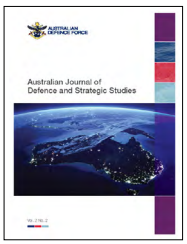

Australian Journal of Defence and Strategic Studies

Volume 2, Number 2 (2020)

ISSN 2652-3728 (PRINT) 2652-3736 (ONLINE)
https://www.defence.gov.au/ADC/Publications/AJDSS

Westmoreland's dream and Perrow's nightmare: two perspectives on the future of military command and control

Shane Halton

Published online: 3 December 2020

ADC Publications are produced by the Centre for Defence Research on behalf of the Australian Defence College

PO Box 7917 CANBERRA BC ACT 2610 Tel + 61 02 6266 0352 Email cdr.publications@defence.gov.au
Web www.defence.gov.au/adc/publications/ajdss

Westmoreland's dream and Perrow's nightmare: two perspectives on the future of military command and control

Shane Halton

The near simultaneous introduction of machine-learning technologies into the heart of traditional command and control arrangements coupled with the operational challenges inherent in executing complex missions, such as hypersonic missile defence, poses unique risks and opportunities to today's military commanders. This commentary explores this challenge from two perspectives. The first is the technological positivist perspective of US Army General William Westmoreland, which holds that military command and control functions can and should be automated to the highest degree possible to increase operational efficiency. The second is the more sceptical perspective of Dr Charles Perrow, which holds that interactively complex systems with tightly coupled components are inherently prone to unexpected and often dramatic failure. By incorporating both these perspectives into the design and operation of modern command and control systems, the author hopes these systems can be made to operate safely and more effectively.

In October 1969, standing behind a podium at the Sheraton Park Hotel in Washington DC, Army Chief of Staff General William C Westmoreland presented his vision of the future of warfare to the assembled attendees of the Annual Luncheon Association of the United States Army.

> On the battlefield of the future, enemy forces will be located, tracked, and targeted almost instantaneously through the use of data links, computer assisted intelligence evaluation, and automated fire control … I see battlefields or combat areas that are under 24-hour real or near real time surveillance of all types. I see battlefields on which we can destroy anything we locate through

> instant communications and the almost instantaneous application of highly lethal firepower.[1]

Westmoreland presented this vision, this dream, years before the US Department of Defense (DoD) embarked on its Second Offset Strategy, which was designed to leverage the US's superiority in science and technology to overcome the Soviet advantage in raw troop numbers in Europe, and decades before the US would first operationalise this approach to warfare during the first Gulf War. In his speech, Westmoreland was describing 'network-centric warfare' almost 30 years before the idea would gain broad acceptance in the Pentagon in the late 1990s.

In April 2017, the Pentagon established the Algorithmic Warfare Cross Functional Team, also known as Project Maven, to integrate:

> computer-vision algorithms needed to help military and civilian analysts encumbered by the sheer volume of full-motion video data that DoD collects every day in support of counterinsurgency and counterterrorism operations.[2]

Maven would later begin a second series of initiatives designed to bring not only Silicon Valley's technology but also its approach to developing and deploying software into the heart of the US military. Eventually the whole of Project Maven would be absorbed into the much larger Joint Artificial Intelligence Center, a new organisation with the express goal of bridging the gap between DoD and Silicon Valley. A close collaboration between the brightest minds in academia, the commercial world and national security, this too was Westmoreland's dream.

Though many facets of Westmoreland's dream have since come to pass, the late 1960s were in many ways a high-water mark for this brand of *technological positivism*, the practical philosophy that holds that almost any environmental, technological or social problem can be overcome if you throw enough resources, computing power and engineers at it. The 1970s and 1980s saw a fairly radical paradigm shift in thinking about complex adaptive systems, such as weather patterns, animal populations and human-machine hybrid organisations like air traffic control systems. In the mid-1970s, research in physics and mathematics by Benoit Mandelbrot, Mitchell Feigenbaum and others laid the groundwork for a new way of thinking about complexity, chaos and the basic nature of the

1 Randolph Nikutta, 'Artificial Intelligence and the Automated Tactical Battlefield' in Allan M. Dims (ed), Arms and Artificial Intelligence: Weapons and Arms Control Applications of Advanced Computing, Oxford University Press, Oxford, 1987, p 101.

2 Cheryl Pellerin, 'Project Maven Industry Day Pursues Artificial Intelligence for DoD Challenges', US Department of Defense, last modified 27 Oct. 2017. https://www.defense.gov/Explore/News/Article/Article/1356172/project-maven-industry-day-pursues artificial-intelligence-for-dod-challenges

universe. This vein of research – which eventually entered into mainstream culture with the popularisation of concepts such as fractals, 'sensitive dependence on initial conditions' and the 'butterfly effect' – set limits on what could be reliably known, modelled or predicted about the world at any given time. And, it placed hard limits on Westmoreland's techno-optimistic vision of the future. Engineers designing complex systems, and the technicians and managers responsible for operating them, began to gain a fuller appreciation for the many devious and difficult to predict ways glitches, friction, malfunctions, turbulence, poor design choices and interactive complexity could cause a system to underperform expectations or in certain cases fail all together.

One of the first researchers to incorporate the lessons from chaos and complexity research into the design and operation of complex systems was Charles Perrow. Perrow, in effect, made his career studying disasters. In 1984, he published *Normal Accidents: Living With High Risk Technologies,* which explored the root causes of industrial disasters, such as the partial meltdown of a nuclear reactor at Three Mile Island complex near Harrisburg, Pennsylvania. Perrow identified two factors which, when combined, increase the risk of a system failing catastrophically: tight coupling and interactive complexity. The 'normal' in normal accidents is a synonym for 'inevitable.' Normal accidents in a particular system may be rare ('it's is normal for us to die, but we only do it once') but the system's design and configuration make it more likely such accidents will occur. Perrow identifies systems at risk of normal accidents as 'high risk systems.'

Interestingly, Perrow released his book two years before the 1986 Soviet nuclear disaster at Chernobyl but it subsequently became the normal accident *par excellence,* providing students of industrial design with an easy shorthand to reference normal accident risk. Today, it is chilling to read Perrow's description of a normal accident knowing what happened in Chernobyl a mere two years later.

> We need two or more failures among components that interact in some unexpected way. No one dreamed that when X failed, Y would also be out of order and the two failures would interact so as to both start a fire and silence the fire alarm. Furthermore, no one can figure out the interaction at the time and thus know what to do. The problem is just something that never occurred to the designers… This interacting tendency is a characteristic of a system, not of a part or an operator; we will call it the "interactive complexity" of the system.
>
> …But suppose the system is also "tightly coupled" that is, processes happen very fast and can't be turned off, the failed parts cannot be isolated from other parts … operator action or the safety

> system might make it worse, since for a time it is not known what the problem really is.[3]

When the reactor crew at Chernobyl disabled the automatic shutdown mechanisms in preparation for a test and a previously undiscovered flaw in the control rod design caused hot nuclear fuel to rapidly mix with reactor cooling water which led to a rapid increase in pressure within the reactor, this was Perrow's nightmare.

Chernobyl isn't the only example from the late Soviet Union where an interactively complex and tightly coupled system catastrophically malfunctioned, causing near-instant death and destruction. In the early morning hours of 1 September 1983, Korean Air Lines Flight 007 (hereafter KAL007) departed Anchorage for Seoul. At the start of the flight, the flight crew made a fateful error; instead of selecting the Inertial Navigation System, which would have steered the plane on the proper route, the autopilot was instead set at a constant magnetic heading. This may have been caused by the failure to twist a knob one position further to the right. KAL007 drifted off course, unnoticed by the flight crew or any civilian air traffic controllers, eventually entering into Soviet air space near Kamchatka.

Ground-based Soviet air defence operators in the region had previously been tracking an American RC-135 spy plane (a converted Boeing 747) that had been tasked with observing a Russian missile test. The missile test was postponed and the RC-135 was told to return to base. As the RC-135 began its return trip to Alaska, Soviet air defence operators confused the two aircrafts' radar tracks and began tracking KAL007 as though it were the RC-135. Eventually, as KAL007 unknowingly moved closer to Russia, Soviet air defence operators scrambled three interceptor aircraft in order to visually identify the wayward aircraft and attempt to communicate directly with the aircrew and guide the trespassing plane down onto a Soviet airfield. However, once aloft none of the three interceptors were able to visually confirm whether the aircraft was an RC-135 or a civilian aircraft, nor were they able to make radio contact with KAL007's aircrew. At 3:25am local time, the pilot of one of the interceptor aircraft, an Su-15, was given the order to shoot down the non-responsive aircraft. He launched two air-to-air missiles which struck the KAL007 and caused it to crash into the sea, killing all on board, including a sitting member of the US House of Representatives.[4]

3 Charles Perrow, Normal Accidents: Living with High Risk Technologies - Updated Edition, Princeton University Press, Princeton, 2011, pp 4–5.

4 David Hoffman, The Dead Hand: The Untold Story of the Cold War Arms Race and its Dangerous Legacy, Random House Inc, New York, 2009, pp 50–52.

Perrow studied whether these types of air traffic control or air defence systems should be considered 'tightly coupled.' His conclusion was yes, though less so than automatic mechanical systems such as those found in nuclear reactors.

> Tight coupling reduces the ability to recover from small failures before they expand into large ones. Loose coupling allows recovery. Time constraints are tight; the (air traffic control) system is … moderately tightly coupled.[5]

Should the Soviet air defence system in this scenario be considered interactively complex? Certainly. The land based military air traffic controllers (ATCs) relied on a combination of radars and interceptor aircraft to gather information on what was happening in the air. These inputs would be delivered to the ATC in a variety of different formats – radar tracks appearing on a screen, interceptor updates relayed via radio – and the ATCs had to convert them into a workable approximation of reality in their heads.

The Soviet air defence example above highlights the essentially dualist nature of modern military command and control; it is a mission, something commanders do, but it is also 'a thing' – a set of modern communication technologies without which it would be impossible for the commander to do anything. Exercising command and control (C2) therefore is as much about aligning responsibilities and functions within a command hierarchy as it is about utilising digital technologies to gather information about one's operating environment and to maintain clear lines of communication and feedback between the different nodes within the chain of command. This challenge – to construct and maintain a robust, resilient information architecture that can keep everyone informed and 'in the loop' about what's happening in the battlespace – gets more difficult, perhaps exponentially so, as we get closer to achieving Westmoreland's dream.

The battlespace of the 2020s is one in which the United States, its allies and its competitors will field hypersonic munitions, robust offensive cyber and electronic attack capabilities, as well as autonomous lethal weapons systems. To direct these forces quickly and effectively militaries across the world are investing in modernising their C2 systems and associated intelligence, surveillance, and reconnaissance (ISR) capabilities. Military leaders in the US are openly discussing when and how they should integrate machine-learning systems into kill chains. At first glance, all this seems to be the ultimate expression of Westmoreland's dream, a military technopia where cutting edge Made-in-America science and technology relieve command staffs

5 Perrow, Normal Accidents, pp 4–5.

of the grunt work of running the war and allow commanders to focus on their real passion – *strategy*.

However, the battlespace of the 2020s will also be an interactively complex and increasing tightly coupled affair – a hugely scaled up version of Perrow's nightmare. To understand why this trend towards tight coupling is accelerating let us consider military C2 at the most basic, functional level.

To command, a commander must first be able to perceive their operating environment, make decisions about it and finally pass orders back to their subordinates. That is the bare minimum. Today there are tools to assist commanders and their staffs in these tasks such as intelligence satellites, classified networks and information technologies, for example the Windows Office suite. However, the technological state of the art in 2020 poses unique challenges to command as well. The speed of modern weaponry, such as ballistic or hypersonic missiles and cyber attacks, reduces human response time. There is also the uncomfortable fact that many of these weapons, specifically cyber attacks, are optimised to attack command structures directly instead of deployed units i.e. why waste time and resources wiping out an army in the field when you can remotely destroy command headquarters and throw the army into disarray?

These challenges have led military commanders to seek out automated solutions to speed up the different command functions. During the Second World War, few C2 functions would be considered tightly coupled in the modern sense. They were based around humans sharing information with one another and humans inherently lack the ability to transfer huge amounts of complex information quickly. We can only absorb and retain so much, and pay attention for so long. The 'Information Age' (roughly 1970 to 2010) saw the integration of machine-to-human information transfers across military command structures, mostly in the form of classified networks and desktop computers. This changed the calculus, as at least one component in the equation (the machine) could pass large amounts of information instantaneously. Humans, however, still needed time to absorb information and make sense of it. This has kept most processes slow enough to be managed effectively. Military operations in the 2020s, by contrast, will be defined, in part, by increasing reliance on instantaneous machine-to-machine connections to support different command functions, reducing or removing the human component entirely for the sake of speed and efficiency.

How does this look in practice? Consider the evolution from a Second World War scout plane to a modern unmanned aerial system (UAS). A scout plane would report back what it was seeing – ideally via radio – to the command staff. In many cases however, radio was not a viable option (it could be broken or the

pilot's plane could be out of radio range) so the pilot would have to land back at base first and be physically debriefed about what they had seen during their flight. This introduced at least two information transfer challenges, the first was that the information was time-late and the second was that the commander did not see exactly what the pilot saw. Instead, the commander received a report of what the pilot thought they saw. The pilot and the commander were forced to construct a common mental picture of events via dialogue, based on the pilot's recollection.

Today, the drone pilot 'sees' what the UAS sees, and sees it instantaneously – even if they are half the world away from one another. This connection is still subject to constraints, however. The video feed from the UAS is bandwidth heavy, which requires its ground station to be equipped with special gear to receive the feed. The feed can also be disrupted which, in many cases but not all, forces the UAS to land, effectively ending its mission. There is also the requirement that a human being constantly watch the video feed to generate a report for the commander.

Advances in machine-vision technology are such that it is now possible to pre-program a UAS so that it is able to see and understand the environment it is operating in (i.e. identify the difference between different types of buildings and vehicles, read licence plates, figure out if a person is holding a weapon, etc.) using software installed directly on the UAS. In this scenario, there would then be no need to pass a constant, bandwidth-intense video feed back to a human operator. Instead, the UAS could fly on autopilot, collect all the information it needed to and send that information directly into a battle management network via bit-sized chunks of text data so that the commander's picture of the world could be updated instantly.

A UAS configured in this manner could also pass that information to a second, third or fourth UAS thereby allowing multiple units to automatically share information about the battlespace without the need for a human to facilitate that sharing. Different UAS could be outfitted with different types of sensors, one UAS collects imagery while another collects electronic signals intelligence (SIGINT). Algorithms on board each UAS could merge this information via multi-sensor fusion so that each UAS had a layered, complex picture of the battlespace. Some UAS could be equipped with weapons so that they could automatically utilise this robust picture to deliver effects on the battlespace. This of course is what Westmoreland meant when he dreamed 'we can destroy anything we locate through instant communications and the almost instantaneous application of highly lethal firepower.' The only difference is that, at this stage in technological development, human action is no longer required beyond the mission planning stage.

While this proliferation of artificial intelligence throughout C2 structures can positively impact the ability of commanders to perceive and understand their operating environments, it also becomes a major driver of tight coupling, highlighting the possibility that realising Westmoreland's dream risks simultaneously birthing Perrow's nightmare. However, before pressing this point any further it is important to differentiate between two different types of artificial intelligence, the latter of which is a more serious risk driver of interactive complexity in C2 systems.

Expert systems are attempts to reproduce human decision-making in mechanical form. The 'decision trees' that form the backbone of expert systems are based on the types of if/then propositions a human mind goes through when completing a complex task (i.e. if the ball is red then put it in the bucket, if not drop the ball on the floor). Most military systems that incorporate artificial intelligence today (such as the US Navy's AEGIS combat system) are expert systems.

Machine-learning (ML) algorithms are not generally concerned with replicating human thought patterns, they just want to find the 'right' answer. ML algorithms are fed a data and then instructed to complete tasks. If they complete the task successfully, they are rewarded, if not they are punished. Over time these systems can become very good at completing tasks but the 'thought patterns' that led them to the right answer over and over again are often completely foreign to human beings.

In the last two decade, systems built around ML have displaced expert systems as the artificial intelligence approach of choice in the commercial world. It is cheaper and easier to generate a solution that simply works than to spend time trying to replicate human behaviour and thought patterns. Most language translation programs are based around a type of ML, as are most visual recognition technologies and fraud detection systems. However, ML continues to present challenges in human endeavours where the stakes are literally life and death, such as military options.

This is because the information that feeds a military commander's decision-making process is should be *traceable, verifiable* and *intelligible* (though the 'fog of war' ensures that is rarely the case in practice). Verifiability is the ability to ascertain whether or not information is correct. Traceability is the ability to identify where information came from (Which UAS detected this?). Intelligibility is the ability to understand the thought process that led to a decision (Why did the UAS classify this wi-fi signal as a cell phone?). The challenge of obtaining reliable, traceable information will be exacerbated as ISR sensors based on ML proliferate throughout the world's militaries. This will in turn drive normal accident risk.

A single ML-based ISR sensor in a broader information architecture (one sensor on a single UAS for example) is unlikely to be a significant driver of normal accident risk by itself. Intelligibility may be a challenge as a human being may never know exactly 'why' a specific sensor is providing erroneous data but it should be easy enough to trace the error to that the one sensor that is known to have a mind of its own. Verifiability too may be easier than it initially appears because when ML-based systems fail they tend to do so in unexpected and occasionally dramatic ways that do not mimic human failure modes.

The challenge is when multiple ML-based sensors are linked to one another within an architecture to facilitate multi-sensor fusion, as in the UAS example above. While the goal of the fusion process is to instantly provide a detailed, multi-faceted picture of the operating environment to the commander, it also creates a complex mini-system with several tightly coupled components.

The question becomes: if ML-based multi-sensor fusion introduces so much normal accident risk into a C2 system, why would a military commander ever choose to rely on it? There are several understandable and perhaps justifiable, if not altogether comforting, reasons why this might be the case. It might be because:

- the system's designers or operators have insufficient understanding of normal accident risk
- there is no one 'designer' of the C2 system, as multiple designers contribute components that, when combined, create a system of systems with a high level of risk
- the system never 'failed' in testing so the risk has never been identified
- the commander inherited a C2 structure dependent on multi-sensor fusion and is not aware of it
- designers, operators and military commanders are aware of the risk but feel they need to rely on the system to accomplish their mission.

The last reason – that the risk is known but it is balanced against the advantage offered by multi-sensor fusion – is worth further consideration. It points to a broader challenge for military commanders in the 2020s; critical modern warfighting functions, such as defensive cyber operations and high-speed frequency hopping to avoid communication jamming, occur faster than human perception and therefore *must* be automated to a high degree. Commanders, consciously or not, will be forced to make trade-offs between possessing a battlespace awareness based on verifiable, traceable and intelligible information on the one hand and operational speed and efficiency on the other.

Recent work by Olivia Garand and B.A. Friedman has explored how modern information technologies are driving commanders towards over-centralisation and depriving subordinates of the ability to exercise 'mission command' (effectively 'command at the lowest possible level'). They note that this is particularly dangerous in a world where lower echelon units may be cut off from higher headquarters and be forced to act on their own.[6]

While this assessment is broadly accurate (and should be adopted where applicable), re-emphasising the advantages of mission command in military operations is not a panacea for the challenges of C2 in the age of modern era. Certain missions, like hypersonic missile defence, manoeuvrer warfare in the electromagnetic space or the synchronising cyber and physical attacks in real time, rely on the ability of commanders to coordinate the activity of multiple actors spread out over time and space. In cases like these, subordinate units will very likely be subjected to a very high degree of centralised C2.

In closing, the military technologies coming fully online in the 2020s (hypersonics, cyber and electromagnetic warfare) are so fast that in many cases they prevent human operators from acting 'in the loop'. These capabilities will therefore be reliant on the use of non-human intelligent agents, likely powered by ML, to coordinate effectively. The full effects of the increased use of these intelligent agents across C2 structures is unknown at this stage but Perrow's research shows us that when it comes to interactively complex, tightly coupled systems, systemic failure is a question of when, not if. Militaries the world over are engaging in a titanic struggle to build the largest, most capable and wide-ranging battle management systems they can while defending against adversary cyber and physical attacks designed to directly target the heart of those systems. Additionally, mitigations will need to be put in place to ameliorate the normal accident potential inherent in the systems themselves. Exercising effective command and control in the modern era will therefore be a delicate balancing act, poised between the *yin* of Westmoreland's dream and the *yang* of Perrow's nightmare. Both perspectives will need to be considered and constantly revisited if we are to successfully navigate this challenge.

6 BA Friedman and Olivia Garand, 'Technology-Enabled Mission Command', *War on the Rocks*, last modified 09 April 2020. https://warontherocks.com/2020/04/technology-enabled-mission-command-keeping-up-with-the-john-paul-joneses/