Early maritime applications of particle filtering

Henry R. Richardson[a], Lawrence D. Stone[*b], W. Reynolds Monach[c], Joseph H. Discenza[d]

[a] Consultant, 247 Prospect Bay Drive West, Grasonville, MD 21638;
[b] Metron Inc, 11911 Freedom Drive, Suite 800, Reston VA 20190;
[c] Daniel H. Wagner Associates, 2 Eaton Street, Suite 500, Hampton, VA 23669;
[d] SmartCrane, 2 Eaton Street, Suite 500, Hampton, VA 23669

ABSTRACT

This paper provides a brief history of some operational particle filters that were used by the U. S. Coast Guard and U. S. Navy. Starting in 1974 the Coast Guard system provided Search and Rescue Planning advice for objects lost at sea. The Navy systems were used to plan searches for Soviet submarines in the Atlantic, Pacific, and Mediterranean starting in 1972.

The systems operated in a sequential, Bayesian manner. A prior distribution for the target's location and movement was produced using both objective and subjective information. Based on this distribution, the search assets available, and their detection characteristics, a near-optimal search was planned. Typically, this involved visual searches by Coast Guard aircraft and sonobuoy searches by Navy antisubmarine warfare patrol aircraft. The searches were executed, and the feedback, both detections and lack of detections, was fed into a particle filter to produce the posterior distribution of the target's location. This distribution was used as the prior for the next iteration of planning and search.

Keywords: Particle Filters; Monte Carlo; Bayesian; Optimal Search; Non-Gaussian

1. INTRODUCTION

This paper is largely an historical account of the development of filtering and prediction software for maritime search applications beginning in 1970. The methods employed are now known as "particle filtering," although this term was not in use at the time the work was carried out. Stone[1] and Stone, Corwin, and Barlow[2] discuss optimal search and tracking in more detail.

Section 2 describes the Bayesian mathematical formulation of the filtering and prediction problem, and Section 3 presents the particle filter solution methodology. The early maritime applications of this approach are discussed in Section 4. Acknowledgements are given in Section 5.

1.1. The search problem

The particle filters to be described were developed to enhance the effectiveness of search for an object (the target) that was lost or was attempting to conceal itself from detection. From the beginning, the objective was to develop algorithms and software that could be used in actual operations by Navy and Coast Guard personnel who were not computer experts. The computers available during the 1970's were primitive by today's standards and the design of filtering software was driven to a large extent by memory and processing speed limitations.

It is worth noting that the approaches taken were focused on the end results to be achieved, not filtering and prediction as an end in itself. For example, in many cases it was not necessary to use large numbers of particles to compute highly precise posterior probability distributions for target location, when the available amount of search effort was small. An interesting analysis of this and related issues is provided by J. R. Weisinger[3].

In the typical search problem, one has some initial information concerning the location and detection characteristics of the target of interest. For example, in a search for a lost fisherman one usually has some idea about the intended fishing area as well as a description of the fishing boat. The location of the fishing area provides the basis for

[*] stone@metsci.com; phone 703 787 8700; fax 703 787 3518

Signal and Data Processing of Small Targets 2003,
edited by Oliver E. Drummond, Proceedings of SPIE Vol. 5204
(SPIE, Bellingham, WA, 2003) · 0277-786X/03/$15 · doi: 10.1117/12.505233

postulating an initial probability distribution for target location, and the description of the fishing boat provides an input for calculating the probability of detection by visual, radar, or other sensors.

When the target is moving, the initial target-location distribution is modified by the dynamics of the particular problem. In the case of a drifting fisherman, the winds, currents, and tides displace and distort the target-location distribution. In the case of a target submarine, the target's tactics and significant underwater features such as seamounts have an important influence on the target's motion.

Once search takes place, certain areas will become more or less likely to contain the target depending upon the results achieved. Both negative and positive search results modify the shape of the posterior target-location probability distribution.

Realistic treatment of some search problems requires a state space of high dimensions. At the minimum, target position and velocity are required in most cases. In addition, target detection characteristics, start times associated with motion, survival times, and other random variables add to the dimensionality of the state space. This "curse of dimensionality" makes purely analytic approaches to filtering and prediction difficult, but does not significantly add to the difficulty of particle-filtering approaches.

1.2. Non-Gaussian structure of the search problem

From the above, one can see that the search problem will often have non-Gaussian characteristics even beginning with the initial target-location distribution. The search for the U. S. nuclear submarine Scorpion lost in 1968 involved consideration of non-Gaussian target-location probability distributions as described by Richardson and Stone[4]. Even in cases where the initial target-location probability distribution is Gaussian, target motion will generally change it into a form which may defy explicit mathematical description.

Until the target is found, search provides what is often called "negative information." This is information conveyed by the fact that the target has not yet been detected. It decreases the target-location probability in the areas searched and raises it in areas that have not been searched.

If contact is made with the target, it may be associated with a sensor that gives non-Gaussian position information, such as a line of bearing with a wedge shaped region of uncertainty.

Figures 1 and 2 illustrate the type of non-Gaussian target-location probability distributions that can be produced by updating for negative and, respectively, positive information. In Figure 1, there is a passive acoustic sensor located at $(0,0)$. The acoustic conditions produce convergence zones (regions of high detection probability) at roughly 30 and 60 nautical miles from the sensor. In this case the sensor has failed to detect the target, so the target is unlikely to be located in regions of high detection probability and more likely to be in regions of low detection probability. The resulting target location distribution is circularly symmetric about the location of the sensor at $(0,0)$. Figure 1 shows the distribution in only one quadrant. Figure 2, shows the posterior target location distribution resulting from a line of bearing detection on a passive acoustic sensor located at $(70,0)$. In this case the target is likely to be located in the regions of good detection probability (i.e., in the convergence zones). Within those regions the target is likely be near the bearing of the detection. These two effects produce the multiple modes in Figure 2. As one can see the distributions in Figures 1 and 2 are highly non-Gaussian and even would be difficult to represent by a combination of Gaussian distributions.

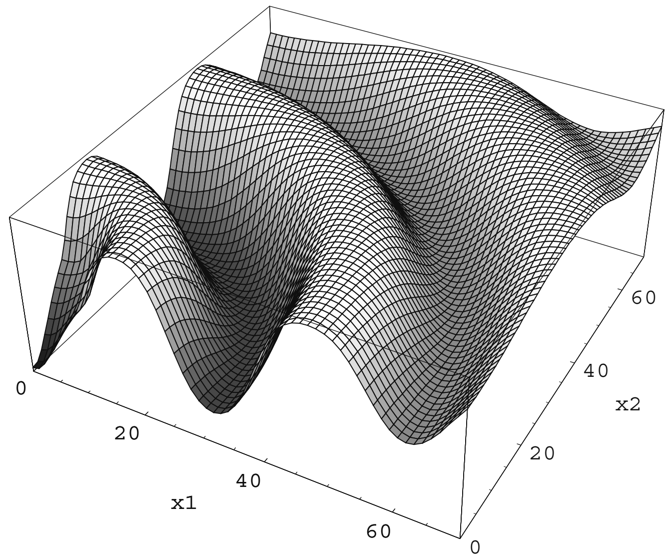

Figure 1: Negative information – posterior distribution resulting from failure to detect.

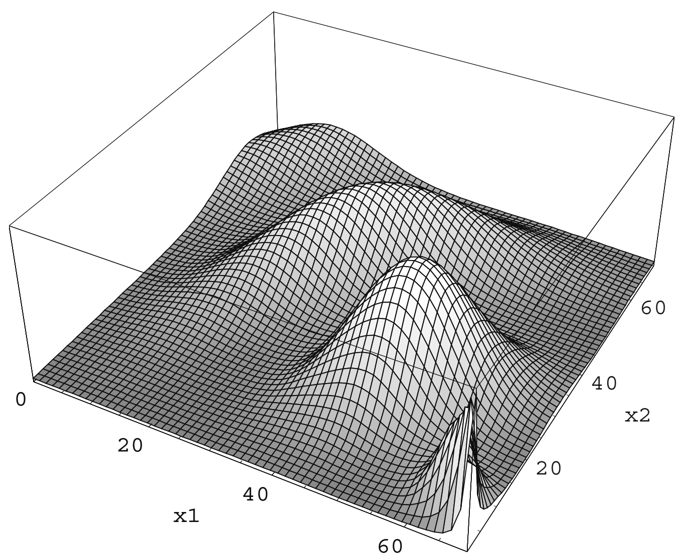

Figure 2: Positive information – line of bearing detection for a passive sensor at (70,0)

2. BAYESIAN MATHEMATICAL FORMULATION

The particle filters implemented in the operational systems described in Section 4 followed the basic Bayesian mathematical formulation given below.

Bayesian filtering is based on the mathematical theory of probabilistic filtering described by Jazwinski[5]. Bayesian filtering is the application of Bayesian inference to the problem of tracking a single target. For the applications in this paper, the target moves in continuous time, but the observations are received at discrete (possibly random) times. This is called continuous-discrete filtering by Jazwinski.

2.1. Target state space

Let S be the state space of the target. Typically, the target state will be a vector of components with some of these components being position and velocity. There can be additional components such as one that specifies target type, which may determine motion and detection characteristics.

2.2. Prior target state-space stochastic process

Let $X(t)$ be the (unknown) target state at time t. We start the problem at time 0 and are interested in estimating $X(t)$ for $t \geq 0$. The prior information about the target is represented by a stochastic process $\{X(t); t \geq 0\}$. Sample paths of this process correspond to possible target paths through the state space, S.

2.3. Sensor information and likelihood functions

There is a set of sensors that report observations at an ordered, discrete sequence of (possibly random) times. These sensors may be of different types and report different information. The set can include radar, sonar, visual, or other types of sensors. The sensors may report both detection and non-detection information. We assume that we know the probability distribution of each sensor's response conditioned on the value of the target state s. This relationship is captured in the likelihood function for that sensor. The relationship between the sensor response and the target state s may be linear or nonlinear, and the probability distribution representing measurement error may be Gaussian or non-Gaussian.

Suppose that by time t we have obtained observations at the set of times $0 \leq t_1 \leq \ldots \leq t_K \leq t$. To allow for the possibility that we may receive more than one sensor observation at a given time, we let Y_k be the set of sensor observations received at time t_k. Let y_k denote a value of the random variable Y_k. We assume that we can compute the likelihood function

$$L_k(y_k \mid s) = \mathbf{Pr}\{Y_k = y_k \mid X(t_k) = s\} \text{ for } s \in S\,. \tag{1}$$

The computation in equation (1) can account for correlation among sensor responses if that is required.

2.4. Bayesian filtering and prediction

Let $\mathbf{Y}(t) = (Y_1, Y_2, \ldots, Y_K)$ and $\mathbf{y} = (y_1, \ldots, y_K)$. Define

$$q(s_1, \ldots, s_K) = \mathbf{Pr}\left\{X(t_1) = s_1, \ldots, X(t_K) = s_K\right\} \tag{2}$$

to be the prior probability (density) that $\{X(t); t \geq 0\}$ passes through the states $s_1, \ldots, s_K$ at times $t_1, \ldots, t_K$. Let

$$p(t_K, s_K) = \mathbf{Pr}\left\{X(t_K) = s_K \,|\, \mathbf{Y}(t_K) = \mathbf{y}\right\}. \tag{3}$$

The function $p(t_K, \cdot)$ is the posterior distribution on $X(t_K)$ given $\mathbf{Y}(t_K) = \mathbf{y}$. The goal of Bayesian filtering is to compute this posterior distribution. From the point of view of Bayesian filtering, the posterior distribution on target state represents our knowledge of the target state. All estimates of target state derive from this posterior.

If we wish to predict the target's location distribution at some time $t > t_K$, then we must compute

$$p(t, s) = \mathbf{Pr}\left\{X(t) = s \,|\, \mathbf{Y}(t_K) = \mathbf{y}\right\}$$

which is the posterior distribution on $X(t)$ given $\mathbf{Y}(t_K) = \mathbf{y}$. for $t > t_K$.

2.5. Recursive method of computing the posterior

Two additional assumptions permit recursive computation of $p(t_K, s_K)$. First, the stochastic process $\{X(t); t \geq 0\}$ must be Markovian on the state space S. Second, for $i \neq j$ the distribution of $Y(t_i)$ must be independent of $Y(t_j)$ given $(X(t_1) = s_1, \ldots, X(t_K) = s_K)$ so that

$$L(\mathbf{y} \mid s_1,\ldots,s_K) = \prod_{k=1}^{K} L_k(y_k \mid s_k)\,. \tag{4}$$

The assumption in equation (4) means that the sensor responses (or observations) at time t_k depend only on the target state at the time t_k. This is not automatically true. For example, if the target state space is position only and the observation is a velocity measurement, this observation will depend on the target state over some time interval near t_k. The remedy in this case is to add velocity to the target state space. There are other observations, such as failure of a sonar sensor to detect an underwater target over a period of time for which the remedy is not so easy or obvious. This observation may depend on the whole past history of target positions and perhaps velocities too.

Define the transition function

$$q_k(s_k \mid s_{k-1}) = \mathbf{Pr}\{X(t_k) = s_k \,\big|\, X(t_{k-1}) = s_{k-1}\} \text{ for } k \ge 1\,, \tag{5}$$

and let q_0 be the probability (density) function for $X(0)$. By the Markov assumption

$$q(s_1,\ldots,s_K) = \int_S \prod_{k=1}^{K} q_k(s_k \mid s_{k-1})\, q_0(s_0)\, ds_0\ . \tag{6}$$

We can now write the recursion for Bayesian filtering.

Recursion for Bayesian Filtering

Initial Distribution: $$p(t_0,s_0) = q_0(s_0) \text{ for } s_0 \in S \tag{7}$$

For $k \ge 1$ and $s_k \in S$, perform

Motion Update: $$p^-(t_k,s_k) = \int q_k(s_k \mid s_{k-1})\, p(t_{k-1},s_{k-1}) ds_{k-1} \tag{8}$$

Compute Likelihood Function L_k from the observation $Y_k = y_k$ (9)

Information Update: $$p(t_k,s_k) = \frac{1}{C} L_k(y_k \mid s_k)\, p^-(t_k,s_k)\,. \tag{10}$$

Prediction: $$p^-(t,s) = \int q_t(s \mid s_k)\, p(t_k,s_k) ds_k \text{ for } t > t_k \tag{11}$$

The motion update in equation (8) accounts for the transition of the target state from time t_{k-1} to t_k. Transitions can represent not only the physical motion of the target but also changes in other state variables. The information update in equation (10) is accomplished by pointwise multiplication of $p^-(t_k,s_k)$ by the likelihood function $L_k(y_k \mid s_k)$ and division by C to normalize the product to a probability distribution. The prediction step in (11) uses the target motion model represented by the transition function $q_t(s \mid s_k) = \mathbf{Pr}\{X(t) = s \mid X(t_k) = s_k\}$ to project the probability distribution ahead to time *t* beyond the time t_k of the last observation. This prediction step is useful for planning a search allocation for a time $t > t_k$.

Likelihood functions replace and generalize the notion of contacts in Bayesian tracking. Likelihood functions can represent sensor information such as detections, no detections, Gaussian contacts, bearing observations, measured signal-to-noise ratios, and observed frequencies of a signal.

In the special case where the target motion model is Gaussian and the measurements are linear functions of target state with Gaussian measurement errors, the above recursion can be computed using the standard formulas for Kalman Filtering. These assumptions do not hold for the maritime applications described below. For this reason we used a particle filter approach to perform the numerical calculations in equations (7) - (11).

3. PARTICLE FILTER SOLUTION METHODOLOGY

The first subsection provides a general overview of the particle filter solution methodology and the second subsection explains the method of resampling. Since the methodology was highly dependent on the computers available in the early 1970's, two of these host computers are discussed in the final subsection.

3.1. Particle Filter Solution

Doucet et al[6] provide an in depth discussion of particle filters. A particle filter is a sequential, Monte Carlo method of performing the computations in the recursion in (7) - (11). In outline particle filtering is simple. One approximates the continuous time and space stochastic process that models the target's motion through its state space with a finite number N of sample paths drawn from this process. Each path is called a particle. Initially each particle is given equal weight (probability). The paths are computed recursively, so that they are generated up to the time t_k at which the last sensor observation was obtained. The distribution of states of the particles at time $t_0 = 0$ produces a discrete approximation to the initial distribution q_0 in (7).

When the first observation is obtained at time t_1, each particle is advanced to its state at time t_1. This is done recursively using the (stochastic) motion model for the target. This accomplishes the motion update step in (8). The likelihood function for the observation(s) at time t_1 is calculated. This implements (9) in the recursion. Next the weight (probability) of each particle is multiplied by the likelihood function evaluated at the state of the particle. The resulting weights are renormalized to add to 1. This process implements (10) in the recursion.

3.2. Resampling

In particle filtering, another step is required called regeneration or resampling. As time progresses and we receive sensor measurements, the distribution on target state will tend to concentrate on a relatively small number of particles. Those few particles will have high probabilities and the other particles will have very small probabilities. In order to preserve the resolution of the filter, resampling splits the large probability particles into many separate particles and deletes the low probability particles. The result is a more balanced distribution of probability over the particles. This process is accomplished in a manner that maintains, at least roughly, the number of particles at N. There are many resampling methods. Doucet et al[6] provide a description of several of these. The technique[7] used in the systems reviewed in Section 4 was motivated by the "Russian roulette and splitting" method described by Herman Kahn[8]. Kahn attributes both the idea and the name of the method to Stanislaw Ulam and John von Neumann.

For $n = 1,\ldots,N$, let $X_n(t)$ denote the n^{th} Monte Carlo sample for the target state space stochastic process advanced to the time t at which resampling is to occur. Also let w_n denote the likelihood weight associated with the n^{th} sample $X_n(t)$. In this formulation, it is not necessary that the sum of the weights be unity.

The next step is to compute the average number of offspring u_n that will be cloned from the n^{th} sample $X_n(t)$. This number is given by

$$u_n = w_n N / \sum_{j=1}^{N} w_j \ .$$

Since u_n is generally not an integer, let K_n be the integer part of u_n and let R_n be the fractional part of u_n. One now plays Russian roulette by drawing a random number z uniformly distributed between 0 and 1. If $z > R_n$, then one kills off the fractional part of u_n and makes K_n copies of state $X_n(t)$. If $z \le R_n$, then one makes $K_n + 1$ copies of state $X_n(t)$. The sum of the number of copies of all particles will be a random variable with expected value N. Some additional steps are taken to prevent the number of particles from drifting away from N over time. After the particles are resampled, all likelihood weights are set to one in preparation for future updating.

3.3. Computer implementation

The basic concepts for the algorithms and for their software implementation were formulated around 1970. By today's standards, computers were primitive in terms of processing speed, RAM, and disk capacity. Time-sharing systems were used via remote teletype networks. Graphical interfaces were still in the future and visual displays had to be constructed on paper from symbols available on a teletype keyboard.

After a few analytic approaches were attempted, it soon became evident that the best way to deal with the problem was to use what is now called particle filtering. In order to minimize the use of RAM, the state-space samples were stored off-line on disk and read into the central processor one record at a time. The sequential motion and information recursions in (7) - (11) were then carried out and the updated records read back onto disk. Richardson and Corwin[9] provide a further discussion of this procedure.

The following subsections briefly describe two representative computer systems that were used to develop and host the particle filtering algorithms.

The CDC 3300 main frame[†]. The CDC 3300 mainframe computer was used to develop the CASP system for the U. S. Coast Guard and was located at Coast Guard headquarters in Washington, D. C. The CDC 3300 computer was obsolete even by the standards of the time. Its RAM memory consisted of 128K 6-bit bytes for program code and data plus another 128K for additional data. Double-precision calculations required four-byte words and consequently allowance for program code reduced processing memory to fewer than 64K double-precision words.

The Data General Nova 800 minicomputer[‡]**.** The Data General Nova 800 was used to develop the MEDSEARCH system and was located in Submarine Group Eight headquarters in Naples, Italy. The computer had four boards of ferrite core memory that consisted of wires wrapped around iron rings. Each board had 8k of addressable 16 bit words. The memory had to accommodate the operating system, the executing program, the call stack, and all program variables. Integers were 16 bits which limited the range to ±32,000. There was no floating point processor so floating point operations were carried out with software. The disk storage consisted of two "platters," one fixed and one removable. Debugging was carried out by stepping through program instructions following panel lights.

4. EARLY MARITIME APPLICATIONS

We now describe applications of particle filtering developed in the 1970s and early 1980s. All of these systems were used operationally by the Coast Guard or the Navy. The search and rescue planning system CASP is still in use today by the Coast Guard. The section is organized chronologically since each new system used ideas and code from previous developments.

4.1. U. S. Coast Guard CASP[10]

The Computer-Assisted Search Planning (CASP) system was the first particle filtering system in the series developed by the authors and their colleagues. Work began in the summer of 1970 and the basic CASP programs were completed in mid-1971. The programs were turned over to the Coast Guard for testing and evaluation using a CDC 3300 mainframe. Coast Guard personnel modified the programs so that they could be used for operations over the Atlantic and Pacific teletype networks.

The state-space for the particle filter consisted of the two-dimensional position of the target. The computer record for each particle consisted of the latitude and longitude of the target together with its current likelihood weight.

The user could model the prior information about target location in three different ways. When a position report was available, a bivariate normal distribution could be generated for location of the initial distress event. The user could then displace each target particle by a random distance and direction. Using the "exponential mapping" from differential geometry, the bivariate normal distribution was "wrapped" around the Earth to eliminate possible distortions that could be introduced when the probability distributions covered large regions.

Prior information about target location could also modeled by a uniform distribution within a user-specified polygonal region. This was often used when a fisherman was lost within a known fishing area. Finally, in certain cases, such as when a float plan or flight plan was available, a target location probability distribution could be generated about

---

† J. R. Frost contributed information for this section.
‡ M. C. Brennan provided details pertaining to the Nova 800.

the intended track line specified in the plan. If required, all three types of distributions could be combined to treat different "scenarios" concerning the target's loss.

After the initial distress event, winds and currents would cause the target to drift, perhaps for days or weeks. Navy oceanographic data were used to estimate drift and random variables were use to represent uncertainties in the estimated parameter values.

Standard manuals and tables were used off line to compute target probability of detection (POD) in the area covered by the search units. When a period of search was unsuccessful, the particles in the search area had their likelihood weights multiplied by the search failure probability (1-POD). Visual, radar, and other types of sensors could be treated in this way.

Search ended with detection and hence there was no update procedure to account for the positive information provided by target contact. Resampling was not used in the original CASP system.

As an illustration of the versatility of the particle filtering approach, T. L. Corwin[11] developed on short notice an innovative multiple-target version of the CASP system. It was used to estimate the location of 520 cyanide containers lost in the Gulf of Mexico as a result of a collision between a freighter and a tanker in 1973. One of the goals was to predict when and where the containers would come ashore in the future, given that they were not destroyed in the collision or subsequently sunk. Observation data included the location of containers already found.

4.2. U.S. Navy CAST[7]

The development of the Computer-Assisted Search and Tracking (CAST) system was undertaken for the Naval Air Development Center. Work on the core programs began in 1972 and they were used in anti-submarine warfare operations the same year. The software design was a derivative of the U. S. Coast Guard CASP system, with the major differences being the target-motion and sensor models, the inclusion of a tracking feature that could incorporate positive information from contacts, and the use of resampling.

The state-space used for the Coast Guard was enlarged to include target velocity as well as position. The computer record for each particle consisted of five numbers, the two-dimensional position and velocity of the target and the current likelihood weight.

Two options were provided for modeling target location and motion. The first used a bivariate normal distribution to represent the initial position. Subsequent target motion was modeled by using "truncated triangular" probability distributions to represent target course and speed. The truncated triangular probability distributions were defined using upper, lower, and best estimates of the various parameters. The user also specified the ratio of the highest to the lowest probability density. When the ratio was one, the distribution was uniform. When the ratio was very much larger than one, the distribution was nearly triangular.

The second method for modeling target motion consisted of specifying a target track using probability distributions for the target's postulated location at various times measured as offsets from a random start time.

Air anti-submarine search used sonobuoys (passive acoustic listening devices) dropped from an aircraft and aided by fixed underwater surveillance systems. Off-line computers were used to estimate target detection probability for all systems under consideration. If the target was not detected, then the target particle likelihood weights were updated to reflect this negative information in the same manner as the Coast Guard CASP system.

If the submarine was detected, then the target particle likelihood weights were updated to reflect the uncertainty in the position report. Since the operations continued on in a tracking mode, a resampling method (Section 3.2) was used to generate additional points in the high probability areas and reduce the number of points in the low probability areas.

4.3. U. S. Navy MEDSEARCH

The development of MEDSEARCH ("Mediterranean Search") began in 1976 at the request of Commander, Submarine Group Eight with headquarters located in Naples, Italy. The particle filtering algorithms built upon previous work on CASP and CAST. The targets of interest were submarines transiting and patrolling the Mediterranean Sea.

Modifications of the particle filtering algorithms were required to adapt to the confined geography of the Mediterranean Sea, with its many islands and restricted passages. In order to address these issues, a network of target transit lanes was established, and motion scenarios were formulated in order to describe various ways that the target might move through the network. Intelligence data were used to assign prior probability weights to each of the target motion scenarios

The MEDSEARCH state space was enlarged to seven dimensions in order to accommodate these new complexities. Three more components were added to the four position and velocity components. The first of these was a component that specified the target transit scenario being followed. The second was a component that specified the average transit speed of advance, and the third new component specified the time that the target entered the Mediterranean and began its transit through the network.

As with the other search systems, observations provided positive location information when the target was detected and negative search information when the target was not detected. New features were added to incorporate negative and positive observations associated with search "barriers" located in various restricted passages.

4.4. U. S. Navy PACSEARCH

PACSEARCH was a particle filter developed from 1984-1987 and used in the Pacific. Its purpose was to evaluate and optimize anti-Submarine warfare operations against Soviet nuclear submarines with Third Fleet resources. These resources consisted primarily of the *Sound Surveillance System* and the *Surveillance Towed Array Sensor System* along with air, submarine, and surface ship search assets.

Although based on MEDSEARCH (and a related system for optimizing the use of anti-submarine warfare patrol aircraft known as VPCAS) PACSEARCH incorporated a number of significant innovations, facilitated by the availability of an early Unix desktop workstation, the Hewlett Packard 9020, with one megabyte of RAM and a fifty-five megabyte hard drive. Some of these innovations included the use of (1) what is now known as sampling importance-resampling, in which N unequally weighted particles are mapped into a new set of N equally weighted samples with very similar statistical characteristics, (2) very detailed non-homogeneous data to describe sensor effectiveness, and (3) complex target motion models with hundreds of possible target motion scenarios.

A critical source of information for PACSEARCH was acoustic transmission loss and directional ambient noise data provided by the *System for Prediction of Acoustic Response of Sensors* installed at Commander, Oceanographic System Pacific (COSP). These data were used by fleet personnel and operations research analysts working on-site at COSP and Third Fleet to generate accurate sensor effectiveness estimates. This capability was essential for processing negative information and optimizing search resources.

Fleet personnel and operations research analysts working on-site at COSP and Third Fleet also generated target motion models utilizing automated PACSEARCH tools. These target motion models, together with all available positive and negative information, were then utilized by PACSEARCH to predict the estimated locations of targets of interest. Given these estimates of current and future target location, PACSEARCH was used to determine the optimal placement and orientation of towed array systems and to optimize the allocation of surveillance processing resources. This significantly increased the operational effectiveness of Pacific ASW search and surveillance operations.

5. ACKNOWLEDGEMENTS

The work described in this paper was carried out while the authors were with the firm Daniel H. Wagner, Associates. We wish to acknowledge the stimulation provided by our many colleagues at this firm during the period covered in this paper. In particular, we wish to thank T. L Corwin for his many substantial contributions to computer-assisted search (CAS) development, E. P. Loane for his suggestion to use Russian Roulette and Splitting as a resampling method, and M. C. Brennan for his tireless development of CAS software which was often undertaken in the field using computers with very limited capabilities.

REFERENCES

1. L. D. Stone, *Theory of Optimal Search*, Second Edition, INFORMS, Linthicum, MD, 1989

2. L. D. Stone, C. A. Barlow, and T. L. Corwin, *Bayesian Multiple Target Tracking*, Artech House, Norwood, MA, 1999
3. J. R. Weisinger, "Influence of Cell Size and Number of Monte Carlo Replications on Search Effectiveness," *Daniel H. Wagner, Associates Memorandum to Office of Naval Research (Code 230)*, 1978
4. H. R. Richardson and L. D. Stone, "Operations Analysis During the Underwater Search for Scorpion," *Naval Research Logistics Quarterly*, Vol. 18, No. 2, pp.141-157, 1971.
5. A. H. Jazwinski, *Stochastic Processes and Filtering Theory*, Academic Press, New York., 1970
6. A. Doucet, N. de Freitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice*, Springer, New York, 2001.
7. H. R. Richardson, "The Processing of Positive and Negative Information obtained from Multiple Detection Systems with Application to the Computer-Assisted ASW Search and Tracking (CAST) System," *Daniel H. Wagner, Associates Memorandum to Naval Air Development Center,* 1973.
8. Herman Kahn, Use of *Different Monte Carlo Sampling Techniques*, The Rand Corporation (P-766), 1955.
9. H. R. Richardson and T. L. Corwin, "An Overview of Computer Assisted Search Information Processing," *Search Theory and Applications,* edited by K. Brian Haley and Lawrence D. Stone, NATO Conference Series, Series II: Systems Science, Volume 8, pp.25-43, Plenum Press, New York, 1980.
10. H. R. Richardson and J. H. Discenza, "The United States Coast Guard Computer-Assisted Search Planning System (CASP)," *Naval Research Logistics Quarterly*, Vol. 27, pp.141-157, 1980.
11. T. L. Corwin, H. R. Richardson, and B. D. Wenocur, "Final Report on the Gulf of Mexico Cyanide Barrel Spill," *Daniel H. Wagner, Associates Memorandum to Commandant (DST), United States Coast Guard*, 1974.