ICAS-94-8.6.1

THE HELICOPTER AIR-TO-AIR VALUE-DRIVEN ENGAGEMENT MODEL (HAVDEM)

Nathaniel E. Bent
Director of Rotorcraft Simulation
Decision-Science Applications, Inc.
Arlington, Virginia

Abstract

HAVDEM is a non-real-time engagement level simulation of helicopter air-to-air combat. It addresses the Army's need for a constructive analysis tool that not only models helicopter hardware systems, but also models pilots' decisions about how those systems are used. Terrain utilization by the pilot decision logic is one of the primary objectives of the HAVDEM program. The combination of government-supplied computer software, a pilot decision methodology first used in fixed wing simulation, and an original algorithm for terrain utilization, HAVDEM is used primarily to compare the mission effectiveness of proposed system improvements to their existing counterparts. A careful comparison study requires that a statistically significant number of simulation runs be performed. For the analysis of conceptual systems, constructive simulation is more cost effective than virtual reality and its analyses are reproducible.

Introduction

Current interest in potential air-to-air missions for the attack helicopter has generated a requirement for simulations representing a realistic combat environment. Several existing simulations contain good representations of the machines used in combat but require a man-in-the-loop (virtual reality) or extensive input data to direct these assets. While virtual reality has its distinct benefits, constructive simulation of the pilot decision process offers a relatively inexpensive tool for understanding air-to-air missions. HAVDEM provides this capability by combining the physical models from the Helicopter Piloted Air Combat Model (HELIPAC) rotary wing simulation with the value-driven pilot decision methodology used in the TAC BRAWLER fixed wing simulation. Additionally, HAVDEM contains a selection of six degree of freedom (6-DOF) helicopter kinematics models and an innovative decision logic that utilizes terrain shielding during air-to-air engagements and waypoint following.

HAVDEM was originally funded by the Army via Small Business Innovative Research (SBIR) Phase I and Phase II contracts which ended December 1991. A simulation allowing an engagement of up to twenty helicopters and SAMs was delivered at this time. The US Army Aviation Applied Technology Directorate (AATD) provided government contract management. An SBIR Phase III effort was then sponsored by AATD to study the mission effectiveness of a modified Apache airframe. HAVDEM was initially developed under the Unix operating system on a MIPS 2030 Workstation using the FORTRAN programming language. Currently the development platform is the Silicon Graphics Indigo. The source code, object code, executables, and on-line documentation currently require about 20 megabytes of disk memory although a larger partition is required for temporary files during the reorganization of digitized terrain elevation data. The current simulation runs a three-versus-three engagement in real time. Recent tests and graphical displays have shown that the innovative methodology for terrain utilization generates flight paths that appear very reasonable given the mission and the terrain. Preliminary criticism of the simulated behavior by experienced helicopter pilots and analysts has been favorable.

HAVDEM Features

Algorithms for HAVDEM physical models are taken from HELIPAC in accordance with customer recommendations. This includes weapons and sensor models. The flight dynamics mathematics models include high quality 6-DOF representations used in virtual reality simulations and efficient medium fidelity 6-DOF kinematics algorithms. The value-driven pilot decision logic methodology is similar, but not identical, to TAC BRAWLER. Several algorithms involving access to the terrain data base and pilot utilization of terrain are unique to HAVDEM.

Physical Models

The flight dynamics are calculated from one of four available 6-DOF point mass models. Two of them, a baseline and a modified version, are taken from the AH-64A Apache math model used in a virtual reality simulation at the Ames Research Center, Moffett Field, CA. The other two are medium fidelity kinematics models. The first uses linear angular acceleration responses to the cyclic and pedal controls and a quadratic thrust response to the collective that peaks at the user specified "bucket" speed. The second calculates forces and moments from perturbations about trim flight conditions and control settings. A control flow diagram of the flight dynamics appears in Figure 1. Desired velocity pilot commands

Copyright © 1994 by ICAS and AIAA. All rights reserved.

(climb angle, heading, speed) are first converted into desired velocity vector rates (climb angle rate, heading rate, and longitudinal acceleration). The desired accelerations account for limits on power output in the current flight regime, angular rate limits, and the difference between the current and commanded velocity vectors. The desired translational accelerations are converted into desired angular accelerations and vertical force in the body frame of reference. Note that an optionally explicit command of body orientation is also part of this model. The desired accelerations are converted into control inputs using control response information obtained from the equations of motion. The control inputs are fed to the 6-DOF equations of motion to produce the resulting average thrust and angular accelerations. These values used to calculate the new position, velocity and orientation.

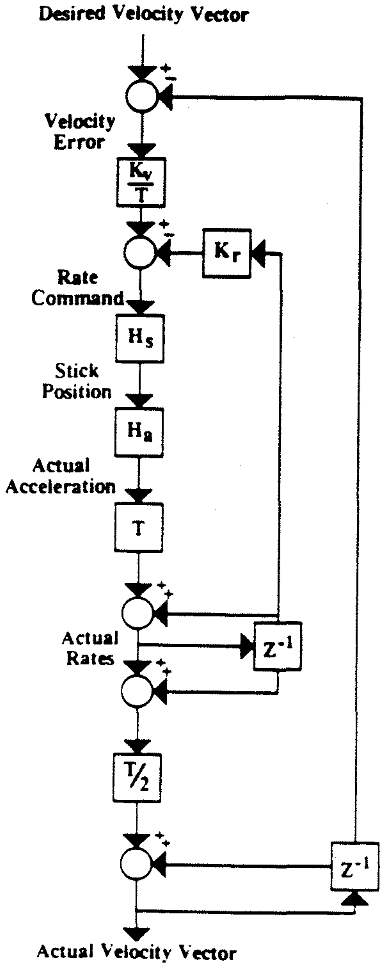

K_v = Velocity Error Gain
T_r = Integration Time Interval
K_r = Rate Gain
H_s = Rates to Stick Position Transfer Function
H_a = Stick Position to Accelerations Transfer Function

Z^{-1} = Discrete Time Delay

Figure 1. HAVDEM Flight Dynamics

The weapons in HAVDEM, like HELIPAC, are guns and missiles. The guns score a kill as a function of the off-boresight angle, the dispersion pattern of the bullets, the range, and the time on target. Missile models allow active radar, semi-active and IR seekers. These selections reflect HELIPAC's origins as a simulation of fixed wing air combat. The radar seekers are subject to field-of-view (FOV) restrictions and clutter. Guidance laws for the missiles include pure pursuit, proportional navigation, and beam guidance. Missile flight dynamics is a 3-DOF model that models translational but not rotational acceleration. The fuzing model calculates a probability of kill as a function of the fragment spray pattern, fuzing geometry, and circular error probable (CEP). Fire control constraints are imposed on the deployment of weapons. An example of a minimal constraint is the requirement of detection of the target by a specific sensor before weapon deployment is permitted. In general, detections by a combination of sensors are required. More than one sensor combination can apply to the same weapon. Fire control constraints also include G, angular rate, and track duration limits.

The sensors in HAVDEM are similar to those in HELIPAC. They include radar, IR, optical, passive radar receiver, and identification systems (IFF). Helicopters and missiles have tables of cross section values for radar, optical, IR, and fuzing. These tables are indexed by the azimuth and elevation of the target body axis with respect to the observer. Detection ranges for all sensors are specified in the input data assuming a standard target cross section. The radar may be scan only, single-target-track, or track-while-scan. Radar detection can be suppressed by main lobe or side lobe ground clutter, although the clutter calculations assume a flat earth at sea level. The IR and optical sensors are vulnerable to being blinded by the sun. Optical sensors are used to model the pilot's eyes. The passive radar receiver detects all radar that is tracking the observer regardless of range. The IFF has a characteristic range and either fails to identify or identifies targets perfectly. Partial or erroneous identification is not part of the HAVDEM IFF model. All sensors require clear line of sight (LOS) between the observer and target for successful detection.

The terrain elevation data for HAVDEM can come from a variety of sources including the Defense Mapping Agency (DMA) and the data base for the ARTBASS ground combat simulation. The raw elevation data is extracted by a preprocessing program and reorganized to increase the efficiency of the LOS calculations. Preprocessing also adds the height of cultural features to the terrain elevation to produce the values used for intervisibility calculations and minimum flight altitude. The preprocessed elevation data enables the LOS module to check for interference at each crossing of the LOS and lines connecting adjacent grid points (Figure 2). A height-over-terrain (HOT) module gives the difference between the altitude of a given point and the elevation of the terrain directly beneath it.

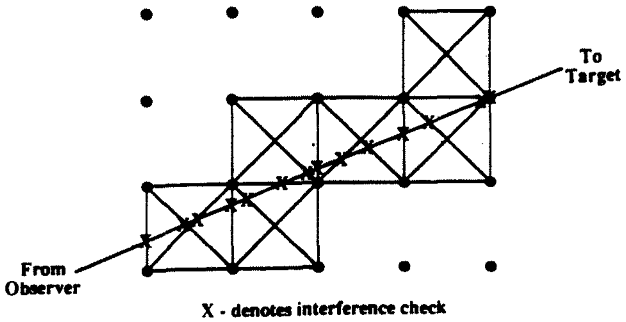

Figure 2. LOS Interference Checking

Pilot Decision Methodology

The pilot decision logic consists of a mental model data structure and modules for observation of sensor displays, situation assessment, weapon and target selection, weapon firing, and the maneuver decision (Figure 3). The mental model contains a description of the engagement status based on the information the pilot has received from his sensors. The discrete nature of computer simulation causes the continuous pilot activities to be modeled as periodic events in HAVDEM. The interval between consecutive "consciousness events" is normally one second. At the time of each consciousness event, all tasks that would have occurred during the preceding second are performed. Although sensor information is assimilated continuously over this interval, it is assumed that the pilot can make use of all the accumulated sensor information for making decisions. After viewing the sensor displays, the pilot makes decisions based on the contents of the mental model and not ground truth. The mental model is the appropriate reference since it represents the pilot's perception of the engagement. In general, it contains information similar to ground truth but its accuracy and completeness depends on the quality of the sensors and their displays. This structure allows for the simulation of surprise and confusion, since pilot reactions will be based on a perception of the engagement that might be incomplete or faulty.

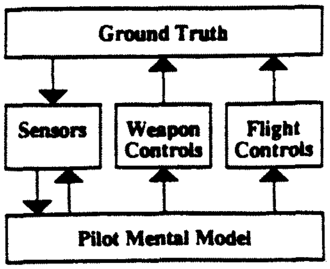

Figure 3. HAVDEM Information Transfer Architecture

Sensor display modules allow the pilot to assimilate information about the combat environment since direct access to ground truth information is prohibited. There is one display for each sensor. Note that the display for the optical sensor, used for the pilot's eyesight, represents the targets visible in a section of the sky. In the case of more than one sensor of a similar type on the helicopter, the displays are treated independently. The simulation assumes that the complete kinematic state of a displayed target is available to the pilot mental model. In the future, the display should give only a subset of the complete target state. Sensor modules are coded so that the addition of a Kalman filter estimation algorithm can be installed to maintain track banks based on stochastic target measurements.

A simple visual workload model prioritizes the sensor displays and schedules the most important ones for examination. The priority of each display depends on the completeness of the kinematic information, the time required for the pilot to read the display, the time since the pilot last looked at the display, and the number and importance of the tracks the pilot expects to see. The pilot examines displays in order of decreasing priority until the fraction of the consciousness event interval allocated for sensor search is used up. Normally two to four displays can be seen in one second. In the absence of any known targets, the displays tend to be searched in a fixed sequence with some preference given to the ones with higher quality and easy-to-read formats. As targets become known, the displays containing tracks on those targets take precedence, although the pilot may also favor a display which may imminently show a detection.

After the sensor information has been stored in the mental model, the situation assessment module classifies known aircraft and establishes the priorities of mission and tactical goals. The relative geometries of the known aircraft are considered with their identities to estimate the likelihood of the various pairs of helicopters attacking each other. A list of attractive targets is based on the conscious pilot's weapons envelope and the perceived threat of hostiles to other friendly platforms. A list of immediate threats is also formed. These lists allow detailed consideration of maneuvers to reference only the important platforms in the engagement, thereby increasing the efficiency of the simulation. Mission priorities and tactical goals can be greatly influenced by the presence or absence of hostiles on these lists. Certain tactical goals are important only if hostiles are present. The placement of hostiles and friendlies can also influence the relative priority of these goals. Situation assessment algorithms determine numerical values corresponding to each tactical goal in HAVDEM. These quantities are called "importance multipliers." The importance multipliers are used extensively in the maneuver decision module. The situation assessment module also evaluates the terrain and performs path planning. An inner module forms a coarse grid on the terrain and forms paths by connecting grid points. Paths are scored for altitude avoidance, concealment from the enemy, and time of transit. High-scoring paths are constructed between waypoints and also for attacking and evading important hostiles. These paths are considered during maneuver selection.

Target selection is performed in conjunction with the choice of weapon. Weapon and target combinations are compared so that eventual weapons employment will produce maximum expected value. Each available weapon is paired with each potential target. The utility for each pair depends on the target location in the weapon envelope, the exposure time required to use the weapon, and the mission value of killing the target. For a given target, a semi-active missile, for example, requires more exposure time to deploy than an IR missile. This would encourage selection of the IR. The range of the semi-active missile might be much greater than that of the IR missile, however. These two factors would compel the pilot to choose the semi-active for long range shots and the IR for close shots. The exact target range estimate at which the selections change is derived from the pilot's knowledge of missile specifications, and is not explicitly considered by the selection algorithm. This is an example of how value-driven decisions automatically adapt to the changing situation.

The selected weapon is fired if the target is sufficiently within the weapon envelope or if the target maneuvers to leave the envelope. A decision to sweeten the shot is permitted, but only under rather ideal conditions. This is because of the potential for the engagement situation to rapidly change as compared to fixed-wing combat. New enemies can emerge from behind the hills at any moment and ruin an overly prepared firing opportunity.

The maneuver decision selects the velocity vector that is perceived to lead the airframe into a position and orientation that best satisfies mission and tactical goals. The maneuver decision consists of three components: candidate maneuver selection, projection, and scoring. The selected maneuver is the one which results in the best scoring projected engagement geometry.

A discrete set of maneuvers is considered in the selection process. A maneuver is defined by a desired velocity vector. Most of the maneuvers in this set are formed by applying a standard positive or negative change in speed, climb angle, or heading to the current velocity. Note that because the maneuver decision is reconsidered each second, the maneuver set needs to reflect velocity vectors achievable in no more than one second. A sequence of these maneuvers may form the longer duration maneuvers that usually are referenced in training manuals, but it is not necessary that they do. The actual sequence will be determined by the needs of the engagement. This method opens up the possibility that the HAVDEM pilots can discover superior long-term maneuvers even for unusual situations. Other maneuvers in the set relate to particular goals that may require precision flying. Examples of these maneuvers are pointing at the selected target and waypoint path following.

The maneuver projection module calculates the future kinematic state of the engagement using the candidate maneuver and assuming that other known platforms do not maneuver. The state projection algorithm is simpler for this module than for the helicopter flyouts. The approximation needs to be good for only 5 seconds. The projection is based on the mental model estimates of the engagement geometry which should be expected to tolerate nominal errors. Limits on angular rates based on helicopter performance are applied to the projection of each helicopter. The projection time of 5 seconds is reduced for situations where the angular rate of change of the LOS to the target is high.

The projected state of the engagement is scored for its satisfaction of the tactical goals according to their relative importance. Associated with each goal is an importance multiplier and a value component. The score of a projected state is the sum of the product of the importance multiplier and the value component for each goal. The importance multipliers are calculated during situation assessment. The value components are defined on the interval [0,1] and represent the degree to which the projected state satisfies the given goal. Goals in HAVDEM include terrain collision avoidance, altitude avoidance, attack selected target, evade threatening hostiles, energy management, and waypoint following. The maneuver leading to the highest scoring projected state is selected. Typically, the best maneuver is among those that score very well against the most important survival goals such as terrain collision avoidance and hostile evasion. Further differentiation of high-scoring maneuvers comes from the satisfaction of mission objectives. Occasionally, two maneuvers can post high scores because they each score very well for separate goals such as waypoint following and pointing at the target. The higher scoring one is, of course, chosen. The value-driven method of selecting the best maneuver determines which course to pursue without the need for special consideration of why these different actions are similarly attractive.

Terrain Utilization Module

The terrain utilization algorithm that was developed in Phase I is called *bugrace*. This name was chosen because the algorithm can be visualized as a swarm of bugs moving between adjacent points on a grid and racing each other to the intended destination. If travel time is the only figure of merit for a path, the bug that reaches the destination first has defined the optimal path. An extended version of *bugrace* is used in HAVDEM to find the best route between waypoints and the best path along which to attack a target or evade a threat.

A demonstration program was delivered to the Army at the end of Phase I. It consisted of a small grid with values assigned to each point representing terrain elevation. Movement of an attacker and a target was constrained to allow step transitions between grid points only. The attacker used *bugrace* to plan the approach while the target was forced to maintain a constant velocity. A favorable attacking position was defined by range and off-boresight angle thresholds, and a clear LOS. Exposure to the enemy was defined as an attacking position for the target against the attacker. A stationary target was assumed to be able to point at the attacker. Paths were scored according to rewards for time an attacking position was maintained and penalties for time exposed to the target while not attacking and time used to set up the attacking position. The program forced an attack even at the risk of exposure. The demonstration showed that *bugrace* will direct the attacker to hide behind a hill, wait for the target to pass, and assume an attacking position behind the target. For a case where the target was stationary on the other side of a hill, the attacker first moved away from the target, the hill between them, to make room to accelerate. Then the attacker flew over the hill at maximum speed to attack the target. The motivation for the maneuver was to minimize the time exposed.

When *bugrace* is invoked in HAVDEM, a 7x7 square grid is generated from the terrain map (Figure 4). The conscious pilot is over one of the grid points. The grid spacing depends on the distance and time to the next waypoint or, in the case of the attacking and evasion modules, the current speed of the conscious pilot. Transitions between adjacent grid points approximate maneuver options. Each transition is scored according to exposure to the enemy, attacking opportunity, altitude avoidance, and time invested. If several "bugs" reach the same grid point at the same time, only the best scoring one is continued. Paths terminate at the waypoint or after 30 seconds flying time. The score for best path is saved for each of the distinct transitions from the initial conscious pilot location. In the maneuver decision scoring, the heading of each candidate maneuver is compared to the headings associated with these initial transitions. Because the candidate maneuver heading will fall between two of the *bugrace* transition headings, the maneuver scoring for the path-following component is a weighted average of the scores for each transition.

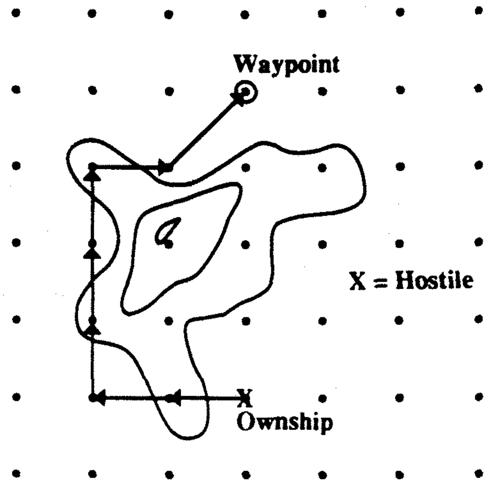

Figure 4. Path Planning

Graphics

HAVDEM engagements can be replayed using a color graphics display that includes a top view and cockpit views of the battle. These views show terrain contours, missile, SAM, and helicopter positions, and helicopter ground trails. The terrain contours are formed by assigning a color to a region of terrain elevation (Figure 5). The colors are discrete selections chosen from a green scale with darker green representing lower terrain. An almost continuous variation in green scale with elevation has been tried, but the resulting picture gave a blurred impression as if the features were all under water. A data window on the side of the terrain map gives the coordinates of each helicopter and missile at the current simulation time. Descriptive narrative accompanies important events such as target detection, missile launching and fuzing, and target identification. The helicopter icon is either a red or blue filled circle with an appropriately oriented tail boom. Missiles in flight are depicted as yellow squares. SAM sites are red or blue squares. A black line extends from the helicopter or missile icon to denote the direction of the velocity vector. The length of the velocity line is proportional to the speed. Pressing the <Enter> key advances the display forward one second. The response interval between displays is sufficiently fast that repeated pressing of the <Enter> key produces a rough animation of the engagement. An animation command is available to produce this effect automatically. A side view of the battle can optionally be displayed showing rotorcraft positions over a mesh rendering of the landscape. The user can return to the beginning of the engagement at any time during the display sequence. The graphics program interfaces with X-Windows library utility functions.

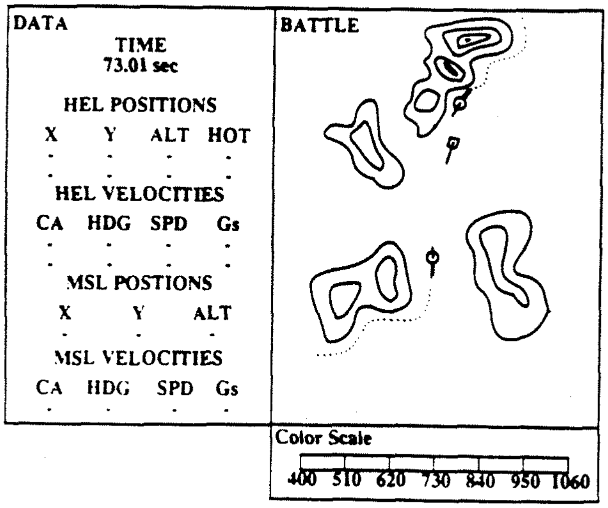

Figure 5 Diagram of Graphics Display

HAVDEM Usage

HAVDEM has been used in the evaluation of the air-to-air effectiveness of helicopter configurations. While the final results of this study are undergoing the government review at the time of writing, some of the experimental methods can be discussed here. The study included many one-versus-one (1v1) engagements and the traversal of about a 15Km course, at various altitudes and speeds, populated by four threat SAM sites.

The 1v1 engagements were conducted over rolling hills with elevations ranging from 500ft to 950ft. The initial speed of each was 60 knots and the rotorcraft were positioned along an east-west line. Four different initial headings were used: 1) both heading north, 2) one north and the other south, 3) one east and the other west nose-to-nose, and 4) one west and the other east tail-to-tail. Two different initial separations were used: 3000m and 1000m. Three different initial altitudes were used: 6000ft, 1500ft, and 950ft. These initial variations give a total of 72 different scenarios. The 950ft scenarios were designed so that the adversaries could initially detect each other and easily interpose terrain between them. The ensuing engagement involved terrain utilization tactics. Of particular interest in these engagements were the weapon employment timelines as compared to the maneuvering timelines. Excessive use of uncoordinated flight to bring weapons to bear minimized the effects of maneuvering. This behavior caused weapon employment timelines to be the determining factors in the engagement outcome.

The set of engagements requiring a single rotorcraft to cross an area populated by four enemy SAM sites employs a different terrain utilization profile for each engagement, but uses the same set of waypoints. The terrain utilization profiles used are low level, contour, and nap-of-the-earth (NOE). The Apache and the modified system are run through this course. The simulated pilots are not given knowledge of the threat laydown and must react to their presence if they are detected. As expected, evasive tactics were more effective at slower speeds for both types of aircraft.

Statistical Significance

"Conventional wisdom" in constructive simulation is that 30 runs provides a statistically significant sample of the distribution of relevant data. This wisdom is likely to be based on the common formula for the confidence interval of the mean of a normal distribution with unknown variance. In this formula, the "student's t" distribution term can be approximated by a constant value after 30 samples are collected when calculating the 95% confidence interval. One problem with this approach is that while the 95% confidence interval can be calculated, there is no guarantee that the width of the interval is acceptable to the experimental design. A correction to this approach would be to continue sampling until the 95% confidence interval width is below the value established in the experimental design. The other problem with the conventional procedure is that the distribution of data is likely to be something other than normal. In that case, the common formula for calculating confidence intervals is invalid. In fact, the distribution of data for many important experiments is not known beforehand and textbook formulas for calculating the confidence intervals are either obscure or unobtainable.

HAVDEM analysis did not assume any knowledge of the distribution of data contributing to the results of the study. Instead, a procedure was adopted that used the sample data as an estimate of the underlying distribution. This estimate of the distribution is used to generate a distribution of the mean values one constructs by averaging randomly selected samples from the estimated distribution. If the mean values are sorted, the central 95% of the values define an estimate of the 95% confidence interval. One needs to construct a large number (the study used 200) of mean values to get a reliable estimate. As the number of measured samples increases, the constructed mean values move closer together and the estimated confidence interval decreases. HAVDEM continued to generate data until the estimated confidence interval decreased below a specified threshold. For simple symmetric scenarios HAVDEM needed as few as 50 repetitions to meet the specified 95% confidence interval threshold. For scenarios involving complex terrain, 500 to 600 repetitions were often needed. One scenario required 1600 repetitions. Each repetition generated a single data measurement. In other studies, such as sensor and weapon effectiveness, many data measurements could be performed in a single repetition. This experimental design can greatly reduce the required number of repetitions.

The following discussion describes the algorithm that determines if a statistically significant number of HAVDEM runs has been performed for a given scenario. In general, this determination depends on the number and distribution of the data samples used to compute the stopping condition statistic (SCS). If many data samples are available in a single run, a relatively few number of runs needs to be performed. The number of runs also depends on the desired confidence interval of the SCS. The number of runs increases as the interval of a given percent confidence decreases. The general stopping condition used in this study is to perform runs until a given percent confidence interval of the SCS is less than a predetermined value.

HAVDEM runs of a given scenario differ in the assignment of the initial random number seed. The parts of HAVDEM sensitive to the random number sequence are sensor detections and the relative timing of sensor and pilot consciousness events. Sensor models ordinarily calculate probabilities of detection for the current geometry. A random number between 0 and 1 is compared to the probability of detection to determine if the detection is successful. A change in the random number seed, and therefore the rest of the sequence, produces a different set of successful detections on which the pilots make decisions and form situation awareness. This causes surprisingly large changes in the course of the engagement. The random number sequence also determines the time within the first second of the engagement for execution of the first event in the event streams of each sensor and pilot decision entity. Engagement behavior is very sensitive to the relative timing of decision events because a HAVDEM pilot reacts to the motions of all other known aircraft. Simply exchanging the order of the pilots' first consciousness events changes who reacts to whom and can produce a completely different set of maneuvers. Change in the random number seed can therefore change the engagement maneuvering in a very chaotic way.

The SCS for the 1v1 scenarios is the difference in the average cumulative probabilities of survival (dAVG_CPS). The convention is to subtract the average CPS of aircraft 2 from that of aircraft 1. A value of zero for dAVG_CPS means that battle is equally lethal for both aircraft. The stopping condition is reached if the 95% confidence interval of the dAVG_CPS statistic is less than 0.1 for the entire duration of the engagement. This means that there is a 0.95 probability that the dAVG_CPS value calculated from the samples is within 0.05 of the dAVG_CPS value of the distribution from which the samples are taken. Note that only one sample of this statistic can be produced from a single HAVDEM run assuming that weapon fire has been exchanged. In some of the runs in the NOE scenarios, the aircraft lose contact among the hills and weapons are never used. These runs contribute nothing the study of the advantages of **relative** aerodynamic performance and are discarded. This effect, along with the chaotic nature of air combat, means that hundreds, sometimes over a thousand, runs are required to reach the stopping condition. To speed up the runs, the engagements are terminated if the CPS of one of the adversaries falls below 0.05. Below this value, the weapons can inflict only negligible damage and the effects of hostile weapons have a small absolute effect on CPS.

The SCS for the rally course scenarios is simply the average cumulative probability of survival (AVG_CPS). The stopping condition is reached if the 95% confidence interval of the AVG_CPS statistic is less than 0.1 for the entire duration of the engagement. This means that there is a 0.95 probability that the AVG_CPS value calculated for the samples is within 0.05 of the AVG_CPS value of the distribution from which the samples were taken. Only one sample of this statistic can be taken from a single HAVDEM run. Runs where the SAM sites do not fire on the aircraft are not discarded because they contain very significant information about the aircraft's **absolute** ability to evade threats. The chaotic nature of path selection still causes from near one hundred to several hundred runs to be performed before the stopping condition has been satisfied. Rally course scenarios also terminate if the CPS falls below 0.05.

For all scenarios, the appropriate stopping condition algorithm calculates the 95% confidence interval at all times in the engagement and terminates the run generation process if all the intervals are less than 0.1. The difficult part of this procedure is finding all of the 95% confidence intervals. There are formulas for finding the interval if the underlying distribution is known. In the case of these studies, however, the distribution is not known. One method for determining the 95% confidence interval of the SCS for a set number of samples n from an unknown distribution is to repeatedly take n samples and calculate the SCS for each set. If enough sample sets are taken, the resulting SCS values (one for each set) can be sorted and the 95% confidence interval is given by the 2.5th and 97.5th percentile values. For example, if 200 sets of n samples are measured, then the 95% confidence interval is given by the 6th and the 195th SCS values of the sorted list of SCS values for each set. Generating confidence intervals in this way is very expensive since a HAVDEM run is required for each sample in each set.

Because the underlying distribution of the SCS is not known, we must rely on the samples themselves for any information about the distribution. The samples each have equal (infinitesimal) probability of being measured. If the probabilities of the samples are scaled so that their total is 1, they form a probability distribution of a discrete random variable. As the number of samples increases without limit, the discrete distribution approaches its corresponding continuous distribution. Similarly, the average of the samples approaches the mean of the continuous distribution. For a small number of samples, 30 for example, there is a 0.9 probability that there is at least one sample from 9 of the 10 one-tenth percentiles of the continuous distribution. This implies that the discrete distribution for more than 30 samples can be used as an approximation to the continuous distribution for calculating the confidence interval of the SCS since samples from the extreme percentiles of the distribution are likely to be represented. This approximation gets better as the number of runs increases.

We have stated that the samples that have been measured are used as a discrete random variable approximation to the actual continuous distribution. If there are n measurements, each has a probability of 1/n. A calculation of the SCS can be formed by selecting n independent measurements on the approximated discrete distribution and combining them appropriately. For the 1v1 runs this involves independently selecting one run from the n runs performed so far and repeating the selection n times. This set of n runs, some of them duplicates, is used to calculate a value for dAVG_CPS. The selection of n runs and calculations of dAVG_CPS is repeated 200 times. Note that this procedure will give a different value of dAVG_CPS for each set. The 6th and 195th values of the sorted list of dAVG_CPS values gives the 95% confidence interval. If this interval is less than 0.1, runs for the 1v1 scenario are terminated. A similar technique is used for the rally course scenarios. We have found that hundreds of samples (runs) are usually required to reach a 95% confidence level of less than 0.1.

Cost Savings

The cost of using constructive simulation is significantly less than that of virtual reality. The cost reduction is directly associated with the reduced requirement to present a highly detailed visual scene for the human operator. Motion inputs to the operator are also unnecessary. This cost savings must always be weighed against the greater realism a human operator can provide. The benefits of constructive simulation come from thorough and repeatable experiments which either don't require the realism of a human operator, or guide the application of the more expensive research methods.

Motion based virtual reality studies can be an order of magnitude more expensive than constructive simulation experiments. For some kinds of research, this ratio can be much higher. One contributor to this cost savings is the reduced computer processing requirements. HAVDEM, for example, needs only a work station to perform a detailed study. Visual based virtual reality requires several work stations (or a reality engine) and support for real-time visual displays. Motion based virtual reality requires complex machinery in addition to real-time visual displays. Additional cost savings is realized by the ability of constructive simulation to operate using approximated, and therefore less expensive, models for physical processes like sensors, weapons, and flight dynamics. For virtual reality, such models must have sufficient fidelity and complexity to suspend the disbelief of the human operators so that they behave realistically. Such fidelity costs more. A larger equipment and complexity requirement also demands a block of engineering support personnel that are not needed by a constructive simulation. Finally, the coordination of this infrastructure supporting virtual reality requires a greater amount of time devoted by each member of this larger support staff.

While it costs less to use, constructive simulation does not replace virtual reality in addressing certain human factors issues. In a recent HAVDEM study, the constructive simulation gave no indication of the magnitude of the handling qualities issues that were evident in corresponding virtual reality tests. Only a human operator can completely identify the aspects of the handling qualities that contribute to or detract from a mission. In general, constructive simulation is useful in identifying inferior or promising concepts independent of human factors. These results can be used to conduct virtual reality investigations either at reduced cost or more thoroughly.

One last comparison between constructive simulation and virtual reality is that the involvement of a human pilot in an experiment precludes experimental reproducibility. A human never reacts to stimulus exactly the same way as in a previous experience. This means that in an attempt to reproduce events that occurred in a previous experiment, the human pilot will fly differently and cause a divergence of events from the desired sequence. It could require several trials, and additional expense, to recreate a similar situation. Constructive simulations are more conducive to research investigations because their experiments are exactly reproducible. The researcher can select an interesting interaction in the mission and immediately replay it while collecting high resolution data at the point of interest.

Future Development

Long-term development of HAVDEM includes cooperative pilot behavior, countermeasures, advanced aerodynamic concepts, and air-to-ground engagements. Value-driven decision logic is especially well suited to the simulation of cooperative behavior because the decisions of the individuals can be based on group goals as well as individual goals. Individual decisions can also influence the importance of specific group goals. A designated flight leader, for example, can make posture decisions for the entire flight and make target assignments to the flight members. This is accomplished by allowing the flight leader's decision subroutine to alter the values of other flight members' importance multipliers for goals relevant to posture selection and target selection. The flight members will continue to consider individual survival goals, so it is possible for these considerations to temporarily override their orders.

Current efforts are improving the fidelity of sensor modules so that the presence of jamming, chaff, or flares will degrade their performance. Pilot decision logic improvements determining the efficient employment of countermeasures will accompany the introduction of these new features to HAVDEM. The flight dynamics module can be replaced with other software that meets the same interface specifications. Such software can model the response of advanced airframes such as tilt rotor and vectored thrust vehicles. Models for artillery and air-to-ground missiles can be incorporated into HAVDEM to provide a comprehensive combat environment for the helicopter. The current terrain implementation and the terrain utilization decision logic would be easily adaptable for these important aspects of helicopter missions.

Conclusions

The development of HAVDEM through the SBIR program has produced a constructive simulation that supports the analysis of air-to-air helicopter combat that is being extended to support attack missions. HAVDEM contains models for physical systems and pilot behavior that are of sufficient fidelity to produce realistic engagements without sacrificing efficiency. Recent experience indicates that hundreds of samples, not just 30, are needed to produce a result for which a high percentage confidence interval is small. This fact motivates the researcher to design experiments where several data samples, instead of only one, are generated by each simulated mission.

The cost of using constructive simulation is significantly less than the cost of virtual reality. Favorable cost comparisons include the computer hardware, visual scene software, human factors support, avionics simulation, engineering maintenance support, engineering operation support, and pilot time. Constructive simulation does not address certain human factors issues, but it is useful in identifying inferior or promising concepts independent of human factors. Constructive simulations are valuable to research investigations because their experiments are reproducible.

References

1. Bent, N., "The Helicopter Air-to-air Value-Driven Engagement Model (HAVDEM)", Military Computing Conference Proceedings, May 1987.

2. Dloogatch, M., HELIPAC Analyst's Manual, Sikorsky Aircraft Division, United Technologies Corporation, October 1990.

3. Bent, N., Helicopter Air-to-air Value-Driven Engagement Model (HAVDEM) Analyst's Manual, Decision-Science Applications Report No. 58/1253/R1, September 1993.

4. Bent, N., Air-to-air Combat Engagement Analysis Utilizing Terrain Shielding Phase III Interim Report, Decision-Science Applications Report No. 58/1470, March 1994.

5. Bowker, A., Lieberman, G., Engineering Statistics, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1972.