Proceedings of the 2004 Winter Simulation Conference
R .G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, eds.

ASSESSING OBSTACLE LOCATION ACCURACY IN THE REMUS UNMANNED UNDERWATER VEHICLE

Timothy E. Allen

Submarine Officer Advanced Course
Naval Submarine School
Groton, CT 06349, U.S.A.

Arnold H. Buss
Susan M. Sanchez

Naval Postgraduate School
Monterey, CA 93943, U.S.A.

ABSTRACT

Navy personnel use the REMUS unmanned underwater vehicle to search for submerged objects. Navigation inaccuracies lead to errors in predicting the location of objects and thus increase post-mission search times for explosive ordnance disposal teams. This paper explores components of navigation inaccuracy using discrete event simulation to model the vehicle's navigation system and operational performance. The simulation generates data used, in turn, to build statistical models of the probability of detection, the mean location offset given that detection occurs, and the location error distribution. Together, these three models enable operators to explore the impact of various inputs prior to programming the vehicle, thus allowing them to choose combinations of vehicle parameters that reduce the offset error between the reported and actual locations.

1 INTRODUCTION

Naval vessels have always been susceptible to mine warfare. Whether the threat is real or perceived, the end result is the same: mine warfare disrupts the ability to project and maintain forces away from home waters. The technology involved in the construction and employment of mines has not changed appreciably in the past few decades. However, the United States Navy has placed a high priority on developing technology and tactics designed to counter the threat of mine warfare. Since 1988, three U. S. warships have encountered mines in the Persian Gulf. Mine warfare is affordable and available to any country willing to invest the time required to lay the mine-field.

In March 2003, during the second Gulf War, Navy ships intercepted three tug boats in the vicinity of the Iraqi port of Umm Qasar loaded with over 130 mines intended for the harbor inner reaches (Eisman 2003). This incident underscores the relative ease with which a country could employ this rather primitive, yet effective, tactic.

Mineman (U.S. Navy 1994) lists the following advantages that mines have over other conventional weapons:

- Mines lie in wait for the enemy with no reasonable threat of counter-detection;
- Mines may win a conflict passively by causing the enemy to alter tactics;
- Mines may force ships to travel longer, less reliable routes to deliver troops and materials;
- Mines are cost effective due to their relatively primitive technology.

Our research focus is the accuracy with which the vehicle estimates an object's location. Operators are concerned with the area that may have to be searched to clear identified mines. Greater uncertainty results in more time for divers to clear obstacles because of the increased search effort required. The time allotted for a search effort is limited by several factors, including bottom time, the amount of air carried, and diver fatigue, so using search area to estimate the time required to complete a mission is important. The insight gained can help operators choose the best configuration for the vehicle prior to putting it into the water. Decreasing the location time by shrinking the search area would save time, money and personnel.

Many components of the navigation system contribute to uncertainty about the vehicle's true location and thus the location of the contact. Deciding which components affect the error most severely, and which might be influenced the most by human operators, is our primary interest. Time and cost considerations prohibit gathering data from actual vehicle runs. It is also difficult to isolate individual components of the navigation system in an ocean environment. Instead, we use simulation to explore how individual inputs contribute to position prediction error.

We seek to provide insight into three critical areas of vehicle operation. First, what operating conditions provide the largest probability of detection of a mine-like object? Second, given that detection has occurred, how can the predicted mean offset location error be minimized? And third, once a predicted mean error is determined, how does the operator establish a probability that the mine is actually inside of a desired range? Insights into these three questions are provided by building models that predict the

probability of detection and the mean location offset given that detection occurs, and using the offset error distribution to establish a probability that the mine is inside a desired range.

2 THE REMUS VEHICLE

The Remote Environmental Measuring UnitS (REMUS) vehicle (Figure 1) is built by Hydroid Technologies and is designed to collect hydrographic data in relatively shallow water. The overall system is comprised of the vehicle, various auxiliary equipment necessary to support its mission, and software programs designed to conduct pre-mission planning and post-mission data analysis. During a mission the vehicle collects side-scan images that can be viewed, post-mission, and used to identify and locate objects (such as mines) on the ocean floor. Other data are collected and saved that are important in assessing the accuracy of the ostensible location of the object. Theoretically, returning to an object is an easy task for the explosive ordnance disposal (EOD) teams because the vehicle always "knows where it is," but in reality the task is made more difficult due to inaccuracies in the navigation system that result in errors in the reported location of the objects.

Figure 1: The REMUS in its Transportation Container

The vehicle is 7.5 inches in diameter, 40 inches in length and weighs about 70 lbs. It is equipped with internal batteries and is capable of diving to depths up to 100 meters for as long as 22 hours on one charge. It has side scan sonar for its search function. Navigation is accomplished primarily by the vehicle ranging with in-stratum transponders. Secondary navigation is done by dead reckoning with an internal compass and Doppler velocity log. The navigation system is discussed in more detail in Section 2.1. Onboard computers store mission-essential parameters and collect data for post-mission analysis, including salinity, temperature, depth, optical backscatter, side-scan sonar images, and vehicle location information.

2.1 REMUS Navigation Methods

The REMUS vehicle operates autonomously while performing its mission and thus navigation error is an important consideration. Understanding the vehicle's navigation methods and the potential sources of error provides insight into the output from mission playback and assists operators in interpreting collected data.

REMUS may switch between two navigation methods during the course of a field survey: dead reckoning and long baseline. (Ultra-Short Baseline is used when the vehicle is finished with its primary mission and preparing for recovery.) Mission design parameters may dictate that one method is preferred, e.g., due to the vehicle's proximity to the transponders or by the ocean bottom composition. We now describe the methods.

If an acoustic fix is not available, the vehicle navigates by *dead reckoning* (DR). DR position is computed using onboard speed and heading information. Speed is calculated using propeller rpm input and Doppler acoustic signals. The Doppler acoustic signal may only be used if the vehicle is within 20 meters of the sea floor. Heading is computed using an onboard magnetic compass. DR position is computed by determining how far and in what direction the vehicle has traveled since the last update. This distance traveled is added to the last known position and the new position is updated in the computer. Precise DR navigation relies on accurate speed and heading information.

With *long baseline* (LBL), the most common method of navigation, REMUS interrogates pre-positioned transponders and receives a return signal. The vehicle then triangulates its position based upon the length of time the signal takes to make the round trip. Many variables affect the speed of sound in water, so REMUS collects real time data on water temperature, salinity and depth to calculate the speed of sound and then calculates a distance based upon the time. Position is then fixed based on the triangulated range. LBL requires a minimum of two transponders. If a signal is not received from either one of the two transponders then the fix is assessed as "bad" and no position updates occur. The vehicle continues to navigate by DR between good fixes. The vehicle may be set to interrogate at specified time intervals. The time interval affects the maximum range because a signal must depart and return prior to the next interrogation for the range to be accurate. The maximum range for the transponders is specified at 1500 meters, but ranges as long as 1700 meters have been observed to work in favorable water conditions.

Triangulation of position is achieved when the lengths of the three sides of a triangle are known (Stewart 1999). As Figure 2 shows, since the distance between transducers (C) is known, the REMUS can determine its position once it calculates its distance from each transducer (A and B). REMUS may switch to DR if the fixes it receives are perceived to be bad.

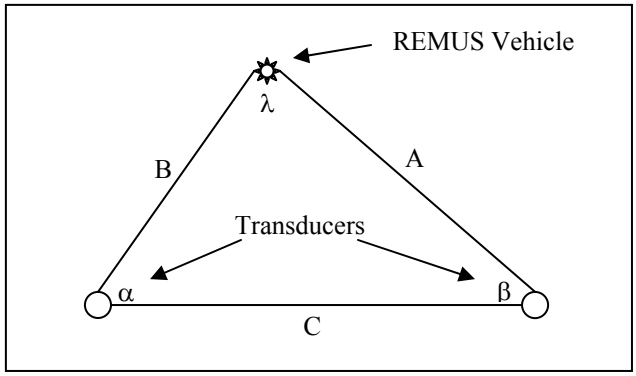

Figure 2: Triangulation of Position

2.2 Sources of Error

There are five sources of navigation error in the REMUS vehicle: compass error, effects of current, errors in transducer placement, errors due to motion of the transducers after placement, and errors due to the uncertainty of the vehicle's attitude relative to the transducers.

Compass error is important because it is the primary tool utilized during DR navigation. REMUS combines an ordinary magnetic compass that is fairly accurate over long distances with yaw-rate sensors that are more accurate over short distances. This produces a heading input with more accuracy than either would have produced individually. The manufacturer claims that heading errors are reduced to ± 0.1%, or about 0.36 degrees. This claim is difficult to verify and not supported by data collected in support of this work (Allen 2004); in the short term, heading errors can reach 3% or up to 10.8 degrees to either side.

Current effects are another source of error. REMUS operates primarily in ocean environments where the current can push it off track or cause it to inaccurately assess its position on the planned track. The net result is error in the prediction of an object's location. When the vehicle approaches from the same orientation on consecutive runs over a short period of time the error is predictable because the current does not change direction or strength appreciably. Current is not constant, however, nor is it easily determined at a specific location for use by the vehicle. Currents are often reported in a location in the vicinity of some important navigation aid, but generally not in specific locations as might be necessary in order to apply correction to the vehicle ahead of time, thus the vehicle's position is subject to uncertainty due to the effects of current.

Of all the variables, the operator has the most control over *transducer placement*. The transducer is transported to the drop area and positioned at the surface using GPS, then dropped and allowed to sink to the bottom. It is held in place by an anchor attached to a neutrally buoyant line. The transducer is allowed to float in the same stratum as the vehicle. The transducer's position is programmed into the onboard computer and subsequently used to triangulate REMUS' position. Generally the position is programmed into the vehicle prior to making the drop, but there is no restriction that would prohibit inputting the position after the drop occurs. Assuming that once the device is dropped overboard it neither moves nor drags its anchor on the bottom, a one-time placement error occurs if the operator is not at the exact position for the drop. The accuracy of a handheld GPS unit is quantifiable and contributes to the transducer misplacement. Placement errors can also occur because the anchor and transducer do not drop straight down to the bottom—current and hydrodynamic effects acting on the device are unpredictable as it descends.

Transducer motion is also a concern Once a transducer is placed on the ocean floor, it is affected by current and wave action surges. Under best case conditions, a small displacement of the transducer results in a small error in position fix. However, errors are magnified when the vehicle operates with a small angle between the baseline (an imaginary line joining the two transducers) and the vehicle's track. Current has the effect of displacing the transducer from a vertical position off to one side with a fairly constant magnitude. Wave action tends to displace the transducer from side to side in a repeating manner. The overall effect is that the transducer is rarely in the programmed position that the vehicle uses for calculations. In the long run the alternating behavior mitigates wave effects, but in the short-term there can be substantial error.

The final consideration for position uncertainty involves the *vehicle attitude* relative to the location of the transducers. Consider first the distance from the vehicle to the transducer. As the distance increases the likelihood that the transponder "hears" the interrogation of REMUS or that REMUS "hears" the return ping degrades, because the intensity of sound decreases as range increases. The second consideration is the vehicle's relative aspect to the transponders. The location and construction of the sensors in the vehicle's nose mean it is more likely to receive a signal when it is pointing toward the transponder than when it is traveling away from the transponder. Recall that a good fix is obtained only if a response is received from both transponders. The longer the vehicle navigates by DR without a good fix, the higher the position uncertainty.

3 SIMULATION MODEL

We briefly describe the REMUS model developed to assess navigational accuracy. For more detail, see Allen (2004). The REMUS model utilizes Event Graph methodology (Schruben 1983) and the LEGO component framework (Buss and Sanchez 2002) to design the simulation. The REMUS model was implemented in Simkit, a Java-based API for implementing DES models using the LEGO framework (Buss 2002).

The components making up the REMUS model represent the major pieces of the navigation system plus a few objects that "listen" for prompts so that they may collect data and conduct housekeeping functions. The full functionality of the model is not separated into classes based upon their use by the vehicle. Initially, we build a simple model which incorporates only factors affecting DR, such as compass error and current effects. A more complex model incorporates factors related to the LBL navigation method. Finally, the complete model is exercised in a typical mission involving sweeping a set area for a randomly placed mine. Later models incorporate the earlier components and add functionality. Names of components are used to reference the name of the JAVA™ class in the model (e.g., RemusMover is a class name in JAVA™).

3.1 Dead Reckoning Model

Dead reckoning is the default navigation method for the vehicle when it is not receiving good fixes, and thus plays a part in every version of the model. The event graph displayed in Figure 1 represents the model used to analyze the vehicle's navigation system while dead reckoning. The DR version models vehicle movement as uniform linear motion using the RemusMover class. A CookieCutterSensor models the vehicle's side-scan sonar capability, and data are generated using the RangeFinder component.

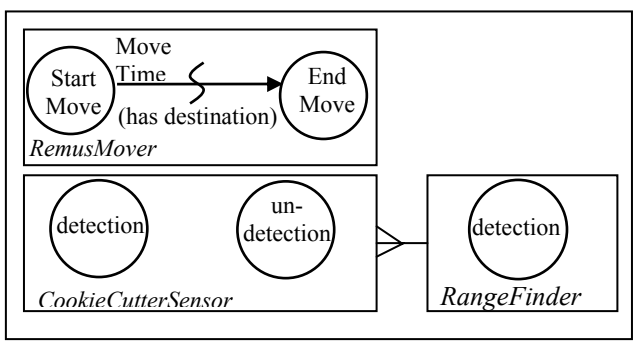

Figure 3: Event Graph for Dead Reckoning Model

The RemusMover component monitors the vehicle's perceived position vs. ground truth while detecting obstacles. As the vehicle transits it keeps track of its perceived position in the operating space, which may differ from the actual position. During mission playback the position information is accessible to the user. While collecting data from the model, it is useful to be able to ascertain the magnitude of the error. We accomplish this by keeping track of the vehicle's actual position and its perceived location.

Movement is initiated by the StartMove event in the RemusMover class. All movement in this model is uniform linear motion, so the vehicle is assumed to travel from point to point in a straight line. Recall that in DES, changes in state (e.g., position, etc.) are performed when the time reaches the next scheduled event. This presents a modest modeling challenge in that the REMUS' position is constantly changing as it travels, but this challenge is overcome by maintaining the positions ("real" and "perceived") implicitly rather than explicitly. The RemusMover maintains three state variables—position when movement began, time movement started, and velocity—which can be used to calculate its position at any time between events.

The RemusMover component executes orders to move from one point to another, and only the most recently departed point and destination point are stored. Functionality is added by using Simkit's PathMoverManager class to coordinate the vehicle's path. Waypoints are passed to the manager in the order in which the vehicle is expected to transit, allowing complicated paths that reflect desired behaviors for REMUS. The PathMoverManager listens for an EndMove event from the RemusMover, and then issues orders to the RemusMover to move. Movement occurs if the order is accompanied by an actual destination.

Compass error is passed to the RemusMover via the test program as a maximum percent error. Each time the mover is ordered to move, the compass error is calculated and the real velocity of the vehicle is adjusted. Ideal velocity remains unchanged. The effect of current is implemented through an adjustment to the real velocity while leaving ideal velocity unchanged. The current adjustment is passed to the RemusMover class via the test class in the form of a vector consisting of x and y components, altered by adding some randomly generated noise.

The sensor for the model is assumed to be a simple cookie cutter sensor, meaning that as soon as the obstacle is inside the range of the sensor detection occurs with probability one. The detection range is adjustable. The detection and undetection events from CookieCutterSensor in Figure 3 are scheduled by the RemusSensorMediator. Detection events are scheduled when the mover reaches the closest point of approach with the target; undetection events are scheduled to occur when the mover exits the sensor's detection range. This behavior models the way operators locate obstacles in the side-scan images during post-mission analysis. A property change is fired when a detection or undetection event is processed on the event list, and can be used to gather statistics on the number of detections or times of detection, for example.

Finally, the RangeFinder class gathers data on the error in target location by listening for a detection event in the CookieCutterSensor. The RangeFinder detection event calculates the offset in predicted vs. actual target location and outputs the data to a user-designated text file.

3.2 Transducer Model

The preferred navigation method utilizes transponders for triangulation of position with DR between fixes. This reduces the error in predicting obstacle location, but errors still exist due to transducer placement errors and move-

ment. Figure 4 shows the event graph. Dashed lines joining two events represent interrupt actions that remove the event on the head of the edge from the event list.

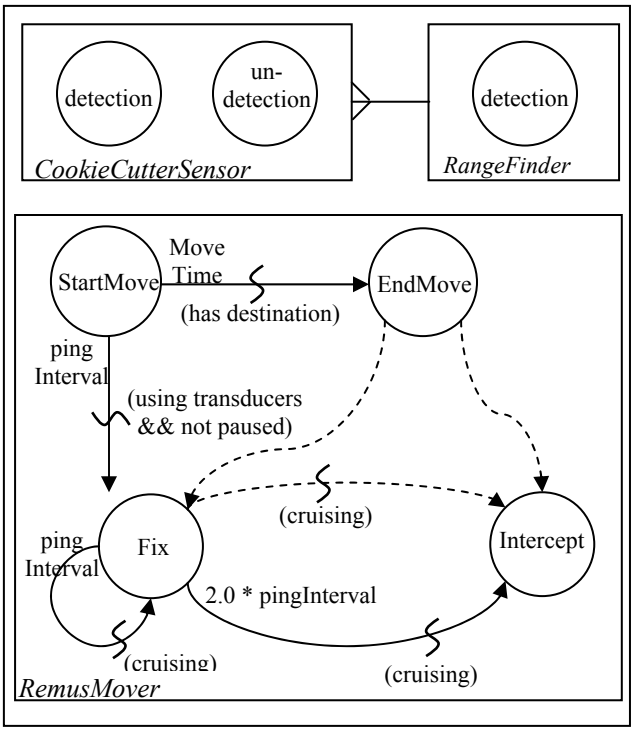

Figure 4: Event Graph for Transducer Model

Many details of the RemusMover component remain unchanged, but position fixes using the transducers are implemented. The StartMove event schedules the first Fix event if the vehicle is not paused and transducers are used for navigation. The Fix event is scheduled after a delay determined by the ping interval variable. A longer ping interval means more time elapses (and more error accumulates) between fixes. A shorter ping interval implies less error accumulates, but at the expense of more computing time. Additionally, a shorter ping interval means that the vehicle cannot operate as far away from the transducers because it would have too long a wait for the return signal. The StartMove event also schedules the EndMove event provided that the vehicle has an assigned destination.

The EndMove event represents the end of the vehicle's current leg. The vehicle either will stop or will be assigned a new destination. The execution of the EndMove event interrupts any scheduled Intercept or Fix event, but the vehicle fixes its position immediately after the turn.. In actuality the REMUS never stops attempting to fix its position, even while turning or maneuvering to the next waypoint, but our simplification is reasonable since the time to execute a turn in the simulation is essentially zero.

Position fixes are developed and implemented in the Fix event, which is initially scheduled by the StartMove event. The first logic gate determines whether a good signal is received from each of two transducers. As the distance between the vehicle and a transducer increases, the probability that a good signal is returned decreases (see Allen 2004). The Fix event implements this by comparing the range between the vehicle and each transducer to random draws from an exponential distribution. Both transducers must return good signals for the triangulation technique to work, allowing the vehicle's position to be calculated and the perceived position to be updated. In either case the next Fix event is scheduled with a delay equal to the ping interval.

The vehicle drives back toward the original track once a position fix is executed. The vehicle's position is projected onto the original track and the vehicle alters course to intercept the track at a point two fix intervals farther down the track. An Intercept event is scheduled by the Fix event execution. Any prior Intercept events are interrupted by the Fix event, as well. Eventually the vehicle will be in a position where the ultimate destination is closer than the point calculated for intercept. In this case, the vehicle is given orders to move to the final destination. All these tasks are executed only if the vehicle is actually moving

The vehicle can travel any amount of time after a good fix prior to executing another good fix. In the interim, it uses DR to navigate. A course alteration is executed if the vehicle reaches the track prior to the next fix occurring and the vehicle is ordered to move to the final destination. This behavior is handled by the Intercept event.

The vehicle uses transducers to triangulate its position. Their behavior is implemented by use of the Transducer component. Like the RemusMover component, the Transducer maintains a real and an perceived location. The perceived location is the spot where the vehicle thinks the transducer is located based upon pre-programmed information. The real location reflects position uncertainty generated by drop error, current, and wave action effects.

Drop error, set at the beginning of each run, accounts for uncertainties in operator use of GPS to position each transducer. It is implemented by adjusting the perceived location and storing the new location as the real location. The current causes the transducer to be offset from its perceived location proportional to the intensity of the current and in the current's direction. Hydrodynamic theory is complicated and the detail possible in modeling this phenomenon exceeds the intent of this work. Instead, we simplify the transducer location errors to a bivariate normal distribution, where the X axis is aligned with current direction. This distribution is chosen for its ability to provide coverage in two dimensions and for the ease of adjusting the variance in only one direction at a time.

As Figure 5 shows, increasing current intensity tends to displace the transducer from its perfect vertical tending position in the direction of current flow. Increased current intensities shift the mean more, and reductions in current intensity shift the mean less. Wave action is modeled via

the standard deviation of the X component of the bivariate normal. Higher wave height tends to cause more movement in the transducer and the standard deviation is larger. A bivariate normal with X and Y means and standard deviations of 0.0 feet and 1.0 feet, respectively, is used to model a calm sea with no current. The current speed, current direction, and wave height are fixed prior to each run.

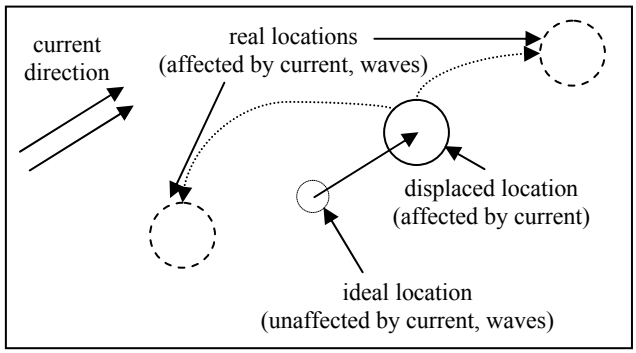

Figure 5: Transducer Current and Wave Action

4 EXPERIMENTAL DESIGN

Understanding how the model inputs and outputs are related is important, but improperly designed experiments can lead analysts to draw erroneous conclusions. Exploring every possible combination of input factors exhaustively would require prohibitively large amounts of computing time. However, limiting the number of inputs factors or their levels might prevent complex interactions from being discovered. Instead, we make use of Latin Hypercube Sampling (LHS). These efficient designs are detailed enough to uncover complex behaviors, and have been used successfully in a number of simulation experiments (Sanchez and Lucas 2002; Kleijnen et al., 2004).

Implementing LHS in the REMUS model is fairly straightforward. We randomly produce multiple LHS matrices, and discard any for which the maximum pairwise correlation exceeds 0.25. Typically, about 60 such matrices are found from 10,000 candidates. The resulting matrices can then be stacked. For example, if ten matrices are stacked, each with ten design points, then the experiment will have 100 total design points. This leads to lower pairwise correlations (usually < 0.05). The near-orthogonality of the stacked designs also aids in interpreting the results.

5 ANALYSIS OF DEAD RECKONING MODEL

We begin by presenting results for the DR model. The sensor detection range is set to 35 meters, or approximately 115 feet. This reflects the actual limitation of the side-scan sonar onboard the vehicle. The vehicle will start from the origin, $(x, y) = (0, 0)$ with orders to move to the point $(x, y) = (1000, 0)$. All dimensions are in feet. The vehicle has a speed of 5 knots. An obstacle is placed at $(x, y) = (500, 0)$. These dimensions were picked in an attempt to mimic actual employment of the vehicle. No transducer field was utilized for this model.

5.1 Probability of Detection

Of primary interest is whether or not the vehicle will detect an object given a set of operating conditions over which the operator has control. These include current direction and speed to the extent that the vehicle can be operated to take advantage of existing conditions. 100 LHS matrices were produced and stacked to obtain 800 design points. The four factors (each at eight levels) were the current direction, current speed, and current noise as well as the maximum compass error. Current directions were then combined to three indicator variables corresponding to non-zero current offsets.

A linear regression is not appropriate for modeling the probability of detection because predicted probabilities could be less than zero or greater than one. Logit regression provides a more realistic model for probabilities than linear regression (Hamilton 1992). Using the JMP™ software, we fit the logit L as a function of the input terms $X_1, \ldots X_k$, estimating the β coefficients in:

$$L = \beta_0 + \beta_1 X_1 + \cdots + \beta_{K-1} X_{K-1}.$$

The expression $\hat{P} = 1/(1 + \exp(\hat{L}))$ can then be used to predict the probability of detection as a function of the input factors. Several models of L were fit, involving the current speed, current noise, three current offset values, and interactions as potential explanatory variables. Details are provided in Allen (2004). The best model for predicting the probability of detection under DR navigation involved all main effects and eight interactions. It was statistically significant (p-value < .0001) with R-squared(U) = 0.683. Here, R-squared(U) is not not the typical R-squared value, but it does provide a measure of how the uncertainty explained by the model.

The fitted coefficients made intuitive sense, but graphical displays make it easier to jointly examine their impact. For example, Figure 6 is a contour plot of the predicted probabilities as a function of current speed and current offset. The detection probability drops as current speed increases, and the most rapid decrease occurs when the current offset is approximately 90 degrees—in short, when the current is perpendicular to the vehicle track and acting to push the vehicle off track. Predicted probabilities are acceptable with current speeds as high as one knot under most current directions. It is obvious that operating the vehicle in low intensity currents is desirable, but if current intensity cannot be minimized then the vehicle should be operated with the current acting to push the vehicle from behind or from ahead. Deciding between these two alternatives is best explored by considering a model of the magnitude of location error given that detection has occurred.

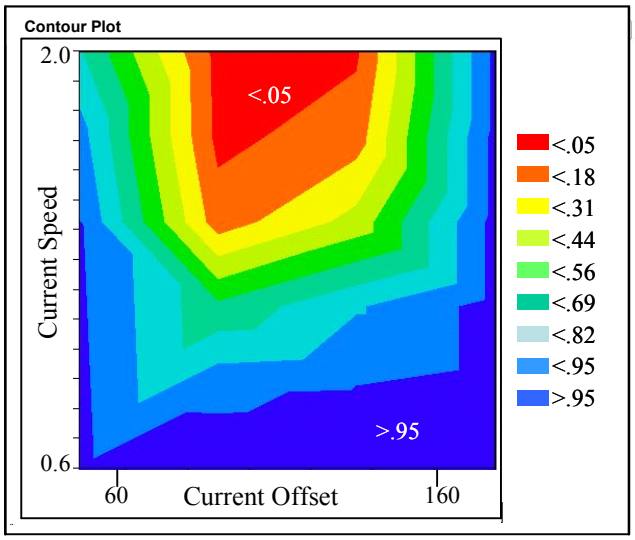

Figure 6: Predicted Detection Probability Given Current Speed and Current Offset for DR Navigation

5.2 Magnitude of Error Given Detection Occurs

A subset of the output data, corresponding to only those repetitions in which detection occurs, are utilized for modeling the magnitude of error. This model seeks to predict mine location offset errors given that mine detection has occurred. As Figure 7 illustrates, the data reveal some interesting patterns in the perceived X and Y directions (MineX and MineY, respectively) that must be considered before attempting to build the model.

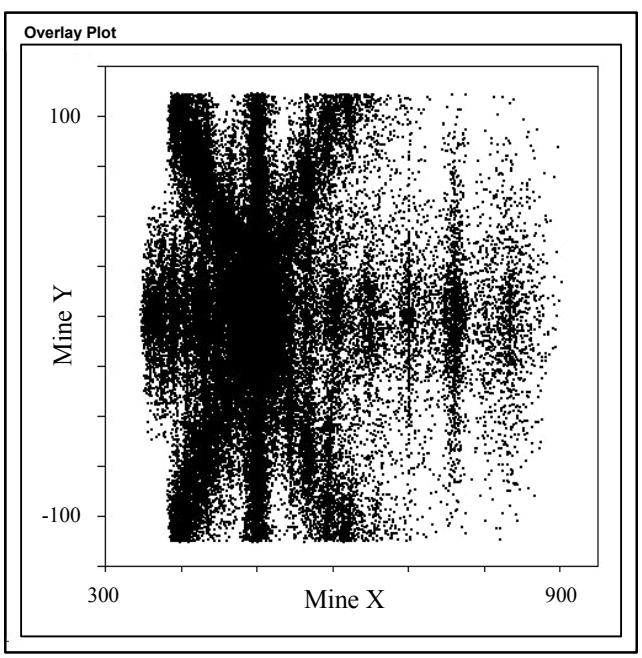

Figure 7: Perceived Mine Location for DR Navigation

First, note that the upper and lower limit for the MineY prediction is 115 feet. This corresponds to the maximum range of the cookie cutter sensor and represents a truncation point in data collection. Secondly, note that "stripes" in the plot result from the way design points were constructed. The factor levels were varied in discrete increments for simplification, but in reality these levels are continuous. Introducing more granularity in the model design points would fill the input space more evenly, but at the cost of additional computational complexity. Recall that the mine was placed at $(x, y) = (500, 0)$ for the DR model runs. The MineY distribution is effectively symmetric about the Y axis but the MineX distribution is skewed with a heavy right tail. See Figure 8 for more detail.

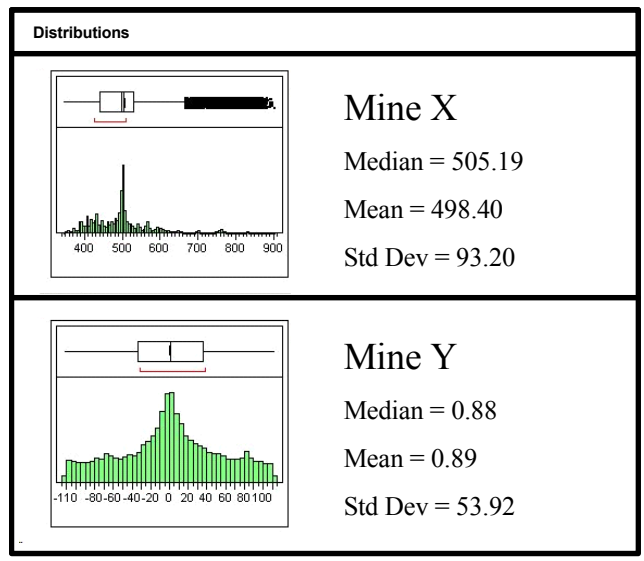

Figure 8: Summary Statistics of Perceived Mine Location for DR Navigation

We fit several regression models (detailed in Allen 2004) to the data. A model containing main effects for current speed and current offset indicators, a quadratic term for current speed, and interactions between speed and offsets provides an extremely good fit (p-value = .0000, adjusted R-squared = 0.98). The base case, represented by the intercept condition, is when current direction is zero degrees relative to the vehicle, or pushing directly from behind. The current offset of 45 degrees is the only variable that was not statistically significant. All of the coefficients make sense. For example, the positive coefficient for current speed implies that as speed increases the predicted offset will also increase. Additionally, when current offset is at 180 degrees—pushing directly from ahead—the predicted offset increases. A contour plot (Allen 2004) indicates that the gradient in predicted mean offset is quite modest when the current offset is at most 45 degrees. The gradient at 180 degrees is much steeper, and the predicted offsets can be quite large for currents as low as one knot.

6 ANALYSIS OF TRANSDUCER MODEL

We use a similar approach, incorporating factors for the transducer drop error, ping interval, and sea state into the design. Transducer drop error is dependent on GPS accuracy and operator proficiency. We model drop error with an equivariant bivariate normal distribution, and implement the impact of wave action via the standard deviation of the X component of the bivariate normal used for current effects, according to the sea state level (zero to four). (The REMUS does not operate in higher sea states.) This creates a larger ellipse of uncertainty for the transducer's real location. Nominally, operators use a five second ping interval because of a perceived tradeoff between navigation accuracy and computing time onboard the vehicle. We consider ping intervals of 2-9 seconds since we are interested in determining whether reducing (or increasing) the interval would significantly improve (or degrade) accuracy. Our design involves 50 stacked LHS matrices.

The transducers are nominally placed at the locations (0, -100) and (1000, -100). Other characteristics, such as the sensor detection range, vehicle starting and ending positions, obstacle location, and vehicle speed remain the same as in the DR experiment.

6.1 Probability of Detection

The output reveals that current speeds less than 0.86 knots, or currents pushing directly ahead or behind the vehicle, always detect the mine regardless of other factors. We remove these design points before continuing our analysis.

Once again, we use logistic regression to fit models for predicting the probability of detecting the mine. A simple model with seven main effects has an R-squared(U) of 0.55, but it can be improved slightly (R-squared(U)=.59) by the addition of six interaction terms and the remaining main effect. (Both models are statistically significant with p-value = .0000.) The results are intuitive: the probability of detection decreases as current speed increases, or as the interval between fixes increases. We also find that high sea states decrease the detection probability, but the impact is mitigated if the ping interval is short.

Contour plots and other graphics again assist in interpreting the results. They show that when the REMUS uses transducers, it can operate with reasonably high detection probability in higher current conditions. As before, the best performance occurs when the vehicle operates with current either directly opposing or directly aiding its path.

6.2 Magnitude of Error Given Detection Occurs

Once again, the question about which direction is best is addressed more effectively by modeling the mean offset as a function of the input factors the operator can control. Overall, the bivariate distribution of perceived location is lighter-tailed and more symmetric for the transducer experiment than for the DR experiment. This means that simple search area calculations are possible. Ideally, the vehicle should be operated in low current conditions with very little wave action. Smaller ping intervals improve detection. The mean drop error is perhaps the one factor most under the operator's control, and pains must be taken to accurately place each transducer. Detection is best when the transducers are placed so the current direction and vehicle path direction are the same. Overall, using transducers for navigation improves the vehicle's performance.

7 CONCLUSIONS AND FURTHER WORK

By building and exploring models of the REMUS vehicle and navigation methods, we found that operating the vehicle with lower current speed improves both the prediction probability of detecting a mine and the offset error given the mine is detected. If higher current speeds cannot be avoided, the model clearly illustrates that operating the vehicle with the current pushing the vehicle either directly from behind or ahead are the most advantageous for improving prediction probability and offset error. Therefore, operate the vehicle with current speeds as low as possible and with current pushing the vehicle from behind to minimize offset errors and maximize detection probability.

The results also suggest that dead reckoning simply does not provide enough predictive power or accuracy to be relied upon as the vehicle's primary navigation method. Operating the REMUS vehicle with transducers is clearly the desired mode of operation. Our models show that detection probability improves when transducers are utilized, and experience in the field backs this up.

Our investigation provides additional insights, since understanding the impact of factors under the operator's control can improve prediction accuracy and thus reduce the amount of time necessary to locate the obstacle post mission. Ideally, the vehicle should be operated in low current conditions with very little wave action. The operator should endeavor to place the transducers with maximum accuracy, and use the smallest possible ping interval as determined by the maximum distance the vehicle is expected to operate from the transducers. If noticeable currents are present, then the vehicle should be operated with the current acting directly from ahead or behind to maximize probability of detection.

While the results of our study make intuitive sense, little data exists for validation. Therefore, the model may be used to gain insight into the effect of combinations of inputs on prediction accuracy, but caution should be maintained about the output being the "right" answer. As field data become available over time, this type of analysis may be the basis for a tool operators can use to predict the time needed to search and relocate mines, even for more complex missions like the "area sweep" model (Allen 2004).

The REMUS vehicle is expected to have the capability to obtain GPS fixes in the near future. The impact of this technology has been hypothesized but not fully explored. It is thought that GPS fixes would significantly improve the accuracy of the vehicle, but since obtaining fixes requires the vehicle to surface, hence lengthening the search operation, GPS is not a panacea. Implementing operational behavior of GPS into a simulation model such as the REMUS model in this paper would go a long way toward answering questions about how best to employ unmanned underwater vehicles, and how much improvement can be expected.

REFERENCES

Allen, T. E. 2004. Using discrete event simulation to assess obstacle location accuracy in the REMUS unmanned underwater vehicle. M.S. thesis, Naval Postgraduate School, Monterey, California. Available online via <http://library.nps.navy.mil/uhtbin/hyperion/04Jun_Allen.pdf> [accessed July 14, 2004].

Buss, A. 2002. Component based simulation modeling with Simkit. In *Proceedings of the 2002 Winter Simulation Conference*, ed. E. Yücesan, C.-H. Chen, J. L. Snowdon, and J. M. Charnes, 243-249. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

Buss, A. H. and P. J. Sanchez. 2002. Building complex models with LEGOs (listener event graph objects). In *Proceedings of the 2002 Winter Simulation Conference,* ed. E. Yücesan, C.-H. Chen, J. L. Snowdon, and J. M. Charnes, 732-737. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

Eisman, D. 2003. Navy ships seize boats carrying mines in Iraqi port. *The Virginian-Pilot*. Available online via <http://www.globalsecurity.org/org/news/2003/030322-mineboats01.htm> [accessed July 14, 2004].

Hamilton, L. 1992. *Regression with Graphics, A Second Course in Applied Statistics*. Belmont, California: Duxbury Press.

Kleijnen, J. P. C., S. M. Sanchez, T. W. Lucas, and T. M. Cioppa. 2004. A user's guide to the brave new world of simulation experiments. Working paper, Department of Information Management/Center for Economic Research (CentER), Tilburg University, Tilburg, The Netherlands.

Sanchez, S. M. and T. W. Lucas. 2002. Exploring the world of agent-based simulations : simple models, complex analyses. In *Proceedings of the 2002 Winter Simulation Conference,* ed. E. Yücesan, C.-H. Chen, J. L. Snowdon, and J. M. Charnes, 116-126. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

Schruben, L. 1983. Simulation modeling with event graphs. *Communications of the ACM* 26 (11): 957-963.

Stewart, J. 1999. *Calculus.* 4th ed. Pacific Grove, California: Brooks/Cole.

United States Navy. 1994. NAVEDTRA 14152, *Mineman, Volume. 01*. Available online via <http://www.tpub.com/content/combat/> [accessed July 14, 2004].

AUTHOR BIOGRAPHIES

TIMOTHY E. ALLEN is a Lieutenant in the United States Navy. He received a B.S. in Applied Mathematics from the University of Idaho in 1997 where he was commissioned as an Ensign in the U.S. Navy. LT Allen completed 12 years of enlisted service as a nuclear operator and achieved the rank of E-7 prior to accepting his commission. Upon completion of the Nuclear Power training pipeline and the Submarine Officer Basic Course in December 1999, he was assigned to the USS ALASKA (SSBN 732) as the Electrical Assistant. Here he qualified in submarines and completed the necessary qualifications for assignment as an Engineer Officer. LT Allen also filled billets as Main Propulsion Assistant, Communications Officer and Assistant Engineer while assigned to ALASKA. In July 2002, LT Allen reported to the Naval Postgraduate School. He earned his M.S. in Operations Research in June 2004. He is currently assigned to the Submarine Officer Advanced Course in Groton, CT, and will report as the Engineer on USS OHIO later this year. LT Allen's e-mail is <tnjallen@mac.com >.

ARNOLD H. BUSS is a Research Assistant Professor in the MOVES Institute at the Naval Postgraduate School, where he has been teaching and doing research in Discrete Event Simulation and military applications since 1994. His e-mail is <abuss@nps.edu>.

SUSAN M. SANCHEZ is a Professor of Operations Research at the Naval Postgraduate School, where she holds a joint appointment in the Graduate School of Business and Public Policy. She received her B.S. in Industrial and Operations Engineering from the University of Michigan, and her M.S. and Ph.D. in Operations Research from Cornell University. She is a member of INFORMS, DSI, ASA, and ASQ. She is the ASA representative to the WSC Board of Directors, and just completed a term as president of the INFORMS College on Simulation. She is the Simulation Area Editor for the INFORMS Journal on Computing, Her research interests include experimental design, data-intensive statistics, and robust selection. Her e-mail is <ssanchez@nps.edu> and her web address is <http://diana.or.nps.navy.mil/~susan>.