



NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

**OPTIMIZING VLS FIRING POLICY:
AN ENUMERATION OF HETEROGENEOUS
SEQUENCES TO INFORM EXPENDITURE**

by

Steel J. Templin

December 2021

Thesis Advisor:

Michael P. Atkinson

Co-Advisor:

Moshe Kress

Second Reader:

Javier Salmeron-Medrano

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 December 2021		3. REPORT TYPE AND DATES COVERED Master's thesis
4. TITLE AND SUBTITLE OPTIMIZING VLS FIRING POLICY: AN ENUMERATION OF HETEROGENEOUS SEQUENCES TO INFORM EXPENDITURE			5. FUNDING NUMBERS	
6. AUTHOR(S) Steel J. Templin				
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 U.S. Navy (USN) utilizes the Vertical Launch System (VLS) to store and launch both their offensive and defensive missiles. Since the number of VLS silos on a given ship is fixed, to maximize offensive capability, the USN needs to minimize the number of interceptors required to combat incoming anti-surface missiles. Current firing policies may be overly conservative and expend too many interceptors per incoming threat, which results in a substantial fraction of VLS silos dedicated to defensive missiles. Decision makers need an analysis tool to explore the trade-off between missile consumption and probability of raid annihilation (PRA) for various firing policies and would also benefit from a prescriptive algorithm to help inform missile expenditure. This thesis provides a model to optimize VLS firing policy using a set of multiple interceptor types while accounting for range limitations, travel time, multi-interceptor salvos, battle damage assessment, and range dependent probability of kill. Additionally, the thesis derives analytical results for the optimal, lowest-cost allocation of interceptors in the single interceptor case, which, in turn, generates insight into how to structure sequential salvos.				
14. SUBJECT TERMS VLS, Vertical Launch System, missile, interceptor, BDA, battle damage assessment, defense, loadout, interceptor, expenditure, mix, pk, salvo, range, counterfire, SLS, shoot look shoot, SLSS, SSLSS, ASCM, anti-ship cruise missile, BMD, DMO, distributed maritime operations, IAMD, integrated air and missile defense, PEK, PRA, probability of engagement kill, probability of raid annihilation, optimization, low cost, PIP, probable impact point, IAMD, ship, defense, hard kill			15. NUMBER OF PAGES 111	
			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.

**OPTIMIZING VLS FIRING POLICY: AN ENUMERATION OF
HETEROGENEOUS SEQUENCES TO INFORM EXPENDITURE**

Steel J. Templin
Ensign, United States Navy
BS, United States Naval Academy, 2020

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN OPERATIONS RESEARCH

from the

**NAVAL POSTGRADUATE SCHOOL
December 2021**

Approved by: Michael P. Atkinson
Advisor

Moshe Kress
Co-Advisor

Javier Salmeron-Medrano
Second Reader

W. Matthew Carlyle
Chair, Department of Operations Research

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

The U.S. Navy (USN) utilizes the Vertical Launch System (VLS) to store and launch both their offensive and defensive missiles. Since the number of VLS silos on a given ship is fixed, to maximize offensive capability, the USN needs to minimize the number of interceptors required to combat incoming anti-surface missiles. Current firing policies may be overly conservative and expend too many interceptors per incoming threat, which results in a substantial fraction of VLS silos dedicated to defensive missiles. Decision makers need an analysis tool to explore the trade-off between missile consumption and probability of raid annihilation (PRA) for various firing policies and would also benefit from a prescriptive algorithm to help inform missile expenditure. This thesis provides a model to optimize VLS firing policy using a set of multiple interceptor types while accounting for range limitations, travel time, multi-interceptor salvos, battle damage assessment, and range dependent probability of kill. Additionally, the thesis derives analytical results for the optimal, lowest-cost allocation of interceptors in the single interceptor case, which, in turn, generates insight into how to structure sequential salvos.

THIS PAGE INTENTIONALLY LEFT BLANK

Table of Contents

1	Introduction	1
1.1	Background	1
1.2	Motivation	4
1.3	Thesis Contents	6
2	Literature Review	7
2.1	Hughes' Salvo Equations	7
2.2	Armstrong's Stochastic Salvo Model.	9
2.3	Armstrong's Iron Dome Model.	10
2.4	Weapon Target Assignment Problem	11
2.5	Firing with Feedback.	12
3	Model	15
3.1	Terminology and Assumptions	16
3.2	Notation.	19
3.3	Single Salvo Probability of Kill.	20
3.4	Firing Sequence Probability of Kill	26
3.5	Firing Sequence Example	27
3.6	Best Firing Sequence.	30
4	Results for One Interceptor Type	41
4.1	Case $M \geq N$	42
4.2	Case $M = 2$	44
4.3	Case $2 < M < N$	49
4.4	Derived Results	53
4.5	Comparison of SLS Policies	59
5	Case Studies	65

5.1	Baseline	65
5.2	Dropped Interceptor	71
5.3	Dominant Interceptor.	73
5.4	Monotonicity in the Heterogeneous Case	75
5.5	Performance	75
6	Conclusion	79
6.1	Practicality.	79
6.2	Recommendations and Insights.	80
6.3	Future Work	80
	List of References	85
	Initial Distribution List	89

List of Figures

Figure 3.1	Efficient recursion example	38
Figure 4.1	Q vs $x^*/\tilde{N}$	47
Figure 4.2	Q vs x^* for $p \in \{0.3, 0.5, 0.7, 0.9\}$	48
Figure 4.3	p vs x^* for $Q \in \{0.8, 0.9, 0.96, 0.996\}$	48
Figure 4.4	i vs $\frac{x_{M-i}^*}{N}$ for salvos with $M \in \{2, 3, 4, 5\}$	51
Figure 4.5	Q vs $\frac{x_M}{N}$ for salvos of size $M \in \{2, 3, 4, 5\}$	52
Figure 4.6	Salvos of size M vs $\frac{x_M}{N}$ for $Q \in \{0.8, 0.9, 0.96, 0.996\}$	52
Figure 4.7	M vs p^* for various Q s in the integer case	57
Figure 4.8	Q vs p^* for $M \in \{2, 3, 4, 5\}$ in the integer case	57
Figure 4.9	$N - M$ vs p^* for various Q s in the integer case	58
Figure 4.10	E[<i>cost</i>] vs PRA of firing policies for $p = 0.3$	62
Figure 4.11	E[<i>cost</i>] vs PRA of firing policies for $p = 0.5$	62
Figure 4.12	E[<i>cost</i>] vs PRA of firing policies for $p = 0.8$	63
Figure 5.1	Range dependent SSPK for four different interceptors	66
Figure 5.2	Best, worst, and average E[<i>cost</i>] across the top ten cases for $t_{BDA} \in \{30, 55, 80\}$	67
Figure 5.3	Q vs min(E[<i>cost</i>]) and max(E[<i>cost</i>]) of top ten sequences	69
Figure 5.4	Range dependent probability of kill including the SM+	74
Figure 5.5	Sequences generated for various t_{BDA}	76
Figure 5.6	Model runtime in seconds for various t_{BDA}	76

THIS PAGE INTENTIONALLY LEFT BLANK

List of Tables

Table 3.1	Velocities v for each missile in $Z \cup K$	24
Table 3.2	$MaxPIP_{k,r}$ and $MinPIP_{k,r}$ with their respective $SSPK_{k,r}(x)$. . .	24
Table 3.3	Salvo index s , the interceptor types k in that salvo, and $n_{f,s,k}$ for $f = [\{A\} \text{ look } \{B,C\} \text{ look } \{B,D\}]$	28
Table 3.4	Combination with replacement for salvos of size two given the set of feasible interceptors K of $\{A, B, C, D\}$	32
Table 3.5	$ C_r $ given a sampling of $ K $ interceptor	33
Table 3.6	$ F $ given J and $ K $	35
Table 3.7	All firing sequences for $k = \{B,C,D\}$ and $D = 200$ km	37
Table 3.8	Salvo index s , the interceptor types k in that salvo, and $n_{f,s,k}$ for each salvo in f with $[\{A\} \text{ look } \{B,C\} \text{ look } \{D\}]$	39
Table 3.9	Value V_k for each interceptor in K	40
Table 4.1	Fraction of interceptors $x_i^*/\tilde{N}$ for $Q = 0.996$ and $p \in \{0.8, 0.5\}$. .	49
Table 4.2	PRA and $E[\text{expenditure}]$ given $p \in \{0.3, 0.5, 0.8\}$ for various two and three salvo firing policies	60
Table 5.1	Parameters of four interceptors that can be used against a C802 threat	66
Table 5.2	Best and tenth best sequences and their respective $E[\text{cost}]$, probability of raid annihilation (PRA), and looks for each t_{BDA} at $Q = 0.96$. .	68
Table 5.3	Best and tenth best sequences and their respective PRA for each t_{BDA} at $Q = 0.9996$	70
Table 5.4	Top sequence(s) for each scenario given $Q = 0.95$	71
Table 5.5	Top sequence for each scenario given $Q = 0.9996$	72
Table 5.6	Parameters of the SM+ against a red C802 threat	73

THIS PAGE INTENTIONALLY LEFT BLANK

List of Acronyms and Abbreviations

ASCM	Anti-ship cruise missile
BDA	Battle damage assessment
BMD	Ballistic missile defense
DCR	Detect/classify Range
DMO	Distributed maritime operations
IAMD	Integrated air and missile defense
IT	Interceptor-threat
LR	Launch range
PEK	Probability of engagement kill
PIP	Probable impact point
PRA	Probability of raid annihilation
SLS	Shoot-look-shoot
SLSS	Shoot-look-shoot-shoot
SSLSS	Shoot-shoot-look-shoot-shoot
SSPK	Single shot probability of kill
USN	U.S. Navy
VLS	Vertical Launch System
WTA	Weapon target assignment problem

THIS PAGE INTENTIONALLY LEFT BLANK

Executive Summary

Modern U.S. Navy (USN) surface ships and submarines are outfitted with a common launching “mount” known as the Vertical Launch System (VLS). VLS allows for missile compatibility regardless of the specific platform that carries them. Ships can be outfitted to fit mission needs and can utilize future missile that will be developed later in the ship’s life-cycle. VLS inventory is divided into separate mission sets such as anti-submarine, anti-surface, land-attack (commonly known as “strike”), integrated air and missile defense (IAMD), ballistic missile defense (BMD), etc.

Depletion of IAMD loadouts will result in a total degradation of combat effectiveness as a ship is rendered defenseless and — assuming it survives — must return to port to reload. This VLS replenishment is a multi-day voyage to and from a shore installation with additional time potentially spent waiting in port if the reloading infrastructure is preoccupied with other ships. During this duration, the vessel is not on station and therefore is not conducting any mission-set.

This thesis focuses on the employment of defensive (counterfire) interceptors which falls under the IAMD mission set. Traditionally, a USN ship’s captain decides how to deploy and employ missiles against incoming threats, and this firing policy is usually the same for all threat types. A firing sequence consists of a series of salvos fired at the incoming threat. Each salvo is a collection of interceptors which may include duplicates of a single interceptor type or a combination of different types of interceptors. After each salvo, the blue (defending) force performs battle damage assessment (BDA) to determine the outcome of the salvo. There is a cost to fire each interceptor; this cost may be financial or opportunity.

Current policies such as shoot-look-shoot do not consider situation specifics such as detect/classify Range (DCR), inventory, or threat. We focus on improving the engaging sequence by exploring policies regarding the interceptor composition of each salvo relevant to a specific situation. Our model is a heuristic based on simplifying assumptions that generates effective heterogeneous firing sequences against a given threat. It factors in constraints of minimum and maximum interceptor ranges, interceptor transition time between launch and impact, range-dependent probability of kill, BDA, and time to conduct BDA. From

these considerations, we build a recursive algorithm that generates sequences of salvos. We then seek salvos that satisfy a certain effectiveness measure — defined by the probability of successful interception — while minimizing expected cost.

Key assumptions are:

- Salvos are fired as early as possible without violating range constraints
- Target tracking is never lost and BDA is accurate
- Interceptor outcomes (and consequently salvo outcomes) are independent

First we analyze one salvo and this includes examining and determining the probable impact point (PIP) of each interceptor. PIP corresponds to a given launch range which is the distance of the interceptor from the firing platform at the time of launch. The conversion between PIP and launch range depends upon the velocities of the interceptor and threat. This PIP is the input to range-dependent single shot probability of kill (SSPK). Once we have the SSPK of each interceptor in the salvo, we combine them to compute the salvo's probability of engagement kill (PEK).

We next examine a firing sequence which is a sequence of salvos. The first salvo in the sequence proceeds as above. After the last interceptor in the salvo impacts the threat, blue conducts BDA which results in a delay. After BDA that returns a non-successful interception, the second salvo is analyzed in the same fashion as above. Blue continues this process with each salvo in the firing sequence, calculating each salvo's PEK as we go. Once all salvos are evaluated, the PEKs of each salvos are aggregated to compute the probability of raid annihilation (PRA) of the firing sequence.

To determine best firing sequence, we develop a recursive algorithm to generate all feasible firing sequences. This recursion provides heuristic approximations of the lowest cost sequences that surpass a certain PRA. We limit each salvo to three interceptors and remove sequences with PRA below a threshold Q . The remaining sequences are sorted by minimum expected cost and returned in order of lowest expected cost.

In addition to our model concerning heterogeneous mixes of interceptors, we explore a simplified case with homogeneous interceptors distributed across multiple salvos. This work does not include BDA delay or range-dependent SSPK. The key finding of this analysis

is that sequences should be structured such that they are monotonically non-decreasing with salvo index. If the optimal sequence — from an expected cost perspective — has two interceptors in a salvo, all subsequent salvos will have two or more interceptors.

Additional findings for the homogeneous case include: the operator wants to maximize the number of salvo firing opportunities in order to best minimize the expected cost. Second, the optimal fraction of interceptors allocated to each salvo does not depend on SSPK and instead depends on the number of salvo opportunities and PRA threshold Q . This implies SSPK only impacts the number of interceptors required, not the optimal structure of sequence. Finally, above a certain SSPK the optimal firing sequence has one interceptor in all salvos except the last one, which fires the remaining interceptors. This suggests that a sufficiently high SSPK simplifies allocation as only the final salvo needs adjustment.

Our work suggests that the key to decreasing the expected number of interceptors expended per threat is maximizing the number of salvo opportunities in a firing sequence. This can be done through early detection — improving DCR — and having the requisite interceptors to match this expanded capability. Firing opportunities can also be increased through improved BDA time and having high-velocity interceptors with improved minimum range. After maximizing salvo opportunities, the next important aspect is interceptor allocation. We prove that salvos should be monotonic, non-decreasing in size for the homogeneous case however our preliminary results show this broadly holds for the heterogeneous case as well. Finally, our work demonstrates that each interceptor type plays a role in generating optimal sequences and that model inputs of DCR and BDA delay have second-order effects on the importance of a given interceptor.

Ultimately, optimal, low-cost firing sequences are dependent on the threat type, interceptor availability, and DCR. Exploring the interaction between these factors as well as different interceptors is a complicated task, but our model provides a streamlined way to modify interceptors parameters and generate the subsequent near-optimal sequences. Examples of how our model can be used include preemptively deriving a firing policy for a given threat scenario, evaluation of new or upgraded interceptors, and feeding into a loadout planner to provide end-to-end optimization from logistics to firing.

THIS PAGE INTENTIONALLY LEFT BLANK

Acknowledgments

As much as I hate cliches, the proverb “It takes a village to raise a child” perfectly illustrates the amount of sustained effort it takes for a backwoods Florida boy to make it through to finishing his master's. For some individuals, it was but a brief push in the right direction at the right time, while others provided persistent effort to support and shape my academic career. I could never express my full gratitude, but I can attempt to give credit for those who helped shape my success.

The people who were there the entire time, my parents, both biological and step. I could not have made it this far without the four of you and your constant encouragement to strive towards the next step. I would be remiss if I did not thank Mom and Kent who — with great push back — forced me to apply myself throughout elementary and middle school until I eventually cultivated my own standalone academic drive.

From my pre-college education, three names stand out: Star Olson, Carolyn Bristo, and Anna Williams: one taught me to reach for the stars, one helped me focus my wildness into bettering myself and others, and one cultivated my love for math. For the two that had me in high school, I can only apologize for how difficult I was.

From NAWCTSD, CDR Robert Salvia, Jennifer Fowlkes, Kelly Neville, and Hal Isen. The quad who took a clueless intern and dared to give him some autonomy. Many blunders were made that summer, and falling asleep on the first day was just the start. Granted, eight hours of meeting with NAVAIR PAX was perhaps the most brutal introduction to bureaucracy one can experience. Yet I came out of the summer with countless lessons learned, a subliminal push toward operation research, and a heavy-handed push towards aviation (thank you, Diesel).

From my two times at Sandia; Brad Steinfeld, Jennifer Gilbride, Sita Mani, Angie Shia, Gary Ashcraft, Corey Barret, and Staci Dorsey. Through both my experiences, I had the opportunity to pursue projects that interested me, even if they did not perfectly align with my teams' existing research efforts.

From the academy, CAPT Jay Foraker, CAPT David Ruth, and William Travis, who provided top-notch education in the field I came to love and pushed me academically.

From Lockheed, Steve Notarnicola and David Prothero, who showed me insight into the sub-specialty of operations research that is wargaming.

To my advisors on this thesis, you occasionally had to crack the whip on me, but you helped guide me over the last year. The contribution made through this thesis will hopefully lend insights that will help my peers who are or will be deployed across the globe.

Finally, the group of friends I made out here at NPS. Be it study nights, backpacking trips to the most remote parts of California, or “sending it” down the slopes, I could not have made it through this without you. Thanks for somehow keeping me sane.

CHAPTER 1:

Introduction

The U.S. Navy (USN) utilizes the Vertical Launch System (VLS) on surface and sub-surface vessels to store and launch both their offensive and defensive missiles. To combat incoming anti-ship cruise missile (ASCM), VLS should contain a substantial amount of defensive missiles. However, the capacity of a given ship's VLS is fixed so an excess of defensive missiles limits the offensive capability of the ship. If the USN is more judicious with the firing its defensive missiles against incoming threats, the USN can put more offensive missiles in VLS, which allows for more missions. This thesis examines defensive firing policies to efficiently expend defensive missiles.

1.1 Background

In this thesis we denote the two sides of the conflict as red and blue. Red refers to an attacking force that launches missiles to attack blue assets (ships), whereas blue is the defending force attempting to destroy said threat by launching interceptors.

1.1.1 Vertical Launch System

Modern USN surface ships and submarines are outfitted with a common launching “mount” known as the VLS. The first of two defining features for this launcher is that it allows for missile compatibility regardless of the specific platform that carries them. Ships can be outfitted to fit mission needs and can utilize future missiles that will be developed later in the ship's life-cycle. The key to this modularity is that VLS is composed of 80-122 individual cells, each acting as their own allocation slot. Any missile can be loaded into any cell on a ship. Most cells only contain one missile, but a cell can contain four missiles for certain missile types.

Additionally, VLS is integral to the fire control process. Each cell serves as the storage and launch silo for the missile(s) it contains. If a ship has 100 VLS cells and all cells are loaded, one could fire all 100 missiles independently (assuming one per cell). This means the ships entire magazine is pre-loaded while in port and the ship does not reload before firing any

missile it has on board. According to Janes Database: Aegis (2021), VLS integrates with AEGIS radar and fire control systems to provide a self-contained, beginning-to-end process that selects what missiles to fire, decides what cells to launch from, choose the launch order, guides each missile to the target, and conducts battle damage assessment (BDA) upon impact.

VLS inventory is divided into separate missions such as anti-submarine, anti-surface, land-attack, integrated air and missile defense (IAMD), ballistic missile defense (BMD), etc. Once a particular complement is depleted, the ship is no longer combat-effective in that area and must return to port to restore capability. While offensively oriented mission can hypothetically be depleted while maintaining a degree of overall combat efficiency, depletion of defensive loadouts, particularly IAMD, will result in a total degradation of combat effectiveness.

1.1.2 Detect-to-Engage Sequence

This thesis is focused on the expenditure of defensive (counterfire) interceptors which falls under the IAMD missions. Payne (2010) elaborates how individual IAMD engagement is traditionally broken into three distinct phases: detecting, tracking, and engaging.

Detection starts with surveillance either through active means like radar or through passive ones like infrared sensors. Unknown contacts will be identified and information such as range, elevation, velocity, and heading can be deduced from sensor input. As more of this input is gathered, operators classify the threat as a specific type of aircraft or missile. Some explanations of this sequence have classification as its own standalone phase as opposed to a step in the detection phase.

Tracking is its distinct phase of the sequence but overlaps with both detection and engaging phases. Once classification is made, operators want to maintain the identity and position of the contact. For contacts identified as adversarial, interpretation of the intent is important for blue to properly enact rules of engagement. Once engagement begins, tracking continues in that blue attempts to extrapolate the threat's future positions.

Once hostility is deduced and RoE allows it, blue engages the threat. This engagement first involves selecting the weapon(s) to engage with, which is followed by firing of said

weapons. If the selected weapons involve active sensors, blue must maintain a track on the target and blue's own weapon. Upon impact, blue performs BDA to confirm if the threat was successfully killed. If not, these steps are repeated until the target is destroyed or it impacts the vessel. Although understanding the general IAMD process is relevant, the engaging portion of the detect-to-engage sequence is what our efforts are concerned.

1.1.3 Firing Policy

Traditionally, a USN ship's captain decides how to deploy missiles against incoming threats, and this firing policy is usually the same for all threat types. A firing policy consists of a sequence of *salvos* fired at the target. Each salvo is a collection of interceptors; the salvo may include different types of interceptors. After each salvo, blue performs BDA to determine the outcome of the salvo. The term "look" is often used as the action performing BDA.

Common firing policies include: shoot-look-shoot (SLS), shoot-look-shoot-shoot (SLSS), or shoot-shoot-look-shoot-shoot (SSLSS). Each of these examples consists of two salvos separated by a BDA look. The number of "S" on the left side of the "L" dictate how many missiles are fired on the first salvo, and the number of "S" to the right of "L" are the number of missiles fired in the last salvo. Note: "SLS" refers to a distinct firing policy, where "SLS framework" refers to this firing policy notation regarding how salvos are structured.

When moving from SLS to SLSS to SSLSS, the probability of raid annihilation (PRA) increases. However, the benefit of increasing the PRA comes with a potential increase in the number of interceptor expended from two (SLS) up to four (SSLSS). Decreasing the number of interceptors fired per threat, while maintaining a satisfactory PRA can lead to loadouts with fewer missiles allocated for defense.

The BDA look in these firing policies analyzes the outcome of previous salvo(s) to see if the target is destroyed. If BDA reports the target is killed, the next salvo is not fired. BDA tends to be close to perfect at evaluating whether or not a hard-kill measure such as a missile is successful. BDA is less reliable with evaluating soft-kill measures such as jamming, decoys, and chaff. This thesis is scoped to only study hard-kill measures.

1.2 Motivation

1.2.1 Logistics of Vertical Launch System

USN ships have different missiles for different missions and when the missiles for one mission type is depleted, the ship may continue its operation performing other missions. However, once all of its defensive IAMD missiles on the VLS have been fired, the ship is defenseless and must return to port to reload. This VLS replenishment is a multi-day voyage to and from a shore installation with additional time potentially spent waiting in port if the reloading infrastructure is preoccupied with other ships. During this duration, the vessel is not on station and therefore is not conducting any mission.

Stires (2017) reported that CNO John Richardson was working towards restoring the capability to reload VLS at sea; however, these require dedicated ships referred to as tenders. Such tenders do not exist yet and will still require off-station time to transit to-and-from the combat area. It is important to limit the number of these replenishment, and improving firing policy efficiency lengthens the time between reloads.

A contemporary example that illustrates how long the transition takes even after a limited engagement was the 2011 Syrian military intervention. The USS PORTER and ROSS together launched a total of 60 tomahawk cruise missiles. After this single salvo the PORTER had a six day transit back to Naval Station Rota in Southwestern Spain to replenish its depleted VLS complement. This scenario was against a static target with no air threats present, yet required replenishment (Stires 2017). Against a near-peer adversary, the consequences of this consumption are magnified as our forces will be subject to retaliatory missiles and will need to defend.

Logistics challenges are further magnified in hot conflicts such as potential South China Sea scenarios. Such a conflict would likely see the targeting of port infrastructure and hypothetical VLS tenders. Should such assets be denied, sea lines of communication would have to stretch across half the Pacific, degrading the USN's ability to conduct surface operations in the Western Pacific.

1.2.2 Distributed Maritime Operations

American adversaries have focused on anti-access, area denial strategies and subsequent weapons acquisitions to exploit weakness inherent to USN carrier strike groups and surface action groups. Rather than meet the USN surface forces at sea, adversarial nations can prevent USN surface forces from operating in certain areas through amassing long-range shore based anti-surface missiles. Recognizing the vulnerability, the USN began to explore a shift towards distributed allocation of surface forces known as distributed maritime operations (DMO). Although this employment strategy has different interpretation for different entities throughout the U.S. defense complex, Popa et al. (2018) provides the best summary. The broad idea is to have destroyers, littoral combat ships, frigates, and future unmanned systems deployed individually or in small groups.

The advantages of DMO have two primary points. First, survivability is increased. Adversaries must detect, track, and engage each individual USN ship as opposed to finding a formation of multiple ships and flooding missiles at the target. Should a ship be lost, it is only a single loss as opposed to losing a five (or more) ship formation. Although a formation of ships has the advantage of overlapping IAMD and firepower, ship captains defend their ability to select their targets independently of formation, leading to redundant engagements of some targets.

Second, detection and lethality are improved as sensors and payloads are distributed across a larger geographic area (Popa et al. 2018). In the detect-to-engage sequence, this distribution of sensors helps improve the detection, classification, and tracking portions. The trade-off is that although payloads are distributed, the VLS loadouts of an individual ship must be increasingly self-sufficient. Each ship must contain enough offensive firepower to conduct one or more missions such as ASW, ASuW, or Strike while also being able to defend itself and conduct IAMD. This trade-off between offensive and defensive loadouts is the driving question of this thesis, with focus being on minimizing the defensive loadout through more efficient IAMD.

1.3 Thesis Contents

1.3.1 Modeling Approach

This thesis focuses on improving the engaging sequence by exploring policies regarding the interceptor composition of each salvo. Improving earlier phases of detect-to-engage sequences focuses on sensors and computation power and is outside our scope of the thesis. We assume our model begins immediately after detection and tracking, and that blue maintains tracking the entire duration of the engagement.

Our model is a heuristic that generates effective heterogeneous firing sequences against a given threat. It factors in constraints of minimum and maximum interceptor ranges, interceptor transition time between launch and impact, range-dependent probability of kill, and time it takes to conduct BDA. From these considerations, we build a recursive algorithm that generates sequences of salvos and evaluates their effectiveness and expected cost in terms of expected interceptors expended.

1.3.2 Roadmap

In Chapter 2 we perform a literature review to describe and compare related work. In Chapter 3 we formulate our model. In Chapter 4 we derive propositions that provide tactical insight into how optimal sequences are structured. Chapter 5 provides case studies using the model defined in Chapter 3 that incorporate real interceptor parameters and provides examples of generated sequences. In Chapter 6 we present our conclusions, discuss the practicality of our work, and recommend future extensions to our work.

CHAPTER 2:

Literature Review

2.1 Hughes' Salvo Equations

It is common knowledge within naval expertise circle that naval warfare is more offensively oriented than its land counterpart. Vego (2010) writes, "A strategic defensive at sea should always includes elements of offensive. A weaker side at sea should not be passive; just the opposite— it should try to exploit all the opportunities and act offensively at the tactical level in order to create preconditions to eventually go on the strategic offensive." Modern advancement in ASCM from both air, land and sea have accelerated this historical trend. Hughes (1995) explored various attributes in his salvo equations in an attempt to model unit-on-unit and force-on-force engagements and apply it to historical battles. The result is a model that caters to finding out what attributes in naval warfare are important and providing moderate predictive capability that can be backed by historical data. The key attributes of Hughes' model are:

- Force Strength - Number of units in either force
- Fighting Power - Number of well-aimed missiles by a force
- Staying Power - Number of hits that can be absorbed by a force
- Counterfire - Number of incoming well-aimed missiles that could be intercepted

Although force numbers are an important part of the equations, they are effectively multipliers that modify the impact of other parameters. The offensive capability (fighting power) has increased with the advent of modern ASCM. The proportion of a ship's magazine that is "well-aimed" for a given range has increased, and ships carry more missiles. This increase in fighting power further emphasizes the offensive nature of naval warfare.

From a defensive perspective, staying power relates to how many hits a ship can sustain and remain in the fight. Although improving staying power would be important, the practicality of increasing staying power is unrealistic as a single impact will leave most ships combat ineffective since most key systems are non-redundant. The practical way to improve defense is through improving counterfire. Hughes (1995) main equation only includes hard-kill

measures to intercept incoming threats. However, that same paper introduces extended equations with soft-kill measures such as distraction and seduction. Soft-kill is becoming increasingly relevant but time constraints limited our model's scope to focusing on hard-kill counterfire capability.

The beauty of Hughes salvo equations are their simplicity; the equations are easily understandable with minimal background, transparent, and they can be utilized for design of experiment and statistical analysis (Lucas and McGunnigle 2003). The trade-off is that tactical level policy is not explored directly as simplifying assumptions are made, e.g., only one homogeneous salvo is launched at time. Our model is able to account for heterogeneous salvos and examine different types of interceptors across a series of counterfire salvos.

An aspect of naval warfare that the original salvo equations do not account for is the ability for a force to exploit *first fire advantage*. This means one side fires and the opposing force may not be able to respond before that first salvo impacts. These salvos are referred to as *unanswered salvos*. If staying power is low, then it might not matter how strong blue is offensively and defensively: a much weaker red that is able to deliver unanswered salvos can still defeat blue.

Hughes' refers to the impact of unanswered salvos as tactical instability and describes it as: "[when] the combat power of either force grows relative to the staying power of that force, the risk of unstable circumstances grows with it." (Hughes 1995). Mathematically, this is a result of the fighting power growing as staying power remains constant. Increasing fighting power combined with the idea of unanswered salvos generates increasing variability, which is not captured in the deterministic model.

The original model does not allow for unanswered salvos, so Hughes' created an embellished model that introduces sensing variables for both fire and counterfire salvos (Hughes 1995). Despite capturing the importance of sensors, this second model is still deterministic. Modeling tactical instability comes as a result of variability, for which we need a stochastic model.

2.2 Armstrong's Stochastic Salvo Model

Building off of Hughes' equations, Armstrong (2005) developed a stochastic variation of the salvo model. This model still emphasizes the ability for either side to fire well-aimed missiles at the adversary while trying to intercept incoming threats. However, it swaps the parameters of Hughes' model for random variables with specified probabilities. Simplicity is sacrificed, but the model is still understandable and allows for probabilistic effects to emerge. The end goal for both models is to predict number of surviving ships for either side, and Armstrong's model also informs the variation of those survivors. We also utilize stochasticity in our model as the outcome of a firing sequence is random.

An example of stochastic effects is that Hughes' model exhibits abrupt shifts in who wins the battle. Armstrong's model gradual change between probability of win as the capability and numbers of either force engages are adjusted. At the crossover point when an identical force of equal numbers engages, the deterministic model would result in mutual annihilation of either force. The stochastic model still allows for mutual defeat; however, it has a probabilistic chance of victory for either side due to variability. Armstrong also allows for a fourth outcome in finite time where neither side is able to destroy all opposing ships if the lethality of a side is less than the opposing force's counterfire capability. Our model presents insight into how a force might improve counterfire in context of their current defensive inventory.

Armstrong's Stochastic Salvo model presents some broad tactical insight in that superior forces wants to balance offensive and defensive capability and fight in conditions of low risk such as good weather and traditional tactics. In contrast, the inferior force wants to prioritize either offense or defense and fight in high risk conditions such as utilizing experimental tactics and technology (Armstrong 2005). This potentially relates back to the idea of tactical instability provided by Hughes (1995) but with the ability to discuss in context of variability. Superior forces will probabilistically win the engagement and want to establish a scenario with as little variation as possible. Meanwhile inferior forces will never be "favored" and must exploit variability to create situations that present opportunities for victory.

Various derivatives of Armstrong equations exist such as a variations with an area fire model (Armstrong 2013) and a sequential exchange of fire (Armstrong 2014b). For both Hughes and Armstrong, Kesler (2019) ran a data farming experiment utilizing nearly orthogonal and balanced designs. A finding included the confirmation of regions of instability as well

as the advantage of numerical superiority. All of the equations and models mentioned have been two-sided whereas our model is one sided.

2.3 Armstrong’s Iron Dome Model

Applying the Stochastic Salvo Model to an applied example, Armstrong (2014a) adjusted his model to explore expected number of missiles intercepted by Israel’s Iron Dome system and the subsequent effects of non-interceptions. This model focuses on the fire and counterfire aspect of missile engagements with no reduction of force effects (losing a ship or battery). Hostile salvos can contain one or more missiles, of which a certain percentage are “accurate.” The defending force counter fires with a salvo of one or more interceptors, and each of these hit the accurate missiles according to a specific probability.

The Iron Dome model focuses on long term expenditure rather than individual engagements. In the long term, there is a trade-off in the cost of firing interceptors versus the financial and casualty damage caused by allowing threats to impact. The paper asserts that the defending force wants to minimize the damage value on both non-hardened and hardened targets, and primarily does this through adjusting the number of interceptors fired (Armstrong 2014a). Other secondary actions include increasing the number of hardened targets — akin to increasing the staying power in the salvo models — and conducting preemptive strikes to reduce the number of red missiles. The model seeks to minimize the expected value of damage but also considers the cost of overkill expenditure in the counterfire salvos. Our model focuses on the expected cost of the interceptors expended in a single raid as opposed to the cumulative expected damage over multiple raids. We require that the firing sequences satisfy a minimum PRA to ensure the solution is operationally effective.

Armstrong’s Iron Dome extension introduces salvos where interceptors are fired sequentially with BDA in between launches (Armstrong 2014a). Our model also incorporates this; however, we provide the additional fidelity of pairing interceptors to a single threat. The Iron Dome model includes an assessment portion of the model where the defender attempts to classify whether the target will hit the desired target: if the assessment determines the threat will miss, the threat is ignored by the Iron Dome. No time element is present in the assessment like in our model, but it does lead to false-positives and false-negatives that affect outcomes. The main insight from this work is that a sufficient number of salvos with

sufficient volume can overwhelm the defender’s systems.

2.4 Weapon Target Assignment Problem

Dating back to Manne (1958), the weapon target assignment problem (WTA) is a class of problems that seeks to maximize value by optimally allocating weapons to targets. This problem has become more important as militaries have switched from firing large volleys of “dumb weapons” to firing strategies with a smaller number of “smart weapons” with greater precision, including ASCM and naval surface combat.

Kline et al. (2019) lays out much of the most recent work on this problem and includes optimization formulations as well as various algorithmic approaches that have been developed. Most of the research is oriented on the defensive perspective, and the problem breaks down into a static and dynamic WTA sub-categories, the latter of which incorporates the dimension of time. Most approaches attempt to be “fast enough to provide an engagement solution before the oncoming targets reached their goals” (Leboucher et al. 2013). Even in the simpler static WTA, exact solutions can be difficult to obtain as complexity increases as number of missiles, missile types, interceptors, interceptors types, and other factors are considered.

Our model has diverted away from being a WTA as we made the simplifying assumption to only target a single threat. However, for that given threat our model considers detected/classified range, combinations of interceptors, combination of sequential salvos, and number of sequential salvos given the entry range. From these consideration we provide a heuristic estimate of the best counterfire sequences and our results could potentially feed into a multi-salvo WTA that allows for heterogeneous salvos.

Summers et al. (2020) formulates two dynamic WTA algorithms designed to allocate expenditure from various bases for a certain number of threats. “[These] algorithms outperformed both baseline policies for 14 and 16 instances, respectively, at varying and often high (i.e., $\geq 90\%$) confidence levels” Summers et al. (2020). However, they do not implement SLS or BDA which could improve firing policy further at the trade-off of additional computational complexity. A distinction of this model is that the algorithms are for the BMD case, which is a sub-specialty of missile defense. BMD along with hypersonic scenarios have nuances

due to the physical characteristics of the missiles flight path, making analysis and models not perfectly transferable to traditional missile defense.

2.5 Firing with Feedback

Most of the works in this section utilize the SLS framework introduced in Section 1.1.3 where salvos are separated by a look. A look is synonymous with BDA and provides feedback so that the operator evaluates the state of the target after the previous expenditure.

Perhaps the most similar to our work is Kalyanam and Clarkson (2021), who studied the structure of sequential salvos over a finite time horizon with distinct decision epochs (looks). Our result in Chapter 4 about monotonic, non-decreasing salvos was proven by Kalyanam and Clarkson (2021) through a Markov decision process as opposed to our nonlinear optimization approach.

A multi-threat version of this work is illustrated in Soland (1987) which explores a threat salvo with a random number of threats being intercepted by a finite number of interceptors. The defenders have a finite number of opportunities to engage each threat salvo, with probability of kill depending on which salvo interceptors are fired in. Their result is that distribution of interceptors between the threats should be uniform. Both Kalyanam and Clarkson (2021), and Soland (1987) utilize recursions to generate optimal salvo allocation; however, both are for a simplified case with only one interceptor type as opposed to our heterogeneous solution. Soland (1987) does introduce the idea of range dependent probability of kill.

Dutta (2014) also explores allocation of various engagement policies and our work in Section 4.5 mirrors some of the analysis. Dutta's model factors in range considerations and explores the maximum number of salvos that can be fired sequentially given an entry range. The paper also considers sensors effectiveness through radar range equations and probability of detecting threats. However, Dutta's model only considers homogeneous salvos.

Another SLS work is by Glazebrook and Washburn (2004) which looks for the optimal firing strategy using a Markov Decision process for various assumptions such as perfect versus imperfect information, infinite versus finite time horizon, and number of targets versus number of interceptors. A takeaway is the importance of BDA in scenarios where the operator has

imperfect information. Obtaining information about the state of the threat is critical for the selection of a proper firing policy, and this is also an important part of our model. The value of information obtained by the shooter is also emphasized in Washburn and Kress (2009) which provides formulations for various SLS extensions including heterogeneous salvo cases. Glazebrook and Washburn (2004) shows that even when homogeneous shots face a moderate number of threats in a constrained time window, the problem is difficult to solve. Our model allows for heterogeneity of shots but assumes perfect information for hard kill interceptors.

Glazebrook and Washburn (2004) also introduces the idea of framing the SLS problem as a multi-armed bandit (or just bandit) problem. Glazebrook et al. (2007) compliments this idea in that it focuses on index policies for determining the allocation of interceptors. This work explores extensions of the base problem including learning information over several engagements, partial and cumulative damage to the incoming threat, and the idea of the shooting platform disengaging from the threat (Glazebrook et al. 2007). The work proves that an index policy dominates other firing strategies for all extensions.

The basics of our model rely on substituting time for range by using velocities. This has been incorporated in numerous sources on missile defense, however it is worth noting that much of this literature focuses on BMD. One notable paper for our purpose is by Wang et al. (2017), as it introduces an optimization model that has range considerations and multiple interceptors. However the focus of this paper is on sensors exchanging data and the model allowing staggered salvos such that a salvo can be fired before the previous salvo impacts the target. The hypothetical advantages of such a tactic would be decrease expenditure; the tactic results in an increased number of salvos while decreasing the number of interceptors per salvo. This results in more frequent BDA assessments, and in cases where the interceptor can be steered post launch, interceptors are potentially utilized more effectively.

Despite potential merit, the tactic of staggered salvos is not a part of current USN tactics and is not included in our model. The prospects of this strategy is also modeled in Pryluk et al. (2013) which explores a shoot-shoot-look policy when not enough time remains after the BDA to fire a second sequence. The finding suggested that “the tactic increases the probability of the defended target surviving the salvo attack, compared to using the same number of defenders in a single shoot stage. This makes sense as PRA per-defender

can be increased while keeping the expected expenditure constant. Alternatively, using the proposed algorithm can reduce the amount of defenders being used” (Pryluk et al. 2013).

CHAPTER 3:

Model

In this chapter we develop a model to generate low cost firing sequences that satisfy a minimum effectiveness requirement regarding mission success. The official term for this kill probability is PRA. The model is based on a scenario where a blue asset defends against a single incoming red threat. The blue asset is simultaneously the detection platform, firing platform, and the target of the threat. The blue asset has a complement of interceptors that it can expend in an attempt to kill the red threat.

Blue engages with the threat by firing a sequence of salvos at the threat. A salvo is a collection of interceptors fired simultaneously. Each salvo is followed by a BDA where blue determines if the salvo killed the threat. We assume BDA is accurate and firing will cease if blue detects a kill. Blue can fire salvos at an alive threat as long as the threat is between the maximum and minimum range for all blue interceptors in a salvo. There is a cost to fire each interceptor; this cost may be financial or opportunity. The user specifies a minimum overall kill probability that the firing sequence must meet or exceed, for which the model heuristically generates a low cost firing sequence that satisfies the kill probability condition.

While we focus on USN surface vessels and their VLS complements, the model applies to other scenarios as long as the relative velocities of the asset and the threat are such that the target is static relative to the incoming threat. Other applications include a ground installation with missiles or an adversarial vessel trying to defend against a United States (U.S.) missile.

We first introduce terminology and assumptions in Section 3.1. Next, we define our notation in Section 3.2. We compute the kill probability of a single salvo in Section 3.3. Section 3.4 computes the probability of a single firing sequence. Section 3.5 illustrates how to work through evaluating a single firing sequence. Section 3.6 demonstrates how we generate all firing sequences, apply a cost function, and choose the best sequence.

3.1 Terminology and Assumptions

3.1.1 Definitions

Throughout the remainder of this document we shall use the following terminology:

- **Threat** - Red missile fired at the blue asset. The key characteristic is velocity. The defensive mission of the asset fails if the threat is not intercepted.
- **Interceptor** - Blue missile fired to intercept and eliminate the threat.
- **Type** - Designation of the threat or interceptors such as SM-6 or DF-21. Each type has an associated velocity, ranges, and probability of kill.
- **Interceptor-threat (IT) pair** - The combination of blue and red types involved in an interception attempt.
- **Single shot probability of kill (SSPK)** - Probability that a single interceptor will eliminate the threat. The SSPK depends on the IT pair since it is a function of range.
- **Effective range** - The interval for which the SSPK is positive for a given interceptor type. The effective range depends on the IT pair.
- **Min/Max range** - Minimum and maximum ranges of the effective range interval.
- **Launch range (LR)** - Distance between the firing platform and the threat at the time when an interceptor is launched.
- **Probable impact point (PIP)** - Point at which an interceptor is projected to impact a target. This value is measured as a distance from the firing platform. We compute PIP based on the LR and velocities of the interceptor and threat.
- **Salvo** - A collection of interceptors fired at a threat. Salvos may contain different types of interceptors and may have multiple interceptors of the same type. All interceptors in a salvo have the same LR, which means that the launches of interceptors in a salvo are assumed to be simultaneous.
- **Initial range** - For a specific salvo, the maximum range at which blue can fire a salvo. This is the upper bound on the LR of the salvo.
- **Detect/classify Range (DCR)** - Initial range for the first salvo of a firing sequence such that the target is detected, classified, and blue initiates target engagement. This is an input to the model. The initial range for later salvos are calculated based on the outcome of previous salvos.
- **Battle Damage Assessment (BDA or simply look)** - The process in which blue

determines whether or not a salvo was successful in killing the target.

- **BDA delay** - Amount of time to conduct BDA. This delay starts when the last interceptor in the salvo reaches its PIP and finishes after a deterministic amount of time. Blue cannot launch the next salvo until this delay concludes.
- **Feasible salvo** - A salvo for which there exists an LR such that the PIP of each interceptor in the salvo is within its corresponding effective range.
- **Valid Salvo** - A salvo where the initial range is greater than the minimal effective range for all interceptors.
- **Firing sequence** - A series of salvos each followed by a BDA and its subsequent delay.
- **Probability of engagement kill (PEK)** - Probability that a salvo will eliminate the threat. Defined in accordance with existing USN vocabulary.
- **Probability of raid annihilation (PRA)** - Probability that a firing sequence will eliminate the threat. Defined in accordance with existing USN vocabulary.

3.1.2 Assumptions

Our model also makes the following assumptions:

1. **The firing platform is static relative to the threat** - Although surface ships move, their velocity is negligible compared to that of the incoming threat, thus they are treated as stationary targets in trajectory calculations.
2. **Detection and classification are instantaneous and the target is never lost** - Missile defense is characterized by a three-part process of detecting, tracking, and engaging a target. We are interested in the engagement portion of this sequence.
3. **A salvo can contain up to three interceptors** - In practice, standard operation procedure will not see more than three interceptors fired at a given threat at a time. This limitation reduces the computational complexity for solving practical situations.
4. **All interceptors in a salvo are fired at the same time** - Standard operating procedure involves firing the interceptors of a salvo in immediate succession with minimal delay.

5. **We only consider feasible salvos in firing sequences** - By definition, an infeasible salvo contains at least one interceptor that produces a SSPK of 0. We could remove that interceptor and create a new firing sequence that generates the same overall PRA at a lower cost. We could have chosen a positive threshold (e.g., 0.1) to define feasibility to ensure all interceptors have a non-negligible SSPK. However, in practice firing sequences where interceptors have degraded SSPK are not usually chosen.
6. **A salvo is fired as early as possible** - This ensures blue can fire the maximum number of salvos, however this policy does not optimize individual salvos' PEK and might affect overarching PRA and minimum cost. Future work should examine trade-offs between optimal timing of salvos and number of salvos launched.
7. **Missiles have constant velocity** - Threats and interceptors both travel with constant horizontal velocity and do not have acceleration phases.
8. **An interception always results in a kill** - We treat interception as an event with a binary response regarding whether or not the threat has been killed. We do not consider the mechanism to achieve this kill, and ignore partial kills.
9. **Interceptors are independent of each other** - Regardless of what type and the number of interceptors fired in a salvo, we assume the outcome of each interceptor is not influenced by other interceptors fired.
10. **Salvos outcomes are independent of each other** - This is a consequence of interceptor independence.
11. **BDA is accurate** - We do not consider potential false positive or negative assessments. Also, BDA delay is the same deterministic time for all interceptors - The delay begins after the salvo's last interceptor reaches its PIP.

3.2 Notation

Indexes and Sets

- Red - Set of red threat types
- r - Threat index, denotes type of threat such that $r \in Red$
- K - Set of blue interceptor types
- k - Interceptor index, denotes type of interceptor such that $k \in K$
- S - Set of salvo indices. Contains integers from 1 to $|S|$
- s - Salvo index, denotes the serial number of a salvo in a firing sequence such that $s \in S$
- F - Set of firing sequences indices. Contains integers from 1 to $|F|$
- f - Firing sequence index, denotes serial number of the firing sequence such that $f \in F$
- $K_{f,s} \subset K$ - Subset of interceptors k included in salvo s of firing sequence f

Input Parameters

- $n_{f,s,k}$ - Number of interceptors of type k in salvo s of firing sequence f
- S_f - Number of salvos in firing sequence f
- D - DCR of the threat [km]
- v_k - Velocity of interceptor k [km/s]
- v_r - Velocity of threat r [km/s]
- t_{BDA} - Time to perform BDA [s]
- $SSPK_{k,r}(x)$ - Probability that an interceptor of type k kills a threat of type r at PIP x

S_f and $n_{f,s,k}$ are inputs in sections 3.3.1 and 3.3.1, when we focus on one salvo or one firing sequence. However S_f and $n_{f,s,k}$ are calculated parameters in later sections when we determine the best firing sequence via enumeration.

Calculated Parameters

- $R_{f,s}$ - First firing opportunity for salvo s in firing sequence f . $R_{f,1} = D$ for all f and $R_{f,s}$ depends upon the outcome of salvo $s - 1$ for $s \geq 2$ [km]
- $PIP_{k,r}(x)$ - Point at which an interceptor of type k is projected to intercept a threat of type r for a given LR of x [km]

- $MaxPIP_{k,r}$ - Maximum effective range (furthest out range with non-zero SSPK) for interceptor-threat pairing (k, r) . Function of $SSPK_{k,r}(x)$ [km]
- $MinPIP_{k,r}$ - Minimum effective range (closest in range with non-zero SSPK) for interceptor-threat pairing (k, r) . Function of $SSPK_{k,r}(x)$ [km]
- $MaxLR_{k,r}$ - Range of threat r such that when interceptor k is fired, the PIP is at its farthest effective range $MaxPIP_{k,r}$ [km]
- $MinLR_{k,r}$ - Range of threat r such that when interceptor k is fired, the PIP is at its closest effective range $MinPIP_{k,r}$ [km]
- $MaxSLR_{f,s,r}$ - Smallest value of $MaxLR_{k,r}$ across all interceptors $k \in K_{f,s}$ [km]
- $MinSLR_{f,s,r}$ - Largest value of $MaxLR_{k,r}$ across all interceptors $k \in K_{f,s}$ [km]
- $\hat{R}_{f,s,r}$ - Actual LR of salvo s in firing sequence f [km]
- $\widehat{PIP}_{f,s,k,r}$ - PIP of interceptor k of salvo s in firing sequence f against threat r , assuming the salvo is fired at LR $\hat{R}_{f,s,r}$ [km]
- $\widehat{SSPK}_{f,k,r,s}$ - SSPK that corresponds to $\widehat{PIP}_{f,s,k,r}$ of interceptor k of salvo s in firing sequence f against threat r
- $PEK_{f,s,r}$ - The probability of engagement kill for all interceptors in salvo s of firing sequence f against threat r , assuming the salvo is fired at LR $\hat{R}_{f,s,r}$
- $\widehat{MinPIP}_{f,s,r} = \min_{k \in K_{f,s}} \{\widehat{PIP}_{f,s,k,r}\}$ - The closest PIP for salvo s in firing sequence f against threat r given the salvo's LR $\hat{R}_{f,s,r}$ [km]
- $PRA_{f,r}$ - Probability of raid annihilation for firing sequence f against threat r

Multiple parameters such as $R_{f,s}$, $\hat{R}_{f,s,r}$, $PEK_{f,s,r}$, $\widehat{MinPIP}_{f,s,r}$, and $PRA_{f,r}$ depend on initial detection range D . For notation succinctness, we opt to suppress this functional relationship.

3.3 Single Salvo Probability of Kill

This section describes how we compute the kill probability of a single salvo s of firing sequence f against one red threat r . This requires determining the launch range LR of the salvo and the PIP of each interceptor in the salvo against the threat. We show the mathematical details in Section 3.3.1 and illustrate the concepts with examples in Section 3.3.2.

3.3.1 Description

The inputs to this component of the model are:

- r Threat type
- s Salvo
- $n_{f,s,k}$ Number of interceptors of type k in salvo s of firing sequence f
- $R_{f,s}$ Maximum firing range of the salvo

$R_{f,s}$ is the earliest firing opportunity for the salvo, meaning that if it is within the effective LR for the salvo, we would launch immediately. If this is the first salvo in the firing sequence $s = 1$, $R_{f,s}$ is set equal to the DCR denoted as D . Any $R_{f,s}$ with $s > 1$ depends upon the outcome of previous salvos in the firing sequence. However, the actual LR of the salvo is the first point when all interceptors in the salvo have a positive SSPK. Consequently, the LR may be less than $R_{f,s}$. If no LR exists such that all interceptors in the salvo have a positive SSPK, we deem the salvo infeasible.

Based on the threat type r and salvo sizes $n_{f,s,k}$, we can examine all IT pairings for the salvo. The key inputs needed from the IT pairs are the velocity of the threat v_r , the velocity of an interceptor v_k , and SSPK as a function of PIP: $SSPK_{k,r}(x)$ where x is the PIP. Based on this SSPK function, we determine the closest PIP to generate a positive SSPK and the farthest PIP to generate a positive SSPK. Formally;

$$MaxPIP_{k,r} = \max\{x : SSPK_{k,r}(x) > 0\} \quad (3.1)$$

$$MinPIP_{k,r} = \min\{x : SSPK_{k,r}(x) > 0\} \quad (3.2)$$

We assume that the SSPK is positive within the interval: $SSPK_{k,r}(x) > 0, \forall x \in [MinPIP_{k,r}, MaxPIP_{k,r}]$. We define the interval $[MinPIP_{k,r}, MaxPIP_{k,r}]$ as the PIP interval. The PIP interval represents the effective operational range of the interceptor k against threat r . Firing the interceptor such that the PIP is outside of this interval would be wasteful.

We next determine the LR $\hat{R}_{f,s,r}$ for the salvo. This calculation requires several preliminary steps. First we convert the PIP interval to a LR interval. Recall the PIP is the range of impact

point and the LR is the range when the salvo is fired. The formulas for the bounds of the LR interval use the PIP interval and velocities:

$$MaxLR_{k,r} = \frac{MaxPIP_{k,r}}{v_k}v_r + MaxPIP_{k,r} \quad (3.3)$$

$$MinLR_{k,r} = \frac{MinPIP_{k,r}}{v_k}v_r + MinPIP_{k,r} \quad (3.4)$$

We now have LR intervals for each interceptor type k in the salvo s . A salvo is feasible if an LR exists such that all interceptors in the salvo have positive SSPK. A necessary and sufficient condition for feasibility is the intersection of all LR intervals across the interceptor types in the salvo is nonempty. To determine the intersection (and hence feasibility) status we define:

$$MaxSLR_{f,s,r} = \min_{k \in K_{f,s}} \{MaxLR_{k,r}\} \quad (3.5)$$

$$MinSLR_{f,s,r} = \max_{k \in K_{f,s}} \{MinLR_{k,r}\} \quad (3.6)$$

To check if the intersection is not empty, we need the upper bound on LR given by equation (3.5) to be greater than the lower bound in equation (3.6). This feasibility condition is expressed as:

$$\text{Salvo is feasible iff } MaxSLR_{f,s,r} \geq MinSLR_{f,s,r} \quad (3.7)$$

We next check the validity of the salvo. Validity is a separate condition from feasibility. Feasibility checks whether the interceptors within a salvo are compatible with each other: can all interceptors impact the target with positive probability? Validity checks whether it is too late to effectively fire the salvo due to the threat being too close to the target. Mathematically we check for validity by comparing the salvo's initial range with the salvo's LR. The salvo is valid if:

$$R_{f,s,r} \geq MinSLR_{f,s,r}$$

We define $[MinSLR_{f,s,r}, MaxSLR_{f,s,r}]$ as the launch window as any LR in this interval will yield a feasible salvo. Blue launches the salvo at the earliest opportunity that generates a feasible salvo. To determine $\hat{R}_{f,s,r}$ we compare the initial firing opportunity $R_{f,s}$ to the launch window. If $R_{f,s} \geq MaxSLR_{f,s,r}$, then the threat is too far out and a fire delay must be introduced by setting $\hat{R}_{f,s,r} = MaxSLR_{f,s,r}$. If $R_{f,s}$ is within the launch window, the salvo is currently valid and is fired immediately such that $\hat{R}_{f,s,r} = R_{f,s}$. Finally, if $R_{f,s} < MinSLR_{f,s,r}$, the threat is invalid as it is below the minimum effective range of the salvo.

$$\hat{R}_{f,s,r} = \begin{cases} MaxSLR_{f,s,r} & \text{if } R_{f,s} \geq MaxSLR_{f,s,r} \\ R_{f,s} & \text{if } MaxSLR_{f,s,r} > R_{f,s} \geq MinSLR_{f,s,r} \end{cases} \quad (3.8)$$

For the remainder of this section we assume the salvo is feasible and valid. Because all interceptors in the salvo are launched at the same time, the PIP for each interceptor follows immediately from $\hat{R}_{f,s}$:

$$\widehat{PIP}_{f,s,k,r} = \hat{R}_{f,s,r} - \frac{\hat{R}_{f,s,r}}{v_k + v_r} v_r \quad (3.9)$$

The notation $\widehat{PIP}_{f,s,k,r}$ implies the PIP is computed based on $\hat{R}_{f,s,r}$. The SSPK for interceptor type k is:

$$\widehat{SSPK}_{f,k,r,s} = SSPK_{k,r}(\widehat{PIP}_{f,s,k,r}) \quad (3.10)$$

Again context implies that $SSPK_{f,k,r,s}$ is computed based on $\hat{R}_{f,s}$.

With the SSPK of each interceptor type known, we can now combine them to calculate to Salvo's PEK. First, we denote $1 - SSPK_{f,k,r,s}$ as the complement of $SSPK_{f,k,r,s}$ to represent the miss probability. Next, recall that there are $n_{f,s,k}$ interceptors of type k in the salvo. Finally, by independence across interceptors, the probability the salvo results in a kill is:

$$PEK_{f,s,r} = 1 - \prod_{k \in K_{f,s}} (1 - \widehat{SSPK}_{f,s,k,r})^{n_{f,s,k}} \quad (3.11)$$

3.3.2 Single Salvo Example

To illustrate this process, we define threat $r = Z$ and interceptors set $K = \{A, B, C, D\}$. We decide to launch a salvo with three types $\{A, C, D\}$ with $n_{f,s,k} = [1, 1, 1]$ for $k \in \{A, C, D\}$. The relevant velocities appear in Table 3.1:

Table 3.1. Velocities v for each missile in $Z \cup K$.

	Z	A	B	C	D
v	1	1	2	0.8	2.5

Additionally the $SSPK_{k,r}(x)$ for $x = PIP_{k,r}$ is provided in Table 3.2. For simplicity we assume $SSPK_{k,r}$ is a constant between $MaxPIP_{k,r}$ and $MinPIP_{k,r}$, and 0 outside these bounds. However, our method easily handles arbitrary functions for $SSPK_{k,r}(x)$. Since this is our first (and only) salvo, we set $R_{f,s} = D$, which is 180 km for this scenario.

Table 3.2. $MaxPIP_{k,r}$ and $MinPIP_{k,r}$ with their respective $SSPK_{k,r}(x)$. $MaxLR_{k,r}$ and $MinLR_{k,r}$ are also included. However, they are dependent on their respective $PIP_{k,r}$.

	A		B		C		D	
	Max	Min	Max	Min	Max	Min	Max	Min
$PIP_{k,r}[km]$	90	40	70	25	65	25	30	10
$LR_{k,r}[km]$	180	80	105	37.5	146.25	56.25	42	14
$SSPK_{k,r}(x)$	0.8	0.8	0.65	0.65	0.5	0.5	0.9	0.9

While $MaxLR_{k,r}$ and $MinPIP_{k,r}$ are included in the table, these are not given inputs. We use interceptor D with equations (3.3) and (3.4) to illustrate how the LR interval for interceptor D is calculated:

$$MaxLR_{D,Z} = \frac{30 \text{ km}}{2.5 \text{ km/s}} \cdot 1 \text{ km/s} + 30 \text{ km} = 42 \text{ km}$$

$$MinLR_{D,Z} = \frac{10 \text{ km}}{2.5 \text{ km/s}} \cdot 1 \text{ km/s} + 10 \text{ km} = 14 \text{ km}$$

Given that we are using the salvo {A, C, D}, we generate the $MaxSLR_{f,s,r}$ and $MinSLR_{f,s,r}$ using equation (3.5) and (3.6):

$$MaxSLR_{f,s,r} = Min\{180, 146.25, 42\} = 42 \text{ km}$$

$$MinSLR_{f,s,r} = Max\{80, 56.25, 14\} = 80 \text{ km}$$

With these values and equation (3.7), we have an infeasible salvo. Looking at our interceptors, we see that the upper bound of D's LR ($MaxLR_{D,Z} = 42 \text{ km}$) is less than the lower bound of the LRs of both A and C ($MinLR_{C,Z} = 56.25 \text{ km}$ and $MinLR_{A,Z} = 80 \text{ km}$). If we remove D from the salvo, the salvo becomes feasible. We replace D with a second C such that $n_{f,s,k} = [1, 2]$ for $k \in \{A, C\}$. Our new $MaxSLR_{f,s,r}$ and $MinSLR_{f,s,r}$ are:

$$MaxSLR_{f,s,r} = Min\{180, 146.25\} = 146.25 \text{ km}$$

$$MinSLR_{f,s,r} = Max\{80, 56.25\} = 80 \text{ km}$$

When we evaluate these values in (3.7), we can confirm the salvo is now feasible as $MaxSLR_{f,s,r} \geq MinSLR_{f,s,r}$. Because this is the first salvo, we set $R_{f,s} = D = 180 \text{ km}$. We determine $\hat{R}_{f,s,r}$ according to equation (3.8). Of our three conditions, the first one is true as $R_{f,s} > MaxSLR_{f,s,r}$ meaning the threat is out of range for at least one interceptor in the salvo. We must introduce a delay, and assign $\hat{R}_{f,s,r} = MaxSLR_{f,s,r} = 146.25 \text{ km}$.

With $\hat{R}_{f,s,r}$, $\widehat{PIP}_{f,s,k,r}$ is now calculated for each interceptor in the salvo according to equation (3.9):

$$\widehat{PIP}_{f,s,A,Z} = 146.25 \text{ km} - \frac{146.25 \text{ km}}{1 \text{ km/s} + 1 \text{ km/s}} 1 \text{ km/s} = 73.13 \text{ km}$$

$$\widehat{PIP}_{f,s,C,Z} = 146.25 \text{ km} - \frac{146.25 \text{ km}}{0.8 \text{ km/s} + 1 \text{ km/s}} 1 \text{ km/s} = 65.89 \text{ km}$$

These values are then substituted into $SSPK_{f,k,r,s}(x)$ in equation (3.10) to give us $\widehat{SSPK}_{f,s,A,Z} = 0.8$ and $\widehat{SSPK}_{f,s,C,Z} = 0.5$. Finally we use equation (3.11) to give us $PEK_{f,s,r} = 1 - (1 - 0.8)^1(1 - 0.5)^2 = 0.95$.

3.4 Firing Sequence Probability of Kill

We now expand our single salvo model of Section 3.3 into a sequence of salvos. The salvos are fired sequentially and the validity of a salvo is determined by its predecessor salvo(s). Key components of our analysis include the LR of each salvo, the closest PIP in a salvo, and the BDA time delay. If any individual salvo is infeasible or invalid, the entire sequence is designated as such and we assign a PRA of zero, which precludes the sequence from further consideration in determining the best firing sequence.

3.4.1 Description

The inputs to this portion of the model are:

- D DCR of the threat
- r Threat type
- f Firing sequence
- $n_{f,s,k}$ Number of interceptor of type k in salvo s of the firing sequence f

We start with the first salvo, $s = 1$ and set the initial firing opportunity $R_{f,s}$ equal to the DCR denoted by D . We next check for feasibility and validity using equations (3.7) and (3.8). If the salvo is infeasible or invalid, we do not evaluate any subsequent salvos in the sequence and set $PRA_{f,r} = 0$. If the salvo is feasible and valid, our LR is $\hat{R}_{f,sr}$ given by equation (3.8), which yields $\widehat{PIP}_{f,s,k,r}$, $\widehat{SSPK}_{f,s,k,r}$, $PEK_{f,s,r}$ from equations (3.9) - (3.11).

To move to salvo $s = 2$ requires the initial firing range $R_{f,s}$. We compute $R_{f,2}$ based on the PIPs from salvo $s = 1$. Specifically, the closest PIP of salvo $s = 1$ is:

$$\widehat{MinPIP}_{f,s,r} = \min_{k \in K_{f,s}} \{\widehat{PIP}_{f,s,k,r}\} \quad (3.12)$$

At this closest PIP, all of the interceptor in salvo $s = 1$ have engaged with the threat and

blue can perform BDA. After a BDA delay of t^{BDA} , blue is cleared to fire its second salvo. Hence, the earliest firing opportunity for this salvo satisfies:

$$R_{f,s+1} = \widehat{MinPIP}_{f,s,r} - t^{BDA} \cdot v_r \quad (3.13)$$

While we illustrate the process of the first salvo of a sequence, $s = 1$, the steps utilized apply for any salvo s . The only difference between salvos 1 and s is that salvo 1 uses D for R_{f1} and salvo s uses equation (3.13). Each salvo is checked for feasibility and validity. Assuming all salvos are feasible, a global condition on validity is:

$$\text{Firing sequence } f \text{ is valid iff } R_{f,s} \geq MinSLR_{f,s,r} \quad \forall s \in S_f \quad (3.14)$$

If any salvo is deemed infeasible or invalid, the $PRA_{f,r}$ of the entire sequence f is set equal to 0 so that it is not considered for further analysis. Assuming all salvos are feasible and valid, we calculate the firing sequence $PRA_{f,r}$ by multiplying the probability of engagement miss, denoted by $1 - PEK_{f,s,r}$, together for each $s \in S_f$:

$$PRA_{f,r} = 1 - \prod_{s \in S_f} (1 - PEK_{f,s,r}) \quad (3.15)$$

3.5 Firing Sequence Example

Returning to the salvo example utilized in Section 3.3.2, we use the same interceptors $K = \{A, B, C, D\}$, threat $r = Z$, velocities as in Table 3.1, and $SSPK_{k,r}(x)$ as in Table 3.2. We want to put together a sequence of salvos to eliminate the threat. Interceptor A is a long range, highly capable; C has mid-range capability; B is cheaper with reasonable range; and, D has short range but is highly capable. We define our firing sequence f as $[\{A\} \text{ look } \{B, C\} \text{ look } \{B, D\}]$ and the associated $n_{f,s,k}$ for each salvo is provided in Table 3.3: For our sequence, we select a DCR of $D = 200$ km and $t_{BDA} = 18$ s.

Looking at Table 3.2 and using equations (3.1)-(3.5) we can generate the $MaxSLR_{f,s,Z}$ and $MinSLR_{f,s,Z}$ for each salvo. To illustrate this process, we first calculate the $MaxLR_{k,Z}$ and

Table 3.3. Salvo index s , the interceptor types k in that salvo, and $n_{f,s,k}$ for each salvo in f with $[\{A\} \text{ look } \{B,C\} \text{ look } \{B,D\}]$. Note $MaxSLR_{f,s,r}$ and $MinSLR_{f,s,r}$ are not inputs but calculated quantities.

s	1	2	3
k	{A}	{B,C}	{B,D}
$n_{f,s,k}$	[1, 0, 0, 0]	[0, 1, 1, 0]	[0, 1, 0, 1]
$MaxSLR_{f,s,r}$ [km]	180	105	42
$MinSLR_{f,s,r}$ [km]	90	56.25	37.5

$MinLR_{k,Z}$ for each interceptor type. Below are the equations for type B:

$$MaxLR_{B,Z} = \frac{70 \text{ km}}{2 \text{ km/s}} \cdot 1 \text{ km/s} + 70 \text{ km} = 105 \text{ km}$$

$$MinLR_{B,Z} = \frac{25 \text{ km}}{2 \text{ km/s}} \cdot 1 \text{ km/s} + 25 \text{ km} = 37.5 \text{ km}$$

Repeating the process for interceptor type C, we obtain $MaxLR_{C,Z} = 146.25$ and $MinLR_{C,Z} = 56.25$ km. Finally, using equations (3.5)-(3.6) to find the salvos minimum and maximum LR, salvo $s = 2$ consisting of {B,C} yields:

$$MaxSLR_{f,1,Z} = \text{Min}\{105, 146.25\} = 105\text{km}$$

$$MinSLR_{f,1,Z} = \text{Max}\{37.5, 56.25\} = 56.25\text{km}$$

From equation (3.7), we can confirm all of these salvos are feasible as $MaxSLR_{f,s,Z} \geq MinSLR_{f,s,Z} \forall s$.

We set $R_{f,1} = D = 200$ km, and since $MaxSLR_{f,1,Z}$ is 180 for A, equation (3.8) dictates a delay and $\hat{R}_{f,1} = MaxSLR_{f,1,Z} = 180$ km. Using equation (3.9) to calculate the impact point, we acquire:

$$\widehat{PIP}_{f,s,A,Z} = 180 \text{ km} - \frac{180 \text{ km}}{1 \text{ km/s} + 1 \text{ km/s}} 1 \text{ km/s} = 90 \text{ km}$$

Inserting this into equation (3.10) and referencing the probabilities in Table 3.2, we attain $\widehat{SSPK}_{f,1,A,Z} = 0.8$ and since A is the only interceptor in the salvo, $PEK_{f,1,Z} = 0.8$. To move on to salvo two, we need the $\widehat{MinPIP}_{f,1,Z}$ according to equation (3.12), which is just $\widehat{MinPIP}_{f,1,Z} = 90$ km as A is the only interceptor. We now apply our BDA length of $t_{BDA} = 18$ s. To obtain the earliest firing point of salvo two, we apply equation (3.13):

$$R_{f,2} = 90 \text{ km} - 18 \text{ s} \cdot 1 \text{ km/s} = 72 \text{ km}$$

Given $s = 2$ and $k = \{B,C\}$, we look up the effective LR of salvo two in Table 3.3 and find $MaxSLR_{f,2,Z} = 105$ km, $MinSLR_{f,2,Z} = 56.25$ km. We are already within the effective LR at a range of 72 km, and equation (3.8) sets $\hat{R}_{f,2} = R_{f,2} = 72$ km, confirming the salvo is valid. We use equation (3.9) to calculate impact points:

$$\begin{aligned}\widehat{PIP}_{f,s,B,Z} &= 72 \text{ km} - \frac{72 \text{ km}}{2 \text{ km/s} + 1 \text{ km/s}} 1 \text{ km/s} = 48 \text{ km} \\ \widehat{PIP}_{f,s,C,Z} &= 72 \text{ km} - \frac{72 \text{ km}}{0.8 \text{ km/s} + 1 \text{ km/s}} 1 \text{ km/s} = 32 \text{ km}\end{aligned}$$

Combining this into equation (3.10) yields $\widehat{SSPK}_{f,1,B,r} = 0.65$ and $\widehat{SSPK}_{f,1,B,r} = 0.5$. Equation (3.11) shows that our $PEK_{f,2,Z} = 1 - (1 - 0.65)(1 - 0.5) = 0.825$. To advance to the next salvo, (3.12) dictates $\widehat{MinPIP}_{f,s,r} = \min\{48, 32\} = 32$ km. Applying t_{BDA} with equation (3.13), we get:

$$R_{f,3} = 32 \text{ km} - 18 \text{ s} \cdot 1 \text{ km/s} = 14 \text{ km}$$

For $s = 3$ with $k = \{B,D\}$, we look up the effective salvo LR of salvo 3 in Table 3.3 and find $MaxSLR_{f,3,Z} = 42$ km, $MinSLR_{f,3,Z} = 37.5$ km. While $\{B,D\}$ is a feasible salvo, equation (3.8) demonstrates that this is invalid as $R_{f,3} < MinSLR_{f,3,Z}$. This means our entire sequence f is invalid as currently defined. However, if we were to drop the last salvo, and define f as $[\{A\} \text{ look } \{B,C\}]$, this salvo is valid according to equation (3.14) and can use equation (3.15) to find $PRA_{f,r}$:

$$PRA_{f,r} = 1 - (1 - 0.8)(1 - 0.825) = 0.96$$

Since type B causes the invalidity according to Table 3.2, we can drop B from salvo 3 and explore $f = [\{A\} \text{ look } \{B,C\} \text{ look } \{D\}]$. This new sequence is similar to Table 3.3, with the last column $s = 3$ switched to:

$$k = \{B, D\}, n_{f,s,k} = [0, 0, 0, 1], MaxSLR_{f,s,r} = 42 \text{ km}, MinSLR_{f,s,r} = 14 \text{ km}$$

Resuming at the third salvo and $R_{f,3} = 14$, according to our new values we find $MaxSLR_{f,3,Z} = 42 \text{ km}$, and $MinSLR_{f,3,Z} = 14 \text{ km}$. This means the threat is within effective range and $\hat{R}_{f,3} = R_{f,3} = 14 \text{ km}$ according to equation (3.8). Although the standard process would use equation (3.9) to calculate the impact point, since we are exactly at $MinLR_{f,s,D,Z}$ we highlight that $MinLR_{f,s,D,Z}$ and $MinPIP_{D,Z}$ are related through equation (3.4). Instead of using equation (3.12) to calculate $\widehat{PIP}_{f,s,D,Z}$ we instead assign $\widehat{PIP}_{f,s,D,Z} = MinPI_{D,Z} = 10 \text{ km}$.

Utilizing equation (3.11) provides $SSPK_{f,3,D,Z}$ of 0.9 for interceptor D and since it is the only interceptor, $PEK_{f,3,Z} = SSPK_{f,3,D,Z} = 0.9$. There are no more salvos and we use equation (3.15) to calculate our $PRA_{f,r}$:

$$PRA_{f,r} = 1 - (1 - 0.8)(1 - 0.825)(1 - 0.9) = 0.9965$$

3.6 Best Firing Sequence

Now that we know how to evaluate a given firing sequence, we can determine the best one. We do this by generating all feasible and valid firing sequences and finding the lowest cost sequence. We define a cost function in Section 3.6.1; generate the set of all possible salvos in Section 3.6.2; enumerate all firing sequences in Section 3.6.3; improve our enumeration in Section 3.6.4; and summarize the analysis in Section 3.6.5 to specify the best sequence.

3.6.1 Cost Function

We utilize the following additional notation:

- $K_f = \cup_{s \in S_f} K_{f,s}$ - Set of interceptors across all salvos s in firing sequence f
- V_k - Cost of interceptor k
- $X_{f,s,k}$ - Random variable denoting the number of interceptors k expended in salvo s of firing sequence f
- $Y_{f,k}$ - Random variable denoting number of interceptors k expended in firing sequence f
- $cost_f$ - Random variable denoting cost of firing sequence f

To determine the best firing sequence, we only consider firing sequences that have surpass a certain PRA threshold Q such that $PRA_{f,r} > Q$. From the sequences that with the requisite PRA, we choose the sequence with the minimum expected cost. Each interceptor has a cost V_k which could be a financial cost, opportunity cost, or operational cost. To calculate the expected cost, we examine the actual number of interceptors of type k fired in salvo s of sequence f . We denote this quantity $X_{f,s,k}$, which is a random variable because the salvo s is only fired if the threat has not been killed on previous salvos. $X_{f,s,k}$ takes on two possible values and has PMF:

$$X_{f,s,k} = \begin{cases} n_{f,s,k} & \text{with probability } \prod_{j=1}^{s-1} (1 - PEK_{f,j,r}) \\ 0 & \text{otherwise} \end{cases} \quad (3.16)$$

Hence the expected number of interceptors of type k fired in salvo s of sequence f is the product of the number of type k interceptors in salvo s ($n_{f,s,k}$) and the probability that all previous salvos have missed.

$$E[X_{f,s,k}] = n_{f,s,k} \left(\prod_{j=1}^{s-1} (1 - PEK_{f,j,r}) \right) \quad (3.17)$$

In order to obtain the expected number of type k interceptors fired in the firing sequence, we define:

$$E[Y_{f,k}] = \sum_{s \in S_f} E[X_{f,s,k}] \quad (3.18)$$

We now apply the cost of each interceptor, V_k to $E[Y_{f,k}]$. Although it can be the monetary value of the interceptor, V_k is a weighting to assign how interceptors are valued when compared to one another. This weighted value summed across all interceptors gives us our overall cost of the firing sequence:

$$E[cost_f] = \sum_{k \in K_f} E[Y_{f,k}] \cdot V_k \quad (3.19)$$

If we set $V_k = 1$ for all types k , then we minimize expected missiles expended. If we want to minimize a particular type j , we can set $V_k = 0$ for $k \neq j$ and $V_k = 1$ for $k = j$.

3.6.2 Salvo Combinations

We utilize the following additional notation:

- C_r - Set of feasible salvos against threat type r

Before enumerating all firing sequences, we first generate the set C_r of all feasible salvos for a given threat r . Each salvo in C_r salvo can contain one, two, or three interceptors and we allow for repeated elements while ignoring ordering. For example, if K has four types $\{A, B, C, D\}$ and we want to generate the salvos in C_r of size two, Table 3.4 provides the 10 possible salvos.

Table 3.4. Combination with replacement for salvos of size two given the set of feasible interceptors K of $\{A, B, C, D\}$.

AA	AB	AC	AD
	BB	BC	BD
		CC	CD
			DD

Notice the blank elements of the table, where elements such as BA and CA would exist if they were not duplicates of AB and AC respectively. This type of counting is slightly more complicated than the standard combination and is called a combination with replacement (or repetition). Stanley (2012) defines this function its own unique notation, denoted by double parenthesis. Using this established notation, the equation to generate all salvos up to size three is:

$$C_r = \bigcup_{i=1}^3 \left(\binom{K}{i} \right) \quad (3.20)$$

The number of elements in set k is denoted by $|K|$, the number of elements in a combination with replacement of $|K|$ choose i is $\binom{|K|+i-1}{i}$. Summing across salvos of size one, two and three gives us the total elements in C_r :

$$|C_r| = \sum_{i=1}^3 \binom{|K|+i-1}{i} \quad (3.21)$$

In Table 3.5 we present values of $\text{size}(C_r)$ for different values of K generated by the equation 3.21:

Table 3.5. $|C_r|$ given a sampling of $|K|$ interceptor

$ K $	5	10	15	20
$ C_r $	55	285	815	1770

However, this number does not account for feasibility and is an upper bound on the number of salvos we may have. We determine and remove infeasible salvos using equations (3.5)-(3.7) before doing any other analysis. The resulting C_r is the set of all feasible salvos up to size three.

3.6.3 Enumeration

From the set of all feasible salvos, we now illustrate how to generate all feasible and valid firing sequence using a brute force approach. Recall that F is the set of all firing sequences; We now denote F_j as the set of all firing sequences that fire j salvos. The first of these, F_1 , contains all firing sequences f with only one salvo and is equal to the set of all salvo combinations:

$$F_1 = C_r \quad (3.22)$$

To build firing sequences of two salvos, we introduce a new element not included in C_r : “look”. This element denotes separation between the different salvos s within a firing sequence f . The set F_2 contains all firing sequences with two salvos and is generated according to $F_2 = C_r \cdot C_R$. Generalizing this formula, where F_j contains all salvos of size j , we get:

$$F_j = (C_r)^j \quad (3.23)$$

From this equation, we generate new F_j as long as it contains at least one valid sequence. We denote J as the smallest value such that F_J contains no valid sequences. This allows us to combine all F_j for $j < J$ to procure F , the set such that contains all the valid firing sequences:

$$F = \bigcup_{j=1}^{J-1} F_j \quad (3.24)$$

This brute force method generates all valid sequences, but it also generates many invalid sequences and becomes computationally infeasible rapidly. Recalling equation (3.21), salvos of size three in C_r “grow” the quickest as they are $\binom{|K|+2}{3}$. This combination can be transformed into:

$$\frac{|K|(|K| + 1)(|K| + 2)}{6}$$

This is known as the tetrahedral sequence which grows as $O(n^3)$ where n is the number of elements in set K . However the firing sequence F_j is the set product C_r with itself j times and $\text{size}(F_j) = O(|K|^{3j})$. Finally $|F| = O(|K|^{3J})$, where $J - 1$ is the maximum number of salvos that can be fired.

Note J depends upon the velocities and ranges of the interceptors and threats. While J tends to be relatively small (≤ 5), if one interceptor is very fast and has a small minimum range, J can become large. To visualize the complexity, the resulting $|F|$ for a collection of $|K|$ and j are displayed in Table 3.6.

Table 3.6. $|F|$ given J and $|K|$. Note this is the actual size and not the $O(\cdot)$ approximation.

	$ K $				
	1	3	5	7	9
1	3	19	55	119	219
2	9	361	3025	14161	47961
3	27	6859	166375	1685159	1.05E+07
4	81	130321	9150625	2.01E+08	2.30E+09
5	243	2476099	5.03E+08	2.39E+10	5.04E+11
6	729	4.70E+07	2.77E+10	2.84E+12	1.10E+14

3.6.4 Efficient Enumeration Through Recursion

To improve the computational efficiency of our enumeration, we utilize recursion to generate the set F without any invalid sequences. With the recursive approach described here, once we determine a firing sequence f is invalid, we never consider appending additional salvos to the end of f . The naïve approach described in Section 3.6.3 includes these (obviously invalid) firing sequences, which increase the size of F with no benefit.

Our recursion approach generates a tree structure where we continue moving down a branch

by adding salvos onto a firing sequence until the sequence is invalid. At that point, we move back up the current branch one level, and then head down another branch by appending the next salvo s in C_r to the firing sequence. Once we have iterated through all salvos s in C_r , we move back one further level of the original branch, and continue the process. In this fashion we include all valid firing sequences, while minimizing the amount of exploration done with invalid sequences.

Below is pseudo code for the recursive function to generate all valid firing sequences:

Inputs

- f - current firing sequence
- F - list of valid firing sequences

Function

GenerateFiringSequence(f, F)

 If f is invalid:

 Return

 Add firing sequence f to sequence list F

 For salvos s in C_r :

$f' \leftarrow$ append salvo s to end of sequence f

 GenerateFiringSequence(f', F)

 Return

To generate the complete final firing sequence list, we call the function GenerateFiringSequence($f = [], F = []$), which sends as inputs the empty firing sequence (i.e., no salvos) and an empty list of firing sequences. Our actual implementation is more sophisticated in a couple ways to help improve efficiency, primarily by keeping track of ranges separate from the firing sequences. This allows us to only check whether the newest salvo is valid as opposed to checking an entire firing sequence. It also allows us to only consider a smaller and smaller subset of C_r as we descend down the tree by tracking and removing any invalid salvos.

We demonstrate this in Figure 3.1 with a subset from earlier examples in Sections 3.3.2

Table 3.7. All firing sequences for $k = \{B, C, D\}$ and $D = 200$ km with salvo separation denoted by “L”. This figure only includes sequences starting with a salvo consisting of a single “B”, and the salvos are all single salvos from the set C_r . $\widehat{MinPIP}_{f,s,r}$ is impact point of the last salvo in the sequence.

Sequence	$\widehat{MinPIP}_{f,max(s),r}$
[B]	70.0
[B, L, B]	46.67
[B, L, B, L, B]	31.11
[B, L, B, L, B, L, D]	22.22
[B, L, B, L, B, L, D, L, D]	15.87
[B, L, B, L, B, L, D, L, D, L, D]	11.34
[B, L, B, L, D]	30.0
[B, L, B, L, D, L, D]	21.43
[B, L, B, L, D, L, D, L, D]	15.31
[B, L, B, L, D, L, D, L, D, L, D]	10.94
[B, L, C]	31.11
[B, L, C, L, D]	22.22
[B, L, C, L, D, L, D]	15.87
[B, L, C, L, D, L, D, L, D]	11.34
[B, L, D]	30.0
[B, L, D, L, D]	21.43
[B, L, D, L, D, L, D]	15.31
[B, L, D, L, D, L, D, L, D]	10.94

and 3.5; interceptors $K = \{B, C, D\}$, threat $r = Z$, velocities as in Table 3.1, $SSPK_{k,r}(x)$ as in Table 3.2, $D = 200$ km, and $t_{BDA} = 18$ s. We only look at the single salvos for our recursion to simplify the illustration and use “L” to denote salvo separation. Table 3.7 shows all valid sequences that start with salvo “B” along with the $\widehat{MinPIP}_{f,s,r}$ of the last salvo of a sequence. Figure 3.1 illustrates how all of these sequences were generated.

Efficient recursion example

```

-> f([B], 70.0)
  -> f([B, L, B], 46.67)
    -> f([B, L, B, L, B], 31.11)
      -> f([B, L, B, L, B, L, D], 22.22)
        -> f([B, L, B, L, B, L, D, L, D], 15.87)
          -> f([B, L, B, L, B, L, D, L, D, L,
                D], 11.34)
            <- No Valid Salvos
          <- No Valid Salvos
        <- No Valid Salvos
      <- No Valid Salvos
    -> f([B, L, B, L, D], 30.0)
      -> f([B, L, B, L, D, L, D], 21.43)
        -> f([B, L, B, L, D, L, D, L, D], 15.31)
          -> f([B, L, B, L, D, L, D, L, D, L,
                D], 10.94)
            <- No Valid Salvos
          <- No Valid Salvos
        <- No Valid Salvos
      <- No Valid Salvos
    <- No Valid Salvos
  -> f([B, L, C], 31.11)
    -> f([B, L, C, L, D], 22.22)
      -> f([B, L, C, L, D, L, D], 15.87)
        -> f([B, L, C, L, D, L, D, L, D], 11.34)
          <- No Valid Salvos
        <- No Valid Salvos
      <- No Valid Salvos
    <- No Valid Salvos
  <- No Valid Salvos
-> f([B, L, D], 30.0)
  -> f([B, L, D, L, D], 21.43)
    -> f([B, L, D, L, D, L, D], 15.31)
      -> f([B, L, D, L, D, L, D, L, D], 10.94)

```

Figure 3.1. Recursion demonstrating how a subset of sequences are generated for $k = \{B, C, D\}$ starting D (DCR) = 200 km. This recursion only includes single salvos for simplicity. Only sequences starting with salvo "B" are displayed, however this recursion continues until all sequences are generated.

3.6.5 Evaluation

We can now determine the best firing sequence in F . First recall we only want to consider firing sequences where $PRA_{f,r} \geq Q$, where Q is defined by the user. When generating F , we have only focused on feasibility and validity and did not consider Q . Unfortunately we cannot use Q to improve the efficiency of the enumeration because during the recursion a sequence f that has $PRA_{f,r} < Q$ might have “child” sequences with $PRA_{f,r} \geq Q$. We mention in 6.3.1 the idea of using Q or modified version $\hat{Q}$ to serve as an upper bound to the PRA. This feature would cut off generation to prevent overkill firing sequences and improve model efficiency. We now filter F through equation (3.25) to acquire F_Q :

$$F_Q = \{f \in F \mid PRA_{f,r} \geq Q\} \quad (3.25)$$

For every sequence in F_Q , we have their sequence index f , $PRA_{f,r}$, $n_{f,s,k} \forall s \in f$. We return to the valid sequence utilized in Section 3.5 of $[\{A\} \text{ look } \{B,C\} \text{ look } \{D\}]$ with the relevant information we need stored in Table 3.8.

Table 3.8. Salvo index s , the interceptor types k in that salvo, and $n_{f,s,k}$ for each salvo in f with $[\{A\} \text{ look } \{B,C\} \text{ look } \{D\}]$. It also contains the $PEK_{f,s,r}$ for each salvo as well as the cumulative $PRA_{f,r}$. Note $\prod_{s=1}^3 PEK_{f,s,r} = PRA_{f,r}$.

s	1	2	3
k	{A}	{B,C}	{D}
$n_{f,s,k}$	[1, 0, 0, 0]	[0, 1, 1, 0]	[0, 0, 0, 1]
$PEK_{f,s}$	0.8	0.825	0.9
$\prod_{j=1}^s PEK_{f,s}$	0.8	0.965	0.9965

We first compute the expected expenditure for each interceptor type, which according to equation (3.18) is the sum of $E[X_{f,s,k}]$, the expected number of interceptors used per salvo. This is found by multiplying the number of interceptor in the salvo by the cumulative probability of engagement miss. Although we do not explicitly perform the calculation from equation (3.17), it is captured in the calculation $E[Y_{f,k}]$:

$$E[Y_{f,A}] = 1 + (1 - 0.8) \cdot 0 + (1 - 0.965) \cdot 0 = 1$$

$$E[Y_{f,B}] = 0 + (1 - 0.8) \cdot 1 + (1 - 0.965) \cdot 0 = 0.2$$

$$E[Y_{f,C}] = 0 + (1 - 0.8) \cdot 1 + (1 - 0.965) \cdot 0 = 0.2$$

$$E[Y_{f,D}] = 0 + (1 - 0.8) \cdot 0 + (1 - 0.965) \cdot 1 = 0.035$$

With $E[Y_{f,k}] \forall k \in K$, we are now ready to apply our value parameters displayed in Table 3.9:

Table 3.9. Value V_k for each interceptor in K

	A	B	C	D
V_k	2	0.5	0.75	1

To calculate our expected cost, we apply $E[Y_{f,k}] \cdot V_k \forall k \in K$ according to equation (3.19):

$$E[cost_f] = 1 \cdot 2 + 0.2 \cdot .5 + 0.2 \cdot 0.75 + 0.035 \cdot 1 = 2.285$$

After calculating $E[cost_f] \forall f \in F$, we sort $f \in F_Q$ by increasing $E[cost_f]$. The first entry of this sorted F is the lowest cost f that is feasible, valid, and has $PRA_{f,r} > Q$.

CHAPTER 4:

Results for One Interceptor Type

Before delving into the output of our main model, we first explore a simplified scenario. The goal in this chapter is to deduce insights about the allocation of interceptors across different salvos for this simple case. This simplified scenario only has a single interceptor type with constant SSPK independent of range. This implies that interceptors can be allocated to different salvos with no impact on their ability to kill the incoming threat. We denote this constant SSPK as p . BDA delay is not considered in these calculations and instead we provide the number of salvo opportunities M as an input.

We derive analytic results for the two-salvo allocation problem and illustrate numerically that many of the two-salvo results generalize to more sequences with more salvo opportunities. We then show different firing policies effects on PRA and E[Expenditure] for various SSPKs.

To solve for the optimal firing policy, we first determine the total number of interceptors to fire, which we denote by N . We need to expend enough interceptors to surpass the pre-determined minimum probability of interception criterion Q . To ensure PRA is at least Q , N satisfies:

$$Q \leq 1 - (1 - p)^N \quad (4.1)$$

Rearranging yields:

$$N \geq \frac{\log(1 - Q)}{\log(1 - p)} \quad (4.2)$$

Taking the ceiling results in:

$$N = \left\lceil \frac{\log(1 - Q)}{\log(1 - p)} \right\rceil \quad (4.3)$$

Taking the ceiling of the right hand side ensures we satisfy the Q-condition without over expenditure. Since $Q > p$ (otherwise $N=1$ by definition) N in equation (4.3) is greater than 1.

With N known, we examine two special cases based on M , the number of salvo opportunities available to blue. The first is “many salvo opportunities” case where $M \geq N$ as blue has enough time to fire at least N salvos, while the second is the two-salvo case where $M = 2$. We conclude the analysis by studying a general case $2 < M < N$ with an intermediate number of salvos. We denote n_j as the number of interceptors fired in salvo j for an arbitrary firing sequence and use n_j^* to denote the optimal firing sequence.

4.1 Case $M \geq N$

We first consider scenarios where M , the number of salvo opportunities blue has, is large. Specifically, blue detects the threat at a great enough distance, that blue can fire at least $M \geq N$ separate salvos. Recall N is the number of interceptors that satisfies the minimum Q requirement, as defined in equation (4.3). The optimal allocation when more firing opportunities M exist than N required to satisfy Q is one interceptor per salvo. We now show that the optimal firing policy is to fire N salvos, each with one interceptor. If we denote n_s^* as the optimal number of interceptors in salvo s , then $n_s^* = 1$ for $1 \leq s \leq N$. We formalize this in the following proposition:

Proposition: $n_j^* = 1$ for $1 \leq j \leq N$

Proof: We proceed by contradiction. Assume there exists some k such that $n_k^* \geq 2$. Hence, if M^* denotes the number of salvos fired for this sequence, then $M^* < N$. The expected cost of this firing sequence is:

$$C^* = \sum_{j=1}^{M^*} n_j^* (1 - p)^{\sum_{i=1}^{j-1} n_i^*} \quad (4.4)$$

A variation of this equation where the distinct parts $j < k$, $j = k$, and $j > k$ are summed is provided by:

$$C^* = \sum_{j=1}^{k-1} n_j^* (1-p)^{\sum_{i=1}^{j-1} n_i^*} + n_k^* (1-p)^{\sum_{i=1}^{k-1} n_i^*} + \sum_{j=k+1}^{M^*} n_j^* (1-p)^{\sum_{i=1}^{j-1} n_i^*} \quad (4.5)$$

The first term in equation (4.5) is the expected cost for interceptors fired prior to salvo k , last term is expected cost for interceptors fired after salvo k , and middle term is expected cost of salvo k .

We next show that if we split salvo k into multiple, smaller salvos, we will have a new firing sequence with an expected cost smaller than equation (4.4). Specifically, let us define a new firing sequence $\hat{n}$ that is equivalent to n^* , except it splits salvo k of n^* into two separate salvos:

$$\begin{aligned} \hat{n}_j &= n_j^* \quad \text{for } j < k \\ \hat{n}_k &= 1 \\ \hat{n}_{k+1} &= n_k^* - 1 \\ \hat{n}_j &= n_{(j-1)}^* \quad \text{for } j > k + 1 \end{aligned}$$

In this process, we reallocated n_k^* interceptors across two salvos in sequence $\hat{n}$, with a single interceptor in salvo k of $\hat{n}$ and the remaining $n_k^* - 1$ in salvo $k + 1$ of $\hat{n}$. The firing sequence defined by $\hat{n}$ fires $\hat{M} = M^* + 1 \leq N$ salvos. $\hat{C}$ is the cost using sequence $\hat{n}$ and is identical to equation 4.4 except the usage of $\hat{\cdot}$ instead of $\cdot^*$. Breaking the cost equation down for $\hat{C}$ based on k in a similar fashion to equation (4.5) yields:

$$\hat{C} = \sum_{j=1}^{k-1} \hat{n}_j (1-p)^{\sum_{i=1}^{j-1} \hat{n}_i} + \hat{n}_k (1-p)^{\sum_{i=1}^{k-1} \hat{n}_i} + \hat{n}_{k+1} (1-p)^{\sum_{i=1}^k \hat{n}_i} + \sum_{j=k+2}^{M^*+1} \hat{n}_j (1-p)^{\sum_{i=1}^{j-1} \hat{n}_i} \quad (4.6)$$

The first term is again the expected cost for interceptors prior to k while the last term is again expected cost for interceptors fired after k . The difference is that the middle term is now contains two distinct parts for k and $k + 1$. If we compare the two completed equations

(4.5) and (4.6), accounting for relationship between $\hat{n}$ and n^* we see that the first portion where $j < k$ is the same for $\hat{C}$ and C^* , and the last portions of C^* and $\hat{C}$, where $j > k$ in C^* and $j > k + 1$ in $\hat{C}$, are the same. Thus, we can cancel out these portions and only compare the middle sections according to:

$$n_k^* (1 - p)^{\sum_{i=1}^{k-1} n_i} \quad \text{vs} \quad \hat{n}_k (1 - p)^{\sum_{i=1}^{k-1} \hat{n}_i} + \hat{n}_{k+1} (1 - p)^{\sum_{i=1}^k \hat{n}_i} \quad (4.7)$$

Replacing $\hat{n}_k$ with 1 and $\hat{n}_{k+1}$ with $n_k^* - 1$ while factoring out $(1 - p)^{\sum_{i=1}^{k-1} \hat{n}_i}$ produces the inequality:

$$n_k^* \geq 1 + (n_k^* - 1)(1 - p) \quad (4.8)$$

The inequality follows because $n_k^* \geq 1 + (n_k^* - 1)(1 - p)$ for all $p \in [0, 1]$. At $p = 0$ the two sides are equal, however the right hand side decreases as p increases. Thus, if there is a single salvo k with an interceptor count $n_k^* > 1$, a more cost effective solution exists by splitting that salvo. This proves that if blue has sufficiently many salvo opportunities such that $M \geq N$, blue should optimally fire one interceptor per salvo.

4.2 Case $M = 2$

When blue only has time to fire 2 salvos $M = 2$, the only decision variable is the number of interceptors to fire in the first salvo, which we denote as n_1 . Blue fires the remaining $N - n_1$ interceptors in the second salvo, where N is defined in equation (4.3). The expected expenditure of a firing sequence can be written as a function of n_1 :

$$E[\text{expenditure}] = n_1 + (1 - p)^{n_1} \cdot (N - n_1) \quad (4.9)$$

n_1 interceptors are always fired in the first salvo; the remaining $N - n_1$ interceptors are fired in the second salvo only if all n_1 interceptors miss in the first salvo. If $n_1 = 0$ or $n_1 = N$ the expenditure equals N . It is trivial to check all integers between 1 and $N - 1$ to determine n_1^* . For the remainder of this section, we focus on a continuous approximation to equation

(4.9), where we denote x as the decision variable instead of n_1 . This allows us to solve the problem without the integer constraint. The following Proposition provides the optimal solution x^* to the continuous approximation. This proposition also reveals that the relative allocation fired in the first salvo is independent of p and strictly decreases with Q .

Proposition: $\frac{n_1}{N} \perp p$ (continuous case: $\frac{x}{\tilde{N}} \perp p$)

$$x^* = \frac{W\left(\frac{-e}{Q-1}\right) + \log(1-Q) - 1}{\log(1-p)}$$

$$\frac{x^*}{\tilde{N}} = \frac{W\left(\frac{-e}{Q-1}\right) + \log(1-Q) - 1}{\log(1-Q)}, \text{ where } \tilde{N} = \frac{\log(1-Q)}{\log(1-p)}$$

Proof: When blue fires $\tilde{N}$ interceptors (which may be a fractional quantity), the PRA is exactly Q . Next we assume x is the continuous quantity for n_1 . For any $x \in (0, \tilde{N})$, the expected expenditure is strictly less than $\tilde{N}$, and the continuous cost function is:

$$g(x) = x + (\tilde{N} - x)(1-p)^x \quad (4.10)$$

This function is convex over $[0, \tilde{N}]$ and hence has a minimizer x^* in that interval. To show convexity, we examine the second derivative:

$$g''(x) = -(1-p)^x \log(1-p) \left((x - \tilde{N}) \log(1-p) + 2 \right) \quad (4.11)$$

$g''(x)$ is non-negative, because $(x - \tilde{N}) \ln(1-p)$ is non-negative, making $g(x)$ convex. To solve for x^* , we examine the first derivative:

$$g'(x) = 1 + (1-p)^x \left((\tilde{N} - x) \log(1-p) - 1 \right) \quad (4.12)$$

We set the derivative equal to zero and solve for x^* , Substituting $\frac{\log(1-Q)}{\log(1-p)}$ for $\tilde{N}$ yields:

$$1 - (1 - p)^x \left(\left(\frac{\log(1 - Q)}{\log(1 - p)} - x \right) \log(1 - p) - 1 \right) = 0 \quad (4.13)$$

From here we rearrange so that we solve for x . This is a non-trivial transformation and Wolfram (2021) was used to find the final solution. The result utilizes a Lambert function $W(z)$, which for a given real number z satisfies the following relationship; $z = W(z) \exp(W(z))$. The optimal number of interceptors x^* fired in the first salvo for a given p and Q is given by:

$$x^* = \frac{W\left(\frac{-e}{Q-1}\right) + \log(1 - Q) - 1}{\log(1 - p)} \quad (4.14)$$

Finally, combining $\tilde{N} = \frac{\log(1 - Q)}{\log(1 - p)}$ and equation (4.14), we divide $x^*/\tilde{N}$ to determine the fraction of total interceptors fired in the first salvo.

$$\frac{x^*}{\tilde{N}} = \frac{W\left(\frac{-e}{Q-1}\right) + \log(1 - Q) - 1}{\log(1 - Q)} \quad (4.15)$$

End Proof

Equation (4.15) reveals the surprising result that the fraction of interceptors fired in the first salvo is independent of p . Intuitively one might think that a higher p would push a larger fraction of missiles to the second salvo in order to hopefully avoid using them. While p certainly dictates the total number of interceptors fired $\tilde{N}$ and the optimal number in the first salvo, it does not affect the fraction $x^*/\tilde{N}$.

In Figures 4.1 - 4.3 we explore the relationship between $\tilde{N}$, x^* , p , and Q . We vary $p \in \{0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$ and $Q \in \{0.8, 0.9, 0.96, 0.996\}$ to generate the resultant x^* , $\tilde{N}$, and $x^*/\tilde{N}$. Various combinations of these results and inputs are displayed; however, these calculations are for the continuous approximation. Figure 4.1 shows the relationships of how $x^*/\tilde{N}$ asymptotically approaches 0.5 and 0 as Q approaches 0 and 1 respectively. Thus, blue always fires fewer interceptors in the first salvo than in the second salvo, and this fraction in

first salvo decreases as Q increases. Figures 4.2 and 4.3 confirm intuition that the number of interceptors fired in first salvo grows as Q increases and p decreases. These Figures show how p and Q impact $x^*/\tilde{N}$ relative to each other.




Figure 4.1. Optimal fraction of number of interceptors fired in first salvo x^* over total interceptors $\tilde{N}$ versus PRA threshold Q .




Figure 4.2. Optimal number of interceptors fired in first salvo x^* versus PRA threshold Q for various values of SSPK p .




Figure 4.3. Optimal number of interceptors fired in first salvo x^* versus SSPK p for various values of the PRA threshold Q .

4.3 Case $2 < M < N$

We now consider the case where we have M salvo opportunities but need to fire $N > M$ interceptors to surpass Q . The number of interceptors in a salvo is denoted by n_i with i being the salvo index. $n_i > 0$ for all $1 \leq i \leq M$ using similar logic to the $M \geq N$ case. In Section 4.4 we prove that the salvo size is monotonically non-decreasing in the salvo index for the integer case. That is, $n_j \leq n_{j+1}$.

This section focuses on the continuous approximation case, so we use x_i and $\tilde{N}$ to represent the continuous approximation. The main result of this section is that p has no effect on the fraction $\tilde{x}_i^*/\tilde{N}$ for any i . Allocation instead depends on Q and M . We have a conjecture for the solution to the continuous approximation in Section 4.4. However, prior to this derivation, we solve the following optimization problem using Excel's generalized reduced gradient algorithm provided by FrontlineSolvers (2011):

$$\begin{aligned} \min & \left(\sum_{i=1}^M (1-p)^{\sum_{j=1}^{i-1} x_j} \cdot x_i \right) \\ \text{s.t.} & \sum_{i=1}^M x_i \geq N \end{aligned} \quad (4.16)$$

Solving this for $p \in \{0.8, 0.5\}$ and $Q = 0.996$ yields Table 4.1:

Table 4.1. Fraction of interceptors $x_i^*/\tilde{N}$ for $Q = 0.996$ and $p \in \{0.8, 0.5\}$. Note that both p 's have the same results.

	p = 0.8					p = 0.5				
	$\frac{x_{M-4}^*}{\tilde{N}}$	$\frac{x_{M-3}^*}{\tilde{N}}$	$\frac{x_{M-2}^*}{\tilde{N}}$	$\frac{x_{M-1}^*}{\tilde{N}}$	$\frac{x_M^*}{\tilde{N}}$	$\frac{x_{M-4}^*}{\tilde{N}}$	$\frac{x_{M-3}^*}{\tilde{N}}$	$\frac{x_{M-2}^*}{\tilde{N}}$	$\frac{x_{M-1}^*}{\tilde{N}}$	$\frac{x_M^*}{\tilde{N}}$
5 salvos	0.084	0.107	0.146	0.223	0.440	0.084	0.107	0.146	0.223	0.440
4 salvos		0.111	0.152	0.239	0.498		0.111	0.152	0.239	0.498
3 salvos			0.161	0.260	0.579			0.161	0.260	0.579
2 salvos				0.289	0.711				0.289	0.711

Note that both p values in Table 4.1 produce the same results, providing numerical experimentation to support the two-salvo conjecture about independence of p on the optimal relative allocation $x_j^*/\tilde{N}$. Keep in mind this is the continuous approximation and if we did

enforce integer constraints, p would have an indirect effect on the relative allocation. This is because it affects the number of interceptors in the model, and fewer interceptors makes it harder to reach the optimal fractions.

Optimal Solution to Continuous Approximation

Below we provide a formula for the solution to the optimization problem in Equation (4.16). We developed this methodology after completing the work in Section 4.3 and did not have time to rigorously prove this solution. It should however provide the same results as the GRG-Nonlinear method.

Conjecture: The optimal solution x_j^* for the continuous approximation problem satisfies the following recursive relationship for $j > 1$:

$$x_k = \frac{-((1-p)^{-x_{k-1}} - 1)}{\log(1-p)} \quad (4.17)$$

Since we do not know the first salvo size x_1 , we guess an initial x_1 , compute the rest of x_j using recursion in equation (4.17). If the resulting sequences has $\sum x_j = N$, we have found the optimal sequence. If it does not, we adjust x_1 and repeat until the sum equals N .

Related, if we define $\tilde{z}_k = x_k/\tilde{N}$ then:

$$\tilde{z}_k = \frac{-((1-Q)^{-\tilde{z}_{k-1}} - 1)}{\log(1-Q)} \quad (4.18)$$

Analysis:

For all result and graphics in this section, we look at salvos in reverse order with $i = 0$ being the very last salvo in a sequence while $i = 1$ is the second to last salvo. The results from Table 4.1 illustrates that salvos should be larger than their preceding salvo, with the first salvo being the smallest and the final salvo the largest. This monotonic result is proven with mathematical rigor in Section 4.4 and implies that salvos are the same or larger than preceding salvos. Intuitively, it allows blue the opportunity to limit the number of missiles fired by putting a larger allocation in later salvos. Figure 4.4 graphically demonstrates this with the last salvo $i = 0$ having largest allocation. However, as M increases the fraction for

a given i decreases.




Figure 4.4. Optimal fraction for allocation of interceptors for the last 5 salvos for salvos of sizes $M = [2,3,4,5]$. $Q = 0.996$.

In the more general version of our model with multiple interceptor types and range dependent SSPK, it is not guaranteed that the solution will have this structure with more interceptors fired in later salvos. However, we expect that in many scenarios for the general case, we should observe this general structure of salvo size increasing with salvo index.

Factoring the combined effects of M and Q on the allocation fractions, Figure 4.5 shows that the fraction of interceptors in the final salvo increases as Q increases, regardless of M . Conversely, Figure 4.6 shows that the final fraction decreases as M increases, and this make intuitive sense as there are more opportunities to disperse interceptors between salvo.




Figure 4.5. Optimal fraction for the allocation of interceptors for the last salvos based on Q for salvos of sizes $M \in \{2, 3, 4, 5\}$




Figure 4.6. Optimal fraction for the allocation of interceptors for the last salvos based on M for $Q \in \{0.8, 0.9, 0.96, 0.996\}$

4.4 Derived Results

In this section, we present three results and conjectures that have operational relevance. This is for the integer case n_i as opposed to the continuous approximation x_i . The first supports the idea that optimal sequences should be monotonic and non-decreasing in interceptor count, the second deals with simplifying the ability to generate optimal salvos if the interception probability p is above a certain threshold, and the third is a method to calculate n_i .

4.4.1 Monotonically, Non-Decreasing Sequence

For the single interceptor case with SSPK independent of range, we prove that salvo sizes n_i are monotonically non-decreasing in accordance with the salvo index i .

Theorem: $n_i \geq n_{i-1}$

Proof: The logic in this section follows the proof for the optimal firing sequence for the $M \geq N$ case in Section 4.1. If we denote n_s^* as the optimal number of interceptors in salvo s , then $n_k^* \geq n_j^*$ for $k > j$. We proceed by contradiction: assume there exists some k such that $n_k^* = x > n_{k+1}^* = y$. The cost of this firing sequence is:

$$C^* = \sum_{j=1}^{k-1} n_j^* (1-p)^{\sum_{i=1}^{j-1} n_i^*} + x(1-p)^{\sum_{i=1}^{k-1} n_i^*} + y(1-p)^{x + \sum_{i=1}^{k-1} n_i^*} + \sum_{j=k+2}^{M^*} n_j^* (1-p)^{\sum_{i=1}^{j-1} n_i^*} \quad (4.19)$$

The first term is the cost prior for $j < k$, the last term is for $j > k + 1$, and the middle terms are respectively the cost for k and $k + 1$. We next show that if we move one interceptor from salvo k to salvo $k + 1$ we will have a new firing sequence with an expected cost smaller than equation (4.19). Specifically let us define a new firing sequence $\hat{n}$ that is equivalent to n^* , except for salvos k and $k + 1$:

$$\begin{aligned}
\hat{n}_j &= n_j^* \quad \text{for } j < k \\
\hat{n}_k &= n_k^* - 1 = x - 1 \\
\hat{n}_k + 1 &= n_{k+1}^* + 1 = y + 1 \\
\hat{n}_j &= n_j^* \quad \text{for } j > k + 1
\end{aligned}$$

$\hat{C}$ is the cost using sequence $\hat{n}$ and is identical to equation (4.19) except the usage of $\wedge$ instead of $*$. Breaking the cost equation down for $\hat{C}$ based on k in a similar fashion to equation (4.19) yields:

$$\hat{C} = \sum_{j=1}^{k-1} \hat{n}_j (1-p)^{\sum_{i=1}^{j-1} \hat{n}_i} + (x-1)(1-p)^{\sum_{i=1}^{k-1} \hat{n}_i} + (y+1)(1-p)^{(x-1) + \sum_{i=1}^{k-1} \hat{n}_i} + \sum_{j=k+2}^{M^*} \hat{n}_j (1-p)^{\sum_{i=1}^{j-1} \hat{n}_i} \quad (4.20)$$

If we compare the two completed equations (4.19) and (4.20), accounting for relationship between $\hat{n}$ and n^* we see that the first portion where $j < k$ is the same for $\hat{C}$ and C^* and the last portion where $j > k + 1$ is the same for C^* and $\hat{C}$. Thus, we can cancel out these portions and only compare the middle sections according to:

$$x \cdot (1-p)^{\sum_{i=1}^{k-1} n_i^*} + y \cdot (1-p)^{x + \sum_{i=1}^{k-1} n_i^*} \quad \text{vs} \quad (x-1) \cdot (1-p)^{\sum_{i=1}^{k-1} n_i^*} + (y+1) \cdot (1-p)^{(x-1) + \sum_{i=1}^{k-1} n_i^*} \quad (4.21)$$

Factoring out the $(1-p)^{\sum_{i=1}^{k-1} n_i^*}$ yields:

$$x + y \cdot (1-p)^x \quad \text{vs} \quad (x-1) + (y+1) \cdot (1-p)^{x-1} \quad (4.22)$$

Canceling x 's from both sides:

$$y \cdot (1 - p)^x \quad \text{vs} \quad -1 + (y + 1) \cdot (1 - p)^{x-1} \quad (4.23)$$

To complete the proof we must show the left hand side is greater than the right hand side, or equivalently:

$$1 > (1 - p)^{x-1} \cdot (yp + 1)$$

Focusing on the right hand side we have:

$$(1 - p)^{x-1} \cdot (yp + 1) < (1 - p)^{x-1} \cdot (p + 1)^y \quad (4.24)$$

Where $(yp + 1) < (1 + p)^y$ follows immediately from the binomial expansion. The inequality is strict because we assume $p > 0$ and $y > 0$ by similar logic to Section 4.1. The right hand side of the above decreases with x and since $x > y$, we have:

$$(1 - p)^{x-1} \cdot (yp + 1) < (1 - p)^{x-1} \cdot (p + 1)^y \leq (1 - p)^{y+1-1} \cdot (p + 1)^y \quad (4.25)$$

To complete the proof we have:

$$(1 - p)^{y+1-1} \cdot (p + 1)^y = [(1 - p)(p + 1)]^y = [1 - p^2]^y < 1 \quad (4.26)$$

The implication of this conjecture is that it provides context to the structure of a firing sequence so that we can more quickly design the optimal allocation of interceptors between multiple sequences. Tactically, this conjecture provides insight that the optimal sequence attempt to “save interceptors” by firing more in later salvos.

This result was independently derived by Kalyanam and Clarkson (2021) through a Markov decision process, however we believe our proof to be more concise.

4.4.2 M-1 Single Shot Salvos According to p^*

For a given sequence with M firing opportunities, there exists a p^* such that for any $p > p^*$ the first $M - 1$ salvos contains a single interceptor and the last salvo contains the remaining $N - (M - 1)$ interceptors. Although we have preliminary results, we did not have time to rigorously prove this conjecture.

Conjecture: For $p > p^*$, $n_s^* = 1$ for $s < M$ and $n_M^* = N - (M - 1)$.

$$p^* = \inf \left\{ p : p > \frac{N(p, Q) - M - 1}{N(p, Q) - M} \right\} \quad (4.27)$$

Since N decreases as p increases, we have to find the point at which the optimal sequence switches from having two interceptors in the second-to-last salvo to having only a single interceptor. We iterative do this and find the maximum p where equation (4.27) is true.

Note that $N(p, Q)$ is in reference to function (4.3). p^* is dependent on Q and M . Any sequence with the same Q and M and $p > p^*$ should be allocated such that $M - 1$ salvos are filled (assuming $N > M$) with a single interceptor and the remaining $N - M$ are allocated to the last salvo of the sequence.

The results of equation (4.27) are plotted in Figures 4.7 - 4.9. Note that since this is the integer case, rounding up causes the points to be more distinct. Figure 4.7 and 4.8 both demonstrate that p^* increases with a rise in Q but decreases as M increases. The increase with Q is due to Q indirectly increasing N , whereas the decrease with M occurs as the salvos can be more efficient with their expenditure even as Q rises. Figure 4.9 shows us that as Q increases, the number of interceptors in the final salvo, the only one with more than one interceptor, increases.




Figure 4.7. For a given M , the resulting $p = p^*$ according to equation (4.27) is displayed. This is plotted across varying levels of Q .




Figure 4.8. For a given Q , the resulting $p = p^*$ according to equation (4.27) is displayed. This is plotted across varying levels of M .




Figure 4.9. For a given $N - M$, the number of interceptors in the last salvo, the resulting $p = p^*$ according to equation (4.27) is displayed. This is plotted across varying levels of Q . The small difference between $Q = 0.8$ and 0.9 is due to rounding.

4.4.3 Solving x_i for a given M and N

The solutions are nested as N increases for fixed M . That is, if n_i is solution for N , and $\tilde{n}_i$ is solution for $N + 1$, then $\tilde{n}_i$ is equal to n_i except for one $\tilde{n}_j$ where it is one greater.

Conjecture: A nested solution exist for solving n_i iteratively.

$$\begin{aligned}\tilde{n}_i &= n_i \\ \tilde{n}_j &= n_i + 1 \quad \text{for exactly one } j\end{aligned}$$

Generating the solution is now just an iterative $O(NM)$ algorithm. We know solution for $N = M$ and $N = M + 1$. We work our way up for $M + 2, M + 3, M + 4, \dots$ until we reach the desired N . We check each salvo i to find the allotment that generates the smallest cost increase. At that salvo i , we set $n_i + 1$ and have our new solution. Since we know monotonicity must be maintained, we only need to check i where this condition holds.

4.5 Comparison of SLS Policies

We conclude our single-interceptor analysis by examining the performance of specific firing policies as a function of p . We look at $p \in \{0.3, 0.5, 0.8\}$ and various SLS firing policies up to 6 interceptors for the two-salvo and three-salvo case. For each firing policy, we compute the PRA and expected expenditure which are provided in Table 4.2.

Table 4.2. PRA and E[expenditure] given $p \in \{0.3, 0.5, 0.8\}$ for various two and three salvo firing policies.

		p= 0.3		p = 0.5		p = 0.8	
		PRA	E[expend]	PRA	E[expend]	PRA	E[expend]
1	SLS	0.510	1.70	0.750	1.50	0.960	1.20
2	SSL	0.657	2.49	0.875	2.25	0.992	2.04
3	SLSS	0.657	2.40	0.875	2.00	0.992	1.40
4	SSSL	0.760	3.34	0.938	3.13	0.998	3.01
5	SSSLSS	0.832	3.69	0.969	3.25	0.9997	3.02
6	SSSLSSS	0.882	4.03	0.984	3.38	0.9999	3.02
7	SLSSS	0.760	3.1	0.938	2.5	0.998	1.6
8	SSLSSS	0.832	3.47	0.969	2.75	0.999	2.12
9	SLSLS	0.657	2.19	0.875	1.75	0.992	1.24
10	SSLSLS	0.760	2.83	0.938	2.38	0.998	2.05
11	SLSSL	0.760	2.74	0.938	2.13	0.998	1.41
12	SLSSLSS	0.760	2.68	0.938	2.00	0.998	1.28
13	SSLSSL	0.832	3.22	0.969	2.56	0.9997	2.08
14	SLSSLSS	0.832	3.09	0.969	2.25	0.9997	1.42
15	SSLSSLSS	0.832	3.18	0.969	2.50	0.9997	2.06
16	SSLSSLSSS	0.882	3.46	0.984	2.63	0.9999	2.08

This table provides context into salvos that are underkill and overkill for a given p . If our goal is $Q = 0.996$, only $p = 0.8$ can surpass this PRA within six interceptors (it requires four). Of the four policies with four interceptors, SLSLSS provides the lowest cost and this supports the findings that salvo sizes should be strictly increasing in the optimal case.

Plotting PRA vs E[cost], we see trends dependent on p . Note that the numbers in the first column of Table 4.2 correspond to the numeric labels in the plots. Color in the plot refers to the number of firing opportunities presented, with black assigned to $M = 2$ while red is assigned to $M = 3$. We see that as p increases, PRA vs E[cost] changes from a linear relationship to some other concave one. Between the different p , the starting PRA is higher

as p increases and this makes intuitive sense. Subsequently, the gap between different levels of N decreases as p increases.

Breaking policies down by N between Figures 4.10 - 4.12, Figure 4.10 demonstrates that for lower p there is a roughly linear trade-off between increased PRA and higher $E[\text{cost}]$. Conversely, Figure 4.12 shows what appears to be diminishing returns as N increases. However, these changes are minimal compare to those seen in the smaller p .

Within all the figures, we see distinct levels at different PRAs where N is equivalent. An example is that at the $N = 4$ level, the PRA for SSLSS, SSSLs, SLSSS, SSLSLS, SLSSLS, and SLSSLSS are all the same. For a given PRA/ N level, the optimal policy involves minimizing the earlier salvo while the least optimal has interceptors fired in earlier salvo. The best policies for their respective N are consistently SLSLS (9), SLSSLSS (12) and SLSSLS (14), which maximize the number of salvos for a given N while saving larger salvos towards the end. Salvos SSSLs (4), SSSLSS (5), and SSSLSSS(6) are consistently the worse as they are limited to two salvos and fire the larger salvo first.

Looking at red (three salvos) versus black (two-salvos) in Figures 4.10 - 4.12, we see that three salvo sequences always outperform two-salvo sequences for a given PRA. However, three salvos will not always be feasible due to ranging requirements. Of the two-salvo sequences, SLSS (3), SLSSS (7), and SSLSSS (8) perform well at their respective level, and this performance is due to firing larger salvos in the second salvo of the sequence.




Figure 4.10. Expected cost and the respective PRA for various firing policies given an interceptor with $p = 0.3$. Labeled numbers are policies referenced in Table 4.2. Black denotes double salvo sequences while red denotes triple salvo sequences.




Figure 4.11. Expected cost and the respective PRA for various firing policies given an interceptor with $p = 0.5$. Labeled numbers are policies referenced in Table 4.2. Black denotes double salvo sequences while red denotes triple salvo sequences.




Figure 4.12. Expected cost and the respective PRA for various firing policies given an interceptor with $p = 0.8$. Labeled numbers are policies referenced in Table 4.2. The PRA range is significantly narrower than the previous graphs. Black denotes double salvo sequences while red denotes triple salvo sequences. Additionally, point 1 is excluded from this graph.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 5: Case Studies

We derive multiple insightful results in Chapter 4. However, Chapter 4 revolves around the simplifying assumption of a single interceptor type with constant single shot probability of kill regardless of range. When these two assumptions are removed, deriving any analytic results about the optimal sequence likely becomes impossible due to the sheer number of factors influencing optimality. These factors include minimum and maximum range of each interceptor type, *SSPK* of each interceptor type, and tracking the range of impact for subsequent salvos. These factors also interact, further complicating analysis.

To inspect this space, we utilize the algorithm laid out in Chapter 3 that provides a heuristic approximation of the optimal sequence that incorporates all the aforementioned factors. From this model, we ultimately choose the top ten lowest-cost sequences that satisfy the operational constraint Q . The reason is that we want to provide options for the decision maker but not to overwhelm them. An additional reason for multiple sequences is that when inventory constraints are considered, some sequences may become infeasible logistically.

In this chapter we are intercepting an incoming C802 missile with a velocity of 0.8 mach (0.274 km/s). We first look at the baseline case in Section 5.1 with four interceptors and a constant entry-range, and explore the effect of adjusting Q and the length of time to complete a BDA, t_{BDA} , has on the top ten lowest cost. The parameters for these interceptors are provided in Table 5.1 and Figure 5.1. Section 5.2 looks at dropping each of the four interceptors one at a time and exploring the effects on cost. Finally, Section 5.4 looks at adding a fifth “high-performance” interceptor and explore the subsequent min-cost sequences.

5.1 Baseline

For the baseline results, we first explore the effects of t_{BDA} and Q on our model. The three BDA lengths we consider are $t_{BDA} \in \{30, 55, 80\}$ seconds. Immediately apparent when running the model for various t_{BDA} is the run time; lower t_{BDA} effectively increases the number of salvos (i.e., the size of M), creating more combinations of interceptors that would be valid. For reference, $t_{BDA} = 80$ s generates 2,637 sequences prior to applying our Q filter,

Table 5.1. Parameters of the four interceptors that can be used against a given C802 threat. Velocity is in km/s and ranges are in km. These values are provided by the sponsor of this thesis N-81 as well as acquired through open source research in Janes Database: C802 (2021).

Interceptor	v	PIP_{max}	PIP_{min}	LR_{max}	LR_{min}
SM-2	0.857	145.00	10.00	191.400	13.200
ESSM	1.372	36.00	9.99	43.200	11.998
RAMD	0.857	7.00	2.00	9.240	2.640
CWIS	1.029	3.99	1.00	5.065	1.266




Figure 5.1. Range dependent SSPK for four different interceptors

whereas $t_{BDA} = 30s$ generates 108,909 sequences.

By decreasing t_{BDA} , we're effectively increasing M as blue has the capability to fire more salvos. Greater M results in more ways to combine salvos, providing more sequence combinations that may generate lower costs as demonstrated in Figure 5.2. Assuming $Q > \max(SSPK)$ such that $N > 1$, the lower bound of $E[\text{cost}]$ for our interceptors is going to be $2 - \max(SSPK) = 2 - 0.9 = 1.1$. This is because we must always fire an interceptor in the first salvo, and if we fire the best interceptor in that first salvo, the probability of firing additional interceptors is $1 - \max(SSPK)$.




Figure 5.2. Best, worst, and average E[cost] across the top ten cases for $t_{BDA} \in \{30, 55, 80\}$

Setting $Q = 0.96$, we can look at the ten best sequences for each t_{BDA} in Table 5.2. Notice the best sequences for each t_{BDA} has the same E[cost] but PRA is decreasing when it should be the same or improved. Even though $t_{BDA} = 30$ should have the same or better PRA than $t_{BDA} = 55$, our assumption about firing immediately does not provide a perfectly optimal answer and is a heuristic approximation. If we introduced delays, we would be able to get closer to the optimal PRA for a given sequence.

Also of note in Table 5.2 is for the tenth best sequence, the E[cost] decreases with t_{BDA} , supporting Figures 5.2 and 5.3. Still looking at the tenth best sequences for all t_{BDA} , $t_{BDA} = 30$ stands out. Despite having two additional expenditures than both $t_{BDA} = 55$ and $t_{BDA} = 80$, the tenth best sequence in $t_{BDA} = 30$ is more efficient with its expenditure (compared to $t_{BDA} = 55$ and 80) as it is using two salvos of ESSMs; double utilizing the highest SSPK interceptor. Since it has a reduced BDA delay, $t_{BDA} = 30$ can launch four different salvos despite only launching the first salvo at 43.2 km as opposed to the SM-2's maximum LR of 191.4 km.

Table 5.2. Best and tenth best sequences and their respective E[cost], PRA, and looks for each t_{BDA} at $Q = 0.96$.

t_{BDA}		Sequence	E[Cost]	PRA	salvos
80	best	[ESSM, look, ESSM]	1.200	0.96934	2
80	tenth best	[ESSM, look, RAMD, RAMD]	1.4	0.99800	2
55	best	[ESSM, look, ESSM]	1.200	0.96714	2
55	tenth best	[SM-2, look, ESSM, look, CWIS]	1.360	0.97000	3
30	best	[ESSM, look, ESSM]	1.200	0.96494	2
30	tenth best	[ESSM, look, ESSM, look, ESSM, look, CIWS, CIWS]	1.246	0.99864	4

The other parameter to adjust is Q , which is the minimum required PRA and serves as a restriction on the problem. As Q is increased, more sequences will be filtered out and E[cost] will gradually increase as seen in Figure 5.3. As $Q \rightarrow 1$, the maximum number of salvo opportunities M that can be generated remains the same requiring interceptors to eventually be placed in the earliest salvo. Figure 5.3 also demonstrate parameters sensitivity to each other and our model is capturing the complexity of these relationships.

To force sequences containing more than 3 interceptors, we set $Q = 0.9996$ and get Table 5.3. Each sequence is broken down by salvo, and within each salvo we see the PIP range of the salvo as well as the SSPK of each interceptor. Immediately fascinating is the diversity of sequences as our algorithm is successfully exploring all the ways to develop sequences. Sequences with higher t_{BDA} are relying more heavily on the SM-2, whereas shorter t_{BDA} make use of the shorter range but higher SSPK ESSM.




Figure 5.3. Minimum $E[\text{cost}]$ for the best and tenth best sequence across Q . As Q increases, $E[\text{cost}]$ begins to converge regardless of t_{BDA} .

Since all interceptors have $\text{cost} = 1$ for our initial exploration, the best firing sequence must contain all homogeneous salvos. Otherwise we could replace a lower SSPK interceptor in a salvo with a higher SSPK interceptor to generate a higher PRA with the same cost. The only time we see heterogeneous salvos is in the tenth best sequences in Table 5.3, which occurs twice. These heterogeneous expenditures occur in the last salvo and only exist because their homogeneous alternatives are ranked higher. In the case of $t_{BDA} = 80$, the following sequences surpass Q and provide either a better $E[\text{cost}]$ or an improved PRA:

[SM-2, look, SM-2, look, SM-2, look, ESSM, look, RAMD, RAMD]

[SM-2, look, SM-2, look, SM-2, look, ESSM, look, RAMD, RAMD, RAMD]

Table 5.3. Best and tenth best sequences and their respective PRA for each t_{BDA} at $Q = 0.9996$. Each sequence is broken down into its' individual salvos, with each salvo's PIP (top number in cell) included as well as the SSPK of each interceptor (bottom number).

$t_{BDA} = 80$, best					PRA: 0.99980	E[cost]: 1.4428	Salvos: 4
191.4 SM-2 0.7		123 SM-2 0.66		43.2 ESSM 0.8		9.2 RAMD 0.9	RAMD 0.9
$t_{BDA} = 80$, tenth best					PRA: 0.99981	E[cost]: 1.4675	Salvos: 5
191.4 SM-2 0.7		123.0 SM-2 0.66	71.3 SM-2 0.5846	32.0 ESSM 0.8179		4.7 RAMD 0.9	CIWS 0.5
$t_{BDA} = 55$, best					PRA: 0.99980	E[cost]: 1.4248	Salvos: 4
191.4 SM-2 0.7		43.2 ESSM 0.8		20.9 ESSM 0.8357	ESSM 0.8357	2.3 CIWS 0.5	CIWS 0.5
$t_{BDA} = 55$, tenth best					PRA: 0.99961	E[cost]: 1.4542	Salvos: 5
191.4 SM-2 0.7		129.9 SM-2 0.67	83.3 SM-2 0.6022	43.2 ESSM 0.8		5.1 RAMD 0.9	CIWS 0.5
$t_{BDA} = 30$, best					PRA: 0.99995	E[cost]: 1.2459	Salvos: 4
43.2 ESSM 0.8		27.8 ESSM 0.8247		14.9 ESSM 0.8453		4.2 RAMD 0.9	RAMD 0.9
$t_{BDA} = 30$, tenth best					PRA: 0.99999	E[cost]: 1.2718	Salvos: 4
43.2 ESSM 0.8		27.8 ESSM 0.8247		14.9 ESSM 0.8453	ESSM 0.8453	4.2 RAMD 0.9	RAMD 0.9

5.2 Dropped Interceptor

Continuing with the interceptors in Table 5.1, we now look at four scenarios where each has one of the interceptors unavailable for sequence generation. This illustrates the impact of depleting one type of interceptor during a prolonged battle, and allows to explore expenditure of the remaining types and see if they provide enough capability and cost-effectiveness to successfully conduct IAMD. For all scenarios, the t_{BDA} is set as 55 seconds and the scenarios are:

- Scenario 0: Baseline, all interceptors available
- Scenario 1: No SM-2 available
- Scenario 2: No ESSM available
- Scenario 3: No RAMD available
- Scenario 4: No CIWS available

We set $Q = 0.95$ to filter out single interceptor sequences. The resulting best sequences for each scenario are displayed in Table 5.4. Some of the scenarios have two entries and this is because they have a tie in $E[\text{cost}]$. We left in ties to demonstrate how redundancy is lost as interceptor types are depleted. Note that each interceptor expended has a cost of “1” and are valued the same for this analysis.

Table 5.4. Top sequence(s) for each scenario given $Q = 0.95$. Given our interceptor parameters, this Q only requires $N = 2$. Note that while some scenarios have two best cost sequences, one of these sequences has a better PRA.

Scenario	Sequence	E[Cost]	PRA	Salvos
0 (Baseline)	[ESSM, look, ESSM]	1.2	0.96714	2
0 (Baseline)	[ESSM, look, RAMD]	1.2	0.98000	2
1 (No SM-2)	[ESSM, look, ESSM]	1.2	0.96714	2
1 (No SM-2)	[ESSM, look, RAMD]	1.2	0.98000	2
2 (No ESSM)	[SM-2, look, RAMD]	1.3	0.97000	2
3 (No RAMD)	[ESSM, look, ESSM]	1.2	0.96714	2
4 (No CIWS)	[ESSM, look, ESSM]	1.2	0.96714	2
4 (No CIWS)	[ESSM, look, RAMD]	1.2	0.98000	2

Scenarios 0, 1, and 4 have the same optimal sequences and actually have redundancy in

that they have the option of two different sequences to fire, albeit one has marginally lower PRA. Scenario 3 has one optimal sequences and retains the best cost of 1.2, but its inability to use RAMD means that said redundancy is lost. Without the ESSM available, Scenario 2 instead utilizes the SM-2 which causes E[cost] to marginally increase. While the results of dropping an interceptor are fairly robust as all scenarios except scenario three are still able to produce the best E[cost] of 1.2, these scenarios show the versatility of ESSM. It has a high velocity and an intermediate range to increase the number of potential salvos, and it has the highest PK of our interceptors.

Setting $Q = 0.9996$, we are able to force more complicated and variable sequences that are shown in Table 5.5. Within these complex scenarios, we first look at the highest E[cost] that appears in scenario 1. Without the SM-2, blue's maximum range is degraded resulting in the number of salvos M being reduced. This is evident in the salvos column, as all other scenarios have four or five looks as opposed to Scenario 1's utilization of two looks. This re-emphasizes that M is a major driver of reducing cost while maintaining PRA above a given Q , and longer range missiles play a role in reducing expenditure.

Table 5.5. Top sequence for each scenario given $Q = 0.9996$, and this higher Q forces a larger N .

Scenario	Sequence	E[Cost]	PRA	Salvo
0 (Baseline)	[SM-2, look, ESSM, look, ESSM, ESSM, look, CWIS, CWIS, CWIS]	1.4248	0.999798	4
1 (No SM-2)	[ESSM, look, RAMD, RAMD, RAMD, end]	1.6000	0.999800	2
2 (No ESSM)	[SM-2, look, SM-2, look, SM-2, look, SM-2, look, RAMD, RAMD]	1.4738	0.999823	5
3 (No RAMD)	[SM-2, look, ESSM, look, ESSM, ESSM, look, CWIS, CWIS, CWIS]	1.4248	0.999798	4
4 (No CIWS)	[SM-2, look, SM-2, look, ESSM, look, RAMD, RAMD]	1.4386	0.999802	4

Having the same best sequence for both Scenarios 0 and 3 initially suggests that RAMD is the least beneficial interceptor, however RAMD is utilized in all other scenarios as well as throughout Table 5.4. This demonstrates how the design space in the multi-interceptor model is significantly more nuanced as different interceptors have individual capability

that can be suited to unique roles. Some interceptors tend to be more valuable but each additional interceptor added to the model provides increased number of opportunities for effective sequences, which may include lower E[*cost*] sequences.

5.3 Dominant Interceptor

We now model an interceptor that dominates the other interceptors in Table 5.1 across several categories. This allows us to demonstrate the flexibility of adding interceptors to our model, explore the effects that an interceptor with superior parameters have on a less capable interceptor, and explore the interactions with other interceptors. We call this new interceptor the SM+ and its parameters are provided in Table 5.6 and Figure 5.4. The parameters are designed such that it has slightly more range than the existing longest (SM-2), as fast as the fastest interceptor (ESSM), minimum range of the closest interceptor (CWIS) and a Range-SSPK that is better than the SM-2 and CIWS. We set $Q = 0.95$ to remove single interceptor salvos as this is above the $\max(SSPK)$ of our interceptors.

Table 5.6. Parameters of the fictional superior interceptor SM+ that can be used against a red C802 threat. Velocity is in km/s and ranges are in km. Values and parameters are made up for illustration purposes.

Interceptor	v	PIP_{max}	PIP_{min}	LR_{max}	LR_{min}
SM+	1.372	175.00	1.00	191.400	1.20

The first item of note is that 775167 sequences were generated. This is because the improved velocity in conjunction with improved PIP_{min} and PIP_{max} has a drastic effect on LR_{min} and LR_{max} . On top of the expanded feasible LR interval, we discussed in Section 3.6.3 how one interceptor with a high velocity and a small minimum range can cause M to increase drastically, resulting in an exponential expansion in the number of sequences. However, this appears to be partially mitigated with the $t_{BDA} = 55$ seconds as the maximum M is equal to seven.

The effects of SM+ depend on the Q selected. With $Q > 0.96$ such that we use at least two interceptors per sequence, the top four sequences utilize the ESSM followed by various non-SM+ interceptors and the first SM+ is not seen until sequence 5. This is because the




Figure 5.4. Range dependent probability of kill for five different interceptors; four original plus the SM+ theoretical interceptor.

ESSM and RAMD have higher SSPKs than the SM+, meaning the SM+ is not totally dominate, otherwise SM+ results would not be interesting). When Q is lower, this decreases this value of the SM+. Only when Q is increased and M subsequently rises does the SM+ become valuable. Sequences 5-9 all have $E[\text{cost}]$ of 1.2905 and feature two single SM+ salvos, with the third salvo alternating between the five different interceptors. Of those five sequence, the lowest PRA is 0.991895 for:

[SM+, look, SM+, look, ESSM]

As Q increases the sequences begin to be overwhelmingly dominated by SM+ and ESSM interceptors. For $Q > 0.9996$, the top three sequences consist of the following sequences with $E[\text{cost}]$ being 1.3026, 1.3030, and 1.3030:

[SM+, look, SM+, look, SM+, look, SM+, look, ESSM, look, ESSM]

[SM+, look, SM+, look, SM+, look, SM+, look, SM+, look, ESSM]

[SM+, look, SM+, look, SM+, look, SM+, look, SM+, look, SM+]

In Table 5.5, we saw that the sequences were utilizing SM-2 for longer shots, as well as shots in between the feasible ranges of the RAMD and ESSM. The under-reliance on the SM-2 after the SM+ was introduced demonstrates that interceptor can dominate each other with significant consequences on the optimal sequences.

5.4 Monotonicity in the Heterogeneous Case

Continuing with the SM+ case, the larger value of M presented by introducing the SM+ allows us to explore how interceptors are distributed among the salvos as N increases. Although not as mathematically rigorous, this allows us to see if our findings about monotonicity in Section 4.4 applies to the heterogeneous interceptor case as it has tactical implications on how sequences should be structured.

Expanding our scope beyond the top ten sequences, we switch from minimizing $E[\text{Cost}]$ to PRA, and limit our scope to a certain number of salvos and total interceptors. We set $M = 7$ and total interceptor equal to $M + 2 = 9$ as we want to have a maximum number of interceptor salvo in the final opportunity. Filtering out all interceptor counts below 9, we minimize PRA and obtain the following sequence with $E[\text{cost}]$ of 1.3030:

[SM+, look, SM+, look, SM+, look, SM+, look, ESSM, look, ESSM, look, SM+, SM+, SM+]

We next increase the interceptor count by one which results in the following minimum PRA sequence with $E[\text{cost}]$ of 1.3048:

[SM+, look, SM+, look, SM+, look, SM+, look, ESSM, look, ESSM, ESSM, look, SM+, SM+, SM+]

The best sequences for interceptors counts nine and ten are both monotonic, suggesting monotonicity holds in the heterogeneous case. There are certainly cases where it will not hold such as having an interceptor that 1) only works at extremely long ranges and 2) has a phenomenal *SSPK*. However, based on our theoretical and numerical results, it appears that monotonically increasing sequences usually occur in the heterogeneous case. This also illustrates the importance of maximizing the salvo opportunities M .

5.5 Performance

From a computational perspective, the run time and the number of sequences generated are connected. For our given parameters with four interceptors, the run-time remains manageable as all of the analysis was run on a personal computer with an Intel i7-4710HQ quad core CPU. However, more interceptors can quickly increase the computational burden

and this can be magnified with minor parameter changes. A good highlight of this is in the baseline scenario with variable t_{BDA} . As t_{BDA} decreases, Figures 5.5-5.6 demonstrate that the sequences generated and run time increases rapidly as t_{BDA} approaches 0. This is because M is effectively increasing, allowing for more combinations.




Figure 5.5. Number of sequences generated for various t_{BDA} with no sequences filtered out by Q in the baseline scenario.




Figure 5.6. Runtime in seconds for various t_{BDA} in the baseline scenario.

An additional illustration of the model run time sensitivity to parameters is evident in the

SM+ scenario as 775,167 sequences were generated. The resulting run-time was roughly 36 hours. Although the SM+ is especially bad as it has higher speed and a wider LR interval resulting in a large boost to M , it demonstrates that more interceptors in the mix drastically increases the number of sequences.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 6:

Conclusion

6.1 Practicality

This thesis has generated insights into how VLS salvos should be structured to better defend against incoming ASCM threats from a cost effectiveness standpoint. We first explored a simplified case which assumes homogeneous interceptors with constant SSPK. For this case we derived general provable properties. We showed that a policy for launching interceptors depends on the properties of the threat (velocity), acceptable level of risk, and DCR. With these factors in mind, the resulting sequence of launched salvos should be monotonically non-decreasing and should generally maximize the number of sequential salvos in a sequence before increasing the number of interceptors in a salvo.

We also built a model that accounts for multiple interceptor types while obtaining a minimum cost firing sequence that satisfies a given minimum PRA threshold Q . Within this model, we incorporated additional factors such as range dependent probability of kill and battle damage assessment. An end user inputs a characterization of the threat and DCR and the model returns a specified number of the lowest cost firing sequences that satisfy $PRA > Q$. All of this is performed in a matter of tens of seconds with an average performance computer.

This model can be used by analysts for a variety of cases. In the simplest case, we can preemptively derive a firing policy for a given threat scenario. For acquisition purposes, this model can be used to evaluate new interceptors — as well as upgrades of existing interceptors — and the subsequent improvements of effectiveness and efficiency on firing policy. Exploring the interaction between different interceptors is a complicated task, but our model provides a streamlined way to do it as adding or modifying interceptors is a simple task. A final use of this model is for it to be fed into a loadout planner to provide end-to-end optimization from logistics to firing.

6.2 Recommendations and Insights

This thesis illustrates that VLS firing policy should be threat, range, and inventory dependent. This insight could be instituted Navy wide or ships can be individually outfitted with the capabilities of our model so that individual inventory constraints can be considered.

Our results suggest that maximizing the number of salvo opportunities M is the critical driver for reducing the expected cost of IAMD while maintaining a certain risk level defined by the relationship between Q and PRA. Maximizing the number of salvos can be done through early detection, that is, improving DCR, as well as having interceptors with appropriate ranges. It can also be done by improving the BDA response time and having high-velocity, short-range interceptors.

After maximizing the number of salvos, next relevant task is to structure sequences such that they are monotonically increasing. This was mathematically proven for the case with homogeneous interceptors with constant SSPK, and exploration of our model results suggest that the concept of monotonic salvos also applies to the heterogeneous cases. This implies that firing policies such as SSLS or SSSLs are sub-optimal except in very specific circumstances such as when the first salvo has a much greater SSPK.

The importance of the type of an interceptor varies on parameters such as t_{BDA} and DCR due to second-order interactions of those inputs, however engagement should generally utilize multiple interceptor types. Long range interceptors, especially if they have a high velocity, play an especially important through their ability to increase the number of salvo opportunities. Interceptors can dominate other interceptors such that the dominated interceptor is worse in all situations, however we did not explore the effects of inventory and cost on the dominated interceptor.

6.3 Future Work

6.3.1 Model Extensions

Soft-kill Measures

The next step for our model would be introducing soft-kill interceptors. Our model already introduces the initial first step towards this extension in that weapons can have imperfect

BDA. For soft-kill, we believe false negatives would be the primary concern. A false negative occurs when blue believes it has not killed the incoming threat despite already having a successful interception. False negatives can lead to wasteful expenditure of interceptors. False-positive would have serious implications for firing policy but such assessment is very unlikely; a missile traveling at mach 0.5 towards a ship is hard to mistake as being killed.

The difficulty with soft kill is analysis is due to it being such a wide area with each measure having drastically different mechanisms of successfully scoring a “kill”. Some of these methods include but are not limited to jamming, directed energy, chaff and decoys. The first three provide an instantaneous effect when “deployed”. Chaff is a one-time expenditure but can be applied to multiple threats. While jamming and directed energy “impact” instantaneously, they require a certain amount of time on target, after which jamming leads to a soft-kill while energy weapons are more akin to a hard-kill. Decoys are the only non-instantaneous soft-kill measure as they require transition to the target, after which it remains active for some defined time parallel to additional salvos.

Multiple Incoming Threats

Due to time limitations, we focused on a single threat case to better explore how the salvos should be structured. Single threats are not how real world scenarios may occur and instead the USN may face salvos with simultaneous threat groupings. This is magnified with the advent of loitering munitions (Atherton 2021).

How to implement multiple threats would require more insight from subject matter experts on what simplifying assumptions can be made. The main considerations are inventory limitations, channel limitation as a finite amount of active/semi-active seekers can be airborne at a time, and expenditures such as chaff that effect multiple threats at once. One approach could be to use our model in order to generate a diverse collection of sequences for each threat which feed into an optimizer with constraints applied.

In the literature review, the idea of staggered salvos was brought up as an advanced tactics not currently in the USN’s repertoire (Wang et al. 2017). Although it could be applied to the single-threat case as an extension, the real value would be in the multi-threat case where you hedge against the number of threats still incoming. If the interceptors fired have active or semi-active homing and can reacquire targets, the outcome of the first salvo’s BDA will

be utilized and interceptors will be utilized more efficiently.

Variable Timing

One limiting assumption is “A salvo is fired as early as possible”. Our model provides near optimal-solutions, and there are potentially lower cost solutions that can be found optimizing the firing time of each salvo. This would significantly increase the computational complexity of the problem. One approach to incorporate timing into the model in a computational tractable way would be to introduce a “waiting” action into the design. The waiting action would for a set length of time (e.g., wait 10 seconds before firing a salvo) that would be operator adjustable. The implementation of such feature is that a given firing sequence will have improved launch timing and our heuristic model provides an even closer approximation.

The trade-off to this is that the smaller the waiting length, the more computationally intensive this feature would become. A potential solution would be to make Q a part of the feasible generation process. Currently we generate all possible firing sequences — even ones that are overkill — and retroactively filter and sort the best sequences. Instead, we introduce the idea of an upper bound to PRA called $\hat{Q}$ that is utilized during the sequence generation process. Once a firing sequence’s PRA surpasses $\hat{Q}$, we terminate our recursion algorithm so that no more branching occurs and that sequences remains a leaf to the recursion model. $\hat{Q}$ could be manually set by the operator or be assigned through the formula $\hat{Q} = 1 - (1 - Q)^F$, where F is a factor for how far over Q we want to explore.

6.3.2 Features Not Explored

Some of the additional features that could be explored in Chapter 5 are:

Variable Interceptor Cost

Although built into the model for the end user to utilize, assigning different cost value to the interceptors was not included in our analysis. We focused on sequence structuring, which is a predecessor step to introducing interceptor cost. Traditionally viewed from a financial perspective, cost can be implemented in other ways such as remaining interceptor inventory.

Validity of Heterogeneous Salvos

For the parameters explored, we have demonstrated the value of having heterogeneous firing sequences. However, for the top sequences generated each salvo only contains one type of interceptor. This suggests that the optimal allocation of interceptors within a salvo is homogeneous. This makes sense as each interceptor has its niche where it is the optimal choice. When the model is scaled up to the multi-threat case or when variable cost is introduced, this might not hold. Further investigation is required and if the optimal solutions retain homogeneous salvos which would allow our model's computational complexity to be drastically reduced.

Entry Range and Threat

We only explored one threat at a single DCR in Chapter 5. The threat we utilized was the C802 and the DCR was based on the AN/SPY-1's radar range of roughly 324 km (175 NM). The model does not inherently account for factors such as radar system, ship, atmospheric conditions, and flight behavior of the target such as sea-skimming, traditional flight, or ballistic trajectory. DCR should be explored for a variety of cases to discover how such situations affect firing sequences.

THIS PAGE INTENTIONALLY LEFT BLANK

List of References

- Armstrong MJ (2005) A Stochastic Salvo Model for Naval Surface Combat. *Operations Research* 53(5):830–841, <http://dx.doi.org/10.1287/opre.1040.0195>.
- Armstrong MJ (2013) The salvo combat model with area fire. *Naval Research Logistics (NRL)* 60(8):652–660, 1520-6750, <http://dx.doi.org/10.1002/nav.21559>.
- Armstrong MJ (2014a) Modeling Short-Range Ballistic Missile Defense and Israel’s Iron Dome System. *Operations Research* 62(5):1028–1039, <http://dx.doi.org/10.1287/opre.2014.1309>.
- Armstrong MJ (2014b) The salvo combat model with a sequential exchange of fire. *Journal of the Operational Research Society* 65(10):1593–1601, <http://dx.doi.org/10.1057/jors.2013.115>.
- Atherton K (2021) Loitering munitions preview the autonomous future of warfare. Technical report, Brookings, <https://www.brookings.edu/techstream/loitering-munitions-preview-the-autonomous-future-of-warfare/>.
- Dutta D (2014) Probabilistic Analysis of Anti-ship Missile Defence Effectiveness. *Defence Science Journal* 64(2):123–129, 0976-464X, <http://dx.doi.org/10.14429/dsj.64.3532>, number: 2.
- FrontlineSolvers (2011) Excel Solver - Algorithms and Methods Used. <https://www.solver.com/excel-solver-algorithms-and-methods-used>.
- Glazebrook K, Washburn A (2004) Shoot-Look-Shoot: A Review and Extension. *Operations Research* 52(3):454–463, 0030-364X, <http://dx.doi.org/10.1287/opre.1030.0086>.
- Glazebrook KD, Kirkbride C, Mitchell HM, Gaver DP, Jacobs PA (2007) Index Policies for Shooting Problems. *Operations Research* 55(4):769–781, 0030-364X, <http://dx.doi.org/10.1287/opre.1070.0444>.
- Hughes W (1995) A Salvo Model of Warships in Missile Combat Used to Evaluate Their Staying Power. *Naval Research Logistics (NRL)* 42(2):267–289, <https://calhoun.nps.edu/handle/10945/60793>.
- Janes Database: Aegis (2021) Janes: Command and Weapon Control Systems - Aegis. <https://customer.janes.com/Janes/Display/JNWS0111-JC4IM>.
- Janes Database: C802 (2021) Janes: Weapons: Naval - CSS-N-4 ‘Sardine’ (YJ-8/YJ-8A/C-801); CSS-N-8 ‘Saccade’ (YJ-82/YJ-83/C-802/C-802A/Noor/Ghader). https://customer.janes.com/Janes/Display/JNWS0132-JNW_.

- Kalyanam K, Clarkson J (2021) Sequential attack salvo size is monotonic nondecreasing in both time and inventory level. *Naval Research Logistics (NRL)* 68(4):485–495, 1520-6750, <http://dx.doi.org/10.1002/nav.21967>.
- Kesler G (2019) *A DATA FARMING ANALYSIS OF MULTIPLE SALVO EQUATIONS*. Ph.D. thesis, Naval Postgraduate School, Monterey, CA, <https://calhoun.nps.edu/handle/10945/62263>.
- Kline A, Ahner D, Hill R (2019) The Weapon-Target Assignment Problem. *Computers & Operations Research* 105:226–236, 0305-0548, <http://dx.doi.org/10.1016/j.cor.2018.10.015>.
- Leboucher C, Shin HS, Ménéci S, Tsourdos A, Kottenkoff A (2013) Optimal Weapon Target Assignment Based on an Geometric Approach. *IFAC ACA 2013*, <http://dx.doi.org/10.3182/20130902-5-DE-2040.00037>.
- Lucas TW, McGunnigle JE (2003) When is model complexity too much? Illustrating the benefits of simple models with Hughes’ salvo equations. *Naval Research Logistics (NRL)* 50(3):197–217, <http://dx.doi.org/https://doi.org/10.1002/nav.10062>.
- Manne AS (1958) A Target-Assignment Problem. *Operations Research* 6(3):346–351, 0030-364X, <http://dx.doi.org/10.1287/opre.6.3.346>.
- Payne CM, ed. (2010) *Principles of naval weapon systems* (Annapolis, Md: Naval Institute Press), 2nd ed edition, 978-1-59114-667-4.
- Popa CH, Stone SP, Aw EHA, Teo CPJ, Cai LE, Chong WH, Cline R, Hong JJ, Koh CKR, Lee WL, Lim JJ, Lin X, Liraz SP, Mok KBE, Ryan AW, Teow BHA, Whitmer AR, Winstead PJ, Chia YJC, SEA CS (2018) DISTRIBUTED MARITIME OPERATIONS AND UNMANNED SYSTEMS TACTICAL EMPLOYMENT. <https://calhoun.nps.edu/handle/10945/59587>.
- Pryluk R, Shima T, Golan O (2013) Guidance - Based missile allocation algorithm against a salvo attack. *52nd IEEE Conference on Decision and Control*, 5391–5396, <http://dx.doi.org/10.1109/CDC.2013.6760737>.
- Soland RM (1987) Optimal Terminal Defense Tactics When Several Sequential Engagements Are Possible. *Operations Research* 35(4):537–542, 0030-364X, <https://www.jstor.org/stable/170499>.
- Stanley RP (2012) *Enumerative combinatorics. Volume 1*. Number 49 in Cambridge studies in advanced mathematics (Cambridge, NY: Cambridge University Press), 2nd ed edition, 978-1-107-01542-5 978-1-107-60262-5.

- Stires H (2017) CNO Announces the Return of Vertical Launch System At-Sea Reloading. <https://nationalinterest.org/feature/exclusive-cno-announces-the-return-vertical-launch-system-21425>.
- Summers DS, Robbins MJ, Lunday BJ (2020) An approximate dynamic programming approach for comparing firing policies in a networked air defense environment. *Computers & Operations Research* 117:104890, 0305-0548, <http://dx.doi.org/10.1016/j.cor.2020.104890>.
- Vego M (2010) On Naval Warfare. *Tidskrift i Sjöväsendet* 2010(1):73–92, 0040-6945, <https://apps.dtic.mil/sti/pdfs/AD1014487.pdf>.
- Wang L, Yao Y, He F, Liu K (2017) A novel cooperative mid-course guidance scheme for multiple intercepting missiles. *Chinese Journal of Aeronautics* 30(3):1140–1153, 1000-9361, <http://dx.doi.org/10.1016/j.cja.2017.01.015>.
- Washburn A, Kress M (2009) Shooting with Feedback. Washburn A, Kress M, eds., *Combat Modeling*, 47–64, International Series in Operations Research & Management Science (Boston, MA: Springer US), 978-1-4419-0790-5, http://dx.doi.org/10.1007/978-1-4419-0790-5_3.
- Wolfram (2021) Wolfram Alpha Calculation. <https://www.wolframalpha.com/input/?i2d=true&i=1+-+Power%5B%281-p%29%2Cx%5D%28%28Divide%5Blog%281-Q%29%2Clog%281-p%29%5D-x%29log%281-p%29-1%29%3D0>.

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