Probabilistic analysis of a passive acoustic diver detection system for optimal sensor placement and extensions to localization and tracking

Rustam Stolkin, Ionut Florescu
Stevens Institute of Technology

Abstract— **Our previous work describes a simple algorithm for automated detection of diver presence using a single passive hydrophone. This technique is based on extracting a single feature, the "Swimmer Number", from the hydrophone signal, which correlates with diver presence. At any point in time, diver presence can be automatically determined by thresholding the incoming Swimmer Number at an appropriate level. For this system (and other threshold based detection systems) this paper explains how to calculate the probability of detecting a diver at various ranges from the hydrophone. This function is then used to evaluate the probability of detecting a diver at any point in a region, given an arbitrary number of hydrophones which are scattered in arbitrary positions over the region. We next show how non-linear optimization techniques can be used to find the optimal set of sensor positions, which maximize the detection probability over a region of interest, for a given number of sensors. Lastly we show how this theory can be incorporated into tracking systems, which estimate the location of a moving diver at any point in time, given the outputs of an arbitrarily positioned set of hydrophones.**

I. PASSIVE ACOUSTIC DIVER DETECTION

One of the most challenging aspects of port security is providing the means to protect against threats from under the surface of the water, [1]. In particular, it is felt that a significant terrorist threat might be posed to domestic harbors in the form of an explosive device delivered underwater by a diver using SCUBA apparatus, [2]. Although active sonar systems exist which can detect and track moving targets, e.g. [3], the problem of automatically recognizing which, if any, moving entities are human divers is less well understood. This recognition problem lends itself to a passive acoustic approach, since these techniques can make use of prior knowledge of the specific sounds generated by a diver.

Our previous work, [4], [5], describes a simple algorithm for automated detection of the presence of a diver, using a single passive hydrophone. SCUBA divers emit sounds in a characteristic high frequency range which are associated with breathing. Hence a useful feature, the "Swimmer Number" can be extracted from a hydrophone signal which evaluates to what extent an entity is present which emits regular pulses of sound in the characteristic high frequency range, such that these pulses are repeated at a low frequency lying in the range of typical human breathing rates (0.3-1Hz). This feature takes large values when a diver is present and small values otherwise, even in the presence of severe background noise. This is because most background sound sources are unlikely to share both characteristic frequencies (low and high) with the diver.

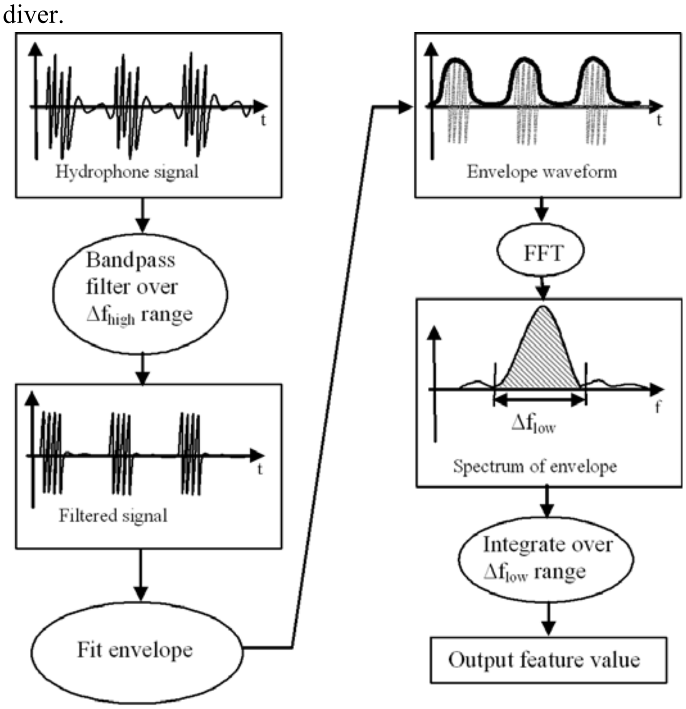

Figure 1. Procedure for extracting a discriminatory feature from hydrophone signals.

The Swimmer Number is calculated as shown in figure 1. Firstly the raw hydrophone signal is narrow bandpass filtered over a small range of frequencies, Δf_{high}, about the characteristic high frequency associated with a SCUBA diver's breathing. Secondly, an envelope is fitted to the filtered signal by connecting peaks and smoothing. Thirdly, the envelope is Fast Fourier Transformed to produce a spectrum for the envelope waveform. Lastly, the spectrum is integrated over the

Manuscript received August 10th, 2007. This work was supported in part by the U.S. Office of Naval Research, Swampworks, grant #N00014-05-1-00632.
R. Stolkin is a Research Assistant Professor at the Center for Maritime Systems, Stevens Institute of Technology, USA. E-mail: RStolkin@stevens.edu. I. Florescu is an Assistant Professor at the Dept. Mathematical Sciences, Stevens Institute of Technolog, USA. Email: ifloresc@stevens.edu.

range, Δf_{low}, representing the typical range of human breathing rates. This procedure yields a single characteristic number. This number is discriminating in that it takes high values when a diver is present and low values otherwise, even in severe conditions of background noise. The number is also “general” in that it works equally well for individuals with different breathing rates.

It is observed that the log(Swimmer Number) falls off approximately linearly with range (figure 2).

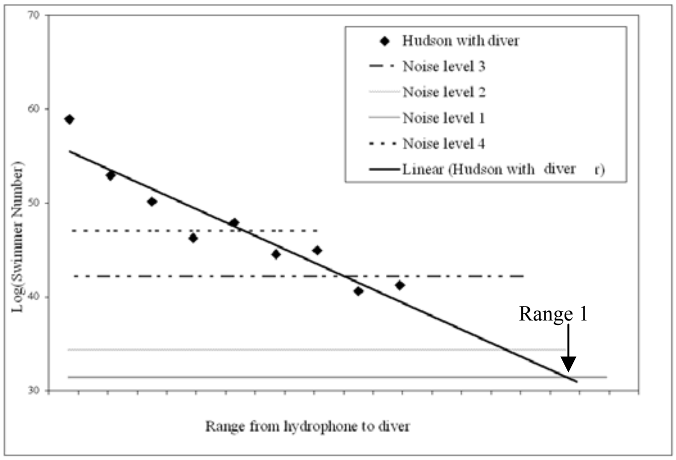

Figure 2. (Reproduced from [5]). Drop off in log(swimmer number) value with range. Comparison with log(swimmer number) calculated for various ambient noise conditions. Noise level 1: River noise with low traffic levels, at night time. Noise level 2: River with ferry and helicopter noise. Noise level 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.

II. Probability of Detection

In practice, as can be seen in figure 2 the swimmer number is not purely deterministic (otherwise it would lie on a straight line), instead it can be treated as a random variable with a distribution which depends on the range to the sensor. A simple approach to automatic detection of object presence, in the case of a single sensor, is to employ a threshold. Any log(Swimmer Number) values above the threshold are taken to indicate diver presence. Thus, the probability, $P(D \mid R)$, that a diver at range, R, is detected, is the probability that the log(Swimmer Number), S, exceeds the threshold value, K:

$$P(D \mid R) = P(S > K \mid R) \tag{1}$$

We use a simple regression of the swimmer number with range which in effect means that the swimmer number at any particular range is assumed to be normally distributed, i.e.

$$S \sim N(\mu_R, \sigma_R) \tag{2}$$

where the mean, μ_R, and standard deviation, σ_R, are themselves dependent on range. These parameters, and their variation with range, can be estimated from experimental measurements. For example, our laboratory is currently undertaking extensive measurements of this kind, deploying expert divers at various known ranges from a hydrophone, both in laboratory tanks and the Hudson River by Manhattan, under a variety of background noise conditions and using a variety of different diving apparatus.

Mean log(Swimmer Number), μ_R, can reasonably be modeled as linearly decreasing with range according to the regression line in figure 2. Similar linear relationships will apply to many kinds of sensor signal which are often modeled as decaying exponentially with range, e.g. [6]. Note that the probability of detecting a diver equals one minus the probability of failing to detect:

$$P(D \mid R) = P(S > K \mid R) = 1 - P(\bar{D} \mid R) = 1 - P(S < K \mid R) \tag{1}$$

Hence, expressing in terms of a normalized Gaussian function:

$$P(D \mid R) = 1 - P\left(\frac{S - \mu_R}{\sigma_R} < \frac{K - \mu_R}{\sigma_R}\right) = 1 - \Phi\left(\frac{K - \mu_R}{\sigma_R}\right), \tag{2}$$

where Φ denotes the distribution function of a standard normal random variable. A convenient (though not strictly true) simplification is to assume that the standard deviation is constant for all ranges, i.e. $\sigma_R = \sigma$. In this case, the probability of detection will simply decrease with range according to the tail of a *cumulative* normal distribution (figure 3).

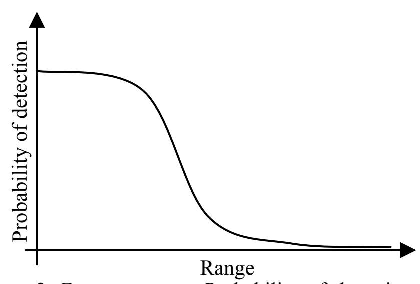

Figure 3. For constant σ, Probability of detection drops off with range according to the tail of a cumulative normal distribution. Similar drop off curves can be computed for more complex functions of σ which themselves vary with range, see [7].

More complex (and more realistic) functions of σ_R can also be handled numerically, as described in [7]. Often standard deviation will itself decrease with range in a similar fashion to the underlying value.

This description of probability of detection has an interesting implication for the maximum detection range. Conventionally, the maximum range might be regarded as the distance at which the Swimmer Number due to a diver drops below the threshold value, e.g. “Range 1” in figure 2. In contrast, the above analysis indicates that there is still a 50% probability of detection at this “maximum” range and significant probabilities of detection at even greater ranges.

III. Optimal Sensor Placement Along a 1D Line

Consider the question: how do we optimally place a certain number of sensors of the type described thus far. As a simple example consider a line of diver detection sensors forming a protective boundary. How far apart can any two sensors be placed such that the minimum probability of detecting a diver, who crosses the boundary at any location between the sensors, exceeds a desired minimum probability of detection? Since the cumulative normal curve drops off very rapidly, the contributions of any other sensors can often be neglected. For the two sensors which bound the point of crossing, the total probability of detection, $P(D^T \mid R)$, is then one minus the probability that both sensors fail to detect:

$$P(D^T \mid R) = 1 - P(\overline{D}_1 \mid R) P(\overline{D}_2 \mid R) \tag{4}$$

i.e. $$P(D^T \mid R) = 1 - \Phi\left(\frac{K - \mu_R}{\sigma_R}\right)\Phi\left(\frac{K - \mu_{L-R}}{\sigma_{L-R}}\right) \tag{3}$$

where L is the distance between two consecutive sensors and R is the range from one of them (figure 4).

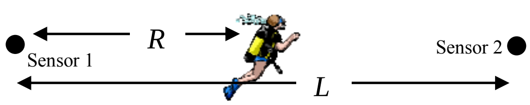

Fig. 4. A diver crosses a boundary at a point somewhere between two diver detection sensors.

For practical purposes, an engineer may wish to determine the minimum number of sensors required in order to achieve a desired minimum detection probability anywhere along this boundary. This can be achieved by preparing a graph, figure 5, showing how detection probability varies with position between the two sensors, for various different sensor spacings.

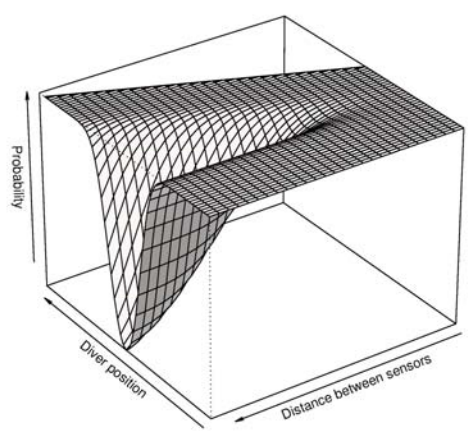

Fig. 5. Probability of detection at all positions between two sensors for various different sensor separation distances. Optimum sensor separation is that distance for which the central minimum is equal to the minimum desired probability of detection at any point.

IV. Optimal Sensor Placement on a 2D Plane

In the previous section the sensors were constrained to a linear protective boundary. Accordingly the influences of other sensors except the two closest to each point along the boundary were considered negligible. For the extension to optimal placement of sensors over a 2D surface we have to consider all the sensors as a system. More specifically, given n sensors, how should they be positioned in order to maximize the probability of detecting a diver in some region of interest? For an arbitrary arrangement of n independent sensors on a plane, the total probability (due to the combined efforts of all sensors) of detecting a diver at a particular position $\mathbf{x}$ is given by:

$$P(D^T \mid \mathbf{x}) = 1 - P(\overline{D}^T \mid \mathbf{x}) = 1 - \prod_{i=1}^{n} \left\{1 - P(D^i \mid \mathbf{x})\right\} \tag{6}$$

i.e. one minus the probability that none of the sensors detect the diver. The probability that the i^{th} sensor detects the diver is given by equation 5, i.e.:

$$P(D^i \mid \mathbf{x}) = 1 - \Phi\left(\frac{K - \mu_{R_i}}{\sigma_{R_i}}\right), \qquad R_i = |\mathbf{x} - \mathbf{x}_i| \tag{7}$$

where $\mathbf{x}_i$ is the position of the i^{th} sensor and μ_{R_i} is found from the regression line of figure 2.

Given a criterion for the "net sensor coverage" of the region, optimal sensor positions can be found as those which maximize this criterion. One such criterion could be the minimum probability of detection anywhere in the region. However this choice of criterion is hard to maximize because it causes many optimization strategies to converge on sub-optimal local maxima. The reason for this is that, for many sensor arrangements, there will be sizeable areas of almost zero detection probability. Gradient based optimizers which arrive at one of these arrangements will be unable to escape as they have become trapped on a region of local maximum.

A better optimization criterion is the expected (i.e. mean) probability of detection over the region. We optimize this second criterion, noting that this simultaneously improves the first criterion.

Many standard non-linear optimization strategies could be used. We prefer Powell's non-linear least squares method, [8], due to its strong performance in high dimensional spaces, since N sensors on a 2D plane require a $2N$ dimensional search space.

Figure 6 shows an example of optimizing the positions of fifteen sensors in order to best protect a square region. These simple techniques can also be used to best position an arbitrary number of sensors over an arbitrarily shaped region. For example, a captain may wish to deploy diver detection sensors to best monitor an exclusion zone around his ship, figure 7.

Note that the optimal sensor positions do not simply lie on an equi-spaced regular square lattice, figure 8.

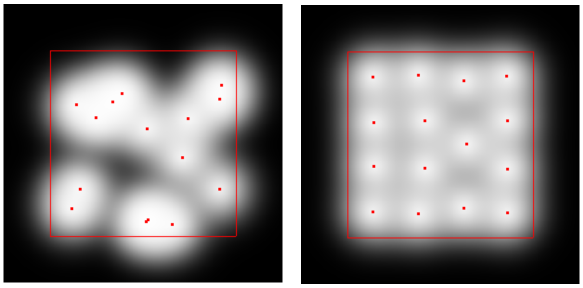

Fig. 6. Fifteen diver detection sensors, initialized with random positions (left) and after optimization (right). Square box denotes the region to be protected. Brightness denotes probability of detection. Minimum probability of detection = 0.005 (left) and 0.28 (right). Expected probability of detection = 0.67 (left) and 0.78 (right).

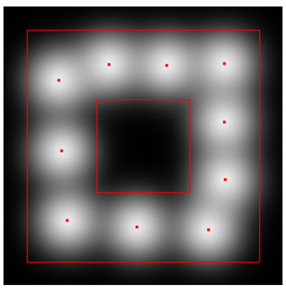

Fig. 7. Optimized sensor positions for odd shaped regions – e.g. monitoring an exclusion zone.

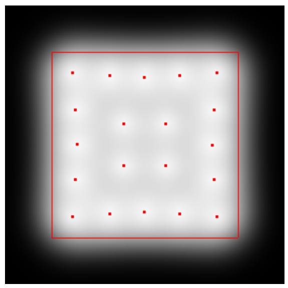

Fig. 8. Optimal positions for 20 sensors to protect a square region. This is not a square lattice, the lines of sensors are distinctly curved.

It is useful to determine the minimum number of sensors required in order to achieve a desired level of protection over a region. This can be achieved by preparing a graph, figure 8, showing how minimum detection probability (after optimization) varies with the number of sensors.

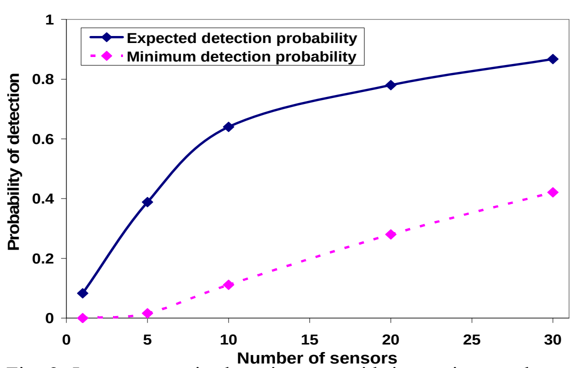

Fig. 8. Improvement in detection rate with increasing numbers of optimally positioned sensors, for the square region shown in figures 5 and 7.

Note that the above technique for optimal sensor placement finds the set of sensor positions for which False Negative errors are minimized. False Negative (FN) errors occur when the system fails to detect a diver that is present. In contrast, we have not addressed the problem of minimizing false positive (FP) errors.

FPs occur when no diver is present, hence no "range" exists. FP errors do not depend on range but on the choice of threshold and the level of background noise. In contrast, given a particular choice of threshold, this paper has shown how to optimize sensor positions to minimize false negative (FN) errors. Since FNs depend on range and FPs do not, during position optimization, FNs are reduced *without increasing FPs*. This is interesting, since FNs and FPs usually present an unavoidable tradeoff – decreasing one usually comes at the expense of increasing the other.

V. Extension to Localization and Tracking

We now briefly outline a way in which the function of probability of detection versus range, derived above, can be incorporated into a probabilistic tracking scheme. This kind of tracking will be a focus of future work.

Consider a set of N sensors, positioned arbitrarily across a region of water to be protected. Each sensor consists of a hydrophone with additional hardware and software for calculating and thresholding Swimmer Number values to determine diver presence in the region local to each sensor.

At any given instant, the i^{th} sensor produces a binary output, O_i, which takes one of two Boolean states ("true" or "false"), indicating whether or not it detects diver presence (i.e. whether or not the Swimmer Number measured at that sensor exceeds the designated detection threshold). The tracking problem is to estimate the location, $\mathbf{x}_t$, of the diver at any time, t, given the set of all sensor outputs, $\mathbf{O}_t = \{O_1, O_2, \dots, O_i, \dots O_N\}_t$ at that instant.

The probability that the diver is at a particular candidate location, $\mathbf{x}_t$, can be evaluated using recursive Bayesian filtering. From Baye's law:

$$P(\mathbf{x}_t \mid \mathbf{O}_t) \propto P(\mathbf{O}_t \mid \mathbf{x}_t) P(\mathbf{x}_t) \quad (8)$$

Assuming that the sensors all perform independently, the conditional term can be evaluated as:

$$P(\mathbf{O}_t \mid \mathbf{x}_t) = \prod_{i=1}^{N} P(O_i \mid \mathbf{x}_t) \quad (9)$$

where

$$P(O_i \mid \mathbf{x}_t) = P(D^i \mid \mathbf{x}_t) \quad \text{if} \quad O_i = \text{"true"} \quad (10)$$

and

$$P(O_i \mid \mathbf{x}_t) = P(\overline{D}^i \mid \mathbf{x}_t) = 1 - P(D^i \mid \mathbf{x}_t) \quad \text{if} \quad O_i = \text{"false"} \quad (11)$$

Equations 10 and 11 can be evaluated using equation 7.

The prior probability term, $P(\mathbf{x}_t)$, can be evaluated from the posterior term at the previous time step via a motion and diffusion model as in Kalman filtering or particle filtering techniques. We suggest a particle filtering approach as in [9], [10], since the distributions are highly non-gaussian and may often be oddly shaped and/or multi-modal.

The following figures show how the binary output of diver detecting sensors can be used to infer diver location based on the conditional probability term, $P(\mathbf{O}_t \mid \mathbf{x}_t)$, described above (equations 9, 10, 11). In each case, green dots represent sensors whose output is "false - no diver detected" and red dots represent sensors whose output is "true - diver detected". Based on these sensor outputs, all locations are assigned a probability of diver presence. White denotes high probability, black denotes low probability, with continuous shades of grey in between.

Figure 10 shows how the binary states of a single sensor place probabilistic constraints on diver location.

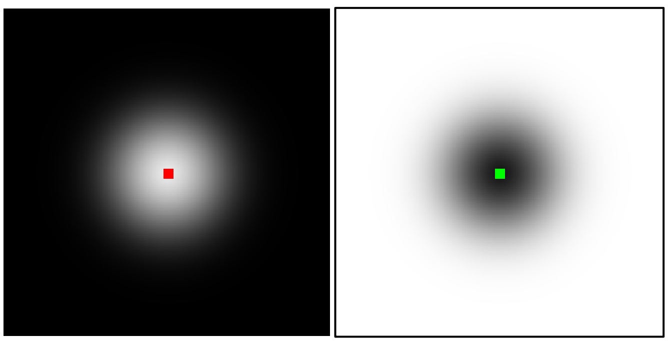

Fig. 10. Binary state of a sensor probabilistically informs diver location.
Left: "diver detected" implies diver is near to sensor.
Right: "diver not detected" implies diver is far from sensor.
Brightness denotes likelihood of diver presence.

Figure 11 shows a simulation of a diver moving through a field of sensors. At each time step, the binary outputs of all sensors are combined to infer the likely location of the diver.

Note, that by using only the binary (thresholded) output of each sensor, we are arguably discarding information. Therefore, future work may also examine techniques which make use of the actual Swimmer Number value itself to determine location. In this case we would seek, for each candidate diver location, to evaluate:

$$P(\mathbf{x}_t \mid \mathbf{S}_t) \propto P(\mathbf{S}_t \mid \mathbf{x}_t) P(\mathbf{x}_t) \qquad (12)$$

where $\mathbf{S}_t = \{S_1, S_2, \ldots, S_i, \ldots, S_N\}_t$ is the set of Swimmer Numbers output by the sensors. It is possible that this alternative approach might outperform the simple approach described above, however the Swimmer Number versus range function is highly noisy. The non-linearity of the previously described approach (above), which utilizes only binary thresholded outputs, may serve to provide added robustness against noise. Future work will implement, test and compare these two potential tracking methods.

VI. CONCLUSIONS AND FUTURE WORK

This paper has derived a simple function of probability of detection versus range for passive acoustic diver detecting sensors. This analysis may also be applicable to many other detection systems which also rely on thresholded feature values derived from sensor signals.

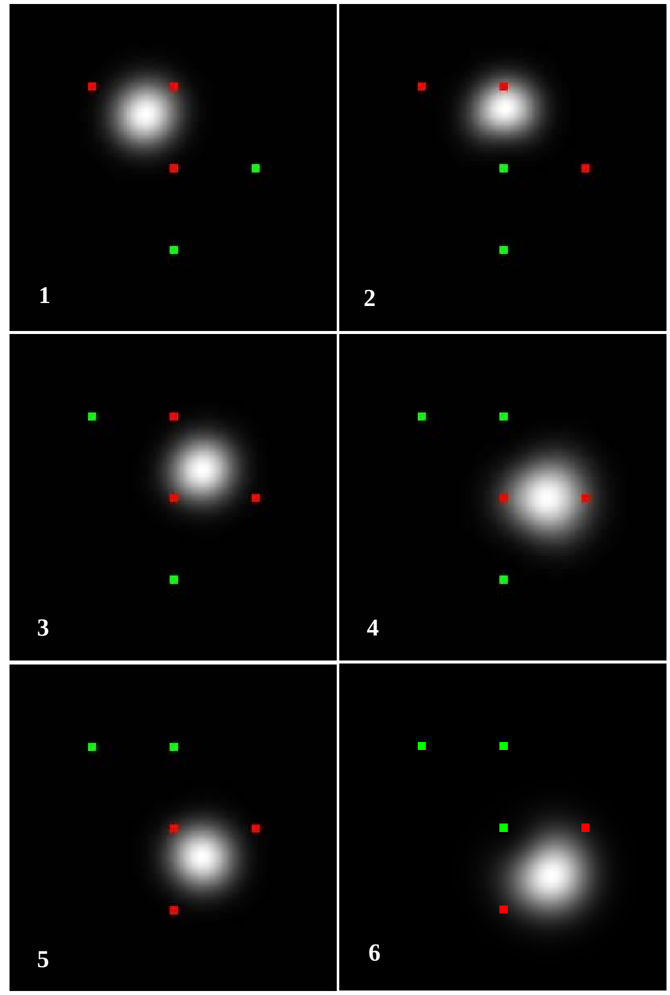

Fig. 11. A diver moves across an arbitrarily spaced field of sensors. At each time step, each sensor either detects diver presence in its local vicinity (red) or fails to detect diver presence (green). Based on this combination of sensor responses, and a curve of detection probability versus range for each sensor, it is possible to reason about the diver's location. For all points, brightness denotes likelihood of diver presence. Note that diver location distributions are not symmetrical and are significantly non-gaussian. Such odd shaped distributions are best represented by "factored sampling" and can be tracked using particle filtering techniques.

The probability of detection versus range function is important because it enables optimal positioning of multiple sensors over a region of interest or along a protective boundary. To perform this optimization with respect to additional knowledge of the detected object or its environment, see [11] and [12], in which diver detection sensors are optimally positioned with respect to river currents forecast by a computational estuarine model.

The probability of detection versus range function is also informative about diver location, and can be used to localize and track a diver which moves across an arbitrarily arranged field of sensors.

Future work will examine more complex models of divers and their corresponding sensor signals. For example, passive acoustic signals from a diver may actually vary with orientation as well as with range. In this case we will examine probability distributions of Swimmer Number versus diver position and orientation, such as $P(\mathbf{S}_t \mid x_t, y_t, \theta_t)$ and $P(x_t, y_t, \theta_t \mid \mathbf{S}_t)$.

REFERENCES

[1] R. Hansen. "Underwater Port Security Signal Processing Challenges," A book of abstracts for the IEEE Underwater Acoustic Signal Processing Workshop, page 9. 2005.

[2] R. Manstan. "Measurements of the Target Strength and Radiated Noise of Divers Wearing SCUBA Equipment," A book of abstracts for the IEEE Underwater Acoustic Signal Processing Workshop, page 13. 2005.

[3] K. Shaw, R. Scott., G. Holdanowicz. 2005. "Sonar sentinels on guard for submerged swimmers", Jane's Navy International, September 2005.

[4] R. Stolkin, A. Sutin, S. Radhakrishnan, M. Bruno, B. Fullerton, A. Ekimov, M. Raftery, "Feature based passive acoustic detection of a diver." SPIE Defense and Security Symposium, 2006.

[5] R. Stolkin, S. Radhakrishnan, A. Sutin, R. Rountree, "Passive acoustic detection of modulated underwater sounds from biological and anthropogenic sources." Proc. IEEE MTS OCEANS conference, 2007.

[6] 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.

[7] R. Stolkin, I. Florescu. Probability of detection and optimal sensor placement for threshold based detection systems. IEEE Sensors Journal, *in review*, 2007.

[8] M. Powell. A Hybrid Method for Nonlinear Equations. In P. Rabinowitz, editor, Numerical Methods for Nonlinear Algebraic Equations, pages 87–144. Gordon and Breach Science, London, 1970.

[9] Dellaert, F., Fox, D., Burgard, W., Thrun, S. Monte Carlo localization for mobile robots. Proc. IEEE International Conference on Robotics and Automation, 1999.

[10] K Nummiaro, E Koller-Meier, L Van Gool. An adaptive color-based particle filter. Journal of Image and Vision Computing, Elsevier, 2003.

[11] R. Stolkin, L Vickers and J. Nickerson, "Using Environmental Models To Optimize Sensor Placement", *IEEE Sensors Journal*, 2007.

[12] L. Vickers, R Stolkin and J. Nickerson "Computational environmental models aid sensor placement optimization", *Proc. Workshop on Situation Management (SIMA), MILCOM*, 2006.