NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

MISSION-BASED UAV SWARMS: BASE DEFENSE

by

Casey G. Escamilla

September 2020

Thesis Advisor:	Duane T. Davis
Co-Advisor:	Kathleen B. Giles

Approved for public release. Distribution is unlimited.

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE

Form Approved OMB No. 0704-0188

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 September 2020

3. REPORT TYPE AND DATES COVERED Master's thesis

4. TITLE AND SUBTITLE
MISSION-BASED UAV SWARMS: BASE DEFENSE

5. FUNDING NUMBERS

6. AUTHOR(S) Casey G. Escamilla

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)

Forward Operating Base (FOB) defense is a manpower-intensive mission that takes valuable resources away from the operational mission. While increasingly capable unmanned aerial vehicles (UAVs) might perform many of the mission tasks, current doctrine does not adequately address their inclusion. In particular, the assumed one-to-one ratio of operators to vehicles does not account for increasing UAV autonomy. This thesis describes the development and testing of an autonomous FOB defense capability using the Advanced Robotic Systems Engineering Laboratory (ARSENL) swarm system. Development leveraged the Mission-based Architecture for Swarm Composability (MASC) for development of complex swarm behaviors in a mission-focused, top-down manner. This approach enabled the development of a doctrinally grounded base-defense tactic in which arbitrary mixes of fixed-wing and quadrotor UAVs autonomously assigned and performed all required FOB defense roles: perimeter surveillance, key area search, contact investigation, and threat response. The tactic was extensively tested in a software-in-the-loop simulation environment and demonstrated during live flight field exercises. Experimental results are discussed using measures of effectiveness and measures of performance that were developed over the course of this research.

14. SUBJECT TERMS
unmanned air vehicle, UAV, swarm, Mission-based Architecture for Swarm Composability, MASC, unmanned, autonomous, Advanced Robotic Systems Engineering Laboratory, ARSENL, software in the loop

15. NUMBER OF PAGES
77

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

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. 239-18

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release. Distribution is unlimited.

MISSION-BASED UAV SWARMS: BASE DEFENSE

Casey G. Escamilla
Captain, United States Marine Corps
BS, U.S. Naval Academy, 2014

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN COMPUTER SCIENCE

from the

NAVAL POSTGRADUATE SCHOOL
September 2020

Approved by: Duane T. Davis
Advisor

Kathleen B. Giles
Co-Advisor

Gurminder Singh
Chair, Department of Computer Science

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

Forward Operating Base (FOB) defense is a manpower-intensive mission that takes valuable resources away from the operational mission. While increasingly capable unmanned aerial vehicles (UAVs) might perform many of the mission tasks, current doctrine does not adequately address their inclusion. In particular, the assumed one-to-one ratio of operators to vehicles does not account for increasing UAV autonomy. This thesis describes the development and testing of an autonomous FOB defense capability using the Advanced Robotic Systems Engineering Laboratory (ARSENL) swarm system. Development leveraged the Mission-based Architecture for Swarm Composability (MASC) for development of complex swarm behaviors in a mission-focused, top-down manner. This approach enabled the development of a doctrinally grounded base-defense tactic in which arbitrary mixes of fixed-wing and quadrotor UAVs autonomously assigned and performed all required FOB defense roles: perimeter surveillance, key area search, contact investigation, and threat response. The tactic was extensively tested in a software-in-the-loop simulation environment and demonstrated during live flight field exercises. Experimental results are discussed using measures of effectiveness and measures of performance that were developed over the course of this research.

THIS PAGE INTENTIONALLY LEFT BLANK

Table of Contents

List of Figures

THIS PAGE INTENTIONALLY LEFT BLANK

List of Tables

THIS PAGE INTENTIONALLY LEFT BLANK

List of Acronyms and Abbreviations

ARSENL	Advanced Robotic Systems Engineering Laboratory
ATP	Army Technical Publication
CPG	Commandant’s Planning Guidance
DARPA	Defense Advanced Research Projects Agency
DOD	Department of Defense
EABO	Expeditionary Advanced Base Operations
FOB	forward operating base
FM	Field Manual
FX	Field Exercise
JP	Joint Publication
MAGTF	Marine Air-Ground Task Force
MASC	Mission-based Architecture for Swarm Composability
MASINT	measurement and signature intelligence
MCDP	Marine Corps Doctrinal Publication
MCRP	Marine Corps Reference Publication
MEO	Mission Execution Ontology
MOE	measure of effectiveness
MOP	measure of performance
MOS	Military Occupational Specialty

NPS	Naval Postgraduate School
OffSET	Offensive Swarm Enabled Technology
ROS	Robot Operating System
SITL	software in the loop
SRWBR	short range wide band radio
TCP	Transmission Control Protocol
UAS	unmanned aerial system
UAV	unmanned aerial vehicle
UDP	User Datagram Protocol
USG	United States government
USMC	United States Marine Corps
USN	U.S. Navy

Acknowledgments

First and foremost, I want to thank my advisors, Dr. Davis and CDR Giles. Your patience with translating my “trucker’s English” to academic thought was invaluable. Your wisdom and encouragement as a team made the road an enjoyable one. I also want to thank the ARSENL team for all of the support they gave me, which provided the opportunity for two live flights during the course of this research. Finally, I would like to thank my wife, Erin, for holding the line with our three boys, Orion, Wesley, and Seamus, the latter two of whom were born during this research. Your strength is an inspiration and truly something to see.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 1:
Introduction

1.1 Background and Motivation

In 2019, the Commandant of the Marine Corps, General David H. Berger, released his planning guidance as a way of shaping the force for the next four years. In it he stated, “As good as we are today, we will need to be even better tomorrow to maintain our warfighting overmatch” [1]. Taken from Commandant of the Marine Corps General David H. Berger’s 2019 Commandant’s Planning Guidance (CPG), this quote calls for focused action to meet the ever-changing challenges that the Marine Corps is expected to face in the coming wars. Additional guidance within the CPG, which sets priorities and direction for the Marine Corps for the next four years, calls for a “robust family of unmanned systems suitable for reconnaissance, surveillance, and the delivery of lethal and non-lethal effects” [1]. General Berger further called for exploiting new technologies in support Expeditionary Advanced Base Operations (EABO). EABO will require agile systems capable of both effective offensive operations and independent and sustainable defensive operations. In short, realizing EABO will require maximized use of each system and Marine.

In essence, General Berger is calling for a change in the employment of unmanned vehicles. Meeting this objective will be facilitated through the use of large systems of cooperating autonomous unmanned vehicles, or swarms. Unmanned vehicle swarms provide the ability to exponentially increase battlefield capabilities with minimal increases in manpower requirements and logistical burden. As General Berger references the “next battlefield,” the Marine Corps will have to take advantage of technologies that maximize the use autonomy and the impact that each war-fighter has on the battlefield.

Current doctrine for employment of unmanned systems is centered around systems with little or no autonomy. Also, current systems rely on remote piloting of individual vehicles; that is, one operator for each vehicle. This lack of autonomous systems in the force has created a gap in operational employment ability for both surveillance and direct action. Furthermore, doctrine for unmanned systems focusing on one-to-one operator-vehicle management requires the number of operators to scale linearly with the number of vehicles. This will not be sufficient for the "next battlefield." Rather, the Marine Corps will require systems that disencumber the operators or increase their ability to simultaneously control multiple vehicles [2].

With these objectives in mind, the Naval Postgraduate School (NPS)'s, Advanced Robotic Systems Engineering Laboratory (ARSENL) has developed and demonstrated a system for control of large, autonomous, multi-vehicle systems that leverages the strengths of distributed computing and minimizes the cognitive requirements of piloting. ARSENL demonstrated the efficacy of their system during field experiments in which 50 autonomous unmanned aerial vehicles (UAVs) were successfully launched, simultaneously controlled by a single operator, and safely recovered [3].

1.2 Research Objectives

This primary objective of this research was to demonstrate a proof of concept for employment of UAV swarms in support of forward operating base (FOB) defensive operations. In particular, this entailed the autonomous generation, assignment, and execution of the subtasks required for effective, doctrine-compliant base defense. This component of the research focused on the development of state-based descriptions of surveillance, investigation, and threat response tasks; implementation of decision mechanisms supporting vehicle-level task assignment; and on-vehicle control during task execution.

Subsidiary research objectives included demonstration of the Mission-based Architecture for Swarm Composability (MASC) process to develop complex swarm behaviors in a top-down, mission-focused manner, exploration of distributed approaches to autonomous swarm control and decision-making, and implementation of general swarm algorithms and plays that will prove useful in a broad array of potential swarm tactics. Collectively, these objectives were part and parcel to the primary objective and served as the means by which it was achieved.

1.3 Methodology

The development of the `Base Defense` tactic began with a review of existing base defense doctrine. This review served as the basis for identifying essential tasks and sub-tasks that were to be accomplished by the behavior. We then reviewed current doctrine regarding the employment of UAVs in the Marine Corps to determine whether it accounted for the employment of these systems in the base defense mission.

Following characterization of the mission-required tasks, we developed a high level state diagram for the overall mission of base defense. We then developed state diagrams for all of the states of the high-level diagram. The subtask-level state diagrams equated to plays within the MASC hierarchy.

Existing algorithms and plays within the ARSENL codebase and new algorithms and plays that were developed over the course of this research were used to implement subtask-level state diagrams on the ARSENL system. Finally, the plays were combined according to the high-level state diagram to complete the `Base Defense` tactic's implementation.

Following play and tactic development, doctrinally based measures of effectiveness (MOEs) and measures of performance (MOPs) were devised. These were used to assess the `Base Defense` tactic through extensive experimentation in a software in the loop (SITL) simulation environment. The tactic and all included plays were also demonstrated during live-flight experiments conducted at Camp Roberts, California.

1.4 Results

Ultimately, this research was successful in its primary objective and demonstrated a base defense tactic incorporating perimeter surveillance, key area search, contact investigation, and threat response plays. Further, development relied heavily on the MASC hierarchy as an approach to developing mission requirements and breaking these requirements down into manageable tasks that could be implemented on the ARSENL swarm system. This tactic was tested in both live-flight and simulation environments and evaluated using mission-focused MOPs and MOEs. Final results were encouraging and the tactic that was developed over the course of this research was assessed as satisfactory proof of concept.

1.5 Thesis Organization

This thesis is divided into six chapters. Chapter 1 provided the motivation for this research, described the capability gap that this proof of concept seeks to close, and provided a short background on ARSENL and the research objectives that were pursued. Chapter 2 discusses the current Marine Corps doctrine for rear area operations as described in Marine Corps and Joint Publications. It also provides an overview of current UAV employment within the Marine Corps and describes the levels of autonomy that various current systems are capable of. Chapter 3 provides an overview of previous work in the field of behavior based architectures for autonomous systems, the ARSENL multi-vehicle unmanned aerial system (UAS), and the MASC hierarchy. Chapter 4 provides a state-based description of the overall design of the `Base Defense` tactic as well as the plays upon which the high-level tactic relies. This chapter also details the methodology used to create, test, and evaluate this proof of concept. In doing this, the focus is primarily on the MOPs and MOEs that each tactic and play was evaluated against. Chapter 5 details the live-flight and simulation experiments that were conducted and discusses the test results in regard to the associated MOPs and MOEs. Finally, Chapter 6 presents the conclusions from this proof of concept. This chapter also provides recommendations for future work relating both to the `Base Defense` tactic itself and also to autonomous swarm capability and control more broadly.

CHAPTER 2:
Base Defense Doctrine and Current UAS Employment in the USMC

The Marine Corps employs FOBs throughout the globe to enable Marines to fight in any clime and place. These bases provide a launching point for operations and a relatively safe location for planning, rest, and resupply. These bases are only able to securely provide these services through the employment of defense-in-depth which has always been a drain on manpower. UASs are currently employed in the Marine Corps to provide surveillance throughout areas of operation to include FOBs, but the current utilization is not sufficient to the objective. This chapter provides an overview the doctrine surrounding base defense and UAS utilization, respectively.

2.1 Base Defense Doctrine

The primary publications for base defense are Joint Publication (JP) 3-10.1, *Joint Tactics, Techniques, and Procedures for Base Defense* [4] and Marine Corps Reference Publication (MCRP) 3-30C.1, *Marine Air-Ground Task Force (MAGTF) Rear Area Security Operations* [5]. Both of these publications outline the planning considerations, essential tasks, and primary objectives of base defense. These two publications serve as the guideline for determining swarm tasks used by the defense tactic developed over the course of this research.

Both JP 3-10.1 and MCRP 3-30C.1 define the primary mission of the base as the provision of support. The base in turn is a supporting establishment rather than a supported establishment. This means that base functions, to include security, must not interfere with the operations that they support. Unfortunately, as bases grow, the required minimal force to support them also grows. The largest element of this support structure is the base defense force, which grows linearly with the size of the supported combat unit that the base supports.

MCRP 3-30C.1 states that the objective of base defense is to “maintain a secure base and minimize disruptions to primary support missions.” Furthermore, a base is defined as “an area or locality from which operations are projected or supported and contains installations which provide logistic or other support" [5]. In other words, bases serve as an enabling structure for forces to complete missions. As such, effective base defense prioritizes the following objectives:

1. Provide all-round security with a minimum of available forces
2. Maintain a secure base
3. Minimize disruptions to primary support missions [5]

JP 3-10.1 classifies the three levels of threats to the joint rear area as:

- Level 1: Agents, saboteurs, sympathizers, terrorists
- Level 2: Small tactical units, unconventional warfare forces, guerrillas
- Level 3: Large tactical force operations, including airborne, heliborne, amphibious, infiltration, and major air operations [4]

The doctrinal expectation for defense of FOBs is that they are self-sustainable against Level 1 and 2 threats but require reinforcements to repel Level 3 forces.

MCRP 3-30C.1 further specifies four main tasks for base defense:

1. **Securing the Base**: Establishing defensive measures within a defensive plan that is disseminated throughout the Rear Area.
2. **Detecting the Enemy**: Surveillance and dissemination of information from friendly units that provides early warning of attack. Patrolling, listening posts/observation posts (LPs/OPs), and perimeter surveillance serve as the primary implementations of this task.
3. **Delaying the Enemy**: Slowing and disrupting the enemy's progress. This is accomplished through the use of both natural and man-made obstacles under cover of direct and indirect fires.
4. **Destroying the Enemy**: Repelling enemy attacks through the use of direct and indirect fires as well as close air support [5].

Any system that is intended for use in base defense must support at least one of these tasks while minimizing manpower and logistical burden.

2.1.1 Securing the Base

The task of securing the base is assigned to the base commander and is accomplished through the promulgation and execution of a defensive plan. In this order, defensive measures are prescribed and prioritized. These measures consist of constructing barriers around a defined perimeter, establishing entry control points, and establishing perimeter security. Entry control points allow for vetting of the mounted and unmounted traffic through the base and must be manned. Perimeter security requires an armed force to deter and repel enemy infiltration. This is typically provided through the use of manned guard posts along the base perimeter and a variety of stationary surveillance assets. Each of these tasks are manpower intensive and increase as the size of the base expands.

2.1.2 Detecting the Enemy

The task of detecting the enemy is intended to provide advanced warning of attempts by the enemy to infiltrate the base, disrupt operations, or destroy MAGTF forces. Prior to occupation of the FOB, an analysis of the area is conducted to identify key areas in the vicinity of the base. A *key area* or key terrain is defined by MCRP 2-3A as “any locality, or area, the seizure or retention of which affords a marked advantage to either combatant” [6]. These areas include terrain features that provide a vantage point, depressions where enemy movements can be hidden, heavily trafficked routes and intersections. In short, surveillance of key areas is a priority for base defense. Base commanders primarily rely on manned patrols and listening posts/observation posts (LPs/OPs) to detect the enemy.

Patrols consist of a force of any size, from fire team to platoon, which provide surveillance and disrupt enemy activity within a given area, usually around a central point such as a base or command post. Patrols serve as an unpredictable disrupting force that requires the enemy to constantly adapt. Patrols, mounted and unmounted, pose a large logistical and manpower burden on the base’s occupying force. Not only does it remove patrol personnel from their primary missions, it adds an additional burden and mission that contributes to battlefield fatigue.

LPs/OPs consist of an inherently small force of as small as a single buddy pair, which provide limited surveillance in exposed and cut-off areas or terrain features. LPs/OPs serve as stationary posts that are located in the vicinity of areas of interest to serve as an early warning of enemy activity. LPs/OPs therefore are the one of the primary methods for detecting the enemy beyond line of sight at most FOBs. LP/OP missions require units to provide an inserting and extracting force. This is usually accomplished during patrols. Thus, the manpower requirements for this form of surveillance is substantial.

2.1.3 Delaying the Enemy

Marine Corps Doctrinal Publication (MCDP) 1 describes speed, or more precisely, relative speed, as a crucial factor in warfare. The purpose of delaying the enemy is to slow the enemy's attack and increase the base's relative speed [7]. In this context, delaying the enemy increases friendly force relative speed and provides base security forces time to respond and gain the advantage. The delay task is commonly implemented through the use of obstacles to include wire barriers, minefields, tank berms and ditches, and natural terrain. MCWP 3-20.5 further defines countermobility as the use of "obstacles integrated with fires to inhibit the maneuver of an enemy force, increasing the time for target acquisition, and increase the effectiveness of friendly force weapons" [8]. Thus, once the enemy has been detected, the delaying objective is to slow or halt his progress long enough for friendly weapons engagement. The enemy must then be tracked until his presence is neutralized.

2.1.4 Destroying the Enemy

The destroy task is the culminating event in the defensive cycle. This task is typically implemented through the use of fires and security forces to eliminate the enemy threat. Two types of fires are utilized to complete this task, direct and indirect.

Direct fires include small arms such as the M4 carbine or M16 rifle organic to every individual Marine and crew served weapons such as the M240B medium machine gun and M2 .50 caliber heavy machine gun that are commonly employed at guard posts and entry control points. A fundamental weakness of direct fires is that they are incapable of engaging targets that occupy dead space. *Dead space* is defined as that area in which direct fire weapons have no effect. Usually, dead space is found in a depression or behind a terrain feature [9].

Indirect fires consist primarily of mortars and artillery and provide cover of dead space. Indirect fire units require a forward observer who has to identify and report back the location, size, and type of target [10].

2.2 Current UAS Employment in the USMC

Marine Corps Warfighting Publication 3-20.5 (MCWP 3-20.5) *Employment of UAS within the Marine Corps* defines UAV systems and their employment within Marine Corps operations and in the joint environment. The Marine Corps defines an UAS as a "system whose components include the necessary equipment, network, and personnel to control an unmanned aircraft." These systems are recognized separately from munitions and decoys in that they are generally intended for recovery and reuse. UASs are categorized into five groups based on their weight, operating altitude, and maximum flight speed as indicated in Table 2.1.

Table 2.1. This table describes the flight categories as defined within MCWP 3-20.5 by the Federal Aviation Administration, Department of Transportation, and the Department of Defense. Source: [2].

UA Category	Maximum Gross Takeoff Weight (pounds)	Normal Operating Altitude (feet)	Speed (knots indicated airspeed)
Group 1	0-20	< 1,200 AGL	< 100
Group 2	21-55	< 3,500 AGL	< 250
Group 3	56-1,320	< 18,000 MSL	< 250
Group 4	> 1,320	< 18,000 MSL	Any Airspeed
Group 5	> 1,320	> 18,000 MSL	Any Airspeed

MCWP 3-20.5 asserts that increased endurance and reduction of human factors are the primary strengths of UASs. The principle factor of endurance is that the dwell time of UASs and the ability to "hot seat" UAS operators provides a marked advantage over manned systems limited by pilot endurance. Longer endurance leads to fewer aircraft required, a more complete surveillance picture, increased flexibility for unit commanders and faster reaction times. Also with regard to human factors, the significantly reduced risk to crew makes a UAS a more employable system in contested environments. This is not to say, however, that no human factors are present in the operation of UAVs, simply that these factors are more easily overcome or spread across multiple operators. Further, MCWP 3-20.5 points out that the reliance on sensors and vulnerability to meteorological effects are the primary weaknesses of UASs.

A principal concern with the implementation of UAS surveillance is the lack of "cockpit observation." This possible loss of information due to the reliance on sensors presents a possible vulnerability that must be accounted for during execution. Meteorological phenomena such as precipitation and wind have a greater effect on UASs due to the tendency for UASs to be smaller than their unmanned counterparts. This weakness means that the use case for perpetual reconnaissance and early warning systems must be augmented by manned systems or that UASs must improve to serve in all weather environments [2].

Currently, the Marine Corps employs three Group 1 systems and two systems above Group 1. The Group 1 systems include RQ-11B Raven DDL, RQ-12A Wasp All-Environment, and the RQ-20A Puma All-Environment. Each of these systems is designed for employment at the battalion echelon or lower and within Marine Special Operations (MARSOC) Companies and Teams. Two critical attributes of these systems are that they are considered "fully autonomous" and that they are capable of identifying man-sized targets at their operating altitude.

Group 1 UAV systems are relatively simple and are operated as a collateral duty rather than as a primary Military Occupational Specialties (MOSs) (i.e., operators do not receive special training and their UAV roles are in addition to their primary responsibilities). These systems can be operated by a single operator but are optimized with two operators, one serving as the mission operator and the other as the vehicle operator. This relatively simple system provides surveillance in an operational radius of ten kilometers and a runtime of 60 to 90 minutes. Thus, the manpower tradeoff for surveillance or reconnaissance is a two to one ratio for our most forward units.

2.3 Unmanned Aerial System (UAS) Utilization for Base Defense

Use of UAS for base defense in the Marine Corps is currently limited to the surveillance space, or "Detecting the Enemy," due to the lack of armed UASs in the current inventory. Also, the low level of autonomy that these systems are capable of paired with the manpower requirements for each system create a gap in our technical abilities and thus a gap in our defenses. Further, the requirement of two Marines to operate a single surveillance aircraft for Marine Corps Group 1 systems represents a high cost to FOBs in austere environments, where manpower is at a premium.

To close this gap and reduce manpower requirements, the Marine Corps must invest in systems that scale more efficiently in surveillance and response capabilities when compared to current manpower requirements. These capabilities must be able to replace or reduce current FOB defense tasking with as little additional operator overhead as possible. In other words, the United States Marine Corps (USMC) must find a way to maximize the use of our most prized asset, the individual Marine, by leveraging low-level autonomy. This research outlines a solution which follows base defense doctrine and reduces or replaces tasks conducted by live Marines to meet base defense objectives.

2.4 Summary

In summary, the primary objective of base defense is to provide the most secure point at the lowest manpower cost to the occupying unit. Every Marine that is required to protect the base not otherwise available to the occupying unit and his or her absence is an impediment to mission success. Any tool that maximizes the capability of the individual while minimizing the cognitive complexity should be leveraged. Swarming systems of autonomous UAVs provide one such capability that can maximize the potential of individual Marines while minimizing non-mission workload.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 3: Robotic Swarm Systems

This chapter presents an overview of robotic swarm system architectures, the MASC hierarchy, and previous work completed by the ARSENL that contributed to the development of the `Base Defense` tactic.

3.1 Swarm System Architectures

Multiple approaches to multi-vehicle system behavior have proven successful. One of the first of these architectures was developed by Rodney Brooks in 1985 and presented in his paper "A Robust Layered Control Systems for a Mobile Robot." The goal of what Brooks dubbed the Subsumption Architecture was to "reduce complex tasks to their simplest parts and enable modularity and control of complex systems" [11]. The Subsumption Architecture accomplished this by dividing complex tasks into their simplest parts, called modules, and modelling these modules as finite state machines. This layered approach to abstraction allowed for modularity because the simpler modules could be reused for multiple complex tasks and produced "increasing levels of competence" for the autonomous robots [11]. Though Brooks initially proposed the Subsumption Architecture for control of individual robots, its reliance on a simplified world model and its building block approach to behavior development have made it applicable to multi-robot systems as well [3], [12].

An expansion of Brooks' architecture with modern applications was developed in 2018 by Brutzman et al. in [13]. Their paper proposes a solution to the problem of ensuring ethical control of unmanned systems through a mathematically rigorous controller and human oversight of mission orders. The research team posits that in order to ensure ethical autonomous execution of a mission, which would include not only mission objectives but also operational constraints, a common mission definition and syntax must be understood by both operator and vehicle. This common syntax is implemented through the use of a Mission Execution Ontology (MEO), a human-readable form that enables mathematical validation of a described mission and execution by a target vehicle. The MEO provides the ability to divide tasks into a hierarchy that is modelled after the command organization of naval vessels. The hierarchy consists of strategic, tactical, and execution layers, which correspond to the duties of a commanding officer, watch officers, and watchstanders, respectively. This layered model simplifies the overall complexity of the problem and allows for the creation of a mathematical, state defined, approach to ensuring ethical constraints are met. The primary takeaway for this research, is that a mission-based approach to swarm control generates greater understanding of mission flow [13]. The MEO approach to mission definition and control has recently been tested with large multi-UAV systems in simulation and live-fly experiments [14].

Additional research that yielded real world results was completed by Brett Browning et al. in their paper "STP: Skills, Tactics and Plays for Multi-robot Control in Adversarial Environments" [15]. The goal of their research was to develop plays and tactics which capitalize on multi-robot teaming and quick decision making. They tested their control methods by competing in the RoboCup robot soccer tournament in which a vehicle team was required to make real-time decisions to defend against adversarial attacks and mount their own offensive maneuvers [16].

Browning's team broke the problem into two primary parts: defense and offense. For defensive plays the team was programmed to recognize patterns at the individual and team layers. Once certain patterns were identified, the vehicles reoriented into appropriate blocking formations [15]. For offensive plays the team developed sequences of action that could be combined to create a full attack [15]. This research resulted in a reactive multi-vehicle control architecture [17] that allowed for timely decision making and simple testing and augmentation.

Despite being initially developed in 2004, the "skills, tactics, and plays" approach to multi-vehicle system control continues to find adherents. In 2017, for instance, the architecture was augmented with a voting system to better overcome distributed system challenges that often hamper multi-vehicle teams. Also utilized in a RoboCup competition, this augmented approach was noted to significantly improve offensive plays in particular. [18].

Miller et al. [19] proposed a way to simplify vehicle control in their research which complements the hierarchical approaches developed by Brooks and Browning. Miller's research sought to develop a UAV control system that was both simple and useful. Miller's work was based on the assertion that current systems are either too complex and therefore require extensive training and constant user control input, or too simple to provide any real utility.

To overcome this, Miller equated the control of a UAV system to a quarterback calling a play in a football game. In this approach he simplified behaviors to pre-designed plays requiring only simple parameters. This served to limit the interaction required from the user and to maximize the efficiency of the user's time. The research team was able to demonstrate UAV control with little to no specialized training and less than 15 seconds to complete each request for support. [19]

A final example of hierarchical vehicle control is found in an open source framework for UAVs, AEROSTACK. AEROSTACK is a "system architecture and open-source multi-purpose software framework for autonomous multi-UAS operation" [20]. The creators of AEROSTACK sought to create a system for UAV control with a large scope of use cases that was also hardware independent. The creators segmented the control in to five layers: reactive, executive, deliberative, reflective, and social. Through the use of abstraction and compartmentalization these layers created a hierarchy for control. The lower layers–reactive, executive, and deliberative–create a framework that allows for code reuse. These layers address tasks such as motion control and visual perception that are widely reused across multiple applications. The higher layers–reflective and social–accomplish more specific tasks and provide greater degrees of autonomy. These layers provide greater flexibility and autonomy than the lower layers.

The AEROSTACK platform proved effective over the course of three years during which it was used in international competitions and exhibitions to showcase a wide range of applications. For the final test case, AEROSTACK successfully executed a simulated indoor search and rescue mission. The fundamental concept of this system is that utilizing layers of abstraction is a successful approach for simplifying control and expanding capability. [20]

The Defense Advanced Research Projects Agency (DARPA)'s Offensive Swarm Enabled Technology (OffSET) program is another example of the playbook approach. The goal of this program is to "create highly capable, heterogeneous swarm systems with upwards of 250 collaborating autonomous swarm elements, across multiple spatial and temporal scales of tactical interest" The OffSET program uses a top-down approach to behavior implementation that enables simplified swarm control for the warfighter. The OffSET program manager, Dr. Timothy Chung, asserts that the use of the hierarchy creates a level of abstraction that enables the focus on commander's intent.

OffSET's hierarchy begins with a mission which is decomposed into a set of tactics, then primitives, and finally algorithms. This hierarchy allows for complex behaviors to be adjusted to emerging technology and enables teams to compartmentalize individual complex problems [21].

3.2 Mission-based Architecture for Swarm Composability (MASC)

The Mission-based Architecture for Swarm Composability (MASC) is a behavior based top-down approach to swarm control developed by CDR Kathleen Giles for use with the ARSENL repository. MASC addresses a crucial problem with current systems: the "single robot parenting" paradigm. This problem is summarized in CDR Giles' dissertation with the observation that current systems rely on at least a one-to-one ratio of operators to UAVs; that is, one remote piloting operator and possibly others for each UAV. This notion is specifically assumed by current Marine Corps doctrine [2], [7]. Thus, as the number of UAVs within a system grows, so does the number of operators. This linear growth is what MASC is specifically designed to combat.

CDR Giles suggests a series of system architecture goals to which a viable swarm-control architecture should aspire and proposes a hierarchical approach to meet them. The goals of this system are as follows [12]:

- Remote supervision
- Offload of tasks previously accomplished by humans to the vehicles at the system level
- Reduced operator attention to activity at the individual vehicle level
- Reduced pilot workload
- Modularity in the form of rearrangeable and reusable elements applicable to a variety of missions and conditions
- Intuitiveness
- Composability of reusable elements and components in various combinations to support high-level mission requirements

In accomplishing these goals, the MASC approach is intended to allow a higher vehicle-to-human ratio without taxing the limits of human cognition. These goals also serve as the basis for the measures of effectiveness against which the `Base Defense` tactic developed in this research was assessed.

MASC uses integrated behaviors at the vehicle level that react to triggers. These behaviors allow individual vehicles and UAV groups to accomplish tasks with minimal operator oversight. This prevents task saturation of the human operator while improving reaction time. Behaviors are arranged hierarchically within the MASC architecture as depicted in Figure 3.1.

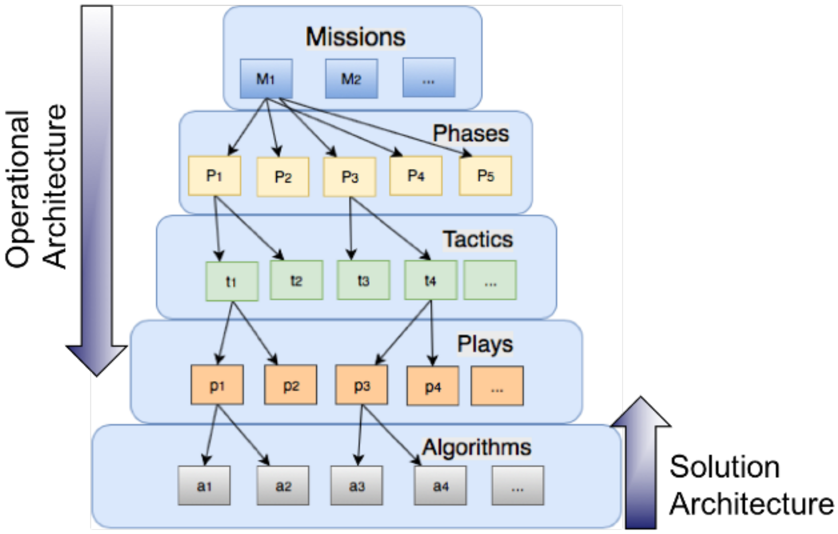

Figure 3.1. Mission-based Architecture for Swarm Composability (MASC) hierarchy demonstrating the relationships between missions, phases, tactics, plays, and algorithms [12].

The highest level of this hierarchy is the mission layer. A mission is defined as the "overall task and purpose delineating the action assigned to the UAV swarm" [12]. CDR Giles' MASC hierarchy breaks a mission down in to predefined phases. The five phases within a typical mission are Pre-flight, Ingress, On Station, Egress, and Post-flight. In this sort of mission, the UAVs will progress sequentially through the phases, however MASC supports other phase-to-phase progressions as well. Each phase is composed of one or more tactics. A tactic is defined as an "employment and ordered arrangement of agents in relation to one another for the purpose of performing a specific task." Finally, each tactic is comprised of one or more plays. Plays are lower-level behaviors through which autonomy is realized. They are the foundational building blocks upon which complex autonomy is based. Plays achieve their functionality through one or more algorithms that support autonomous maneuver and decision making.

The strength of the MASC architecture is that it supports top-down capability development in a mission-focused manner that abstracts decision making away from the lowest levels. This approach alleviates the linear increase in complexity as swarm size increases by allowing human operators to focus on high-level system performance rather than individual vehicle oversight. Further, MASC maximizes the capability of individual agents through the use of a repository of behaviors that can be parameterized to accomplish mission-level goals in an operational environment. This in turn increases the ability to reuse and adapt previously developed behaviors to changes in the operational environment.

3.3 The Advanced Robotic Systems Engineering Laboratory (ARSENL)

The ARSENL swarm system is a multi-UAV architecture that enables the autonomous control of a heterogeneous swarm. The ARSENL Swarm Architecture is comprised of three primary elements: the vehicles, the network, and the ground station. Currently, four airframes, depicted in Figures 3.2 and 3.3, are employed by ARSENL:

- Zephyr II fixed wing
- Zephyr II Gen3 fixed wing
- Penguin fixed wing
- Mosquito Hawk quadrotor

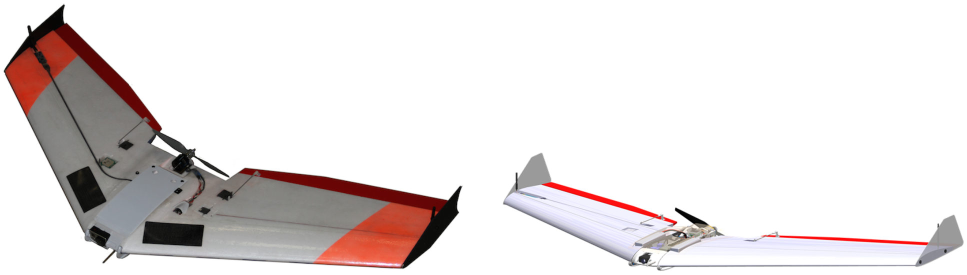

Figure 3.2. ARSENL Vehicles: Zephyr II and Zephyr II Gen 3 [3].

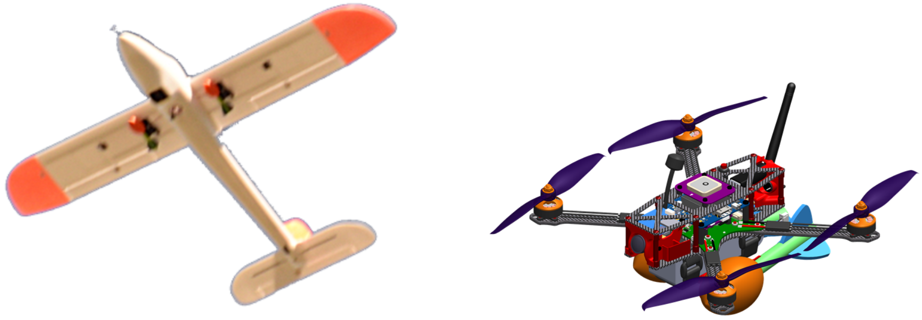

Figure 3.3. ARSENL Vehicles: Penguin and Mosquito Hawk [14].

These vehicles and their capabilities are described in [3] and [14]. The ARSENL system has been showcased in live-fly experiments with individual swarms of up to 50 UAVs and multiple swarms with up to 40 UAVs each [3], [22].

MASC algorithms, plays, and tactics are implemented for each vehicle on an onboard companion or payload computer running an Ubuntu Linux LTS release with the open source Robot Operating System (ROS) [23] installed. The ARSENL autonomy package on the companion computer directs vehicle maneuvering and sensor operation through a PixHawk-family autopilot running open-source Ardupilot [24] firmware with NPS-specific failsafes. The operator utilizes the NPS-developed Swarm Commander application to parameterize and initiate tactics and monitor fleet performance.

The ARSENL communications architecture, depicted in Figure 3.4, relies on an 802.11n ad hoc network for both ground station-to-UAV and UAV-to-UAV communication using an NPS-designed application layer message protocol. All messages are transmitted via User Datagram Protocol (UDP) broadcast. The use of UDP intentionally prioritizes low latency over high reliability and necessitates robust behaviors that do not rely on reliable message delivery [3], [25].

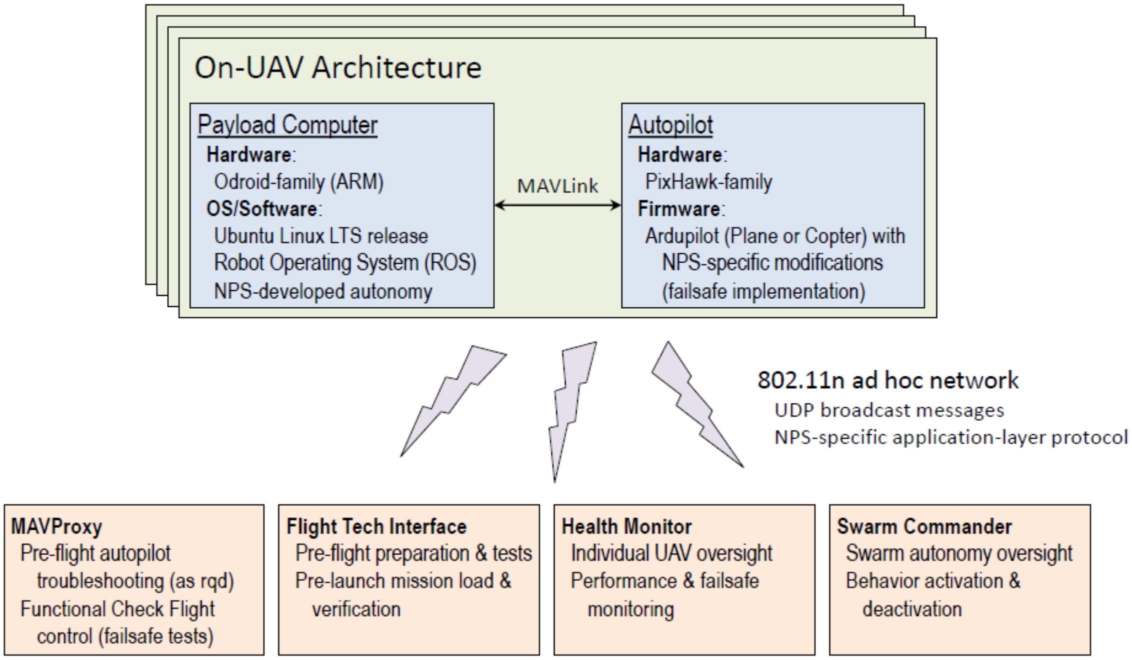

Figure 3.4. The ARSENL multi-UAV architecture [14].

Synchronous messages transmitted by ARSENL vehicles include status messages containing autopilot and system information and state messages containing position, velocity, and behavior-specific information. These messages are transmitted at two and ten hertz, respectively. A heartbeat message is transmitted from the Swarm Commander application at a rate of one hertz to ensure satisfactory ground-to-air communications (a failsafe is initiated if communications are interrupted for more than two minutes). This approach results in an "eventually right protocol" that enables a loss-tolerant system [3].

Asynchronous messages from the Swarm Commander application to airborne UAVs are used to initiate parameterized tactics. Asynchronous messages are also utilized for UAV-to-UAV communication as tactics are executed to facilitate coordination, subtask assignment, and fleet-wide situational awareness.

3.4 Previous Work: ARSENL

In autonomous multi-vehicle systems, determining optimal role assignments for individual vehicles is inherently difficult. This is due in part to the communications-limited environments in which multi-vehicle systems operate. Even in the best of circumstances, however, role assignment is subject to the computational complexity of assignment problems. The generalized assignment problem (i.e., the number of roles and vehicles differ and each vehicle can be assigned zero or more roles), for instance, is applicable to many distributed decision processes and has been shown to be NP-hard [26]. Solutions for finding near-optimal solutions to the generalized assignment problem have been demonstrated with worst-case performance of $O(n^2 \log n)$ and near-linear performance in many situations [26]. Nevertheless, the problem remains daunting for a system of widely distributed and communications limited vehicles.

One approach to these problems that has proven applicable in lossy-communications environments the use of auction algorithms. ARSENL has demonstrated two auction algorithms that have been effective in role assignment within their system. In his 2018 thesis, Major Hopchak demonstrated and evaluated a single item auction for assigning search cells to individual UAVs tasked with a large area search [27]. His findings concluded that auctions are inherently useful in "restrictive, lossy-communications environments of various size and complexity" [27]. Lieutenant Britt Campbell extended this work to account for heterogeneous emergent tasks in the form of search-identified contacts requiring further investigation [28]. More recently, ARSENL has demonstrated a multi-item auction more suitable to the generalized assignment problem than single-item auctions. The multi-item auction utilizes a Prim Allocation approach that requires exactly n auction rounds to assign n roles to an arbitrary number of agents [29].

Both auction algorithms are executed as a series of rounds, each of which consists of the same basic sequence: bid submission, winner determination, role or task assignment. At the beginning of each round, each agent submits a bid for a single task or role. In the single-item auction, roles are tentatively assigned to the high bidders, but final assignments are not made until all agents have submitted a high bid for their "desired" role. The Prim auction makes a final assignment of exactly one role to an agent in each round. The individual focus of both algorithms mitigates the need for a broad consensus and allows each agent to determine for itself the roles for which it is most suited [27]. Further, they allow a system of heterogeneous agents to allocate roles effectively with little or no shared understanding or knowledge of individual vehicle capabilities [28].

ARSENL's experience with the use of auction algorithms in their multi-UAV systems was particularly relevant to this research. As described, the base defense mission entails numerous roles that occur simultaneously or arise unpredictably. Effective assignment of those roles to suitable UAV agents was critical to the success of the autonomous swarm in this mission. As will be discussed in Chapter 4, single-item auctions were typically used when the number of roles or tasks was equal to or less than the number of available UAVs, and the multi-item Prim auction was used when the number of roles or tasks exceeded the number of UAVs.

In addition, previous ARSENL research into reactive and pheromone-based behaviors was leveraged during this work. In his 2017 thesis, Captain Andrew Hietpas developed a pheromone-based behavior in which UAVs were required to serve as data ferries by regularly "marking on top" of communications-limited ground stations. Upon each overflight of a ground station, the UAV would transmit any messages destined for that station and receive any messages transmitted by that station. Hietpas devised a method for dynamically controlling UAV transit speed using digital pheromones to achieve uniform visit intervals [30]. The requirement to maintain even UAV spacing over a repeating circuit proved analogous to the requirements noted in Chapter 2 for perimeter surveillance. The approach was therefore utilized in the `Perimeter Surveillance` play described in Chapter 4.

3.5 Summary

In summary, much research suggests that the complex problem of multi-vehicle control can be simplified through the use of hierarchical behaviors. From the initial work of Miller to the implementation of MASC through ARSENL, breaking down complex behaviors into their atomic parts enables both modularity and utility. Furthermore, MASC's top-down approach to swarm control enables the creation of doctrine-compliant behaviors by way of developing mission requirements and fulfilling them through combining individual behaviors. This approach provides an inherently useful method for rapid creation and adaptation of complex behaviors to the ever changing tactical landscape.

Chapter 4 describes the `Base Defense` tactic in detail. In particular, the application of the MASC approach is discussed through the description of the algorithms and plays, many of which were drawn from or adapted from existing ARSENL capabilities, of which the tactic is comprised. The testing approach, MOEs, and MOPs that were utilized to assess the tactic's final implementation are also discussed.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 4:
Tactic Design and Specifications

This chapter describes the design of the MASC tactic and plays that were developed to support the base defense mission. In addition, the reasoning behind algorithmic and control decisions is discussed. Each tactic and play is implemented as a state machine whose depiction shows the transitions and reactions associated with various triggering events that occur over the course of the mission. All plays explained in this section were specifically developed for use with the `Base Defense` tactic with the exception of the `Pounce Drop` play which was re-used from the existing ARSENL play library. Plays were developed in such a way as to be useful in tactics and missions that may be developed in the future.

The design for the `Base Defense` tactic centered around doctrinally-defined essential tasks. The two essential tasks for base defense were determined to be base surveillance and base response (described in Section 2.1). As implemented, the `Base Defense` tactic serves as a proof of concept for the MASC approach to complex swarm behavior development.

The Base Surveillance and Base Response tasks are summarized as follows.

- Base Surveillance: The primary purpose of this task is to monitor the perimeter and key areas surrounding the base and in turn to alert the occupying force to enemy activity. This task must be executed constantly regardless of the enemy's posture. This task remains the priority until a threshold is reached at which the base requires intervention at a higher level. This threshold will depend on the situation and posture of each base and thus should be thought of as variable.
- Base Response: The purpose of this task is to repel the enemy's assault by fire and close combat. This task's initiation is always in response to enemy activity, therefore the severity of the threat determines the level of response. This task is augmented through observation for fire control measures and direct action by the vehicles.

In other words, the Base Surveillance task is steady state while the Base Response task supersedes surveillance with kinetic tasks that are augmented by ongoing surveillance and contact tracking until the threat has been eliminated. The following is an explanation of the tactic and plays that were developed and used during this research.

4.1 The Base Defense Tactic

Tactics serve as intermediary behaviors to be used in the conduct of phases as part of a larger mission. Within the MASC hierarchy, multiple plays and algorithms are typically combined to comprise a single tactic. Tactics can also be embedded within other tactics to foster increasing capability with a limited increase in overall complexity. Tactics are designed for use at the operational level of system employment. From a practical standpoint, this means that within the MASC hierarchy tactics serve as the lowest level of operator control [12].

The `Base Defense` tactic is designed for defense of a base with an established perimeter in accordance with current doctrine as outlined in Chapter 2. The tactic is comprised of four plays, each of which is associated with a different tactical role. The top-level state transitions are depicted in Figure 4.1.

The objective of the `Base Defense` tactic is to detect enemy presence, track contacts, and repel verified threats through fire and observation. The `Base Defense` tactic accepts the following parameters: a list of key area coordinates in the form of latitude-longitude tuples, an ordered list of coordinates for the vertices of the base perimeter, and a list of contacts. A list of contacts is used upon initialization to maintain the state of the tactic when the behavior is restarted, which occurs after every contact generation. Each contact is stored as an object containing the contact identifier, the status of the contact, and the location of the contact.

The `Base Defense` tactic begins in the role assignment state. During this state, a single item auction is used to assign tasks for each individual vehicle. The tasks are specified as one of the following four individual plays: `Perimeter Surveillance`, `Key Area Search`, `Contact Investigation`, and `Threat Response`. Once the task is determined, the vehicles execute that task until either a contact is reported or updated by the swarm or a contact is detected by the individual vehicle. The vehicles conducting the `Perimeter Surveillance` and `Key Area Search` tasks report possible contacts to the rest of the swarm. The vehicles conducting the `Contact Investigation` tasks investigate possible contacts to determine the contact status, **Contact Confirmed** or **Contact Denied**, and send a contact report update throughout the swarm. The vehicles conducting the `Threat`

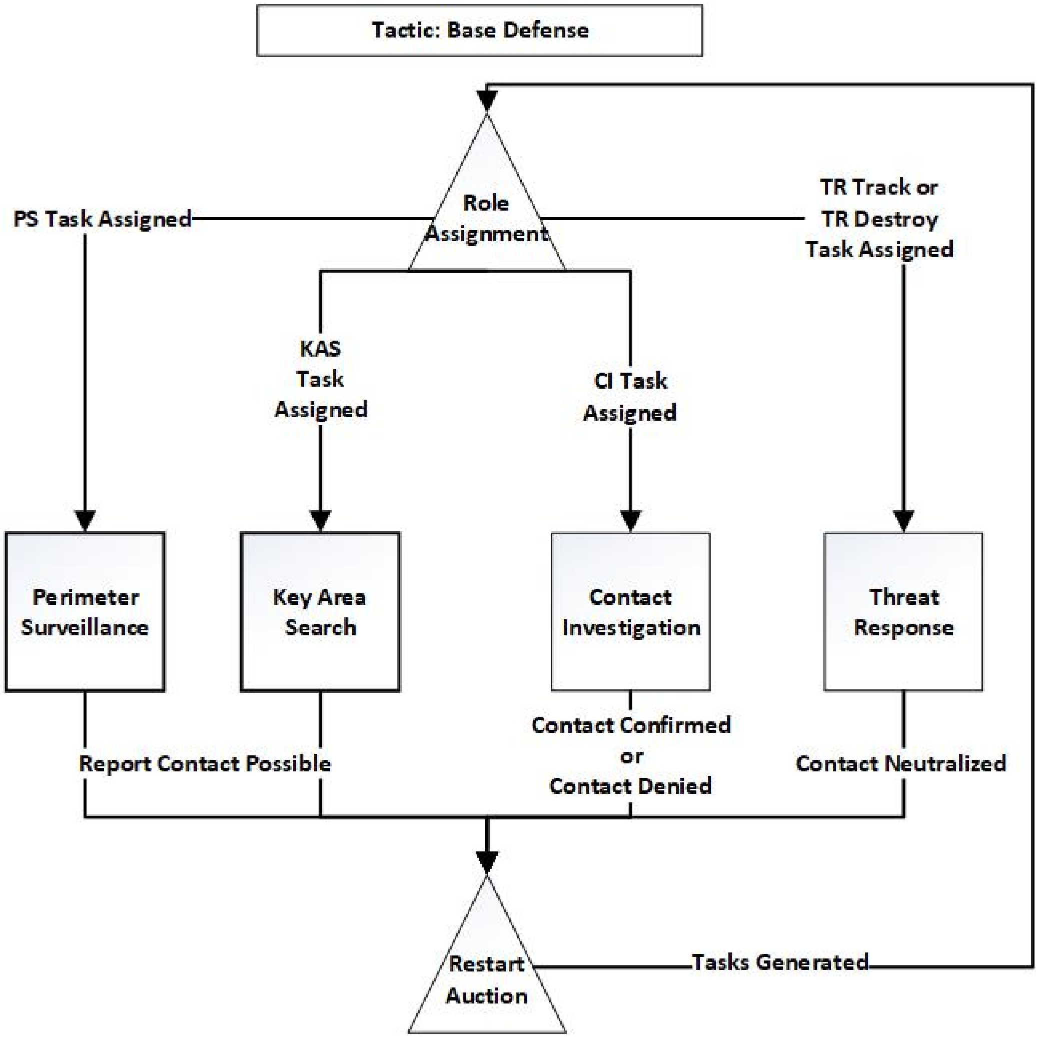

Figure 4.1. The `Base Defense` Tactic is composed of four plays and two decision points that govern reactions to the presence of contacts within the Area of Operations.

`Response` task neutralize the **Confirmed Contact** through observation of fires or direct action and then send a contact report message throughout the swarm. When a contact is observed or an update to a contact is determined, an **Auction Restart** message is sent to initiate role assignment for the new tasks.

The `Base Defense` tactic utilizes a single-item auction for assignment of surveillance, investigation, tracking, and destruction tasks to achieve a complete defensive posture. Tasks for the auction are determined through a deterministic algorithm based on the number of

UAVs and the tactic parameters to ensure appropriate coverage of the required roles. The critical points within this behavior include task generation, task evaluation, and synchronization of the contact list across the swarm.

Single item auctions are simplified when the number of tasks equals the number of participants (i.e., when they are used to solve the basic assignment problem). The generation of tasks for the `Base Defense` tactic, therefore, starts with the number of participants and generates exactly that many tasks. Tasks are generated according to the following priority:

1. Minimum surveillance requirements
2. Confirmed contact tracking
3. Confirmed contact destruction
4. Possible contact investigation
5. Remaining surveillance

The first tasks generated accomplish the minimum acceptable surveillance requirement, ensuring that at least one participant is assigned to perimeter surveillance and one to key area search. This ensures that as long as two vehicles are assigned to the tactic, all surveillance areas will be at least minimally monitored. After minimum surveillance tasks are generated, confirmed contact tracking tasks are generated. For each confirmed contact within the contact list a threat response task with an observation role as described in Section 4.2.4 is generated. The next priority is destruction of confirmed contacts designated for engagement. A threat response task with a destruction role is generated for each of these. This enables teaming as discussed further in Section 4.2.4. Once these tasks have been generated, contact investigation tasks are generated for all possible contacts requiring confirmation. Finally, tasks are generated alternatively for key area search and perimeter surveillance until exactly one task is available for each UAV.

Task evaluation is conducted by each vehicle individually upon initialization of the auction, and tasks are re-evaluated each time the auction is restarted. This enables the vehicles to bid for tasks based on their most recent proximity to the task locations (i.e., it more accurately accounts for the transit time cost for each task).

Vehicles evaluate each task as a function of anticipated transit time to the task location and a task-type multiplier that facilitates preference for tasks to which the vehicle is more suited. Investigation and tracking task evaluations also take into account whether or not the evaluating vehicle was the original observer of the contact (i.e., the initial observer is prioritized in the assignment of these tasks). The formulas for vehicle-specific evaluation of investigation, threat response for tracking, perimeter surveillance, and threat response for destruction are shown as Equations 4.1 through 4.4, respectively, where $t_{transit}$ is expected transit time, $K_{transit}$ is a scaling constant, $K_{observer}$ is a constant associated with a contact's initial observer, and C is a vehicle-type-specific constant indicating preference or aversion for a particular type of task.

$$\beta_{investigate} = \frac{K_{transit}}{t_{transit}} \times C_{investigate} \times K_{observer} \tag{4.1}$$

$$\beta_{track} = \frac{K_{transit}}{t_{transit}} \times C_{track} \times K_{observer} \tag{4.2}$$

$$\beta_{perimeter} = \frac{K_{transit}}{t_{transit}} \times C_{perimeter} \tag{4.3}$$

$$\beta_{destroy} = \frac{K_{transit}}{t_{transit}} \times C_{destroy} \tag{4.4}$$

These equations take into account the anticipated time required to transit to the objective and vehicle-suitability for the task so that better suited vehicles with shorter transit times will bid more aggressively. The value or benefit for key area search tasks are set at a fixed base value due to the difficulty in quantifying transit time or task difficulty for the potentially wide key area disposition (this is discussed in more detail in Section 4.2.2). This benefit calculation ensures that the tasks are evaluated according to the priorities previously discussed.

4.2 Plays

As discussed in Chapter 3, plays in the MASC architecture are developed as combinations of one or more low level algorithms to accomplish a singular objective. Plays are the lowest level of multi-vehicle control in the MASC hierarchy. Each tactic uses a combination of plays to realize the desired behavior. The strength of this hierarchy is the ability to aggregate simple behaviors to create more robust behaviors. The `Base Defense` tactic is composed of five plays: `Perimeter Surveillance`, `Key Area Search`, `Contact Investigation`, `Threat Response`, and `Pounce Drop`. The first four of these plays were created during the course of this research specifically in support of the `Base Defense` tactic. The fifth play, `Pounce Drop`, was developed during previous research but integrated due to the modularity of the MASC hierarchy.

The primary surveillance plays supporting base defense are `Perimeter Surveillance` play and `Key Area Search`. The `Contact Investigation` and `Threat Response` plays provide reaction to potential and confirmed contacts. The `Pounce Drop` play is implemented as a sub-play within the `Threat Response` play. Each of these is discussed in detail in the following sections.

4.2.1 Perimeter Surveillance

The `Perimeter Surveillance` play provides continual surveillance of the area immediately surrounding a base's perimeter. The overall goal of this play is to provide complete surveillance of the perimeter while minimizing the time that any given sector of the perimeter is not observed. The input parameter for this play is a single ordered list of latitude and longitude tuples that represent the vertices of the base perimeter. The state transitions for this play can be seen in Figure 4.2.

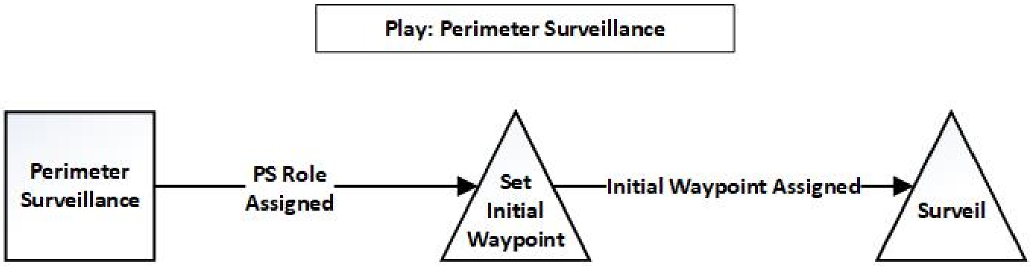

Figure 4.2. The `Perimeter Surveillance` play provides surveillance along the perimeter of any polygonal shape and uses a digital-pheromone-based spreading algorithm to evenly space the vehicles along this perimeter.

The `Perimeter Surveillance` play follows a three state diagram. When the `Perimeter Surveillance` role is assigned, the *Set Initial Waypoint* state is entered. During this state, surveillance speed and pheromone spacing are determined, and an unique initial waypoint is set. Once the initial waypoint is assigned, the *Surveil* state is entered. The vehicle remains in this state until termination of the behavior.

The surveil state of the `Perimeter Surveillance` play is implemented through the use of the ARSENL's repeating waypoint sequence driver algorithm. This driver directs the vehicle through a series of waypoints, the vertices of the base perimeter in this case, and then resets to the first waypoint upon reaching the last waypoint. Following initialization, the algorithm will direct the vehicle's continuous transit around the base perimeter until the behavior is terminated.

A shortcoming that was discovered during initial testing of the play was that the vehicles remained in close interval along the perimeter, leaving large gaps in surveillance. Ideally, assigned vehicles will be evenly spaced around the perimeter in order to minimize these surveillance gaps. A digital pheromone approach similar to that explored in [30] was utilized to mitigate this issue.

Digital pheromones mimic the chemical pheromones utilized by insects and other animals to provide indicators to others of where they have been. Pheromones, both chemical and digital, degrade over time, and it is this characteristic that is leveraged by the `Perimeter Surveillance` play. During parameterization, vehicles calculate an optimal revisit interval based on the number of participating vehicles, the length of the perimeter, and a "nominal" speed at which the circuit is to be conducted. Digital pheromones are placed by each vehicle as it passes evenly spaced points along the perimeter. Pheromones are initially set to a maximum strength value and decay linearly according to Equation 4.5 where s_t is the strength of the pheromone at time t, S_{max} is the maximum pheromone strength, t_0 is the time at which the pheromone was last set, and $T_{optimal}$ is the desired revisit interval. Note that the pheromone value will be positive or negative, respectively, if the vehicle arrives at the pheromone point earlier or later than desired.

$$s_t = S_{max} - \frac{t - t_0}{T_{optimal}} \tag{4.5}$$

When a vehicle passes over a pheromone point, the speed command, v_{order}, is updated using Equation 4.6 if it arrived early or 4.7 if it arrived late. In both equations, s_t and S_{max} are from Equation 4.5 and V_{nom}, V_{max}, and V_{min} are the nominal, maximum and minimum commandable speeds respectively. The minimum and maximum commandable speeds are based on the inherent vehicle capabilities as determined by the controller and aircraft type.

The nominal speed is set upon initialization, and was set to 17 meters per second for this work (this value is the typical base speed for the Zephyr II in ARSENL experiments). Evidently, if the vehicle reaches the pheromone point early, the pheromone strength will be positive and the vehicle's speed will be decreased, and the negative pheromone value will cause the vehicle's speed to be increased if it arrives late.

$$v_{order} = V_{nom} - \frac{s_t}{S_{max}} \times (V_{max} - V_{nom}) \tag{4.6}$$

$$v_{order} = V_{nom} - \frac{s_t}{S_{max}} \times (V_{nom} - V_{min}) \tag{4.7}$$

In addition to updating the ordered speed, a reset message is sent to other participating vehicles upon pheromone point passage to direct the reset of the pheromone to the maximum strength value. This minimizes the overall communication required to synchronize movement and maximizes spread over time [30].

4.2.2 Key Area Search

The `Key Area Search` play provides surveillance over key areas as described in Chapter 2. The parameters for initialization of the `Key Area Search` play include a single list of latitude and longitude tuples representing the key area locations to be surveilled. The state transitions for the `Key Area Search` are depicted in Figure 4.3. The play utilizes ARSENL's reversing waypoint driver algorithm which is similar to the repeating waypoint driver algorithm utilized in the `Perimeter Surveillance` play. The key difference between these algorithms is that the reversing waypoint driver algorithm reverses the order of waypoints upon reaching the final waypoint rather than restarting the sequence from the initial waypoint. This algorithm, therefore, directs the vehicle back and forth through the waypoint sequence. It is tailored for lists of waypoints that are to be transited linearly rather than in a circuit and is intended to minimize wasted transit time returning to the initial waypoint unnecessarily.

The `Key Area Search` play follows a three state diagram. When the `Key Area Search` role is assigned the play enters the *Key Area Assignment* state. During this state, a multi-item auction is conducted to assign key areas to individual vehicles. Once all key areas have been assigned and the auction concludes, the *KAS Search* state is entered. During this state,

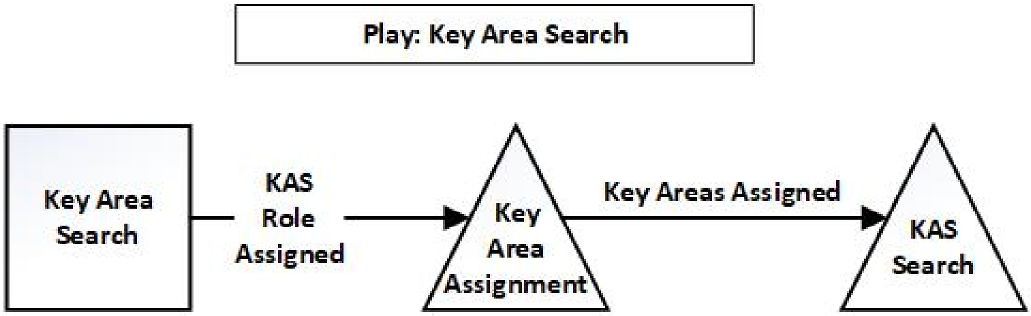

Figure 4.3. The Key Area Search play accomplishes the task of providing surveillance over multiple key areas spread across the area of operations. This play uses a multi-item auction for assignment of key areas across the vehicles within the play.

the vehicle transits to each key area that it was assigned until the behavior is terminated.

Two optimizations were developed for this play to minimize the transit time and in turn maximize time spent surveilling assigned key areas. The first is associated with the assignment of tasks while the second is associated with transit between assigned key areas. Both are based on initial test observations.

The assignment of an arbitrary number of key areas to an arbitrary number of UAVs is far from trivial (i.e., it reduces to the NP-hard general assignment problem). The first approach to this problem utilized a variation of the single-item auction that yielded clearly sub-optimal results. Significantly better results were obtained through the use of a Prim allocation auction based on graph spanning forests. Utilization of a Prim auction facilitates key area assignments in linear time with a guaranteed relationship to the optimal assignment [29].

The Prim auction was implemented using ARSENL's code base for multi-item auctions. Implementation required two essential methods to be developed: one to generate tasks and a second to calculate the costs of the tasks for each vehicle. A separate task is generated for each key area, and each task's cost is evaluated based on the estimated transit time to the key area from its currently assigned tasks.

At auction commencement each UAV has no assignments. This is represented locally by a root-only assignment tree containing only the UAV's starting state as the root. In each round, every UAV identifies the unassigned task that can be added to its assignment tree at the lowest cost as calculated by Equation 4.8 where C_{key_area} is a fixed cost associated with each area, $t_{transit}$ the shortest transit time from a node in the local assignment tree, and d_{task0} is a discount factor (discussed below). After identification of the lowest cost task, each vehicle submits a bid for that task at the calculated cost. The globally low bidder is assigned the task for which it bid. Thus, a single task is assigned to a UAV in every round, and the auction will complete after a number of rounds equal to the number of key area tasks to be assigned. Once the auction is complete, each UAV uses a depth-first traversal of its assignment tree to generate a key area visit order for use in the reversing waypoint driver algorithm.

$$\beta = C_{key\_area} \times t_{transit} \times d_{task0} \tag{4.8}$$

To encourage assignment of at least one task to each UAV, a discount of 0.5 is applied to the bid of every UAV that does not have a key area task in its assignment tree. Further, the use of transit time rather than proximity to the key area, the individual vehicle's capabilities encouraged assignment of closer key areas to the slower quadrotor vehicles. [29]

The second `Key Area Search` play optimization was based on the assumption that time in the vicinity of a key area is well spent, while time transiting between key areas is not. Transit and surveillance speeds are therefore set based on proximity to a key area. When a vehicle crosses within an observation threshold, the vehicle slows to a surveillance speed near its minimum speed. Once the vehicle is no longer within observation distance, the faster transit speed is ordered.

Both of the play's optimizations are intended to minimize transit time and maximize observation time. The Prim auction does this by assigning slower UAVs (e.g., Mosquito Hawks) to fewer, nearer, and more closely spaced key areas while faster UAVs (e.g., Zephyrs) to more, farther, and more widely dispersed ones. Variable speed, on the other hand, maximizes time spent on top of assigned key areas relative to time spent in transit.

4.2.3 Contact Investigation

The `Contact Investigation` play was the first of the response plays created for the `Base Defense` tactic. The objective of this play is to transit to a possible contact's location following initial detection to determine the contact's status. The parameter for this play is a record object containing the contact's ID, location, and current status. The state transitions for this play are depicted in Figure 4.4.

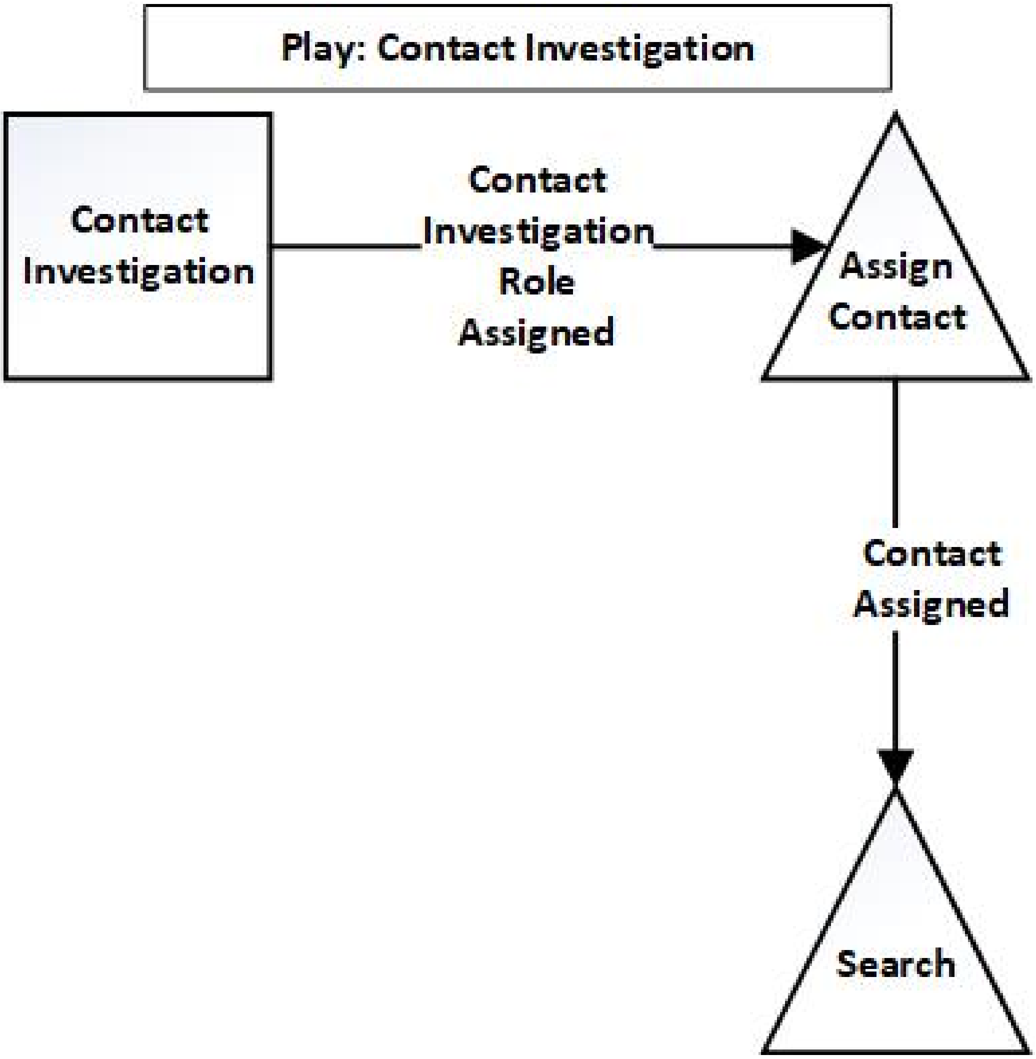

Figure 4.4. Contact Investigation is a straight forward behavior that directs transit to the location of a possible contact for investigation. Following the investigation, the updated contact status.

The `Contact Investigation` play follows a three state diagram. Once the `Contact Investigation` role is assigned, the *Assign Contact* state is entered. During this state an initial waypoint is set for the contact's location. Once the location is set, the vehicle enters the *Search* state, wherein the vehicle transits via the most direct route to the possible contact location. Once there, the vehicle maintains position over the contact until it is either confirmed or denied. Following the determination, a contact update message is transmitted to the rest of the swarm.

Rather than utilizing onboard sensors, the evaluation of which was beyond the scope of this work, contact confirmation was randomly determined in experiments. In real-world implementation this would likely be a situation in which the swarm operator would determine the contact's status.

4.2.4 Threat Response

The `Threat Response` play is the final play that was created for the `Base Defense` tactic. The objective of this play is to direct the response to confirmed threats within the Area of Operations by tracking contact movement and possibly neutralizing the contact utilizing secondary plays. This play utilizes two parameters: the contact record containing the contact ID, status, and location and the assigned role.

The `Threat Response` play follows a five state diagram. Upon assignment of the `Threat Response` role, the last known location of the contact is set for the initial waypoint of the vehicle. When the track role is assigned, the vehicle enters the *Track* state and remains in that state until the behavior is terminated. When the vehicle is assigned the *track* role, it maintains position on top of the contact and sends updates on the contact location to the rest of the swarm whenever movement is detected. When the *Destroy* role is assigned, the *Destroy* state is entered. Once a contact location update has been confirmed, the vehicle initiates the embedded `Pounce Drop` play to control the contact engagement.

Based on the task assignment priorities discussed in Section 4.1, two vehicles will be assigned to each threat if possible: one in the *track* role and one in the *destroy* role. While either role can be executed individually, the play was intentionally designed to take advantage of teaming. When a vehicle is assigned to the *track* role in support of a vehicle assigned to the *destroy* role, continuous surveillance and movement updates are provided to increase accuracy of the engagement.

In the spirit of the MASC, this play incorporates the `Pounce Drop` play from the existing ARSENL behavior library. This enabled faster implementation through abstraction and

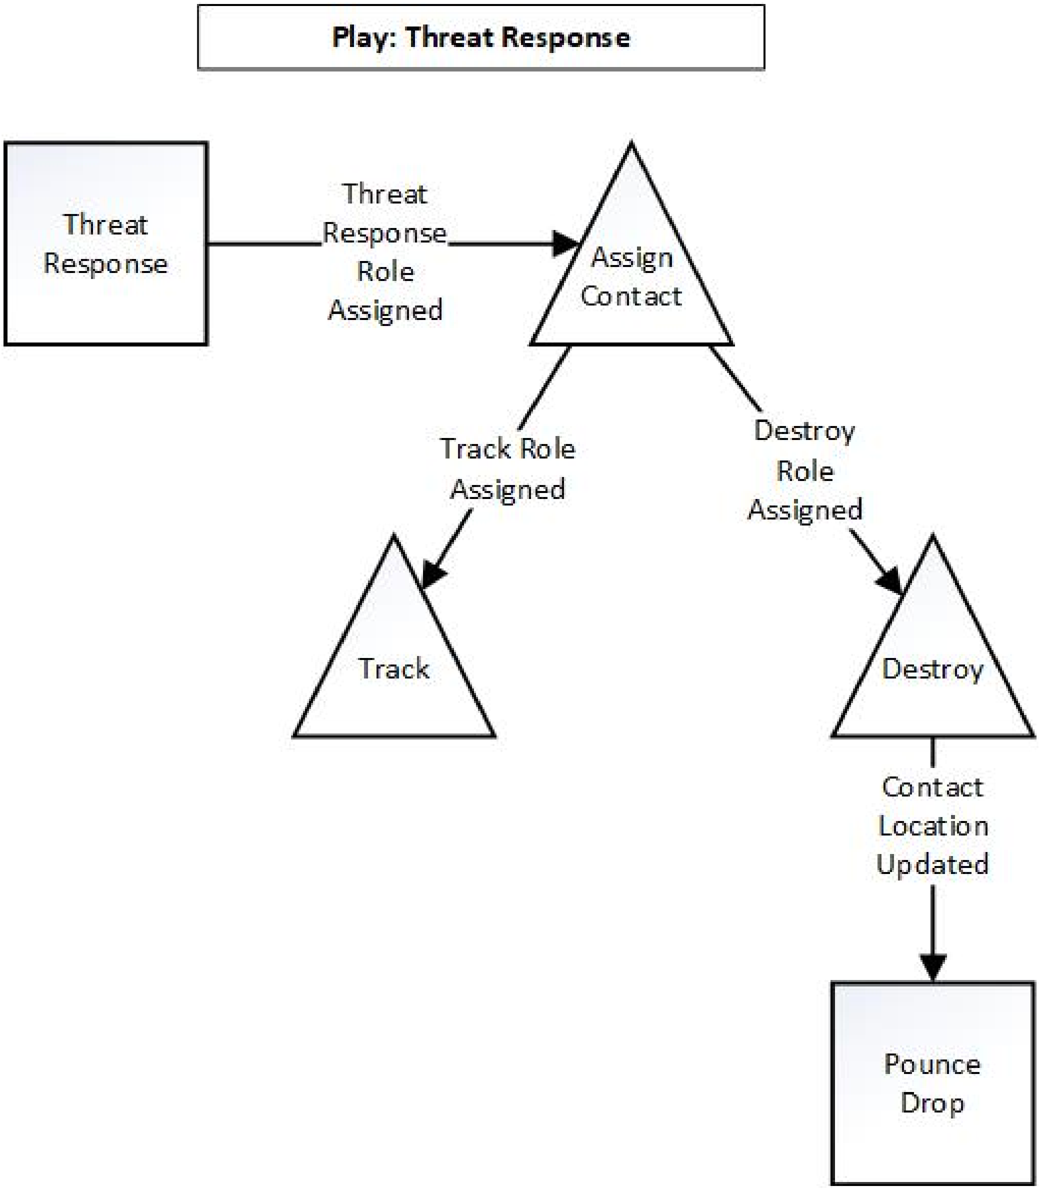

Figure 4.5. The Threat Response play accomplishes tracking or destruction of a confirmed threat. If set to a tracking role, continuous contact with the threat will be maintained. In a destroy role, the contact will be engaged through an air-to-ground drop.

code reuse. Also, this approach provides additional flexibility in mission design as more engagement-oriented plays are developed.

4.2.5 Pounce Drop

The final play utilized in the `Base Defense` tactic is the `Pounce Drop` play. The pounce drop play was developed by ARSENL and incorporated into the repository to support tactics requiring the immediate expenditure of air-to-ground expendables at a designated location [14]. The play directs the vehicle to a specific location and releases a single expendable store. Within the `Base Defense` tactic, the expendable is assumed to be ordnance, but the `Pounce Drop` play is defined more generally to support various mission requirements. This play requires a single parameter in the form of a latitude and longitude tuple specifying the location of the stores drop.

When initiated this behavior directs the vehicle along the most direct path to the specified drop location. The play utilizes the direct transit driver algorithm to provide waypoint planning along the most direct path. Upon arrival at the drop location, a single expendable store is released. The play maintains position at the drop location following release until termination or retasking.

4.3 Tactic and Play Development

Tactic and play development was conducted using the ARSENL SITL simulation environment, an extension of the open source ArduPilot capability [24]. The SITL environment provides a realistic environment in which actual on-vehicle software can be exhaustively tested [31]. The ARSENL implementation allows an arbitrary number of UAVs to interact in the same simulated environment making the system well-suited for development and testing of swarm algorithms. Further, exhaustive testing in the SITL environment facilitates more efficient live-fly experimentation through early identification of algorithmic flaws and bottlenecks.

Since the SITL environment provides for testing under best case circumstances, it served as an effective proving ground for individual plays. Plays could be quickly prototyped and an iterative development used to incorporate changes and add functionality. Each play went through at least two versions: initial and improved. The details of each are explained in Chapter 5.

The SITL environment facilitated a similar approach to tactic development. Tactic development was approached by first ensuring correct functionality according to the state diagram and proper assignment of tasks to vehicles and then by incrementally incorporating and improving plays.

Initial tactic and play implementations often showed ineffective state transitions, suboptimal assignments, and suboptimal task performance (e.g., the inefficient perimeter surveillance spacing discussed in Section 4.2.1). Improved versions were created after significant SITL testing was conducted to characterize the play or tactic's performance.

4.4 Testing Approach

Each play was initially tested independently in the SITL environment with a homogeneous swarm of three Zephyr II fixed-wing UAVs. Once each play's state transitions and tasking decisions were verified for the small homogeneous swarm, Mosquito Hawk quadrotors were added to verify that tasking decisions correctly accounted for dissimilar aircraft capabilities. Finally, the size of the swarm was increased to ten UAVs with various combinations of fixed-wing and quadrotor aircraft. This iterative approach enabled agile development of plays and allowed for testing of multiple state transition approaches to identify those yielding the best results.

Despite its utility, the SITL environment does not accurately capture a number of key real-world variables, network performance, weather effects, and aerial maneuvering constraints for example. Therefore, live-fly experimentation was also required to fully demonstrate successful tactic implementation and characterize its performance.

In order to demonstrate this tactic on the ARSENL vehicles (which do not have target detection sensors), sensors were simulated using a contact generator class. The contact generator class was used by the `Perimeter Surveillance` and `Key Area Search` plays to randomly generate contacts at a rate of approximately three contacts per minute.

The higher-than-realistic contact generation rate facilitated testing by ensuring that all tactic and play states would reached all states could be reached in the limited live flight time. Further, the use of simulated contact generation does not constitute a shortcoming of this research since the objective was to demonstrate the functionality. It is understood that sensors must be characterized and accounted for in any real-world implementation. This, however, is left to future work for now.

4.4.1 Measures of Effectiveness (MOEs) and Measures of Performance (MOPs)

MOPs and MOEs are utilized to describe how well an implementation meets fundamental requirements specified by the capabilities development documentation. MOPs are defined by the Technical Measurement Guide as "the physical or functional attributes relating to the system operation" [32]. Stated differently, MOPs measure how well the essential tasks that must be accomplished in order to produce a functional system are performed. MOEs, on the other hand, are defined as "measures of success that are closely related to the achievement of mission operational objectives" [32]. MOEs, then, measure the extent to which the larger objectives are accomplished. Thought of together, MOPs assess how well the system is doing the things that it was designed to do while MOEs assess how well the system accomplishes its objective. In other words, MOPs support MOEs. The remainder of this section describes the MOPs and MOEs that were used to evaluate the `Base Defense` tactic.

Evaluation of the `Base Defense` tactic focused on effectiveness and efficiency, that is, how well did it accomplish the mission and how efficiently were the objectives achieved. The most important MOEs for the `Base Defense` tactic in this context are *Detection* and *Efficiency of Force*.

The *Detection* MOE describes the effectiveness of the tactic in successfully detecting a contact within the area of operations. The MOP associated with this MOE is the *elapsed time from contact observation to new task assignment throughout the swarm*. This is measured as the time it takes the swarm to detect a contact and transition from the current state, through the auction states, and complete task assignment.

The goal of the *Efficiency of Force* MOE is to maximize the overall productivity of the swarm. Productivity, in this sense, is the engagement of the swarm in tasks that directly support base defense. The *Efficiency of Force* MOE is assessed through MOPs measuring the proportion of time spent by each vehicle in the swarm in active states compared to loss states. Active states are those states in which active detection, delaying, or destruction of the contact is being conducted. Conversely, loss states are those that do not actively engage in those activities (e.g., transit time between key areas, states in which an assignment auction is being conducted, and initializing states).

Table 4.1. This table provides an overview of the associated MOE and MOP pairs used for evaluation of the `Base Defense` tactic.

MOE	MOP	Evaluation Metric
Efficiency of Force	Proportion of Time Spent in active versus loss states	$\frac{\sum ActiveStates}{\sum ActiveStates+\sum LossStates}$
Detection	Time to react to detection of contact	Elapsed time from contact observation to new task assignment throughout the swarm

4.5 Summary

The `Base Defense` tactic and the plays from which it is comprised, `Perimeter Surveillance`, `Key Area Search`, `Contact Investigation`, `Threat Response`, were designed to accomplish the mission of base defense as described in this chapter. The `Base Defense` tactic and `Key Area Search` play utilize auctions to assign a single or multiple tasks respectively. The generation and evaluation of the associated tasks, therefore, proved to be a critical point in the development of these behaviors. Using an iterative design process, they were the focus of many optimization efforts.

As discussed in Section 4.1, `Base Defense` tactic utilizes a single item auction to allocate tasks and generates these tasks based on the number of participating UAVs, the base perimeter, number of key areas, and the number and status contacts. This recursive implementation allows the swarm to respond to contacts as they are detected with minimal loss of time or efficiency.

The surveillance plays, `Perimeter Surveillance` and `Key Area Search`, were the first behaviors developed during the research. They provided baseline functionality that ensured continuous surveillance of the base perimeter and designated key areas. The implementation of these plays was optimized through multiple iterations of testing in both the SITL simulation environment and during live flight testing at Camp Roberts, California.

The response plays, `Contact Investigation` and `Threat Response` provide the capability to respond to contacts generated by the surveillance assets. Once a contact is generated and detected, the detecting vehicle sends a contact report to the rest of the swarm, and this prompts a restart of the task assignment auction. This approach enables the swarm as a whole to respond to individual contacts in a manner that accounts for individual vehicle strengths. The `Contact Investigation` play directs a vehicle to the point at which a possible contact was detected so that it can make a determination of whether to confirm or deny the contact. If the contact is confirmed, a threat response task is generated, and a new tasking auction is initiated.

The `Threat Response` play combines two roles: *track* and *destroy*. If the *track* role is set, the vehicle transits to and follows a moving contact and reports its position at intervals. If the *destroy* role was set, the vehicle maneuver is governed by an embedded `Pounce Drop` play that directs the vehicle to the contact's location to deploy a single expendable store (presumably ordnance) against the target.

Finally, this chapter describes two MOEs that were used to evaluate the `Base Defense` tactic: *Detection* and *Efficiency of Force*. These MOEs are supported by a number of MOPs which were also described in this chapter. These MOEs and MOPs were used through multiple tests in the SITL simulation environment and live flight experiments. The results of these tests are discussed in Chapter 5.

CHAPTER 5:
Test Design and Results

This chapter summarizes SITL environment and live flight experiments conducted during ARSENL field exercises. Both field exercises were conducted at McMillan Field at Camp Roberts, California and provided opportunities for the ARSENL team to test and evaluate new behaviors and changes to previously implemented behaviors. The SITL environment used in simulation experiments was also aligned with the McMillan Field operating area.

5.1 Live Flights

ARSENL field exercises presented opportunities to test and refine the design of the `Base Defense` tactic. The first event supporting this research took place in February 2020. Initial implementations of the `Perimeter Surveillance` and `Key Area Search` plays were tested during this event.

The February 2020 experiment served as a crucial introduction to the effects that weather would have on the control of the vehicles and their associated algorithms. Each play was initiated from a standby position once all vehicles were airborne and in the ready state. The plays were tested with both homogeneous and heterogeneous swarms and tested in isolation and also in tandem with subswarms executing both plays simultaneously.

The `Perimeter Surveillance` play performed successfully during the first live flight tests. The initial requirement of this play was to ensure that all surveillance targets along the perimeter of the base were reached and that the vehicles followed the perimeter. The vehicles met these goals. Inter-vehicle spacing, however, was irregular and the repeat time between surveillance targets was clearly not optimal. In particular, multiple vehicles overflew each point in quick succession followed by long periods with no overflights. This observation led to the implementation of the pheromone-based spacing approach that was discussed in Section 4.2.1.

The `Key Area Search` also met the minimum requirements for initial implementation in that each vehicle successfully reached each key area. The primary takeaways from the experiment, however, were that waypoint sequencing was inefficient and that vehicles spent the majority of the time in transit rather than in observation of a key area. To improve the observed assignment and sequencing shortcoming, the Prim auction approach was implemented. To further improve the ratio of on-top versus in-transit time, speed adjustments were implemented. Both of these improvements were discussed in detail in Section 4.2.2.

The second ARSENL field exercise took place in July 2020. The entire `Base Defense` tactic was tested during this field exercise. The primary tasks of allocation, surveillance and response were successful overall, however, adaptations to the live environment were required before the tactic would execute to completion.

During the initial flight, nine Zephyr IIs and one Mosquito hawk were utilized. Initial tasking and role assignment via auction were successful, however the geographic positions of the key areas were outside of the standard operating area at McMillan Airfield, which triggered geographic fence failsafe [3] conditions causing the test to terminate prematurely. A similar geographic fence breach occurred for some `Contact Investigation` UAVs during the follow-on event due to contacts generating outside of the operating area. A final change was implemented to decrease the Zephyr II transit speed in order to reduce battery drain and turn radius (which was also causing occasional geographic fence breaches). It is worth noting that none of these issues were observed during SITL environment testing. In fact, the issues were not even directly associated with the tactic implementation. Rather, they were simply reminders of the difference between experimentation in a simulated environment and experimentation in the real world.

Once these issues were addressed, a final flight was conducted with a heterogeneous swarm of five Zephyr IIs and four Mosquito Hawks. This test successfully completed initial assignment of roles and transitioned through all states to include `Contact Investigation` and `Threat Response`. Following the initial state transitions in response to contact generation, the vehicles reported errors due to a lack of synchronized contact lists across the swarm. This lack of synchronization effectively broke the deterministic tactic-level task-generation algorithm (i.e., vehicles did not all generate the same task set). Inspection of the flight logs, revealed that this was due to the asynchronous and lossy communications environment in which the swarm operates [3], [33] and the temporal proximity of contact reporting and auction re-initiation. Since a new contact causes the auction to restart immediately, if a vehicle is at the edge of communications range and does not receive a message for some reason, it may receive the auction restart message prior to the contact report updating the contact list (if it is received at all). Subsequent to the July 2020 field exercise, a semaphore was incorporated to delay the restart of the auction until receipt of an updated contact report. This resulted in additional time spent in loss states (defined in Section 4.4) as described in Section 5.2.1 to ensure synchronization of the contact list across the swarm.

5.2 Test Results and Analysis

Following the July 2020 field exercise, final updates were implemented on the `Base Defense` tactic to make improvements based on lessons learned and to incorporate logging for the capture of quantitative data. The following subsections document results obtained in SITL simulations with regard to the *Efficiency of Force* and *Detection* MOEs.

5.2.1 Evaluation: Efficiency of Force MOE

To test the efficiency of force for the `Base Defense` tactic, the total time within each state was monitored. Each vehicle was monitored individually to track the amount of time each vehicle spent in active states versus loss states. The following states from the diagrams in Chapter 4 were deemed loss states:

1. Base Defense AUCTION IN PROGRESS
2. Base Defense INIT AUCTION
3. Key Area Search ALLOCATION
4. Contact Investigation INIT
5. Threat Response INIT
6. Key Area Search TRANSIT

The following states were deemed active states:

1. Base Defense DETECT
2. Base Defense DELAY
3. Base Defense DESTROY
4. Contact Investigation INVESTIGATE

5. Key Area Search SURVEIL
6. Key Area Search SEARCH
7. Threat Response TRACK
8. Threat Response DESTROY

The final implementation of the Base Defense tactic was evaluated with seven simulations. Each simulation and a sum of the time in each state is displayed in the graphs of Figures 5.1 through 5.3. The graphs in Figures 5.1 and 5.2 present the homogeneous flight tests with five vehicles of either fixed wing or quadrotor per test. The graphs in Figure 5.3 present the heterogeneous flight tests with five vehicles in total, three of which were fixed wing and two of which were quadrotors.

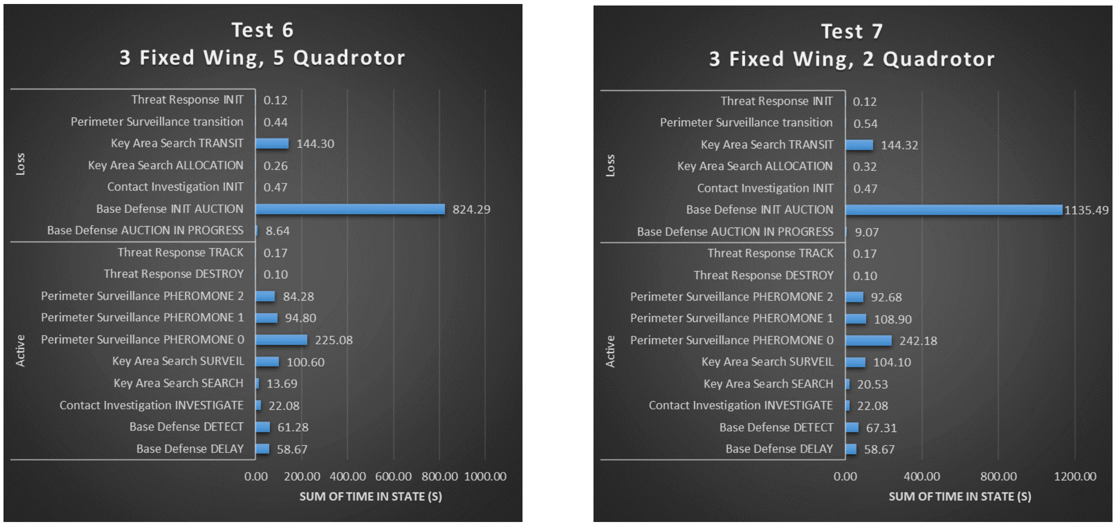

Figure 5.3. Tests 6 and 7 were conducted with heterogeneous swarms of three fixed wing and two quadrotor vehicles. This figure depicts the amount of time in seconds spent in each state.

The average active state to loss state proportion across all simulations was 0.50. This is a significant achievement which serves as a testament to the flexibility of the MASC hierarchy and the previously implemented plays and algorithms. The maximum and minimum active state to loss state ratios for each test across the simulations were 0.65 and 0.22, respectively.

The highest contributing loss state across all simulations was the Base Defense INIT AUCTION state. This state caused the most delay due to the requirement to synchronize

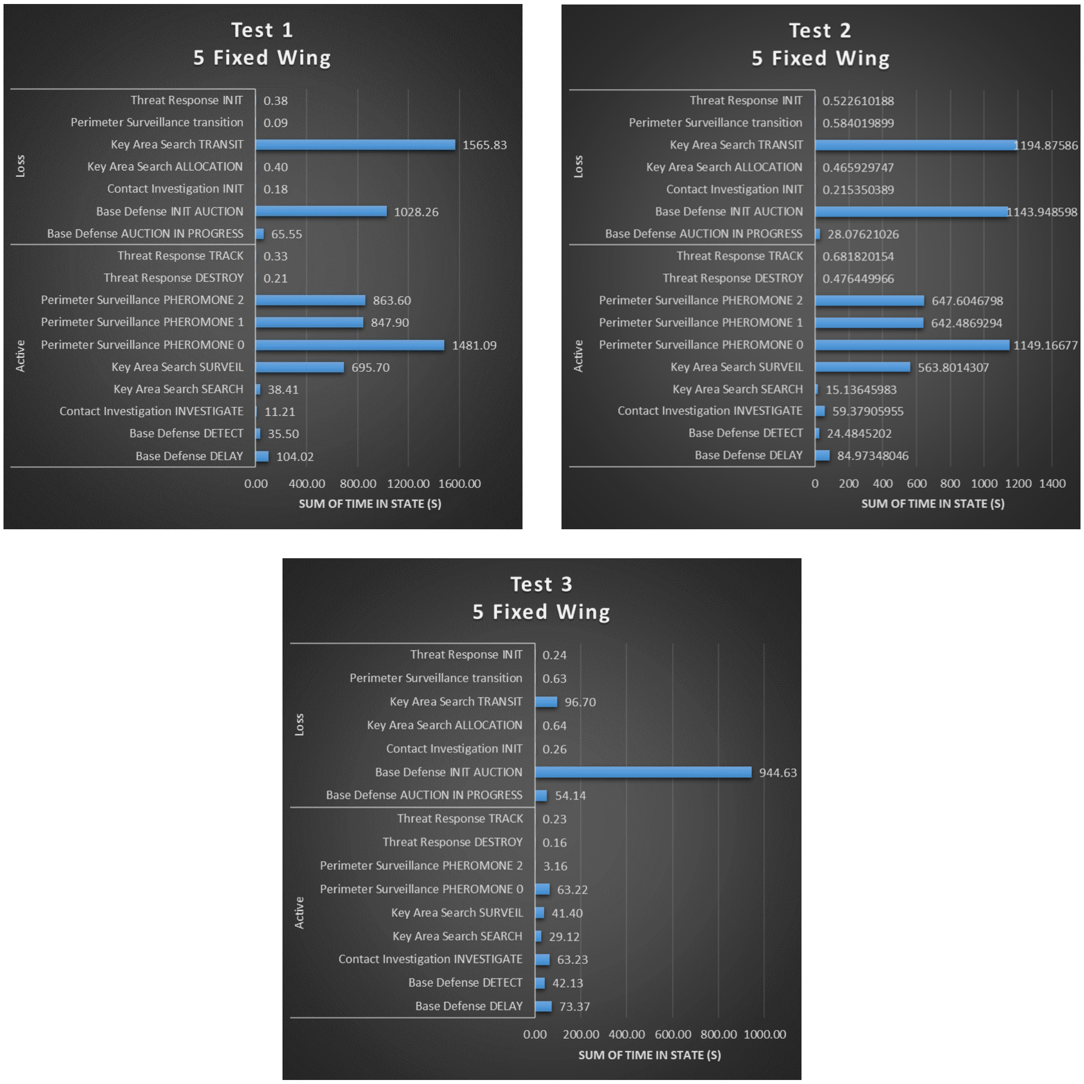

Figure 5.1. Tests 1 through 3 were conducted with swarms of five fixed wing vehicles. This figure depicts the amount of time in seconds spent in each state.

the contact list prior to beginning the auction and allocation process. Due to the ARSENL system's communications environment, each message has to be sent explicitly to each vehicle within the swarm to ensure delivery. Because of the delay this can impose, it is not a given that all vehicles will have completely accurate contact lists during the auction. As noted during live-flight experiments, this is a prerequisite for generating an accurate auction task list. In order to assure the accuracy of this critical task, a flag was set on

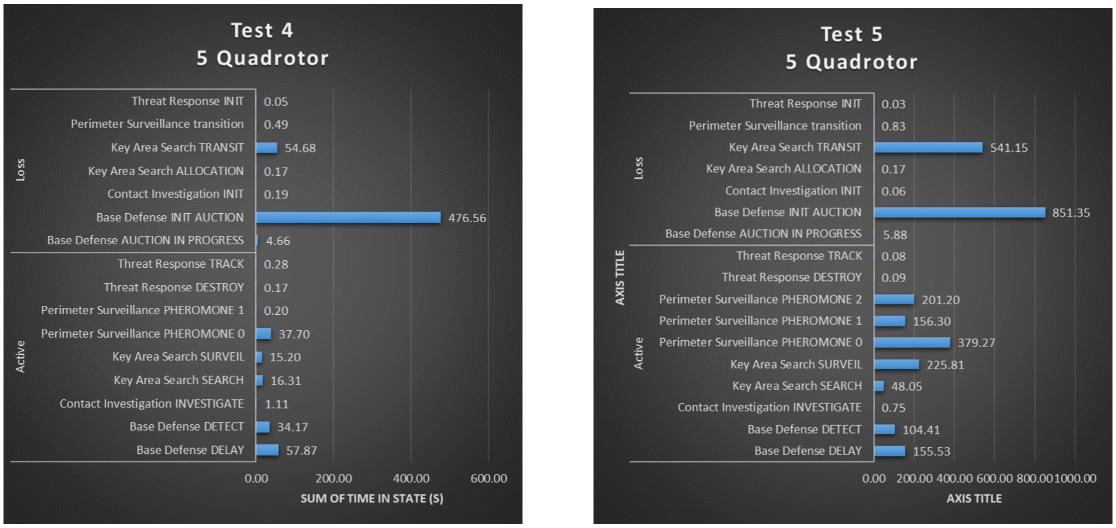

Figure 5.2. Tests 4 and 5 were conducted with swarms of five quadrotor vehicles. This figure depicts the amount of time in seconds spent in each state.

each vehicle whenever the auction was restarted to wait until a contact update was received before moving forward. This delay inherently added loss time but ensured the fleet-wide accuracy of the auction. A more efficient approach to ensuring auction synchronization will significantly increase the overall efficiency of the tactic. This, however, is left to future work.

5.2.2 Evaluation: Detection MOE

In order to evaluate the *Detection* MOE, we measured the total time spent in the INIT AUCTION and AUCTION IN PROGRESS states. These are the primary states of the auction that generated the tasks and assigned the roles across the swarm. Whenever a contact was detected, confirmed, denied, or neutralized the entire swarm was notified, an updated task set was generated, and the auction was restarted. Completion of role assignment process indicates that the swarm has reacted and adjusted to the change in the environment. Thus, completion serves as a competent measure of detection and reaction across the swarm. Table 5.1 summarizes the average time to react over the course of the seven tests described in Section 5.2.1.

Table 5.1. Summary of the average time to react to changes in contact status. Of note, the final column represents the average time to react across all auctions during the respective test.

	Swarm Configuration	Total Time to Respond (s)	Number of Auctions	Average Time to React (s)
Test 1	5 Fixed Wing	1093.80	51	21.44
Test 2	5 Fixed Wing	1172.02	60	19.53
Test 3	5 Fixed Wing	998.76	55	18.15
Test 4	5 Quadrotor	481.21	45	10.69
Test 5	5 Quadrotor	857.23	55	15.58
Test 6	3 Fixed Wing, 2 Quadrotor	832.92	60	13.88
Test 7	3 Fixed Wing, 2 Quadrotor	1144.56	65	17.60

Overall, the maximum of the averages across all tests was found in Test 1 which was run for the longest time and contained only fixed wing vehicles. It is hypothesized that the response time of homogeneous fixed-wing swarms is slower because no vehicle is inherently more suited to the task than any other and the higher transit speed makes distance to the contact less important. Further research would be required, however, to confirm or refute this hypothesis. Regardless, task assignment auctions completed reliably in under twenty-five seconds for all swarm configurations and enabled the system to react appropriately when responding to up to ten contacts.

5.3 Summary

The field exercises discussed in this chapter were used to verify and validate the behaviors implemented in the `Base Defense` tactic. Overall, the live-flight tests proved to be a valuable tool for testing conditions that the SITL environment was unable to simulate. The successful field experiments validated the `Base Defense` tactic as a successful proof of concept for a top-down, mission-centric approach to swarm behavior development.

SITL simulation experiments provided a controlled environment in which the quantitative data could be detected to evaluate the *Efficiency of Force* and *Detection* MOEs. As documented in this chapter, evaluation of the *Efficiency of Force* MOE showed room for improvement in the implementation of the `Base Defense` tactic. In particular, loss of efficiency was shown to be auction initialization requirements in the lossy communications environment proved inefficient. The most obvious way to improve the overall *Efficiency of Force* is clearly to reduce the time spent allocating tasks, and the most significant part of the task allocation process proved to be auction initialization.

Evaluation of the *Detection* MOE showed that the time to respond to changes in contact status remained acceptable across both homogeneous and heterogeneous swarms and that the response time due to task allocation was at its highest during tests with homogeneous fixed-wing swarms.

In summary, the `Base Defense` tactic as described was evaluated in SITL experiments and validated in live flight exercises. The tactic demonstrated successful allocation of tasks initially and following changes in the status of contacts. Further, both SITL and live flight experiments showed that the tactic was able to accomplish surveillance, investigation, and response tasks regardless of the composition of the swarm.

CHAPTER 6:
Conclusion

6.1 Summary of Findings

The `Base Defense` tactic was implemented with the MASC hierarchy and incorporated into the ARSENL repository. This research demonstrated the potential for using autonomous swarm systems in support of base defense operations and showed the utility and functionality of heterogeneous swarms for both surveillance and response tasks.

The `Base Defense` tactic accomplished the primary research objective of proving that an autonomous swarm can be tasked in a way that provides effective, doctrine-compliant base defense. We showed that the basic components of base defense can be decomposed into tasks which can be augmented or completely subsumed into an autonomous swarm. In addition, the development process showed that a top-down, mission-centric approach is effective for developing swarm capabilities.

The `Base Defense` tactic is also the first behavior within the ARSENL repository to utilize the Prim Auction algorithm, which showcased the utility, simplicity, and efficiency of this type of multi-item auction. Expanding upon this algorithm and utilizing it within this behavior showcased the potential for use in other ARSENL swarm behaviors. This reuse and expansion of the repository demonstrated the core concepts of MASC, most notably modularity and composability, to develop robust swarm behaviors to meet complex mission requirements.

Finally, by focusing on the mission, and deconstructing the requisite tasks and their subtasks we were able to produce an end result that met doctrine based MOEs and MOPs. The *Detection* and *Efficiency of Force* MOEs provided a baseline metric upon which future implementations can expand and improve. Overall, test results displayed that the `Base Defense` tactic as developed meets all essential tasks, however, multiple areas for potential optimization and improvement were identified.

6.2 Future Work

Future research suggested by this work can reasonably be divided into two categories. First, multiple ways in which the `Base Defense` tactic might be improved were identified. Second, more generalizable research into the algorithms and processes underpinning the tactic and its development is warranted.

6.2.1 `Base Defense` Tactic Iteration and Improvement

A number of ways in which the `Base Defense` tactic might be improved have been previously discussed. There are, however, many additional ways in which this behavior can be expanded.

One improvement that might be incorporated into the `Base Defense` tactic would be a more robust capability-based task assignment. In the current implementation, contact investigation and response tasks are automatically assigned to the vehicle that identified them, regardless of the vehicle's capabilities. The identification and inclusion of additional metrics to account for vehicle specific information such as battery life, sensor capabilities, and expendable configuration would improve the effectiveness of the tactic. Granted, this improvement may come at the expense of speed of response and efficiency of force; however, this is likely mitigated by the ability to more reliably assign the "right" vehicle to the "right" role.

A second improvement to the `Base Defense` tactic might be the incorporation of the unmanned ground vehicles (ARSENL is currently working to add unmanned ground vehicles, called rovers to their system). Generation of ground surveillance tasks for autonomous rovers can improve behavior lifespan and overall capability of the tactic. In order to adapt the `Base Defense` tactic to take advantage of ground vehicle incorporation, the generated tasks and their supporting plays must account for the maneuverability limitations of ground vehicles. Specifically, since ground vehicles are incapable of traversing or reaching certain areas, the task allocation process must invariably assign tasks only to vehicles capable of accomplishing them. In the present system, this means that vehicles would be unable to bid for tasks that they cannot accomplish. Successful implementation might also effectively account for disparate capabilities of vehicles operating in the same domain. For example, the assignment of `Threat Response` tasks with a *Destroy* setting to vehicles not capable of delivering air-to-ground ordnance might be precluded. It is important to note, that this constraint does not require all vehicles to accomplish a task the same way; different types of vehicles may very well use different plays to satisfy requirements of the same role.

Another possible focus for future work is the implementation of the `Base Defense` tactic within other behaviors. Examples might include mutually supporting defense and convoy defense. A mutually supporting defense behavior could be implemented at the tactic level with the `Base Defense` tactic being modified to support the defense of multiple bases with integrated fires and priorities. Another approach might be to embed the `Base Defense` tactic into a more robust tactic in much the same way that the `Pounce Drop` play was incorporated into the `Threat Response` play. Significantly more robust task assignment capable of dealing with competing priorities would be required for either approach. A second behavior to which the `Base Defense` tactic might be applied is convoy defense. This is a tactic that is already under development by ARSENL researchers as part of their rover-incorporation efforts, and extension or inclusion of the `Base Defense` tactic seems appropriate.

In general, the MASC approach calls for tactics to be utilized as building blocks for more complex phases making up even more complicated missions. This is an area that this research did not address, but one that would benefit from future work as the preceding examples indicate.

6.2.2 Multi-Vehicle System Capability and Development Research

In addition to enhancements to the `Base Defense` tactic itself, this research identified a number of areas for more broadly applicable future research. A number of examples are associated with the task assignment processes. As the *Efficiency of Force* MOE results indicate, inefficient task assignment has significant impact. Auction synchronization, particularly when dealing with dynamic or emergent tasks, was observed to have the most impact. Further research in this area might include the development of improved synchronization mechanisms or approaches to distributed decision making that are more robust in the face of inconsistency or lack of synchronization.

Similarly, the task assignment process itself might be improved. The approach used in this work accounted for different vehicle capabilities only to a limited extent. The identification of mechanisms for more granular representation of vehicle capabilities is worth further exploration. This would necessarily include accounting for vehicles that are not capable of completing certain tasks as noted in the previous section, but it would be significantly broader. For instance, two vehicles might both be capable of accomplishing a task, but one might do so more effectively than the other. This work would largely focus on quantitative metrics suitable for distributed negotiation algorithms.

Another potential area for future work is the development of plays and algorithms within the ARSENL repository for broad use. The `Base Defense` tactic required the bottom up development of three plays and only re-used one existing play. The MASC approach, however, is intended to foster code re-use. The identification and development of more algorithms and plays that are likely to be applicable to a broad array of tactics will foster tactic development and more effectively showcase the strengths of the MASC approach.

A loiter waypoint driver algorithm might, for instance, prove useful. This algorithm would enable a timed loiter at each waypoint prior to proceeding to the next. This sort of algorithm would allow for increased flexibility and might be useful in plays such as the `Contact Investigation` and `Key Area Search`. Additional approaches to waypoint navigation such as a providing a way to skip waypoints that are unreachable or take too long might also prove broadly applicable.

Among plays that might be useful is the notion of prioritized observation. This play might function similarly to the `Key Area Search` play but would allow for the prioritization of observation targets. Higher priority locations would be visited more often or observed for longer periods. This would enable higher level behaviors to modify the search patterns based on the given situation.

Finally, increased capability development using the MASC approach is warranted. The `Base Defense` tactic provides a proof of concept, but only scratched the surface of what can be developed using the architecture. A number of NPS efforts have developed notional tactics and missions using MASC [14], [34]. Few have actually been implemented on an actual UAS. The ARSENL swarm system is an ideal platform for further development and vetting of MASC processes and capabilities.

List of References

[1] D. H. Berger, *38th Commandant's Planning Guidance*, United States Marine Corps, 2019, pp. 1–23.

[2] United States Marine Corps, *3-20.5 Unmanned Aircraft Systems Operations*, Marine Corps Warfighting Publication, Washington D.C., 2016, pp. 1–1 : 5–5.

[3] T. H. Chung, M. R. Clement, M. A. Day, K. D. Jones, D. Davis, and M. Jones, "Live-fly, large-scale field experimentation for large numbers of fixed-wing UAVs," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, A. Okamura, Ed. Stockholm, Sweden: IEEE, 2016, pp. 1255–1262.

[4] Chairman Joint Chiefs of Staff, *3-10.1: Joint Tactics, Techniques, and Procedures for Base Defense*, Joint Doctrinal Publication, Washington D.C., 2019, pp. I–1 – II–3.

[5] United States Marine Corps, *3-30 Rear Area Security Operations*, Marine Corps Reference Publication, Washington D.C., 2018, pp. 1–2 – 2–19.

[6] United States Marine Corps, *2-10B.1 Intelligence Preparation of the Battlefield/Battlespace*, Marine Corps Reference Publication, 2014, pp. 4–8.

[7] United States Marine Corps, *1 - Warfighting*, Marine Corps Doctrinal Publication, 2018, pp. 1–18.

[8] United States Marine Corps, *3-20.5 Unmanned Aircraft Systems Operations*, Marine Corps Warfighting Publication, Washington D.C., 2016, pp. 3–8.

[9] United States Marine Corps, *3-10A.4i, Marine Rifle Squad*, Marine Corps Interim Publication, Washington D.C., 2020, p. 82.

[10] United States Marine Corps, *3-10F Fire Support Coordination in the Ground Combat Element*, Marine Corps Technical Publication, Washington D.C., 2018, pp. P–2.

[11] R. A. Brooks, "A robust layered control system for a mobile robot," MIT, Boston, MA, Tech. Rep. 864f, 1985.

[12] K. Giles, "Mission-based architecture for swarm composability," Ph.D. dissertation, Naval Postgraduate School, Monterey, CA, 2018.

[13] D. Brutzman, C. L. Blais, D. T. Davis, and R. B. Mcghee, "Ethical mission definition and execution for maritime robots under human supervision," *IEEE Journal of Oceanic Engineering*, pp. 1–17, 2018.

[14] D. Davis, K. Jones, M. Jones, and K. Giles, “Advanced swarm UAV capabilities through collaborative field experimentation,” Naval Postgraduate School, Monterey, CA, Tech. Rep., Apr 2018.

[15] B. Browning, J. Bruce, M. Bowling, and M. Veloso, “Stp: Skills, tactics, and plays for multi-robot control in adversarial environments,” *Proceedings of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering*, vol. 219, no. 1, pp. 33–52, 2005.

[16] A. Bredenfeld, A. Jacoff, I. Noda, and Y. Takahashi, *Robocup 2005: Robot Soccer World Cup IX*. Springer, 2006, vol. 4020.

[17] D. Kortenkamp, R. Simmons, and D. Brugali, *Robotic Systems Architectures and Programming*. Cham: Springer International Publishing, 2016, pp. 283–306. Available: https://doi.org/10.1007/978-3-319-32552-1_12

[18] L. de Koning, J. P. Mendoza, M. M. Veloso, and R. van de Molengraft, “Skills, tactics and plays for distributed multi-robot control in adversarial environments,” in *RoboCup*, 2017.

[19] C. A. Miller, R. P. Goldman, H. B. Funk, P. Wu, and B. B. Pate, “A playbook approach to variable autonomy control: application for control of multiple, heterogeneous unmanned air vehicles,” in *Proceedings of FORUM 60, the Annual Meeting of the American Helicopter Society*, 2004.

[20] J. L. Sanchez-lopez, H. Bavle, and C. Sampedro, “AEROSTACK : an architecture and open-source software framework for aerial robotics,” in *International Conference on Unmanned Aircraft Systems*, IEEE, Ed. Arlington, VA: IEEE, 2016, pp. 332–341.

[21] T. H. Chung. OFFensive Swarm-Enabled Tactics (OFFSET) Proposers Day. Defense Advanced Research Projects Agency Tactical Technology Office. [Online]. Available: https://www.darpa.mil/work-with-us/offensive-swarm-enabled-tactics. August 25, 2020.

[22] D. T. Davis, T. H. Chung, M. R. Clement, and M. A. Day, *Multi-swarm Infrastructure for Swarm Versus Swarm Experimentation*. Cham: Springer International Publishing, 2018, pp. 649–663. Available: https://doi.org/10.1007/978-3-319-73008-0_45

[23] About Robot Operating System (ROS). (n.d.). Robot Operating System (ROS). [Online]. Available: ros.org/about-ros. Aug 3, 2020.

[24] ArduPilot Documentation. (n.d.). ArduPilot Dev Team. [Online]. Available: https://ardupilot.org/ardupilot/index.html. Aug 3, 2020.

[25] D. T. Davis, T. H. Chung, M. R. Clement, and M. A. Day, "Consensus-based data sharing for large-scale aerial swarm coordination in lossy communications environments," in *2016 IEEE/RSJ Proceedings of the International Conference on Intelligent Robots and Systems (IROS)*. Daejeon, Korea: IEEE/RSJ, Oct 2016, pp. 3801–3808.

[26] B. Schwartz, "A computational analysis of the auction algorithm," *European Journal of Operational Research*, vol. 74, no. 1, pp. 161 – 169, 1994. Available: http://www.sciencedirect.com/science/article/pii/0377221794902143

[27] M. S. Hopchak, "Autonomous decision in multi-robot systems," M.S. thesis, Naval Postgraduate School, 2018.

[28] B. J. Campbell, "Autonomous cueing within heterogeneous robot swarms," M.S. thesis, Naval Postgraduate School, 2019.

[29] M. G. Lagoudakis, M. Berhault, S. Koenig, P. Keskinocak, and A. J. Kleywegt, "Simple auctions with performance guarantees for multi-robot task allocation," in *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2004, vol. 1, pp. 698–705.

[30] A. Hietpas, "Stigmergic control of dual direction communication ferry nodes for denied communication environments," M.S. thesis, Naval Postgraduate School, 2018.

[31] M. A. Day, M. R. Clement, J. D. Russo, D. Davis, and T. H. Chung, "Multi-UAV Software Systems and Simulation Architecture," in *2015 International Conference on Unmanned Aerial Systems*. Denver, CO: IEEE, 2015, pp. 426—-435.

[32] G. Roedler and C. Jones, "Technical measurement guide," *International Council on Systems Engineering and Practical Software and Systems Measurement*, 2005.

[33] D. T. Davis, T. H. Chung, M. R. Clement, and M. A. Day, "Consensus-based data sharing for large-scale aerial swarm coordination in lossy communications environments," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. Daejeon, Korea: IEEE, 2016, pp. 3801–3808.

[34] S. Deleon, R. Fernandez, Julio, P. Hunter, David, K. Janicek, and S. Naser, "Mission profile development for the use of swarm unmanned aerial vehicles or unmanned aerial systems," M.S. thesis, Naval Postgraduate School, 2019. Available: https://calhoun.nps.edu/handle/10945/64134

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