

Swarm Counter-Asymmetric-Threat (CAT) 6-DOF Dynamics Simulation

by
James Bobinchak
Weapons and Energetics Department
and
Gary Hewer
Research and Engineering Sciences Department

JULY 2005

**NAVAL AIR WARFARE CENTER WEAPONS DIVISION
CHINA LAKE, CA 93555-6100**



Approved for public release; distribution unlimited.

Naval Air Warfare Center Weapons Division

FOREWORD

This report was prepared for, and funded in fiscal year 2004 by Gil Graff of the Air Weaponry Technology Program (Code 351) at the Office of Naval Research. The task under this program was to investigate the feasibility of utilizing swarming submunitions to counter massed raids of at least 40 asymmetric threats directed against surface ships. Swarming behavior in weapon systems is a new concept borrowed from nature, where certain animals (such as birds, fish, ants, and hornets) are observed to function cooperatively to pursue a common goal by employing network communications and coordinated teamwork. The research was particularly focused on possible swarm models for optimum target-weapon pairing.

This report documents the mathematical models used in the swarm counter-asymmetric-threat (CAT) simulation and the results of extensive Monte Carlo simulations. The swarm CAT simulation is a full six-degree-of-freedom (6-DOF), physics-based simulation that is capable of modeling the dynamics of a swarm of small, low-cost, lightweight, airborne canisters that cooperatively acquire, track, and pursue a plurality of highly maneuverable asymmetric littoral threats.

This report was reviewed for technical accuracy by Sam Ghaleb.

Approved by
W. H. TANAKA
NAVAIR Air Weaponry Technology Program Manager
2 March 2005

Under authority of
W. M. SKINNER
RDML, U.S. Navy
Commander

Released for publication by
K. L. HIGGINS
Director for Research and Engineering

NAWCWD Technical Publication 8593

Published by..... Technical Information Division
Collation..... Cover, 25 leaves
First printing..... 20 paper, 1 electronic media

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 instructions, 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, D.C. 20503.				
1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE July 2005		3. REPORT TYPE AND DATES COVERED Final; fiscal year 2004
4. TITLE AND SUBTITLE Swarm Counter-Asymmetric-Threat (CAT) 6-DOF Dynamics Simulation (U)			5. FUNDING NUMBERS	
6. AUTHOR(S) James Bobinchak and Gary Hewer				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Air Warfare Center Weapons Division 1 Administration Circle China Lake, CA 93555-6100			8. PERFORMING ORGANIZATION REPORT NUMBER NAWCWD TP 8593	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) (U) This report documents the mathematical models used in the swarm counter-asymmetric-threat (CAT) simulation and the results of extensive Monte Carlo simulations. The swarm CAT simulation is a full six-degree-of-freedom (6-DOF), physics-based simulation that is capable of modeling the dynamics of a swarm of small, low-cost, lightweight, airborne canisters that cooperatively acquire, track, and pursue a plurality of highly maneuverable asymmetric littoral threats. The canisters independently collect and share information among the swarm and therefore can achieve large lethality footprints that are not possible with single canisters or with clusters of canisters acting unilaterally. Distributed information sharing is key to achieving cooperation and is essential for maintaining swarm cohesion, avoiding canister collisions, pursuing multiple targets, and optimally assigning canisters to engage targets. In order to provide insight on the feasibility of the swarm concept, Monte Carlo simulations were undertaken and the results are presented in this report.				
14. SUBJECT TERMS 6-DOF simulation Ad hoc wireless network communications Asymmetric raids Autonomous swarms Littoral threats			15. NUMBER OF PAGES 48	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED		18. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	19. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	20. LIMITATION OF ABSTRACT SAR

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE *(When Data Entered)*

CONTENTS

Introduction.....	5
Background.....	5
Swarm Counter-Asymmetric-Threat (CAT) Overview	5
Application.....	7
Swarm Simulation Overview	7
Simulation Coordinates.....	9
Coordinate Systems	9
Inertial Coordinates.....	11
Canister Body Coordinates	11
Seeker Coordinates	11
Target Body Coordinates	11
LOS Coordinates.....	11
Quaternions	12
Operation Analysis.....	13
Network Communication and Message Traffic	13
Virtual Coupling	14
Target–Canister Pairing.....	15
Problem Solving.....	17
Solution to the Asymmetric Multi-Assignment Problem.....	17
Forward Auction	18
Reverse Auction.....	19
Swarm Simulation Description	20
Canister Model	20
Seeker Model	20
Seeker Vibration	20
Kalman LOS Filter.....	20
Maneuver Detection.....	22
Guidance Law	22
Divert Autopilot	23
Divert Thrusters	23
Canister Dynamics	23
IMU.....	24
GPS Receiver	24

Target and Canister Discussion	25
Target Motion Model	25
Resolving Commonality of Targets	25
Canister Hardware Component Mock-Up.....	26
Simulation Noise Sources	27
Simulation Evaluation Using Box Plots	28
Box Plots	28
Simulation Setup	30
Pk Versus Number of Canisters	31
Variation of Canister Drop Height	35
Pk Versus Canister-to-Target G Ratio.....	35
Pk Versus Constrained and Unconstrained Multi-Assignment	36
Pk Versus Canister Spatial Density	37
Canister Spatial Coverage Versus Target Deployment	38
Noise Parameter Trade-Off (Noise Sensitivity)	39
Swarm Drop Versus Baseline Canister Drop	40
Swarm Drop With and Without Virtual Coupling	42
Conclusions	43
Future Work	44
Imaging Seeker Model	44
Atmospheric Effects on Canister Maneuverability	44
Extend 2-D Planar Virtual Coupling to 3-D.....	45
Wind Gust Noise Model.....	45
Wireless Communications Model	46
Computational Delay.....	46
Divert Thruster Model.....	46
References	47

Figures:

1. Cooperating Swarm of Canisters Simultaneously Engaging a Group of Littoral Asymmetric Threats.....	6
2. High-Level Overview of the Essential Elements in Cooperative Multi-Target Tracking and Intercept.	8
3. Swarm Simulation.	9
4. Coordinate Systems and Transformation Matrices Used in the Simulation.	10
5. Orientation of a Rotated Frame Relative to an Inertial Frame.	10
6. Typical Swarm Network Configuration as a Result of Virtual Spring Coupling.	15
7. Example of Reachability Circles and Corresponding Table of Intercept Probabilities.....	16

8. Two Canisters Making Statistically Independent Measurements of the Same Target.....	26
9. Canister Hardware Component Mock-Up.....	27
10. Gaussian and Rayleigh Probability Distributions and Corresponding Box Plots.	29
11. Seven Distributions Compared Using Box Plots.	30
12. Initial Target and Canister Lattice Configuration.	31
13. Pk as a Function of Number of Canisters Per Target for Targets 1 Through 4.	32
14. Pk as a Function of Number of Canisters Per Target for Targets 5, 6, and 7.	33
15. Initial and Final Lattice Configuration When Three Canisters Are Assigned to Each of Seven Targets.....	34
16. Pk as a Function of Canister-to-Target G Ratio, Using 7 Canisters in an Hexagonal Attack Configuration.....	35
17. Pk Versus Constrained and Unconstrained Multi-Assignment of 28 Canisters to 7 Targets.....	36
18. Pk as a Function of Virtual Spring Rest Length for Canister-to-Target G Ratio of 1.0.	37
19. Pk as a Function of Virtual Spring Rest Length for Canister-to-Target G Ratio of 1.5.	38
20. Pk as a Function of Target Range.	39
21. Pk Versus Number of Canisters Per Target.....	40
22. Pk Results for a Swarm Drop Versus Biased Baseline Canister Drop With Inter-Canister Spacing of 19.7 Feet.....	41
23. Pk Results for a Swarm Drop Versus Biased Baseline Canister Drop With Inter-Canister Spacing of 29.5 Feet.....	42
24. Comparison of Pk Results for a Swarm Drop With and Without Virtual Coupling.....	43

Table:

1. Noise Sources Currently Modeled in the Simulation and a Representative Value of the 1-Sigma Uncertainty	28
--	----

This page intentionally left blank.

INTRODUCTION

The use of swarming behavior in weapons systems is a new concept borrowed from observations that certain living beings (such as birds, fish, ants, and hornets) tend to perform swarming behavior. By employing network communications and coordinated teamwork, complicated tasks are mastered by collective behavior. The swarm concept, by analogy, is extended in this research to investigate the following hypothesis: a swarm of low-cost submunitions with inter-submunition communications capabilities can effectively engage a plurality of highly maneuvering asymmetric threats.

The littoral asymmetric threat model for this study includes jet skis, bog hammers, Boston whalers, and zodiacs that engage in massed raids of 40 or more boats. Because the primary objective of the research is to define and prove the efficacy of a swarm model, the research will assume asymmetric raiders confined and engaged in a sector 8 kilometers or less from the target ships, and that the submunition delivery platform has been vectored to the sector by shipboard assets.

To date, the research has focused on possible swarm models for optimum target–weapon pairing. Increasing the size of the cluster lethality footprint as a function of unit density and spatial configuration was also a definite goal. Specific questions included: (1) Can a functional grid of submunitions be synthesized by modifying and augmenting the submunitions' capabilities to counter asymmetric threats? (2) What type of target–weapon pairing is possible from the grid array? and (3) Can tangible benefits be realized at a reasonable cost within the submunition family sizing constraints? The first two questions were answered using extensive Monte Carlo simulation and the results are presented in the simulation results section of this report. Tailored flight tests in conjunction with additional simulations will answer the third question.

BACKGROUND

SWARM COUNTER-ASYMMETRIC-THREAT (CAT) OVERVIEW

Figure 1 illustrates a typical scenario in swarm CAT. The submunition canisters are ejected from a single or multiple delivery platforms, and spread over a wide area to form a large swarm of 500 canisters or more, which engage up to 40 highly maneuverable asymmetric targets. The canisters function cooperatively as autonomous agents that rely on simple instructions to achieve a common goal. They are autonomous in that there is no centralized control, or hub, in the network to direct them. Each canister transmits a message to the other canisters in the swarm concerning its sensor measurements, and likewise receives messages from the other canisters in the swarm. Initially this message traffic is used to assign canisters to targets so as to maximize some objective, such as the global probability of intercepting all targets. Immediately thereafter, the message traffic

is used to compute intercept trajectory and to maintain a safe inter-canister spacing during formation flying. It is also used to dynamically adjust the inter-canister spacing as a function of target maneuver, and time-to-go, in order to increase the probability of killing (Pk) the target. The canisters share information so that all have access to the same knowledge database, stored locally within each canister, thereby creating database redundancy within this robust network. If a few canisters malfunction or are destroyed, the remaining canisters in the network will continue to function and cooperate without problems.




FIGURE 1. Cooperating Swarm of Canisters Simultaneously Engaging a Group of Littoral Asymmetric Threats.

Every canister in the swarm contains a global position system (GPS) receiver and inertial measurement unit (IMU) for measuring its position, velocity, and acceleration relative to some inertial reference, such as the position of swarm deployment. Canister altitude is obtained via a laser altimeter. A low-cost infrared (IR) camera is used for detecting the angular position of targets within the vicinity of, and relative to, the canister. Each canister also possesses wireless local area networking capability, such as IEEE 802.11b (Wi-Fi), or better yet, Bluetooth wireless technology, used to communicate

with other canisters in the network. Measurements from each sensor on the canister are combined to form the message packet transmitted among canisters. The message packet includes canister address, position, velocity, and acceleration, and the positions of any targets that happen to fall within the field of view (FOV) of the IR camera. An on-board processor, in conjunction with a software algorithm, uses the message traffic from all canisters to compute target–weapon assignments and to compute guidance commands for intercepting the assigned target. The message traffic is also used for maintaining swarm cohesion during target pursuit.

Once the canisters are ejected from the delivery platform and assigned to a specific target, they maneuver so that those assigned to the same target form a virtually coupled local swarm network, with each canister acting as a node in the network. Node connectivity is achieved using a potential function of any reasonable shape, so canisters become virtually coupled once they maneuver into the local neighborhood of another canister pursuing the same target. The potential function provides the local guidance and control for formation flying, while divert thrusters provide the necessary maneuver capability.

Robust assignment algorithms provide the means for optimally assigning canisters to targets. The assignment objective may be to maximize the global probability of intercepting all targets, or it may be to maximize the probability of intercepting a specific high-value target at the expense of missing a lower-value target.

APPLICATION

SWARM SIMULATION OVERVIEW

In order to understand the complex interaction of the swarm model capabilities just presented, progressive simulations incorporating varying degrees of network and sensor fidelity and control detail were needed. Embracing this philosophy, a modular simulation incorporating all of the high-level components shown in Figure 2 was created as an initial integration effort. Simple motion models for the threats and canisters were used in the initial feasibility studies, with Stochastic component uncertainties incorporated using parametric noise models. Subsequent refinements included a more extensive Monte Carlo capability, a Gaussian circular lethality model, a GPS model, an IMU model, a Laser Altimeter model, a Kalman filter for tracking pointing angle estimates, and a finite seeker FOV model.




FIGURE 2. High-Level Overview of the Essential Elements in Cooperative Multi-Target Tracking and Intercept.

The interaction among these models within the swarm simulation is shown in Figure 3. Although the diagram shows only one target block and one canister block, the simulation can actually model an indefinite number of canisters engaging an indefinite number of targets, with the only real limit being the amount of computer memory required to store the variables used by the models. For a particular target–canister pair, the kinematic relationship between the pair is used to compute a true (error free) line-of-sight (LOS) angle that is sent to the seeker model. Parametric errors for LOS noise and canister body vibration are injected at the seeker level, and a measured LOS angle is computed relative to the inertial reference frame. The two-state angular Kalman filter makes an estimate of both the true LOS angle and true LOS angular rate, given the noisy measurement, and passes the angular rate estimate to the guidance-and-control computer. At this point, the guidance computer calculates the commands for driving the canister toward the target. Because the canister is only one node in a network of many canisters, the swarm control command required for maintaining inter-canister separation and cohesion is added to the guidance command, and the resultant command is passed to the autopilot. The autopilot commands the thrusters to generate forces that will accelerate the canister toward the target, while simultaneously maintaining swarm cohesion. Although not yet implemented, forces and moments on the canister due to aerodynamic effects would be combined with the thruster force, along with small random forces and moments generated by wind gusts. The resultant forces and moments are used in the equations governing canister motion in order to compute the velocity of the canister relative to the inertial reference. Canister and target velocities are used in the kinematics model to compute a true LOS angle to the target, thereby closing the simulation guidance loop.




FIGURE 3. Swarm Simulation.

SIMULATION COORDINATES

COORDINATE SYSTEMS

Figure 4 is a map of the five coordinate systems (or “coordinate frames”) used in the simulation: the inertial, canister body, seeker, target body, and LOS coordinate frames. All, except for the inertial frame, are referred to as rotated frames because their orientation is expressed relative to the inertial frame axes and defined by yaw, pitch, and roll angles ψ , θ , and ϕ , as shown in Figure 5. The transformation matrices T_{can} , T_{skr} , T_{tgt} , and T_{LOS} allow a mapping of points and vectors from the inertial frame to the rotated frame designated by the subscript. Because the seeker is body-fixed, the seeker frame is aligned with the canister body frame and therefore $T_{\text{can}} = T_{\text{skr}}$. The following subsections summarize each coordinate frame.




FIGURE 4. Coordinate Systems and Transformation Matrices Used in the Simulation.




FIGURE 5. Orientation of a Rotated Frame Relative to an Inertial Frame.

Inertial Coordinates

A non-rotating flat earth is assumed in the simulation. The inertial coordinate frame is aligned with the canister body frame at the beginning of the simulation. The origin of the inertial frame is located at the specified altitude of the swarm. The inertial xy-plane is parallel with the earth and the positive z-axis is directed vertically upward. The xyz-axes form a right-handed set.

Canister Body Coordinates

The canister body coordinate frame is defined with its origin on the canister's longitudinal axis at the center of gravity. The positive x-axis is directed out the nose and the positive y-axis is directed out the left-hand side. The z-axis completes the right-handed set. Canister attitude relative to the inertial frame is defined by the angles ψ_{can} , θ_{can} , and ϕ_{can} , and is adjusted so that the canister nose is pointing downward. Canister attitude relative to the inertial frame is obtained via measurements made by the IMU.

Seeker Coordinates

Because the seeker is fixed to the canister body, the seeker frame is identical to the body frame except for an offset in origin. The origin of the seeker frame is at the position of the seeker along the longitudinal axis of the canister body.

Target Body Coordinates

The target coordinate frame is defined with its origin on the target's longitudinal axis at the center of gravity. The positive x-axis is directed out the nose and the positive y-axis is directed out the left-hand side. The z-axis completes the right-handed set. The initial location of the origin is relative to the inertial frame and is specified by an x,y-coordinate position on the ocean plane. Because the target model is a three-degree-of-freedom (3-DOF) model, target attitude relative to the inertial frame is defined simply by the heading angle ψ_{tgt} .

LOS Coordinates

The LOS coordinate frame is defined so that the x-axis lies along the range vector between the canister and target. The y-axis is perpendicular to the x-axis and is parallel to the inertial xy-plane. The z-axis completes the right-handed set. The orientation of the LOS coordinate frame relative the inertial frame is defined by the angles ψ_{LOS} and θ_{LOS} . The LOS frame does not roll (i.e., $\phi_{LOS} = 0$).

QUATERNIONS

Quaternions are used in the simulation to obtain the matrix for performing transformations between coordinate systems (Reference 1). Once the transformation matrix is known, a set of Euler angles can be determined.

The elements $a_{11}, a_{12}, \dots, a_{33}$ of the transformation matrix A are functions of the quaternion parameters e_1, e_2, e_3 , and e_4 . The elements are computed as

$$a_{11} = e_1^2 - e_2^2 - e_3^2 + e_4^2$$

$$a_{12} = 2 \cdot (e_1 \cdot e_2 + e_3 \cdot e_4)$$

$$a_{13} = 2 \cdot (e_2 \cdot e_4 - e_1 \cdot e_3)$$

$$a_{21} = 2 \cdot (e_3 \cdot e_4 - e_1 \cdot e_2)$$

$$a_{22} = e_1^2 - e_2^2 + e_3^2 - e_4^2$$

$$a_{23} = 2 \cdot (e_2 \cdot e_3 + e_4 \cdot e_1)$$

$$a_{31} = 2 \cdot (e_1 \cdot e_3 + e_2 \cdot e_4)$$

$$a_{32} = 2 \cdot (e_2 \cdot e_3 - e_1 \cdot e_4)$$

$$a_{33} = e_1^2 + e_2^2 - e_3^2 - e_4^2$$

The Euler angles are then calculated as

$$\psi = \tan^{-1} \left(\frac{a_{12}}{a_{11}} \right)$$

$$\theta = -\sin^{-1}(a_{13})$$

$$\phi = \tan^{-1} \left(\frac{a_{23}}{a_{33}} \right)$$

Quaternion parameters are obtained by integrating their time-derivatives, which are functions of the angular rates (p,q,r) about the rotating system's coordinate axes (x,y,z). The quaternion time-derivatives are

$$\dot{e}_1 = \frac{1}{2} \cdot (-e_4 \cdot p - e_3 \cdot q - e_2 \cdot r)$$

$$\dot{e}_2 = \frac{1}{2} \cdot (-e_3 \cdot p + e_4 \cdot q + e_1 \cdot r)$$

$$\dot{e}_3 = \frac{1}{2} \cdot (e_2 \cdot p + e_1 \cdot q - e_4 \cdot r)$$

$$\dot{e}_4 = \frac{1}{2} \cdot (e_1 \cdot p - e_2 \cdot q + e_3 \cdot r)$$

Transformations are initialized by solving for the initial values of the quaternion parameters given the initial values of the Euler angles. The quaternion parameter initial values are

$$e_1(0) = \cos(\psi_o) \cdot \cos(\theta_o) \cdot \cos(\phi_o) + \sin(\psi_o) \cdot \sin(\theta_o) \cdot \sin(\phi_o)$$

$$e_2(0) = \sin(\psi_o) \cdot \cos(\theta_o) \cdot \cos(\phi_o) - \cos(\psi_o) \cdot \sin(\theta_o) \cdot \sin(\phi_o)$$

$$e_3(0) = \cos(\psi_o) \cdot \sin(\theta_o) \cdot \cos(\phi_o) + \sin(\psi_o) \cdot \cos(\theta_o) \cdot \sin(\phi_o)$$

$$e_4(0) = \cos(\psi_o) \cdot \cos(\theta_o) \cdot \sin(\phi_o) - \sin(\psi_o) \cdot \sin(\theta_o) \cdot \cos(\phi_o)$$

where $\psi_o = \frac{\psi(0)}{2}$, $\theta_o = \frac{\theta(0)}{2}$, and $\phi_o = \frac{\phi(0)}{2}$

OPERATION ANALYSIS

NETWORK COMMUNICATION AND MESSAGE TRAFFIC

Data processing delays and time-division radio transmission schemes are currently not implemented in the swarm simulation. Instead, each canister transmits and receives data instantly and synchronously with all other canisters in the swarm network. Subsequent simulation improvements will include a more refined model of the actual network communication system.

A preferred embodiment of the network communication system would employ Bluetooth wireless technology.* The Bluetooth protocol is preferred (over IEEE 802.11b, for example) because it is designed to operate in noisy frequency environments. It uses adaptive frequency hopping to reduce interference between other wireless technologies sharing the 2.4-gigahertz spectrum. Bluetooth uses a baseband layer implemented as a link controller, to carry out low-level routines such as link connection and power control. The baseband transceiver applies a time-division duplex scheme that allows the canisters to alternately transmit and receive data packets in a synchronous manner. Data packets consist of an access code, header, and payload. The access code is used for timing synchronization, offset compensation, paging, and inquiry. The header contains information for packet acknowledgement, packet numbering for out-of-order packet

* <http://en.wikipedia.org/wiki/Bluetooth>. 20 July 2005.


reordering, flow control, slave address, and error checksum. The packet payload contains the combined data from all the sensors on the canister. Data include canister position, velocity, and acceleration, and the positions of any targets detected within the IR sensor FOV. A unique canister identification number (the address) is also needed during target–weapon pairing. This data packet is transmitted to all canisters in the swarm. Without additional power amplification, the maximum transmission power output limits the communication range to about 100 meters. As implemented, Bluetooth wireless technology provides a robust, low complexity, low power, and low cost solution for creating a network of radio links among canisters.

Virtual Coupling

The general understanding is that swarming behavior is a result of the interplay between long-range attraction and short-range repulsion (Reference 2). This behavior is implemented using a piece-wise linear virtual spring having a potential function with a minimum value at some finite distance from the canister. When two or more canisters are within the local neighborhood of one another, they move toward this minimum potential. As an example, the potential function of the piecewise linear virtual spring illustrated in Figure 6a is

$$V = \frac{1}{2}k(r - r_o)^2$$

where r_o is the virtual spring rest length, k is the spring coefficient, and r is the canister separation distance. When the canisters are separated by a distance equal to the rest length of the virtual spring (i.e., $r = r_o$), they are at the minimum value of their neighbor's potential function and form a stable network as shown in Figure 6b. High spring stiffness (i.e., $k = k_1$) is used when $r < r_o$, and low spring stiffness (i.e., $k = k_2$) is used when $r > r_o$. This piecewise linear spring has the effect of quickly forcing canisters to separate if they get too close to one another, and easing them back into position when they are too far apart. A damping term, proportional to the canisters' relative velocity, is used to prevent oscillations within the swarm. The first derivative of the potential function yields the steering command (the commanded force) that is superimposed with the guidance commands from the guidance and control computer. The resultant command is sent to the divert autopilot, which activates the divert thrusters to generate the force required to maneuver the canister to the location of minimum potential among its neighbors, while simultaneously pursuing its assigned target.



(a) Potential function and corresponding force function for a piecewise linear virtual spring. (b) Overhead view of virtually coupled canisters in a swarm network.

FIGURE 6. Typical Swarm Network Configuration as a Result of Virtual Spring Coupling.

TARGET-CANISTER PAIRING

Consider the asymmetric assignment problem, where we want to assign n canisters to m targets ($m < n$). Each canister is capable of intercepting at most one target; however each target may be attacked by more than one canister. The probability that a canister can intercept a target is used as a means of matching canisters to targets. This is illustrated in Figure 7a, where it is assumed that the delivery platform has ejected the canisters C1, C2, ..., C7, widely over the threats T1, T2, ..., T4. Because the canisters are falling, they will hit the ground within some time-to-go interval. Given this time interval, each canister has a finite area—known as its reachability area—within which it can maneuver. Likewise, each target can maneuver within a finite area during that same time interval and so has associated with it a reachability area. For simplicity we assume that the areas are circular, but they need not be. The probability of a canister's intercepting a particular target is simply the ratio of the overlapped area to the total target reachability area. Once a table of probabilities is generated (Figure 7b), an assignment algorithm is used to maximize the global probability of intercepting all targets. This is accomplished using the reverse auction algorithm (Reference 3) for the solution of unconstrained multi-assignment problems. For constrained multi-assignment problems, target-canister pairing is accomplished using the algorithms proposed by Castañon (Reference 4) and Kennington (Reference 5). The latter two algorithms have the advantage of allowing the number of canisters per target to be specified during the assignment process. This enables one to allocate more canisters to high-valued targets and fewer to low-valued targets, or to

balance the number of canisters per target while maximizing the global probability of intercept.

Another way of generating the probability table is to simply use inverse range, or any monotonically decreasing function of range, as the cost benefit for matching a canister to a target. This is possible because range is a pretty good indicator as to whether a canister can intercept a target. Targets that are closer to a canister are easier to detect, track, and therefore intercept, than targets at a distance.

Whether using reachability circles or a monotonically decreasing function of range as the cost benefit used in target–canister assignment, a problem occurs if a canister is directly over a target, or nearly so, but the two are moving in opposite directions. Because the canister is very small and lightweight, it might not possess enough impulse to change its direction of motion to coincide with that of the target. Even if the required impulse were available, there may not be enough time to make such a drastic course change because the canister is dropped from a relatively low altitude (500 to 1000 feet) and is falling due to gravity. To overcome this problem, canister–target closing velocity should be incorporated into the cost benefit when matching canisters to targets.




FIGURE 7. Example of Reachability Circles and Corresponding Table of Intercept Probabilities.

PROBLEM SOLVING**SOLUTION TO THE ASYMMETRIC MULTI-ASSIGNMENT PROBLEM**

Once a table of canister–target intercept probabilities is generated, an assignment algorithm is used to maximize the global probability of intercepting all targets. The actual linear programming problem to be solved is

$$\text{maximize } \sum_{(i,j) \in A} a_{ij} x_{ij} \quad (\text{maximize global probability of intercept})$$

subject to

$$\begin{aligned} 1 \leq \sum_{j \in A(i)} x_{ij} &\leq \alpha_i & \forall i = 1, \dots, m & \quad (\text{multiple canisters assigned to each target}) \\ \sum_{i \in B(j)} x_{ij} &= 1 & \forall j = 1, \dots, n & \quad (\text{one target assigned to each canister}) \\ 0 \leq x_{ij} & & \forall (i, j) \in A & \quad (0 \text{ or } 1 \text{ for all possible pairs}) \\ \sum_{i=1, \dots, m} \alpha_i &\geq n & & \quad (\text{due to the equality constraint}) \end{aligned}$$

where

x_{ij} = decision variable (0 or 1)

$A(i)$ = set of canisters to which target i can be assigned

$B(j)$ = set of targets to which canister j can be assigned

A = set of all possible pairs (i, j)

a_{ij} = probability of canister j intercepting target i

α_i = upper bound on the number of canisters to which target i can be assigned

m = total number of targets

n = total number of canisters

This problem simply states that the global probability of intercept must be maximized, while ensuring that every target i is assigned to at least one canister, but no more than α_i canisters, and every canister j is assigned to exactly one target. Because α_i is an upper limit on the assignment, this is a constrained multi-assignment problem. To generate an unconstrained multi-assignment problem, let $\alpha_i \rightarrow \infty$.

Using duality theory, the unconstrained multi-assignment problem becomes

$$\text{minimize } \sum_{i=1}^m \pi_i + \sum_{j=1}^n p_j + (n-m)\lambda \quad (\text{minimum cost network flow})$$

subject to

$$\begin{aligned} \pi_i + p_j &\geq a_{ij} & \forall (i, j) \in A & \quad (\text{complementary slackness}) \\ \lambda &\geq \pi_i & \forall i = 1, \dots, m & \quad (\lambda = \pi_i \text{ for multi-assigned target } i) \end{aligned}$$

where

π_i = profit of target i

p_j = price of canister j

λ = maximum profit

One method of solving the unconstrained multi-assignment problem is the forward/reverse auction algorithm proposed by Bertsekas (Reference 6). The algorithm is implemented as follows.

Forward Auction

Bidding Phase. For each target i that is unassigned under the assignment S , find the best canister j_i having best value v_i

$$\begin{aligned} j_i &= \arg \max_{j \in A(i)} \{a_{ij} - p_j\} \\ v_i &= \max_{j \in A(i)} \{a_{ij} - p_j\} \end{aligned}$$

and find the second best value

$$w_i = \max_{j \in A(i), j \neq j_i} \{a_{ij} - p_j\}$$

If j_i is the only canister in $A(i)$, then define w_i to be $-\infty$. Compute the bid of target i

$$b_{ij_i} = p_{j_i} + v_i - w_i + \varepsilon$$

where $\varepsilon < 1/n$.

Assignment Phase. For each canister j , let $P(j)$ be the set of targets from which j received a bid during the bidding phase of the iteration. If $P(j)$ is nonempty, increase p_j to the highest bid

$$p_j = \max_{i \in P(j)} b_{ij}$$

and remove from the assignment S any pair (i, j) and add to S the pair (i_j, j) , where i_j is the target in $P(j)$ attaining the maximum above.

Reverse Auction

For each canister j that is unassigned under the assignment S (if all canisters are assigned, the algorithm terminates), find best target i_j having best value β_j

$$i_j = \arg \max_{i \in B(j)} \{a_{ij} - \pi_i\}$$

$$\beta_j = \max_{i \in B(j)} \{a_{ij} - \pi_i\}$$

and find the second best value

$$\omega_j = \max_{i \in B(j), i \neq i_j} \{a_{ij} - \pi_i\}$$

If i_j is the only target in $B(j)$, then define ω_j to be $-\infty$. Let

$$\delta = \min\{\lambda - \pi_{i_j}, \beta_j - \omega_j + \varepsilon\}$$

where $\lambda = \max_{i=1, \dots, m} \pi_i$ and $\varepsilon < 1/m$. Add (i_j, j) to the assignment S and set

$$p_j = \beta_j - \delta$$

$$\pi_{i_j} = \pi_{i_j} + \delta$$

If $\delta > 0$, then remove from the assignment S the pair (i_j, j') , where j' is the canister that was assigned to i_j under S at the start of the iteration. Continue iterating until all canisters are assigned.

Note that the forward auction proceeds up to the point where each target is assigned to a single distinct canister. Because some canisters are still unassigned, the reverse auction is used to assign the remaining unassigned canisters.

SWARM SIMULATION DESCRIPTION

CANISTER MODEL

In this section the model used to define the forces and moments acting on the canister is described (see Figure 3). The canister model includes the body-fixed seeker, Kalman LOS filter, guidance control, swarm control, divert autopilot, divert thrusters, attitude control system, canister dynamics, and IMU. The various blocks are described below.

Seeker Model

The body-fixed seeker samples the LOS tracking angles in azimuth and elevation, which are measured from the seeker centerline to the target. The sampled values are produced through the use of Gaussian noise sources to degrade the true LOS angles obtained from the kinematics module.

Seeker Vibration

In addition to LOS angular measurement noise, another noise source degrading the seeker measurement is the angular deflection of the seeker due to canister body vibration. The vibration is modeled as band-limited random noise with a specified standard deviation, obtained by passing random noise through a shaping filter.

Kalman LOS Filter

A Kalman filter model is used in the swarm simulation because it yields a tracking algorithm that provides optimal performance against manned maneuvering targets (Reference 7). Specifically, the Kalman filter is implemented as an angular filter having two states: LOS angle and LOS angular rate. As a pair they form a state vector. Measured seeker LOS angle, the filter input, is used to smooth the state vector at the current sample time and to estimate the state vector at the next sample time. LOS angular rate, the filter output, is used in the guidance law to compute acceleration commands for steering the canister toward the target. Because filtering takes place in the inertial reference frame (which is a stable frame of reference) and the attitude of the canister body is currently fixed relative to the inertial frame, the filtered angular-rate output can be used directly in the guidance law without having to perform any matrix transformations. This will not be the case once an atmosphere model is included in the simulation. The atmosphere will exert forces and moments on the canister, causing it to yaw, pitch, and roll. Seeker measurements will then need to be transformed from seeker frame to inertial frame using IMU measurements of canister body attitude.

To reiterate, the purpose of the two-state angular Kalman filter is to provide estimates of the true LOS rate, given measurements of the LOS angle to the target. The algorithm is implemented as follows. The angular form of the target equations of motion are given by the state transition equation

$$X(k+1) = \Phi X(k) + q(k)$$

where $X(k)$ is the true state vector at time k , $q(k)$ is the process noise representing the uncertainty in the target dynamics, and Φ is the state transition matrix given by

$$\Phi = \begin{bmatrix} 1 & T \\ 0 & 1 \end{bmatrix}$$

where T is the time between measurements (i.e., the sample time). Measurements at time k are given by the equation

$$Y(k) = HX(k) + r(k)$$

where $H = [1 \ 0]$ is the measurement matrix and $r(k)$ is the measurement noise representing the uncertainty in the measurement of $X(k)$. The Kalman gain at time k is calculated as

$$K(k) = S(k|k-1)H^T [HS(k|k-1)H^T + R]^{-1}$$

where $S(k|k-1)$ is the predicted covariance matrix at time k given measurements up to and including time $k-1$, and R is the covariance of the measurement noise $r(k)$. State smoothing (i.e., filtering) is done using

$$\begin{aligned} S(k|k) &= [I - K(k)H]S(k|k-1) \\ X(k|k) &= X(k|k-1) + K(k)[Y(k) - HX(k|k-1)] \end{aligned}$$

where $S(k|k)$ is the smoothed covariance matrix at time k given measurements up to and including time k , and $X(k|k)$ is the corresponding smoothed state vector estimate.

State prediction is accomplished using

$$\begin{aligned} S(k+1|k) &= \Phi S(k|k)\Phi^T + Q \\ X(k+1|k) &= \Phi X(k|k) \end{aligned}$$

where $S(k+1|k)$ is the predicted covariance matrix at time $k+1$ given measurements up to and including time k , and $X(k+1|k)$ is the corresponding predicted state estimate.

The recursive nature of the Kalman filter permits the smoothed state vector estimate $X(k|k)$ to be improved over time, based upon additional seeker LOS measurements $Y(k)$. The improved state vector estimate, like all the previous estimates, has minimum variance and hence maximum probability of representing the true state vector $X(k)$.

Maneuver Detection

Detecting target maneuvers is important because maneuvering targets are inherently more difficult to intercept than non-maneuvering targets. When a maneuver is detected, the virtual spring rest length is increased so the canisters are forced to spread out over a wider area, thereby increasing the probability that one of them will intercept an unpredictably maneuvering target. In the absence of a maneuver, the virtual spring rest length is decreased as a function of time-to-go, ensuring that all canisters in the swarm are closely clustered at time of target intercept.

A simple fading memory average of the innovations (i.e., measurement residuals) is used to detect if a target maneuver has taken place and is given by (Reference 8)

$$u(k) = \alpha u(k-1) + d(k)$$

with

$$d(k) = v(k)^T S(k)^{-1} v(k)$$

where $0 < \alpha < 1$, $v(k)$ is the innovation vector at time k , and $S(k)$ is the corresponding covariance matrix that was calculated during the Kalman filtering process. If $u(k)$ exceeds a threshold, determined empirically, then a maneuver has occurred. Assuming $v(k)$ is Gaussian distributed, $d(k)$ will have a chi squared distribution with n-DOF, where n is the dimension of the measurement vector.

Guidance Law

A proportional navigation guidance law is used in the swarm simulation. The Kalman filter estimate of LOS rate is used in the calculation of the desired acceleration to be applied to the canister, perpendicular to the LOS to the target. The commanded acceleration (in Gs) perpendicular to the LOS is

$$G_{cmd} = \left[\frac{\dot{K} \dot{\lambda} \dot{R}}{g} + G_x \sin(\lambda) \right] \frac{1}{\cos(\lambda)}$$

where

- K = proportional navigation gain (unitless)
- $\dot{\lambda}$ = LOS angular rate (rad/sec)
- λ = LOS angle (rad)
- $\dot{R}$ = closing velocity (ft/sec)
- g = gravity constant (ft/sec²)
- G_x = acceleration of canister along longitudinal axis (Gs)

Divert Autopilot

The simulation uses a propulsive divert autopilot that computes commanded thrust for use by the divert thrusters, mounted at the canister center of gravity, to achieve the desired acceleration. Using the G command from the guidance law, G_{cmd} , and canister mass, m_{can} , the divert autopilot calculates the commanded proportional thrust, F_{cmd} , as

$$F_{cmd} = G_{cmd} m_{can} g$$

where g is the constant of gravity.

Divert Thrusters

Because a model of the divert thrusters (e.g., the pulse-width-modulation of thruster force) is not currently implemented in the simulation, the proportional thrust command, computed in the previous section, is used “as is” in the canister equations of motion detailed in the next section.

Canister Dynamics

Linear and angular accelerations of the canister are the result of forces and moments acting on the canister. Within the swarm simulation, the canister linear accelerations are computed via the canister equations of motion

$$\dot{u} = f_x / m + g_x + v \cdot r - q \cdot w$$

$$\dot{v} = f_y / m + g_y - u \cdot r + p \cdot w$$

$$\dot{w} = f_z / m + g_z + u \cdot q - p \cdot v$$

where $\dot{u}$, $\dot{v}$, $\dot{w}$ are the x,y,z components of canister linear acceleration; p, q, r are the canister Euler angular rates about the x,y,z axes; g_x, g_y, g_z are the components of gravity; and f_x, f_y, f_z are the components of external force acting on the canister of mass m , all expressed in the canister frame of reference. Likewise, the angular accelerations of the canister are given by

$$\dot{p} = m_x / I_{xx}$$

$$\dot{q} = m_y / I_{yy} + k \cdot p \cdot r$$

$$\dot{r} = m_z / I_{zz} - k \cdot p \cdot q$$

where $\dot{p}, \dot{q}, \dot{r}$ are components of canister angular acceleration about the x,y,z axes of the canister frame; I_{xx}, I_{yy} are the canister's principle moments of inertia about the x and y axes ($I_{yy} = I_{zz}$ due to canister symmetry about the y and z axes); m_x, m_y, m_z are the components of external moments acting about the canister frame; and $k = (I_{yy} - I_{xx}) / I_{yy}$.

Given the equations above, it is clear that the canister model is a full 6-DOF model. Compare these equations to those of the simpler 3-DOF target motion model to be discussed.

IMU

The IMU model includes the effects of rate sensor drift and random noise errors on the IMU measurement of canister attitude. The model sums the error rates due to gyro drift and random walk noise terms, integrates the sum to obtain the attitude error, and then adds this error to the canister attitude.

GPS Receiver

Errors due to inaccuracies in GPS measurements of canister position are modeled in the swarm simulation by adding a small Gaussian noise term to the canister's true position. The position measurement is used to compute the G commands required to maintain virtual coupling between canisters. It is also used to compute the xy-position of targets falling within the seeker FOV via seeker measurement of target angular position and the laser altimeter measurement of canister altitude. The target positions are transmitted to all canisters in the swarm and aid in the assignment of canisters to targets using the multi-assignment algorithms previously described in the Target-Canister Pairing section.

TARGET AND CANISTER DISCUSSION

TARGET MOTION MODEL

Because target motion is confined to the ocean surface (the xy-plane), a simplified planar point-mass model (Reference 6) can be adopted as

$$V_x = V \cos(\theta)$$

$$V_y = V \sin(\theta)$$

$$\dot{\theta} = \frac{G_{cmd} \cdot g}{V}$$

where V is the target velocity, θ is the heading, $\dot{\theta}$ is the turning rate, G_{cmd} is the commanded lateral acceleration (in Gs), and g is gravitational acceleration. To model the motion of any ocean target, we only require the maximum linear velocity and maximum turning rate of the actual vessel, and use these as upper limits in the equations above. Target position and heading are then obtained via numerical integration of the target velocity and turning rate, respectively.

RESOLVING COMMONALITY OF TARGETS

Because each canister in the swarm transmits information concerning the positions of targets currently within its seeker FOV, if two or more FOVs overlap there is likely to be redundant target information in the global knowledge database. The redundant information must be culled from the database before meaningful target–weapon pairing can be done. For example, consider two canisters that are imaging the same area on the ocean and there is only one target in that area. Both canisters will measure and then transmit the location of the target (relative to some inertial reference frame) to the global database, thereby generating a redundant entry in the database (i.e., there is only one target, but there are two entries for that target). To make matters worse, each canister is making a statistically independent measurement of the target position. Thus the position values that each transmits will not be identical due to systematic errors (e.g., uncertainties in measurement of canister altitude, uncertainties in LOS angle to the target, IMU errors, gyro drift rate and bias, etc.). Fortunately, the solution to this problem is simple: to eliminate redundant entries, the database should be scanned for target positions that are nearly identical (i.e., scan for target positions that are closely spaced, and eliminate all but one). The threshold used to determine whether targets are closely spaced should be a function of the canisters' systematic errors, so that higher measurement uncertainties result in a higher threshold. This is equivalent to placing an error basket (the threshold) of radius R around a measured target position and scanning the database for any other target positions that fall within that basket, as illustrated in Figure 8. If canister #1 measures the target position to be (x_1, y_1) and canister #2 measures it to be (x_2, y_2) , then they are

detecting one target (i.e., the *same* target) within the overlapping areas of their seeker FOV, if the following inequality is true:

$$(x_1 - x_2)^2 + (y_1 - y_2)^2 < R^2$$

This will obviously be true because in our example scenario there is only one target. Consider the case where there are now two targets being detected by two canisters. If the inequality above is *not* true, then each canister is independently detecting a separate target (within a non-overlapping area of its seeker FOV) and therefore each entry will remain intact in the database because two targets are actually present. However, if the two targets are closely spaced such that the above inequality is true, then the targets are unresolvable and only one entry will remain in the database. Database accuracy can be improved slightly if, instead of eliminating the redundant entries for a particular target, they are averaged to form an improved entry, because the average of several independent measurements is better than one measurement.




FIGURE 8. Two Canisters Making Statistically Independent Measurements of the Same Target.

CANISTER HARDWARE COMPONENT MOCK-UP

Given the submunition family sizing constraints, one possible layout of canister hardware components is illustrated in Figure 9. The front end of the canister contains a body-fixed IR seeker consisting of a focal plane array (FPA) sensor and associated image signal processing hardware and a laser altimeter. Next comes the central processing unit (CPU) that executes the tracking, guidance, and target–weapon assignment algorithms. This section also contains the IMU, power supply, and hardware for wireless network communication with other members of the swarm. The next section contains the

ordnance, safe-arm device, and fuse. Divert thruster control and nozzles are next, followed by the thruster propellant. The tail section contains stabilization fins, GPS antenna, and GPS receiver. Placement of the GPS antenna at the aft end of the canister not only increases signal reception but is also a low cost solution to the potential ground-based jamming vulnerability of GPS-aided tactical weapon systems.




FIGURE 9. Canister Hardware Component Mock-Up.

SIMULATION NOISE SOURCES

Experience has shown that no measurement, however carefully made, can be completely free of uncertainties (Reference 9). As the number of measurements increases, their distribution approaches some definite continuous curve, called a limiting distribution. Different types of measurements have different limiting distributions. However many measurements are found to have a Gaussian distribution. Therefore, whenever a measured value is needed in the simulation, the true value is degraded with Gaussian noise of a specific standard deviation in order to reflect the uncertainty in the measured value. Gaussian noise sources are specified in the simulation according to their 1-sigma standard deviation from the true value. Several noise sources are modeled in the simulation, including uncertainties in the GPS measurement of canister horizontal and vertical position, drift rate and random walk of the gyros within the IMU, linear accelerometers, seeker LOS pointing angle, seeker vibration due to divert thruster operation, and Kalman filter noise terms (see Table 1).

TABLE 1. Noise Sources Currently Modeled in the Simulation
and a Representative Value of the 1-Sigma Uncertainty.

Noise Source	Units	Value*	Comments
Canister xy-Position via GPS	ft	12.5	GPS update rate = 1 Hz
Canister Altitude via GPS	ft	24.3	Not used. For comparison only.
Canister Altitude via Laser Altimeter	ft	0.33	Used instead of GPS measurement
IMU Gyro Drift Rate	deg/hr	3.9	3σ drift rate
IMU Gyro Random Walk	deg/ $\sqrt{\text{hr}}$	0.03	12.7 deg/hr rms @ 20ms guidance sample rate
IMU Accelerometer	mg/ $\sqrt{\text{Hz}}$	6.6	Noise in 50 Hz bandwidth = 46.7mg rms
Seeker LOS Pointing Angle	deg	0.08	20 deg / 256 pixels
Seeker Vibration Bandwidth	Hz	10	Vibration modeled as band-limited random noise
Kalman Filter Process Noise	rad/s ²	0.02	Tuned for optimum target maneuver detection
Kalman Filter Measurement Noise	rad	0.014	Seeker LOS pointing angle expressed in radians

* 1σ unless stated otherwise

SIMULATION EVALUATION USING BOX PLOTS

BOX PLOTS

Box plots are useful for identifying outliers and for comparing distributions, particularly when there are a large number of observations. Other techniques for visually representing data are available, such as histograms and scatter plots, but box plots have several important advantages (Reference 10):

1. The chosen percentiles can be compared effectively.
2. By graphing large and small values, unusual values (outliers) are not swept under the rug, as they often are when the summary of the distribution consists of a sample mean and a sample standard deviation.
3. Box plots can be used even when the number of distributions is large.

Figure 10 illustrates two familiar distributions and their corresponding box plots. The Gaussian distribution shown is symmetric about mean value zero, so its corresponding box plot is also symmetric. Note the equal number of outliers below the 10th and above the 90th percentiles. The Rayleigh distribution, on the other hand, is lopsided, which is clearly evident in the corresponding box plot.




FIGURE 10. Gaussian and Rayleigh Probability Distributions and Corresponding Box Plots.

To illustrate the three box plot strengths, seven distributions are compared using box plots in Figure 11. The data used here are taken from actual simulation output (more detail is available in the Simulation Results section). The y axis indicates the Pk of the target and the x axis indicates the ratio of canisters per target as the number of canisters assigned to this target is increased. By comparing these distributions it is clear that in order to achieve an average Pk (50th percentile) of 90% or better, we must assign at least 4 canisters to this target.




FIGURE 11. Seven Distributions Compared Using Box Plots.

SIMULATION SETUP

Unless otherwise stated, all of the Monte Carlo simulations that were run over the course of this study used the target formation shown in Figure 12a. Targets were arranged to form a hexagon, with a target at each vertex and one at the hexagon center. At the start of the simulation all targets move rightward as a group at 10 m/s (32.8 ft/s). Two seconds into the simulation, the two upper targets break away from the group and move upward, and the two lower targets break away and move downward. Thus, we have 4 targets that maneuver, by undergoing constant lateral acceleration, and three targets that do not maneuver. We chose this arrangement because it allows comparison of Pk results for maneuvering targets against those that do not maneuver (our “control group”). Because the non-maneuvering targets should be easier to intercept, we expect higher Pk for those targets. For a single-assigned target the goal was to maneuver the canister as close as possible to the target. For a multi-assigned target (i.e., a swarm of multiple canisters assigned a target) the goal was to maneuver the swarm center-of-mass as close as possible to the target.

The canisters were arranged in a lattice configuration directly over the targets (Figure 12b). At the start of the simulation the canisters were dropped, at rest, in this configuration from an altitude of 500 feet. When a particular study required more canisters, they were added to the outer edge of the configuration so that the lattice was

maintained. When fewer canisters were required, they were removed from the configuration starting at the outer edge of the lattice furthest from the center (the canister above target #1 is at the center of the lattice). Next, a target–weapon assignment algorithm was used to assign the canisters to the targets such that the global probability of intercepting all the targets was maximized. A color-coded example of a 3 canister-per-target assignment is presented below.




FIGURE 12. Initial Target and Canister Lattice Configuration.

Pk VERSUS NUMBER OF CANISTERS

It seemed reasonable to expect that the Pk should increase as a function of the number of canisters assigned to each target. To verify this, a 1000-run Monte Carlo simulation was performed for a canister-to-target ratio of one (i.e., 1 canister per target). The ratio was incremented and another 1000 runs were made. This continued until the ratio reached 7. In each case, the canisters were dropped from a height of 500 feet onto the hexagonal group of targets. The goal was to get each swarm center-of-mass as close as possible to its assigned target. The maneuvering targets began pulling constant lateral acceleration of one-third G 2 seconds into the simulation. Figures 13 and 14 are box plots showing the results of Pk as a function of the canister-to-target ratio.




FIGURE 13. Pk as a Function of Number of Canisters Per Target for Targets 1 Through 4.




FIGURE 14. Pk as a Function of Number of Canisters Per Target for Targets 5, 6, and 7.

From these results we can conclude three important facts.

1. The non-maneuvering targets (T1, T4, and T7) are indeed easier to intercept than those maneuvering with constant lateral acceleration, and it does not matter how many canisters are assigned to those targets because the Pk is essentially 100% for any canister-to-target ratio.

2. Concerning the maneuvering targets, in order to achieve an average Pk (50th percentile) of 90% or better, we must assign at least 5 canisters to each target. We now have a lower limit on the required number of canisters to achieve a particular Pk, and it is determined using the targets that maneuver. Incidentally, for air-to-air missile encounters the rule of thumb for a one-on-one encounter is that the weapon G capability should be at least 3 times that of the target G capability in the endgame to ensure a high Pk. Because our canisters have a one-half G capability and the targets have one-third G capability, this leads to a canister-to-target G ratio of 1.5. Even though our G ratio is low, we still achieve high Pk as long as we assign at least 5 canisters per target.

3. The initial lattice configuration of the canisters has an effect on the target Pk. Note that when there are three canisters per target, we get a Pk of 100% on T5, but T2,

T3, and T6 have Pk (again, 50th percentile) in the 70 to 80% range. This is entirely due to the initial configuration of the canister lattice, as shown in Figure 15. When the simulation ends (i.e., when the canisters reach zero altitude), note that in the final configuration there is a canister very near target T5, which results in very high Pk. This is not the case for the other maneuvering targets, which are located near the 70% Pk region of their nearest canister. The final configuration is a direct result of the initial configuration as seen in Figure 15, and thus the initial configuration affects the Pk. However the initial configuration has less impact on the Pk as the number of canisters per target is increased. For our canister-to-target G ratio of 1.5, the optimum attack configuration is a hexagonal pattern wherein 7 canisters are assigned per target, which results in a target Pk above 95%. The hexagonal configuration is optimal because it provides the most spatial coverage around the swarm center-of-mass for a given number of canisters.




FIGURE 15. Initial and Final Lattice Configuration When Three Canisters Are Assigned to Each of Seven Targets.

VARIATION OF CANISTER DROP HEIGHT

We increased the canister drop height from 500 to 1000 feet and repeated the previous study. The results indicate that higher Pk is achievable due to the additional altitude, which provides about 2.5 seconds more maneuvering time for the canisters. There is, however, a drawback to increasing the drop height to 1000 feet. At this height the maneuvering thrusters must be capable of delivering about 8 seconds of impulse, which may not be possible given the small size constraint of the canister.

Pk VERSUS CANISTER-TO-TARGET G RATIO

In the Pk Versus Number of Canisters section of this report we came to the conclusion that for a canister-to-target G ratio of 1.5, the optimum attack configuration is a hexagonal pattern wherein 7 canisters are assigned per target. This attack configuration resulted in a target Pk above 95%. The questions now are: (1) Just how low can the canister-to-target G ratio go before the Pk drops to an unacceptable level? and (2) Do we gain or lose anything by using a higher G ratio? To answer these questions, 7 canisters were arranged in an hexagonal attack configuration, and the canister-to-target G ratio was varied from 0.5 to 3 (the value used as a rule of thumb for one-on-one air-to-air missile encounters). The results are shown in Figure 16, where each box plot represents a 1000-run Monte Carlo. It is clear from this figure that a G ratio of 1.5 *is* the lowest acceptable ratio that will maintain a high target Pk (e.g., above 95%). Lower G ratios force the Pk rapidly to zero, while higher G ratios are unnecessary but welcome.




FIGURE 16. Pk as a Function of Canister-to-Target G Ratio, Using 7 Canisters in an Hexagonal Attack Configuration.

Pk VERSUS CONSTRAINED AND UNCONSTRAINED MULTI-ASSIGNMENT

To investigate the effect of constrained and unconstrained multi-assignment on target Pk, 28 canisters were assigned to 7 targets. In the constrained case, the number of canisters per target was balanced so that 4 canisters were assigned to each of the 7 targets. The assignment algorithm then maximized the global probability of intercept while simultaneously maintaining the constraints. In the unconstrained case, the multi-assignment algorithm was allowed to assign canisters to targets such that the global probability of intercept was maximized while at least one canister was assigned per target. The results are shown in Figure 17, where each box plot represents a 1000-run Monte Carlo. The gray box plots represent the Pk distribution for the constrained case, and the blue box plots represent the Pk distribution for the unconstrained case. The figure shows that unconstrained assignment increases the Pk slightly on targets 1, 2, and 3 at the expense of a large (14%) decrease on target 5. This large drop is due to target 5's (a maneuvering target) being assigned one less canister while target 4 (a non-maneuvering target) is assigned an additional canister. The current unconstrained multi-assignment algorithm does not take into account the fact that maneuvering targets are harder to intercept, and therefore should be assigned *more* canisters than non-maneuvering ones. Therefore, we can achieve better overall Pk by constraining a balanced number of canisters per target.




FIGURE 17. Pk Versus Constrained and Unconstrained Multi-Assignment of 28 Canisters to 7 Targets.

Pk VERSUS CANISTER SPATIAL DENSITY

The goal of this study was to determine the effect canister spatial density had on target Pk. Canister spatial density is simply the number of canisters per unit of area, and is a direct result of the virtual spring rest length that defines the quiescent separation distance between canisters. To carry out this study, 7 canisters were arranged in a hexagonal configuration and dropped from 500 feet onto one target traveling at 10 m/s. Two seconds into the simulation, the target began pulling constant lateral acceleration of 0.5 G. Because the canister also has a 0.5-G maneuver capability, the canister-to-target G ratio is 1.0. The virtual spring rest length was stepped from 2 to 12 meters in increments of 2 meters. At each step a 1000-run Monte Carlo was performed and the results plotted (see Figure 18). Figure 18 indicates that average Pk is maximized when the rest length is set to 8 meters. Pk drops as the rest length deviates from 8 meters because of the relatively low G ratio. This is a very challenging scenario in which the target has the same maneuver capability as the canisters, so as the rest length is decreased, the canisters become more tightly clustered about the swarm center-of-mass (i.e., high spatial density) and the probability that one of the canisters will intercept, and thereby kill, the maneuvering target decreases because the target can out-maneuver the canisters. On the other hand, as the rest length is increased from 8 meters, the canisters become loosely clustered about the swarm center-of-mass (i.e., low spatial density) and the target can slip between the individual canisters, again resulting in low Pk.




FIGURE 18. Pk as a Function of Virtual Spring Rest Length for Canister-to-Target G Ratio of 1.0.

The study was repeated for a slightly less challenging scenario in which the target can pull a constant lateral acceleration of one-third G (i.e., canister-to-target G ratio = 1.5). The results are shown in Figure 19. Note that in this case the target Pk is relatively constant as the virtual spring rest length is increased from 2 to 12 meters. This is due entirely to the canisters' higher maneuverability.




FIGURE 19. Pk as a Function of Virtual Spring Rest Length for Canister-to-Target G Ratio of 1.5.

CANISTER SPATIAL COVERAGE VERSUS TARGET DEPLOYMENT

Another important question to be answered concerns canister spatial coverage, that is, how far from the target can the canisters be dropped and still intercept it with reasonable Pk (e.g., above 90%)? In other words, what is the maximum canister-to-target initial range? To answer this question, 7 canisters were arranged in an hexagonal configuration and dropped from 500 feet onto one non-maneuvering target traveling at 32.8 feet/second perpendicular to the LOS between canister and target (i.e., the target was traveling cross-range). The initial range between canister and target was stepped from 200 to 300 feet in increments of 20 feet. At each step a 1000-run Monte Carlo was performed, and the results plotted in Figure 20. The figure indicates that the initial range to target should be no greater than 260 feet if the canister is to intercept it with Pk greater than 90%. This is not bad considering that the canisters are dropped from 500 feet, *at rest*. In reality, the deployment vehicle will provide an initial forward momentum to the canisters that will ultimately increase their range (i.e., spatial coverage).




FIGURE 20. Pk as a Function of Target Range.

NOISE PARAMETER TRADE-OFF (NOISE SENSITIVITY)

Because LOS noise is the dominant noise source in the simulation to date, we restricted the noise sensitivity study to this particular type of noise. Figure 21 illustrates how increasing the magnitude of the LOS noise affects the endgame performance of the swarm. Gray box plots represent the Pk distribution using nominal LOS noise (1-sigma = 0.08 degree) while blue box plots represent the Pk distribution for a 10% increase in LOS noise. Note that performance is only slightly degraded because the worst-case shift is a 6% decrease in Pk, and the decrease becomes insignificant once the canister-to-target ratio rises above 4. Also, note that there is essentially no performance degradation when the swarm is intercepting non-maneuvering targets. Therefore we can conclude that the simulation is quite insensitive to the LOS noise source.




FIGURE 21. Pk Versus Number of Canisters Per Target. Gray box plots represent the Pk distribution using nominal LOS noise while blue box plots indicate the Pk distribution after a 10% increase in LOS noise.

SWARM DROP VERSUS BASELINE CANISTER DROP

The purpose of this study was to compare the performance of a network of swarming canisters against a baseline canister drop consisting of a lattice of canisters that do not collect any information—there is neither a seeker nor a GPS receiver—and there is no communications link between the canisters (i.e., these are “dumb” canisters). Using the same canister–target engagement scenario as described earlier, 49 canisters were dropped from an altitude of 500 feet, at rest, directly over 7 targets. A 1000-run Monte Carlo was performed, first in “swarm mode” where guidance and virtual coupling were enabled, and then rerun with those capabilities disabled (i.e., “baseline” or “dumb” mode). The results of these two runs indicated that the dumb canisters were at a distinct disadvantage because they were dropped over the initial position of moving targets. By the time the dumb canisters hit the ocean (5.6 seconds after release) the targets were far outside the area of impact, and as a result each target received a Pk of zero. This result shows the poor performance of the dumb canisters, but it does not allow us to gage the sensitivity of the performance when varying a particular parameter. For example, how does a change in canister density affect the Pk? Given the current scenario, there would be no change in

the dumb canisters' probability of killing the targets: the Pk would remain zero. To remedy this situation, we decided to bias the drop in favor of the dumb canisters. This was accomplished by moving the canister lattice directly over the area the targets would occupy when the lattice hit the ocean, and we increased the number of canisters in the lattice until the entire target area was covered (this required 75 canisters). It is important to note that this was only done for the lattice of dumb canisters; the swarming canisters received no such bias. Also, in order to produce a meaningful Monte Carlo run using dumb canisters, which do not have guidance-related noise sources, we had to randomly adjust the lattice location at the beginning of each run by a small amount equal to the canister lethality radius. The results of a 1000-run Monte Carlo in the "swarm mode" are shown in Figure 22a and the corresponding results in "dumb mode" are shown in Figure 22b. Note the Pk (50th percentile) is essentially 100%, with almost no variance, on every target for the swarm drop, but around 90% with wide variance on every target for the baseline drop. A Pk of 90% might first appear to be acceptable, but keep in mind the fact that 75 canisters were required to achieve that value, and the distribution of Pk ranges from 60 to 100%.

The results of Figure 22 were achieved using an inter-canister spacing of 19.7 feet. How does increasing the inter-canister spacing (i.e., decreasing canister density) affect performance? This question was answered by repeating the study exactly as before except that the inter-canister spacing was increased to 29.5 feet (a 50% increase). The results are shown in Figure 23, where it is immediately obvious just how poorly the dumb canisters perform. The results for the swarming canisters, Figure 23a, are insensitive to the increase in inter-canister spacing while, conversely, the results for the dumb canisters, Figure 23b, are extremely sensitive to inter-canister spacing.




FIGURE 22. Pk Results for a Swarm Drop Versus Biased Baseline Canister Drop With Inter-Canister Spacing of 19.7 Feet.




FIGURE 23. Pk Results for a Swarm Drop Versus Biased Baseline Canister Drop With Inter-Canister Spacing of 29.5 Feet.

SWARM DROP WITH AND WITHOUT VIRTUAL COUPLING

One of the basic premises of Swarm CAT is the need for virtual coupling, which results in formation flying, but what exactly do we gain by flying in formation? Why not allow the canisters to unilaterally pursue a target once they are assigned to it, thereby eliminating the need to maintain a communications link among canisters during the entire duration of the drop? Would this adversely affect the Pk? The answer to this question is yes. Figure 24 shows a comparison of the results of 1000-run Monte Carlos for swarm drops with and without virtual coupling. The engagement scenario was as described previously, where 49 canisters were dropped onto 7 targets. In Figure 24a (a repeat of Figure 22a) the canisters were virtually coupled and in Figure 24b the canisters were not virtually coupled. Note that the non-maneuvering targets (targets 1, 4, and 7) receive high Pk whether or not virtual coupling is used. However, the maneuvering targets (targets 2, 3, 5, and 6) are harder to kill without the use of virtual coupling.




FIGURE 24. Comparison of Pk Results for a Swarm Drop With and Without Virtual Coupling.

CONCLUSIONS

This section recaps the results of the Monte Carlo studies reported here. These studies included

1. Pk Versus Number of Canisters
2. Variation of Canister Drop Height
3. Pk Versus Canister-to-Target G Ratio
4. Pk Versus Constrained and Unconstrained Multi-Assignment
5. Pk Versus Canister Spatial Density
6. Canister Spatial Coverage Versus Target Deployment
7. Noise Parameter Trade-Off (Noise Sensitivity)
8. Swarm Drop Versus Baseline Canister Drop
9. Swarm Drop With and Without Virtual Coupling

Having completed these studies, several conclusions can be reached given the current state of grid simulation. Endgame performance is only slightly degraded by a 10% increase in nominal noise terms. Thus the grid simulation is rather insensitive to current noise sources, but more noise modeling and latency need to be included (detailed examples are discussed in the next section). For air-to-air missile encounters, the rule of thumb for a one-on-one encounter is that the G capability of the weapon must be 3 times that of the G capability of the target in the endgame to ensure a high probability of intercept. Because our canisters have a one-half G capability and the targets have a one-

third G capability, this leads to a canister-to-target G ratio of 1.5. Even though our G ratio is low compared to the rule-of-thumb value, we still achieve high Pk as long as we assign at least 5 canisters per target. Arranging 7 canisters in an hexagonal attack configuration provides the highest possible Pk (at least 95%) for the given canister-to-target G ratio. The hexagonal configuration is optimal because it provides the most spatial coverage around the swarm center-of-mass for a given number of canisters. Should the center canister miss the target, the outer canisters have a good chance of killing it. The relationship of G ratio, canister assignments, and Pk will be revisited after more latency, additional noise models, and an atmosphere model are added to the simulation in future studies.

FUTURE WORK

The goal of the following potential tasks is to extend the simulation to further quantify the swarm hypothesis. These tasks correspond to various modules that provide an additional level of realism within the simulation.

IMAGING SEEKER MODEL

An imaging seeker model is needed because it will facilitate the implementation and testing of detection and tracking algorithms using both simulated imagery and real imagery (collected via helicopter) of various targets. The imaging seeker model degrades the image by a series of transforms that introduce optical blurring and transmission effects, platform jitter (due to the divert thrusters and wind gusts), and detector quantization and non-uniformity. The degraded image is then passed to the detector/tracker, which detects the target and establishes a track on the angular position of the target relative to the seeker centerline.

ATMOSPHERIC EFFECTS ON CANISTER MANEUVERABILITY

Currently, the simulation consists of point-mass canisters maneuvering within a vacuum. While this was adequate for closing the guidance loop so that various multi-assignment algorithms could be investigated, a more realistic model is needed to study the effect the atmosphere will have on swarm formation and canister maneuverability. With this goal in mind, the point-mass model will be extended to allow full 6-DOF rigid-body motion with the additional inclusion of an atmospheric model. Aerodynamic drag forces and pitching moments will affect canister maneuverability and endgame performance. The main effort of this task is not to develop the 6-DOF model because the 6-DOF equations of motion are already in the simulation, but rather to allow the canister to yaw, pitch, and roll in response to atmospheric effects. The atmospheric model under consideration has been widely used in air-to-surface missile simulations as a first-order

approximation, while the canister aerodynamic coefficients will be obtained from standard equations for cylindrical bodies.

Because the ultimate shape and aerodynamic performance of the submunition canister are currently unknown, the computer program Missile Datcom may also be used to generate the aerodynamic coefficients that are used by the simulation to compute normal and axial forces and pitching moments acting on the canister. Missile Datcom is an aerodynamic design tool (a wind-tunnel simulator) that has the predictive accuracy suitable for preliminary design, and will allow quick and economic estimation of the aerodynamics of various submunition configurations.

EXTEND 2-D PLANAR VIRTUAL COUPLING TO 3-D

The canisters in the swarm simulation are currently confined to a planar grid (i.e., we have a grid simulation) in which virtual coupling between canisters takes place. Once the atmosphere model is included in the simulation, the canister swarms will no longer be confined to a two-dimensional (2-D) grid-like configuration but will instead form a 3-D cluster (or “cloud”). Therefore, the simple 2-D virtual coupling currently being implemented must be extended to handle 3-D swarm clusters.

WIND GUST NOISE MODEL

Wind gusts acting on the lightweight canisters are another noise source that should be included in the simulation. Three models for implementing this type of noise are available in a published paper by Gawronski (Reference 11). Two of the models are directly applicable for simulating the disturbances on a canister due to wind, and they simply need to be introduced into the swarm simulation as additional sources of unknown errors. The first model represents wind forces on the canister, and it closely represents actual wind action. Wind forces are obtained by applying white noise of unit standard deviation to a Davenport filter (filtered on the spectrum of the Davenport wind gusts), and then multiplying by an appropriate scale factor. The second model represents wind gusts as time-varying torque disturbances applied to the canister attitude control system (ACS) input. The Davenport filter is again used to shape white noise into wind gusts, but a different scale factor is now used. In both models, the total wind velocity is computed as a combination of a steady-state (or mean) velocity and a turbulence (gust) velocity. The gust component is a random process with zero mean and a “Davenport” spectrum. Because this spectrum depends on average wind speed and terrain roughness, the wind gust model can be adapted to any type of terrain.

WIRELESS COMMUNICATIONS MODEL

Modeling data processing delays and time-division radio transmission schemes are crucial in our understanding of the timing constraints placed on each canister. For example, if each canister is allotted a 1-millisecond window in which to transmit data to other canisters in a swarm of 500 canisters, then the update rate (the sampling rate) of the entire swarm is 2 hertz. That is, the entire swarm is completely updated with fresh information, such as position and velocity of all canisters and targets, twice every second—even though the sampling rate of each canister may be 20 hertz. Therefore, subsequent simulation improvements will include a more refined model of the actual network communication system in order to investigate whether this type of time-division transmission scheme is appropriate.

COMPUTATIONAL DELAY

Currently, the simulation assumes that all computations within the guidance loop can be completed at the guidance sample rate frequency. This may not be the case in the actual hardware. The LOS rate filter output should be delayed in order to model a data processing delay in the guidance loop. The delay would effectively encompass the time required for calculating the Kalman LOS rate filter equations, the swarm control equations, and the guidance equation. An array of 20 elements, for example, could be used to store the previous LOS rate values. The value of the delay, which must be less than or equal to 20, represents the number of guidance sample periods that the data are to be delayed.

DIVERT THRUSTER MODEL

Divert thrusters operate in a pulse-width modulation (PWM) mode and are driven by thrust commands (i.e., average commanded thrust) from the divert autopilot. Solid-fueled thrusters, for example, operate in pairs and are connected by a diverter valve. Once ignited, the thruster continues to expend impulse in order to avoid over-pressurizing the combustion chamber. During each PWM period, a square-wave signal is used to toggle the diverter valve from one thruster to the opposing thruster, resulting in maximum thruster output with a duty-cycle proportional to the commanded thrust. A model of the PWM operation will provide additional simulation realism in the form of thrust latency.

REFERENCES

1. Wright-Patterson Air Force Base. *On The Use of Quaternions in Simulation of Rigid-Body Motion*, by A. C. Robinson. Ohio, W-P AFB Aeronautical Research Laboratory, December 1958. (WADC TR 58-17, publication UNCLASSIFIED.)
2. V. Gazi and K. Passino. "Stability Analysis of Swarms," *IEEE Trans. Automatic Control*, Vol. 48, No. 4 (April 2003). Pp. 692-697.
3. D. P. Bertsekas. *Network Optimization: Continuous and Discrete Models*. Belmont, Massachusetts, Athena Scientific, May 1998.
4. D. P. Bertsekas, D. A. Castañón, and H. Tsaknakis. "Reverse Auction and the Solution of Inequality Constrained Assignment Problems," *SIAM J. Optimization*, Vol. 3, No. 2 (May 1993). Pp. 268-297.
5. J. L. Kennington and R. V. Helgason. *Algorithms for Network Programming*. New York, John Wiley and Sons, 1980.
6. J. A. Roecker and C. D. McGillem. "Target Tracking in Maneuver Centered Coordinates," *Proc. IEEE National Radar Conf.* (1988). Pp 68-73.
7. R. A. Singer. "Estimating Optimal Tracking Filter Performance for Manned Maneuvering Targets," *IEEE Trans. Aerospace and Electronic Systems*, Vol. AES-6, No. 4 (July 1970). Pp. 473-483.
8. A. Bicchi and L. Pallottino. "On Optimal Cooperative Conflict Resolution for Air Traffic Management Systems," *IEEE Trans. Intelligent Transportation Systems*, Vol. 1, No. 4 (December 2000). Pp. 221-232.
9. J. R. Taylor. *An Introduction to Error Analysis*. Mill Valley, California, Oxford University Press University Science Books, 1982.
10. AT&T Bell Laboratories. *The Elements of Graphing Data*, by W. S. Cleveland. Murray Hill, New Jersey, AT&T, 1985.
11. California Institute of Technology Jet Propulsion Laboratory. *Three Models of Wind-Gust Disturbances for the Analysis of Antenna Pointing Accuracy*, by W. Gawronski. Pasadena, California, JPL, 15 May 2002. (JPL IPN Progress Report no. 42-149, report UNCLASSIFIED.)

This page intentionally left blank.