Spatial Correlation of Detections for Impulsive Echo Ranging Sonar

Andrew C. Coon

Impulsive Echo Ranging systems are air-deployed, active, multistatic sonar systems designed to operate in both littoral and open-ocean areas. Sources ensonify a target, and multiple receivers detect resulting short-duration echo signals. Because the source illuminates ocean bottom features, particularly in shallow littoral waters, and the ocean is rich in independently generated transients, the system requires automation dedicated to reducing operator load and enhancing target recognition. The APL-designed Multiple Return Association and Localization (MRAL) algorithm serves as part of this automation. MRAL is an efficient data-fusion algorithm that groups sets of spatially consistent energy detections, collected across pings and receivers, into single operator alerts. Operator classification decisions based on grouped detections are less frequent and more informed than those otherwise made for each detection independently. Also, combined detections provide the means for automation to harness the joint information inherent in the individual detections.
(Keywords: Data fusion, Sonar, Sonobuoy.)

INTRODUCTION

Impulsive Echo Ranging (IER) systems are the Navy's newest generation of air-deployed active sonar systems. The envelope of operations for these systems includes shallow littoral areas where APL's contribution in the design of a spatial correlation algorithm, referred to as the Multiple Return Association and Localization (MRAL), provides the means to reduce operator loading and enhance target recognition. The IER systems were born out of a need to detect increasingly quiet submarines emerging from the Soviet Union in the late 1970s and 1980s.[1] Before that time, passive sonar systems successfully relied on the target to provide detectable signals originating from the submarine's main propulsion system, auxiliary systems, and hotel systems such as air-conditioning. Active sonar systems emerged with increasing importance to ensonify the quiet submarine to provide a detectable target echo. However, after the collapse of the Soviet Union in the early 1990s and the initial retreat of Soviet submarines from the open ocean, the new threat became quiet, third-world, diesel-electric submarines in littoral areas. Quiet diesel-electric submarines in shallow coastal waters typically challenge active sonar systems because of the increased reverberation levels caused by energy returning from the ocean bottom. The increase in reverberation reduces the probability of detection and, in many circumstances, increases the false alarm rate.

BACKGROUND

As Fig. 1 indicates, during an IER operation, an aircraft deploys a sensor field by dropping a number of sonobuoy receivers and sources. The aircraft then sequentially commands each source to sink to transmit an impulsive signal. For a given transmission (or ping), the aircraft has the opportunity to detect a target when sound energy propagates to the target and a subset of sonobuoys (perhaps only one, if any), dependent on channel characteristics and the target's scattering strength pattern, senses the reflected target echo. All acoustic data (i.e., acoustic pressure transformed to electrical signals in the sonobuoy) are transmitted via a radio link to the aircraft for processing and operator evaluation. After each ping, sonar operators are tasked to identify possible target echoes. This is a critical aspect of the mission because chasing a false alarm is costly in a time-critical mission. Unfortunately, not only does the source acoustically illuminate the target, but features among broader bottom topology can provide target-like echoes as well. Equally detrimental are independently generated transients such as whale whistles or system-induced electrical glitches. Whereas the system attempts to automatically eliminate signals that do not behave like target echoes, enough target-like signals pass the screening algorithms to require further aural and visual evaluation by the operators.

Figure 2 illustrates how the IER system combats false alarms. Using signal feature attributes, an automated parametric pattern recognition algorithm screens non-target-like signals flagged by the energy detector.[2] All signals that pass the screening algorithm, referred to as "returns," feed the MRAL data-fusion algorithm. The MRAL algorithm operates on returns collected across multiple sonobuoys, accumulated across multiple pings, and uses arrival times and bearings to spatially correlate the returns. The set of associated returns, referred to

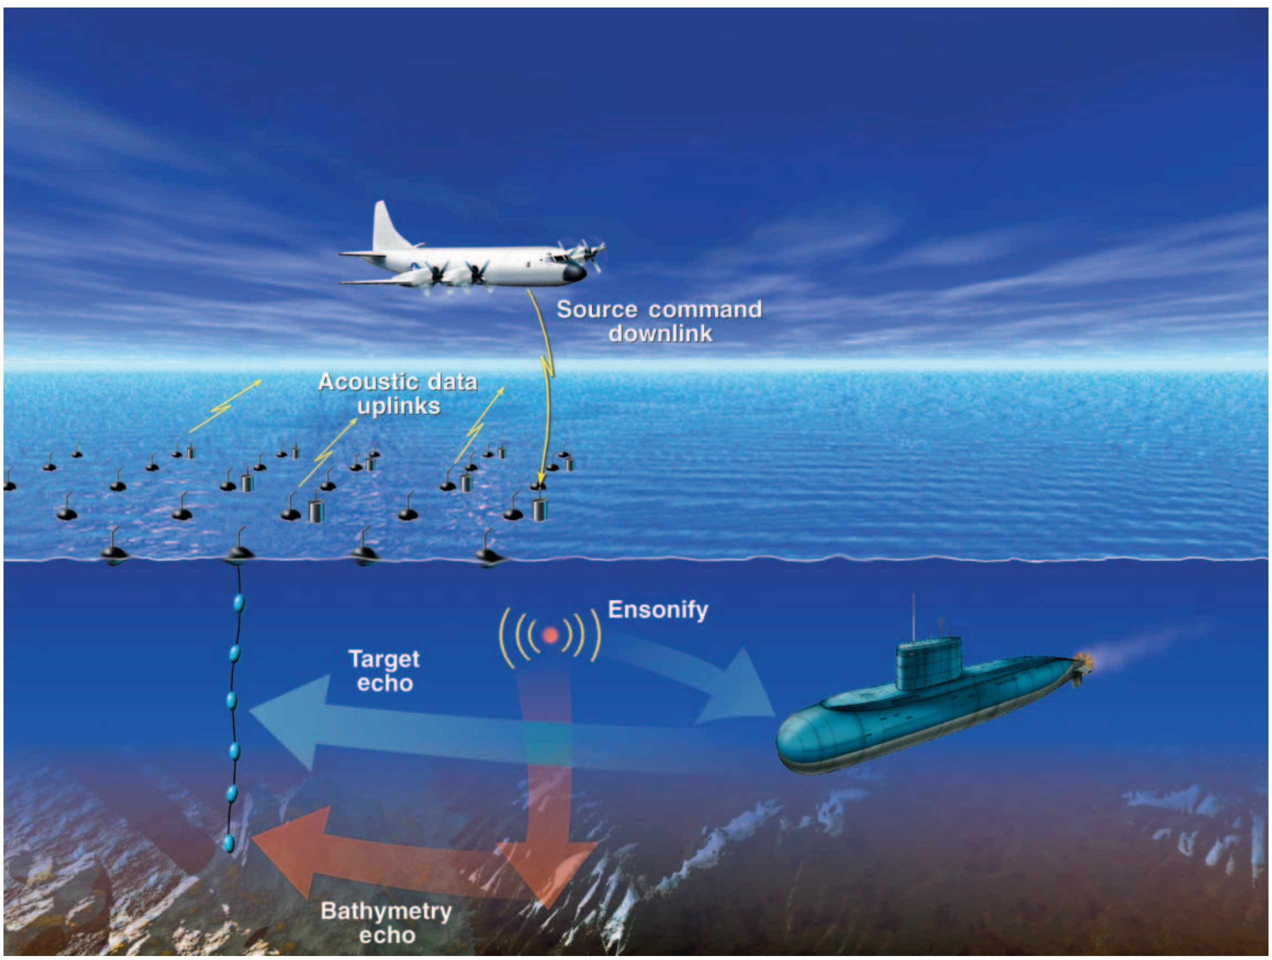

Figure 1. Impulsive Echo Ranging systems are air-deployed, active, multistatic sonar systems. The aircraft monitors multiple sensors to detect a submarine echo.

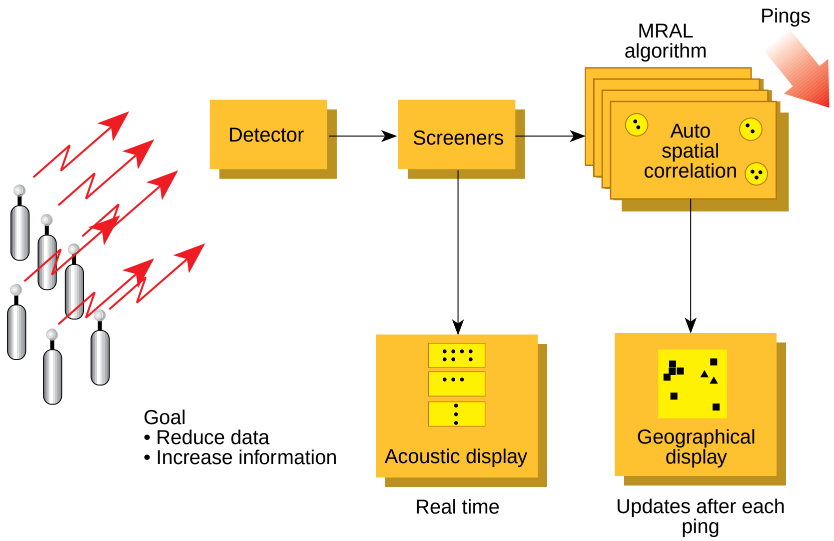

Figure 2. Spatial data correlation applies after application of a classifier that screens individual detections.

The MRAL approach is preferable to clustering the bearing ellipse intersections defined by each return. Inherent in each bearing-ellipse intersection is the error in both the bearing and time, which manifests itself differently as a function of the source-to-target-to-sonobuoy geometry. Because of this, a fixed proximity tolerance for associating bearing ellipse intersections is not ideal. A proximity tolerance variable with geometry, however, adds complexity to the algorithm. The time and bearing uncertainties of a reflector's return are independent of the reflector's position. Avoiding the nonlinear transformation of time and bearing to a Cartesian coordinate allows MRAL to operate with association tolerances ΔT and Δb, which relate to independent errors in time and bearing. For example, with the MRAL approach, perfect time measurements enable a set of returns to associate with a point nearest the reflector position, regardless of the bearing errors, given that sufficiently large Δb is used. In fact, if we set $\Delta b = 360°$, the algorithm can function without any bearing information at all. The argument is easily reversed for perfect bearing measurements and poor or no time measurements. The implication is that dominant errors in either the bearing or time measurements do not drive the overall error in associating a set of returns.

as a "contact," acts as a single operator alert. MRAL reduces an operator's load by mapping multiple nontarget returns (clutter), each of which would otherwise require independent evaluation, into single alerts. The number of contacts generated by the algorithm, by definition, must always be less than or equal to the number of single-return alerts that feed the algorithm.

Beyond this simple concept of data reduction, the MRAL algorithm adds a dimension to classification. Whereas MRAL is not a classifier, the additional information it generates, through grouping spatially consistent returns, provides measurable attributes unavailable from any single return. When harnessed by an operator or a computer, the multireturn information may bolster the system's ability to prioritize the search and classify the target.

THE MRAL ALGORITHM

Association Criteria

A set of returns associates to form a contact when it has consistent bistatic range times and bearings. Figure 3 illustrates this concept. The bistatic range time is the travel time of sound from the source to the target to the sonobuoy. (With the source and sonobuoy positions acting as focal points, the bistatic range is a time-of-arrival ellipse that defines all possible reflector points). A group of returns associates with a point if all the returns define 1) bistatic range times within a bistatic range tolerance ΔT of the bistatic range time implied by the point, and 2) bearings within a bearing tolerance Δb of the bearing to the point. This two-parameter criterion allows for separate treatment of independent errors in bearing and time.

Algorithm Design

The MRAL input includes each return's detection time, bearing, bearing error (standard deviation based on signal-to-noise ratio), score from the screening algorithm, detecting sonobuoy, and detecting ping (all ordered by an index k). The algorithm input also includes the positions of the sonobuoys, indexed by sonobuoy i, and the sources, indexed by ping p.

As Fig. 4 illustrates, the MRAL algorithm first generates sample points of the module search area: (X_m, Y_n), $m = 1$, M, $n = 1$, N. A predefined sample resolution, δ, and the size of the search area determine M and N. The algorithm presumes all reflector positions to be quantized to one of the sample points. Ideally, a reflector between sample points is identified with its nearest sample point. The sample resolution is set as small as possible without computationally overloading the computer. The MRAL algorithm proposes the points as reflector positions and, for each point, searches for returns that are consistent in both time and bearing. An alternative approach of hypothesizing combinations of returns and assessing consistency with

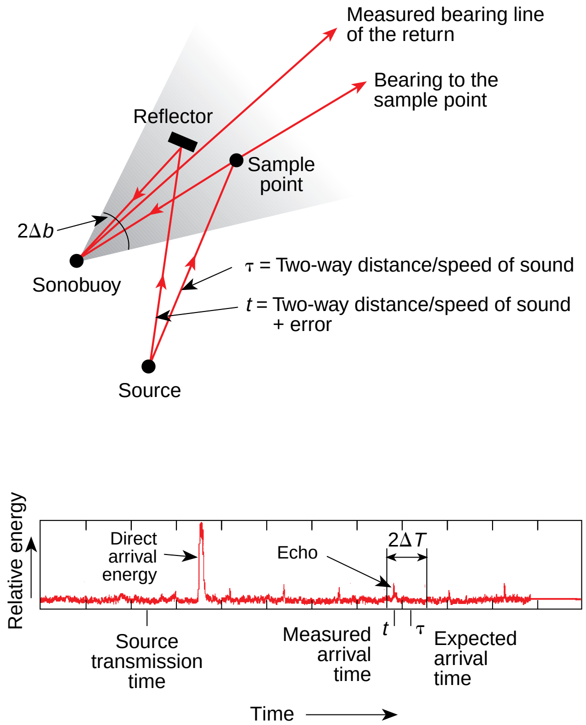

Figure 3. Illustration of the concept that for a return to associate with a point, it must carry a bearing and time consistent with those of the sample point. Here, the measured bearing is within Δb of the bearing from the sonobuoy to the point, and the measured arrival time is within ΔT of the time it takes sound to travel from the source to the point to the sonobuoy.

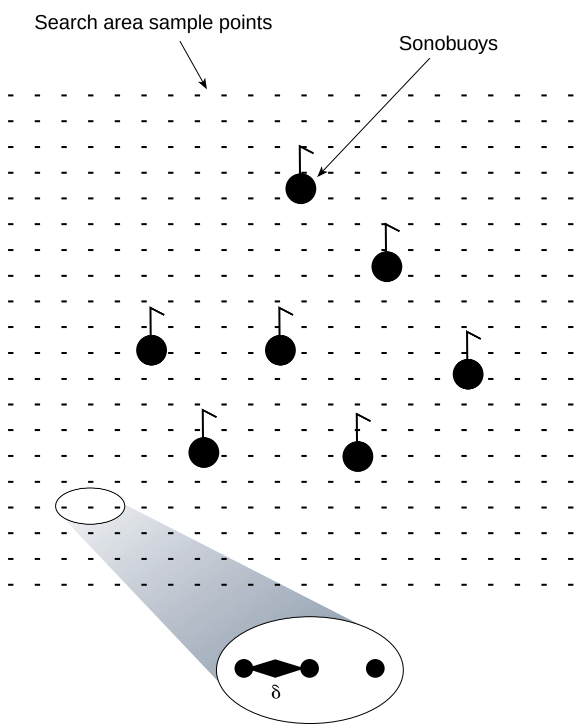

Figure 4. The search area is quantized into sample points evenly spaced by δ.

the residual error of their implied localization point proved less efficient and more complex.

Before the search for consistent returns begins, two steps take place to make the search more efficient. First, many of the sample points can never associate with any returns. To prevent association testing for *NM* points, the MRAL algorithm applies a simple prescreening test to tag each point with an "on" or "off" status. The default status is "off"; however, if a point falls within a box with diagonal corners, defined by the endpoints of an elliptical arc, subtended by the bearing tolerance for a return (expanded to the nearest sample points), then the status is switched to "on." Figure 5 illustrates how the box of sample points is defined. Only points with an "on" status are considered for association testing.

The second step before the search involves reorganizing the input to be accessible by ping p, ping's sonobuoy q, and time bin b. (The ping's sonobuoy q is an index that identifies one of the sonobuoys monitored for a ping. The same sonobuoy can be monitored on different pings and therefore may use a different index q depending on the ping.) To understand how this reorganization improves efficiency, we first recognize that there are at most $P \cdot B$ unique source-sonobuoy pairs (given P pings consisting of B monitored sonobuoys per ping), which implies that there are at most $P \cdot B$ unique expected bistatic range times for a given sample point. To test if any returns from a specific source-sonobuoy pair are consistent with a point, the algorithm only needs to compute a single expected arrival time. The expected arrival time for a given point (m, n) and a given source–sonobuoy pair (p, q) needs to be compared to all the returns from the source–sonobuoy pair. The majority of these returns, however, are far outside the expected arrival time

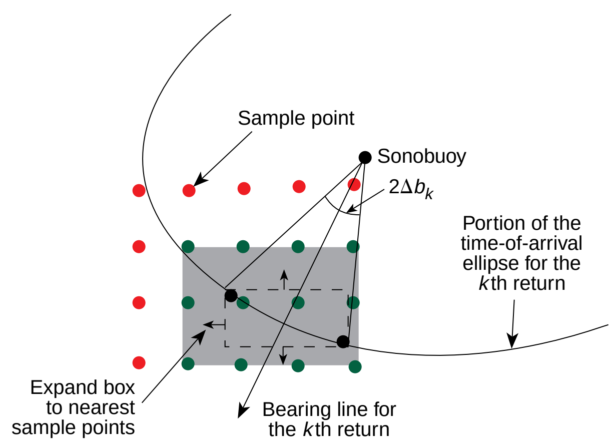

Figure 5. MRAL flags points in and on the solid box (shown in green) as "on." After all returns are tested, points still "off" cannot be consistent with any return and are excluded from association testing.

window implied by the point. If we pare down the testing to include only returns that are close to the bistatic range time, we can eliminate even more redundant calculations. The solution is to map returns from a particular source–sonobuoy pair into time bins of width $2\Delta T$. Efficiency is increased by testing only those returns within the time bin that contains the predicted arrival time τ. Because we do not know in advance where within a bin τ will fall, all returns are also mapped to their nearest bin. The fact that the bins have a width of $2\Delta T$ guarantees that all returns not mapped into the bin containing τ will be farther than ΔT from τ.

After the algorithm applies sample point screening and reorganizes the returns, the contact search proceeds by cycling through the sample points. For each sample point tagged with an "on" status, each source-sonobuoy pair is considered in turn. For each source-sonobuoy pair and sample point, the algorithm defines an expected bistatic range time to the point with

$$\tau = (rs_p + rb_i)/c \ , \tag{1}$$

where rs_p is the source-to-point range, rb_i is the sonobuoy-to-point range, c is the speed of sound in water, p is ping, and i is sonobuoy. As Fig. 6 illustrates, for a given point and source–sonobuoy pair, the algorithm determines a bin index,

$$b = \mathrm{INT}(\tau/2\Delta T) + 1 \ , \tag{2}$$

and cycles through all returns within the bin. (The INT function truncates its argument.) The algorithm compares each return time t to the bistatic range time τ. The return is consistent in bistatic range time if

$$\Gamma_t = |t - \tau| < \Delta T \ . \tag{3}$$

Subsequent bearing consistency checks are conditioned on successful cases of bistatic range time consistency and are executed less frequently.

When a return is consistent in bistatic range time, the algorithm then determines if the return is also consistent in bearing. This process proceeds in two steps. First, a coarse bearing test computes which of eight possible 45° sectors contains the sonobuoy-to-point bearing, and checks if the return's measured bearing is within the sector or just outside the sector by no more than Δb_k. The bearing tolerance is indexed by k since the bearing accuracy depends on a return's signal-to-noise ratio, which is different for each return. An eight-element look-up table provides the sector containing the sample point's (x, y) position. The table is accessed by the three-bit value,

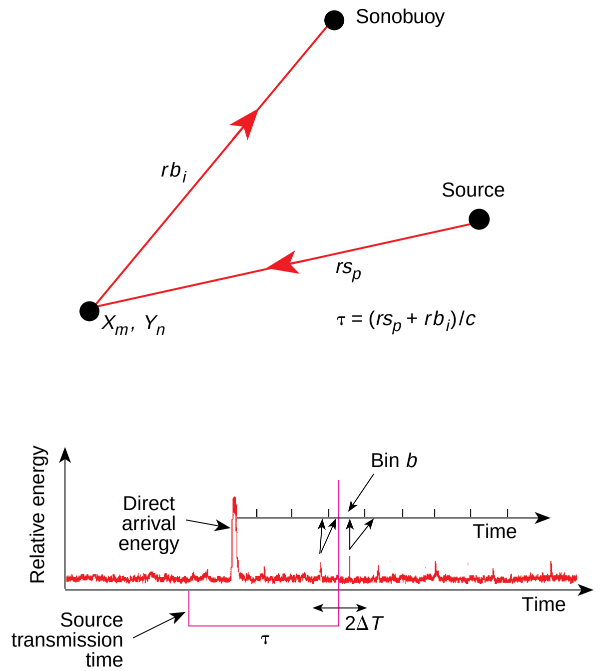

Figure 6. Two returns mapped to bin b are within ΔT of τ, where τ is the expected travel time of sound energy from the source to the sample point to the sonobuoy. If each return is consistent in time, subsequent testing will verify if either is consistent in bearing.

$$\begin{aligned} bit_1 &= \{x \cdot y > 0\} \ , \\ bit_2 &= \{x > 0\} \ , \\ bit_3 &= \{|x| > |y|\} \ , \end{aligned} \tag{4}$$

which quickly determines the 45° sector containing the sample point. As an example, if $(bit_1, bit_2, bit_3) = (1, 1, 1)$, the sample point is in the lower 45° sector of the first quadrant. This first step finds the majority of returns to be inconsistent in bearing.

For those returns roughly consistent in bearing, a second bearing test computes the actual bearing to the sample point via an arctangent function (a time-consuming task for the aircraft's tactical computer). If the return's bearing is within Δb_k of the actual bearing to the point, the return is considered to be consistent in bearing. If we define Γ_b as the angular difference between the sample point bearing and the return's bearing, then bearing consistency occurs when

$$\Gamma_b < \Delta b_k \ . \tag{5}$$

Because the return is also consistent in range (otherwise the bearing test is never executed), the return is consistent with the sample point.

For each return found consistent with the sample point, the MRAL algorithm temporarily stores the return's index k. After all source–sonobuoy pairs are considered, all of the returns temporarily stored for the point under consideration constitute a "candidate contact." Stored with the candidate contact are its size (number of returns), position (sample point indices m and n), and consistency parameter (some Γ_t or Γ_b). For a multireturn contact, there are multiple measures of Γ_t—one for each return. The return least consistent with the point (with the greatest Γ_t) defines the consistency parameter for a multireturn contact. For single-return contacts, Γ_b defines the consistency parameter.

The candidate contact is stored in a contact list if it is not a subset of any currently stored contact. For example, a candidate contact with returns $k \in \{7, 19, 97\}$ is a subset of a stored contact with returns $k \in \{7, 19, 22, 97\}$ and is not added to the contact list. Conversely, if the algorithm adds a candidate contact to the list, any previously stored contacts that are subsets of the new contact are removed. The example in Fig. 7 illustrates a geometric interpretation of the logic. The figure shows the segments of elliptical arcs subtended by the bearing tolerance for each return. Each arc segment is labeled with its return index k. The presumption for Fig. 7 is that returns $k = 7$, $k = 19$, and $k = 97$ are consistent in time and bearing to points 1, 2, and 3, but that they are only mapped to point 2. This is because the return $k = 22$ is only consistent with point 2, and the contact defined with $k = \{7, 19, 97, 22\}$ is a superset of candidate contacts found for points 1 and 3. From examining Fig. 7, it is not surprising that the algorithm chooses point 2 because it visually appears to best approximate the center of all the ellipse intersections.

The positions of candidate contacts that are subsets of stored contacts are typically located near the contact that finally gets stored. The MRAL algorithm scans the sample points in an order that takes advantage of this fact to rapidly find a stored superset of a candidate contact and move on to the next point more quickly. Same-sized candidate contacts compete via their consistency parameters to allow the candidate contact most consistent with a sample point to survive. Recall that the consistency parameter is Γ_t for multireturn contacts and is Γ_b for single-return contacts. The term Γ_t geometrically represents the tightness of the time-of-arrival ellipse intersections and Γ_b geometrically represents the closeness of a single return's bearing to a sample point.

Although no contact in the list of contacts is a subset of any other, some may intersect, i.e., share returns. After the algorithm tests all sample points in the search area, the contact list is modified a final time to eliminate any intersections. Figure 8 shows how a single return might exist in more than one contact. As shown in the figure, return $k = 22$ may have originated from a reflector near point 2 or near point 4. The algorithm forces a choice and puts the return into the contact with the most returns. By defining the size of a contact by its number of returns, algorithmically, the scheme starts with the largest contacts (then second largest, then third largest, etc.) and removes returns shared by any smaller contacts. The choice is random for equal-sized contacts.

The forced-choice logic reduces the total number of contacts as smaller contacts lose their returns to larger ones. Also, as an operator selects different contacts from a geographical display, a return is never repeated. Simulation results indicate that the forced-choice logic also reduces the probability of false association. A false

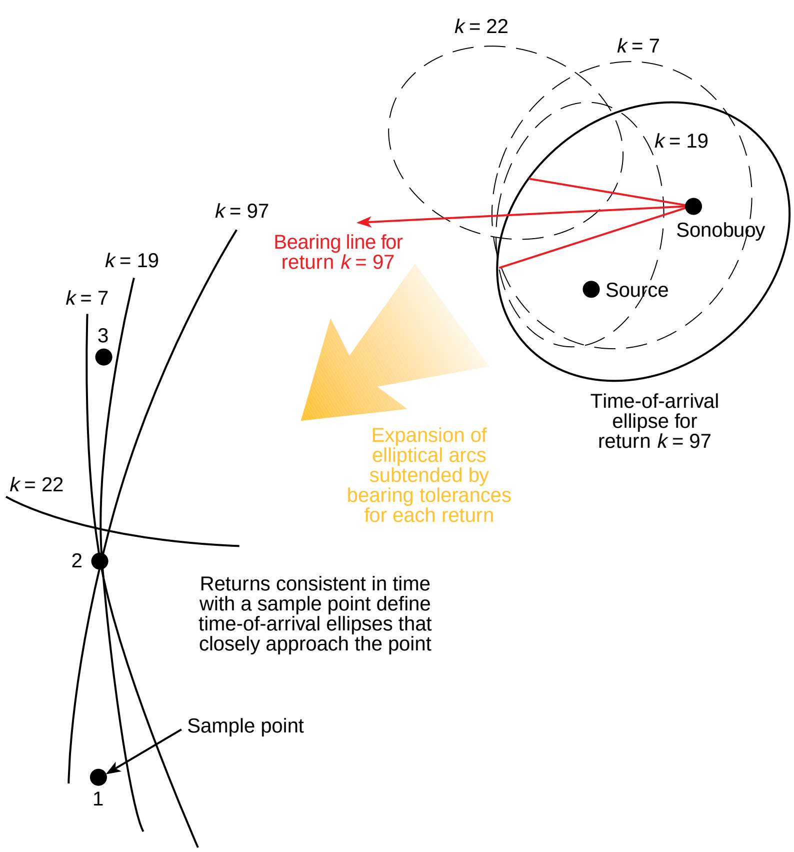

Figure 7. MRAL finds all three returns labeled $k = 7$, $k = 19$, and $k = 97$ consistent with points 1, 2, and 3. Candidate contacts formed for points 2 and 3 are dropped in favor of the contact at point 2, which is also consistent with return $k = 22$.

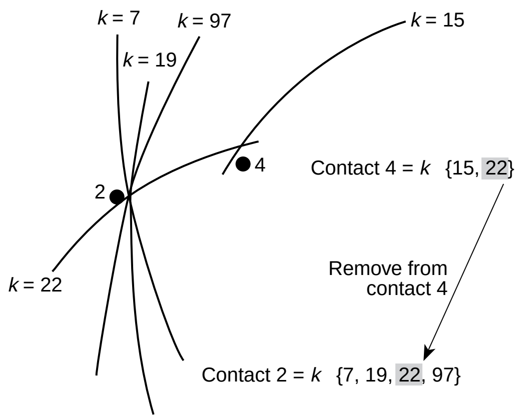

Figure 8. Given the return $k = 15$ added to the situation depicted in Fig. 7, the candidate contact at point 4 is not a subset of the contact at point 2 and gets added to the contact list. After the initial search for contacts is complete, returns in multiple contacts are forced into the contact that can hold the most returns. Return 22 is therefore forced to reside in contact 2.

association occurs when a return from one reflector associates with returns from another reflector. The probability can be parameterized by the spacing between the two reflectors, and the forced-choice logic improves performance at all spacings.

Automatic Target Recognition

As suggested earlier, the joint information in a set of spatially consistent returns provides measurable clues for target recognition, which are unavailable from any single return. When harnessed by an operator or a computer, the multireturn information may bolster the system's ability to prioritize the search and classify the target. For the operator, a target-like detection that is difficult to classify alone may more readily be dismissed if it aurally and visually correlates with a number of more distinct clutter events. In addition, contact position information on a geographical display reveals existing clutter patches. This information enables operators to prioritize their searches.

For the computer, a number of measurable multireturn features exist that an automated classifier might use to screen or rank returns. Examples of these features include the following:

- Median score (MRAL score)
- Greatest time (ping) separation between returns in a contact
- Number of returns in a contact
- Local contact density

The MRAL score is the only multireturn feature currently used. The algorithm that screens energy detections prior to MRAL ranks each return with a score. The MRAL algorithm is therefore able to rank a contact with the median score of all returns within the contact. This ranking is referred to as the MRAL score.

Much like an operator might cope better with a target-like clutter return when examined with other spatially correlated returns, the median likelihood measure assists the computer by harnessing multiple confidence measures. Other multireturn feature measures of potential value capitalize on the fact that, unlike many bottom features, the submarine is a moving, point-like reflector with a specific scattering pattern. Bottom reflectors do not move and detections may persist over many pings. If a contact has returns from pings spaced widely in time, the measured time separation may provide a useful classification clue. In addition, the scattering strength of a bottom feature, unlike that for a target, might produce a large number of returns, even over a few pings. Finally, if clutter sources exist in clutter patches, the MRAL algorithm is likely to generate areas dense in contacts, and a local contact density measure may serve as a useful clue. The field testing of the MRAL algorithm conducted off the coast of New Jersey is described in the boxed insert.

FUTURE WORK

The MRAL algorithm is incomplete without specification of the bistatic range time tolerance ΔT and the bearing tolerance Δb_k. Current tasking involves simulation analysis dedicated to establishing an operating point for ΔT and Δb_k with $\Delta b_k = \alpha\sigma_k$, where α is a scale factor and σ_k is the standard deviation of the bearing error for a return. Choosing a ΔT and an α that are too large increases the probability of associating returns from different reflectors and lowers the resolution of the system. Choosing a ΔT and an α that are too small increases the probability of not associating returns originating from the same reflector. The operating point that best balances the probability of false association with the probability of missed association depends on the sample point resolution and the system errors.

In terms of enhanced target recognition, APL plans to research and develop automated classification algorithms that exploit MRAL-enabled multireturn feature measures. After the collected data are better understood, multidetection feature measurements may add a new dimension to IEER automated classification currently restricted to environmentally dependent signal waveform clues.

REFERENCES

[1] Tyler, G. D., "The Emergence of Low-Frequency Active Acoustics as a Critical Antisubmarine Warfare Technology," *Johns Hopkins APL Tech. Dig.* **13**(1), 145–159 (1992).
[2] Shin, F., Kil, D., and Wayland, R., "IER Clutter Reduction in Shallow Water," in *Proc. of ICASSP*, Vol. 6, Atlanta, GA (May 1996).

ACKNOWLEDGMENTS: The author would like to express appreciation for the support and dedication from PMA-264's B. Spiotta, S. Schreiner, and R. Loyer, along with the Naval Air Warfare Center Aircraft Division, PMA-264's subcontractors. The author would also like to thank Julius I. Bowen, Ronald W. Mitnick, and Alan D. Ravitz for their guidance in putting this article together.

SAMPLE MRAL ALGORITHM TEST RESULTS

Sample results from a test conducted in the Hudson Canyon off the coast of New Jersey are shown in part (a) of the figure. The target is the ex-USS *Salmon* (SS 573) moored to the ocean floor. The geographical display is for research and development rather than operator use. On the display, contacts are numbered, and the color coding indicates the size of the contact. Returns collected across five pings from three closely spaced sonobuoys provide the input to MRAL. Contact 30 is the target contact. The target contact is defined with a large number of returns because the geometry of the test was geared to produce a large number of target detections. However, the value of MRAL is linked more to the association of clutter events.

The results indicate that a significant fraction of the clutter returns map to a few distinct points. A return's arrival time provides ranging information that geometrically determines a time-of-arrival ellipse with the source and sonobuoy positions acting as the focal points. MRAL combines the arrival times with bearing measurements to map returns to one of a finite number of equally spaced points in the search area. Part (b) of the figure shows the time-of-arrival ellipses and bearings for Contact 30 (target) and Contact 25 (clutter) overlaid on the MRAL sample points from a portion of the search area.

The MRAL algorithm is sufficiently general to be incorporated by any active multisensor sonar system. Measures of bearing and time are treated independently to naturally map the measurements into a single Cartesian reference system.

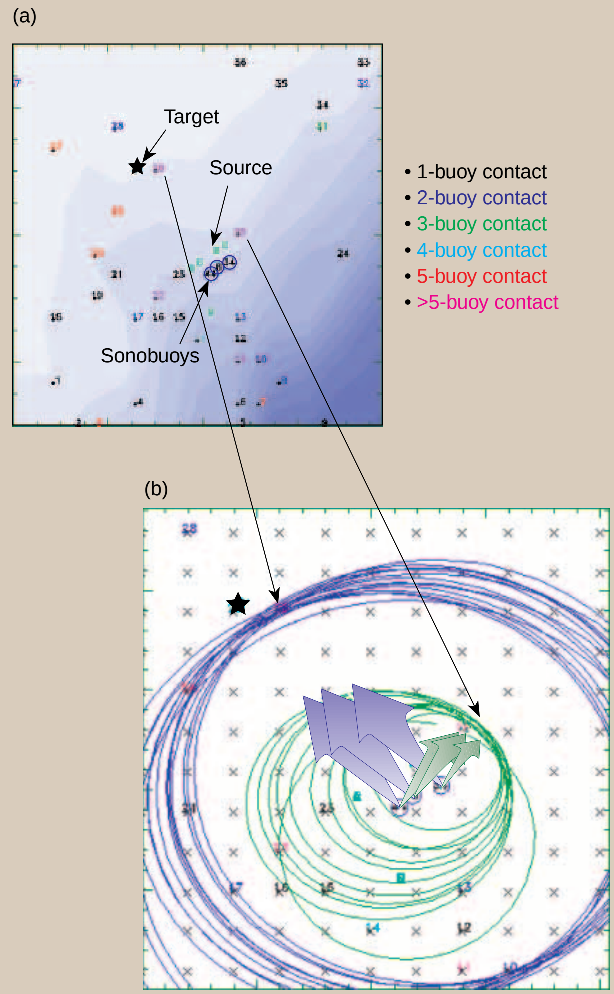

Sample MRAL test results from the Hudson Canyon: (a) Contact positions generated from five pings of data collected on three sonobuoys. Contacts are color coded to indicate the number of returns in the contact and are labeled with identification numbers. (b) Time-of-arrival ellipses and notional bearings for returns associated with Contact 30 (target) and Contact 25 (clutter).

THE AUTHOR

ANDREW C. COON received a B.S. degree in electrical engineering from the University of Maryland in 1987 and an M.S. degree in electrical engineering from the JHU G.W.C. Whiting School of Engineering in 1992. In 1988, he joined APL's Submarine Technology Department and currently works in the Signal and Information Processing Group. Mr. Coon's work has focused on Impulsive Echo Ranging systems and specifically has addressed multistatic data fusion and signal processing techniques to improve system performance. Mr. Coon also developed a desktop training system designed to train signal classification to a user population of over 200 Navy sonar operators. His more recent internal research and development work has focused on development of small-aperture processing techniques for battlefield acoustics including sniper detection and localization. Mr. Coon is a member of APL's Senior Staff. His e-mail address is Andy.Coon@jhuapl.edu.