COMPUTATIONAL ENVIRONMENTAL MODELS AID SENSOR PLACEMENT OPTIMIZATION

Lucas Vickers
Rustam Stolkin
Jeffrey V. Nickerson
Stevens Institute of Technology, NJ
{lvickers, rstolkin, jnickerson}@stevens.edu

ABSTRACT

Sensors used in protective applications are typically placed on perimeters or over areas in an evenly distributed pattern. However, such patterns may actually be suboptimal, since environmental factors may make some forms of attack more or less likely than others. We describe a protective application of sensors for detecting underwater threats in an urban estuary environment. We demonstrate that environmental information, derived from a computational river current model, can be utilized to optimize sensor placement, increasing detection rates and decreasing the number of required sensors. Simulation results show a significant improvement in detection for a given number of sensors; alternatively, fewer sensors can be used while still maintaining the detection rate of a conventional approach.

INTRODUCTION

Detecting and tracking moving objects is difficult. Sensory data and other information are often sparse or incomplete and are always associated with a degree of uncertainty (Heeger 2003) which may, itself, be hard to estimate.

Sometimes it is possible to incorporate additional information into decision making by combining prior knowledge of the environment in which the object is moving with a statistical model of the object's behavior in response to that environment. In Musman, Lehner, and Elsaesser (1997) mobility and terrain analysis are used to predict possible movement plans for targets. Bayesian nets are used to generate logistical plans for target searches. The search plans focus on organizing a limited number of mobile sensor systems to detect targets which are believed to have specific patterns of behavior and preferred terrains. Other techniques exist in the literature for designing sensor networks around environmental obstacles such as walls or cliffs (Dhillon, Chakrabarty, and Iyengar 2002).

An interesting example is the case of detecting and tracking intentionally moving objects in a river, since rivers contain distinctive and varying currents which impose constraints on the objects' motion. Forecast river current data can be derived from a computational estuarine model (Blumberg and Bruno 2003).

Our previous work has explored different configurations of both moving (Stolkin and Nickerson 2005) and stationary (Olariu and Nickerson 2005) sensors. To our knowledge, no previous work considers the effects of current speeds on intentionally moving objects, and in turn how sensor networks should be designed to take advantage of those effects.

Our research question is the following: using environmental information, by how much can we improve the detection of intentionally moving underwater entities? We focus on the detection of a diver swimming in the Hudson River.

The contribution of this paper is related to the optimal placement configuration of a group of sensors. We will show that the use of environmental information, in the form of a computational estuarine model, can lead to unusual-looking sensor patterns that are more effective than the conventional, equally distributed patterns which are typically deployed.

We first discuss environmental information derived from our computational estuarine model for the New York Hudson River, and then analyze a simple model of diver behavior in response to this environment. We next present a simple, probabilistic sensor model and use this to reason about the detection rates of an arbitrary arrangement of sensors.

We show how environmental data can be used to optimize the positions of a group of sensors, and evaluate the benefits of this optimization, in terms of reduced numbers of required sensors and increased detection rates, as compared with conventional configurations.

ENVIRONMENT MODEL

We examine the problem of optimizing sensor placements with respect to environmental data in the form of current speed forecasts for the Hudson River, derived from our computational estuarine model, The New York Harbor Observing and Prediction System (NYHOPS) (Blumberg and Bruno 2003). In this paper, we look at a simple one dimensional cross section of the Hudson River. However, the approach and formulae described in this paper extend easily to two and three dimensions, which will be the focus of future research.

We focus on a segment of the Hudson River located adjacent to New York City. The latitude is approximately 40.8811 degrees north, and the longitude is approximately 73.9383 degrees west, close to where the Harlem River meets the Hudson River. The current data is from March 15, 2004, 8:30AM. The surface currents in this section of the river flow north. This may seem counter intuitive; however the Hudson River is a tidal estuary. Tides in the Atlantic Ocean will cause water in the Hudson to rise and fall causing complex currents to occur and making optimal sensor placement an interesting and dynamic problem.

The computational estuarine model forecasts current speeds for locations on a grid of cells, spanning the lower Hudson River and New York harbor. At the time of writing, a high resolution model is being developed, however for proof of principle we have generated a curve, $v_x = f(x)$, showing variation of current speed, v, with position, x, across the river by least-squares fitting a polynomial function to current values from a relatively low resolution forecast for a line of cells across the river (figure 1).

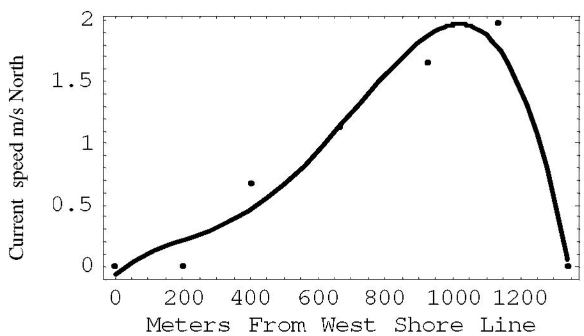

Figure 1. Current speeds for cross section of Hudson River. Polynomial interpolation of forecast current speeds from the NYHOPS computational estuarine model.

DIVER CHARACTERISTICS

We model the diver's interaction with the environment using a "diver preference" curve, a simple function of current speed based on advice from our expert divers. This curve can be viewed as a probability density function, describing the likelihood of a diver choosing or managing to swim in water of various current speeds.

We assume that a diver has the intention of moving forward, will therefore avoid negative currents (i.e. those moving opposite to the diver's intended direction of motion) and is unable to sustain progress against negative currents of greater than 0.5 m/s (approximately 1 mile per hour).

Since moving with a current requires less effort, we assume that a diver has a preference for currents which move in his desired direction of travel (denoted as positive current speeds). However, for reasons of safety and control, the diver will also avoid very fast currents (greater than 4 meters/second). Even if these are in his preferred direction, moving too fast underwater causes a danger of entanglement in obstacles, or can disorient a diver reducing his ability to navigate accurately, especially in turbid and debris strewn urban waterways.

Diver current speed preference can be conveniently modeled using a Log-Normal distribution (figure 2), which is able to represent zero probability of fighting currents beyond the divers maximum sustained swimming speed (0.5m/s). In contrast, a conventional Gaussian distribution is not able to model zero probabilities for these negative current speeds.

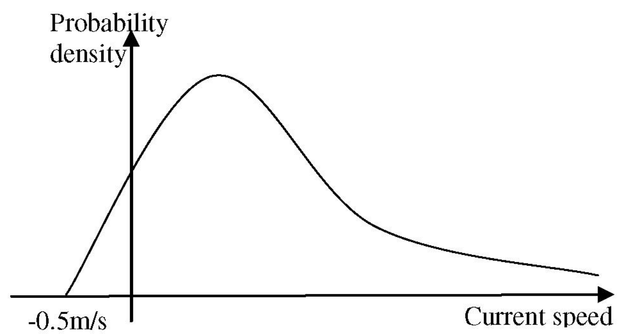

Figure 2. Log-Normal diver preference curve.

PROBABILITY OF DETECTION

We consider the case of detecting a diver who crosses a linear array of sensors stretched across a river. We denote the conditional probability that the i^{th} sensor from this array will detect the diver, given that he crosses the line of sensors at a particular point, x, by $P(D^i \mid x)$. Assuming that the performance of each sensor is independent of the others, it is easy to show that the total probability that a diver, crossing at point x, be detected by the sensor array is given by:

$$P(D^T \mid x) = 1 - P(\overline{D}^T \mid x) = 1 - \prod_i \left\{ 1 - P(D^i \mid x) \right\} \qquad \textbf{(1)}$$

The terms, $P(D^i \mid x)$, are calculated using a simple sensor model based on recent work on passive acoustic diver detection (Stolkin et al., 2006). Stolkin et al. suggest that divers can be detected by thresholding a feature value, the "swimmer number", derived from a passive acoustic hydrophone signal. Their work also experimentally measures the drop off in this feature value with range to the diver (figure 3) and compares with feature values obtained for various levels of background noise. Although probability of detection can be shown to fall off with range as a more complex function, for the purposes of this paper it is reasonably approximated with a simple model (figure 4) in which probability of detection decreases linearly with range, starting from a maximum value of 95% to reflect the fact that detection is never guaranteed.

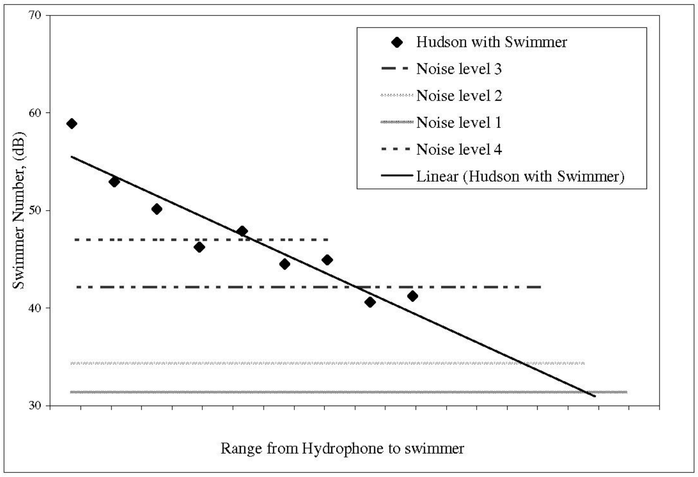

Figure 3 (reproduced from Stolkin et al., 2006). Drop off in log Swimmer Number value with range. Comparison with Swimmer Number calculated for various ambient noise conditions. Noise level 1: River noise with low traffic levels, at nighttime. Noise level 2: River with ferry and helicopter noise. Noise level 3: Rough surface conditions, large waves and two helicopters present. Noise level 4: Severe background noise sources including airplane and helicopter traffic, speed boat and ferry.

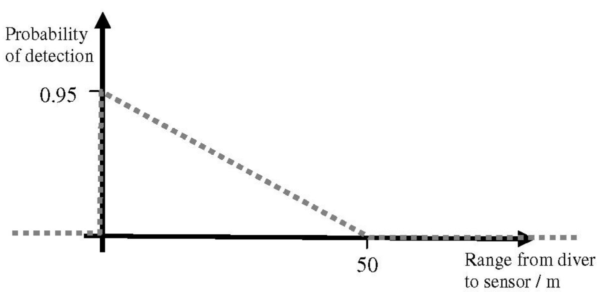

Figure 4. Simple sensor model. Probability of detection decreases linearly with range.

Hence we can define:

$$P\left(D^i \mid x\right)=\begin{cases} 0.95\left(1-\dfrac{\left|x-x_i\right|}{50}\right) & \text{if } 0 \le \left|x-x_i\right| < 50 \\ 0 & \text{all other } \left|x-x_i\right| \end{cases} \qquad \textbf{(2)}$$

where x_i denotes the position of the sensor.

Note that other authors suggest alternative sensor models, and obviously different sensor models are appropriate for different kinds of sensor system. Dhillon et al. (2002), use a probability of detection function which decays exponentially with range. Their exponential model is used to describe a wide variety of different kinds of sensors and is chosen somewhat arbitrarily. Note that the techniques described in this paper will work equally well with any function of detection probability versus range and thus will be useful for reasoning about a wide range of different sensors and sensor models. We demonstrate proof of principle using the simple linear model in this paper.

We now examine the joint probability that a diver chooses to cross the sensor array at a particular point, x, and is detected when he does so. Again, assuming that sensor performance is independent of the diver's decision, we have:

$$P\left(D^T, x\right) = P(x)P\left(\overline{D}^T \mid x\right) = P\left(D^T \mid x\right)p(x)\delta x \qquad \textbf{(3)}$$

where $p(x)$ is the probability density function which describes the likelihood that locations, x, along the array line will be the site of an attempted crossing by the diver.

The term, $p(x)$, is defined by the "diver preference" curve (see figure 2) which is a function of current speed, i.e. $p(x)=f(v_x)$, where the river current speed (v_x) at location x, is found from the current profile data output by our computational estuarine model (see figure 1). The "diver preference" curve (figure 2) is chosen to be a log-normal distribution to reflect the fact that it is impossible for the diver to swim against strong negative currents, and also that the diver will not wish to travel at very high speeds in the highly turbid and cluttered river environment. Hence:

$$p(x) = f(v_x) = \frac{1}{(v_x+0.5)\sigma\sqrt{2\pi}} e^{-\{\ln(v_x+0.5)-\mu\}^2/2\sigma^2} \qquad \textbf{(4)}$$

Using our model derived current information (figure 1), it is now possible to construct the curve (figure 5), $p(x)=f(v_x)$, showing the likelihood of a diver choosing to cross the sensor line at any point x:

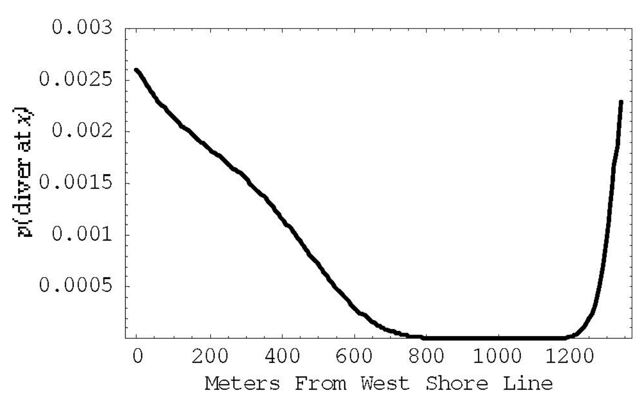

Figure 5. Probability density curve showing likelihood that diver will choose to cross the sensor line at each point in the river.

Equation 1 is useful in that it gives us a measure of sensor "coverage" at any point across the river. Equation 3 can be thought of as a weighted "coverage" term that has been adjusted to take into account the different likelihoods of attack at different locations.

We now seek a single term that describes the effectiveness of a particular arrangement of sensor placements. This term can then be used to optimize sensor placement positions. For this term, we seek to evaluate the total probability that a diver will be detected, should he attempt to cross the sensor array line. This is readily found by integrating equation 3 over all locations, x, across the river:

$$P\left(D^T\right) = \int_{x \text{ on East\_bank}}^{x \text{ on West\_bank}} P\left(D^T \mid x\right) p(x).dx \qquad \textbf{(5)}$$

RESULTS

First we test a conventional, equally distributed arrangement using 35 sensors. Our x axis spans a 1340 meter cross section of the river, so that each sensor will be located every 38.29 meters. This creates a large sensor overlap, providing a relatively high probability of detection. Figure 6 represents this sensor layout.

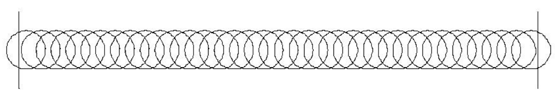

Figure 6. Equally distributed arrangement of 35 sensors across 1340 meter wide of river. Vertical lines represent shore line. Circles represent maximum detection range of sensors.

Figure 7 shows the probability of detection at each point in the river based on this sensor layout given no further information, i.e. the probability of detection assuming (erroneously) that the diver is equally likely to attempt an attack in any part of the river. Based on Figure 7, the total probability of detection (found by integrating equation 4 across the river) is estimated to be 0.875.

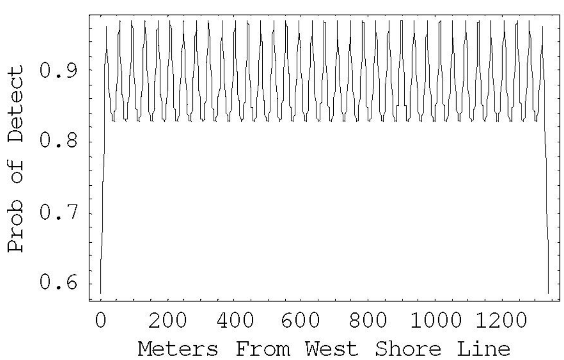

Figure 7. An estimate of the probability of detection at any point in the river, which ignores environmental information.

Figure 8 shows a new estimate of detection probability, which does take into account environmental information about the current: given that the diver crosses the sensor line, the figure shows the likelihood that he crosses at any point x and is detected when he does so. Based on Figure 8, the total probability of detection for a diver traveling south is re-estimated to be the less effective value of 0.869.

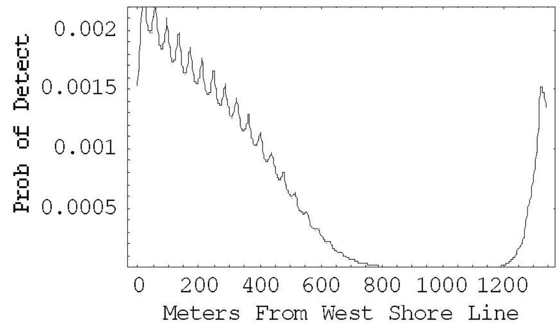

Figure 8. Estimate of the probability of detection at any point in the river, taking into account environmental information. Note that the vertical axis readings are probability densities.

We next examine the results of optimizing the placements of all sensors with respect to the additional environmental information. It is straightforward to employ a standard non-linear optimization technique to optimize all sensor positions by using equation 5 as a fitness function, i.e. modify sensor positions in order to maximize the total probability of detection. Note that a variety of well known non-linear optimization techniques could be used at this stage, e.g. Press et al. (1992), Nocedal and Wright (1999). Broadly speaking these different algorithms trade off speed of convergence against robustness against local minima. The particular choice of algorithm becomes more critical when the number of sensors is very large or when the sensors must be reconfigured in real time. In this paper, for proof of principle, we have used a simple gradient ascent technique. Each optimization was repeated several times, starting from different initial arrangements to prevent local maxima convergence (which does not appear to cause problems with the functions described here).

Figure 9 represents the layout of sensors when optimized with respect to the prior environmental information. The total probability of detection for this arrangement is 0.971 (compare 0.869 for un-optimized arrangement, figure 6).

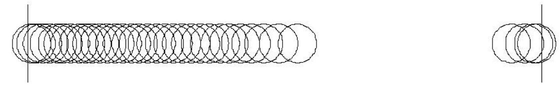

Figure 9. Optimal arrangement of 35 sensors to maximize total probability of detection, given prior knowledge of the environment derived from our computational current model. Vertical lines represent shore line. Circles represent maximum detection range of sensors.

Figure 10 shows, for each point across the river, the probability that the diver crosses at that point and is detected when he does so.

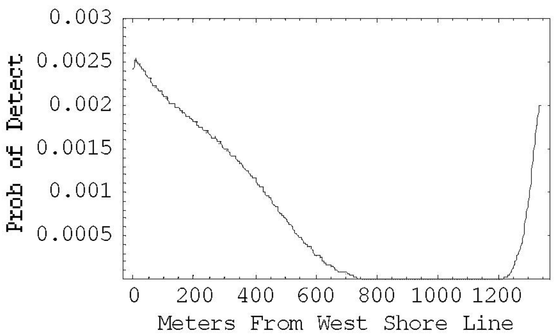

Figure 10. Estimated probability of detection at any point in the river, for environmentally optimized sensor positions. Note that the vertical axis readings are probability densities

Using environmental modeling data to optimize sensor positions, a 97% probability of detection was obtained. This high percentage was obtained by clustering the sensors in areas of high diver probability. The lack of sensors from roughly 800 to 1200 meters may seem counterintuitive, however, due to high current speeds, it is nearly impossible for a diver to cross the sensor line in this area.

For comparison, we ask how many sensors in a conventional, equally distributed arrangement would be required in order to obtain the same 97% detection rate and find the required number to be 58. Hence, our environmentally determined sensor arrangement has saved us 23 sensors; in other words, for the same detection rate, we can reduce the number of sensors by 40% in this example.

For further comparison, optimization was run from 5 to 60 sensors, in increments of 5 sensors. The optimized configurations were compared against equally spaced conventional arrangements created using the same number of sensors.

Figure 11 shows the total probability of detection rates which were obtained using environmentally determined arrangements. These results are compared to those obtained using equally distributed arrangements. The increase in coverage tapers off for very high detection rates.

Figure 12 shows the difference in total probability of detection between an environmentally optimized sensor arrangement and a conventional, equally distributed arrangement. As in Figure 11, the benefits of the optimization taper off with very large numbers of sensors.

Figure 13 illustrates the value of environmental optimization when high probabilities of detection are desired. The number of sensors that can be saved by our technique increases rapidly with desired detection rate, and large numbers of sensors can be saved for high detection rates by using environmental factors as compared to a conventional equally distributed arrangement.

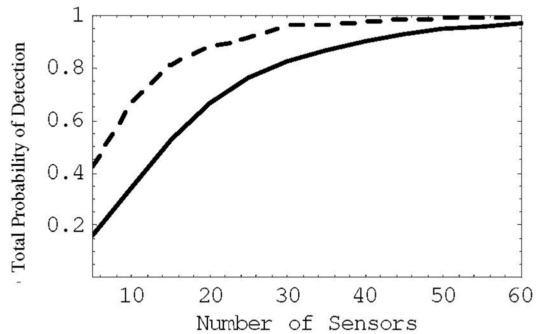

Figure 11. Variation in total probability of detection with number of sensors. Top dashed line represents environmental optimization. Bottom solid line represents conventional equi-spaced arrangement..

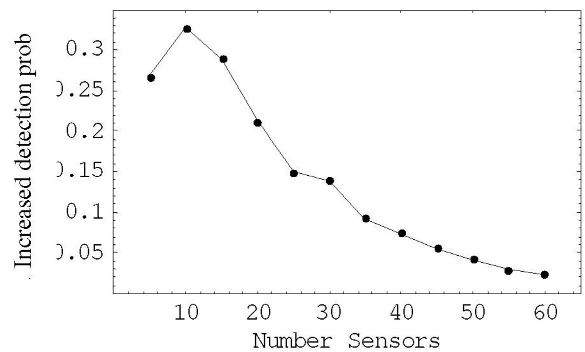

Figure 12. Increase in total probability of detection due to environmental optimization, as compared to a conventional, equally distributed arrangement using the same number of sensors.

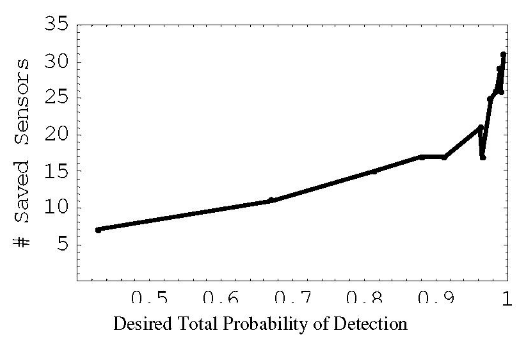

Figure 13. Number of sensors saved using environmental optimization, while maintaining same detection rate as conventional arrangement, for various levels of desired detection rate.

CONCLUSIONS

We have provided a framework for probabilistically reasoning about the location of a scuba diver, taking into consideration data from a computational environment model and a simple model of diver behavior.

We have shown how this leads to a method for optimizing the placement of sensors. This method substantially reduces the number of sensors necessary to achieve a given rate of detection. Alternatively, much greater rates of detection can be achieved with a given number of sensors. The benefits of this environmental optimization appear to decay for very large numbers of sensors but, in contrast, the benefits (in terms of saved sensors) increase dramatically with the required rates of detection.

FUTURE WORK

The techniques described in this paper extend very readily to two, three or four (time) dimensional spaces and this will be a focus of future work. Additional questions which arise in the case of higher dimensional spaces include: will hotspots of high diver probability exist? Can certain areas of high diver probability be ignored if they are surrounded by areas of low probability? Will concurrently optimizing sensor placement for divers traveling in different directions lead to unusual sensor placement structures? Can structures be found which will adequately protect against divers during an entire year, or will different tidal seasons require separate sensor layouts?

An important alternative use of environmental reasoning about diver location is to provide prior probability terms for Bayesian detection, localization and tracking algorithms. We hope to incorporate these ideas into ongoing work exploring probabilistic algorithms for tracking moving targets through distributed arrays of sensors of various kinds.

In addition, it may be possible to use environmental information to guide the path of moving sensors in response to varying environmental conditions, and future work may examine control rules for mobile sensors in response to varying real time current information.

ACKNOWLEGEMENTS

This research was supported by the Office of Naval Research, grant #N00014-05-1-00632.

REFERENCES

A. Blumberg, M. Bruno. The New York Harbor Observing and Prediction System (NYHOPS). http://onr.dl.stevens-tech.edu/webnyhos3/. 2003.

S. Dhillon, K. Chakrabarty, S. S. Iyengar. Sensor placement for grid coverage under imprecise detections. *Proc. 5th ISIF Intl. Conf. on Information Fusion*, pp. 1581-1587. July 2002.

D. Heeger, "Signal Detection Theory." New York University, http://www.cns.nyu.edu/~david/sdt/sdt.html, 2003.

S. A. Musman, P. E. Lehner, C. Elsaesser. Sensor Planning for Elusive Targets. *J. Mathematical and Computer Modeling,* Vol. 25, No. 3, pp. 103-115. 1997.

J. Nocedal, S. Wright. Numerical Optimization. Springer. 1999.

S. Olariu, S. and J. V. Nickerson. Protecting with Sensor Networks: Perimeters and Axes. *MILCOM Workshop on Situation Management*, 2005.

W. Press, S. Teukolsky, W. Vetterling, B. Flannery. Numerical Recipes in C. 2nd Edition. Cambridge University Press. ISBN 0 521 43108 5. 1992.

R. Stolkin, J. V. Nickerson, Combining multiple autonomous mobile sensor behaviors using local clustering. *MILCOM Workshop on Situation Management*, 2005.

R. Stolkin,. A. M Sutin, S. Radhakrishnan, M. Bruno, B. Fullerton, M. Raftery. *Feature-based passive acoustic detection of a diver*, SPIE Defense and Security Symposium, 2006.