the possibilities offered by digital networking and programming. At the same time, there is a growing recognition of the threat posed by hackers to all types of computer systems, including those that control critical national infrastructure. The Stuxnet attack on the Iranian enrichment facility at Natanz discovered in 2010 is perhaps the best-known example, but cyberattacks have become a regular occurrence and never far from the minds of military planners. In fact, most nations now have units in their militaries and associated doctrines dedicated to offensive cyber operations, and some have even spoken of cyber warfare. Taken together we stand at a point today where all nations' nuclear weapons could be vulnerable to a cyberattack. A fact recognised by, amongst others, the US Defence Science Board in a 2013 report.[3]

---

[1] *Find more details in the author's most recent book: Futter, A, 2018, 'Hacking the Bomb', Georgetown University Press. Available from: http://press.georgetown.edu/book/georgetown/hacking-bomb*

[2] *Kaplan, F, 2016, ''WarGames' and Cybersecurity's Debt to a Hollywood Hack', The New York Times, February 19. Available from: https://www.nytimes.com/2016/02/21/movies/wargames-and-cybersecuritys-debt-to-a-hollywood-hack.html*

[3] *'Task Force Report: Resilient Military Systems and the Advanced Cyber Threat', 2013, United States Department of Defense, Defense Science Board, January. Available from: http://www.acq.osd.mil/dsb/reports/ResilientMilitarySystems.CyberThreat.pdf*

The good news is that this threat is still to some extent in its infancy, and there is time to 'get ahead' and perhaps mitigate its worst aspects before they fully materialise and become normalised. The bad news is that US–Russia relations and prospects for arms control are at their nadir for a generation, and both (and maybe others too) may actively be pursuing the ability to hack into an adversary's nuclear weapons systems. This paper is a call for renewed cooperation in the nuclear realm and makes the case for moratoria between the US and Russia, and hopefully others, that prohibits cyberattacks against nuclear systems. As will be explained below, all states – and everyone on the planet – would be better off without hackers messing around inside the systems that control nuclear weapons.

The Emergence of a Norm

The incorporation of Computer Network Operations (a more precise label than 'cyber'[4]) into military planning can probably be traced back at least 30 years, certainly to the late 1980s and the so-called Revolution in Military Affairs of the early 1990s. However, it is perhaps only in the last decade or so that such thinking – and the required technological capability – has percolated up to the strategic, nuclear level. Specifically, it can be traced back to the George W. Bush administration's plans to diversify nuclear deterrence thinking to include a greater role of non-nuclear systems in its global strike plans in the early 2000s, the decision to use cyber capabilities against the Iranian nuclear programme, and then more recently as part of Pentagon proposals for new 'full spectrum missile defence' and global prompt strike missions.

The idea for 'full spectrum missile defence' is fairly simple. New methods to prevent missiles being launched should be included alongside traditional methods such as ballistic missile defence systems based on kinetic intercept. But instead of waiting for a missile to be fired, the idea is to prevent the missile from being launched at all by interfering with key control systems, or the weapon itself, either electronically (by targeting its telemetry) or digitally (by targeting its software and hardware or its support systems). To achieve this, hackers would break into nuclear control systems prior to the missile being fired, lace systems within the missile or associated infrastructure with malware, or interfere in normal operations in another way. This is known as 'left-of-launch'. In theory, combining kinetic and non-kinetic methods of missile defence in this way makes

[4] *See, Futter, A, 2018, 'Cyber Semantics: Why We Should Retire the Latest Buzzword in Security Studies', Journal of Cyber Policy. Available from: https://www.tandfonline.com/doi/full/10.1080/23738871.2018.1514417*

the system more comprehensive and reduces the reliance on in-flight interception (which even today remains a very difficult task[5]). As Brian McKeon, Principal Deputy Under Secretary of Defence for Policy, explained in a 2016 testimony to Congress,

> [W]e need to develop a wider range of tools, and that includes the efforts underway to address such threats before they are launched, or 'left-of-launch'. The development of left-of-launch capabilities will provide US decision-makers additional tools and opportunities to defeat missiles. This will in turn reduce the burden on our 'right-of-launch' ballistic missile defence capabilities. Taken together, left-of-launch and right-of-launch will lead to more effective and resilient capabilities to defeat adversary ballistic missile threats.[6]

The most obvious target for the US full spectrum defence mission is North Korea, and it is at least possible that US hackers were responsible for a series of recent missile test failures.[7] It is conceivable that similar plans are also afoot against Iran as a hedge against a future Iranian nuclear capability. It also seems likely that the Donald Trump administration's forthcoming *Missile Defense Review* could include reference to greater 'full spectrum' capabilities in addition to upgrades to existing systems.

In the past two decades, what has essentially happened is that missile defence has met precision strike; and offense and defence have become the commingled in military and nuclear planning. This might even be interpreted as a slow rejection of the idea of deterrence through mutual vulnerability, the cornerstone of Mutual Assured Destruction (MAD), and a move towards more active measures of defence and deterrence. This shift has been driven primarily by changes in the 'demand side' of nuclear deterrence – that is who or what needs to be deterred and how, i.e. a shift from preventing a massive nuclear strike from a peer competitor to dealing with nuclear threats from smaller 'rogue' states and maybe terrorists who may not 'play by the same rules' or behave as 'rationally' as peer competitions. But this has now been shifted again thanks to changes in the 'supply side' dynamics: that is, the enormous developments in the technologies and weapons systems that might be used to achieve this,

---

[5] *For example, Larter, D, 2018, 'Reality Check: Failures Happen, Even in Missile Defense Testing', Defense News, February 1. Available from: https://www.defensenews.com/naval/2018/02/01/reality-check-failures-happen-even-in-missile-defense-testing/*

[6] *McKeon, BP, 2016, Statement before the Senate Armed Services Subcommittee on Strategic Forces', April 13. Available from: http://www.armed-services.senate.gov/imo/media/doc/McKeon_04-13-16.pdf*

[7] *Sanger, DE & Broad, W, 2017, 'Trump Inherits Secret Cyberwar Against North Korean Missiles', The New York Times, March 4. Available from: https://www.nytimes.com/2017/03/04/world/asia/north-korea-missile-program-sabotage.html*

themselves a direct product of the latest information or computer revolution. Digital weapons, Computer Network Operations and other capabilities that fall under the cyber moniker are perhaps the best example of this. But it also includes a range of other Advanced Conventional Weapons able to augment and in some case replace nuclear weapons in strategic thinking and policy. The result is growing interest in deterrence by denial (that is preventing an attack from happening) in addition to or perhaps instead of deterrence through retaliation (the threat punishment after an attack).

New Problems and Dynamics

The problem is that unlike kinetic missile defence interceptors, which can be deployed, seen, and quantified, left-of-launch cyber capabilities are by their very nature nebulous and can even be ephemeral. It would therefore be natural for Moscow and Beijing to be suspicious of these developments and to assume (much like with conventional missile defence programmes) that these capabilities might also be used against them in some future scenario. The difference is that there is no way to monitor the extent of the threat and thus react accordingly (by building more missiles, developing new penetration aids, etc) to maintain rough strategic parity or at least to guard against one side (in this case the US, gaining a strategic advantage or even superiority).

For example, the 44 Ground-Based Interceptors that the US currently has deployed in Alaska and California (even when added to other Ground-Based Midcourse Defence deployments elsewhere) are probably not a threat to Russian or Chinese assured retaliation at the moment. But, if the number of interceptors and the required sensors were to be expanded significantly – and if Russian nuclear forces were reduced and Chinese nuclear forces not increased – they could be. The difference is that both Russia and China would react, as arguably they both are already, with new capabilities able to bypass US missile defences before the strategic balance shifted.[8] However, the much more nebulous and intangible nature of left-of-launch technologies would make this far harder to judge, and more difficult to know exactly how to react. Moreover, whereas long-range kinetic missile defence systems are principally designed against land-based missiles, the ability to attack central command

[8] *Roth, A, 2018, 'Putin Threatens US Arms Race with New Missile Declaration', The Guardian, March 1. Available from: https://www.theguardian.com/world/2018/mar/01/vladimir-putin-threatens-arms-race-with-new-missiles-announcement*

nodes (and spoof early warning sensors) makes all nuclear systems vulnerable. Even the nuclear-armed submarine or mobile missiles, key to US and Russian secure second-strike capabilities, could be targeted. As a result, it is difficult to see how a new full spectrum missile defence policy can do anything other than exacerbate the concerns of strategic competitors and lead to greater uncertainty.[9]

There are a number of other concerns with full spectrum missile defence that warrant further unpacking. The first is that pursuing left-of-launch options against potential missile or nuclear threats transforms the missile defence mission, and security policy more generally, from largely passive into one of prevention. This is because systems will almost certainly have to be breached before a threat fully materialises, and almost certainly before a missile is launched. This is known as 'active defence' and would involve hackers breaching sensitive systems prior to a missile being fired. It might even involve interference in the supply chain or focus on the human element. For sure, some operations could be carried out once a missile or other nuclear delivery system was being readied for firing or use, but to enhance confidence that such operations would work, hackers would surely want to have created backdoor access or laced these systems beforehand.

The second is that even the possibility that nuclear systems could be vulnerable to hackers, and therefore may not work as expected or planned, will decrease trust and stability between nuclear-armed actors. Decreased certainty in these systems might lead to pressure to enhance positive control of nuclear weapons, that is to ensure that they will always work, potentially at the expense of keeping them safe and secure. It would also almost certainly drive other states to develop their own left-of-launch operations and capabilities, making all states feel less secure irrespective of whether there is any real intention to use them. A more fearful environment is unlikely to help with any bilateral or multilateral arms control initiative either.

The third is the increased risks of accidents and inadvertent outcomes, either from interfering with the wrong systems, or from being discovered inside these systems. For example, operations targeting conventional weaponry or support systems (such as satellites) might also impact those managing nuclear systems or might spread to nuclear systems. Likewise, it is at least possible that once inside these systems, hackers might inadvertently cause something

[9] *For a more detailed discussion of this see: Futter, A, 2016, 'The Dangers of Using Cyberattacks to Counter Nuclear Threats', Arms Control Today, July/August. Available from: https://www.armscontrol.org/print/7551*

to happen which they had not intended. It might also be difficult to ascertain the intention of any hacker or malware that was found inside these networks (and to verify their identity), and it would be natural for the victim to assume the worst, especially if discovered during a period of heightened tensions. Discovery could lead to knee-jerk responses, diplomatic tensions, and could even be interpreted as an act of war.

Finally, it is also possible that third-party actors such as terrorists might seek to cause or exacerbate a crisis through 'false-flag' attacks on the computer systems used to manage nuclear weapons. Importantly, non-state actors would be far more likely to seek 'enabling' actions against nuclear systems – i.e. to cause them to be used, as opposed to the 'disabling' goal for nation states. For example, these groups might seek to spoof early warning systems and manipulate the nuclear information space or cause havoc by conducting relatively minor interference during a crisis that may be seen as being carried out by an adversary due to the problems of attribution. All of these scenarios could clearly lead to escalation and increased nuclear risks.

For the moment, the tactic of using digital methods to interfere with nuclear and missile systems is primarily a US-centric idea (much in the same way until recently as it was with kinetic interception). But it is difficult to see why other states will not seek to follow suit. Russia, China, and perhaps others may explore similar possibilities against the US, increasing the risks for all involved. Indeed, the US might be even more vulnerable given its high reliance on complex systems across its nuclear weapons infrastructure and also given its recent plans to modernise all component parts of its nuclear command and control systems.[10]

Getting Ahead of the Threat

There are no easy fixes to this emerging problem, and history does not offer much reassurance when it comes to managing the impact of a new technology on warfare before it fully materialises. Neither does it seem like a particularly propitious time to embark on US–Russia arms control, although the recent meeting between Donald Trump and Vladimir Putin in Helsinki might

[10] *Futter, A, 2016, 'The Double-Edged Sword, US Nuclear Command and Control Modernisation', Bulletin of the Atomic Scientists, June 29. Available from: https://thebulletin.org/2016/06/the-double-edged-sword-us-nuclear-command-and-control-modernization/*

offer some hope.[11] But we do have an opportunity now to potentially mitigate the most worrying aspects of the cyber-nuclear challenge before it gets out of hand. This should start with a discussion of the most pressing threats for both sides; surely, hackers messing around in nuclear control systems linked to hundreds of missiles would be a good place to start. This then might lead into other initiatives of mutual interest.

The first is the development of new constraints in the use of Computer Network Operations against nuclear systems and the development of certain rules of the road. This might involve trying to get ahead of the threat by negotiating new forms of arms control in this space, and specifically through an agreement not to target nuclear weapons systems in this way. This does not necessarily have to look like the nuclear treaties of the past but could simply begin with a statement that the US and Russia recognise the severity and risks of attacking each other's nuclear command and control systems and foreswear the option of doing so. This might involve new declaratory policy about: (1) how such 'attacks' would be interpreted and likely responded to if discovered, and (2) that nuclear systems are off-limits. This might then be broadened to include other nuclear-armed states too. In a way, this could draw upon the ideas at the heart of the 1972 Anti-Ballistic Missile Treaty, which limited missile defences in the hope that this would aid predictability and stability between nuclear armed adversaries. Clearly, these options might not be verifiable in the traditional sense, nor stop non-state actors, but it is a start, and states would be unlikely to want to run the risk of being caught in violation of stated policy or agreements.

The second is better security, policy, and cooperation in this space. In the first instance this can be done unilaterally. For example, reducing alert times of nuclear systems (to minimise the ability of non-state hackers to cause a launch or explosion), working to keep these systems separate from other non-nuclear weaponry and command and control apparatus (to reduce the risk of attackers inadvertently hitting the wrong systems), and keeping the command and control infrastructure as simple as possible (so it is understandable and offers less vulnerabilities for attackers to exploit). This might then provide the basis for more ambitious bilateral and even multilateral endeavours aimed at building confidence and trust. Governments (US–Russia in the first instance, but hopefully others after) might wish to share good practice and possibly data on non-state threats, and even begin to build groups of governmental officials and other stakeholders to 'think outside

[11] *Bender, B, 2018, 'Leaked Document: Putin Lobbied Trump on Arms Control', Politico, July 8. Available from: https://www.politico.com/story/2018/08/07/putin-trump-arms-control-russia-724718*

the box' when it comes to new arms control mechanisms. It might also involve establishing a multinational joint early warning or threat assessment centre where officials and experts would be in regular dialogue and ready to react quickly to third party threats or other pressing issues.

We have a chance now to get ahead of a serious development in international nuclear politics that will likely have negative implications for all nuclear-armed states, and thus by implication, all of us. New arms control agreements may not necessarily look like those of the past, or be quick to design and implement, but this does not make the need any less. It took the best part of two decades to begin to codify the nuclear revolution, and we have arguably been refining this ever since. A dual approach of innovative arms control, constraints, and perhaps new rules of the road twinned with a better understanding of the challenge and a desire to work multi-nationally is one way to begin our response to the next generation of nuclear risks.

Ultimately, threats posed by new, emerging, and 'exotic' technologies in the nuclear realm will have to be incorporated into strategic stability discussions, arms control agreements, and into the broader non-proliferation and disarmament initiatives. We no longer live in a world where nuclear discussions can happen in a technological vacuum, or clear linkages between nuclear and non-nuclear can be ignored. Consequently, we must recognise that the nature of the global nuclear order has been shifted by the latest information and computer revolution, and that missile defences, precision conventional strike, drones, anti-space weapons, Artificial Intelligence as well as 'cyber' have transformed the way in which we must manage and secure the nuclear space.

Conclusion: Pre-Emptive Arms Control

Instead of heady and dangerous nuclear rhetoric and spending vast sums of money designing evermore destructive nuclear weapons, President Trump and President Putin and/or their representatives should sit down and begin a serious discussion about the main nuclear risks facing both countries. For sure, they will not agree on everything, but a mutual recognition that hacking into each other's nuclear systems benefits no one has to be a good place to start. Indeed, it is at least conceivable that taking arms control discussions in this direction might be far more fruitful than the current Strategic Arms Reduction

Treaty-based trajectory. The focus on arms reductions might temporarily perhaps best replaced with arms avoidance. It would also remind those who have suggested that arms control might be 'dead'[12] or that arms control in cyberspace is impossible, that new and different avenues exist to enhance stability; ones that may not necessarily mirror those of the past. The lesson from the Cold War was that even if it did not seem likely that anything would be agreed in the field of nuclear arms, both sides would keep talking because they recognised the high stakes involved.

This then must be multi-lateralised, because unlike many nuclear challenges of the Cold War, this is not principally a two-player game. Indeed, it impacts not just all nuclear-armed states, but all those with offensive cyber capabilities too. Rather than seeking an all-pervasive cyber treaty or grand bargain style nuclear agreement, a tentative first step could be to develop a cyber-nuclear convention or at least some very general rules of the road.[13] Key to this will be establishing and agreeing upon the terminology to be used, and on distinctions between what is and is not regarded as 'nuclear' for each party. This would then potentially provide the international basis to look at these challenges in international fora, and for more holistic discussions of arms control. The global nuclear order is in a period of flux and perhaps transition due in no small part to the myriad new weapons technologies of the latest information and computer revolution.

In the past, new military capabilities have had to be built (usually at enormous costs) and the threat realised before agreements could be made, but we may not be so lucky in the new techno-political context. If we can somehow come together and agree on the things that we as a society – and as nation states – most want to avoid, then perhaps we can begin to piece together frameworks to prevent this and begin to work backwards. Surely, we can all agree that hackers messing around in nuclear control systems primed for quick launch, and a general fear that nuclear weapons might not work if needed but could be launched by terrorists, is not good for anyone.

---

[12] Rumer, E, 2018, 'A Farewell to Arms… Control', Carnegie Endowment for International Peace, US-Russia Insight, April 17. Available from: https://carnegieendowment.org/2018/04/17/farewell-to-arms-.-.-.-control-pub-76088; Arbatov, A, 2016, 'An Unnoticed Crisis: The End of History for Nuclear Arms Control', Carnegie Moscow Center, March 16. Available from: http://carnegie.ru/2015/03/16/unnoticed-crisis-end-of-history-for-nuclear-arms-control-pub-59378

[13] See for example, 'Statement by the Euro-Atlantic Security Leadership Group, Support for Dialogue Among Governments to Address Cyber Threats to Nuclear Facilities, Strategic Warning and Nuclear Command and Control', 2018, February 16. Available from: https://www.europeanleadershipnetwork.org/wp-content/uploads/2018/02/Cyber-Statement-Feb-16-Final-Text.pdf

Council on Foreign and Defense Policy

RIAC
Russian International Affairs Council

MGIMO UNIVERSITY

HIGHER SCHOOL OF ECONOMICS
NATIONAL RESEARCH UNIVERSITY

ValdaiClub
ValdaiClub
ValdaiClub
valdai@valdaiclub.com