



NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

AIRCRAFT CYBER COMBAT SURVIVABILITY

by

Austin K. Weinman

June 2020

Thesis Advisor:
Second Reader:

Christopher A. Adams
Jarema M. Didoszak

Approved for public release. Distribution is unlimited.

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			<i>Form Approved OMB No. 0704-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE June 2020	3. REPORT TYPE AND DATES COVERED Master's thesis		
4. TITLE AND SUBTITLE AIRCRAFT CYBER COMBAT SURVIVABILITY			5. FUNDING NUMBERS	
6. AUTHOR(S) Austin K. Weinman				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A			10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release. Distribution is unlimited.			12b. DISTRIBUTION CODE A	
13. ABSTRACT (maximum 200 words) The aircraft combat survivability (ACS) design discipline has proven effective in producing survivable combat aircraft for over fifty years. Currently, the discipline only focuses on kinetic threats; however, an emerging class of cyber weapons has brought forth a new challenge in the endless fight between attackers and defenders. Cyber is a legitimate anti-aircraft threat, and the recent rise in cyber-related incidents raises major concern for our military and its ability to carry out mission objectives. While the attack vectors and damage mechanisms of cyber weapons are fundamentally different from those of traditional kinetic threats, modifying the fundamental ACS concepts can help produce cyber-survivable combat aircraft. This research lays the groundwork for expanding the ACS design discipline fundamentals to include emerging anti-aircraft cyber threats. In this new aircraft cyber combat survivability (ACCS) design discipline, ACS terms are redefined to address cyber threats and a new cyber kill chain is proposed to help assess an aircraft's cyber-survivability. The development of 12 survivability enhancement concepts aims to assist program managers and engineers in designing platforms that are better equipped to survive in a hostile cyber environment. An approach to modeling cyber-attacks using the ACCS probabilistic kill chain shows how to assess an aircraft's cyber-survivability and demonstrates the effectiveness of survivability enhancement features.				
14. SUBJECT TERMS aircraft combat survivability, ACS, cyber, survivability, vulnerability, susceptibility, cybersecurity, cyberspace, damage mechanism, threat, Monte Carlo simulation, aircraft cyber combat survivability, ACCS			15. NUMBER OF PAGES 129	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UU	

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release. Distribution is unlimited.

AIRCRAFT CYBER COMBAT SURVIVABILITY

Austin K. Weinman
Ensign, United States Navy
BS, United States Naval Academy, 2019

Submitted in partial fulfillment of the
requirements for the degree of

**MASTER OF SCIENCE IN ENGINEERING SCIENCE
(AEROSPACE ENGINEERING)**

from the

**NAVAL POSTGRADUATE SCHOOL
June 2020**

Approved by: Christopher A. Adams
Advisor

Jarema M. Didoszak
Second Reader

Garth V. Hobson
Chair, Department of Mechanical and Aerospace Engineering

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

The aircraft combat survivability (ACS) design discipline has proven effective in producing survivable combat aircraft for over fifty years. Currently, the discipline only focuses on kinetic threats; however, an emerging class of cyber weapons has brought forth a new challenge in the endless fight between attackers and defenders. Cyber is a legitimate anti-aircraft threat, and the recent rise in cyber-related incidents raises major concern for our military and its ability to carry out mission objectives. While the attack vectors and damage mechanisms of cyber weapons are fundamentally different from those of traditional kinetic threats, modifying the fundamental ACS concepts can help produce cyber-survivable combat aircraft. This research lays the groundwork for expanding the ACS design discipline fundamentals to include emerging anti-aircraft cyber threats. In this new aircraft cyber combat survivability (ACCS) design discipline, ACS terms are redefined to address cyber threats and a new cyber kill chain is proposed to help assess an aircraft's cyber-survivability. The development of 12 survivability enhancement concepts aims to assist program managers and engineers in designing platforms that are better equipped to survive in a hostile cyber environment. An approach to modeling cyber-attacks using the ACCS probabilistic kill chain shows how to assess an aircraft's cyber-survivability and demonstrates the effectiveness of survivability enhancement features.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I.	INTRODUCTION	1
A.	DEVELOPMENT OF THE AIRCRAFT COMBAT SURVIVABILITY DESIGN DISCIPLINE.....	2
B.	EVOLUTION OF CYBER CONFLICT AND EXPANDING THE ACS DESIGN DISCIPLINE	4
II.	OVERVIEW OF AIRCRAFT COMBAT SURVIVABILITY DESIGN DISCIPLINE.....	9
A.	THE FUNDAMENTALS OF AIRCRAFT COMBAT SURVIVABILITY.....	9
B.	HOW SURVIVABILITY IS MEASURED: THE PROBABILISTIC KILL CHAIN.....	11
C.	SURVIVABILITY ENHANCEMENT	14
1.	Susceptibility Reduction	14
2.	Vulnerability Reduction	17
D.	WHY DO WE NEED AIRCRAFT COMBAT SURVIVABILITY?	19
E.	THREATS TO COMBAT AIRCRAFT	21
1.	Kinetic Energy Weapons	22
2.	Electronic Warfare Weapons	22
3.	Unconventional Weapons	23
4.	Cyber Weapons.....	24
III.	A NEW THREAT IN CYBER	25
A.	FUNDAMENTALS OF CYBER	25
1.	Cyberspace and Cybersecurity	25
2.	Categorizing DoD Cyber Systems	26
B.	PROSPECT OF CYBER AS AN ANTI-AIRCRAFT WEAPON.....	28
C.	APPEAL AND IMPACT OF CYBERWARFARE.....	29
D.	HOW AN ANTI-AIRCRAFT CYBER-ATTACK MANIFESTS ITSELF	31
IV.	DEVELOPING AIRCRAFT CYBER COMBAT SURVIVABILITY (ACCS) FROM ACS FUNDAMENTALS	35
A.	KEY DIFFERENCES BETWEEN ACS AND ACCS	35
B.	DEFINING KEY TERMS	37
C.	CYBER PROBABILISTIC KILL CHAIN	39

D.	DYSFUNCTION MECHANISM AND POTENTIAL ATTACK VECTORS OF A CYBER WEAPON	45
V.	DESIGNING FOR THE CYBER THREAT USING AIRCRAFT COMBAT SURVIVABILITY FUNDAMENTALS	49
A.	SUSCEPTIBILITY REDUCTION CONCEPTS FOR ACCS	51
1.	Situational Awareness	51
2.	Signature Reduction/Management.....	53
3.	Cybersecurity Hardening	54
4.	Deception and Decoys.....	56
5.	Threat Suppression.....	57
6.	Tactics, System Performance, and Crew Training and Proficiency	58
B.	VULNERABILITY REDUCTION CONCEPTS FOR ACCS	59
1.	System Redundancy with Diversity	59
2.	Component Location and Logical Separation	61
3.	Component Elimination or Replacement.....	62
4.	Component Shielding	63
5.	Dysfunction Suppression	64
6.	Recovery.....	65
VI.	SIMULATING A CYBER-ATTACK USING THE PROBABILISTIC KILL CHAIN.....	67
A.	CHALLENGES TO MODELING A CYBER-ATTACK.....	67
B.	METHODOLOGY.....	68
C.	SIMULATION TEST SCENARIOS	74
D.	RESULTS AND DISCUSSION	75
1.	Test Scenario One: No Survivability Enhancement Features.....	75
2.	Test Scenario Two: One Survivability Enhancement Feature	78
3.	Test Scenario Three: All Survivability Enhancement Features.....	82
VII.	CONCLUSIONS.....	85
	APPENDIX A. SME PROBABILITY DATA	87
	APPENDIX B. MONTE CARLO SIMULATION MATLAB SCRIPT	93
	LIST OF REFERENCES.....	103

INITIAL DISTRIBUTION LIST	109
--	------------

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF FIGURES

Figure 1.	The global risk landscape. Source: World Economic Forum (2019).	5
Figure 2.	ACS probabilistic kill chain. Source: Ball (2003).	12
Figure 3.	Historical losses and loss rates. Source: Ball (2003).	20
Figure 4.	Categorizing threats.....	21
Figure 5.	Lockheed Martin’s cyber kill chain. Source: Lockheed Martin (n.d.).	34
Figure 6.	ACCS probabilistic kill chain. Adapted from Ball & Bryant (2020b).	41
Figure 7.	Communication channels attack surface. Source: Ball and Bryant (2020b).	47
Figure 8.	Triangular distribution.....	69
Figure 9.	Launch phase triangular distribution.....	72
Figure 10.	Triangular probability distributions for each phase	75
Figure 11.	Simulating a cyber-attack on an aircraft without SEFs.....	76
Figure 12.	Probability of survival distribution without SEFs.....	77
Figure 13.	Comparing single SEF effectiveness.....	79
Figure 14.	Probability of survival distribution with one SEF (disproportional)	80
Figure 15.	Probability of survival distribution with one SEF (proportional)	82
Figure 16.	Event outcomes for cases with and without SEFs	83
Figure 17.	Probability of survival distributions for cases with and without SEFs	83

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

Table 1.	ACS and ACCS definitions.	37
Table 2.	ACS versus ACCS survivability enhancement concepts	50
Table 3.	Sample SME probability data for launch phase.....	71
Table 4.	Sample SME data for active phase.....	87
Table 5.	Sample SME data for detect phase.....	88
Table 6.	Sample SME data for launch phase	89
Table 7.	Sample SME data for implant phase	90
Table 8.	Sample SME data for hit phase.....	91
Table 9.	Sample SME data for kill phase.....	92

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS AND ABBREVIATIONS

ACS	Aircraft Combat Survivability
ACCS	Aircraft Cyber Combat Survivability
AIAA	American Institute of Aeronautics and Astronautics
AP	armor piercing
AP-I	armor piercing incendiary
APT	advanced persistent threat
CI	confidence interval
CIWS	close-in weapon system
DoD	Department of Defense
EFB	electronic flight bag
EM	electromagnetic
EMS	electromagnetic spectrum
HE	high explosive
HPM	high-power radio frequency/microwave
IADS	integrated air-defense system
IR	infrared
IT	information technology
JASSM	Joint Air-to-Surface Standoff Missile
JCS	Joint Chiefs of Staff
JDAM	Joint Direct Attack Munitions
KEW	kinetic energy weapon
LFT&E	live fire test and evaluation
MMHE	man-made hostile environment
NIPRNET	Non-Secure Internet Protocol Router network
OT	operational technology
PLC	programmable logic controller
RPG	rocket-powered grenade
SCADA	supervisory control and data acquisition
SEF	survivability enhancement feature
SIPRNET	Secure Internet Protocol Router network

SME	subject matter expert
UAV	unmanned aerial vehicle
WEF	World Economic Forum
WMD	weapons of mass destruction

ACKNOWLEDGMENTS

I would like to give my sincerest thanks and appreciation to the individuals named below. None of this would have been possible without their kind help and support.

Thank you to the faculty and staff of Naval Postgraduate School for making all of this possible. Great work is being done here that can change the world. I would especially like to thank my advisor, Christopher Adams, for his insight, support, and direction on this thesis. I know that this was not an easy problem to tackle, but I appreciate your trust and confidence in me and my work.

I would also like to thank Dr. Bob Ball and Dr. Bill Bryant for their vision, guidance, and expertise. I would have been lost without their help, and I cannot thank them enough for letting me be a part this endeavor.

My greatest heartfelt thanks go out to my family and friends who have supported me in this journey. Those I am closest to have seen all of my ups and downs, and have been with me through them all. To my parents, thank you for always being there for me and for raising me to be the person I am today. To my brother, thank you for all the memories and laughs throughout the years. I constantly look up to you and could not be more proud of the man you are becoming. And to my departed grandmother, I wish I could continue to share my life's moments with you. I would not be where I am without you. You are deeply missed.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

The primary mission of our nation's armed forces is to support and defend the Constitution of the United States against all enemies, foreign and domestic. One of the first steps toward ensuring that this mission is met is ensuring that combat platforms are able to survive in man-made, hostile environments. Adversaries are more combat-capable now than they have ever been. It is therefore critical that today's combat systems are not only mission capable, but resistant to the conventional and emerging threats that make up today's warfighting landscape.

One part of warfighting is achieving our objectives; the other part is denying the enemy theirs. Today's adversaries are constantly scheming ways to prevent our forces from carrying out our missions successfully. For the past century, adversaries have been getting better at firing more powerful, more accurate, and more dangerous weapons at our combat platforms. And for the past century, project managers and engineers have been designing combat platforms to better survive attacks from such weapons.

Hundreds of thousands of aircraft have been lost in combat since the early twentieth century. To address these losses, the Aircraft Combat Survivability (ACS) design discipline was formally established to protect aircraft operating in man-made hostile environments (Ball, 2003). The current ACS discipline was born from historical data and deals almost entirely with protecting aircraft against kinetic energy anti-aircraft weapons (guns and guided missiles). Until relatively recently, kinetic energy weapons (KEWs) have been the only significant threat to combat aircraft.

However, the rapid growth of technology in the past few decades, especially in the cyber domain, has given adversaries new tools to disrupt, degrade, deny, and destroy our combat systems. Our fighting force is becoming increasingly reliant on systems that operate in cyberspace as more capable, complex systems are developed and introduced into our nation's military. Every line of flight hardware, every small maintenance software update, and every bit of data transferred from mission planning computers to onboard smart weapons exposes our combat aircraft to potential cyber-based attacks.

The fundamentals of the ACS discipline were specifically designed with kinetic threats in mind, but these concepts can be more broadly applied to address new, emerging threats in cyber. While the attack vectors and damage mechanisms associated with cyber threats are fundamentally different than those associated with traditional, kinetic energy weapons, designing cyber-survivable systems can be done by applying the same ACS fundamentals to emerging cyber threats.

The ultimate goal of this thesis is to provide program managers and engineers tools for designing combat platforms capable of avoiding or withstanding dysfunctions caused by anti-aircraft cyber weapons. The rest of this chapter will discuss why the ACS design discipline was developed and why there is a need to expand the design discipline fundamentals to evolving cyber threats. Chapter II will provide an overview of the ACS fundamentals and how they are used to calculate and enhance survivability as it relates to kinetic attacks. Chapter III briefly defines the cyber domain and discusses why and how cyber could be used as an anti-aircraft weapon. Chapter IV applies ACS fundamentals to cyber threats. Chapter V introduces cyber survivability enhancement concepts that can be used to design combat platforms that are less susceptible and less vulnerable to cyber-attacks. Chapter VI models a cyber-attack using the probabilistic kill chain introduced in Chapter IV and shows how enhancement concepts can improve aircraft survivability. Chapter VII highlights conclusions and identifies further areas of research.

A. DEVELOPMENT OF THE AIRCRAFT COMBAT SURVIVABILITY DESIGN DISCIPLINE

Combat aircraft have been flying missions since the early twentieth century when the Wright Military Flyer was first used by the U.S. Army in 1909 (Maksel, 2010). World War I introduced the first widespread use of aerial combat and officially made the skies a warfighting battlespace. At the time, pilots would fly at altitudes beyond the range of enemy ground defenses, would sit on metal lids to offer additional shielding, and would carry small firearms to suppress enemy threats (Ball, 2003). The idea of designing and operating aircraft that could survive in a hostile, combat environment was already in practice by the end of World War I in 1918.

It was not until the 1970s, however, that the survivability of combat aircraft became a formal design discipline. The Aircraft Combat Survivability design discipline came as a response to the more than 4,000 aircraft lost in the 1964-1973 Southeast Asia conflict (Ball, 2003). In 1971, the Department of Defense (DoD) established the Joint Technical Coordination Group on Aircraft Survivability, whose mission was to protect aircraft from anti-aircraft guns and guided missiles. In 1985, the American Institute of Aeronautics and Astronautics (AIAA) published Dr. Bob Ball's textbook, *The Fundamentals of Aircraft Combat Survivability Analysis and Design*, officially introducing ACS as a design discipline.

Over the last fifty years, this design discipline has grown and has been successfully implemented for a number of combat aircraft platforms. The number of aircraft combat losses during that time span has dropped significantly (Ball, 2003). The ACS discipline is now an integral part of the DoD acquisition process, survivability requirements are commonplace in project management, and testing and evaluation of combat systems is congressionally mandated (Ball & Bryant, 2020a). It is important to note that, while the ACS discipline is specific to aircraft, the fundamentals of survivability can be applied to any combat or even non-combative platform. Ships and ground vehicles must also be able to survive in hostile, combat environments; while some of the examples that Dr. Ball uses may not be entirely applicable to every type of combat platform, the broader survivability concepts can be very useful.

As combat aircraft have evolved drastically over the last fifty years, so too have the threats to combat aircraft. In 1971, when the ACS design discipline was first imagined, unconventional weapons (biological, nuclear, and chemical weapons), electromagnetic weapons, and cyber-weapons were never considered realistic threats. Design decisions were not made with these threats in-mind. Things have changed since 1971 and these threats cannot be ignored. Attacks on combat platforms in the cyber domain pose a real and credible threat that should be considered when designing and operating these platforms.

B. EVOLUTION OF CYBER CONFLICT AND EXPANDING THE ACS DESIGN DISCIPLINE

The world we live in, and the world that our military operates in, is increasingly dependent on networks of systems that run on a series of 1's and 0's. Unfortunately, many of these cyber-systems were not designed with combat survivability in mind; they were designed to operate in an uncontested, permissive environment (Bryant, 2016). Every year, state-based cyber power and cyber weaponry are growing more advanced and the potential damage from a state-based cyber-attack is becoming more severe (Maness, 2020b).

In 2019, the World Economic Forum (WEF) released a report based on surveys of 750 experts and decision makers (Figure 1). The report found cyber-attacks as one of the most likely and one of the most catastrophic risks to the global landscape, just behind natural disasters which kill about 90,000 people every year according to the World Health Organization (2012).




Figure 1. The global risk landscape. Source: World Economic Forum (2019).

The U.S. DoD has started to address the potential threat that cyber poses. In 2018, Secretary of Defense Patrick M. Shanahan stated at the Air Force Association’s Annual Conference that, “Cybersecurity is, probably going to be what we call the ‘fourth critical measurement.’ We’ve got quality, cost, schedule, but security is one of those measures that we need to hold people accountable for” (Mehta, 2018, para. 3). In 2020, the DoD mandated that Program Managers safeguard their systems from cyber-attacks, design for the cyber threat environment, and manage cybersecurity impacts (Department of Defense).

In just the last decade, U.S. cybersecurity spending has almost doubled, from a reported twenty-nine billion dollars in 2010 to sixty-five billion dollars reported in 2018 (Ashworth, 2020). During that same time frame however, the number of cyber incidents shot up dramatically as well, from less than half a million in 2010 to nearly sixty million in 2018. It is clear that the DoD believes that cyber is a legitimate and powerful threat to our fighting force, but, while some progress is being made to address this emerging risk, more still needs to be done.

New technologies have raised issues, and continue to raise issues, in international civil-military relations. For example, most people had no idea the extent to which the role of nuclear weapons and nuclear power was going to play in international affairs prior to World War II because nuclear power was never previously regarded as a legitimate threat. It was not until the path to chaos and destruction was made clear that nations started to focus on developing, regulating, and safeguarding access to nuclear energy. Cyber seems as if it may be following a similar path. Nations will not take cyber-threats seriously until the path to chaos and destruction in the cyber domain is made clear (Maness, 2020a).

The potential impact of an all-out cyber war (a war in which “unauthorized actions by a nation-state to penetrate another nation’s computers or networks for the purposes of causing damage, disruption, or theft of classified, restricted, or proprietary information or materials” (Jackson, 2012, p. 304) is the goal) could be catastrophic. It may just be a matter of time before the full effects of cyber warfare are apparent. With the millions of pieces of hardware and software components needed to operate technologically complex combat platforms worldwide, the attack surface is so large and the rate of technological development is so fast that it is hard to imagine a scenario where large-scale cyber conflict is not possible (Kollars, 2018).

Ben Buchanan (2016), an expert in the field of international cyber conflict, believed that, given the level of complexity of modern cyber-systems, determined and technically-competent hackers would always find a way to infiltrate our systems. In his professional opinion, cyber defense cannot be assured because “adept programmers will [inevitably] find vulnerabilities and overcome security measures” (p. 108) It is therefore imperative not only to defend our cyber-systems to prevent an adversary from gaining access, but also to

design systems capable of surviving those situations in which an adversary does successfully gain access. Our systems should be designed not only to avoid, but also to withstand a cyber-based attack from an advanced persistent threat (APT).

Attacks like Stuxnet demonstrate how a cyber-attack can covertly have physical effects on protected hardware. Stuxnet was a malicious computer worm launched in 2010 that destroyed nearly one-fifth of Iran's nuclear centrifuges (Cyber Security Forum Initiative [CSFI], 2010). The computer worm specifically targeted the supervisory control and data acquisition (SCADA) systems, which allowed the attacker to access the programmable logic controllers (PLCs) that controlled the centrifuges. The worm's malware then modified the industrial controllers, causing the centrifuges to spin so fast that they broke. The malware was also designed to make the centrifuge's computer monitors appear as if everything were operating normally so that the plant operators remained in the dark. The worm was self-replicating and infected over 200,000 computers in the process of reaching the Iranian nuclear plant (CSFI, 2010). While the cyber weapon is believed to have been developed by the United States and Israel, both nations continue to deny their involvement.

While it is true that cyber-attacks have historically been focused on information technology systems and have had limited physical effects, the possibility of a cyber-attack on a combat aircraft's software or hardware should not be dismissed. Just like the Stuxnet worm targeted a specific piece of hardware, adversaries may be able to target specific aircraft cyber-systems if those systems are not well protected.

The rapid evolution of cyber conflict and the current global landscape suggest that cyber is a real and growing threat that should not be ignored. While no aircraft has been successfully hit by a cyber-attack (any cyber incident with a cyber-attack would likely be highly classified), the possibility that such an attack is possible should be taken seriously and combat platforms should be designed with cyber survivability in mind. There is a need to expand the current ACS design discipline to include the emerging cyber threat to combat aircraft. The survivability design discipline needs to evolve as the threats that combat aircraft face evolve.

THIS PAGE INTENTIONALLY LEFT BLANK

II. OVERVIEW OF AIRCRAFT COMBAT SURVIVABILITY DESIGN DISCIPLINE

A. THE FUNDAMENTALS OF AIRCRAFT COMBAT SURVIVABILITY

Aircraft combat survivability (ACS), as defined by Dr. Robert Ball in his textbook *The Fundamentals of Aircraft Combat Survivability Analysis and Design* (2003), is “the capability of an aircraft to avoid or withstand a man-made hostile environment (MMHE)” (p. 1). The survivability of an aircraft is a function of its susceptibility and vulnerability. In the ACS discipline, susceptibility is defined as the inability of an aircraft to avoid the threat elements of the MMHE (Ball, 2003). These elements include the guns, missiles, exploding warheads, air interceptors, lasers, radars, and all other elements associated with an enemy’s air defense. The more likely a combat aircraft is to be hit by a weapon, the more susceptible the aircraft is. The active threat’s characteristics, the aircraft’s detectable signatures, the defense countermeasures, the aircraft and crew performance, and the combat scenario will influence the susceptibility of an aircraft.

Vulnerability is defined as the inability of an aircraft to withstand the damage caused by the threat elements of the MMHE (Ball, 2003). The more likely that a hit results in an aircraft kill, the more vulnerable the aircraft is. The types of damage-causing mechanisms, the number of warheads that hit the aircraft, the design of the aircraft, and the location and number of critical components on an aircraft influence the aircraft’s vulnerability. Dr. Ball describes critical components as the parts of an aircraft that are either mission-essential or flight-essential (Ball, 2003). Killing a critical component will ultimately result in an aircraft kill. Aircraft are vulnerable because their critical components are vulnerable (Adams, 2019a). The goal of the ACS discipline is to enhance the survivability of combat aircraft by reducing its susceptibility and vulnerability. Susceptibility reduction and vulnerability reduction concepts for KEWs are discussed later in this chapter.

The opposite of survivability is killability. In ACS, killability is the ease at which an aircraft is killed and is the product of susceptibility and vulnerability (Ball, 2003). The survivability of an aircraft is enhanced when the killability is reduced. Killability is reduced

when susceptibility or vulnerability are reduced. In ACS, an aircraft kill can result from either the physical loss of an aircraft (attrition kill) or if that aircraft, due to the loss of critical components, is unable to perform its mission (mission kill). Both flight essential functions, functions needed to sustain controlled flight like lift, thrust, control, and structural integrity, and mission essential functions, functions needed to successfully carry out the intended mission, are important to protect in a MMHE. An enemy's attack is successful if it can take a target aircraft either out of flight or out of the fight.

The type and extent of damage done on an aircraft from an enemy warhead depends on the KEW's damage mechanisms. In ACS, a warhead's damage mechanism is the "physical entity that causes damage to the aircraft" (Ball, 2003, p. 280). The damage mechanisms for high explosive warheads and penetrator type weapons are typically metallic penetrators and fragments, incendiary materials, and blasts. Warheads can, and often do, utilize more than one damage mechanism. Dr. Ball uses the term damage process to describe how the damage mechanisms and the aircraft's physical components interact (2003). In ACS, the terminal effect refers to the physical damage state of an aircraft's components affected by the damage process (Adams, 2019b). Another important and useful term for defining how damage mechanisms may lead to an aircraft kill is kill mode. If an enemy warhead hits an aircraft, the kill mode describes the physical response of the aircraft's components that result in those components becoming inoperable (Ball, 2003). If the killed component happens to be a critical component, the component kill will result in an aircraft kill due to the loss of an essential function.

To clarify the differences between the terms damage mechanism, damage process, terminal effect, and kill mode, let us consider a scenario in which an incendiary round hits a wing's fuel tank. The damage mechanism in this scenario would be the incendiary material in the round. The damage process would be the physical combustion in the wing's fuel tank. The terminal effect would be the large hole in the wing's skin that resulted from the hit. Finally, the kill mode would be the in-tank explosion that rendered the control surfaces of the wing inoperable.

B. HOW SURVIVABILITY IS MEASURED: THE PROBABILISTIC KILL CHAIN

In combat, nothing is known with certainty. For any mission scenario, many random variables will ultimately determine the mission's outcome. Aircraft survivability is measured using probabilities because of these uncertainties. In ACS, the probability that an aircraft survives an encounter is denoted as P_S (Ball, 2003). The complement to an aircraft's probability of survival is its probability of being killed, denoted as P_K . These probabilities range from 0 to 1 and are mutually exclusive and exhaustive outcomes. Either the aircraft survives the encounter or the aircraft is killed; the two probabilities must add to one (Equation 1).

$$P_S = 1 - P_K \quad (1)$$

An aircraft's killability, the ease at which an aircraft is killed, is a product of an aircraft's susceptibility and vulnerability (Adams, 2019a). For an engagement level assessment (that is one weapon versus one aircraft), an aircraft's susceptibility can be broken down into five sequential, mutually exclusive events. For an aircraft to be hit by a KEW, the weapon must perform the following phases in order: 1) the weapon system must be active, 2) the weapon system must detect the aircraft, 3) the weapon must be launched at the aircraft, 4) the weapon must intercept the aircraft, and 5) the weapon's damage mechanisms must hit the aircraft, either by direct hit or proximity fusing (Ball, 2003). For an aircraft to be killed by a KEW, the hit must result in the loss of a critical component and therefore a loss of either a flight or mission essential function. This last part makes up the aircraft's vulnerability section.

The major phases of a kinetic attack and their respective probabilistic outcomes can be represented using a simple tree diagram. The probabilistic kill chain shown below is used in ACS to describe the structure of a kinetic attack and for calculating survivability and killability for a one-on-one engagement (Figure 2).




Figure 2. ACS probabilistic kill chain. Source: Ball (2003).

The numbers in parenthesis represent each of the six phases necessary for an aircraft kill. Each phase is binary and results in either a successful outcome or an unsuccessful outcome. There are two complementary probabilities for each phase of the kill chain. P_x represents the probability that a particular phase will be successful. P^c_x represents the probability that a particular phase will be unsuccessful. Conditional probabilities are used after the first phase to represent the likelihood of outcomes given that the previous phases were successful. Conditional probabilities allow us to view each phase as an independent event. $P_{x|y}$ represents the likelihood of phase 'x' occurring *given* that phase 'y' (the previous phase) was successful. Ball defines the six successful conditional phase probabilities as follows:

1. P_A is the probability that a threat weapon on the vicinity of the aircraft is active, that is, the weapon is searching, actively or passively, and ready to encounter and engage aircraft flying within its defined area.
2. $P_{D|A}$ is the conditional probability that the aircraft is detected, given that the threat is active.
3. $P_{L|D}$ is the conditional probability that the aircraft is tracked, a fire control solution is obtained, and a missile is launched or a gun is fired at the aircraft, given that the threat was active and detected the aircraft.
4. $P_{I|L}$ is the conditional probability that the threat propagator approaches or intercepts the aircraft, given that the propagator was launched or fired at the aircraft.
5. $P_{H|I}$ is the conditional probability that the propagator hits the aircraft, given that the propagator has intercepted the aircraft.
6. $P_{K|H}$ is the conditional probability that the aircraft is killed, given a direct hit by the propagator. (Ball, 2003, p. 12)

Each path of one or more sequential branches is called an event (Ball, 2003). The likelihood of a particular event occurring can be calculated by multiplying each subsequent branch together. For example, the probability of the weapon launching but not intercepting the aircraft (an event) would be: $P_A \times P_{D|A} \times P_{L|D} \times P_{I|L}^c$.

From the attacker's point of view, for an encounter to result in an aircraft kill, each of the subsequent phases must have successful outcomes. An aircraft's susceptibility to a KEW, P_H , is measured by the probability that an aircraft is hit during an engagement (Equation 2).

$$P_H = P_A \times P_{D|A} \times P_{L|D} \times P_{I|L} \times P_{H|I} \quad (2)$$

An aircraft's vulnerability, $P_{K|H}$, is measured by the conditional probability that an aircraft is killed given a hit. Killability, as defined earlier, is the product of susceptibility and vulnerability. Killability is therefore the product of all of the successful phase probabilities listed above (Equation 3).

$$P_K = P_H \times P_{K|H} = P_A \times P_{D|A} \times P_{L|D} \times P_{I|L} \times P_{H|I} \times P_{K|H} \quad (3)$$

If any phase results in an unsuccessful outcome, the chain is broken and the aircraft survives. An aircraft's survivability during a one-on-one engagement can be calculated by subtracting the probability of kill from one (Equation 1). Therefore, to enhance an aircraft's

survivability, one must reduce an aircraft's susceptibility or vulnerability by decreasing the likelihood that one or more of these six events is successful.

C. SURVIVABILITY ENHANCEMENT

Survivability, as defined earlier, is the ability of an aircraft to avoid or withstand an MMHE (Ball, 2003). Therefore, the survivability of an aircraft can be improved by either making the aircraft more capable of avoiding threats present in the MMHE or making the aircraft more capable of withstanding damage caused by threats present in the MMHE. In other words, survivability can be enhanced by either reducing an aircraft's susceptibility (the inability to avoid damage) or by reducing an aircraft's vulnerability (the inability to withstand damage).

A survivability enhancement feature is defined by Dr. Ball (2003) as "any particular characteristic of the aircraft, specific piece of equipment, design technique, armament, or tactic that reduces either the susceptibility or the vulnerability of the aircraft" (p. 34). From the ACS probabilistic kill chain (Figure 2), an aircraft kill results only if all six events successfully occur. The purpose of survivability enhancement is to increase the likelihood that the sequence of events is broken at some point along the kill chain. The current ACS design discipline created by Ball (2003) lays out twelve general concepts that are fundamental to survivability enhancement. The survivability enhancement concepts summarized below function as tools for program managers and engineers as they design and build survivable combat platforms. These concepts focus primarily on KEWs.

1. Susceptibility Reduction

The first six survivability enhancement concepts focus on reducing an aircraft's susceptibility. The susceptibility reduction concepts are designed to reduce the likelihood that an aircraft is hit in combat.

a. Situational Awareness

The purpose of the situational awareness concept is to provide information to the pilot and crew about the location, type, and status of the kinetic threat elements near the aircraft (Ball, 2003). Warning the pilot of an imminent threat gives the pilot a chance to

respond with a course of action that minimizes the likelihood that the KEW successfully hits the aircraft. The use of radar warning receivers or missile launch and approach warning systems are examples of reducing susceptibility via situational awareness.

b. Signature Reduction and Control

The purpose of the signature reduction and control concept is to reduce how observable an aircraft is to an adversary or an adversary's weapon system (Ball, 2003). For a guided missile to successfully hit its target, it must be able to detect and track the position of the target using sensors. Reducing the radar, infrared (IR), ultraviolet (UV), acoustic, visual, or magnetic signature of an aircraft below the weapon system sensor's threshold reduces susceptibility by decreasing the likelihood that the threat can successfully detect and intercept the aircraft. For example, stealth aircraft use the signature reduction concept by minimizing its radar cross-section to avoid detection from anti-aircraft weapon systems.

c. Noise Jamming and Deceiving

The purpose of the noise jamming and deceiving concept is to degrade the effectiveness of various threat elements by interfering with their electronic signals (Ball, 2003). Noise jamming utilizes active electronic equipment that radiates a concentrated signal toward an enemy's detection and tracking equipment. Noise jamming interferes with and blocks the signal received by an enemy's weapon system. Noise jamming reduces susceptibility by reducing the likelihood of successful detection. The objective of deception is to transmit signals designed to confuse or mislead an enemy's weapon system. Susceptibility can be reduced by using deceptive signals to appear as a large number of false targets or as a target with incorrect bearing, range, or velocity information.

d. Expendables

The purpose of the expendables concept is to eject materials or devices from an aircraft to divert a KEW's tracking system (Ball, 2003). Expendables can be used for self-protection or for mutual support between multiple aircraft. Examples of expendables include chaff, retroreflectors, flares, and aerosols. Chaff and retroreflectors are effective against radar-guided missiles, flares are effective against IR-guided systems, and aerosol

is effective against visually-directed or IR-guided missiles. Expendables reduce susceptibility by decreasing the likelihood that a weapon successfully intercepts the aircraft by confusing and diverting anti-aircraft missiles.

e. Threat Suppression

The purpose of the threat suppression concept is to physically damage, destroy, or deny the use of an enemy's air defense system (Ball, 2003). Threat suppression can be done by supporting aircraft, ground elements, or by the target aircraft. While physically destroying an enemy's air defense is obviously an extremely effective strategy for enhancing survivability, effective susceptibility reduction can also be accomplished via deterrence. For example, the presence of a fighter escort might discourage an adversary from launching their anti-aircraft missiles for fear that the fighter escort might fire back. Threat suppression effectively reduces susceptibility by destroying or damaging the threat before the threat can destroy or damage the target aircraft.

f. Weapons and Tactics, Flight Performance, and Crew Training and Proficiency

The last susceptibility reduction concept combines personnel and platform performance (Ball, 2003). The weapons onboard a combat aircraft greatly affect the aircraft's survivability. For example, a shooter aircraft fitted with JASSMs (Joint Air-to-Surface Standoff Missile) that can be fired from over 200 miles away would have a much higher likelihood of survival than a shooter aircraft fitted with only JDAMs (Joint Direct Attack Munitions), which have a maximum range of only 15 miles. Susceptibility is reduced if the aircraft's weapons can be launched outside of the range of enemy air-defenses.

The combat tactics are also important for determining the survivability of an aircraft during a given mission. Survivability will change depending on how many aircraft are used for a mission and the amount of air and ground support they receive. Generally, greater aerodynamic performance is preferred for survivable fighter aircraft, but it is important to consider the design trade-offs (i.e., stealth versus agility) required to design faster, more maneuverable aircraft. Finally, a large part of reducing susceptibility is left to the training

and proficiency of the crew operating and maintaining the aircraft. A more proficient aircrew will undoubtedly have a higher likelihood of survival than a poorly trained or inexperienced aircrew.

2. Vulnerability Reduction

The final six survivability enhancement concepts focus on reducing an aircraft's vulnerability. The vulnerability reduction concepts are designed to reduce the likelihood that an aircraft hit by a KEW is killed.

a. Component Redundancy (with Physical Separation)

The purpose of the component redundancy concept is to have multiple parts, systems, or mechanisms in place that can each perform the same function in the event that the main component is rendered inoperable (Ball, 2003). For example, if the only engine of a single-engine aircraft is killed, the hit will likely result in an aircraft kill. However, if only one of two engines is hit, the aircraft will likely be able to withstand the damage and continue flying using its working engine. However, if the two engines are right next to each other (without physical separation), one hit may kill both engines. Physical separation is critical for component redundancy. Backing up critical components with physically separated redundant components reduces vulnerability by reducing the likelihood that a single hit results in the loss of essential functions.

b. Component Location

The purpose of the component location concept is to position critical components in a manner that reduces the likelihood that a single hit will result in the kill of that critical component (Ball, 2003). This can be done multiple ways. Positioning non-critical components in front of critical components reduces the likelihood that a damage mechanism will produce lethal damage to the critical component. The orientation of critical components can also reduce vulnerability by minimizing the component's presented area. The vulnerable area (and therefore the vulnerability) can also be reduced by overlapping critical components.

c. Component Elimination or Replacement

The purpose of the component elimination or replacement concept is to reduce aircraft vulnerability by either getting rid of a critical component or substituting a critical component with a less vulnerable component that serves the same function (Ball, 2003). For example, UAVs (unmanned aerial vehicles) reduce vulnerability by eliminating the need for onboard life-supporting equipment. Alternatively, replacing a large critical component with a smaller component with the same functionality reduces vulnerability by reducing the presented surface area of that critical component. It is important to note that replacing a component may introduce new kill modes associated with the new component. These new potential kill modes must be taken into account when designing aircraft for survivability.

d. Component Shielding

The purpose of the component shielding concept is to reduce aircraft vulnerability by using coatings or plating to shield critical components from damage mechanisms (Ball, 2003). Strong, armored material can resist or absorb kinetic damage effects so that the critical components behind them can still function if an aircraft is hit by a KEW. For example, a blast or fragment shield in the cockpit reduces vulnerability by reducing the likelihood that an explosion from an anti-aircraft warhead kills or injures the pilot and/or crew. An important distinction is necessary to differentiate the component shielding and component location concepts. The component shielding concept is used when coatings or armored plates are added to a design. The component location concept is used when critical components are positioned behind non-critical components.

e. Damage Suppression

The purpose of the damage suppression concept is to reduce vulnerability by containing damage caused by a damage mechanism or reduce the effects of that damage (Ball, 2003). Damage suppression can be either active (requiring damage-sensing capability) or passive in nature. Damage tolerance and ballistic resistance are just two passive ways of suppressing damage to reduce the effects of damage. An override switch

that, after detecting damage to a certain system, allows a pilot to disengage the damaged system is an example of active damage suppression. Damage suppression techniques reduce vulnerability by decreasing the likelihood that the damage caused by a damage mechanism will ultimately result in an aircraft kill.

f. Recovery

The purpose of the recovery concept is to disengage damaged or disabled components and return the aircraft to responsive control or, if recovering full flight control is not possible, to safely land the aircraft to be repaired by a ground maintenance crew. The ability to recover an aircraft is dependent on both the training of the aircrew operating the aircraft and the engineering design of the aircraft itself. For example, if a KEW were to hit and kill the port engine of an F-18, the pilot may still be able to recover the aircraft by disengaging the damaged engine to return the aircraft to responsive control using only the starboard engine. Aircraft recovery almost always requires knowledge of the location and extent of damage. Therefore, for successful recovery, it is imperative that operators have strong situational awareness of existing threats and threat effects. If full recovery of mission essential functions is not possible, recovering the physical aircraft via a forced landing, while still resulting in a mission kill, is far better than losing the aircraft and its crew.

D. WHY DO WE NEED AIRCRAFT COMBAT SURVIVABILITY?

The DoD mandates that air combat programs conduct thorough survivability assessments for good reason (DoD, 2018b). The goal of survivability enhancement is to improve the aircraft's overall cost effectiveness as a warfighting platform (Ball, 2003). While designing for survivability may increase the initial development and production costs, designing aircraft that are more capable of achieving mission objectives in high-risk combat scenarios and returning home safely will ultimately decrease the lifecycle cost of the program (Adams, 2019a).

History has shown that platforms not designed to survive in their hostile operating environments are less effective at carrying out sustained air operations (Ball, 2003). Survivability, therefore, has a direct impact on our nation's ability to save lives in combat

and win wars decisively. Historically, when aircraft are designed to be survivable, “fewer wars are fought, more missions are accomplished, more battles are won, and more lives are saved” (Ball, 2003). Figure 3 shows historical combat aircraft loss rates throughout the twentieth century.




Figure 3. Historical losses and loss rates. Source: Ball (2003).

Since the 1970s, when the ACS design discipline was formally introduced, combat loss rates have decreased significantly. The focus on survivability over the last fifty years has saved countless lives and although current combat loss rates continue to remain low, the hostile environments in which our aircraft operate are dynamic and always changing. Kinetic air-defense systems continue to evolve and new anti-aircraft threats are beginning to emerge. The survivability of aircraft must continue to improve today to face the increasingly sophisticated and lethal anti-aircraft threats of tomorrow.

E. THREATS TO COMBAT AIRCRAFT

Aircraft combat survivability is needed because combat aircraft operate in hostile environments full of lethal threats. Just how much survivability is needed depends on the types of threats expected in the operating environment and the estimated effectiveness of those threats (Ball, 2003). Threats to aircraft can be broken down into four categories: kinetic energy weapons, electromagnetic (EM) weapons, unconventional weapons, and cyber weapons (Figure 4). While KEWs have historically been the largest threat to combat aircraft, other threats exist and should not be ignored. Today's combat platforms are expected to operate effectively for decades. It is therefore critical that systems are designed not just to survive in the current warfighting environment, but in the projected warfighting environment of the future, when electronic warfare, unconventional, and cyber threats might play a larger role in armed conflicts. Dr. Ball's textbook on ACS, published in 2003, has no mention of cyber threats. The global risk landscape has changed dramatically since 2003. Cyber threats were not even included in the World Economic Forum's global risk report until 2010 (World Economic Forum [WEF], 2010). Now, cyber-attacks are listed as one of the top five risks in both likelihood and impact (WEF, 2019).




Figure 4. Categorizing threats

1. Kinetic Energy Weapons

Kinetic energy weapons (KEW) include the guns, missiles, and rocket-powered grenades (RPGs) traditionally used for air-defense. Anti-aircraft guns are devices that propel rounds, projectiles, or shells via an explosive force. Guns range from 20 mm, hand-held small arms to over 100 mm, target tracking anti-aircraft heavy artillery. Missiles are self-propelled weapons that typically contain a propulsion system, guidance system, tracking sensors, and high-explosive warhead. These weapons are referred to as kinetic energy weapons because the damage mechanisms associated with their warheads rely on kinetic energy (Ball & Bryant, 2020a). The damage mechanisms associated with KEWs typically include metallic penetrators, armor-piercing ballistic penetrators (AP), armor-piercing ballistic penetrators with incendiaries (AP-I), high-velocity fragments, and blasts.

KEWs are the most common threat to combat aircraft. These weapons can be either guided or unguided and can operate independently, in small groups, or as part of a large integrated air-defense system (IADS) (Ball, 2003). KEWs can be fired from stationary or mobile surface-based platforms or airborne platforms. This includes air-to-air missiles, such as the infrared-guided AIM-9 Sidewinder, surface-to-air guns, such as the Phalanx close-in weapon system (CIWS), and surface-to-air missiles, such as the land-based MIM-104 Patriot. Dr. Ball's ACS textbook focuses primarily on these KEWs (Ball, 2003).

2. Electronic Warfare Weapons

The DoD defines electronic warfare as any “military action involving the use of electromagnetic and directed energy to control the electromagnetic spectrum or to attack the enemy” (Joint Chiefs of Staff [JCS], 2020, p. 71). The weapons used to conduct electronic warfare can be divided into two categories: traditional electromagnetic weapons and directed energy weapons.

Traditional electromagnetic weapons include jammers and anti-radiation missiles. Jammers are devices that generate EM energy to blind radars with high-power noise or hide the location of a target by generating false radar targets (LaMarche, 2018). Anti-radiation missiles are missiles designed to seek and destroy an adversary's radar, radio, or communication signals. The damage mechanism associated with traditional

electromagnetic weapons is EM radiation. Traditional electromagnetic weapons are primarily used to disrupt or deny an adversary's use of the electromagnetic spectrum (EMS).

Directed energy weapons use beams of concentrated EM energy or subatomic particles to incapacitate, damage, or destroy target platforms, facilities, or personnel (JCS, 2020). This category of weapons includes neutral or charged particle beams, high-energy lasers, and high-power radio-frequency/microwave (HPM) weapons. Particle beam weapons, as the name suggests, use high-energy beams of atomic or subatomic particles to affect a target's atomic or molecular structure. These particles can be launched at near light-speed and can cause instantaneous superheating and ionization. Lasers (light amplification by stimulated emission of radiation) are devices that emit and direct highly concentrated light through optical amplification. While low-energy level lasers are typically used for weapon guidance, high-energy lasers may be used to take down drones or manned aircraft (Ball, 2003). High-power radio-frequency/microwave weapons are EM sources that direct intense radio-frequency/microwave energy to produce temporary or permanent damage to electronic systems (Tatum, 2019). HPMs can engage multiple targets at speeds near the speed of light, can produce scalable effects with unlimited low-cost ammo, and is non-lethal to humans. While most large-scale directed energy weapons are still in the developmental stage, the scientific theories behind these weapons are sound and experiments backed by decades of research continue to show that direct energy weapons are plausible, if not probable.

3. Unconventional Weapons

Unconventional weapons include nuclear, chemical, and biological weapons. An important distinction between unconventional and conventional weapons is their scope of lethality. Whereas kinetic energy and electromagnetic weapons are likely to attack single platforms or weapon systems, unconventional weapons are capable of killing many targets or causing mass casualties (Ball, 2003). Unconventional weapons are often referred to as weapons of mass destruction (WMD) and their use in armed conflict is illegal.

Nuclear weapons are explosive devices that release extremely high amounts of energy using nuclear fission and/or fusion processes. The primary damage mechanisms associated with nuclear weapons are thermal radiation, nuclear radiation, and blast waves (Ball, 2003). Chemical weapons are munitions that release chemical substances designed to kill, injure, incapacitate, or irritate. Chemical weapons can be categorized by their physiological damage mechanisms and include nerve agents, blister agents, choking agents, and blood agents. Biological weapons release viruses, bacteria, or other toxic microorganisms to cause material deterioration or death and disease in humans, plants, and animals (World Health Organization, 2020). Biological weapons typically contain pathogens (living organisms that cause disease and are capable of multiplying and spreading) or toxins (highly poisonous substances produced by living organisms). Because of the distinct differences between unconventional and kinetic energy weapons, an entirely separate nuclear, biological, and chemical survivability discipline has evolved apart from ACS (Ball, 2003).

4. Cyber Weapons

A cyber weapon is any piece of malware employed by a state or non-state actor intended to degrade, disrupt, destroy, or deny operations in cyberspace. The cyber domain is made up of millions of lines of code and data that control information technologies, computer systems, processors, industrial controllers, and communication networks. While cyber weapons can cause physical damage to cyber systems and hardware, attacks in the cyber domain typically result in functional damage or dysfunctions. A cyber weapon's dysfunction mechanism is the malicious computer code that causes the loss or degradation of system capabilities (Ball & Bryant, 2020a). Unlike other threats, cyber weapons are almost entirely virtual, but their effects are still felt in the real world. Cyber weapons can be launched instantaneously from anywhere in the world and the covert nature of cyber-attacks makes attribution incredibly difficult.

While kinetic energy weapons continue to be the leading threat to combat aircraft, emerging threats in cyber are growing fast. The rest of this thesis will discuss the cyber threat in more depth and offer guidance for designing cyber-survivable combat aircraft.

III. A NEW THREAT IN CYBER

A. FUNDAMENTALS OF CYBER

Many members of the DoD have heard of “cyber,” but few are fully aware of the extent to which cyberspace is foundational to military operations and of the increasing threat levels associated with it. Before attempting to solve the issue of enhancing the cyber combat survivability of aircraft platforms and weapon systems, it is necessary to lay some groundwork about what cyber defense is and what DoD systems are at risk of a cyber-attack.

1. Cyberspace and Cybersecurity

Cyberspace, as defined by the United States Joint Chiefs of Staff (2014), is “a global domain within the information environment consisting of the interdependent network of information technology infrastructures and resident data, including the Internet, telecommunications networks, computer systems, and embedded processors and controllers.” Two important take-aways need to be highlighted from this definition. The first is that cyberspace, and therefore a cyber-attack, is not limited to computers and the Internet. Any piece of hardware that communicates remotely with another piece of hardware is a part of cyberspace and therefore has cyber vulnerabilities. Any processor, controller, or avionics computer is a part of the cyber-domain. The first important take-away leads to the second important take-away. Nearly all mission platforms, weapon systems, and day-to-day logistics and operations used by our armed forces are heavily reliant on cyberspace operations.

Cybersecurity is the capability to restore and/or prevent dysfunctions to those components that comprise cyberspace, including computers, electronic communications systems, and the information contained therein. The Department of the Navy (2020, para. 1) defines cybersecurity as the “prevention of damage to, protection of, and restoration of computers, electronic communications systems, electronic communications services, wire communication, and electronic communication, including information contained therein, to ensure its availability, integrity, authentication, confidentiality, and nonrepudiation.”

Our military has become increasingly reliant on systems that operate in cyberspace and some of the information contained within those systems is highly sensitive. It is therefore imperative that all of our military's cyber systems are secure and well protected.

While the importance of cybersecurity is clear, it is naive to think that our systems are entirely safe from all attempted cyber-attacks, even with the most advanced, modern cybersecurity strategies (Bryant, 2016). Cyberspace operates through millions of lines of code. Advanced nation-state level adversaries have and will continue to look for the smallest chink in our armor and exploit that vulnerability. This is why it is so important to anticipate and prepare for successful cyber-attacks and ensure that the successful implementation and triggering of malware does not take the attacked system or platform out of the fight.

2. Categorizing DoD Cyber Systems

As previously mentioned, a majority of military operations are heavily dependent on systems that operate in cyberspace. Most service members are familiar with information technology (IT) systems and understand how those systems might be vulnerable to cyber-based attacks. These systems, however, makeup only one category of cyber systems that are vulnerable to cyber threats. While IT systems have generally been the most targeted systems, operational technology (OT), and weapon platforms are also vulnerable to cyber threats and should not be overlooked.

a. Information Technology Systems

The first category of cyber systems is traditional IT systems (Bryant, 2018b). IT systems consist of all of the command and control centers used by our military. This includes the Non-Secure Internet Protocol Router network (NIPRNET), Secure Internet Protocol Router network (SIPRNET), and other weapon systems, logistics systems, and personnel networks. The goal of most cyber-attacks against IT systems is to introduce a small amount of false information that, as a result, makes us question all of the information on our systems. This seemingly harmless, small attack has massive ramifications on the readiness and warfighting capabilities of our forces. Just a small amount of doubt in these

mission critical systems can render units unable to carry out mission objectives until the small amount of disguised, false information is found.

b. Operational Technology Systems

Operational technologies makes up the second category of cyber systems used by our Department of Defense (Bryant, 2018b). Operational technology refers to the computer-controlled systems that are necessary for everyday operations such as heating and air-conditioning systems. While these systems are not normally considered as potential targets for cyber-attacks, it has become a growing concern in recent years thanks to a myriad of successful cyber-attacks launched against private banks and major corporations. Imagine a scenario in which an attacker is able to disable a building's air conditioning in the middle of summer. Now imagine if that building is not just any building, but the cryptologic headquarters of tenth fleet, whose computers can no longer run because of the extreme high temperatures. Unfortunately, compared to IT systems, OT systems are largely unprotected and not designed to defend against harmful cyber-attacks because they were not designed with cyber threats in mind.

c. Combat Platforms/Weapons Systems

The last category of cyber systems is made up of the combat platforms and weapons systems themselves. These systems are normally thought to be "air-gapped," or standalone systems that are not directly connected to other networks or the internet (Bryant, 2019). This way of thinking leads many to believe that these systems are not vulnerable to cyber threats. This way of thinking is overly optimistic. These systems have to be routinely connected to maintenance devices for updates and patches, exposing these systems to the larger cyberspace universe and thus to cyber-attacks.

Cybersecurity of IT and OT systems is already widely discussed in literature, and traditional defenses can be put in place to protect these systems from cyber-attacks. The focus of this paper is to address the protection and survival of the third, less discussed category of DoD cyber systems, combat platforms/weapons systems, using the same fundamentals that Dr. Ball lays out in the ACS design discipline.

B. PROSPECT OF CYBER AS AN ANTI-AIRCRAFT WEAPON

The prospect of using a cyber-weapon to attack an aircraft is an idea that is just now beginning to be researched by the aircraft survivability community (Bryant, 2019). Traditionally, cyber-attacks have been almost solely on IT networks and very few of those attempted have had lasting physical effects. Decision-makers still do not fully understand the scope of cyber operations and the potential that the cyber domain will offer adversaries over the next decades. There has been a false assumption that because combat aircraft are “air gapped” or not directly connected to the internet, they are immune to cyber-based attacks. But, as previously discussed, cyber involves more than just the internet.

The cyber domain includes “the internet of networked computers but also intranets, cellular technologies, fiber-optic cables, and space-based communications” (Nye, 2011, p. 19). This means that cyber-attacks are not limited to computers and the Internet. Any piece of hardware that communicates remotely with another piece of hardware is a part of cyberspace and therefore has cyber vulnerabilities. Any processor, controller, or avionics computer is a part of the cyber-domain and is a potential target for a cyber-attack.

Today’s combat aircraft have a very large attack surface (Bryant, 2018a). Mission planning computers, data transfer devices, maintenance systems, smart weapons, pods, and components of software development are all physically connected to the aircraft and operate in cyberspace (Bryant & Young, 2019). Additionally, datalinks, GPS, SATCOM, radio communication, and even radar components operate in the cyber sphere through the electromagnetic spectrum and may be susceptible to a potential cyber-attack (Bryant & Young, 2019). Access to any one of these systems could put a combat aircraft and its mission in jeopardy. A cyber-attack on any of these cyber-physical systems could spread to other cyber components and cause dysfunctions that result in the loss or degradation of flight and/or mission essential functions.

To access these systems, all an adversary would have to do is find an insider or wait for an exploitable maintenance update of a “standalone” aircraft system (Bryant, 2019). Insider threats have been and continue to be a massive problem in our armed forces. The ease at which an insider threat could access an aircraft’s cyber-systems is higher than many

would like to admit and the possibility of accessing and implanting malicious code on an aircraft's systems is not as farfetched as one might think. There have been several examples of either successful cyber-attacks or investigative studies done on the defenselessness of automobiles and other "air gapped" platforms to cyber threats (Anderson et al., 2011). The Iranian nuclear plant that the Stuxnet worm targeted was one very well protected, yet an attacker was able to successfully access the SCADA system and implant malware that caused significant physical damage. Theoretically, the same could be done to an aircraft.

To make matters worse, the potential damage that the cyber threat brings to the table is only starting to be realized. Cyber operations have been getting more creative, more sophisticated, more impactful, and more covert. To nation-states that may not have had the physical resources to compete with the militaries of world superpowers, cyber offers a chance for them to punch above their weight. For our adversaries who already have a strong cyber presence like China, Russia, and Iran, cyber armies backed by nation-state resources provide a new, powerful means of damaging U.S. military assets covertly. Undoubtedly, the U.S. is not the only power pouring time and money into defending and attacking cyber-systems. The scope to which cyber-attacks might manifest themselves will likely only grow.

C. APPEAL AND IMPACT OF CYBERWARFARE

There are many reasons why nation-states would want to use cyber to conduct offensive operations against combat platforms. Sandro Gaycken (2011, p. 79), a cybersecurity expert and advisor to NATO, advocates, "states take cyber warfare seriously as they are viewed as an attractive activity by many nations, in times of war and peace. Offensive cyber operations offer a large variety of cheap and risk-free options to weaken other countries and strengthen their own positions." Gaycken adds that offensive cyber operations have high potential to agitate military conflict and reduce military capabilities. General Keith Alexander, the first USCYBERCOM commander, listed air defense networks and combat weapon systems as potential targets that could be attacked by a cyber-weapon (Shanker, 2010).

Offensive cyber operations are appealing because of the distributed nature of cyber-based attacks. This is one reason why cyber operations have been evolving so rapidly. In the cyber domain, determining the attacking party and the motivation behind the attack is often quite challenging (Ragan, 2010). This makes it increasingly difficult to determine whether a serious cyber-attack is an act of war or not. It is in this gray area between peace and war that cyber warfare operates and thrives in.

To adversaries, cyber provides a powerful platform for causing major disruptions with little to no consequences. It is difficult for a victim of a cyber-attack to respond to an attack if the attacker can easily deny their involvement. This is one of the major motivators for choosing cyber over kinetic operations. There is less need to launch a physical attack if the same outcome can result from a cyber-based attack. As Arquilla and Ronfeldt put it, “we’re no longer just hurling mass and energy at our opponents in warfare; now we’re using information, and the more you have, the less of the older kind of weapons you need” (Duggan, 2015, p. 47).

Cyber is also an appealing weapon because it has potential to pack a large punch at a relatively low cost and with low personal risk. Sophisticated yet inexpensive technologies make it increasingly easier for adversaries to develop substantial offensive capabilities. The low relative cost of developing a cyber-weapon is another big reason why cyberwarfare is appealing (Duggan, 2015). Smaller countries with fewer resources seeking to gain military equivalence to more powerful nation-states can use cyber, an option that is cheap, available, and powerful, to level the playing field. Cyber also has very low personal risk compared to traditional kinetic forms of warfare (Korunka et al., 2019). A cyber-attack aimed at downing a combat aircraft is less likely to result in attacker casualties than if the attacker sent in its own combat aircraft to try to shoot it down.

Lastly, developing offensive cyber capabilities is appealing for adversaries because it can be planted without the victim knowing of its presence and then can be triggered when the adversary sees fit. This opens many doors for coordinating large-scale, well-timed operations. It also keeps the defending nation in a state of uncertainty, always wary of the possibility that their systems have been compromised.

Targeting combat systems in the cyber domain is both possible and appealing to highly skilled, persistent threats and the impact of such an attack could be tremendous. If an actor were to successfully bring down a combat aircraft with a cyber-weapon, it is likely that the conflict would escalate to war. While cyber conflict between nations is currently operating well below the threshold of armed conflict, a cyber-attack that produced kinetic effects and resulted in loss of life could very well shift the entire geopolitical landscape and forever alter the role that cyber plays in it (Olejnik, 2019). Historically, armed forces in close proximity have experienced higher risk of force escalation, but with cyberspace, borders are virtualized, so constant interaction between armed forces and the potential for escalation between actors is always possible.

Lawmakers are still behind the curve in addressing what to do in the event of a serious cyber-attack. There is currently no international stance on what determines whether a cyber-attack is considered an act of war or not (DoD, 2018a). It is therefore almost entirely up to the victim of the attack to determine what course of action, if any, to take. For a successful attack on a combat aircraft, the course of action would almost certainly involve significant retaliation. Of course, for retaliation to occur, the victim would have to know (or claim that they know) who attacked them. Retaliation could happen in the cyber domain or in the physical domain, but a kinetic response would be more likely given the severity of the attack.

More importantly, launching a cyber-weapon capable of downing a combat aircraft would set a precedence for how warfighting is done in the future. No lives have been lost due to a cyber-attack, but it may just be a matter of time before that changes. Cyber is a new field and the customs and norms of cyber operations are still developing. If one actor can successfully bring down a combat platform with nothing but malware, other nations will likely follow suit and try to develop their own anti-aircraft cyber weapons. Cyber could very well be the next great arms race.

D. HOW AN ANTI-AIRCRAFT CYBER-ATTACK MANIFESTS ITSELF

The idea of using cyber-weapon designed to attack combat aircraft systems and cause functional damage to mission or flight performance may be appealing for competent,

determined adversaries, but how might such an attack manifest itself? It has been previously discussed how large a combat aircraft's attack surface is; aircraft have many systems and subsystems that are connected to and operate in the cyber domain. The question remains- how might an adversary gain access to, exploit, and implant malicious code on one or more of these cyber-physical systems?

To gain access to a combat aircraft's highly-secured, well-monitored cyber-systems, the attacker would almost have to be from another powerful nation-state. The term advanced persistent threat (APT) is used to describe the category of sophisticated, precise attackers backed with the resources, intelligence, and skill of a combative nation-state. APTs use precise methods to hit specific targets with malicious intent. Actions from APTs are typically slow and methodical with unmatched complexity to avoid detection (Maness, 2020b).

There are many methods that an APT might use to gain access to an aircraft's cyber systems. An intrusion, such as a Trojan horse, trapdoor, or backdoor, is one potential method of accessing a susceptible system and remotely injecting malicious software (Maness, 2020b). Trojan horses, for example, are designed to hide malicious code in normal data so that the target runs the malware unknowingly. These attacks typically involve an unsuspecting user clicking on an email with the purpose of espionage, but the method could easily be applied to an unsuspecting maintenance worker who downloads malicious code hidden in flight data that is meant to appear normal.

Another method that an adversary might utilize to gain access would be through infiltration techniques. These are used to penetrate a target network and remotely or physically install malicious code on those networks (Maness, 2020b). Infiltration methods include the use of viruses, worms, and keystroke logging. The Stuxnet worm, for example, was able to successfully infiltrate protected machines using stolen digital certificates (CSFI, 2010). If a well-designed attack can target and infiltrate specific hardware used to control a nuclear plant, what is to say that similar tactics could not be used to infiltrate third-party hardware attached to a combat aircraft?

Adversaries could also relatively easily use impersonations to gain physical access into highly-secured cyber-components. An adversary could gain access to and plant malicious code on aircraft hardware by becoming a respected member of a specific community, like an aircraft maintenance loader or an engineer at a major defense company. Long-term, undercover operations are not new concepts and have been known to work in the past. If an adversary could pretend to be a hard-working employee and earn access to a combat aircraft's cyber components, they could modify or add malicious code that could be later triggered to cause a dysfunction.

Lockheed Martin, one of the giants in the aerospace, defense, and technology industry, has developed a useful framework for understanding the steps an adversary must complete to launch a successful cyber-attack (Figure 5).




Figure 5. Lockheed Martin's cyber kill chain. Source: Lockheed Martin (n.d.).

While the Lockheed Martin cyber kill chain model can be used to better understand and defend against a majority of cyber-attacks, its primary focus is on attacks against traditional IT cyber systems. In the next chapter, a kill chain based on the ACS fundamentals is proposed to specifically address cyber-attacks on aircraft.

IV. DEVELOPING AIRCRAFT CYBER COMBAT SURVIVABILITY (ACCS) FROM ACS FUNDAMENTALS

A. KEY DIFFERENCES BETWEEN ACS AND ACCS

Before attempting to integrate anti-aircraft cyber threats into the preexisting ACS model, we must highlight some important differences between cyber weapons and traditional KEWs. Understanding these differences is essential to defending against the cyber threat and creating cyber-survivable combat aircraft.

The first difference between KEWs and cyber weapons is the difficulty of detection. It is relatively easy to know when an aircraft is under attack from a KEW. Today's combat aircraft have special instruments designed to detect incoming missiles and launch last-second countermeasures (Bryant, 2019). If the instruments do not warn you of an impending attack, the blast from the warhead will make it obvious. Cyber-attacks are much harder to detect. Cyber-attacks are often intentionally well-hidden and can be designed to make systems appear as if they are running normally. Differentiating a routine system failure from a hostile cyber-attack may be exceptionally challenging. Unlike with KEWs, there are currently no warning or detection instruments on aircraft designed to warn operators of an incoming cyber-attack.

Second, cyber weapons have a nearly unlimited range and can travel across the world in a matter of seconds (Bryant, 2019). Kinematic relationships govern the range of KEWs. There are limitations as to how far and how fast missiles can travel. Cyber weapons do not have to fight gravity. They can be launched from anywhere in the world and, unlike most KEWs, a single cyber warhead can have multiple targets. Cyber weapons may be capable of targeting not one aircraft, but every aircraft that runs a particular version of software (Ball & Bryant, 2020a). The distributed nature of cyber-attacks gives cyber operations an incredibly large scope and makes attribution exceedingly difficult.

The third major difference between cyber weapons and KEWs is that cyber weapons are much less predictable (Bryant, 2019). The underlying physics behind the damage caused by guns and missiles is well tested and well understood. Analysts have

created incredibly precise computer models to simulate the flight, accuracy, and impact of KEWs and have backed up these models with mandated live-fire testing. Given initial launch conditions, KEWs will have a predetermined and calculated range, trajectory, and launch envelope. Cyber weapons, their method of implantation, and their damage-causing mechanisms can be much more difficult to predict. Cyber weapons expose weaknesses in software made up of millions of lines of code. Furthermore, the attacked software often connects various other cyber components, each composed of millions of more lines of code. The virtual interactions between these immensely complex cyber components are incredibly difficult to predict in even the most carefully controlled lab (Ball & Bryant, 2020a). On top of that, intrusion, infiltration, and impersonation techniques frequently take advantage of human factors, which again, are not easy to predict.

Fourth, unlike KEWs, cyber weapons do not have decades of historical air combat data to back up combat models (Bryant, 2019). Reasonable and credible estimates are typically founded on historical data. Adversaries have been launching KEWs at combat aircraft for over one hundred years and for a majority of that time, analysts have collected and analyzed the combat data to enhance aircraft survivability or improve weapon lethality. There is no historical data on anti-aircraft cyber-attacks because the cyber threat has not yet been fully realized. The lack of historical evidence does not imply that cyber is not a legitimate anti-aircraft threat. As discussed in the previous chapter, a number of experts agree that cyberwarfare will play a major role in future engagements. Today's cyberwarfare experts are in the same position that air warfare experts were in the beginning of the twentieth century, when aerial combat was nothing more than an emerging threat on the horizon (Bryant & Ball, 2020a).

Finally, cyber weapons, once discovered, are easy to render harmless (Bryant & Ball, 2020a). Unlike modern anti-aircraft guided missiles, cyber weapons are generally only effective once. An air-to-air missile capable of killing a combat aircraft one day might still be an effective and lethal anti-aircraft option a month later. Cyber weapons, on the other hand, rely on small changes to data or code to be effective. If the malicious instructions of a cyber-weapon are discovered, uploading software patches or reloading

different versions of the software can easily fix the problem. This is one advantage defenders have against cyber weapons that they do not have against KEWs.

B. DEFINING KEY TERMS

Aircraft Cyber Combat Survivability is intended to be an expansion of the ACS design discipline; it is not to be its own, separate design discipline. As such, many of the terms used in ACS are also used in ACCS with some slight definition modifications. The definitions of key terms for both ACS and ACCS are provided below in Table 1. These terms are essential to the ACS design discipline, and thus, are equally essential to understanding ACCS.

Table 1. ACS and ACCS definitions.

ACS (Kinetic Threats)	ACCS (Cyber Threat)
<i>Aircraft Combat Survivability:</i> The capability of an aircraft to avoid or withstand a man-made hostile environment. To avoid means the aircraft avoids being physically hit by one or more warhead damage mechanisms. To withstand means the aircraft maintains an acceptable level of flight and mission essential functions, while in flight, after being hit by one or more warhead damage mechanisms.	<i>Aircraft Cyber Combat Survivability:</i> The capability of an aircraft to avoid or withstand a hostile cyber environment. To avoid means either the aircraft avoids the implantation of a cyber-weapon's dysfunction mechanism into the aircraft's internal cyber systems or the aircraft avoids the subsequent activation or triggering of a successfully implanted dysfunction mechanism. To withstand means the aircraft maintains an acceptable level of flight and mission essential functions, while in flight, after the activation or triggering of one or more implanted warhead dysfunction mechanisms.
<i>Man-Made Hostile Environment:</i> The threat environment including Command and Control mechanisms, search mechanisms such as radars, warhead delivery mechanisms such as guns and missiles, and warheads such as ballistic projectiles and exploding warheads.	<i>Hostile Cyber Environment:</i> The cyber threat environment including wired and wireless communication pathways, Command and Control mechanisms, cyber search mechanisms such as scanning or sensing, cyber warhead delivery mechanisms such as malware, and cyber warheads such as implanted malicious functionality.

ACS (Kinetic Threats)	ACCS (Cyber Threat)
<i>Aircraft Susceptibility:</i> The inability of an aircraft to avoid being physically hit by one or more warhead damage mechanisms associated with KEWs. The more likely an aircraft is to be hit by one or more kinetic energy warhead damage mechanisms, the more susceptible the aircraft is.	<i>Aircraft Cyber Susceptibility:</i> The inability of an aircraft to avoid having portions of its internal cyber code accessed, modified, and activated or triggered by a cyber-weapon's dysfunction mechanism. The more likely an aircraft's internal cyber system is to be accessed, modified, and activated by one or more cyber-weapons, the more cyber susceptible the aircraft is.
<i>Aircraft Vulnerability:</i> The inability of an aircraft to withstand, while in flight, one or more hits by warhead damage mechanism hits associated with KE weapons. The more likely an aircraft is unable to withstand one or more hits by the kinetic energy warhead damage mechanisms, the more vulnerable the aircraft is.	<i>Aircraft Cyber Vulnerability:</i> The inability of an aircraft to withstand the activation or triggering of a successfully implanted malfunction mechanism. The more likely an aircraft is unable to withstand the activation of an implanted malfunction mechanism, the more vulnerable the aircraft is.
<i>Kinetic Energy Warhead:</i> KEWs include guns, with their ballistic projectiles, and guided missiles, with their tube or rail launchers. Every KEW has a warhead. The warhead consists of, contains, or generates the physical entity(s) that can cause damage to an aircraft's components when they impact on, or hit, the aircraft. For the smaller guns, the warhead is the ballistic projectile itself, which may consist of an armor-piercing penetrator (AP), or it may also contain incendiary particles intended to start an internal fire (API). For both the larger guns and the guided missiles, the warhead contains a high explosive (HE) core with a surrounding metal case. The HE core detonates either upon impact on the aircraft, or in proximity to the aircraft, generating a blast wave and high velocity warhead case metal fragments. The penetrators, incendiary particles, blast wave, and warhead case fragments are the warhead's damage-causing mechanisms.	<i>Cyber Warhead:</i> A cyber-weapon warhead consists of a set of malicious computer instructions that is designed to gain access to the aircraft's internal cyber system and subsequently command one or more aircraft flight critical or mission-critical components to dysfunction. The set of malicious instructions or commands is the cyber-weapon's warhead and the specific component dysfunction(s) contained in the set of malicious instructions is known as the warhead's dysfunction mechanism(s).

ACS (Kinetic Threats)	ACCS (Cyber Threat)
<i>Damage Mechanism:</i> The physical entity that causes damage to the aircraft. The damage mechanisms for high explosive warheads and penetrator type weapons are typically metallic penetrators and fragments, incendiary materials, and blasts.	<i>Dysfunction Mechanism:</i> The malicious computer code that causes the loss or degradation of system capabilities or the output of the cyber-warhead that causes functional damage to the target. The dysfunction mechanisms for anti-aircraft cyber warheads include functionally damaging an aircraft's critical components, disrupting or changing information that will cause operators to make poor decisions, or falsely representing data to cause a forced landing.
<i>Survivability Enhancement Concept:</i> General functions or concepts fundamental to survivability enhancement and reducing either the susceptibility or the vulnerability of the aircraft.	<i>Cyber Survivability Enhancement Concept:</i> General functions or concepts fundamental to cyber survivability enhancement and reducing either the cyber susceptibility or the cyber vulnerability of the aircraft.
<i>Survivability Enhancement Feature:</i> Any particular characteristic of the aircraft, design, design of supporting systems, or operational procedures that reduce either the susceptibility or the vulnerability of the aircraft. These features can improve an aircraft's survivability.	<i>Cyber Survivability Enhancement Feature:</i> Any particular characteristic of the aircraft, design, design of supporting systems, or operational procedures that reduce either the cyber susceptibility or the cyber vulnerability of the aircraft. These features can improve an aircraft's cyber survivability.

Adapted from Ball & Bryant, 2020a; Ball & Bryant, 2020b, and Ball & Bryant, 2020c

C. CYBER PROBABILISTIC KILL CHAIN

The same probabilistic kill chain used in ACS for engagement-level survivability assessment can be adapted to address cyber weapons as well. This modified cyber-survivability kill chain can be implemented in the same way as the previously discussed ACS probabilistic kill chain for kinetic weapons (Ball, 2003). The purpose of this cyber kill chain serves the same purpose as the traditional probabilistic kill chain: “to illustrate

the major phases and outcomes that occur in the one-on-one scenario and to define the terminology and the probability associated with each of the phase outcomes.”

A new ACCS kill chain (Figure 6), adapted from the ACS kill chain (Ball, 2003), can be divided into two sections: a susceptibility section and a vulnerability section. For an aircraft to be susceptible to a cyber-engagement, the adversary must perform the following five tasks (Ball & Bryant, 2020b). A weapon must actively search for the aircraft and its cyber-physical components. The aircraft must then be detected in cyberspace, so that the cyber-warhead can be launched and delivered to the aircraft. The cyber-warhead must then be successfully implanted into the target aircraft’s systems. And finally, the aircraft is hit by the cyber-warhead when the implanted malware is triggered, causing component and system dysfunctions. All five of these steps must take place, in that order, for an aircraft to be hit by a cyber-weapon.

For an attrition or mission kill to result from said hit, the dysfunction caused by the cyber-warhead must cause enough functional damage to the aircraft’s cyber-systems that flight or mission essential functions are lost and cannot be recovered. This last step of the kill chain makes up its vulnerability section.




Figure 6. ACCS probabilistic kill chain. Adapted from Ball & Bryant (2020b).

The schematic of the probabilistic cyber kill chain shows both the susceptibility and vulnerability sections as well as the individual phases associated with them. It also shows the process of the cyber-attack as a linear, time-wise sequence. Following the kill chain from top to bottom, if one event does not occur, then the chain is broken and none of the following phases can occur. The outcome in such a scenario would be the survival of the combat aircraft.

In traditional ACS, probabilities are used for each phase of the kinetic ACS kill chain because the outcomes for each phase cannot be known with certainty (Ball, 2003). This uncertainty is amplified in the cyber domain because of the covert nature of attacks, the rapidly changing level of cyber capabilities, and the lack of historical data of cyber-attacks on combat platforms. This makes determining probabilities much more difficult to develop. Nonetheless, ranges of probabilities can be applied to each phase of the cyber kill chain to help estimate the overall likelihood that an attempted cyber-attack would result in an aircraft kill (Bryant, 2019). These probabilities measure the likelihood of success or failure at each point along the kill chain, and can be multiplied together to determine the

overall survivability or killability of an aircraft. An aircraft will be killed by a cyber-warhead if and only if each of the phases produce successful outcomes. This is an important concept for designing survivable cyber systems. To enhance cyber survivability, the probability that one or more of these phases is successful must be reduced.

The probabilistic outcomes associated with the cyber kill chain can be used in the same way that they are used for the traditional ACS kill chain (Ball, 2003). There are only two possible outcomes, success or failure, for each phase of the kill chain. Each outcome for each phase is represented by a branch of the kill chain and has a probability associated with it. For a given phase, the success and failure probabilities must add to one. Likewise, the probability of survival and the probability of kill must also add to one. Determining the probability that the cyber-warhead kills the aircraft is done by simply multiplying the success branch probabilities.

Because of the higher uncertainty associated with cyber-attacks, confidence intervals on probability distributions can be used instead of single values for each phase (Bryant, 2019). Experts in the cybersecurity field can come up with reasonable and credible estimates for these probabilities. These ranges of probabilities can then be used to model cyber-attacks on combat aircraft using Monte Carlo simulations. This idea is discussed in depth in Chapter VI.

Starting from the top of the kill chain, P_A represents the probability that an adversary has a cyber-weapon that is actively searching for the target (Ball & Bryant, 2020b). This probability includes the likelihood that an adversary is able to develop a cyber-weapon. Through reconnaissance, adversaries can determine how to gain access and exploit one or more of the aircraft's cyber-systems. Without the knowledge of how an aircraft's cyber-systems work, how to gain access to those systems, and how to exploit the data or code on those systems, an adversary would be unable to detect, launch, implement, hit, and kill the combat aircraft with their cyber-warhead. If an adversary is able to successfully develop an anti-aircraft cyber weapon and if that weapon is actively searching for its target, the attack moves down to the next phase of the kill chain. A failed first phase results in the survival of the aircraft.

The second phase of the cyber kill chain is the detect phase. This phase characterizes the likelihood that a developed cyber-weapon is able to detect the aircraft in cyberspace and detect a pathway to access the aircraft's internal cyber-systems. Detecting aircraft in the cyber domain can be more difficult for adversaries because aircraft are not continuously connected to the Internet. While air-gapping cyber systems make detection more difficult for the attacker, it does not guarantee aircraft survivability (Bryant & Young, 2019). As mentioned previously, there have been instances where seemingly secure, air-gapped cyber systems were successfully attacked (CSFI, 2010).

Conditional probabilities are used to more easily associate probabilities with specific phases. $P_{D|A}$ denotes the probability of successful detection *given* that the active phase already took place and was successful. Analyzing the probabilities in this way allows us to address each branch of the kill chain without reliance on the outcomes of previous phases. The overall probability that the cyber-weapon can detect the aircraft and a sufficient pathway (P_D) is the probability of active multiplied with the probability of detection given active. Again, a scenario in which the active reconnaissance phase is successful, but the detection phase is unsuccessful will result in the survival of the aircraft. If all previous phases are successful, we move down the kill chain to the next attack phase.

The third phase of the cyber kill chain is the launch phase. The launch phase occurs after the aircraft and a potential pathway have been detected in cyberspace and the adversary launches their operation. During the launch phase, the malware on the cyber-warhead is sent to the target through cyberspace. There are multiple ways of launching a cyber-warhead and the malware may infect multiple hosts before finding its way to the target system. The probabilities associated with the launch phase and the way that these probabilities are used follow that of the previous two phases. $P_{L|D}$ is the probability of successful launch given successful detection. $P_{L|D}$ is independent of the outcome of other phases of the cyber kill chain. P_L is the probability of a successful launch and is the product of $P_{L|D}$, $P_{D|A}$, and P_A .

The next phase after a successful launch is the implementation phase. During this phase, the cyber-weapon reaches the aircraft and successfully implements the cyber-

warhead's malware. Implementation is a necessary step for cyber-attacks because many attacks do not go into effect immediately. Well-designed cyber-attacks can be installed on a given systems well before they are triggered. The adversary's goal in this phase is to exploit an unknown flaw in the aircraft's system (zero-day attack) and implement a way to maintain access for an extended period of time. $P_{I|L}$ represents the probability of successful implementation given a successful launch. P_I represents the total probability of implementation and is the product of all of the conditional probabilities from the phases that precede it.

The final phase of the susceptibility section of the cyber kill chain is the hit phase. This phase starts, not when the cyber-warhead reaches the target, but when the cyber-warhead is triggered, after the adversary has successfully installed the malware on the aircraft's cyber-system. Implanted malware can go days, months, or even years without being detected before an adversary decides to trigger its effects. The triggering of the cyber-weapon results in component or system dysfunctions, which lead to functional damage of the aircraft. Depending on the aircraft's level of vulnerability to cyber-based attacks, it is this functional damage that may or may not result in a mission or aircraft kill. The susceptibility of an aircraft is synonymous to the probability that the aircraft is hit during an engagement, where P_H is the product of each of the conditional probabilities from previous phases up until this point (Equation 4) (Ball, 2003).

$$P_H = P_A \times P_{D|A} \times P_{L|D} \times P_{I|L} \times P_{H|I} \quad (4)$$

Like each of the other phases, $P_{H|I}$ represents the probability the aircraft is hit given the successful implementation of the cyber-warhead.

The vulnerability phase, or kill phase, of the cyber kill chain is what happens after the aircraft has been hit and the dysfunction mechanisms associated with the cyber-warhead have begun to take effect. If the cyber-warhead makes it this far down the kill chain, this phase determines whether an aircraft hit by a cyber-warhead survives or is killed. Vulnerability in ACCS is synonymous with the probability of an aircraft kill given a hit from a cyber-weapon (Ball, 2003). If the damaged aircraft is able to suppress the dysfunctions and/or recover from the attack, the engagement will result in survival. If the aircraft is unable to withstand this damage, the engagement will result in a kill.

For an aircraft to survive a cyber-attack, all that is needed is for the cyber kill chain to be broken somewhere along the way. If methods are in place to stop an adversary from detecting an aircraft in cyberspace, the engagement will result in survival. If a cyber-weapon is developed and launched, but is unable to be successfully implanted on an aircraft's cyber-system for an extended period of time, the engagement will result in survival. Aircraft survival can result even if the cyber-weapon successfully hits an aircraft, but the resulting dysfunctions are sufficiently managed and the aircraft is able to recover. While it is better to break the kill chain earlier rather than later, it is important to design for any scenario. Today's adversaries are highly capable and persistent; we should not underestimate their ability to launch sophisticated cyber-attacks.

D. DYSFUNCTION MECHANISM AND POTENTIAL ATTACK VECTORS OF A CYBER WEAPON

Damage mechanisms of traditional anti-aircraft KEWs produce kinetic effects. In Chapter II, damage mechanism was defined as the physical entity that causes damage to the target (Adams, 2019b). Historically, the kinetic warheads attached to anti-aircraft weapons used one or more of the following damage mechanisms to take down combat aircraft: 1) metallic penetrators and fragments, 2) incendiary materials, and 3) blasts (Ball, 2003). Cyber weapons, however, rarely cause physical damage. They cause functional damage, or dysfunctions. The result of a successful cyber-attack on a combat aircraft would be the dysfunction of critical cyber-components due to some sort of malware. The dysfunction of these critical components could then lead to either an attrition kill or mission kill.

In the cyber domain, a dysfunction mechanism is a term used to describe the output of the cyber-warhead that causes functional damage to the target. Dysfunction mechanisms for anti-aircraft cyber weapons depend on the type of malware implanted on an aircraft's cyber-physical systems. Theoretically, a cyber-warhead can be designed to 1) functionally damage an aircraft's critical components, 2) disrupt or change information that will cause operators to make poor decisions, or 3) falsely represent data that will force the pilot to land the aircraft (Alexandrovich, n.d.). While these dysfunction mechanisms describe some of the means of producing functional damage, they are not all-inclusive. New, more

sophisticated, and more creative cyber threats are developing every year, introducing more potential dysfunction mechanisms to an already growing list.

For dysfunction mechanisms to cause functional damage to an aircraft, the cyber-warhead must first take advantage of one of an aircraft's many attack vectors. Attack vectors are routes that cyber weapons may take to infiltrate and exploit a target's cyber-systems or network. Attack vectors on combat aircraft include 1) specific, targetable cyber hardware, 2) communication channels, 3) software patches, and 4) the supply chain (Alexandrovich, n.d.).

One example of a specific piece of targetable hardware on an aircraft is an electronic flight bag (EFB). An EFB is an electronic information management device used to help operators manage and perform tasks more easily and efficiently. Adversaries could get access to a pilot's EFB by targeting that specific piece of hardware with advanced infiltration or intrusion techniques. The attacker could then manipulate the information seen on the EFB or degrade its capabilities and render the equipment useless. The terminal effect of an attack on an EFB might be poor decision making by the pilot or crewmembers based on false information. If the attack rendered the system inoperable, the pilot may be forced to land without completing its mission. This attack vector applies not just to EFBs, but to any piece of hardware that is targetable in the cyber domain and is flight or mission essential.

Communication channels between aircraft or between the aircraft and ground crew provide another route for adversaries to launch their cyber-attack. Many communication channels create unprotected, open links that can be intercepted by an adversary. Increasing the use of new communication and data transfer technologies puts operators at a higher risk of those channels being hacked or compromised. The figure below (Figure 7) outlines the attack surface of a general combat aircraft, but the cyber-physical systems and open communication channels will vary depending on the platform and mission.




Figure 7. Communication channels attack surface. Source: Ball and Bryant (2020b).

Aircraft communication channels are not always encrypted or protected and contain highly sensitive information. There is a possibility that, if an attacker were to intercept these signals, they could inject false information or cut off the communication lines altogether (National Business Aviation Association, 2016). Additionally, attackers may not have to send all elements of their weapon through the same access point (Ball & Bryant, 2020b). If an aircraft's network of cyber components is not thoroughly air-gapped, attackers may be able to gain access using one access point and deliver the cyber-warhead on another.

Software updates provide adversaries with another potential attack vector. Software bugs are relatively common and can have disastrous effects. Take for instance, the Boeing 737 Max. A seemingly minute software dysfunction involving onboard sensors and flight controls combined with a lack of sufficient manual redundancy or training resulted in multiple crashes and many lost lives (Disis & Pham, 2019). Now imagine if a highly capable adversary, who knew of the software patch and could reverse engineer what problems the patch was fixing, intentionally planted these bugs. The results could be catastrophic. Software updates are never perfect. They can highlight vulnerabilities in older

software versions and introduce new ones. Additionally, when aircraft software is uploaded, the air-gap breaks and the aircraft becomes open to potential cyber-attacks.

The supply chain also introduces a unique attack vector (Alexandrovich, n.d.). A supply chain attack is an attack that involves physically tampering with cyber components during the handling, distributing, or manufacturing stages of production. The intent of a supply chain attack is to install undetectable malware during the production phase, before the equipment becomes operational (Pipeline & Gas Journal, 2012). Remember, implanted malware can remain dormant for months or years before being triggered.

The cyber-components of a combat aircraft can come from various production facilities. Several people may have access to these systems throughout different stages of production. Attackers can take advantage of less-secure elements of a complex supply network to launch their attack. Insider threats are not uncommon within the DoD. If an adversary could pretend to be a hard-working employee and earn access to a combat aircraft's cyber components, they could modify data or add malicious code that could later be triggered to cause functional damage to an aircraft's cyber-systems.

V. DESIGNING FOR THE CYBER THREAT USING AIRCRAFT COMBAT SURVIVABILITY FUNDAMENTALS

The purpose of Aircraft Combat Survivability is to design and build combat systems that can avoid or withstand hostile actions in a man-made, hostile environment (Ball, 2003). To enhance the survivability of combat aircraft, either susceptibility or vulnerability need to be reduced. For traditional KEWs, the moment the weapon hits (or fuses in close proximity to) the aircraft is the threshold used to differentiate when susceptibility stops and vulnerability starts. A combat aircraft is less susceptible if it can avoid being hit. A combat aircraft is less vulnerable if, given a hit, it can withstand the damage from the weapon and prevent an aircraft or mission kill. In Chapter II, susceptibility and vulnerability reduction concepts were briefly discussed for enhancing survivability against kinetic attacks. In this chapter, those concepts will be modified and redefined to enhance survivability against cyber threats.

The ACS susceptibility and vulnerability reduction concepts are universal to attacks from traditional KEWs, but many of the same concepts can be applied to attacks from cyber-weapons as well. Table 2 highlights the differences between the susceptibility and vulnerability reduction concepts of ACS and the proposed susceptibility and vulnerability reduction concepts for ACCS (Aircraft Cyber Combat Survivability). These concepts, if applied effectively, will help program managers and engineers design more survivable combat platforms.

Table 2. ACS versus ACCS survivability enhancement concepts

ACS	ACCS
Susceptibility Reduction Concepts	
1. Situational Awareness 2. Signature Reduction & Control 3. Noise Jamming and Deceiving 4. Expendables 5. Threat Suppression 6. Weapons & Tactics, Flight Performance, and Crew Training & Proficiency	1. Situational Awareness 2. Signature Reduction/Management 3. Cybersecurity Hardening 4. Deception and Decoys 5. Threat Suppression 6. Tactics, System Performance, and Crew Training & Proficiency
Vulnerability Reduction Concepts	
1. Component Redundancy (with Physical Separation) 2. Component Location 3. Component Elimination or Replacement 4. Component Shielding 5. Damage Suppression 6. Recovery	1. System Redundancy with Diversity 2. Component Location and Logical Separation 3. Component Elimination or Replacement 4. Component Shielding 5. Dysfunction Suppression 6. Recovery

While there are many parallels between the susceptibility and vulnerability reduction concepts for ACS and ACCS, the functions and implementations of the cyber-based susceptibility and vulnerability reduction concepts are much different.

For this discussion, it is essential to define where along the cyber kill chain susceptibility stops and vulnerability starts. A “hit” by a cyber-weapon is defined as the moment the malicious code is triggered by an adversary. This means that a combat aircraft vulnerable to a cyber-attack is unable to withstand the dysfunctions caused after a cyber-warhead has been triggered. Susceptibility, then, is used to define the aircraft’s inability to avoid the detection, launch, implantation, and triggering of a cyber-weapon.

A. SUSCEPTIBILITY REDUCTION CONCEPTS FOR ACCS

1. Situational Awareness

Situational awareness is the ability to be alert and responsive to one's surroundings and the current threat environment. Just as kinetic weapons produce detectable signals that can indicate whether an aircraft has been targeted, cyber weapons produce emissions and unique signals that can indicate if an aircraft is in danger of a cyber-attack (Ball & Bryant, 2020c). However, cyber-attacks differ from kinetic attacks in that attackers put much more effort into hiding the presence and effects of cyber-attacks. Modern combat aircraft also routinely fail and operate unexpectedly, adding additional challenges for determining if the cause of the failure was routine or from a hostile cyber-attack. While this makes it more difficult for operators to be aware of active cyber threats, there are still effective reporting and monitoring techniques that can be implemented to enhance situational awareness.

Before a cyber-attack is triggered, active monitoring, cyber threat warning, and attestation measures can be used to reduce an aircraft's susceptibility to cyber-attacks. Active monitoring is a type of systematic observation in which regular surveillance of cyber-physical systems is maintained over a long period of time. Monitoring systems can operate on exterior IT support systems or on interior combat systems built into the aircraft (Ball & Bryant, 2020c). Active monitoring can be done with or without the use of personnel. Personnel can be used to periodically check networks and observe communication traffic for hostile presence. Experts can use various IT-monitoring techniques to observe data, signatures, and logs to detect changes in code or data. Automated software can also be installed on the aircraft itself to monitor files that should not change and notify operators if and when they do. Cyber threat warning is the urgent communication and acknowledgment that an immediate cyber-threat is present (Bryant & Young, 2019). Warning aircraft operators of an imminent cyber threat gives them time to respond to the threat.

Attestation, the act of showing that something is true, is any method of verifying that code is unaltered and can be a powerful tool to warn operators of potential cyber threats (Bryant & Young, 2019). Attestation on an aircraft works by injecting data into a friendly

cyber-system. The timing of the response and the response itself can then indicate if the code had been changed or altered in any way. Attestation is very similar in practice to hashing in traditional IT cybersecurity. Attestation is a one-way function that uses a complex mathematical algorithm to create a string of values (referred to as hashes) unique to a certain file, program, or software code (Nohe, 2019). No two files can create the same string value. Any alteration of code or data, a comma out of place or a 1 that should be a 0, will produce a very different hash value. If the string that is returned from the attestation algorithm matches what is expected for that cyber-system, it verifies that the data has not been altered, modified, or tampered with. Attestation helps ensure that those who are supposed to be altering code (for software updates or patches) are the only ones doing so. While attestation encompasses not just one, but any number of methods of verifying the integrity of data, the primary purpose of verifying the data is to warn operators of potential cyber threats. A hash that has changed unexpectedly is a sign that foul play from a threat actor may be occurring.

Situational awareness also involves being aware of the global landscape, the threat actors that may be operating in a contested cyber environment, and the methods they typically use to launch cyber-attacks. If operators are properly educated of the current cyber threat environment, they can take necessary precautions to reduce the likelihood of actors successfully implanting malware on protected systems.

One of the first steps to addressing a cyber-attack is being aware of who the threat actors are and recognizing when a threat is present. To be situationally aware, operators must understand how critical cyber systems operate. Operators should be able to identify parts of their cyber systems that may be susceptible to a cyber-based attack and be able to recognize when an imminent cyber-threat may be attempting to access a system. Because cyber-attacks are designed to be covert, the importance of having good situational awareness in the cyber domain cannot be overstated. With malicious code installed, combat systems may appear to operate normally until the enemy chooses to initiate the attack. Even then, it may be difficult to differentiate a cyber-based attack from a typical hardware or maintenance issue. Having situational awareness is critical; the proper response to a cyber-based attack cannot begin to happen if the cyber-attack is never found or recognized.

Having good situational awareness of an aircraft's cyber-physical systems reduces its susceptibility to a cyber-attack by increasing the likelihood that malware is discovered and removed before the cyber warhead has a chance to be triggered by an adversary.

2. Signature Reduction/Management

Signature reduction/management is the ability to reduce or regulate the observability of an aircraft and its cyber-physical components in cyberspace. In traditional ACS, combat aircraft give off visual, IR, radar, and acoustic signatures which can be observed by adversaries and their air-defense systems (Ball, 2003). Reducing an aircraft's signature reduces how observable the aircraft is to an air-defense system and therefore improves survivability by reducing the likelihood that the combat aircraft is detected.

A similar approach can be applied in the cyber domain. By making an aircraft's cyber systems less observable to an adversary, the code on those systems become more difficult to access and exploit and therefore harder to hit in cyberspace. For a cyber-weapon to be effective, it must successfully search for, locate, and deliver malicious code to an aircraft via an attack vector. Managing the cyber signature of combat aircraft reduces the likelihood that adversaries can locate and access a combat aircraft's cyber-systems.

Computers, their networks, and the data traveling across them are observable in cyberspace. An aircraft's cyber emissions include any signal or data transfer that takes place in cyberspace. Adversaries often use active cyber scanning techniques to locate hardware in cyberspace, but cyber defenses can be put in place to block, track, and disrupt an adversary's attempt at locating an aircraft's protected systems.

The most common way of reducing an aircraft's cyber signature is through "air-gapping" (Ball & Bryant, 2020c). Air gapping involves physically separating and isolating a secure network from unsecure networks. Air gapping, however, should not be the only protective measure taken to ensure survivability. Many systems are presumed to be adequately air gapped, when in fact the "protected" systems are merely a well-hidden member of an overall unsecure network. If an aircraft's cyber system connects to a ground system that connects to another, more vulnerable system, then an attacker has a path for launching an attack. The first step toward reducing an aircraft's cyber signature should be

examining the connection points that the aircraft has with external systems (Ball & Bryant, 2020c).

Cyber signature can also be effectively managed by limiting non-essential emissions and hiding essential emissions. More advanced methods for managing an aircraft's cyber signature include changing network and address naming structures so that an aircraft's location in cyberspace changes over time (Ball & Bryant, 2020c). A constantly changing target is much harder to hit than a stagnant one. Periodic reconfigurations of cyber structures might render an adversary's thorough reconnaissance useless. By reducing or managing an aircraft's cyber-signature, the likelihood that the cyber-weapon can detect an aircraft and successfully launch a cyber-warhead at it is reduced.

3. Cybersecurity Hardening

Cybersecurity hardening involves the selective restriction of access to cyber-systems and the information contained therein. The process of restricting access across the cyber domain is done so that only friendly forces can access and/or manipulate an aircraft's systems or data. This concept involves assessing platforms, missions, and network systems and applying firewalls, implementing security tools, tightening policies, and adding security software or hardware to at-risk systems to make them more difficult for an adversary to access. The goal of cybersecurity hardening is to reduce the attack surface by reducing the number of points that the system can be attacked from. Setting up firewalls, limiting administrator privileges, whitelisting and blacklisting, and employing multi-factor authentication are examples of methods that can be used to control access and reduce the attack surface of our systems (Bryant & Young, 2019).

Before securing the aircraft's cyber components, defenders should first focus on securing supporting IT equipment using traditional cybersecurity measures. If an adversary can gain access to the supporting equipment that connects to the combat aircraft, they can potentially gain access to the aircraft itself. Cybersecurity hardening measures that can be taken to protect supporting IT equipment include firewalls and limiting administrative privileges.

Firewalls are network security systems that monitor and regulate incoming and outgoing traffic based on a set of security rules. Firewalls act as barriers between trusted internal networks and untrusted external networks, so that not all information and users on the internet (or other untrusted external networks) can access protected military networks.

Administrator privileges are the highest level of rights granted to a user of a computer or network. Access to administrator privileges would give adversaries the ability to make significant unauthorized changes to an aircraft's cyber-components or large software programs such as flight database management systems. Administrator privileges might also allow an adversary to navigate protected, internal networks to gain access to more hardware and software, increasing the scope and potential damage of a cyber-attack. Administrator privileges should be well protected with multi-factor authorization and the number of people with administrator privileges should be limited to reduce the likelihood that an adversary is able to exploit human vulnerabilities.

Locks and login credentials are powerful mechanisms for controlling access to critical flight hardware or software, but locks and login credentials have to be strong to be effective. Multi-factor authentication is an access control method that requires a user to offer more than one piece of evidence to access a system or network (Bryant & Young, 2019). A combination of knowledge (password or security question that the user knows), possession (a key or government ID that a user holds), and inherence (something that the user is such as facial recognition or fingerprint scanning) creates authentication mechanisms that can be used for multi-factor authentication (Bryant, 2016).

The second part of cybersecurity hardening requires defenders to harden the cyber systems on the aircraft itself. One of the more prominent attack vectors for cyber-attacks is software updates. Hardening an aircraft's cyber systems should include verifying that software updates and code loaded on to the aircraft is well-trusted (Ball & Bryant, 2020c).

Additional methods of securing an aircraft's cyber systems include whitelisting and blacklisting. Whitelisting is the practice of explicitly allowing only known, trusted entities to access a particular system (Bryant & Young, 2019). Whitelisting solutions allow only previously approved and trusted programs and code to run on an aircraft's cyber systems.

If the code comes from another source, it will not run. Blacklisting is the opposite; it prohibits known, suspicious entities from accessing particular systems. Blacklisting solutions can only be applied after malware has been detected; it is therefore not the preferred method for hardening an aircraft's cyber systems, but can be useful in protecting IT systems. Whitelisting and blacklisting are additional measures to control who can and cannot access flight or mission-critical aircraft systems. Cybersecurity hardening reduces susceptibility by decreasing the probability that an adversary is able to successfully access and implant malware on protected cyber systems.

4. Deception and Decoys

Deception is the act of causing an adversary to believe that something false is true in order to gain an advantage. In most cases, deception involves feeding false data back to the enemy to have them believe that they are successfully attacking your systems, when in fact they are attacking a decoy or a system that has already been protected from said attack (Bryant & Young, 2019). With kinetic weapons, it is fairly easy to determine if an attack was successful or not because the extent of the damage can be observed. With cyber weapons, battle damage assessment is much more challenging (Ball & Bryant, 2020c). The perceived success of an attack is generally reliant on information fed back to the adversary. If defenders can return misleading information or inject noise into the feedback signal, attackers will begin to doubt their cyber capabilities altogether, potentially reducing the number of future attacks (Libicki, 2007). Adversaries will begin to doubt the validity of any feedback after being successfully deceived once.

Decoys are fake versions of cyber-systems, strategically placed to trick or confuse an adversary into attacking a false system. One of the most commonly used decoys in the cyber domain is a honeypot. Honeypots are fake computers or networks that are used as bait to attract a cyber-attack. These honeypots have no operational significance, but traffic on these fake networks lets defenders observe the techniques and objectives of attackers. This information is extremely useful for designing more secure systems in the future because defenses can be tailored to specific threats. Trapping the weapon in a honeypot also allows defenders to check their real systems for similar malware. The knowledge of

an attempted cyber-attack that can be attributed to a particular nation-state or group also can be used as leverage.

Deceiving adversaries using false cyber-systems reduces susceptibility by decreasing the likelihood that the cyber warhead is successfully implemented on real systems. Learning how an adversary carries out a cyber-attack gives defenders the ability to tailor countermeasures to specific threats, decreasing the likelihood that future attacks will be successful.

5. Threat Suppression

Threat suppression is an offensive effort that degrades the performance of an opposing forces' cyber-weapon capabilities below what is needed to accomplish their objectives (Ball, 2003). The purpose of cyber threat suppression is to attack an adversary's systems before they can attack yours and the goal is to disrupt, deny, or destroy the capabilities and cyber infrastructure needed for an adversary to carry out a successful cyber-attack. Threat suppression can be done either kinetically or in the cyber domain, though a kinetic response to a cyber-threat may often be ineffective.

Given the covert nature of cyber-attacks and the challenge of locating the exact source of a piece of malware, threat suppression by kinetic means (i.e. strategic bombing of enemy hardware) would be very difficult. It is also likely that an adversary would have multiple computers that they could launch the attack from. However, the threat of a kinetic or cyber-based retaliation may sufficiently deter an adversary from attempting a large-scale cyber-attack in the first place. A rival nation-state might not launch a sophisticated, anti-aircraft cyber campaign against the U.S. for fear that they might be caught. The fear of retaliation through kinetic or cyber means or through influential economic sanctions may be enough to discourage the attack in the first place.

Threat suppression via the cyber domain would require highly skilled personnel. The cyber professionals would have to be able to gain access and exploit enemy systems to prevent hostile cyber-weapons from achieving their intended effects. Successful threat suppression does not necessarily require entirely destroying enemy hardware or software. Just a slight modification to an adversary's cyber weapon could render it completely

useless. Threat suppression reduces susceptibility to cyber-attacks by disrupting, denying, or destroying an adversary's cyber capabilities before the cyber-warhead is able to be triggered. In short, threat suppression aims to break the cyber kill chain by breaking the cyber weapon.

6. Tactics, System Performance, and Crew Training and Proficiency

This susceptibility reduction concept encompasses specific actions or strategies used to address cyber threats, the performance of the cyber-systems attached to a combat aircraft, and the proficiency of the military personnel that handle those cyber-systems. One relatively easy tactic for potentially improving system performance is to periodically update and reconfiguring combat networks. Accessing and exploiting a system in a dynamic cyber domain that is constantly moving and changing is much more difficult. Malicious code designed to operate against an older version of combat software may become obsolete against a newer or different network.

Another tactic that can be used to reduce susceptibility is incorporating friendly “red teams” that act as adversaries (Bryant & Young, 2019). Red teams are groups of white-hat hackers that attack digital infrastructure as an attacker would. One of two things can be learned from this sort of penetration testing of cyber-systems. Either the red team would be unable to access an aircraft's protected cyber systems, verifying a low level of susceptibility, or the red team would be able to access an aircraft's protected cyber components, highlighting previously unknown areas of susceptibility that can now be fixed.

Patching known flaws in susceptible combat software is a way to enhance system performance. Patching is the process of discovering, removing, and fixing areas of susceptibility in software. While patched software can be beneficial, poorly designed, poorly tested, and rushed patching software can open more attack vectors for a highly-skilled adversary to exploit. A poorly designed patch could end up being more susceptible than the original software. Adversaries may be able to reverse engineer the patch to discover the weakness the patch was meant to protect. They can then use this information to attack cyber-physical systems that have not been updated with the new patch yet.

The training and proficiency of operators and maintenance crew members is also a vital concept that should not be overlooked. Regardless of how reliable our weapon platforms are on secure cyber-systems, there is always a human element. Humans, by nature, are not perfect. Training and proficiency of crew members involve improving the overall awareness of every service member's role in keeping cyber-systems secure. Cyber-based weapons are often delivered through web servers or through emails, USBs, or social media sites. Sufficient cyber-awareness training reduces the likelihood that an adversary is able to access password-protected systems or carry out a successful phishing attack. Cyberspace is inclusive of not only the code and data that travels through servers, but also the physical hardware needed for such data links. Protecting this hardware is equally as important as protecting the data on it. Training and proficiency include having security personnel well trained to protect essential aircraft hardware from adversary reconnaissance and insider threats. One annual online cyber awareness training is not sufficient for protecting our military's systems that rely so heavily on the cyber domain.

B. VULNERABILITY REDUCTION CONCEPTS FOR ACCS

Reducing an aircraft's cyber vulnerability presumes that the attacker has already triggered the cyber-warhead's dysfunction mechanisms. Thus, these concepts are in place, not to prevent cyber-attacks, but to detect, respond to, and limit the functional damage of those cyber-attacks on combat aircraft in an attempt to reduce the likelihood that a hit results in a mission or aircraft kill.

1. System Redundancy with Diversity

System redundancy is the inclusion of hardware and software components that are not strictly necessary for an aircraft to function, but are included in the event that one or more flight or mission essential cyber systems fails (Ball, 2003). If one cyber component is compromised, another system with the same functionality can take its place. Redundant cyber systems should already be running, or should be capable of coming up to speed quickly, so that operators can switch over to the redundant system with minimal time lag. The concept of redundancy can be applied at the system, subsystem, or component level and can either be fully redundant or partially redundant (Ball, 2003). Fully redundant

systems are able to entirely perform the same functions as the primary system, while partially redundant components can only perform part of the functions that the primary system could otherwise perform.

Having purely redundant systems is not enough however. If the redundant systems are identical to the primary systems with the same vulnerabilities, a single cyber-weapon could easily take out both out with the same malware (Bryant & Young, 2019). The same attack vectors can be exploited if redundant systems do not implement some level of diversity. Either having different hardware/software or having multiple versions of software available is necessary to reduce an aircraft's vulnerability to a cyber-attack. While redundant systems should not be the same, they should have the same functionality as the components they intend to replace.

One example of system redundancy in the cyber domain is backing up mission-critical files and data, so that if those files are damaged, an older, back-up copy can be used with minimal effect. An advantage of this method of redundancy is that it can be entirely software based. It can be implemented without having to add equipment or hardware to the aircraft. However, there may be a small time delay as back-up files overwrite and reload software.

Another potentially powerful way of implementing system redundancy is through virtualization (Bryant & Young, 2019). Virtualization is the process of running multiple virtual machines on a single piece of physical hardware. All of the virtual computers run as if they were the real computer, so that if the main computer is compromised, another unaffected computer is readily available. Virtual systems require a host operating system, a hypervisor, and a guest operating system (Xing & Zhan, 2012). Unfortunately, access to the hypervisor can give an adversary access to all of the virtual machines and the data associated within them, so protection of the hypervisor is absolutely critical.

The host operating system provides all of the computing resources that enable the host and its virtual machines to run simultaneously and share information with one another. This can be a powerful tool for several reasons. Having multiple systems running simultaneously allows corrupted data and infected computers to be easily discarded and

replaced. It also enables the computers to check each other's work. Theoretically, all of the computers should be running identically. A computer that is producing different results from the network of virtual machines indicates that a cyber-attack may be present. Redundancy of cyber components reduces vulnerability by reducing the likelihood that the dysfunctions caused by an attack on primary cyber systems cannot be restored by redundant systems that offer similar functionalities.

2. Component Location and Logical Separation

Component location is the physical or logical positioning of critical cyber components such that the probability that a dysfunction mechanism will result in a mission or aircraft kill is reduced. In traditional ACS, the concept of component location involves shielding critical components with non-critical components or physically separating critical components such that a single hit is less likely to take out multiple components (Ball, 2003). There is some importance to the physical location of cyber hardware, but the logical separation of the cyber-system's data and software is a more important concept for reducing cyber vulnerability. There have been instances of physical damage caused by cyber-based attacks. In the event of physical damage, physical separation of cyber components is important so that the damage done to one piece of hardware does not also damage nearby hardware.

While the actual location of cyber hardware is important in the event that a cyber-attack causes physical damage, the virtual location of the data and software that runs on that hardware is even more important when designing survivable cyber systems. Cyber logic refers to the infrastructure organization, software architecture, design pattern, and communication networks between cyber-physical systems on an aircraft. Similar to the purpose of physical separation, logical separation can be used to reduce an aircraft's vulnerability by preventing a single cyber-attack from denying, disrupting, or destroying multiple components. The easiest way to implement logical separation is by physically separating an aircraft's cyber networks.

Logical separation of cyber components prevents an adversary from navigating a protected combat cyber network by reducing its level of connectivity. Compartmentalizing

damage is key to the concept of component location. Cyber components do not necessarily need to be air-gapped from one another (although this would be highly effective), but logic software should be in place that prevents an adversary who has gained access to one cyber-component from gaining access to other components.

3. Component Elimination or Replacement

Completely getting rid of a non-critical cyber component or substituting a critical cyber component with a less vulnerable component with similar functionality reduces an aircraft's vulnerability to cyber-attacks (Ball, 2003). While designing physically and logically separated cyber components with sufficient redundancies reduces a combat aircraft's vulnerability to cyber-attacks, the vulnerability is reduced to zero if those vulnerable components are removed altogether. It is impossible to attack hardware or software that is not there.

The key to component elimination is to determine whether or not the component in question is mission-critical and whether or not that component is worth an increased level of risk. For example, downloading maintenance data in flight may be more convenient than waiting until after landing, but having an open communication link operating in flight may not be worth the risk of opening a potential cyber-attack vector (Bryant & Young, 2019). Maintenance personnel are often proponents of sharing more flight and maintenance data across more devices because it improves efficiency and effectiveness; security personnel are often proponents of restricting the amount of data shared across devices for security reasons (Ball & Bryant, 2020c). The level of connectivity between an aircraft and supporting systems should be carefully examined to balance performance and security.

While modern, high-tech combat aircraft may offer attractive additional features, not all of the bells and whistles that come with more technologically advanced aircraft may be worth the larger attack surface. It is important to make these risk decisions early in the program's lifecycle, before millions of dollars are spent on the aircraft's bells and whistles.

An alternative to entirely eliminating a cyber-component is to replace a more vulnerable component with a less vulnerable component that can perform the same functions. Vulnerability is reduced when a component with a larger attack surface (i.e., has

more potential points of access) can be replaced with a component with a smaller attack surface. High-tech cyber components with more sophisticated code and cyber infrastructure might need to be replaced with more robust cyber components that are less likely to dysfunction in the event of a cyber-attack. It is important to keep in mind that a replacement component may introduce new cyber kill modes that were not there previously. The new component might still be less vulnerable, but only if these new kill modes are properly addressed.

4. Component Shielding

Component shielding is the process of guarding an aircraft's critical cyber components from the dysfunction mechanisms of a cyber-based attack using additional protective layers. Component shielding in the cyber domain is analogous to component shielding in the physical domain in that its purpose is to add extra layers of defense to critical components to protect them from damage after a warhead has a successful hit an aircraft (Ball, 2003). Rather than adding a physical armored coating or plating to protect a critical component from KEWs, component shielding in the cyber domain involves adding additional security measures to detect an attack and prevent it from causing a critical component dysfunction. The goal of shielding cyber components is to reduce vulnerability by detecting an attack and preventing it from causing serious damage to an aircraft and its cyber systems through layers of securities, firewalls, and protocols.

Designing hardware with a secure root of trust is one way of shielding components from cyber-attacks. For secure roots of trust, thorough identification and authentication are required before external devices can connect to flight hardware. This helps ensure that communication channels are secure. Another method for shielding an aircraft's cyber components is by adding switches, certificates, and signings to control who can and cannot write to a component's memory (Ball & Bryant, 2020c). These security measures prevent writing system updates unless the update is physically activated by a trusted operator.

An important part of component shielding in ACCS is the ability to sense and detect when an aircraft's cyber systems are under attack. In the physical domain, it is relatively easy to know whether or not an aircraft was hit by a gun or missile. It is harder in the cyber

domain to recognize if dysfunctions are due to a hostile cyber-attack or due to a typical operator or maintenance issue. It is also more difficult to determine which cyber-systems are being attacked. Cyber-weapons, by nature, are meant to be covert. A well-designed cyber-attack will try to hide the dysfunction or make the dysfunction appear as if it were a routine problem. Having active or passive sensors onboard that can see if cyber-systems are running correctly is vital to improving the likelihood of survival from a successfully implemented cyber-attack. Without the effective sensing and detecting of an attack, it becomes much less likely that the correct course of action will be taken to mitigate the dysfunction mechanisms.

5. Dysfunction Suppression

Dysfunction suppression describes any active or passive design technique or feature that reduces cyber vulnerability by containing functional damage or reducing its effects (Ball, 2003). Passive dysfunction suppression features are features that occur automatically in the event of a cyber-attack. Active dysfunction suppression features require some level of operator action, and therefore require a degree of situational awareness.

Passive design techniques include the use of onboard sensors that can monitor cyber-systems and software tools that can automatically reload and reboot systems under attack. Passive features may also be capable of automatically repairing damaged software or data if it recognizes inconsistencies. Passive dysfunction suppression features should be capable of responding to unanticipated commands and irregular data. Examples of passive dysfunction suppression methods include implementing damage-tolerant components, scrambling techniques, and systems designed to postpone failure.

Damage-tolerant components are components that can accept a degree of functional damage without losing their ability to function normally. This can be done by loading multiple paths of communication onto hardware so that if one path is damaged, the component can still send and receive data via an alternative route. Scrambling devices transpose or invert signals so that data is unintelligible without the complimentary unscrambling device. Scrambling techniques could be used on the binary level to better synchronize data between networks and software (Bryant & Young, 2019).

Synchronization is important because it ensures that two or more locations share the same data. Delayed failure techniques can also be used so that elements of cyber components can still operate for a period of time after they have been attacked.

The active dysfunction suppression features might include the pilot restoring software to a previously saved setting or physically shutting down infected hardware. If an operator is able to accurately notice an attack and determine which systems are being affected, they can respond accordingly. Situational awareness is therefore necessary before operators can effectively implement active dysfunction suppression techniques.

The use of highly-skilled cyber teams can also be used to suppress functional damage from an ongoing cyber-attack (Bryant & Young, 2019). Cyber defense teams would have to intimately know the aircraft's mission and cyber-physical systems and would need remote access to those systems. Cyber teams may be able to remotely monitor and locate an active attack, and respond to it accordingly. While cyber teams would provide the knowledge and expertise necessary for recognizing and responding to attacks, an open link would be required to monitor data flow. This open link between the aircraft and support systems would introduce an appealing attack vector for adversaries to exploit. The benefit of having these cyber teams may not be worth the risk.

6. Recovery

Recovery is the ability of an aircraft to either return to a normal state of operation or be safely grounded so that problems can be further assessed. After a cyber-attack has been successfully implemented, if the functional damage caused by the attack has not been entirely suppressed, the next step is to try to recover the aircraft and its information. This recovery step occurs after the cyber warhead has hit the aircraft and the dysfunction mechanism has taken full effect. If normal operations cannot continue, it is then the job of the pilot and supporting ground crew to get the aircraft grounded so that they can take a more in-depth look into what caused the dysfunction and who launched the attack. If the attack and the resulting dysfunctions do not pose an immediate threat to the survival of the physical aircraft, operators may continue to try to carry out the mission.

One design feature that would improve the likelihood of an aircraft could returning to a normal state of operation is the use of fail-safe responses. Overwriting and reloading essential software may get rid of the implanted malware. If it does not, reverting to a safe, older version of software saved on the device's memory could do the trick.

Another way an aircraft could recover from a cyber-attack is by manually overriding non-essential cyber-systems. Combat aircraft have not always relied so heavily on complex cyber systems. It is naïve to believe that today's combat aircraft could be effective without any reliance on cyber-systems, but some of the more complex cyber components could be backed up with either manual components or components with primitive software and basic functionality. In the event that one of these more complex components is attacked, the pilot could switch over to less cyber-reliant components. This would be analogous to a basic "get home" mode (Ball & Bryant, 2020c). The loss of some functionality is better than losing the entire aircraft.

VI. SIMULATING A CYBER-ATTACK USING THE PROBABILISTIC KILL CHAIN

A. CHALLENGES TO MODELING A CYBER-ATTACK

The ACS engagement level probabilistic kill chain (Figure 2) briefly discussed in Chapter II can assess the susceptibility, vulnerability, and overall survivability of a combat aircraft subject to a single shot from a kinetic energy weapon. This probabilistic kill chain was then modified in Chapter IV to address the phases and associated probabilities of an anti-aircraft cyber-attack (Figure 6). While these probabilistic kill chains are not perfect representations of real-life scenarios, they can be useful models for assessing platform survivability. Defenders can better implement survivability enhancement features to reduce an anti-aircraft weapon's effectiveness if the sequence of steps that a weapon must go through to hit and kill a target is well understood (Ball & Bryant, 2020b).

While the math behind these probabilistic kill chains is relatively simple, assigning accurate probabilities to each phase can be quite challenging (Ball & Bryant, 2020b). This is especially true for ACCS. One of the most tried and true methods for determining the phase probabilities of an anti-aircraft weapon is by using historical data. In ACS, a long, detailed historical record of kinetic aircraft engagements is used to develop reasonable and credible probability ranges. There is no such historical data for anti-aircraft cyber weapons.

Another advantage that ACS has over ACCS in this regard is the ability to accurately model the physics of a kinetic attack and back up computer calculations with live fire testing and evaluation (LFT&E) (Bryant, 2019). Testing the effectiveness of an anti-aircraft cyber-warhead in a laboratory is much more difficult because the cyber domain is so extensive. Most kinetic anti-aircraft threats are known and well researched. This is not the case in the cyber domain. Cyber-attacks are generally more interactive, creative, and complex than traditional kinetic energy attacks, making it much more challenging for defenders to determine a cyber-weapon's probability of success or failure.

For these reasons, the ACCS kill chain's phase probabilities are likely to have much higher levels of uncertainty than what would be expected for kinetic energy weapons.

While some uncertainty is unavoidable without more historical or testing data, there are methods that can decrease a model's uncertainty and help produce useful results.

B. METHODOLOGY

This first and most important step to modeling the ACCS probabilistic kill chain was to accurately determine the cyber-attack's phase probabilities (Figure 6). Because of the current lack of historical data on anti-aircraft cyber weapons, the phase probabilities would have to be approximated using subject matter experts (SMEs) and confidence intervals (CI) (Bryant, 2019).

Given that high levels of uncertainty are expected with cyber weapons, instead of asking experts for probabilities as single point values, a more realistic approach would be to ask them for ranges of expected probabilities. For example, a cybersecurity expert could say with 90% confidence that the probability of launch given detection for a cyber-weapon might be 40-75%. This range of probabilities could then be used as bounds to provide the relative likelihood of outcomes using probability distributions.

For modeling and simulation, the shape of the distribution is very important. We cannot simply assume that the phase probability distributions are normal. In this case, given how much uncertainty there is and how little data there is on the topic, assuming normal probability distributions may be inappropriate.

Likewise, assuming a uniform distribution, in which every outcome in a given range is equally likely to occur, might also be unfitting. A uniform distribution assumes that the likelihood of occurrence between upper and lower parameters is arbitrary. If experts could only confidently predict the probability of launch given detection lay somewhere between 40% and 75%, then a uniform distribution might be appropriate. However, if an expert were able to give a realistic estimate as to where along that 40-70% range they think the probability of success is most likely to occur, a triangular prediction would be a better approach.

Triangular distributions are commonly used in risk and uncertainty analysis when relationships between variables are presumed but there is limited sample data available

(Johnson, 2002). Triangular distributions are typically favored for cases when subject matter experts attempt to reflect the overall subjective probability distributions for variables where limited data is available (Book et al., 2016). Given the high levels of uncertainty, the lack of underlying data, and that expert opinions are inherently subjective, triangular distributions were a good place to start for predicting phase probabilities.

These distributions are based on a minimum expected value, maximum expected value, and an “inspired guess” as to the most likely value (Hesse, 2000). From these three data points, a triangular probability distribution can be created according to Figure 8.




Figure 8. Triangular distribution

Triangular distributions are intentionally not smooth because they are based on only a few data points. It is also important to note that many triangular distributions are not symmetrical. The likelihood of launch given detection might range from 40-75%, but the most likely probability could be on either the lower or the higher end of that. The mean, or average, value of a triangular distribution can be easily calculated using Equation 5,

$$\mu = \frac{a+b+b}{3} \quad (5)$$

where a is the minimum, b is the maximum, c is the modal value, and μ is the mean of the triangular distribution.

To further increase the validity (and decrease the uncertainty) of the ACCS kill chain's phase probabilities, instead of asking one SME, a group of SMEs could be questioned. While one expert might say that the probability of launch given detection is 40-75% with a probability of 60% being most likely, another might think that the probability of launch given detection is closer to 70-90% with a probability of 80% being most likely. By asking many SMEs for their minimum, maximum, and most likely probabilities, a comprehensive triangular distribution could be determined by averaging the experts' opinions for each phase of the kill chain.

One potential shortcoming of using subject matter experts to determine probabilities is the human tendency to be overconfident (Ball & Bryant, 2020b). Psychological studies have shown that when people provide 90% confidence intervals on estimates, the actual values lay outside of their confidence interval more than half the time (Moore, 2018). Fortunately, there are proven techniques that can reduce the effect of overconfidence bias and produce predictions that are more accurate (Hubbard, 2014). These techniques should be taught to the SMEs before they are used to determine the ACCS kill chain's phase probability distributions.

Actual experts were not consulted to obtain phase probability data. Instead, minimum, maximum, and most likely estimates were made up merely to represent possible predictions from twenty SMEs. The 'expert' predictions and resulting triangular distribution for the probability of launch given detection are shown below for reference (Table 3) (Figure 9). Similar data tables were produced for the other five kill chain phases (Appendix A).

Table 3. Sample SME probability data for launch phase

Probability of Launch given Detection, $P_{L D}$	"Experts"	Minimum Probability	Most Likely Probability	Maximum Probability
	A	0.4	0.73	0.94
	B	0.5	0.67	0.8
	C	0.75	0.92	1
	D	0.4	0.68	0.95
	E	0.2	0.48	0.8
	F	0.25	0.53	0.77
	G	0.55	0.61	0.75
	H	0.24	0.5	0.83
	I	0.62	0.71	0.97
	J	0.58	0.68	0.92
	K	0.7	0.93	1
	L	0.76	0.84	0.96
	M	0.1	0.31	0.63
	N	0.29	0.45	0.78
	O	0.27	0.53	0.87
	P	0.65	0.86	1
	Q	0.5	0.75	1
	R	0.7	0.88	1
	S	0.66	0.9	1
	T	0.45	0.69	0.95
	mean	0.4785	0.6825	0.896




Figure 9. Launch phase triangular distribution

The simulation for this paper was done purely for proof of concept. For real-life cyber survivability assessment applications, actual, calibrated SMEs should be questioned and that data should be used to determine the appropriate triangular probability distributions.

Because conditional phase probabilities in the ACS kill chain are single point values, the overall likelihood of events can be determined by multiplying subsequent branches. However, unlike the point values used in ACS, the probability distributions used for ACCS cannot simply be multiplied together (Ball & Bryant, 2020b). Instead, a Monte Carlo simulation was used to randomize the phase probabilities based on each phase's triangular distribution. For each trial, a random probability was assigned to each of the phases based on the respective triangular distributions. For example, the probability of launch given detection for trial one could be anything between 0.478 and 0.896, with a

greater likelihood that the assigned probability would be closer to 0.683 than either of the two extremes (Figure 9).

Another random number was generated between zero and one for each phase. This random number was then compared to the assigned phase probability. If the random number was less than the assigned phase probability, then the trial moved on to the next phase of the kill chain. If the random number was greater than the assigned phase probability, then the kill chain was broken and the attack stopped. If the attack moved on to the next phase, a new phase probability was assigned based off its own respective triangular distribution and was compared to a new random number between zero and one to determine if the outcome was successful or not. This was done for each of the six phases, starting from the active phase and ending after the kill phase (Figure 6). The simulation was then repeated for one million trials and the results were tabulated and averaged. The distribution of calculated survival probabilities was also used to determine the overall uncertainty of the final results.

Monte Carlo simulations typically use tens of thousands or hundreds of thousands of trials. For this model, given unusually high levels of uncertainty for each of six independent phases, more iterations were required. The number of trials necessary for this Monte Carlo simulation was determined using Equation 6 (Liu, n.d.).

$$\bar{Z} \frac{s}{\sqrt{N}} = LOP \quad (6)$$

Where $\bar{Z}$ (3.29) was the z-statistic based off a desired 99% certainty in results, s was the standard deviation (0.0198), LOP was the desired level of precision, and N was the minimum number of trials necessary to obtain the desired level of precision. The standard deviation was determined by running the simulation with only 1,000 iterations and calculating the standard deviation from the probability of survival results. Using Equation 6, at least 425,000 Monte Carlo iterations were required to achieve a level of precision of 0.01%. This number was increased to one million for this study because the simulation algorithm used was relatively simple, computing power was available, and run-times remained reasonably short. A more complex cyber survivability assessment model would not require this many iterations.

C. SIMULATION TEST SCENARIOS

The Monte Carlo simulation was run for three test scenarios. The first test scenario was for an aircraft without survivability enhancement features (SEFs). The second test scenario was for an aircraft with only one SEF. Six total simulations were run for scenario two, each with a single SEF reducing the effectiveness of a different phase of the kill chain:

1. Reduce probability of active by using cyber threat suppression techniques (Threat Suppression)
2. Reduce probability of detection by air-gapping systems (Signature Reduction/Management)
3. Reduce probability of launch given detection by threatening retaliation (Tactics, System Performance, and Crew Training & Proficiency)
4. Reduce probability of implantation given launch by using multi-factor authentication (Cybersecurity Hardening)
5. Reduce probability of hit given implantation by using active monitoring (Situational Awareness)
6. Reduce probability of kill given hit by using virtualization techniques (System Redundancy with Diversity)

Each of the six SEFs were assumed to reduce the likelihood of their respective phases by 15-20%. The SEFs for scenario two were implemented two ways: proportionally and disproportionately. The SEFs were implemented disproportionately by simply subtracting a random number between 0.15 and 0.2 from the assigned phase probabilities for each trial. The proportional method multiplied the assigned phase probability by a random number between 0.8 and 0.85. The second test scenario was done to determine if there were an optimal phase to attempt to break the kill chain.

The third and final test scenario was for an aircraft with all six SEFs. While the survivability of an aircraft is expected to increase as more SEFs are added, a cost-benefit

analysis should be done to determine if the addition of more SEFs is worth the monetary or performance cost.

A copy of the MATLAB script used to run these Monte Carlo simulations can be found in Appendix B.

D. RESULTS AND DISCUSSION

1. Test Scenario One: No Survivability Enhancement Features

For the first test case, no survivability enhancement features were used to reduce the effectiveness of the anti-aircraft cyber weapon. Figure 10 below verifies that the Monte Carlo phase probabilities from the one million trials closely fit the expected triangular probability distributions for each phase of the kill chain.




Figure 10. Triangular probability distributions for each phase

Figure 11 summarizes the results for the first test case. The numbers on each branch represent the number of events that occurred over the course of one million trials. The

numbers in parenthesis represent the conditional probabilities for each branch based on the Monte Carlo simulation. The theoretical probability of survival was determined using the means of the probability distributions above (Equation 5). The experimental probability of survival was determined by subtracting the ratio of killed aircraft to total trials by one (Equation 7).

$$P_{S, Experimental} = 1 - \frac{\text{Number of Aircraft Killed}}{\text{Number of Trials}} \quad (7)$$




Figure 11. Simulating a cyber-attack on an aircraft without SEFs

The experimental probability of survival was only one-hundredth of a percent off from what was expected. We can therefore conclude that one million trials were more than sufficient for determining overall probabilities, even though the uncertainty for any single given trial may have been higher. From this simulation, the anti-aircraft cyber weapon killed 3.11% of aircraft and 96.89% of aircraft survived.

While 96.89% represents the most likely probability of survival, it does not mean that every similar engagement will have exactly a 96.89% chance of surviving. It is important to note that there is still a level of uncertainty in the overall probability of survival. This uncertainty was calculated by multiplying the unique conditional probabilities from each trial together, plotting them on a histogram, and fitting a distribution curve to the results. Figure 12 shows the probability of survival distribution based on the Monte Carlo simulation.




Figure 12. Probability of survival distribution without SEFs

While the mean geometric mean of the data set was 0.969, the probability of survival ranged from 0.851 to 0.997. This means that for any given attack, the probability of survival could differ by almost 15%. Only 1% of survival probabilities however fell below 0.926, so while a probability of 0.851 is possible, it is highly unlikely in this case. The standard deviation was 0.014 and the skewness was -1.0605. The negative skewness means that the probability is more likely to fall on the left side of the peak (or mode) than on the right side of the peak. While averaging the results can provide a realistic single value for the probability of survival given a cyber-attack (0.969), the entire probability distribution gives a more complete picture for assessing survivability. The probability distribution shows how much uncertainty is in the survivability estimation.

2. Test Scenario Two: One Survivability Enhancement Feature

For the second test scenario, one survivability enhancement feature was added to the model. For the first round of simulations, each of the SEFs was assumed to decrease the respective phase probabilities disproportionately. Conditional phase probabilities were simply decreased by 0.15 to 0.20, one at a time. The first run assumed that only P_A could be reduced, the second run assumed that only $P_{D|A}$ could be reduced, and so on. The simulation was repeated for each of the six phases. Figure 13 shows how implementing SEFs at different points along the kill chain effected the overall likelihood of events.




Figure 13. Comparing single SEF effectiveness

The numbers above the ‘killed’ bar represent the total number of aircraft killed out of one million when a single SEF was implemented. As a reference, in the simulation without any SEFs (test scenario one) 31,095 (3.11%) aircraft were killed. As expected, a reduction in any of the six phase probabilities resulted in an overall decrease in the probability of kill. Each case with a SEF resulted in less kills than the case without SEFs. Figure 13 also shows that SEFs implemented at one location in the kill chain had no effect on previous phases. This is because the kill chain is time-wise and linear; it only moves in one direction. In this model, SEFs were limited to reducing only one phase probability. In reality, many survivability enhancement features can reduce a cyber-weapon’s effectiveness at many points along the kill chain.

Figure 14 compares the probability of survival distributions for each of the test scenarios.




Figure 14. Probability of survival distribution with one SEF (disproportional)

Although each SEF reduced their respective phase probability by the same amount (0.15-0.20), they did not all have the same effectiveness. Reducing P_A enhanced survivability the least, while reducing $P_{D|A}$ and $P_{H|I}$ enhanced survivability the most. The reason for this had to do with the initial phases probability distributions determined by the SMEs (Figure 10). The active phase had the highest expected probability of success, while the detect and hit phases had the lowest expected probabilities of success.

This result is important because it can inform defenders of where they should focus their attention. More effort should be placed on further reducing the phases of the kill chain that are believed to have the lowest probabilities of success. This, in effect, bottlenecks the attack at one point, stopping the majority of attacks from advancing further along the kill chain. If one phase's probability can be reduced to zero, it would be impossible to kill that aircraft with a cyber-weapon.

The same test was run a second time with proportional reductions in phase probabilities. Instead of subtracting 0.15 to 0.20 from the phase probabilities, the probabilities were multiplied by 0.80 to 0.85. The results of this test are shown below (Figure 15).




Figure 15. Probability of survival distribution with one SEF (proportional)

Unlike the previous case, if each option proportionally reduces the phase's probability of success by the same amount, focusing on one phase of the kill chain over another has no effect on the overall survivability of the aircraft. The distributive property tells us that when taking the product of a set of probabilities, it does not matter if the 0.825 (average amount or reduction in this case) is multiplied first or last; the outcome will be the same. If all options are expected to enhance survivability proportionally, defenders should choose the easiest or least expensive option.

3. Test Scenario Three: All Survivability Enhancement Features

The third test scenario tested the survivability of an aircraft with all six survivability enhancement features. For this simulation, all six of the phase probabilities were reduced by 15-20%. Figures 16 and 17 below compares two aircraft: one without SEFs and one with multiple SEFs. Figure 16 shows where the kill chains were broken. Figure 17 compares the probability distributions.




Figure 16. Event outcomes for cases with and without SEFs




Figure 17. Probability of survival distributions for cases with and without SEFs

Both figures show that SEFs can drastically increase the survivability of a combat aircraft. Decreasing each phase probability by only 15-20% resulted in a 91.3% reduction in the aircraft's killability. Based on the Monte Carlo results, the probability of survival for the aircraft without SEFs was 96.88%, while the probability of survival for the aircraft with SEFs was 99.73%.

Figure 17 also shows that enhancing survivability can decrease uncertainty. The range and standard deviation were each much smaller for the case with SEFs than for the case without SEFs. The ranges for the cases with and without SEFs were 0.9694-1.0000 and 0.8562-0.9968 respectively; the standard deviations for the cases with and without SEFs were 0.0023 and 0.0140 respectively.

While full-scale testing and evaluation may not be possible for predicting the outcomes of cyber-attacks as they are for predicting the outcomes of KEWs, some testing can be done to help validate cyber survivability assessment results. Red teams are composed of friendly cyber experts that assume the role of an attacker to test a system's cybersecurity. Red teams can be used to imitate enemy cyber-attack techniques in an attempt to detect weaknesses and assess a platform's level of security. By simulating attacks with red teams, analysts can validate or modify predictions and defenders can learn what areas of the kill chain need to be better protected.

As a reminder, the purpose of these simulations was simply to prove that anti-aircraft cyber-attacks could be modeled and to quantitatively show how survivability can be enhanced with susceptibility and vulnerability reduction concepts. While the theory and mathematical process is sound, the numbers used for these simulations were largely made up and should not be taken as fact. Actual SMEs should be consulted to help assign ranges of probabilities to each of the six phases of the probabilistic kill chain.

VII. CONCLUSIONS

The Aircraft Combat Survivability design discipline has been tremendously successful in increasing the cost effectiveness of combat aircraft operating in man-made hostile environments. Until recently, the only serious anti-aircraft threat has been in the form of traditional, kinetic energy weapons; however, the rapid growth of technology, especially in the cyber domain, has given adversaries new tools to disrupt, degrade, deny, and destroy our combat systems. Cyber weapons pose a legitimate and dangerous threat to our combat aircraft's ability to accomplish mission objectives and return home safely.

The US military and its combat platforms have become almost entirely reliant on cyber-physical systems that operate in a domain that has yet to be well defined as a legitimate warfighting battlespace. As our combat aircraft continue to evolve with modern technologies, so too do the threats to those aircraft. Integrating aircraft with newer, more sophisticated cyber-components opens potential pathways for determined, highly capable adversaries to launch cyber-based attacks. For our adversaries, the covert nature of cyber operations and the potential to pack a large punch at a low cost and with low personal risk makes cyber-attacks an appealing alternative to kinetic energy weapons.

Program managers, engineers, pilots, and maintenance personnel should always consider the cyber threat when designing and operating combat platforms. While Aircraft Combat Survivability has proven incredibly useful in defending combat aircraft from kinetic attacks, expanding the design discipline to address the emerging anti-aircraft cyber threat is necessary. We can improve an aircraft's ability to avoid or withstand functional damage from a cyber-attack by applying the same ACS fundamentals to cyber weapons.

While the dysfunction mechanisms and attack vectors for cyber-attacks are fundamentally different from those of traditional, kinetic energy weapons, a modified version of the ACS probabilistic kill chain can be incredibly useful in assessing an aircraft's susceptibility and vulnerability to a cyber-attack. Many of the same ACS susceptibility and vulnerability reduction concepts for enhancing survivability against traditional KEWs can be used to enhance cyber-survivability. While there are many

parallels between the susceptibility and vulnerability reduction concepts for ACS and ACCS, the functions and implementations of the cyber-based susceptibility and vulnerability reduction concepts are much different. These concepts, if applied effectively, will help program managers and engineers design more survivable combat platforms.

Finally, methods for assigning kill chain phase probabilities based on expert opinions and evaluating an aircraft's cyber-survivability using Monte Carlo simulation techniques were proposed. The results from the survivability assessment model showed how survivability enhancement features could improve an aircraft's ability to avoid or withstand functional damage caused by a cyber-attack.

While this research specifically addressed anti-aircraft cyber threats, the ACCS concepts can be applied to any combat platform. Ships and ground vehicles also need to be survivable against cyber threats, and although some of the attack vectors and examples might not apply to all combat platforms, most of the survivability assessment and survivability enhancement concepts do. Cyber is an emerging threat, not just to combat aircraft, but to all military cyber systems.

This area of study is relatively new. While this thesis attempts to establish some of the ACCS fundamentals, further classified research should be done to assess the susceptibility and vulnerability of specific combat platforms and/or hardware. The ACS design discipline should also expand further to address how to defend against directed energy weapons, another emerging threat to combat aircraft. Lastly, industry-wide cyber-survivability test and evaluation requirements should be proposed and incorporated to reduce the likelihood of aircraft-related cyber incidents.

APPENDIX A. SME PROBABILITY DATA

Table 4. Sample SME data for active phase

Probability that Cyber Weapon is Active, P_A	"Experts"	Minimum Probability	Most Likely Probability	Maximum Probability
	A	0.8	0.9	1
	B	0.6	0.8	0.95
	C	0.85	0.95	1
	D	0.83	0.85	0.98
	E	0.74	0.89	0.99
	F	0.66	0.87	1
	G	0.85	0.92	1
	H	0.65	0.78	1
	I	0.8	0.9	0.94
	J	0.89	0.95	0.99
	K	0.9	0.98	1
	L	0.77	0.86	0.9
	M	0.85	0.92	0.95
	N	0.9	0.95	1
	O	0.95	0.99	1
	P	0.85	0.94	0.98
	Q	0.93	0.97	1
	R	0.95	0.98	1
	S	0.85	0.94	1
	T	0.72	0.83	0.95
	mean	0.817	0.9085	0.9815

Table 5. Sample SME data for detect phase

Probability of Detection given Active, $P_{D A}$	"Experts"	Minimum Probability	Most Likely Probability	Maximum Probability
	A	0.39	0.48	0.78
	B	0.15	0.45	0.71
	C	0.37	0.5	0.81
	D	0.26	0.44	0.6
	E	0.23	0.39	0.64
	F	0.12	0.21	0.46
	G	0.087	0.23	0.49
	H	0.058	0.27	0.71
	I	0.21	0.23	0.62
	J	0.25	0.31	0.66
	K	0.16	0.3	0.67
	L	0.39	0.49	0.64
	M	0.2	0.49	0.74
	N	0.16	0.32	0.64
	O	0.23	0.38	0.59
	P	0.19	0.24	0.67
	Q	0.11	0.38	0.57
	R	0.16	0.39	0.57
	S	0.33	0.47	0.67
	T	0.16	0.42	0.53
	mean	0.21075	0.3695	0.6385

Table 6. Sample SME data for launch phase

Probability of Launch given Detection, $P_{L D}$	"Experts"	Minimum Probability	Most Likely Probability	Maximum Probability
	A	0.4	0.73	0.94
	B	0.5	0.67	0.8
	C	0.75	0.92	1
	D	0.4	0.68	0.95
	E	0.2	0.48	0.8
	F	0.25	0.53	0.77
	G	0.55	0.61	0.75
	H	0.24	0.5	0.83
	I	0.62	0.71	0.97
	J	0.58	0.68	0.92
	K	0.7	0.93	1
	L	0.76	0.84	0.96
	M	0.1	0.31	0.63
	N	0.29	0.45	0.78
	O	0.27	0.53	0.87
	P	0.65	0.86	1
	Q	0.5	0.75	1
	R	0.7	0.88	1
	S	0.66	0.9	1
	T	0.45	0.69	0.95
	mean	0.4785	0.6825	0.896

Table 7. Sample SME data for implant phase

Probability of Implantation given Launch, $P_{I L}$	"Experts"	Minimum Probability	Most Likely Probability	Maximum Probability
	A	0.06	0.37	0.76
	B	0.3	0.47	0.82
	C	0.01	0.26	0.67
	D	0.11	0.37	0.57
	E	0.32	0.52	0.9
	F	0.17	0.4	0.85
	G	0.35	0.58	0.92
	H	0.27	0.41	0.68
	I	0.2	0.45	0.8
	J	0.11	0.34	0.72
	K	0.24	0.45	0.63
	L	0.05	0.44	0.9
	M	0.02	0.22	0.43
	N	0.29	0.51	0.89
	O	0.07	0.38	0.77
	P	0.2	0.32	0.45
	Q	0.04	0.19	0.65
	R	0.32	0.44	0.59
	S	0	0.15	0.45
	T	0.15	0.43	0.71
	mean	0.164	0.385	0.708

Table 8. Sample SME data for hit phase

Probability of Hit given Implantation, $P_{H I}$	"Experts"	Minimum Probability	Most Likely Probability	Maximum Probability
	A	0.1	0.43	0.7
	B	0.35	0.47	0.76
	C	0.21	0.38	0.55
	D	0.18	0.46	0.85
	E	0.3	0.44	0.55
	F	0.24	0.36	0.61
	G	0.3	0.41	0.75
	H	0.22	0.45	0.64
	I	0.05	0.19	0.55
	J	0.16	0.36	0.5
	K	0.05	0.5	0.83
	L	0.28	0.45	0.58
	M	0.08	0.22	0.36
	N	0.05	0.4	0.61
	O	0.14	0.42	0.7
	P	0.28	0.44	0.55
	Q	0.31	0.47	0.65
	R	0.09	0.38	0.56
	S	0.2	0.39	0.76
	T	0.14	0.28	0.44
	mean	0.1865	0.395	0.625

Table 9. Sample SME data for kill phase

Probability of Kill given Hit, $P_{K H}$	"Experts"	Minimum Probability	Most Likely Probability	Maximum Probability
	A	0.57	0.67	0.75
	B	0.36	0.54	0.7
	C	0.54	0.67	0.92
	D	0.56	0.72	0.93
	E	0.88	0.94	1
	F	0.63	0.74	0.94
	G	0.57	0.77	0.83
	H	0.63	0.75	0.86
	I	0.55	0.6	0.68
	J	0.72	0.81	0.92
	K	0.45	0.57	0.77
	L	0.5	0.65	0.8
	M	0.4	0.59	0.73
	N	0.68	0.76	0.86
	O	0.7	0.85	1
	P	0.61	0.72	0.88
	Q	0.77	0.85	1
	R	0.65	0.8	0.95
	S	0.69	0.8	1
	T	0.59	0.7	0.8
	mean	0.6025	0.725	0.866

APPENDIX B. MONTE CARLO SIMULATION MATLAB SCRIPT

```
%% Monte Carlo Simulation of an Anti-Aircraft Cyber-Attack
% Based on ACCS Probabilistic Kill Chain
% Run with and without Survivability Enhancement Features

% Austin Weinman
% ENS, United States Navy
% Naval Postgraduate School

clear all

%% Read Data into MATLAB

% Probability of Active
PA_min = xlsread('ACCS_Data.xlsx', 1, 'C2:C21');
PA_mode = xlsread('ACCS_Data.xlsx', 1, 'D2:D21');
PA_max = xlsread('ACCS_Data.xlsx', 1, 'E2:E21');

% Probability of Detection given Active
PD_min = xlsread('ACCS_Data.xlsx', 1, 'C24:C43');
PD_mode = xlsread('ACCS_Data.xlsx', 1, 'D24:D43');
PD_max = xlsread('ACCS_Data.xlsx', 1, 'E24:E43');

% Probability of Launch given Detection
PL_min = xlsread('ACCS_Data.xlsx', 1, 'C46:C65');
PL_mode = xlsread('ACCS_Data.xlsx', 1, 'D46:D65');
PL_max = xlsread('ACCS_Data.xlsx', 1, 'E46:E65');

% Probability of Implantation given Launch
PI_min = xlsread('ACCS_Data.xlsx', 1, 'C68:C87');
PI_mode = xlsread('ACCS_Data.xlsx', 1, 'D68:D87');
PI_max = xlsread('ACCS_Data.xlsx', 1, 'E68:E87');

% Probability of Hit given Implantation
PH_min = xlsread('ACCS_Data.xlsx', 1, 'C90:C109');
PH_mode = xlsread('ACCS_Data.xlsx', 1, 'D90:D109');
PH_max = xlsread('ACCS_Data.xlsx', 1, 'E90:E109');

% Probability of Kill given Hit (Vulnerability)
PK_min = xlsread('ACCS_Data.xlsx', 1, 'C112:C131');
PK_mode = xlsread('ACCS_Data.xlsx', 1, 'D112:D131');
PK_max = xlsread('ACCS_Data.xlsx', 1, 'E112:E131');

%% Triangular Distributions

x = 0:.001:1;

% Probability of Active
PA_lower = mean(PA_min);
PA_peak = mean(PA_mode);
```

```

PA_upper = mean(PA_max);
PA_pd = makedist('triangular','a',PA_lower,'b',PA_peak,'c',PA_upper)
y_PA = pdf(PA_pd,x);
pa = (PA_lower+PA_peak+PA_upper)/3;

% Probability of Detection given Active
PD_lower = mean(PD_min);
PD_peak = mean(PD_mode);
PD_upper = mean(PD_max);
PD_pd = makedist('triangular','a',PD_lower,'b',PD_peak,'c',PD_upper)
y_PD = pdf(PD_pd,x);
pdga = (PD_lower+PD_peak+PD_upper)/3;

% Probability of Launch given Detection
PL_lower = mean(PL_min);
PL_peak = mean(PL_mode);
PL_upper = mean(PL_max);
PL_pd = makedist('triangular','a',PL_lower,'b',PL_peak,'c',PL_upper)
y_PL = pdf(PL_pd,x);
plgd = (PL_lower+PL_peak+PL_upper)/3;

% Probability of Implantation given Launch
PI_lower = mean(PI_min);
PI_peak = mean(PI_mode);
PI_upper = mean(PI_max);
PI_pd = makedist('triangular','a',PI_lower,'b',PI_peak,'c',PI_upper)
y_PI = pdf(PI_pd,x);
pigi = (PI_lower+PI_peak+PI_upper)/3;

% Probability of Hit given Implantation
PH_lower = mean(PH_min);
PH_peak = mean(PH_mode);
PH_upper = mean(PH_max);
PH_pd = makedist('triangular','a',PH_lower,'b',PH_peak,'c',PH_upper)
y_PH = pdf(PH_pd,x);
phgi = (PH_lower+PH_peak+PH_upper)/3;

% Probability of Kill given Hit (Vulnerability)
PK_lower = mean(PK_min);
PK_peak = mean(PK_mode);
PK_upper = mean(PK_max);
PK_pd = makedist('triangular','a',PK_lower,'b',PK_peak,'c',PK_upper)
y_PK = pdf(PK_pd,x);
pkggh = (PK_lower+PK_peak+PK_upper)/3;

%% Implementinng ACCS Survivability Enhancement Features

% [all SEF's expected to reduce conditional probability by 15-20%]
SEF = makedist('Uniform','lower',15,'upper',20);

% Reducing Susceptibility
    % Reducing Probability of Active
        SEF_1 = random(SEF,1000000,1)/100; % Cyber Threat Suppression
(Threat Suppression)

```

```

    % Reducing Probability of Detection given Active
    SEF_2 = random(SEF,1000000,1)/100; % Air-Gapping (Signature
Reduction/Management)
    % Reducing Probability of Launch given Detection
    SEF_3 = random(SEF,1000000,1)/100; % Threats of Retaliation
(Tactics, etc..)
    % Reducing Probability of Implantation given Launch
    SEF_4 = random(SEF,1000000,1)/100; % Multi-factor
Authentication (Cybersecurity Hardening)
    % Reducing Probability of Hit given Implantation
    SEF_5 = random(SEF,1000000,1)/100; % Active Monitoring
(Situational Awareness)

% Reducing Vulnerbality
    % Reducing Probability of Kill given Hit
    SEF_6 = random(SEF,1000000,1)/100; % Virtualization (System
Redundancy with Diversity)

%% Monte Carlo simulation
n=1;
trials=1000000;

% initialize all possible events
active=0;
notactive=0;
detect=0;
notdetect=0;
launch=0;
notlaunch=0;
implant=0;
notimplant=0;
hit=0;
nothit=0;
kill=0;
notkill=0;

while n <= trials;
% Phase I (Active)
rA = random(PA_pd);
RA = rand(1,1);
rD = random(PD_pd);
rL = random(PL_pd);
rI = random(PI_pd);
rH = random(PH_pd);
rK = random(PK_pd);

if RA < rA
    active = active+1;
    % Phase II (Detect)
    RD = rand(1,1);
    if RD < rD
        detect = detect+1;
        % Phase III (Launch)

```

```

    RL = rand(1,1);
    if RL < rL
        launch = launch+1;
        % Phase IV (Implant)
        RI = rand(1,1);
        if RI < rI
            implant = implant+1;
            % Phase V (Hit)
            RH = rand(1,1);
            if RH < rH
                hit = hit+1;
                % Phase VI (Kill)
                RK = rand(1,1);
                if RK < rK
                    kill = kill+1;
                else
                    notkill = notkill+1;
                end
            else
                nothit = nothit+1;
            end
        else
            notimplant = notimplant+1;
        end
    else
        notlaunch = notlaunch+1;
    end
else
    notdetect = notdetect+1;
end
notactive = notactive+1;
end

rA_values(n,1) = rA;
rD_values(n,1) = rD;
rL_values(n,1) = rL;
rI_values(n,1) = rI;
rH_values(n,1) = rH;
rK_values(n,1) = rK;

n = n+1;
end

%% Monte Carlo simulation with SEF
n2=1;
trials2=1000000;

% initialize all possible events
active2=0;
notactive2=0;
detect2=0;
notdetect2=0;
launch2=0;
notlaunch2=0;

```

```

implant2=0;
notimplant2=0;
hit2=0;
nothit2=0;
kill2=0;
notkill2=0;

while n2 <= trials2
% Phase I (Active)
rA2 = random(PA_pd)-SEF_1(n2);
RA2 = rand(1,1);
rD2 = random(PD_pd)-SEF_1(n2);
rL2 = random(PL_pd)-SEF_1(n2);
rI2 = random(PI_pd)-SEF_1(n2);
rH2 = random(PH_pd)-SEF_1(n2);
rK2 = random(PK_pd)-SEF_1(n2);

if RA2 < rA2
active2 = active2+1;
    % Phase II (Detect)
    RD2 = rand(1,1);
    if RD2 < rD2
detect2 = detect2+1;
        % Phase III (Launch)
        RL2 = rand(1,1);
        if RL2 < rL2
launch2 = launch2+1;
            % Phase IV (Implant)
            RI2 = rand(1,1);
            if RI2 < rI2
implant2 = implant2+1;
                % Phase V (Hit)
                RH2 = rand(1,1);
                if RH2 < rH2
hit2 = hit2+1;
                    % Phase VI (Kill)
                    RK2 = rand(1,1);
                    if RK2 < rK2
kill2 = kill2+1;
                        else
notkill2 = notkill2+1;
                        end
                    else
nothit2 = nothit2+1;
                    end
                else
notimplant2 = notimplant2+1;
                end
            else
notlaunch2 = notlaunch2+1;
            end
        else
notdetect2 = notdetect2+1;
        end
    else

```

```

notactive2 = notactive2+1;
end

rA2_values(n2,1) = rA2;
rD2_values(n2,1) = rD2;
rL2_values(n2,1) = rL2;
rI2_values(n2,1) = rI2;
rH2_values(n2,1) = rH2;
rK2_values(n2,1) = rK2;

n2 = n2+1;
end

%% Calculating Probabilities
% Probabilities of Success for each Phase (based on MC)
PA = rA_values;
PD = PA.*rD_values;
PL = PD.*rL_values;
PI = PL.*rI_values;
PH = PI.*rH_values;
PK = PH.*rK_values;

% Probabilities of Success for each Phase (based on MC)
PA2 = rA2_values;
PD2 = PA2.*rD2_values;
PL2 = PD2.*rL2_values;
PI2 = PL2.*rI2_values;
PH2 = PI2.*rH2_values;
PK2 = PH2.*rK2_values;

% Theoretical Probabilities w/o SEF
prob_active_t = pa;
prob_detect_t = pa*pdga;
prob_launch_t = prob_detect_t*plgd;
prob_implant_t = prob_launch_t*pi gl;
prob_hit_t = prob_implant_t*phgi;
prob_kill_t = prob_hit_t*pkgh;

prob_survival_t = 1-prob_kill_t;

% Calculating Conditional Probabilities (based on MC)
PA = PA;
PDgA = rD_values;
PLgD = rL_values;
PIgL = rI_values;
PHgI = rH_values;
PKgH = rK_values;

% Calculating Conditional Probabilities (based on MC)
PA2 = PA2;
PDgA2 = rD2_values;
PLgD2 = rL2_values;
PIgL2 = rI2_values;
PHgI2 = rH2_values;

```

```

    PKgH2 = rK2_values;

% Probability of Kill and Probability of Survival (based on MC)
PK = PK;
PS = 1 - PK;

    % Probability of Kill and Probability of Survival (based on MC)
    PK2 = PK2;
    PS2 = 1 - PK2;

%% Plotting Results

% Conditional Probabilities
figure (2)
subplot (3,2,1)
hold on
title('Probability of Active')
yyaxis left
ylabel('No. of Trials')
histogram(PA,20)
yyaxis right
ylabel('Expected Distribution')
plot(x,y_PA)
xlabel('Probability')
legend ('Monte Carlo Results','Expected Triangular PDF')

subplot (3,2,2)
hold on
title('Probability of Detection given Active')
yyaxis left
ylabel('No. of Trials')
histogram(PDgA,20)
yyaxis right
ylabel('Expected Distribution')
plot(x,y_PD)
xlabel('Probability')
legend ('Monte Carlo Results','Expected Triangular PDF')

subplot (3,2,3)
hold on
title('Probability of Launch given Detection')
yyaxis left
ylabel('No. of Trials')
histogram(PLgD,20)
yyaxis right
ylabel('Expected Distribution')
plot(x,y_PL)
xlabel('Probability')
legend ('Monte Carlo Results','Expected Triangular PDF')

subplot (3,2,4)
hold on
title('Probability of Implantation given Launch')

```

```

yyaxis left
ylabel('No. of Trials')
histogram(PiGL,20)
yyaxis right
ylabel('Expected Distribution')
plot(x,y_PI)
xlabel('Probability')
legend ('Monte Carlo Results','Expected Triangular PDF')

subplot (3,2,5)
hold on
title('Probability of Hit given Implantation')
yyaxis left
ylabel('No. of Trials')
histogram(PHGI,20)
yyaxis right
ylabel('Expected Distribution')
plot(x,y_PH)
xlabel('Probability')
legend ('Monte Carlo Results','Expected Triangular PDF')

subplot (3,2,6)
hold on
title('Probability of Kill given Hit')
yyaxis left
ylabel('No. of Trials')
histogram(PKGI,20)
yyaxis right
ylabel('Expected Distribution')
plot(x,y_PK)
xlabel('Probability')
legend ('Monte Carlo Results','Expected Triangular PDF')

% Flow of Events
Events = [active, detect, launch, implant, hit, kill];
Events2 = [active2, detect2, launch2, implant2, hit2, kill2];
% EventNames = {'Active'; 'Detected'; 'Launched'; 'Implanted'; 'Hit';
'Killed'};
figure(3)
hold on
bar(Events)
bar(Events2)
ylabel('Successful Events out of 10,000 Trials')
xticks([1 2 3 4 5 6])
xticklabels({'Active'; 'Detected'; 'Launched'; 'Implanted'; 'Hit';
'Killed'})
legend('Without SEFs','With SEF')

% Probability of Survival/Kill
figure(4)
hold on
histogram(PK)
histogram(PK2)
title('Probability of Kill')
legend('Without SEFs','With SEF')

```

```

figure(5)
hold on
histogram(PS)
histogram(PS2)
title('Probability of Survival')
xlabel('Probability')
legend('Without SEFs', 'With SEF')

Killed = kill
    Killed2 = kill2
Survived = trials - kill
    Survived2 = trials2 - kill2
MC_Prob_Kill = Killed/trials
    MC_Prob_Kill2 = Killed2/trials2
MC_Prob_Survival = Survived/trials
    MC_Prob_Survival2 = Survived2/trials2

Theoretical_Prob_Kill = prob_kill_t
Theoretical_Prob_Survival = prob_survival_t

% Descriptive Probability of Survival Statistics
PS_percentiles = prctile(PS,[1 5 10 25 50 75 90 95 99]);
PS_mean = geomean(PS);
PS_std = nanstd(PS);
PS_skewness = skewness(PS);
PS_range = range(PS);
PS_CV = PS_std/PS_mean;

    % Descriptive Probability of Survival Statistics
    PS2_percentiles = prctile(PS2,[1 5 10 25 50 75 90 95 99]);
    PS2_mean = geomean(PS2);
    PS2_std = nanstd(PS2);
    PS2_skewness = skewness(PS2);
    PS2_range = range(PS2);
    PS2_CV = PS2_std/PS2_mean;

%% Where should we focus our attention?

pd_0 = fitdist(PS, 'Kernel', 'Kernel', 'epanechnikov')
pd_1 = fitdist(PS1, 'Kernel', 'Kernel', 'epanechnikov')
pd_2 = fitdist(PS2, 'Kernel', 'Kernel', 'epanechnikov')
pd_3 = fitdist(PS3, 'Kernel', 'Kernel', 'epanechnikov')
pd_4 = fitdist(PS4, 'Kernel', 'Kernel', 'epanechnikov')
pd_5 = fitdist(PS5, 'Kernel', 'Kernel', 'epanechnikov')
pd_6 = fitdist(PS6, 'Kernel', 'Kernel', 'epanechnikov')
%pd_all = fitdist(PSAll, 'Kernel', 'Kernel', 'epanechnikov')

figure(6)
hold on
histogram(PS)
histogram(PS1)
histogram(PS2)
histogram(PS3)
histogram(PS4)

```

```

histogram(PS5)
histogram(PS6)
xlim([0.93 1])
%histogram(PSAll)
legend('No SEFs', 'Focus on Active Phase', 'Focus on Detect Phase', 'Focus
on Launch Phase', 'Focus on Implantation Phase', 'Focus on Hit
Phase', 'Focus on Kill Phase')

y0 = pdf(pd_0,x);
y1 = pdf(pd_1,x);
y2 = pdf(pd_2,x);
y3 = pdf(pd_3,x);
y4 = pdf(pd_4,x);
y5 = pdf(pd_5,x);
y6 = pdf(pd_6,x);
%yAll = pdf(pd_all,x);

figure(7)
hold on
plot(x,y0)
plot(x,y1)
plot(x,y2)
plot(x,y3)
plot(x,y4)
plot(x,y5)
plot(x,y6)
%plot(x,yAll)
xlabel('Probability of Survival')
xlim([0.93 1])
legend('No SEFs', 'Focus on Active Phase', 'Focus on Detect Phase', 'Focus
on Launch Phase', 'Focus on Implantation Phase', 'Focus on Hit
Phase', 'Focus on Kill Phase')

%% section

figure(99)
plot(x,y_PL)
xlabel('Probability')

```

LIST OF REFERENCES

- Adams, C. (2019a). *Introduction to the ACS Discipline* [Class notes for ME4751: Aircraft Combat Survivability]. Department of Mechanical and Aerospace Engineering, Naval Postgraduate School.
- Adams, C. (2019b). *Threats and Threat Effects* [Class notes for ME4751: Aircraft Combat Survivability]. Department of Mechanical and Aerospace Engineering, Naval Postgraduate School.
- Alexandrovich, T. (n.d.). *Civil aviation cyber threats: Security vulnerabilities in next general air transportation systems* [Presentation]. Cyber Israel. Tel Aviv.
- Anderson, D., Checkoway, S., Cxeskis, A., Kantor, B., Kohno, T., McCoy, D., Roesner, F., Savage, S., & Shacham, H. (2011). Comprehensive experimental analyses of automotive attack surfaces. *USENIX Security Conference*.
- Ashworth, G. A. (2020). *Cyber survivability way ahead*. Office of the Under Secretary of Defense for Research and Engineering. v 2.0
- Ball, R. E. (2003). *the fundamentals of aircraft combat survivability analysis and design (2nd ed.)*. American Institute of Aeronautics and Astronautics.
- Ball, R. E. & Bryant, W. D. (2020a). Developing the fundamentals of aircraft cyber combat survivability (ACCS): Part 1- What is a cyber anti-aircraft weapon, and how can it kill an aircraft? Joint Aircraft Survivability Program, 16-23.
- Ball, R. E. & Bryant, W. D. (2020b). *Developing the fundamentals of aircraft cyber combat survivability (ACCS): Part 2- A cyber attack on an aircraft in a man-made hostile environment*. v.3.2 [Manuscript submitted for publication].
- Ball, R. E. & Bryant, W. D. (2020c). *Developing the fundamentals of aircraft cyber combat survivability (ACCS): Part 3- Design concepts for enhancing cyber survivability*. v.2.2 [Unpublished manuscript].
- Book, S. A., Covert, R. P., & Garvey, P. R. (2016). *Probability methods for cost uncertainty analysis: A systems engineering perspective* (2nd ed.). CRC Press.
- Bryant, B. (2016) Mission assurance through integrated cyber defense. *Air & Space Power Journal*.
- Bryant, B. (2018a) *First steps to cyber defense of aircraft: The big three* [Presentation].
- Bryant, W. D. (2018b). Surfing the chaos: Warfighting in a contested cyberspace environment. *Joint Force Quarterly*, 88, 28-33.

- Bryant, W. D. (2019). *Aircraft cyber combat survivability white paper*, v. 1.2
- Bryant, B. & Young, W. (2019). *Weapon systems cybersecurity tutorial* [Presentation]. Aircraft Survivability Symposium, Monterey, CA, United States.
- Buchanan, B. (2016). *The cybersecurity dilemma: Hacking, trust, and fear between nations*. Oxford University Press.
- Cyber Security Forum Initiative. (2010). *Preliminary STUXNET report*, v1.0.
- Department of Defense. (2018a). *Cyber strategy: Summary*.
https://media.defense.gov/2018/Sep/18/2002041658/-1/-1/1/CYBER_STRATEGY_SUMMARY_FINAL.PDF
- Department of Defense. (2018b, August 31). *The Defense Acquisition System* (DOD Directive 5000.1).
<https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodd/500001p.pdf>
- Department of Defense. (2020, January 23). *Operations of the adaptive acquisition framework* (DODI 5000.02).
<https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500002p.pdf?ver=2019-10-25-134150-283>
- Department of the Navy (2020). *Cybersecurity*.
<https://www.doncio.navy.mil/tagresults.aspx?ID=28>
- Disis, J., & Pham, S. (2019, October 15). Boeing promises to change 737 Max software as Indonesia releases Lion Air crash report. *CNN*.
<https://www.cnn.com/2019/10/25/business/lion-air-crash-boeing-737-max/index.html>
- Duggan, P. M. (2015). Strategic development of special warfare in cyberspace. *Joint Forces Quarterly*, (79).
- Gaycken, S. (2011). *Cyberwar: das Internet als Kriegsschauplatz*. München: Open Source Press.
- Hesse, R. (2000). Triangle distribution: Mathematica link for Excel. *Decision Line*. 12-14.
- Hubbard, D. W. (2014). *How to measure anything: Finding the value of “intangibles” in Business* (3rd ed.). John Wiley and Sons.
- Jackson, G. (2012). *Predicting malicious behavior: Tools and techniques for ensuring global security*. John Wiley and Sons.

- Johnson, D. (2002). The triangular distribution as a proxy for the beta distribution in risk analysis. *Journal of the Royal Statistical Society*, 46(3), 387-398.
- Joint Chiefs of Staff. (2014). *Information operations* (JP 3-13).
https://www.jcs.mil/Portals/36/Documents/Doctrine/pubs/jp3_13.pdf
- Joint Chiefs of Staff. (2020). *DOD dictionary of military and associated terms* (JP 1).
<https://www.jcs.mil/Portals/36/Documents/Doctrine/pubs/dictionary.pdf>
- Kollars, N. (2018, September 6). *Beyond the cyber leviathan: White hats and U.S. cyber defense*. War on the Rocks. <https://warontherocks.com/2018/09/beyond-the-cyber-leviathan-white-hats-and-u-s-cyber-defense/>
- Korunka, K., Sanchez, F. C., & Weilun, L. (2019). Applying irregular warfare principles to cyber warfare. *Joint Forces Quarterly*, (84).
- LaMarche, M. (2018, December 21). Electronic attack: An overview of electronic warfare part 4. *Mercury Systems Blog*. <https://blog.mrcy.com/electronic-attack/>
- Libicki, M. C. (2007). Conquest in cyberspace: National security and information warfare. *Cambridge University Press*, 71.
- Libicki, M. C. (2009). APPENDIX A: What constitutes an act of war in cyberspace? In *Cyberdeterrence and cyberwar* (pp. 179-182). RAND Corporation.
- Liu, M. (n.d.). Optimal number of trials for Monte Carlo simulation. *ValuationResearch*.
https://www.valuationresearch.com/wp-content/uploads/kb/SpecialReport_MonteCarloSimulationTrials.pdf
- Lockheed Martin (n.d.). *Proactively detect persistent threats: The cyber kill chain*.
<https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-kill-chain.html#Resources>
- Maksel, R. (2010, June 14). *The World War II history of the wright military flyer*. Air & Space. <https://www.airspacemag.com/daily-planet/the-world-war-ii-history-of-the-wright-military-flyer-139803059/>
- Maness, R. (2020a). *Cyber Strategy: Deception, Deterrence, Compellence* [Class notes for DA3105: Cyber Conflict]. Department of Defense Analysis, Naval Postgraduate School.
- Maness, R. (2020b). *Nuts and Bolts* [Class notes for DA3105: Cyber Conflict]. Department of Defense Analysis, Naval Postgraduate School.

- Mehta, A. (2018, September 19). *Shanahan: cybersecurity will become new measure for industry*. Fifth Domain. <https://www.fifthdomain.com/digital-show-dailies/air-force-association/2018/09/19/shanahan-cyber-security-will-become-fourth-critical-measurement-for-industry/>
- Moore, D. A. (2018). Overconfidence: The mother of all biases. *Psychology Today*.
- National Business Aviation Association. (2016, July 4). *Cyber security: Top flight department threats*. <https://nbaa.org/aircraft-operations/security/cyber-security-top-flight-department-threats/>
- Nohe, P. (2019, August 27). *The difference between encryption, hashing and salting*. Hashedout. <https://www.thesslstore.com/blog/difference-encryption-hashing-salting/>
- Nye, J. S. (2011). Nuclear lessons for cyber security. *Strategic Studies Quarterly*.
- Nye, J. S. (2012). *The future of power*. New York: PublicAffairs.
- Olejnik, L. (2019, April 2). *Global consequences of escalating U.S.-Russia cyber conflict*. Council on Foreign Relations. <https://www.cfr.org/blog/global-consequences-escalating-us-russia-cyber-conflict>
- Pipeline & Gas Journal*. (2012). *Next generation cyber attacks target oil and gas SCADA*. 239(2). <https://pgjonline.com/magazine/2012/february-2012-vol-239-no-2/features/next-generation-cyber-attacks-target-oil-and-gas-scada>
- Ragan, S. (2010) *The cyber shockwave event and its aftermath*. The Tech Herald.
- Shanker, T. (2010, April 15). Cyber-War Nominee Sees Gaps in Law. *The New York Times*. <https://www.nytimes.com/2010/04/15/world/15military.html>
- Tatum, J. T. (2019). *Aircraft survivability against high power radio frequency/microwave (HPM) directed energy weapons (DEWs)* [Presentation]. Aircraft Survivability Symposium, Monterey, CA, United States.
- World Economic Forum. (2010). *Global risks report 2010*. http://www3.weforum.org/docs/WEF_Global_Risks_Report_2010.pdf
- World Economic Forum. (2019). *Global risks report 2019*. <https://www.weforum.org/reports/the-global-risks-report-2019>
- World Health Organization. (2012, August 24). *Natural events*. https://www.who.int/environmental_health_emergencies/natural_events/en/.
- World Health Organization. (2020). *Biological weapons*. https://www.who.int/health-topics/biological-weapons#tab=tab_1

Xing, Y. & Zhan Y. (2012). Virtualization and cloud computing.
https://doi.org/10.1007/978-3-642-27323-0_39

THIS PAGE INTENTIONALLY LEFT BLANK

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
Ft. Belvoir, Virginia
2. Dudley Knox Library
Naval Postgraduate School
Monterey, California