Correlation Based Testing for
Passive Sonar Picture Rationalization

Garfield R. Mellema

Defence R & D Canada – Atlantic
P.O. Box 1012, 9 Grove Street
Dartmouth, NS, Canada B2Y 3Z7
Garfield.Mellema@drdc-rddc.gc.ca

Abstract - Modern passive sonar systems employ a high degree of automation to produce a track-level sonar picture. Further refinement of the track-level information is normally performed by a human operator. Providing automated assistance would reduce the operator's workload and is a key enabler for semi- and fully automated sonar systems. The nature of the signals emitted by targets and of the underwater environment typically results in each target being represented by multiple track segments. A tool is required which can numerically describe the relationship between pairs of track segments so that those that apparently share a common origin can be identified automatically. The sample correlation coefficient, is a statistical measure of relatedness. This paper describes the application of a test based on that measure to compare tracks produced by a probabilistic data association filter from a set of towed array sonar data.

Keywords: Tracking, passive sonar, towed array, track association.

1   Introduction

In a typical towed array passive sonar system the received data is refined in a number of distinct stages. Many of these stages, such as the detection and following of signals, are either automated, or capable of being automated. One task in particular however, track association, is still heavily reliant on human intervention. [1]

Track association is the collection of groups of tracks, as produced by signal followers, into composite tracks representing the combined features of those multiple track segments. The assembled tracks are believed to share a common attribute, typically a common origin and propagation path from their origin to the receiver [2].

The signals arriving at a receiver are the result of complex interactions among a host of parameters, only some of which are known and many of which can only be estimated. It is unusual therefore in a low signal to noise ratio (SNR) environment, for a received signal or track to be clearly identified, on its own evidence, with a particular

source. This ambiguity, along with the variety of attributes that could be used to build evidence for or against an association decision, makes it difficult to automate the track association task.

As the number of sensors and the processing capabilities of sonar systems increases, one option for dealing with this bottleneck is to discard those track segments that do not meet the criteria for further advancement. A better solution is to investigate methods by which at least some of those tracks could be assembled into a more useful form.

2   The Track-Level Tactical Picture

Tracks are generated from sensor data, typically by following one or more signals on a time-step by time-step basis as they are detected and observed in the environment. The tracks can then be used to represent target information in a tactical picture.

In the passive sonar scenario, tracks are developed from acoustic energy in the underwater environment, typically shed by surface or subsurface vessels, denoted here without prejudice as targets. The acoustic emissions may be intentional or unintentional [3].

If a directional receiver, such as towed array, is used, the tracks can include bearing information relative to the receiver. Due to the lack of travel time information, such as might be found in active sonar scenarios, the range of the target from the receiver cannot be easily determined. This limited format of track data can be assembled into a track-level tactical picture referenced to the receiver.

After sufficient track data of the correct type is acquired, it may be possible to cross-fix a target between two or more tracks, or to apply target motion analysis (TMA) to estimate the position of a target as well as its course and speed. This target localization can then be used to describe the target on a chart display without reference to the receiver. Target localization is a significant improvement to the bearing-level track but it is contingent on the presence of several characteristics, including significant duration, in the track data [4].

Other avenues for track refinement are also available. Tracks representing pairs of signals that originated from the same source but travelled to the receiver along different propagation paths can also be used to estimate the range of a target. In the simplest case, one track

represents the direct propagation path and the other represents a path that includes a single bottom or surface reflection. Knowledge of the bathymetry can then be used to triangulate the range of the target.

Of particular interest is the case where tracks represent pairs of signals that originated from the same target but at different frequencies due to the presence of harmonics or the presence of multiple collocated sources. Although the pair of signals cannot be easily used for localization, they can be used to identify the type or identity of the target. In addition, in a low SNR environment or other situations where one or more of the signals may intermittently disappear and reappear, the tracks could be stitched together to produce a master track of longer duration than any of the component track segments. Assembling composite tracks in this manner would improve continuity in the track-level picture as well as reduce the number of potentially independent track segments.

Tracking in a low SNR environment is also susceptible to the production of spurious tracks due to noise. Typically, these track segments are of very short duration as the tracks are terminated as soon as they lose lock on the fictitious signal. The situation can also cause the production of multiple tracks following the same real signal. This can occur when a track segment that was initiated based on spurious noise is seduced onto a real signal, or when a tracker following a real signal is temporarily distracted by spurious noise and a new tracker initiated to follow the original signal.

The use of reverse-time tracking to augment forward-time tracks can also produce multiple tracks following the same signal, albeit in differing directions and having slightly different track histories and characteristics.

Prior to the application of track association, each track segment in the track-level picture must be assumed to be independent and to represent a potentially independent target. Reducing the number of independent track segments therefore also reduces the number of potentially independent targets and, therefore, the complexity of the track-level picture.

A second benefit in all four cases is the potential to produce a composite track of longer duration than either of the track segments being associated. This is a significant benefit as it will provide greater continuity in following a target and therefore greater opportunity for TMA. Additional associations could also be made to further extend the composite track.

In most of these four cases, identifying and associating pairs of related tracks brings a third benefit as well. In the case of multipath propagation this additional benefit is the potential for triangulation, while in the case of differing frequencies this additional benefit is the potential for target identification.

3 Identification of Related Tracks

The characteristics of an underwater acoustic signal are influenced by many factors including the type,

configuration and loading of source, the source platform, the propagation path and the receiver. Those characteristics that may be fixed, such the source or target type or configuration, can be difficult to measure since only a single realization can be found in each track. Dynamic characteristics, such as variations in the source loading, target motion or propagation path offer multiple state realizations during a typical observation period and can be used effectively for comparisons between potentially related tracks.

All acoustic emissions from a target are influenced to some degree by the motion of that target regardless of their frequency or bearing. Even in the anomalous case where a target is securely immobilized, the lack of target motion is significant. Target motion affects all aspects of a track including amplitude, frequency and bearing, although the ease with which this influence can be discriminated varies significantly depending on the type and degree of the additional influences. This situation is similar to that of communication by FM radio transmission and, not surprisingly, the most easily observed and correlated variations are those in frequency.

The ability to automatically and reliably identify pairs of tracks that originated from a common target would significantly improve the quality and clarity of the underwater picture. Given that pairs of signals originating from the same target share a common influence, e.g. ship motion, is the effect of that common influence sufficiently discernable in the resulting tracks that it can be used to automatically and reliably indicate their common origin?

A human operator looking over a passive sonar track display will pick out pairs or groups of track that are believed to have originated from the same target and identify common characteristics in each pair or group that provide evidence for their association. Most often, these characteristics are based on changes in the bearing or frequency of the tracks, and quite often these changes are either coincident in time or appropriately delayed, according to the difference in bearing between the tracks. Abrupt changes provide especially strong evidence.

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, calculate the two sample means, $\bar{x}$ and $\bar{y}$,

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

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

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

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

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

4. Calculate the sample correlation coefficient,

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

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.

4   Comparison of Tracks

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 at appeared concurrently at a dissimilar frequencies. Both tracks were
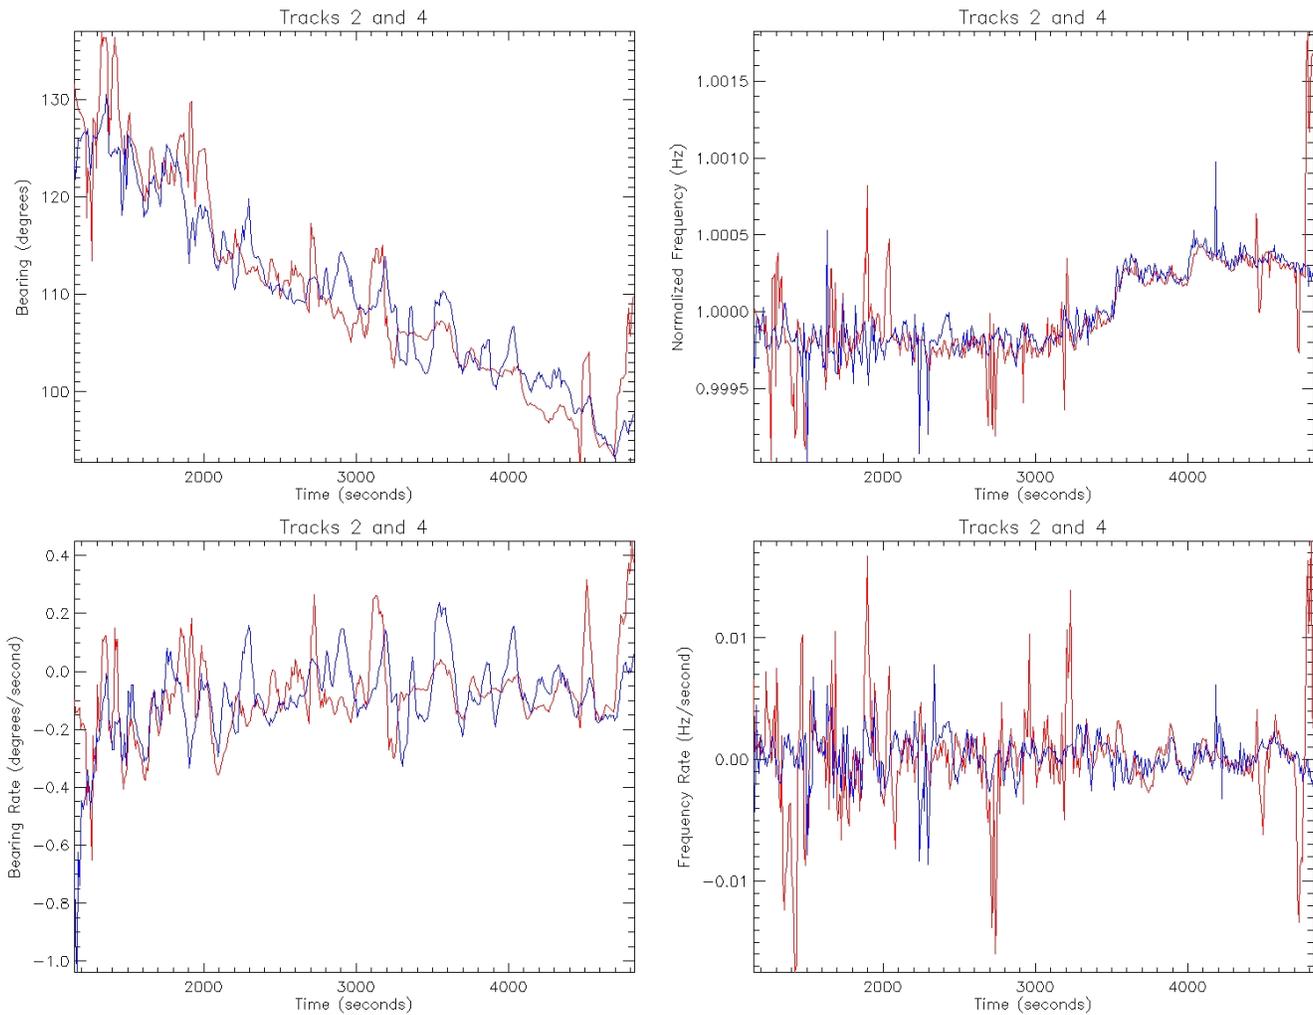

Figure 1 A comparison of tracks 2 (red) and 4 (blue). Only the coincident portions of the tracks are shown.
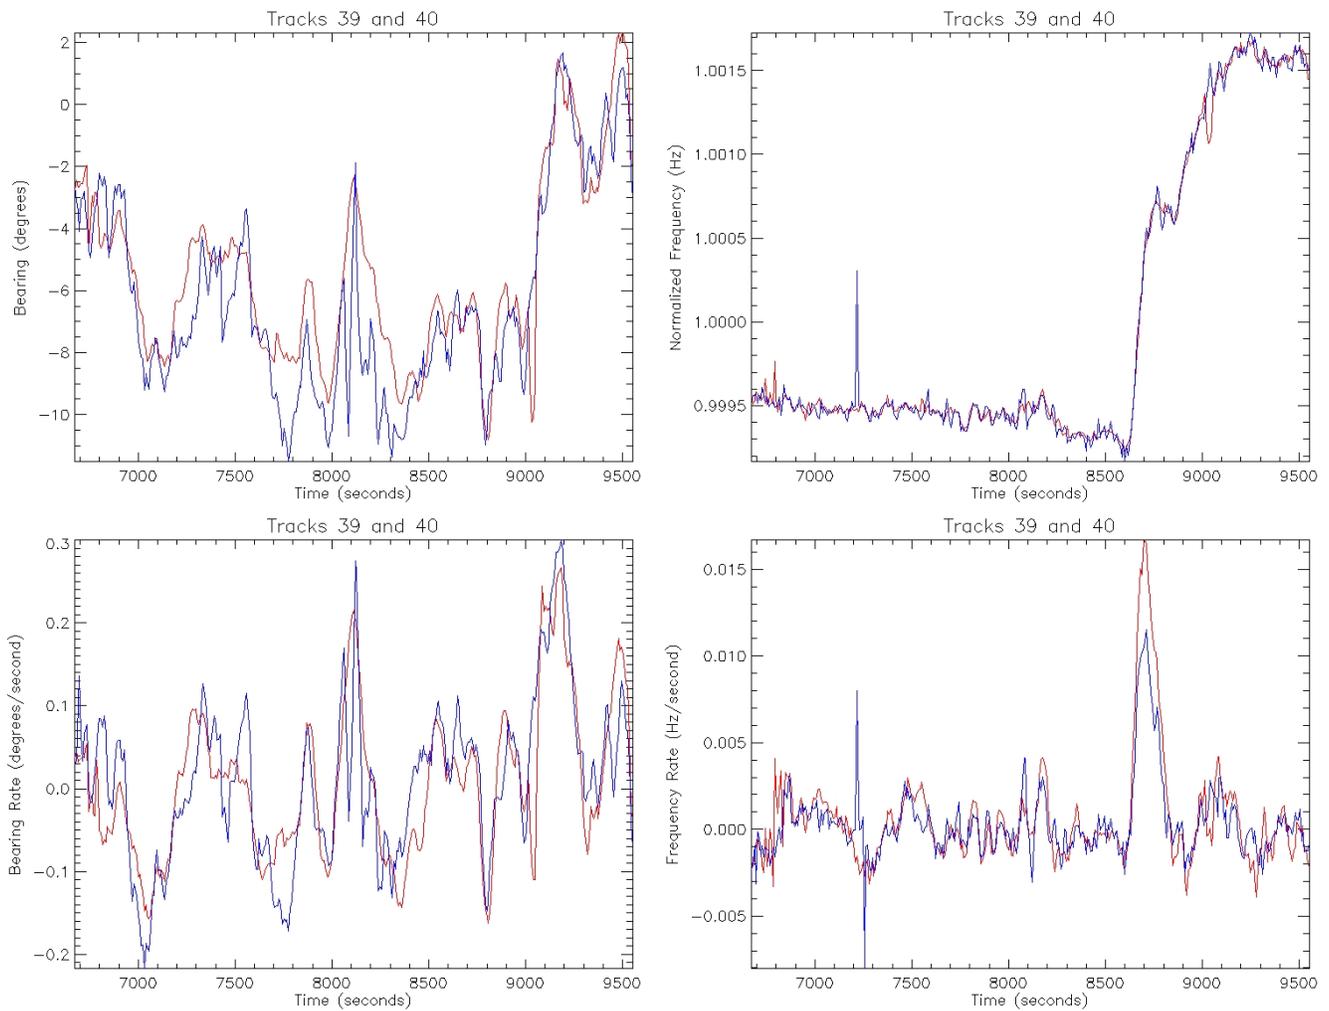
Figure 2 A comparison of tracks 39 (red) and 40 (blue). Only the coincident portions of the tracks are shown.

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 2 and 4 is shown in Figure 1. A comparison of the latest pair of tracks, labelled tracks 39 and 40, is shown in Figure 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 39 and 40. 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 40 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 39 and 40. 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 2 and 4 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.

When the sample correlation coefficient test was applied to tracks 2 and 4 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.

Following the termination of track 2, another track segment was produced at the same frequency, track 24. It continued for only 368 seconds before it was terminated. A third track segment was then initiated, track 37, which continued for 312 seconds before it too was terminated, ending at the same time as track 4. Correlation values for tracks 4 and 24 were r = [0.573, 0.994, 0.597, 0.931]. Correlation scores for track 4 and 37 were r = [0.970, 0.611, 0.959, 0.570]. Both of these sets of values indicated a high degree of correlation between the track pairs.

5 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 an a decision to associate track 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 be useful as a means to identify unrelated tracks.

We have addressed here only the cases of pairs of tracks at different frequencies. Given that track 4 appear to be related to all of tracks 2, 24, and 37 and all of those tracks were at the same frequency, it would be reasonable to associate the three of them together. That association would require only one intermediate association. It is quite possible that comparing three tracks, A, B and C, might show that both A and B, and B and C are related but that A and C are unrelated. A useful formulation of the correlation results should be able to indicate degrees of relationship, which might clarify this problem.

Consider again the problem of track seduction. Over the duration of a track segment, the identity of the signal being tracked changes. In this case the relationship of tracks A and B prior to the seduction and of tracks B and C following the seduction would be accurate. The apparent incongruity of the results would be due to the nature of the underlying problem, not the quality of the test.

This bring up another interesting aspect. 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. This would be the case for track seduction as described above. 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 frequency, then at some point their respective trackers should converge and begin producing identical results. Prior to that 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.

6 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 pairs 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.

References

[1]   R.O. Nielsen, Sonar Signal Processing, Artech House, Boston, 1981.

[2]   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.

[3]   R.J. Urick, Principles of Underwater Sound, 3rd ed., McGraw-Hill, New York, 1983.

[4]   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.

[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.