

A Structured Approach to Passive Sonar Track Association

Garfield R. Mellema, *Member, IEEE*

Abstract— The presence of multiple, apparently independent track segments originating from the same target complicates the track-level passive sonar picture. If some of these segments can be shown to have common origin, they can be associated into a single composite track, simplifying that picture. The process could also provide additional information about the target, such as range or classification. Track association is typically a manual process, relying on the expertise of a human sonar operator. If a method could be found to reliably apply numerical scores to the degree of apparent relationship between pairs of tracks, those scores could be used to assist an operator in the track association process or as inputs to a fully automated track association process. This paper outlines the construction of a test based on the sample correlation coefficient and describes the result of its application to tracks produced by a probabilistic data association filter (PDAF) from a set of towed array sensor data.

Index Terms—Passive sonar, sonar, tracking, track association.

I. INTRODUCTION

Passive sonar can be used to detect, track, localize and identify targets by their own emissions, making it an excellent tool for covert surveillance. The task of surveillance is complicated, however, by the diversity of signals available in the underwater environment, not all of which can be easily distinguished. Additionally, the targets of greatest interest are often those whose emissions are weakest and which are most easily masked by emissions from targets that could be more easily observed by other means [1].

In a typical surveillance scenario the objective is to extract as much information as possible about the surrounding environment, especially about those elements that pose the greatest potential threat. At the lowest level, much of this extraction can be automated through the use of detection and tracking algorithms to convert maps of acoustic intensity into track segments indicating bearing and frequency over time.

At this point the track segment information could be plotted on a chart relative to the receiver to indicate the relative bearing of the source over time. It is of limited value, however, as it is referenced to the receiver, which might or might not be in motion itself.

If a relative track segment meets certain criteria, including

sufficient duration, Target Motion Analysis (TMA) may be used to estimate the position, source and speed of the source. The result is a geo-referenced track that can be plotted independently on the chart and fused with information from other sensors. The identity of the target represented by the track is also more easily estimated [2].

Not all relative track segments, however, can be refined into geo-referenced tracks, typically due to insufficient track duration. A common example is a signal with a very low signal-to-noise ratio (SNR) which produces an intermittent series of track segments. In this case the sonar data cannot be easily advanced from beyond the track-level display onto the chart display and its interpretation is greatly complicated.

A human operator observing the track level display might identify by eye multiple track segments that appear to be related by a common origin and may choose to associate these segments into a single, composite track. In the case of an intermittent series of segments along a single bearing and frequency line, the result of associating the segments might be a composite track of sufficient duration for TMA, in which case the information represented by the relative composite track could be represented as a geo-referenced track. Other types of association are also possible with similar results [3].

When track segments from the track level display are advanced to the chart level display, much of the information they continue to provide on the track level display, such as the presence of an acoustic source, either becomes redundant or diminishes in value due to its duplication on the chart display. Either removing or diminishing the intensity of these segments in the track level display would emphasize the presence and characteristics of the remaining segments.

Since the value of the sonar information increases significantly as it is refined, it is highly desirable to refine relative track data into geo-referenced track data. The remaining track segments in the track-level display would then primarily represent either noise, or the sort of very weak or intermittent signals that might be produced by a distant or intentionally covert source.

While there are automated tools for the refinement of passive sonar data to the level of track segments, the task of identifying apparently related track segments is left to the human operator. Automating this process would reduce the workload of the human operator and increase the capabilities of autonomous sonar signal processing systems.

II. THE PASSIVE SONAR PROBLEM

A. Acoustic Signals in the Underwater Environment

A typical vessel in the marine environment sheds acoustic energy primarily from the operation of its propulsion, housekeeping and other mechanical equipment. These sources may be dispersed throughout the vessel and may not all be synchronized. The result is a flow of signals into the surrounding waterspace in all directions and usually at multiple frequencies.

The signals observed at an underwater receiver are often assumed to have followed a near-direct propagation path from the source. Signals can also arrive along other paths with one or more reflections from the ocean surface or bottom. The longer propagation paths delay the arrival of the signal at the receiver, resulting in constructive and destructive interference if the multiple arriving signals cannot be differentiated.

A directional receiver can be used to differentiate among the signals arriving at different bearings. The towed array receiver, a linear phased array, can be very effective at this task but, due to the conical nature of its radiation pattern, is only able to distinguish the arrival angle of signals relative to its own axis. If the common assumption of purely horizontal arrivals is made, signals that were reflected one or more times, and therefore have a significant vertical component in their arrival angle, appear to arrive closer to broadside.

B. The Sonar Data Refinement Process

In a typical towed array passive sonar system, the received data is refined in a number of distinct stages. As the data is advanced through these stages its information content is increased, although not all data at each level is suitable for advancement [4].

In the first stage, sound pressure levels in the water are converted in to electrical signals and digitized by a series of hydrophone elements in the array. In the second stage, the digitized signals are phased-shifted and summed to steer the directionality of the array along a series of beams. The number of beams, distinct directions in which the array is steered, is determined by the desired minimum sensitivity of the array and the width of each beam is determined by the projection of the array towards the signal source. The result of beamforming is a series of 2-dimensional beammaps of received intensity versus frequency and time on a beam-by-beam basis. Continuity of the received signals in bearing and/or frequency has not been established at this point, although it may be apparent to the eye.

In the third stage, detectors are used to identify sequences of features in the beammaps and tracking algorithms, typically Kalman-type filters, are used to identify and record these as tracks. Each track is a history of the state bearing, frequency and intensity of a signal represented by a time series of the state vectors of the filter. The continuity of each record, however, is broken whenever contact with the signal is lost and so a single intermittent signal may produce a series of independent track segments, although the continuity of these segments may be apparent to the eye. This problem is

especially common in scenarios with a low SNR.

Those signals that were not detected and tracked at this stage continue to be available for analysis as intensities but will not be eligible for any further processing or refinement at higher levels. Conversely, spurious tracks, such as those produced from noise or interference will be eligible for further processing. It is important, therefore, that the detection and tracking process be both sensitive and selective.

Some of the track segments developed in this stage may be due to signals that are harmonics of each other, or which may have originated from the same source but followed different propagation paths. Similarities in the structure of these signals, especially during concurrent portions of the tracks, may also be apparent to the eye.

In the fourth stage, the tracks produced by the detection and tracking process are analyzed and grouped with respect to the source and/or platform from which they are believed to have originated. This is beneficial in several ways. First, the fusing of track segments across time into composite tracks extends the duration of the available tracks. Second, the association of tracks from the same source arriving at differing bearings and therefore along distinct paths can be used to estimate the range of the source by triangulation. Third, since sources and platforms can often be identified by the set of frequencies at which they radiate acoustic energy, the association of track segments across frequency provides a means for target identification.

In the final stage, TMA can be applied to refine those tracks having suitable characteristics into estimates of target position, course and speed. By assuming that the target is following a fixed course at constant speed, an estimate of the target's range can be made by analyzing the bearing rate of the target's track. Observations made before and after a change in the observer's course can be used to cross-fix the target. The accuracy of the estimates improves with the duration of the track. If multiple tracks at different frequencies can be associated with the same target, an estimate of identity can also be made. This highly refined, geo-referenced format greatly facilitates the exchange of target information with other interested parties.

Much of the data originally acquired by the towed array has been left behind by the final stage, primarily due to a lack of sufficient corroborative evidence to advance it to the next level of refinement. Within that data, however, is information about targets that may be useful in the detection of weak or distant targets or to more quickly and accurately localize and/or identify targets. Due to the cascading effect of any lower level enhancements, improvements can be made at all levels to increase the quantity and reliability of the refined underwater picture.

III. PRODUCTION OF RELATED TRACKS

An operator observing the underwater picture at the track level must treat each track segment as if it represented a unique and independent target even though, as we saw in the previous section, it is quite normal for a single platform to produce multiple track segments. If two or more segments

could be reliably identified as sharing a common origin, then they could be associated, or grouped, into a single master track, reducing the number of potentially independent targets. The additional track information would also make the master track easier to follow and would enable additional analysis options. The challenge lies in reliably identifying those related track segments in spite of insufficient knowledge of the sources and the underwater environment.

A. Harmonic Pairs and Multiple Sources

This discussion will address a number of cases that merit consideration for track association. The most obvious case is that of a pair of track segments representing a pair of harmonic signals from a single source that have propagated along a common path to the receiver. In this case it would be reasonable to expect that:

- 1) The track segments would have a similar waveform in bearing due to their common origin and similar propagation path.
- 2) The track segments would have a similar waveform in frequency due to their harmonic nature. Differences due to Doppler shift could be eliminated by normalizing the tracks in frequency.
- 3) The waveforms of the track segments in bearing and frequency would be coincident in time due to their common origin and propagation path.
- 4) The SNRs of the track segments are dependent on the background noise at different frequencies and therefore cannot reliably be expected to have similar characteristics.

The case of multiple sources co-located on a common platform is similar to that of the single harmonic source with the exception that the waveforms cannot be expected to be identical in frequency. Those waveforms will, however, share artifacts due to the motion of their common platform and their common propagation path and in many cases, especially where the source is a relatively pure tone, this may be the dominant source of variations in their waveform.

B. Intermittent Signals

The case of an intermittent signal is especially difficult. Consider a pair of track segments that represent the same signal before and after the signal faded below the threshold for tracking. Straightforward extrapolation of the former segment had already been attempted by the tracker and failed, hence the track has been terminated. Extrapolation may still be merited however, as the association process could expect that the shape of the waveform at the end of the earlier segment would not differ significantly from the shape of the waveform at the start of the later segment.

It is not uncommon for a single target to produce multiple signals at differing frequencies and that may provide a more viable approach to this case. If the intermittent signal is not the only one from this target, there may be shared characteristics among the signals which can be used to associate both the earlier and later track segments with a third track segment that bridges the pair. This too is a problem of extrapolation since the evidence in favour of associating the

track segments may be less than perfect.

C. Track Seduction and Noise

In the case of track seduction, a tracker following one signal is inadvertently drawn onto another nearby, typically stronger signal with the result that multiple trackers are then following the same signal. Similar situations can occur when a tracker is either initiated on noise, or distracted temporarily by noise while the signal it was following is redetected. In this case, the resulting track segments can be expected to become increasingly similar over time. The expectations for this case are similar to that of the harmonic signals case with the primary difference that their frequencies are identical.

D. Multipath Propagation

A pair of track segments may represent a pair of signals which had propagated from the same source along different paths to the receiver. A common scenario includes one or more bottom and/or surface reflections, each of which increases the length of the propagation path. In this case the tracks can still be expected to have relatively similar waveforms in frequency, although their Doppler shift would differ with their apparent closing rates. There would also be a time delay between the two corresponding to the difference in their propagation path. The waveform characteristics may also be affected by the bottom and surface reflection coefficients.

The most significant difference between the pair of track segments is likely to be in their bearing waveforms. Since the difference between the two propagation paths is usually entirely contained in the vertical component of their paths, the effect is to cause the signal following the longer path to arrive at an angle closer to normal on the towed array receiver. While the relationship between the two bearings tracks could be determined from, among other things, knowledge of the relative position of the source and receiver in the environment, that information is not normally available at the time the tracks are produced. On the other hand, that information could be estimated from the bearing waveforms if the pair of tracks is assumed to share a common origin.

IV. IDENTIFICATION OF RELATED TRACKS

A. Testing for Association

The previous section described several ways in which two or more track segments could originate from the same target. In each of these cases there are particular characteristics that both track segments could be expected to share. Those characteristics are not necessarily exclusive to that pair, but the presence of multiple common characteristics in a pair of tracks can provide significant evidence for or against their association.

The simplest characteristic to test for is an appropriate temporal relationship. In the case of multiple sources or harmonics, and in the case of track seduction, the tracks are required to be concurrent. In the case of multipath propagation the signal arriving closest to broadside would

appear later in most cases.

In those cases where the pair of track segments under examination represents the same emission from a target, the segments can be expected to have similar waveforms in normalized frequency and bearing. This is especially true if the segments are concurrent. In the multipath case, where the signals are not likely to be concurrent, each of the waveforms would show effects due to its path, including reflection coefficients for the frequency waveform and increased path length for the bearing waveform.

B. Testing for Waveform Similarity

A useful tool for evaluating coincident variations in a pair of vectors is Pearson's r , also known as the sample correlation coefficient. It removes the steady state bias from each vector and then normalizes their amplitudes prior to evaluating their similarity. Its calculation is a four step process [5].

1. Given a pair of vectors of length n , calculate the two sample means, \bar{x} and \bar{y} ,

$$\bar{x} = n^{-1} \sum x_i. \quad (1)$$

2. Calculate the two sample variances, s_{xx} and s_{yy} ,

$$s_{xx} = \frac{\sum (x_i - \bar{x})^2}{n-1} = \frac{\sum x_i^2 - n\bar{x}^2}{n-1}. \quad (2)$$

3. Calculate the sample covariance, s_{xy} ,

$$s_{xy} = \frac{\sum (x_i - \bar{x})(y_i - \bar{y})}{n-1} = \frac{\sum x_i y_i - n\bar{x}\bar{y}}{n-1}. \quad (3)$$

4. Calculate the sample correlation coefficient,

$$r = \frac{s_{xy}}{\sqrt{s_{xx}s_{yy}}}. \quad (4)$$

Pearson's r represents the mean of the product of the instantaneous normalized amplitudes. Its absolute value can only increase when there are simultaneous excursions from the mean in both vectors and its value can only increase when those excursions share the same sign. Normally distributed random fluctuations average out over time but the time required for this to occur increases with the severity of the fluctuations. Interestingly, it is the presence, not the lack, of fluctuations in the vectors that makes this test effective.

V. WAVEFORM COMPARISONS

A typical sonar track produced by a probabilistic data association filter (PDAF) is a time series of vectors, each of which represents the state of the PDAF at that point in time. Typical components of the state vector include the frequency, and bearing values of the underlying model as well as their rates and all of their variances. Other components may include the SNR and its variance [6].

A set of tracks were produced using a PDAF on passive sonar data received by a towed array. Of particular interest in this set were a pair of tracks that appeared concurrently at a dissimilar frequencies. Both tracks were terminated by the PDAF and later reappeared following a 90 degree change in heading of the towed array. A comparison of the earliest pair of tracks, labelled tracks 1 and 2 is shown in Fig. 1. A comparison of the later pair of tracks, labelled tracks 3 and 4, is shown in Fig. 2. The track frequencies have been normalized for comparison and only the concurrent portions of the tracks are shown. It should be noted that the pair of frequencies was not related by a simple harmonic ratio.

Visual examination reveals sufficient similarities between the tracks in each pair that a human operator might consider them to have likely had shared a common origin and therefore be associable into a composite track. Consider tracks 3 and 4. It is clear that the envelopes of the bearing and bearing rates are similar and that the envelopes of the frequency and frequency rates are very similar. Confidence for their association is increased by the presence of coincident features such as the sudden increase in frequency at about 8600 seconds following the long period of very small variations in the frequency plot. The peak in the frequency rate at about 8700 seconds also adds confidence as do the similarities in phases of the envelopes throughout the two rate plots. The spike in frequency that appears only in track 4 at about 7200 seconds can be attributed to random noise and therefore ignored. Note that the scale of the frequency and frequency rate plots shows detail that would be indistinguishable in a typical display showing multiple tracks without normalization.

In order to evaluate the effectiveness of the sample correlation coefficient test in the identification of related tracks, the test was applied to the concurrent portions of track 3 and 4. To eliminate effects due to the initiation and termination of the PDAF, an additional 40 seconds buffer (which corresponds to 5 time-steps of the PDAF) was excluded from the beginning and end of the concurrent portions. For convenience, the correlation values are described as $r = [r_b, r_f, r_{br}, r_{fr}]$ for bearing, frequency, bearing rate and frequency rate respectively. The result of the test was $r = [0.719, 0.997, 0.804, 0.872]$, which indicates a high degree of correlation.

Tracks 1 and 2 are also sufficiently similar that a human operator might consider them to have likely shared a common origin and therefore be associable into a composite track. Once again the envelopes of the bearing and the bearing rate are similar and the envelopes of the frequency and the frequency rates are very similar. The ridges in the frequency plot at about 3500 and 4000 seconds align well, as do the larger scale features following them. The general trends of the bearing and bearing rate plots are also somewhat well aligned. The differences between the pair of tracks in this case appear to be mostly in the most rapidly changing components of all of the envelopes, suggesting the presence of significant noise or an interfering signal.

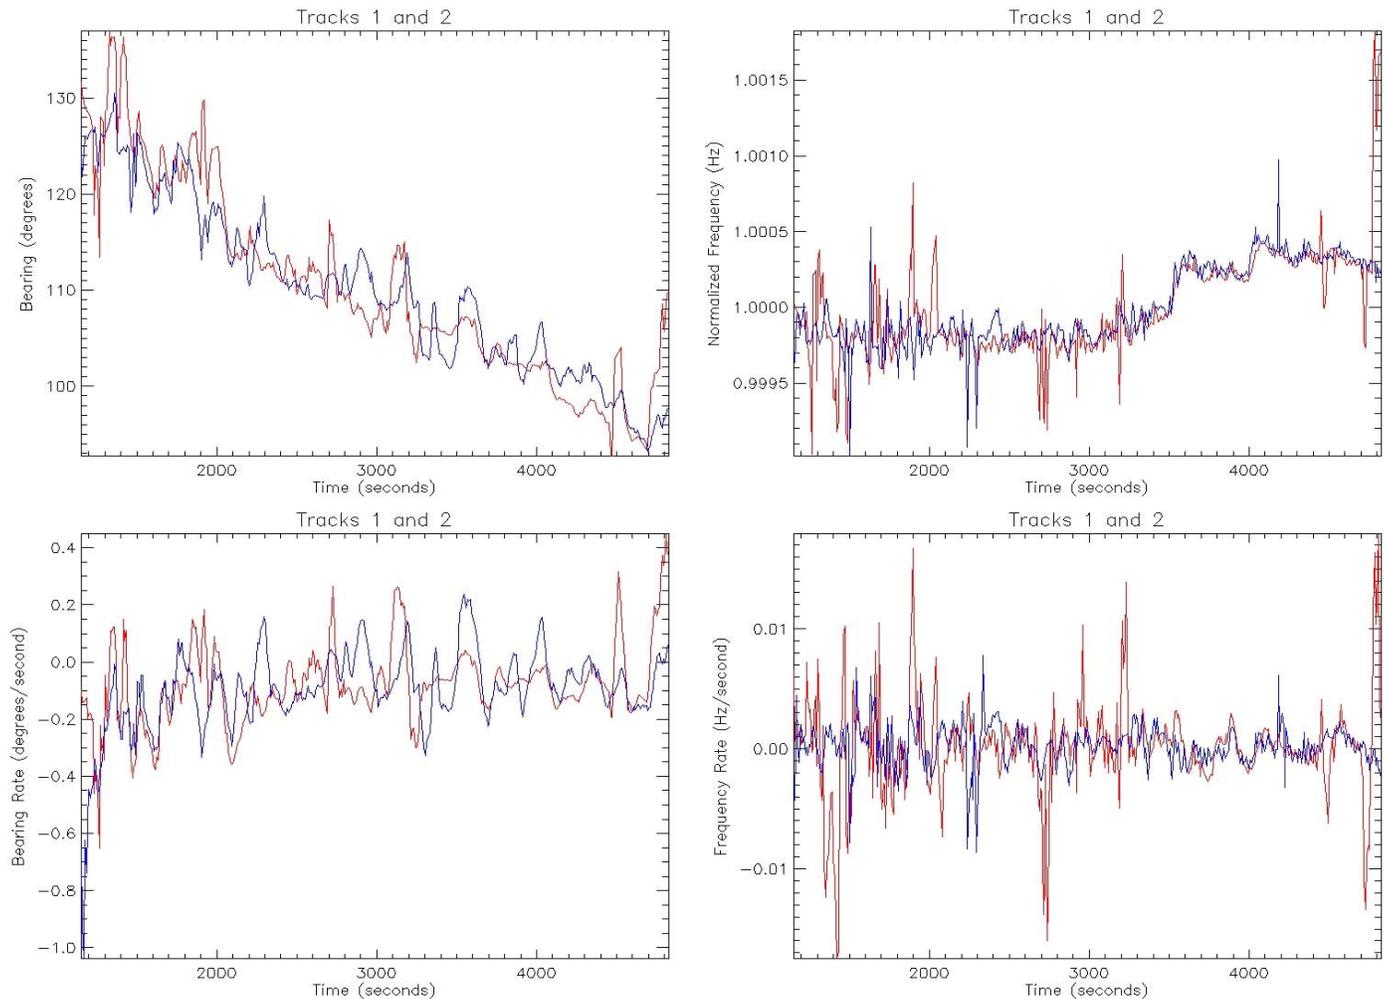


Fig. 1. A comparison of tracks 1 (red) and 2 (blue). Only the coincident portions of the tracks are shown.

When the sample correlation coefficient test was applied to tracks 1 and 2 the resulting correlation values were $r = [0.930, 0.662, 0.383, 0.091]$, which indicates a high degree of correlation in only one of the cases and a medium degree in another. The bearing rate correlation is weak at best and frequency rates could almost be described as effectively uncorrelated. Interestingly, the scale of the differences between the frequency components of the tracks is extremely small, on the order of 0.001 Hz, but these are normalized by the sample correlation coefficient test.

VI. DISCUSSION

The results of the previous section have shown that the sample correlation coefficient can provide numerical values related to the degree of similarity of a pair of sonar tracks. The sensitivity of the test and the range of validity of the results, however, is not yet well defined. The ability to algorithmically identify relationships between pairs of tracks is significant though, since it can be used to automatically build evidence for or against a decision to associate the segments into a composite track.

The choice of 4 correlation values was deliberate, in that the first two, bearing and frequency, were obvious requirements

for the association of signals sharing a common propagation path. The bearing and frequency rates were included as they provide a more sensitive test of the bearing and frequency envelopes. Further tests, such as those involving variances were not found to provide additional useful information. In that light, a useful single-valued correlation score should include the influence of all four sub-scores but not be overly swayed by a single weak score, as might result from an interfering signal. A preferential voting method might be most suitable.

In light of the automated nature of this test, general testing of all pairs of coincident tracks might be used to screen for possible relationships. In this case, more so than others, zero valued correlations would also be useful as a means to identify unrelated tracks.

The case of track seduction is also interesting. In this situation it is quite possible for a pair of tracks to show very good correlation over one part of their coincident duration and very poor correlation over another. A possible solution might be to consider only the latest portion of the tracks. If both tracks are following the same signal at the same bearing and frequency, then at some point their respective trackers should converge and begin producing identical results. Prior to that

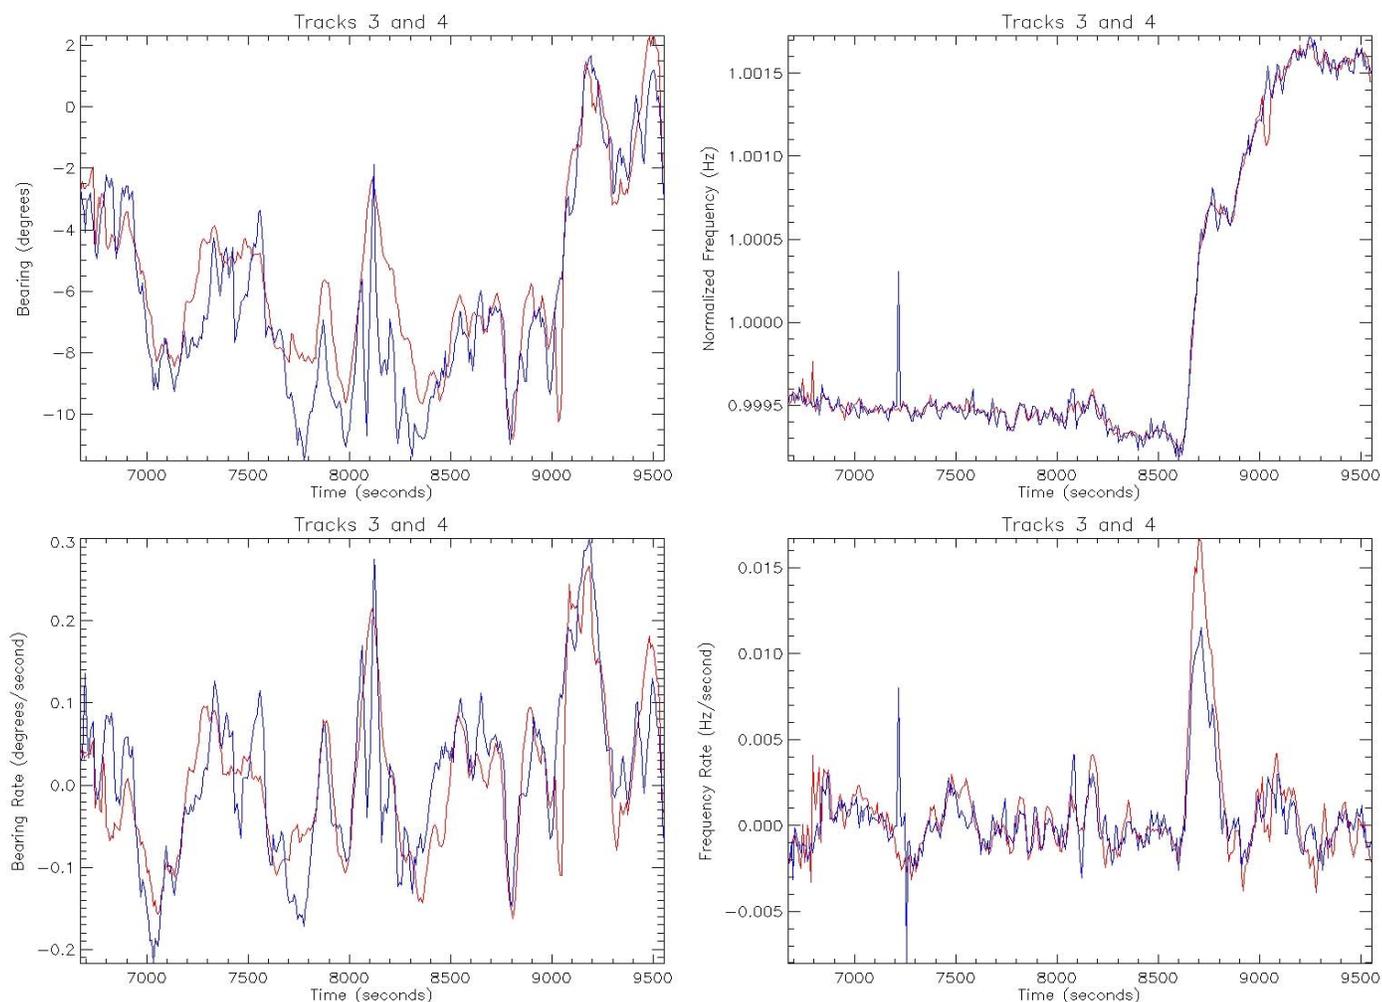


Fig. 2. A comparison of tracks 3 (red) and 4 (blue). Only the coincident portions of the tracks are shown.

time, especially in an environment with a low SNR, it may be possible for one of the trackers to be once again seduced away by noise.

VII. CONCLUSION

Passive sonar track association is a difficult problem. The sample correlation coefficient test is a useful tool to describe the degree of apparent similarity between a pair of track component vectors. The correlation scores of pairs of individual track components, such as bearing, frequency, bearing rate, or frequency rate, are not sufficient to reliably indicate relationships between tracks but the combination of all multiple components appears to be significant.

Correlation scores can be strongly affected if one but not both of the tracks are contaminated by high frequency noise. This could be addressed through improvements in the tracking algorithm, by pre-filtering the tracks or by applying the correlation test in multiple sub-bands of the waveform envelope.

While it is not a turnkey solution for the track association problem, this test can be used to address at least some of the situations described here. That, by itself, is an improvement

from the current situation. As well, the use of this test in combination with one or more other criteria, such as simultaneous initiation or termination, should be able to further reduce the pool of ambiguous track combinations either through the identification of those clearly suitable for association or those clearly unsuitable for association.

VIII. REFERENCES

- [1] R.J. Urick, *Principles of Underwater Sound*, 3rd ed., McGraw-Hill, New York, 1983.
- [2] W.A. Roger and R.S. Walker, *Accurate Estimation of Source Bearings from Line Arrays*, Proceedings of the 13th Biennial Symposium on Communications, Kingston, Canada, June 1986.
- [3] Mellema, G.R., *An Automated Approach to Passive Sonar Track Segment Association*, Proceedings of the 7th International Command and Control Research and Technology Symposium, Québec City, Canada, September 2002.
- [4] R.O. Nielsen, *Sonar Signal Processing*, Artech House, Boston, 1981.
- [5] P.J. Brockwell and R.A. Davis, *Time Series: Theory and Methods* 2nd ed., Springer Verlag, New York, 1996.
- [6] S.S. Blackman, *Multiple-Target Tracking with Radar Application*, Artech House, Boston, 1986.