DATA PROCESSING IN PASSIVE SONAR SYSTEMS

BY

A. M. ROGOYSKI, BSC, PHD, MINSTP, CPHYS
(Logica Defence and Civil Government Ltd)

AND

M. D. JAMES, BENG, MSC, AMIEE, RCNC
(formerly of the Defence Research Agency, Portland)

ABSTRACT

This article discusses the development of the field of automatic data processing, as applied to passive sonar systems. It defines what is meant by data processing, why data processing is important and how automatic data processing can be realized. The subject is discussed in the context of large naval sonar systems, both current and future, where the need for advanced data processing techniques is greatest. Some consideration is given to the individual merits of various techniques; however, detailed descriptions of the algorithms are beyond the scope of this article. Finally, a brief discussion of some of the requirements for implementing automatic data processing systems in operational systems is given.

Introduction

The subject of data processing in passive sonar systems has grown from that describing relatively mundane tasks, for example collation of information obtained by tracking a vessel, to include consideration of a number of very complex tasks which aid or replace those previously performed by human operators. Current data processing techniques use image processing, artificial intelligence methods such as knowledge-based systems and neural networks, together with a wide variety of algorithmic approaches to generate high-level information which is vital to the successful operation of current and future naval vessels.

The Defence Research Agency (DRA, formerly the Admiralty Research Establishment) at Portland has, over the last ten years, been investigating the application of automatic data processing to passive sonar data. The research work has been conducted both within the DRA and through research contracts with industry. The main aim has been to reduce the workload imposed on human operators of sonar systems by undertaking tasks, normally performed manually, with automatic systems. These tasks typically include making initial detections, tracking contacts, classifying them and associating contact information within single systems and across systems.

Future sonar systems, with increased detection ranges and integrated suites of sensors, will greatly increase the load on the operator, necessitating the use of automatic data processing systems if the sonar is to be operable. Current data processing technology is sufficiently far advanced that operational implementations of automatic detection, tracking, classification and association systems may be realized in the very near future.

Data Processing—a Definition

Traditionally, the task of data processing has been carried out by a human operator. Indeed, the definition of what can be called a data processing task suffers from being based upon the rapidly moving dividing line between the tasks performed by machine and those performed by man. Although it is still very common for all tasks undertaken by machine to be referred to as signal processing, this is no longer strictly true.

In the very earliest sonar systems, developed during and after World War I, the operator would listen to acoustic signals amplified from simple microphones. As the systems became more sophisticated, bearing finding became a possibility. This was typically achieved by turning the array mechanically or by having the array of sensors in the sea with variable delay lines attached to each sensor. In this way, a steerable acoustic aperture was realized. The operator was then responsible for initial detection, tracking and, if possible, classification of a target. Additional operator tasks would include extraction of signal parameters such as range or signal strength and monitoring for the presence of weapons such as torpedoes. Such tasks remain essentially unchanged even in current sonar systems, although the methods for achieving these results have changed.

It was not until the advent of digital computers in the 1960s that the methods of processing passive sonar data changed significantly. The advent of digital signal processing technology allowed both spatial and temporal analysis of the data to be undertaken, producing digitally beamformed spectral information from the original acoustic data. Typically, the time-series digitized acoustic data is sampled for a specified length of time and then Fourier transformed. This produces a frequency spectrum. This information is often displayed to an operator in a graphical form known as an 'A-Scan'. By analysing a series of such samples, known as 'updates', the history of frequency information may be used to detect signals at frequencies which are consistently above the level of the background noise. These frequencies, known as 'tonals' are indicative of the presence of a vessel. The information is commonly presented as lofargram data (for Low Frequency Analysis and Recording) in which each frequency spectrum is displayed to the user as intensity information in a single horizontal line. Each new spectrum is stacked vertically on top of the previous one. Signals which appear consistently at a specific frequency appear as vertical lines in the lofargram data. FIG. 1 illustrates this concept.

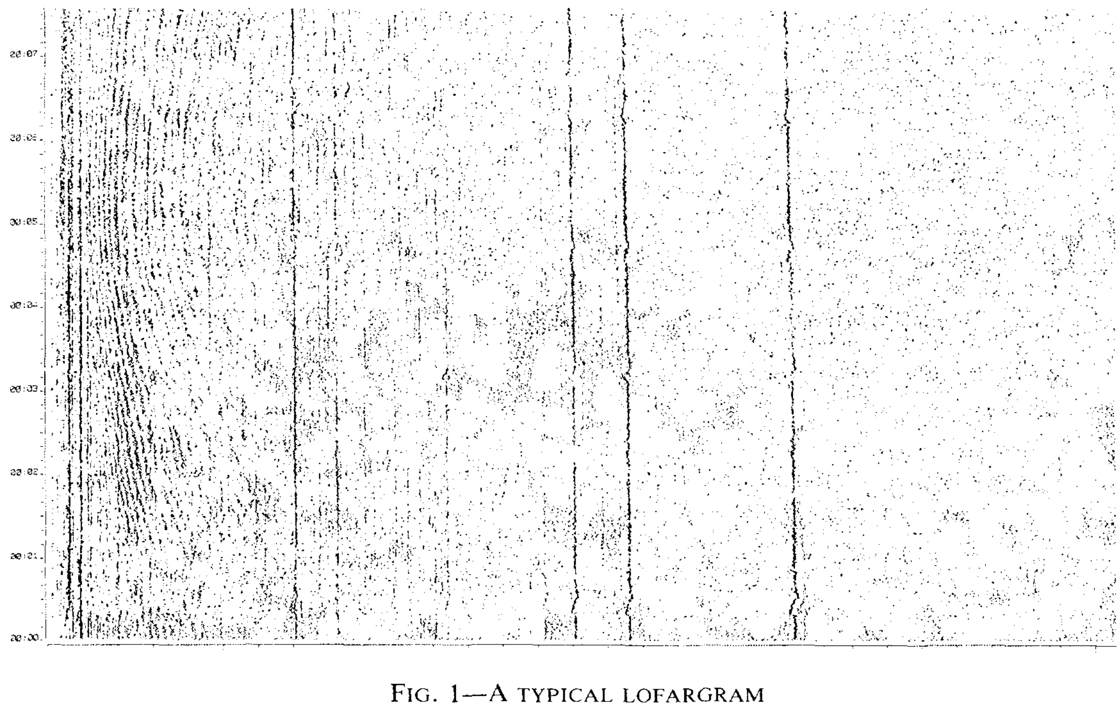

FIG. 1—A TYPICAL LOFARGRAM

Lofargram data, typically consisting of narrowband data at a variety of different spectral resolutions, is supplemented by broadband and demon data. Broadband data can be formed either before or after spectral analysis and describes the total acoustic energy contained within a specified bandwidth and beam. Demon data (for DEMOdulated Noise) provides analysis of the amplitude modulation of broadband noise. Even with sophisticated modern sensors and signal processing systems, there is still a place for analysis of the original acoustic data. Such data is typically used for analysis of transient signals which are too short-lived to be detected by the other processing techniques yet can be used to provide vital tactical information.

The advent of digital signal processing technology has increased the volume and number of different types of data which are available to the operator. However, the tasks which have to be undertaken by the operator using current systems remain essentially unchanged from those which would have been performed in much older systems: detection of signals, extracting bearing, maintaining contact (tracking), association of tracks, extracting parameters such as signal strength, classification and deriving tactical information such as target speed and range. Nowadays, an operator may be required to derive this information from a variety of different sensors, for example, from a hull-mounted sensor array and from a towed sensor array and to associate the results obtained from each.

It is only relatively recently that computing systems have become sophisticated enough to be used to perform some of the tasks which had previously been the preserve of the human operator. It is these systems which undertake data processing tasks that are the subject of this article, rather than signal processing tasks. The distinction made herein between signal procesing and data processing is that signal processing is essentially conventional technology, referring to tasks such as analogue to digital conversion, sensor multiplexing, beamforming and Fourier transforming of data. Data processing refers to the 'smart' technologies which can be used to aid or replace roles previously undertaken by human operators. FIG. 2 illustrates the high-level processing chain of a sonar system, separating the signal and data processing tasks.

Data Processing—the Rationale

In sonar systems currently in use, most data processing tasks are carried out by a human operator. However, this is unlikely to be the case for future sonar systems. Indeed, when considering the possible upgrade paths for current sonars, the performance benefits achieved by the addition of some of the available data processing technology should be compared to those achieved by refitting the sonar.

Current sonar programmes such as Sonar 2031Z already make use of some data processing techniques, while future systems such as Sonar 2057 and Sonar 2076 will not be able to operate at all without making use of these new methods. There are several distinct reasons for this:

- Increasing data volume and operator workload.
- Increasing system complexity.
- Changing operator skill levels.
- Improving performance.

One of the major requirements behind proposed future sonar systems such as Sonar 2057 and Sonar 2076 is that the systems should have better detection performance both in terms of a greater detection range and in terms of detection of weak signals. This requirement is based on the assumption that the threat is becoming stealthier, emitting narrowband and broadband noise at a lower level. The most effective method of achieving this requirement is to

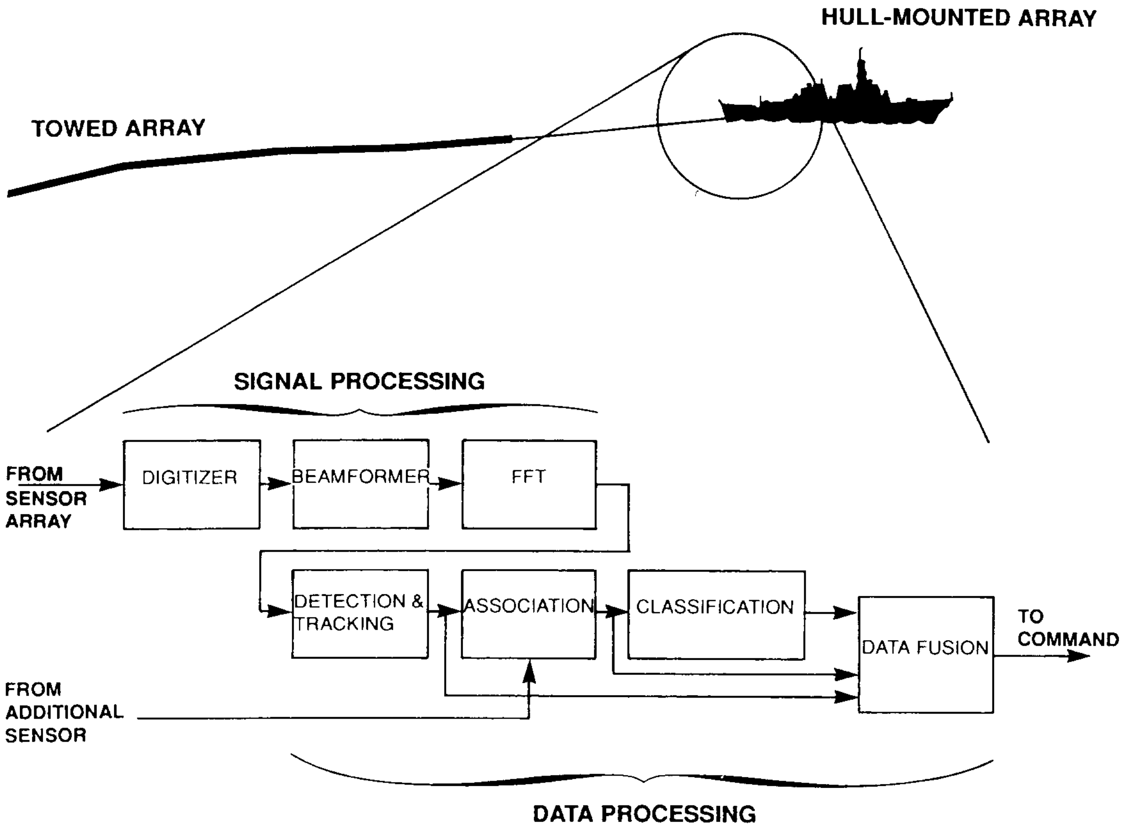

FIG. 2—PROCESSING CHAIN OF A TYPICAL SONAR SYSTEM
FFT: Fast Fourier Transform

increase the acoustic aperture of the sensor array which increases the directionality and thus the detection range of the system. This is the case with the towed array of Sonar 2057. The immediate consequence of this is to increase the number of beams which have to be processed. In addition, if the array is physically long enough, then detection of curvature in the arriving acoustic wavefront is possible. This means rangeformed as well as beamformed data may become available. The requirement for increased detection range is then met but at the cost of greatly increasing the number of beam/range cells (the data volume) which have to be monitored by the operator. The increased detection ranges of the system will also increase the average number of contacts being dealt with at any one time, again increasing the workload of the operator. It is the greatly increased data volume and the associated increase in operator workload of the future sonar systems which has lead to automatic surveillance techniques using automatic detection and tracking becoming essential requirements for such systems.

The second reason for including advanced data processing techniques in future sonar systems is their increasing complexity. Of the Royal Navy sonars, Sonar 2054, equipping the VANGUARD Class SSBNs, is the first integrated sonar suite consisting of the Sonar 2020 bow array, the Sonar 2007 flank array, the Sonar 2019 intercept array and the Sonar 2046 towed array which form a federated sonar system where data from one sonar system may be relayed for processing and display by another. The proposed SWIFTSURE and TRAFALGAR sonar update, Sonar 2076, is required to be a fully integrated system where the information processing is no longer performed on the basis of individual sensors but as an integrated whole. In addition, sophisticated weapons control systems are to be interfaced to these systems. Without very sophisticated data processing techniques, none of these design aims would be possible; the operator would be required to monitor vast amounts of data arriving from several different sensor systems (including Sonar 2057) and to associate contacts detected in more than one sensor system while attempting to undertake all of the normal tasks such as detection, tracking, association and classification.

In addition to the considerations discussed so far, the role of the operator has to be considered—specifically, the required skill and manning levels which are necessary to operate the systems. There is certainly a very strong requirement to keep manning at levels no higher, and preferably lower, than is the case for current operational systems. There is also a movement to reduce the required level of skill needed by the operators, reducing training overheads. It could be argued that reduced manning levels would increase the required level of skill of the operator; however, with staffing costs becoming a major issue in the design of future sonar systems, it seems likely that both manning levels and required skill levels will be reduced. In order to make the systems usable, a significant number of tasks currently undertaken by human operators will have to be performed automatically, freeing the operator for more important tasks.

Finally, it is important to consider what performance may be obtained from future sonar systems. This is linked with the possible change in the required skill and manning levels needed to operate the system. When comparing human operators and automatic systems in terms of the likelihood of successfully detecting a vessel against the likelihood of making a false detection, the human operator would undoubtedly be superior. However, automatic systems have several advantages over human operators: they do not suffer from fatigue or stress and they can be applied to very large data volumes. The latter consideration is probably the most important; without increasing manning levels by a very significant amount, a human operator would simply not have sufficient time to perform successfully tasks such as surveillance on future sonar systems like Sonar 2057. This means that no matter how good the human operator is, an automatic system will outperform him if the data volume is sufficiently large.

Data Processing Techniques

One of the first tasks which should be performed by automatic data processing systems is that of surveillance. This task is probably the most time-consuming for a human operator to perform and requires the least skill. Automatic systems could be used to monitor signal activity in any or all data channels (beam/range cells), making automatic detections of signals, tracking these signals, extracting parameters which relate to the signals and automatically terminating the tracking when the signal is lost. This type of processing can typically be applied to narrowband, broadband or demon data. Each task which is associated with performing surveillance can now be replaced by an automatic system. This section discusses several techniques which are being developed to this end.

Pre-Processing

A data processing task which is often performed immediately after obtaining frequency information is display normalization. In less sophisticated sonar systems, this type of processing is the only data processing performed automatically. It is nonetheless very important for these systems as the operator relies on the displayed data to make detections and track contacts. Display normalization is carried out for two reasons. Firstly, there is normally only a limited dynamic range of grey levels available using conventional display monitors. The dynamic range of the data may be very large, therefore a method for allocating grey levels to signal amplitude information has to be used. A typical example of such a technique is to allocate a specific number of grey levels for every standard deviation about a specified mean. This may be preset or calculated on an update by update basis. Alternative display techniques include histogram equalization, Gaussian binning and logarithmic binning, all common image processing techniques[1]. Processing of this nature is not essential for automatic detection and tracking systems, in fact it can be detrimental because quantising the data for display will cause information, present in the original data, to be lost.

The second reason for normalizing the data is to remove background trends. A common feature of lofargram data is that the average level of background noise is low at frequencies below ≈ 10 Hz, higher at moderate frequencies (≈ 100 Hz) and low again at higher frequencies (≈ 500 Hz). Removing such trends not only improves the display of information to the operator but is vital for subsequent enhancement, detection and tracking algorithms. Methods include averaging, for example using a boxcar filter, or median filtering. Both techniques aim to estimate the local mean of the data which can then be removed from the original data by subtraction or division to produce data with a uniform background noise and standard deviation (assuming the original data has well defined statistical characteristics).

Data Enhancement

When using narrowband, broadband or demon data, the original data may be enhanced to improve the signal to noise ratio of the data. This can aid the operator in making an initial detection and in the subsequent tracking of a signal. Logica has undertaken several studies for DRA Portland (formerly ARE) over a number of years to assess the performance of image processing techniques which can be applied to lofargram data. A common starting point is line integration. At the simplest level, this form of integration produces frequency spectra where each frequency bin is the average of identical frequency bins of a specified number of updates. This method is analogous to that used by operators of existing systems who perform this task manually by viewing hard copy of the lofargram data at an oblique angle. This method has the advantage of being simple but does not perform well for signals/tonals which are unstable over time.

A wide variety of image processing techniques can and have been applied to the task of enhancing lofargram data. Some methods have already been described in the previous section on pre-processing. Generally, the most useful enhancement techniques make use of the known characteristics of the data, i.e. extraction of lines, which tend to be vertical, on a background of noise which has statistics that are typically Rayleigh or exponentially distributed. Standard techniques for line enhancement include gradient operators such as the Sobel operator. The problem of extracting lines from noisy data has similarities with the classical brachistochrone problem of classical mechanics for which there are many different solutions available[2]. One approach has been to describe the frequency spectra using Hidden Markov Models[3]. This identifies the information held in the individual frequency bins of a discrete frequency spectrum as states of a hidden Markov chain. This can then be used to derive a track which is meaningful to the observer.

A slightly different approach is the use of a Bayesian tracker[4] in which the hypothesis of a signal plus noise is compared with that of noise only for each frequency cell. The values used to test each hypothesis are derived from analysis of the statistics of the data, assessing the probability that a data value is part of a noise distribution or part of a signal (at a predefined level) plus noise distribution. Multiple signal levels must be hypothesized if the results are to be used to produce an enhanced lofargram in which the enhanced relative signal amplitudes can be related to the pre-enhanced relative signal amplitudes. The likelihood of a signal can be modified by using prior probabilities, i.e. using the historical information. The technique suffers from being dependent on the assumptions made about the statistical form of the background noise. If no assumptions are made, then the statistics may be measured either locally or globally. The former risks mis-estimating the statistics while the latter risks ignoring local variations in the statistics. There is additional complexity in extending this method to include hypothesis of signals which change frequency.

One of the most successful techniques for enhancement of lofargrams which has been developed to date was developed by Logica and is the path integral[5], based on a graph-theoretical approach to enhancement[1,2,6]. This has similarities with the Viterbi algorithm[7] which systematically examines metrics of all paths to find a path with the greatest metric through a trellis diagram. The trellis in this case is the amplitude of the signal values of each frequency bin as a function of time.

A wide variety of other techniques can be used to enhance lofargram data but they tend to fall foul of one particular constraint; in operational systems, the length of time which elapses between having the raw lofargram data available and having the enhanced data generated (the lag) has to be kept to a minimum. This has been one of the strengths of the path integral technique. The algorithm is not nearly as compute-intensive (i.e. does not require a powerful computing resource) as some of the other approaches discussed here, it can be used to good effect for a wide variety of different lags, it successfully enhances signals which change frequency and it has a natural real-time implementation. Indeed, the path integral technique is now in operational use in the Sonar 2031Z tracker system. This tracker system relies on the operator making the initial detection, based on the information displayed in the original and enhanced (using the path integral) lofargram displays. Once detected, a simple tracking algorithm may be invoked so that the operator is free to continue surveillance in other beams. The technology embodied in the Sonar 2031Z tracker system has since been overtaken by the findings of more recent research programmes, however, the path integral technique is still one of the most attractive algorithms because of its simplicity, robustness and ease of use.

Detection and Tracking

Detection and tracking algorithms may be used on raw or enhanced data. If suitably enhanced data is available, the detection and tracking algorithms should be applied to this rather than raw data. If the data is sufficiently well enhanced, then simple thresholding techniques may be used to make the intial detections and to track signals. The detection and tracking algorithms developed for use with the path integral method have proved to be particularly simple yet effective. This has the advantage of keeping the demands on computational resources as low as possible. Alternatively, detection and tracking algorithms may be an intrinsic part of the enhancement technique. For example, the Hough Transform has been used very successfully in image processing to extract lines in images[1]. Typically, tonals appear in lofargram data as lines and would therefore seem suitable for analysis using a Hough transform method. However, if the line is unstable or diffuse, a common characteristic of real lofargram data, then methods such as the Hough transform tend to perform poorly. Generalizations of the Hough transform may be used to improve the detection of such signals, however, there is a risk that a profusion of different signal templates are developed, making the implementation of such a system clumsy and highly compute-intensive.

Neural networks have been used at the inital detection stage but do not perform well for low signal strengths. If the neural network is presented with an input vector which is derived from the original discrete frequency bin data, and trained to find a peak, then the method is not significantly different from applying a matched filter. A neural network may be trained to recognize different signal types such as discrete and diffuse signals, however, this tends to increase the likelihood that background noise is detected as a signal. By taking the activation level of the network which has been applied in such a way, an enhanced lofargram may be derived. This ensures that the output data can be inspected by a human operator and compared with the original raw lofargram in a meaningful way. Neural networks have been used with greater success for classifying detected line shapes[8]. This work relied on thresholding the signals, cleaning up the resulting data using morphological operators, contour tracing and finally using a chain code representation which is presented to the network. The network then classes the line into one of several groups, for example, stable, oscillating and increasing frequency. The technique is similar in concept to that used by Logica for track editing[5] where false tracks are detected before output (see following paragraphs).

Methods which analyse the statistical nature of the data in the lofargram have been proposed for line extraction; these are essentially the same as the Bayesian methods described earlier. The likelihood estimator may be thresholded to obtain an automatic detection. This technique, while based on a sound theoretical basis, does not perform well for low signal strengths, where the statistics of data containing signal plus noise are virtually indistinguishable from those of noise only. In addition, the technique may perform poorly if the signal is diffuse or unstable in frequency and amplitude.

Track Validation and Parameter Extraction

One of the issues often overlooked when considering enhancement, detection and tracking techniques, is that of validation; it is desirable that the operator is able to confirm that an automatic system is extracting features which can be seen to be consistent with those present in the original data. This is essential if the automatic system is to gain acceptance from the users and is an important part of the consideration of the operability of any system. For example, presentation of data enhanced by Hough transforms may be derived but can be difficult to interpret for an untrained operator. The Bayesian method based on the statistical analysis of the data can provide more understandable information providing multiple signal amplitudes are hypothesized. The path integral produces a display which is identical in format to the original data, simply increasing the signal to noise ratio of the data. This is another characteristic of the path integral algorithm which has lead to its successful use both as a research tool and as a technique used in an operational system.

One of the major constraints which is present in the design of automatic detection and tracking systems is that the false alarm rate should be kept very low. This is of greatly increased importance when considering the large data volumes which future sonar systems may be required to process. All of the detection and tracking techniques which can be applied to lofargram data tend to generate false alarms. One solution to this problem is to allow the operator to validate each track before it is output to other systems. This would be likely to be a very time-consuming task and one which should, if possible, be undertaken automatically. One potentially very powerful technique developed recently by Logica for DRA is track editing[5]. After the original data has had enhancement, detection and tracking processing applied, the track editor assesses the frequency and signal strength characteristics of the tracks, stopping output of tracks that can thereby be identified as false alarms. This has been achieved using an algorithmic approach and, more recently, by using a neural network to analyse the track characteristics. Preliminary work has shown that the use of the neural network for identifying false tracks performs better than the previous manually optimized approach by a factor of two. Further research is under way to develop this technique.

Once signals have been detected and tracked, the next task is to extract characteristic signal parameters. For example, frequency, signal strength, bearing, range and bandwidth can all be extracted automatically. Other parameters can be extracted which are useful to subsequent automatic processing such as association, classification and contact motion analysis.

Frequency may be extracted by reference to the frequency bin of the original spectrum. Sub-cell frequency estimation may be obtained by correlating the data across several frequency bins, centred on the detected track, with a known frequency smearing function. This function would be derived from the function used to window the original time-series data before Fourier transforming it. A typical example of such a window function is the Hanning function. This approach works well for discrete, moderate to high amplitude signals but performs poorly for low signal strengths and diffuse signals. Bandwidth, often associated with the estimation of frequency, may be derived simply by applying matched filters or measuring the full width half maximum of the signal.

A similar technique is applicable for bearing estimation, where the signals measured in adjacent beams at the same frequency can be correlated with a known beam pattern to derive a high resolution bearing estimate[5]. Such techniques have to be replaced by more generalized curve-fitting techniques where the beam pattern is not known accurately. This is likely to be the case for very large towed arrays such as Sonar 2057 where the array movements are likely to cause the beam shapes to change dynamically. Correlation and curve-fitting techniques may be applied to range estimation if range cells are provided by the sonar system. Range estimation techniques such as these may be supplemented by range finding using multipath propagation effects.

Other parameters may also be automatically extracted. Such parameters may be less traditional in nature, for example, identifying harmonic sets and their components. Other features, such as correlated frequency shifts in a set of tonals may be measured and used to provide tactical information. It can also be useful to provide symbolic information about signals. This would attempt to describe the characteristics of a signal in terms used by a human. For example, instead of describing a signal as x dB, it may be more useful to describe the signal as 'weak' or 'strong'. Other, more useful, examples exist. The measurement of such parameters is typically more complex than the more conventional parameters such as bearing and frequency. Symbolic interpretation of signals also requires that the extraction techniques are based upon and validated by human expert's experience.

Association and Classification

The next stage of processing is to associate the various signals. This may be carried out at a number of different levels. At the lowest level, tonals which are smeared across several beams and have been automatically detected in each beam may be associated to give a single tonal. Set of tonals may be associated by considering their harmonic relationship, common frequency shifts, signal strength histories and bearings. Higher level association may involve contact motion analysis, associating contacts in narrowband, broadband, demon and transient information on the basis of the bearing and signal strength history.

Classification of a contact is one of the highest priorities once a vessel has been detected. There are several different systems which are currently being developed to this end. Logica has been engaged, through a research contract with DRA Portland, in designing and building a sonar classification system (SoCS) for use with passive sonar data. The system, using a knowledge-based approach is now complete and expected to undergo assessment trials shortly.

The knowledge base has been built from extensive knowledge elicitation interviews with expert JAAC sonar analysts and contains tasks and rules which have been derived from these interviews. The knowledge base also contains an object-oriented database of vessel characteristics. An inference engine, coupled with a task scheduler is used to apply the tasks and rules to the input data to obtain the classification. One of the most important aspects of the SoCS system is the ability to validate the decisions the system is making. This is supported by an advanced user interface which allows the operator access to all stages of the processing. Validation is also one of the reasons for using a knowledge-based approach; the user may check each task, rule and decision in a way which is intuitive and easy to comprehend because the whole system is based on a human approach to the problem. The system has also been designed to give advice to the operator, based on the knowledge contained within the system and to enable operators to add to the knowledge contained in the system. Because of the amount of knowledge embodied and the ease of use, it is also possible that the system could be used for training purposes.

There are a number of alternative approaches to automatic classification which have also been considered. A straightforward approach uses template matching and is similar to a very crude knowledge-based approach but with little of the power and flexibility of a true knowledge-based system. Vessels rarely appear with exactly the same signature; aspect, propagation and speed are among the reasons for this variability. Template matching can be implemented in the frequency domain, matching predefined tonal sets with those of the input data. Time domain template matching can also be achieved by creating synthetic time series data incorporating the signals due to each tonal and correlating the resultant data with the input data. The frequency domain approach is the more successful of these two methods.

Neural network classifiers have also been developed for passive sonar analysis[8,9]. These systems make use of signal information detected from several passive sonar data sources by extracting salient features, associating them and performing a classification. The neural network is used as a pattern recognition facility, where the pattern of a particular vessel is contained in some combination of signal information present in the input data. The advantage of using a neural network is its ability to recognize an incomplete or slightly modified pattern. This gives the system an element of robustness not present in more 'brittle' template approaches. One of the main disadvantages of such a system is the necessity for the system to be trained on sufficient quantities of representative data so that the system can recognize the vessel. This data is normally of a highly sensitive nature and, in any case, hard to obtain. Although a neural network can be very robust to signal variations, it is generally true to say that the more data that the neural network classifier is trained on, the more reliable will be the action of the classifier. This accentuates the problem associated with validation; it is often very difficult to understand the reason why a neural network has failed to recognize a pattern. For example, if the neural network mis-classifies a vessel, it may do so in a completely spurious fashion. A more structured approach will ensure that there is evidence for every decision made so that even if the classification is not complete it is much less likely to be wrong. There is a role for neural network systems in classification tasks, for example, in areas in which human knowledge about a classification does not exist or is not readily understood. In addition, many supporting tasks such as pattern recognition, which can be used as part of a classification system, can be undertaken by neural network systems very effectively.

An alternative approach to neural networks, which retains some of the robustness to inaccurate or uncertain information, is the use of fuzzy logic. The development of such a system would allow non-binary decisions to be used in the classification process. The method allows concepts used by human classification experts to be included in the system. The difficulty of this approach lies in the derivation of the knowledge necessary to make classifications; tasks and rules are typically derived from consultation with human experts and asking a human to convert concepts into real number scales can be difficult and not always meaningful. Even if this were possible, the potential number of condition combinations explodes as each piece of information is described more 'fuzzily'. A more acceptable approach for implementing this method would be to use a small finite set of values for each piece of information. This approach is, in fact, used in SoCS, allowing the symbolic description of signal characteristics which is closely aligned to the way in which a human operator interprets the characteristic. The number of different values for each piece of information can be derived from analysis of the human classification expert's approach. For example, assuming the classifier is being based on a human expert's knowledge, if a human is only interested in whether a signal is strong or weak, it becomes pointless to store such information in more than these two categories. The above discussion illustrates that there are often alternative ways of dealing with uncertainty to those used by methods such as neural networks or fuzzy logic and that the derivation of the information required to deal with uncertainty and, in particular, its subsequent maintenance is often more straightforwardly accomplished using an approach such as a knowledge-based system. For smaller, more well-defined classification tasks, the use of fuzzy logic may be used to good effect but for larger systems, involving hundreds or thousands of rules and tasks, the additional level of complexity introduced by using fuzzy logic may well be unsatisfactory.

No fully automatic classification system has yet been developed which is capable of performance which exceeds that of a human expert. Various operator tools, somewhat optimistically termed as 'classifiers', are used in operational systems. However, a fully validated, trusted, automatic classification system has yet to be included in an operational system. The knowledge-based system, SoCS, is probably the most sophisticated and successful of passive sonar classifiers to be developed in the UK and with future development programmes, may become the first truly automatic classification system for passive sonar systems.

Data Fusion

At the highest level, all of the original signal information detected, possibly from a number of different sensor systems, should be fused to a very simple representation of a detected vessel: where the vessel is, what the vessel is and what the vessel is doing. There is a large degree of commonality between association and data fusion and the boundary between the two tasks is often blurred. Data fusion will become an essential part of future sonar systems where information from different sensor systems is available. The data fusion algorithms typically have to analyse bearing and range information from different sensor systems and will have to accommodate the likely offsets in bearing and range due to their physical displacement, for example, a towed array and a flank array can be separated by tens or hundreds of meters. The sensors which feed information to the data fusion system are not restricted to passive sonar systems but may include input from active, intercept, radar and electro-optical systems.

Other information such as the classification of a vessel, made using an individual sensor system, may be used as a basis for data fusion when similarly classified vessels are detected in the other sensor systems. Techniques used for associating the information from different sensor systems include statistical approaches, such as Bayesian methods, pseudo-statistical techniques such as Schafer-Dempster techniques, use of fuzzy logic, rule-based approaches, truth maintenance systems and blackboard techniques. Although all these established techniques have been used for data fusion, the area is still the subject of research programmes being carried out by defence establishments and industry and there is currently no approach which stands out as being optimal.

Event Detection

One important use of automatic data processing techniques is to provide the operator with alert mechanisms designed to monitor for specific events. At a very simple level, detection algorithms might notify the operator that a new tonal has been detected or that a signal which is being tracked is fading or has been lost. It is possible to create event detectors which monitor for specific combinations of events, for example a sudden change in the bearing of a contact or a shift in frequency of several tonals may indicate that a manoeuvre has been initiated. A contact which is moving at a high bearing rate may also imply that a weapon has been released. In essence, any detectable event may have an alarm associated with it. Therefore it is simply a matter of identifying the significant events and which sensor system or combination of sensor systems can be used to detect the event.

One instance of a significant event detection system which is worth mentioning specifically is a transient detection system. Traditionally, a human operator would monitor the acoustic signals being received by a sensor system. Facilities such as beam steering and playback allow the operator to record when and where an event occurred. With experience, the operator may also be able to classify the event. Automatic systems have been used to good effect in this area; Logica has been involved, for the Sea Systems Controllerate, in the evaluation of speech recognition algorithms for transient detection and classification. Transient detection using neural networks or wavelets have also been applied to this problem with promising results. Such a system might be used to perform surveillance on a data volume which would otherwise be too large for a human operator to handle, while the provision of playback facilities would still allow a human operator access to the raw data.

Data Processing Tools

In addition to the data processing techniques which can be used to perform tasks previously undertaken by human operators, it is worth considering how data processing systems can be used to reduce the workload of the operator without actually replacing him. Much of this can be achieved by paying more attention to the user interface, supplying tools which make the operator's tasks easier and more efficient. An example of this might be the provision of cursors linked to the data display screens which automatically provide frequency, bearing and signal strength information to the operator. Cursors showing the harmonic relationship between narrowband tonals can also be used to great effect. Multi-level displays may also be very useful. These enable the user to change the type of display, for example, to zoom in on a contact displayed as part of a high level data fusion display, through association and classification displays, right the way through to displays of the raw data. The use of colour displays combined with imaginative ways of presenting data, such as those commonly used for scientific data visualization, may also enable much larger quantities of data to be displayed to, and understood by, the operator.

Implementing Automatic Data Processing Systems

As well as reviewing the techniques available for undertaking data processing tasks automatically, it is necessary to consider the systems design issues associated with implementing such ideas in operational systems. Constraints which may be placed on the design include cost, size, power consumption, heat generation together with availability, reliability and maintainability (ARM). These constraints can be the deciding factor on whether a specific technique will be used in an operational system.

Hardware Architectures

Although most data processing algorithms may be implemented on conventional processors, it is often more efficient to make use of specialized hardware designed for the specific application. However, there is now some momentum behind the use of industry standards for current and future data processing systems, improving the maintainability of the systems and reducing costs.

The most important requirement for most data processing systems is generally computational power. Future systems may be based upon processors such as Intel's i860 or Inmos's T9000 Transputer. Both are currently thought of as very highly powered processors. The i860 processor, although originally designed as a high performance graphics processor, is sufficiently flexible to enable it to be used for a wide variety of compute-intensive tasks. Logica was recently commissioned by DRA Portland to design and build a system which included data processing algorithms for enhancement, detection, tracking, track editing and parameter extraction, together with an advanced user interface to control and analyse the data processing. The system, now delivered, was based on a parallel processing architecture using SPARC and i860 processors to provide a high performance data analysis system. As a by-product, development of the system provided a number of valuable insights into the likely performance of such architectures. An example of this was the requirement to ensure that a good hardware platform is supported by a good compiler. The i860s are vector processors capable of very high throughput. In practice, however, the performance which can be obtained from such processors, while extremely high, is only a fraction of its peak performance. This is due to factors such as the scarcity of vectorizing compilers, requiring that code be hand-tuned to make the best use of the processor's capabilities.

In addition to vectorizing the algorithms, it is desirable to parallelize the algorithms. Vectorizing the code will allow the compiler to make use of the pipelining facilities within the processor. Parallelizing the code will allow a system to be run on several processors. The issue of parallelizing is important in two areas, performance and resource allocation. Parallelizing is necessary if the desired performance of an algorithm or set of algorithms cannot be obtained using a single processor. If the algorithms can be parallelized then the addition of processors together with suitable inter-processor communications should allow the desired performance to be obtained. An important design issue then arises on how to partition an algorithm so that it may be split over several processors. This very much depends on the nature of the task; if it is possible to split the data into several bands, each of which can be processed independently, then the algorithm can effectively be duplicated over several processors and each run with a subset of the complete data. If the algorithm requires access to all of the data in order to make a calculation, then it must be broken down into individual tasks which can be processed separately.

Parallelizing systems can also be extremely useful for resource allocation. For example, if there is suddenly a large number of contacts detected then, if the processing capacity of the system is exceeded, it must respond appropriately. Assuming that the system is to degrade gracefully, it may start by ignoring contacts based upon some criteria such as ignoring contacts which have been identified as non-threat. Alternative strategies might allow some specific processing tasks to be turned off, freeing more processing time to deal with all of the contacts. Probably the most desirable but most complex to implement, is to allocate more processing resource to the tasks. If all processes are designed to be parallelizable, areas of processing not being heavily used may be switched to run on fewer processors allowing additional computing resource to be allocated to the area most in need. Dynamic resource allocation is a very complex problem for system engineers but provides the most acceptable solution to variable loading of a system. This approach also helps manage the system in the event that parts of the computing resource fail.

Hardware Upgradability and Availability

With the present speed of development which most processor manufacturers seem to be able to sustain, it could be argued that designs for future data processing systems based on currently available, or soon to be available, specific processors are not an optimum solution. Such chips include general purpose processors as well as purpose-built ASIC chips. This situation is especially true if one considers that the systems being considered may not be operational for several years to come and that when they are, the systems will remain in service for several years more. The upgradability of chip architectures which form 'families' may provide a better solution. Series of processors such as Intel's 80 × 86 or Motorola's 680 × 0 have proved to have significant staying power, as have Open Systems architectures such as the SPARC architectures. The latest generations of these chip families, together with the advent of new processors such as DEC's Alpha processor, are already starting to make the i860 and T9000 look underpowered. Associated with the issue of upgradability is that of availability. Designs based upon one-off processors which are currently available, are likely to run into problems when, in several years time, the processor is no longer manufactured and there is no direct replacement available. Open architectures like the SPARC, have the additional advantage of being produced by more than one manufacturer, improving the availability and reducing the cost of such processors.

Unique features of specific hardware can make their use invaluable, for example, one of the advantages of the T9000 is that it is designed as a parallel processing engine with communications channels which can use dedicated inter-processor links. This makes it a very attractive processor for use in the type of systems discussed here. However, it should be pointed out that there are now alternatives available in the form of Intel's i-warp processor and Texas Instruments' C40 chip which, unlike the T9000, are available for use immediately. Specialized hardware is also available for applications such as neural networks, fuzzy logic and knowledge-based systems.

The trade-off between the potential performance benefits of using application-specific hardware and the ARM benefits of using common hardware processors has yet to be resolved. Industry standards are used for a wide variety of interfacing needs, for example, the use of VME-based systems is common. This will change as bus bandwidth demands increase, probably switching to a standard such as Futurebus. Data highways are also being built to industry standards such as FDDI, indicating that the use of specialized hardware in the future is likely to be discouraged.

Software

Apart from the hardware of data processing systems, there should also be consideration given to the associated software. Defence systems are commonly implemented in the Ada programming language. Outside of the defence sector, there is an increasingly large number of people using object-oriented approaches because of the more intuitive approach to design which can be used and because of code-reusability. However, many of the tasks which have to be undertaken for automatic data processing are not suitable for object-oriented approaches. There may also be a significant additional cost factor associated with translating the software which exists for some of the more esoteric data processing systems to a standard software language such as Ada. The necessity to translate systems written in languages like Fortran or Lisp, into Ada, arises because researchers tend not to want to be constrained to undertake research solely with the aim of developing an operational system. There is a significant overhead to be incurred by using languages like Ada for simple research programmes which is difficult to justify unless one can be sure the results of the research will be sufficiently valuable. Therefore it seems likely that research will continue to be undertaken in a variety of programming languages only to be translated into languages such as Ada if the research becomes useful to operational systems.

Conclusions

The last decade has seen a significant amount of research being undertaken by defence establishments and industry to try to realize the concept of automatic data processing systems for passive sonar. There are several emerging technologies, some at quite an advanced state of development, which can be used for tasks involving surveillance, detection, tracking, parameter extraction, association, classification and data fusion. However, as the transition between manually operated systems and automatic systems takes place, it is essential that the new technologies are properly validated. This issue will play a significant part in deciding whether advanced data processing systems will be accepted for operational use; systems which are designed to allow the operator to supervise and to confirm their correct operation will have a distinct advantage over those systems which want to replace a human operator with a black box. As the supporting technologies become more diverse and sophisticated, it seems likely that research into automatic data processing will continue.

Overall, if future sonar systems such as Sonar 2057 and Sonar 2076 are to be built, then sophisticated automatic data processing systems will constitute an essential part of their make-up. The pressures to improve system performance, to increase detection ranges, to reduce manning levels and the required skills of the operators, all lead to the conclusion that many of the tasks previously undertaken by a human operator will have to be performed by machines. However, the operator must still be seen as the most important part of future sonar systems. This needs to be addressed by consideration of the operability of the system in question. Operability issues include a very wide and diverse range of topics including environmental analysis, user requirements analysis and scenario analysis together with many other areas. The operability of modern sonar systems is too large and complex an area to be considered herein. However, it will take a great deal of ingenuity on the part of the system designers of future sonar systems to achieve the right balance between the computer and the operator. It is clear that a great many tasks can be, indeed have to be, delegated to automatic systems. However, the technology should not be used for technology's sake but must be carefully targeted at areas where there is a genuine benefit from using such systems.

References

1. Gonzalez, R. C. and Wintz, P.: *Digital image processing;* Addison Wesley, 2nd edn., 1987.
2. Chachra, V. and Moore, J. M.: *Applications of graph theory algorithms;* New York, North-Holland, 1980.

3. Streit, R. L. and Barrett, R. F.: Frequency line tracking using hidden Markov models; *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 38, no. 4, April 1990.
4. Maksym, J. N., Bonner, A. J., Dent, C. A. and Hemphill, G. L.: Machine analysis of acoustical signals; *Pattern Recognition*, vol. 16, 1983.
5. Rogoyski, A. M. and Turner, A.: General research on automated enhancement of Lofargrams (GRAEL); *DRA Technical Memorandum*, 1992.
6. Morris, O. J. and Constantinides, A. G.; Graph theory for image analysis: an approach based on the shortest spanning tree; *Institution of Electrical Engineers Proc.*, vol. 133, part F, no. 2, April 1986, pp. 146–152.
7. Viterbi, A. J. and Omura, J. K.: *Principles of digital communications and coding;* New York, McGraw-Hill, 1979.
8. Russo, A. P.: Constrained feedforward neural network ensembles for recognition of sonar signals using shape; *Proceedings UDT*, Microwave Exhibitions and Publisher, London, July 1992.
9. Sheppard, C. P. and Gent, C. R.: A neural network based sonar classification system; *Underwater Systems Design;* Nov./Dec. 1991.