Linking Sonar Performance to Operational Effectiveness – A Review of the Issues

Dr J S Ryder & Mr S C Johnson,
Defence Science and Technology Laboratory,
Naval Systems Department,
Bldg. A33, Winfrith Technology Centre,
Dorchester,
Dorset,
DT2 8WX,
United Kingdom

E-mail: sjryder@dstl.gov.uk / scjohnson@dstl.gov.uk

ABSTRACT

Many studies of different aspects of Operational Effectiveness in Anti Submarine Warfare have taken place during and since WWII, aimed at supporting real world operations, acquiring new systems and platforms, and also at exercise planning and analysis. The overwhelming majority of these studies have addressed the capabilities of sonar systems of different types, reflecting the ability of the medium of the ocean to transmit sound, as distinct from other types of wave motion, with low levels of attenuation. These effectiveness studies have been supported by the considerable efforts which have been devoted to improving the ability to predict sonar performance. Despite all of these endeavours, the prediction of sonar performance and effectiveness is by no means an exact science, having, as it does, to take due account of the stochastic nature of the various processes involved, and our always incomplete knowledge of the local oceanographic environment.

In recent years, some distinct trends in the analysis have emerged. Firstly, the models used to perform the effectiveness analysis are evolving from classical large Monte-Carlo search and engagement simulations to include simpler types of model, more suited for rapid analysis. Secondly, the aims of the analysis are concentrating more on the absolute, rather than simply comparative, Measures of Effectiveness (MoE), which are required to support Capability Management, the cornerstone of the new Equipment Acquisition process within UK MoD. Finally, the analysis is required to address greater diversity, not only in terms of the oceanographic conditions represented, but also in terms of the military operations addressed.

These trends serve to emphasise the importance of ensuring that there is a smooth link from measures of performance to the effectiveness modelling methodology. This requires that sources of systematic bias in the modelling, arising from the stochastic nature of sonar, do not bias the resulting MoE. This paper discusses some of these sources of possible bias, and the steps which need to be taken to counter them.

1.0 INTRODUCTION

This paper addresses some of the issues involved in using sonar performance data in operational effectiveness modelling to support equipment procurement decisions. It gives an overview of the context in which the modelling is performed and then looks at some of the factors which, if not addressed appropriately can lead to biased or misleading interpretation of the results.

Paper presented at the RTO SET Symposium on "Capabilities of Acoustics in Air-Ground and Maritime Reconnaissance, Target Classification and Identification", held in Lerici, Italy, 26-28 April 2004, and published in RTO-MP-SET-079.

Sonar performance, operational effectiveness and validation are all inter-related. Firstly, sonar performance modelling aims to assess the capability of a sonar in terms of its ability to detect a target as a function of range. Secondly, effectiveness modelling examines the military utility of such a sonar in terms of how it influences the ability to, for example, search an area or barrier, with the aim of detecting a target with a given probability. Finally, there is the need to validate models of these types, and the ambition to do this with reference to measured data from the real world, from trials and exercises.

Because of the relatively low attenuation with which sound, as opposed to other forms of energy, is transmitted through the sea, sonar is of fundamental significance for the detection and prosecution of targets in underwater warfare. The problem for the analyst is to provide robust assessments of the effectiveness of detection systems which use sonar, and whose performance is influenced very strongly by stochastic factors associated with detailed features of the environment, as well as by sensor and target features which are difficult to model explicitly.

2.0 CHARACTERISING SONAR PERFORMANCE

2.1 Detection Range

Detection range is certainly the simplest measure of performance to use as an input to considerations of the operational effectiveness of underwater platforms. As a degree of covertness is required for success in many tasks carried out by underwater platforms, operational effectiveness at least partly depends upon the answers to questions of the form 'what is the probability that I can carry out a specified task whilst remaining outside the detection range of any potential threats in the vicinity?' Detection range is a powerful tool in that: a) it is easy to visualise operational situations in terms of own platform and threat detection ranges, and the dimensions of the operating area. b) a range of simple analytic models exist which, using detection range as an input parameter, yield measures of operational effectiveness. Such models are easy to use and are essentially interactive. They lack, however, the detail of more complex simulations, but they are well suited to initial investigations in the early part of the acquisition process.

The problem with using detection range as a measure of sonar performance is in defining it. If nature were well ordered, so that there would be a probability of unity of detecting a target within the detection range and a probability of zero outside it, i.e. the 'cookie-cutter' approach, there would be no problem. The ocean environment however results in the probability of detection varying both with time and range. How detection range is defined in these circumstances can significantly affect the resultant measure of effectiveness.

A simple definition of detection range is the longest range at which the platform will (with 100% confidence) react to a contact. This is not just the longest range at which there is a signal excess on a sonar channel, it also includes assumptions of persistence of contact and classification.

Problems involved in using a single detection range to characterise a platform's sonar performance tend to arise when sonar transmission loss varies non-monotonically with range. Such conditions give rise to sets of annular in-contact zones rather than a detection 'circle'. Problems also arise when the searcher does not have an overwhelming sonar advantage, so that the inherent variability gives rise to the possibility of the target detecting first and reacting. The example in section 7.1 of this paper illustrates the problem.

Some problems can be alleviated by considering how the sonar is employed, and asking questions such as 'is the search a continuous sweep, is it a dipping (or sprint-drift) search, and is the sonar active or passive?' In the case of a sweep search, the target will pass through the outmost in-contact zone, but the question then arises as to whether it remains in the zone long enough to trigger a reaction on the part of the searcher. Appreciation of the likely encounter geometry can inform this question, but may not yield an unambiguous answer. If the sonar search is a dipping one, the converse problem arises. The target could be between in-contact zones at the time of the dip. However, if the dip is of sufficient duration, and the target has sufficient radial velocity, the geometry could evolve to allow a contact to be made. Again an appreciation of the details of the situation can inform the situation. For dipping sonars the use of so called effective detection ranges, which define an area equivalent to that contained in all the in-contact zones, can be a useful measure of performance provided that the geometry remains essentially unchanged during the duration of the dip.

The use of simple detection ranges in vignettes where the target also has a sonar capability can lead to an erroneous ‘black and white’ view, ignoring the effect of sonar variability. The platform with the greatest detection range is always, and erroneously, predicted to detect first, a situation which may only be true if it has an overwhelming sonar advantage.

2.2 Detection probability

One approach to addressing the effect of variability in the ocean is to consider the variation of detection probability as a function of range for both own and threat systems. Detection probabilities can readily be combined in simple numerical models to ask simple questions of the form ‘what is my search effectiveness against a target which will evade if he detects me first?’ (i.e. the probability that he comes within my detection range given that his detection range is less than mine). Detection probability data for more than one sonar channel or system on a platform can be combined, to give an estimate of the overall effectiveness of a complete sonar system.

Such simple numerical methods can often be used to examine mainly one-sided situations where a simple target reaction may be assumed. They become less applicable when the situation is highly interactive – such as in an ASW melee, in which even the sequences of events can be highly variable. The models are also best suited to situations where the probability of making an initial detection decreases monotonically as range increases. This does not preclude environments where there are annular in-contact zones, provided that the vignette involves a continuous search so that the target closes from long range. For a dipping (or sprint-drift) search the probability of initial detection does not decrease smoothly with range in an environment with annular in-contact zones. For this type of problem, manipulation of the detection probabilities is more involved but can still be achieved; it is more difficult to frame the problem in terms of simple combined probabilities and assumptions.

3.0 EFFECTIVENESS MODELLING

Operational effectiveness (OE) modelling seeks to conduct analysis of systems in terms of Measures of Effectiveness (MoE) which are directly related to military benefit. ‘Systems’ in this context can range from individual systems, such as weapons (where the MoE could be the probability of sinking the opposition) to the full set of forces fighting a campaign (where the MoE could be defined in terms of campaign outcome). This hierarchy is reflected in the progression from low level analysis at the single engagement / platform level, through medium level analysis which considers whole forces and missions, to high level analysis embracing the full scope of activity within the overall campaign. It is the link to military outcome that distinguishes OE modelling from system performance modelling (which addresses issues such as prediction of sonar detection ranges). OE models are, however, dependent, directly or indirectly, on performance models to supply the data on which they depend. An example is where sonar performance predictions are used in an OE model to predict sonar search effectiveness. It does not matter whether the OE model takes performance data as input or has one or more performance models embedded within it (indeed some models adopt both approaches).

Sonar performance modelling has perhaps the greatest impact on OE models when the OE model is at the engagement or 'vignette' level (also referred to as low level OA). A typical vignette could be a submarine carrying out an ASW search and attack task. The scope of possible vignettes is as broad as the range of roles that a platform could undertake; a submarine, for example, is certainly a multi-role platform. This paper concentrates on vignettes based around ASW related scenarios relying on sonar contact.

A signal excess approach to modelling sonar performance (i.e. one in which the sonar contact status is re-evaluated as the engagement evolves) is frequently adopted in Monte Carlo simulations. This type of simulation has been a powerful tool in OE modelling for many years. As the geometry between the protagonists evolves during a replication of the engagement, the signal excess is evaluated, with an embedded sonar performance model which uses the sonar equations. Depending on the detail of the model sonar variability can be represented explicitly, by for example aspect dependence of sources and receivers, or by including a random component into the sonar equation, or by a combination of both.

Typically the simulation data input consists of the elements of the sonar equation and a set of look-up tables of propagation-loss as a function of range, frequency and depth. The data requirements for even this simple approach can be significant if aspect and speed dependence are included for a range of platform sonars. There is also a raft of sonar features, including amongst others, near field effects and persistence of the target in the sonar beam that may also need to be represented. As with the detection probability approach variability in sonar performance and the correlation of sonar variability needs to be considered.

What the OE modeller needs is a set of models, with sonar measures of performance at a level of detail appropriate to the problem being addressed. Three levels are discussed here:

- Analytic models, in which sonar performance is characterised by a simple single-valued detection range
- Semi-analytic or numerical models in which sonar performance is characterised by a detection probability which is a function of range and depth
- Simulations, in which the sonar signal excess is evaluated, and the state of contact determined as the engagement evolves.

In general the simpler analytic models tend to be used at the early stage of an acquisition programme, where there is a wide range of potential solutions, to understand and scope the solution space. More complex simulations tend to be employed on a more limited, and better defined, set of potential solutions later in the acquisition cycle. This is not just [or 'only partly'] because we know more about the system details to be able to model them. Early in the programme the aim of OE modelling is to identify options worthy of further consideration, whereas later in the programme the aim is to discriminate between options before decisions to order and manufacture. In this context, it is important to ensure that the model is capable of discriminating between the options in a sensible manner. It is necessary for the analyst to have a broad understanding of the nature of, and the differences between the different options, and, appreciating how these differences may impact upon effectiveness, to ensure that the model algorithms sensibly reflect the mechanisms involved. It is worth noting that inclusion of too much detail in a model can obscure the clarity of understanding of the results, and so a balance needs to be sought.

As in the real world, Command tactics have an important role to play in OE modelling. They can introduce biases in the modelled sonar (and other) system effectiveness depending upon how the system is used (for example in designing search strategies) and upon the context in terms of other threat and own systems, and the environment. As an example, active sonar systems may be used to detect, or to deter. Different MoEs will apply in either case. An important feature of simulations is that tactical actions are explicitly represented, and can be triggered by changes in sonar contact status (e.g. if a sonar contact is lost, the action of the platform losing the contact might be to turn in the direction of the previously held contact, in the hope of regaining it). A consequence of this is that considerable thought needs to be given to representing the consequences of gaining and losing sonar contact. The problem will be less difficult in an environment where the propagation-loss increases monotonically with range than in one with a complex set of in-contact zones. It is the authors' experience that failing to appreciate fully the effect of the environment on change of sonar contact state, in terms of the Command tactics, can be a fruitful source of anomalous results.

The UK acquisition process is based on identifying capabilities required to produce successful military outcomes in defined scenarios. The capability requirements emerge from an analysis of capability gaps, and should not be system specific. A wide range of potential system solutions can then be examined in the early part of the acquisition cycle, before narrowing down the range of options as the process progresses. This approach, which seeks to identify systems which can fill capability gaps, differs radically from earlier approaches, which sought to identify how much better a replacement system was in comparison to the current one. This approach has implications for effectiveness modelling. Firstly it is usually necessary, at the early stages of the acquisition process to examine the cost-effectiveness of a range of diverse systems, each of which could potentially fill the capability gap, in a variety of scenarios. Secondly, because the question being asked is 'will the system fill the capability gap?' (without over filling it) there is an emphasis on absolute measures of effectiveness. There is a tension between these two implications: the first calls for assessment of a wide range of system/scenario combinations and encourages a rapid, analytic approach to effectiveness modelling. The second, calling for absolute measures, encourages the inclusion of increasing layers of detail, which may lead to complex (and time consuming) simulations. This tension can be alleviated if the analyst has an understanding of the implications of the assumptions made, and which are inherent in the level of detail within the model.

4.0 THE NATURE AND EFFECTS OF SONAR VARIABILITY IN OE MODELS

4.1 Modelling Variability

Variability can be included in the Sonar Equation in the form of a probability distribution function with the properties of standard deviation, and decorrelation time (such that if samples are taken at intervals significantly greater than the decorrelation time, they are essentially independent). The distribution function is usually taken to be a normal distribution. The factors which contribute to variability will depend upon the level of fidelity in the model, in that detailed models will represent certain factors, which are deterministic, explicitly, whereas less detailed models will represent them as contributing to the overall variability.

In a Monte Carlo simulation, the effect of variability can be represented by the use of pseudo random variates. Four different ranges of decorrelation times are of interest, which describe variability in the very short, short, medium, and long timescales. The following paragraphs define the timescales involved in each case and identify some characteristic sources of this variability. They discuss the issues of whether and how these four different categories should be represented in models, and, supplemented by the more detailed examples in Section 7 of this paper, discuss the effect of including the different types of variability in models.

4.2 Very short term variability

Very short term variability (VSTV) implies decorrelation times less than the sonar integration time (i.e. up to a very few minutes). Contributions to VSTV arise mainly from small scale inhomogeneities and scattering in the sea. In the case of VSTV, the variations are integrated over the time it takes for the sonar to make a detection, and will therefore tend to average out, and not give rise to a bias. VSTV therefore tends not to be represented in engagement level models.

4.3 Short term variability

Short term variability (STV) covers decorrelation times greater than the sonar integration time, but less than the typical time over which a sonar search takes place, typically a period from minutes to a few days. Contributions to STV arise, for example, from relative bearing and aspect angle changes, and changes in depth, as well as from multi-path interference. All but the last of these can be regarded as being predictable, and, given sufficiently detailed data, a suitable model would be able to represent their effects explicitly. Multi-path interference is an example of a factor which can be modelled using a coherent sonar propagation loss model, but certainly in many cases, to achieve an exact alignment between the model predictions, and the real world would be extremely difficult. In the case of STV, the hunter and hunted get multiple looks at each other during the period of a search, and the sonar signal excess (SE) will vary from one look to another. (The SE is defined such that a detection occurs when SE $\geq$ 0.) Initial detections will tend to occur on those occasions when the SE is highest, thus detection ranges will tend to be higher than predicted by the 'raw' Sonar Equation (i.e. excluding STV). To compare these 'raw' predictions with real effects from sea, a correction factor will have to be applied. In practice, this is found to be approximately equal to the SD of the STV distribution. Whether or not STV is represented explicitly in effectiveness modelling depends upon whether the sonar equipment provides a continuous or an intermittent cover. If continuous, as in the case of passive sonar, the effect of adding STV, and then applying the correction tend to average out, and therefore STV tends not to be modelled explicitly. If the cover is intermittent, as for active sonar, or sonars which sample at intervals, one may need to invoke STV within the sampling process – with the appropriate correction factor. In such cases, a random variate will be added to the sonar equation within the model (incorporating appropriate time and frequency correlation) whenever a sonar contact status is reviewed within an encounter scenario.

4.4 Medium term variability

In the case of Medium term variability (MTV), the decorrelation time is greater than the time of a search or engagement. Also, the decorrelation time is less than the interval over which, for example, seasonal and geographical changes have a significant and predictable influence over sonar performance, and which therefore need to be modelled explicitly. Intervals of hours or days to weeks are therefore relevant. MTV covers a number of effects which require to be addressed in many cases, in order to provide robustness to the analysis. For example, variations and uncertainties in target sonar characteristics, the day-to-day variation of typical sonar transmission loss, as well as the local geographical variation, together with the effects of shipping and weather on noise levels and noise anisotropy, and a number of other effects all contribute to MTV. It has to be admitted that most of the contributory factors are deterministic, however, the ability to quantify them with precision would require a very considerable programme of measurement. The mechanisms by which MTV influences the relationship between achieved detection ranges, and standard sonar equation predictions are more complicated, and depend upon the following factors:

- Numbers of sonar channels available (to each contender). The variability of these channels may be correlated to differing extents. If the level of correlation is high, detections by good channels will dominate detections by poorer channels, and the 'poor' achieved detection ranges will reduce accordingly. Conversely, if the channels are not well correlated, the poorer channels will make a greater contribution when their performance exceeds predicted values. Overall detection performance will depend upon how well the channels complement each other, and is calculated by considering the overall detection probability, combining that of the individual channels. Note that this effect applies not only to the different sonars / sensors of one contender, but also biases the overall achieved initial detection ranges achieved by a searcher, as compared to his opponent.

- Search time / type of search. If we search for a long time, e.g. searching an area, and the area size is 'sensible'; we expect the search to have a high success rate. This will span the sensor P(d) vs. range distribution, and it is reasonable to expect detection ranges distributed about the median of the predicted distribution. If we search for a short time, or if the area size is unrealistically large, we expect lower probabilities of success, and those searches which are successful will tend to exploit the longer ranges of the distribution, hence there will be an upwards bias in the achieved detection range. This situation can also occur when time is limited by the type of search, e.g. a barrier search, or a datum search.

The example in Section 7.2 shows how search probabilities depend on the assumptions made on medium term variability and on the target's sonar performance.

At a detailed level, a number of sources of what would otherwise be classified as deterministic, or explicit information, may be represented as contributing to the overall variability, for example, if precise knowledge of target radiated noise is not assumed. This will give rise to apparent biases, as described above, in the results; but the important point is that the stochastic variability, as modelled, should exclude any contributions which are explicitly represented in the model. MTV is almost always represented in engagement level simulation models, usually, as with STV, in the form of a random variate added to the sonar equation within the model (incorporating appropriate time and frequency correlation). By contrast, in the case of MTV, the variate is changed from one replication of the simulated encounter to the next, taking a constant value within each individual replication.

4.5 Long term variability

Long term variability (LTV) covers periods of more than a few weeks. It is seldom invoked, as it is intended to cover spatial and temporal intervals over which, for example, seasonal and geographical effects have a significant and predictable influence over sonar performance. In principle, this requires explicit representation, as it usually tends to imply systematic and significant shifts in assumptions, these being in excess of what could ordinarily be regarded as a statistical variation. LTV would almost always be best addressed by the explicit introduction of predicted seasonal or local variations as sensitivity studies, and, in consequence, is not usually modelled as such.

4.6 Discussion

Monte-Carlo models, if they are supported by good data, represent real effects, some of which are fundamentally stochastic in nature, others deterministic, and others again essentially deterministic, but which are represented stochastically, for lack of more detailed information. Thus, in the real world, we cannot always expect achieved detection ranges to match those predicted by the sonar equation. The mis-matches have traditionally been ascribed to semi-empirical 'degradation factors', which are attributed to a number of different causes.

We need, firstly, to establish a level playing field upon which we can compare 'real-world' and 'predicted' detection ranges. This paper indicates that apparent biases, reflecting scenario details, are calculable. Secondly, we need to ensure that if we do not use simulation, we use unbiased measures of detection performance in the effectiveness modelling.

5.0 IMPLICATIONS FOR THE USE OF REAL WORLD DATA

We now come to the third member of our triangle, and the need to relate model data to real data. But what are real data? There may be some artificiality in the results of an exercise, for example. Planners want to make most efficient use of the resources available, and scheduling a high proportion of trials time with no contact could be regarded as wasteful, whereas in the real world of military operations much time may be spent scanning displays in vain for targets which are not detectable. Also the psychological situation is very different if lives are at stake. The situation in a scientific trial is even more artificial

Following previous discussion in this paper, it should be clear that in order to exploit real world data to support validation, it is necessary to measure and understand details of the scenario in which the results were obtained in very considerable detail. It is probably relevant to point out also that the problems of predicting or modelling the sonar characteristics of a specific environment at a specific time are different from, and more demanding than, the problem of modelling nominally the same environment to support OE studies. In the former case we are concerned with knowledge and predictions, almost always incomplete, of actual conditions on the day, and how best to exploit system and platform characteristics in the light of the knowledge and predictions. There may well be some degradation in the nature of the tactical response if there is a mismatch between the understanding and the reality. In the case of OE studies, the need is to represent the nature of a possible (typical) environment in sonar terms, in the expectation that, given an ensemble of such representations of different environments, the degree of robustness of the sonar system can be assessed. In this case, of course, there can be no such mismatch.

There are two caveats to this, firstly in the case of niche systems, which are designed to exploit particular environment factors, in which the identification of such factors is likely to be an issue. The second is that although it may usually be assumed that the Command has a correct understanding of the environment which is the subject of the representation, consideration of particular issues may make it necessary to assess the extent to which this assumption is likely to be justified. An example of such an issue is the question of how much to invest in environmental prediction.

6.0 CONCLUDING REMARKS

This paper has attempted to show how some typical sonar measures of performance can be used or misused in operational effectiveness modelling. It has also attempted to show that increasing the detail of the sonar representation in OE modelling is not necessarily the best way ahead.

Sonar modelling is used to serve different ends, to research and design new sonar systems, to assist in the analysis of exercises and operations, and to conduct effectiveness studies to inform equipment acquisition, and tactical development. All of these purposes can support a virtuous circle, whereby at sea observations can provide validation to support equipment acquisition studies, and studies to ensure that the equipment functions correctly, and is used to best advantage.

Whatever the type of analysis, it is important to use domain understanding to select an appropriate level of fidelity in the modelling; and, in particular, to ensure that critical factors which influence effectiveness are represented in an appropriate way.

There are difficulties in applying real world data to achieve an ‘end to end’ validation of sonar modelling. A part of this arises from the stochastic element, which requires multiple observations to enhance precision. This paper has indicated that many of the sources of bias, which degrade accuracy, can be compensated for given ‘adequate’ knowledge of the overall environment. However, perhaps the most difficult problem is to achieve this state of knowledge whilst preserving credible and realistic sonar system operating conditions. The reality therefore appears to be that validation should be seen as a staged process.

7.0 EXAMPLES

7.1 Detection Ranges

Figure 1 shows two signal excess vs. range curves – an idealised one in blue, in which propagation-loss follows a spreading plus absorption law, and one in red which may be more typical of the real world.

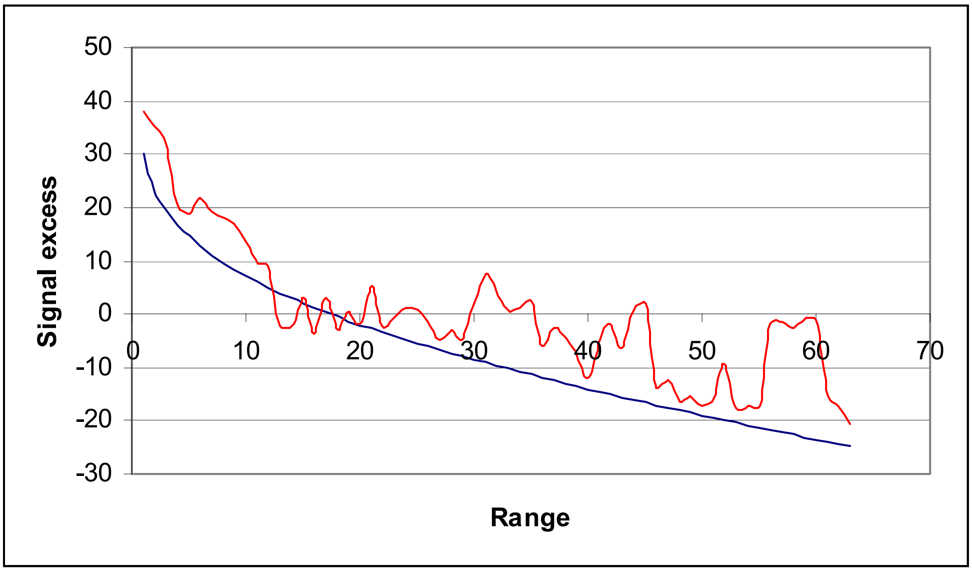

Figure 1: Idealised (blue) and modelled (red) plots of signal excess vs. range

In the idealised case the detection range can be defined without ambiguity as the point where the signal excess becomes zero. For the more practical case there are a number of possibilities and the most appropriate may depend on the type of sonar and scenario being evaluated.

For example consider the case where the sonar is continuously operational, such as a passive sonar on a submarine searching at slow speed, closing on a contact, also at slow speed. At approximately 45km range a positive signal excess will be recorded. This assumes that the target will remain in the signal excess region for sufficient time for the sonar processing gain to be achieved. Possibly more importantly the signal excess is sufficiently persistent for a contact to be recognised and acted upon. Simple arithmetic will indicate the minimum period of signal excess, though of course the case of direct closure on a target is a limiting one, and the actual geometry of the encounter will play a part. However, if the platforms are closing directly at high speed these conditions are less likely to be met and no detection will occur. An analogous situation exists with active sonar where, for a contact to recognised, returns from a finite number of pings are required.

As the range closes further the signal excess falls below zero, only to become positive at about 35km. Here the signal excess is positive over a wider range band, and even on the assumption of direct closure with both platforms at the high speed there is sufficient time for a positive detection to be made. Finally there is the most pessimistic case where the platforms close until the signal excess is continuously above zero.

If the sonar is operated intermittently, rather than continously, then the target could well be in one of the ‘valleys’ of negative signal excess and not be detected. One way to overcome this is to define a detection range that encloses a circle equal in area to the annuli of positive signal excess regions. This is, of course, an artificial construct and although it may be applicable for some search/detection calculations, it may be totally inappropriate when considering sonar detection range in relation to range characteristics of other sensors, including the threat's and other systems (such as weapon system ranges).

A further complication is the effect of variability and uncertainty. The figure above shows that there is a region of just negative signal excess at a range of 60km. If conditions were such, or the sonar performance had been underestimated by a couple of dB then in some scenarios a detection could be made out at that range.

What is the analyst to do? Examining all potential detection channels in the manner described above would be an ideal, but may not be practicable if assessing a multi-channel sonar system in a variety of environments. After all the rationale of using such a simple measure as detection range is to be able to perform quick and simple calculations to scope the problem. There are no 'laws' for what to do – the onus is on the analyst to be aware of the potential biasing factors and act accordingly.

7.2 Simulation Modelling

Detection ranges achieved in Monte-Carlo simulation models of ASW engagements can often be significantly greater than those predicted in performance calculations. This is due to a combination of effects, based on sonar variability, the interaction of sonar channels (of both searcher and target) and on the tactics adopted in the engagement. These result in a filtering effect that tends to favour the occurrence of detection ranges greater than the predicted median. This does not mean that the simulation results are incorrect; it means that these effects must be understood if the results are to be interpreted correctly. The consequences of incorrectly interpreting the results can be significant. An example being the use of achieved sonar detection ranges as a guide to ranges at which a submarine may be safe from enemy attack or a range from which an attack could be launched, with consequences for tactical development or in setting weapon range requirements.

The effect of variability is perhaps the simplest to visualise. Consider a searcher with a single sonar channel searching for a non-reactive target (with no detection capability). Medium term variability is usually represented in the simulation by sampling from a distribution of detection ranges for each replication in the simulation, so allowing for sampling error 50% of replications will have detection ranges greater than the predicted median (and 50% will have less). If the search were to last an infinite time in a finite area, all searchers would detect targets and the overall achieved detection range would be similar to that predicted (provided the median is close to the mean). However, this approach is unsatisfactory, in that a) it does not represent any operational situation, in which the search time will be constrained by the tempo of the operation, and b) achieving 100% success does not yield a meaningful measure of effectiveness to discriminate between candidate searcher systems. Typically a search duration of a few days will be modelled, with area sizes that yield a search effectiveness of somewhere between 20% and 80%. Experience shows this to be aligned with typical operational situations and can yield a useful measure of effectiveness: the probability of making a detection after a finite search time. However, the achieved detection range will be greater than the predicted one simply because detections are more likely in replications with the longer detection ranges. Modelling studies show that this can be significant: medium term variability of typically of up to 8dB (+/-1s.d, assuming a normal distribution) in the sonar equation can result in a wide spread of sonar detection ranges, especially for low frequency narrowband sonars. Figure 2 shows how during a typical area search, the achieved detection range varies with the search time. For short search periods, with low probabilities of detection the achieved detection range is significantly greater than the predicted median (here 10km). As the probability of detection approaches unity, the achieved detection range falls to a value close to that predicted.

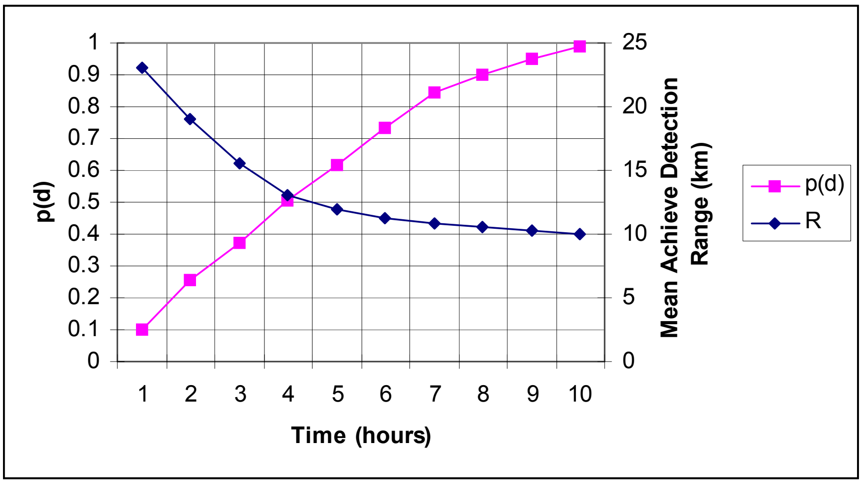

Figure 2: The variation of achieved detection time with search duration

Instead of a single channel, suppose the searcher had several independent channels, all subject to medium term variability. Clearly if one channel were dominant, then this would revert to the single channel case. For simplicity let us assume that a searcher had a number of identical independent sonar channels, each with the same signal excess. In the absence of variability, the number of channels has no effect on the predicted detection range, but in the presence of variability the signal excess of the channels taken together can be significantly higher than for an individual channel. This is illustrated in figure 3.

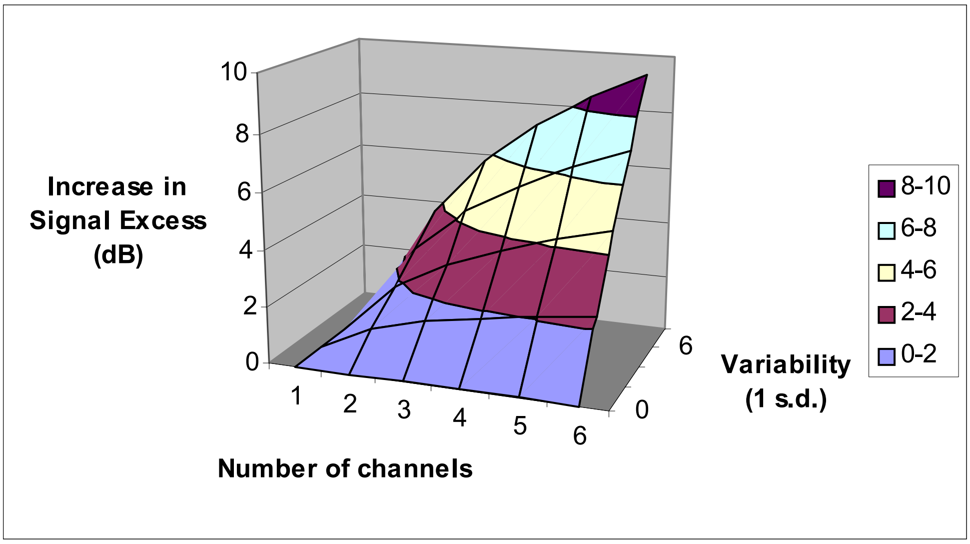

Figure 3: Effect of additional sonar channels on sonar signal excess

Figure 3 shows how additional channels on the searcher can increase the sonar performance (SE) and hence the searcher's detection range. A variation on this is the effect of the target's sonar. So far the target has been assumed to have no effective sonar. Assume the target has a sonar capability similar to that of the searcher, and adopted a tactic of evading every time it detected the searcher. In each replication both the searcher and target have detection ranges sampled from range distributions determined by their respective medium term variabilities. When the searcher's detection range is greater than that of the target, the searcher has the potential to achieve a detection, but when target's detection range is greater than that of the searcher, the target evades and no detection is possible. This means that the searcher's achieved detection ranges are dependent on the sonar performance of the target, and on the tactics adopted by the target once it makes a detection. This in itself will tend to give searcher detection ranges greater than the predicted median for the searcher (searchers with low detection ranges will never make a detection – the target will have evaded). Perhaps more interesting is that by increasing the detection performance of the target, the achieved detection range of the searcher is also increased. The target makes more detections and evades, leaving only those searchers with the longest detection ranges with the ability to achieve sonar detections. The number of searchers that detect decreases, but their average achieved detection range goes up. If the target had adopted different tactics, say either to remain unresponsive or to approach the searcher, this effect would not be manifest.

In summary, simulation modelling has the potential to yield vast amounts of information regarding sonar search effectiveness in operational situations, but the analyst should be aware of the effect sonar variability and the characteristics of the engagement, such as search time and target tactics can have on achieved sonar detection performance