We are IntechOpen, the world's leading publisher of Open Access books Built by scientists, for scientists

6,300

Open access books available

172,000

International authors and editors

190M

Downloads

Our authors are among the

154

Countries delivered to

TOP 1%

most cited scientists

12.2%

Contributors from top 500 universities

WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?
Contact book.department@intechopen.com

Numbers displayed above are based on latest data collected.
For more information visit www.intechopen.com

HF Radar Network Design for Remote Sensing of the South China Sea

S. J. Anderson

Additional information is available at the end of the chapter

<http://dx.doi.org/10.5772/57599>

1. Introduction

HF surface wave radar (HFSWR) is a highly cost-effective technology for remote sensing of ocean surface conditions and monitoring of ship traffic; several hundred radars of this type are in operation around the world. While an individual radar, operating alone, is able to provide a great deal of useful information, the integration of multiple radars into a network results in a system capability which is far more than the sum of its parts. For example, an estimate of a ship's velocity vector can be obtained in seconds, not tens of minutes or hours as is the case with a single radar. As another example, ocean currents can be estimated unambiguously, even in the presence of eddies and upwelling. Apart from these well-known considerations, there is a class of benefits which has special significance for long range HFSWR systems, namely, the potential for bistatic operations. As shown later in this paper, the fusion of monostatic and bistatic measurements enhances radar performance in a number of ways, a gain which is especially important for very long range operations.

While some HFSWR systems have been designed and deployed with a single mission in mind, it is increasingly recognised that the versatility of this technology supports a variety of applications. For instance, one might wish to detect and track shipping but also to measure surface currents so that risks of collision or grounding can be minimised and any transport of pollution predicted. In addition, information on sea state is of considerable economic value for ship routing, planning for offshore wave energy extraction, coastal development, port operation scheduling, search and rescue, fishing, tourism and recreational activities, so extraction and dissemination of environmental data would be welcomed by a wide range of user communities. Of course, these various applications will have relative priorities which vary with location, time of day and season, as will the radar's ability to accomplish them.

The physics of HFSWR dictates that many site-dependent factors contribute to the accuracy, reliability and availability of the various radar products. Moreover, the sensitivity to network configuration varies according to the type of measurement (or ‘mission’) being undertaken. Thus the choice of geographical sites which together comprise the network must reflect not only the family of radar outputs required but also their relative priorities. The resulting optimisation problem is extremely challenging.

An effective methodology for optimising HFSWR network design for the case where multiple missions must be addressed has been developed recently and demonstrated in the context of a hypothetical two-radar system deployed in the Strait of Malacca (Anderson, 2013). The results of that study demonstrated that quite disparate criteria can be accommodated within a genetic algorithm framework and confirmed that the method yielded the true optimum site configurations. Yet that study left a key question unanswered. In practice we are unlikely to be satisfied knowing that our choice of sites is the best for a given budget; we want to know that the network will meet prescribed levels of performance. This could well mean that, in a particular situation, a mix of quite different radar types would be required, adding another dimension to the network design problem.

In this chapter we review the genetic algorithm methodology for multi-objective optimisation in the HFSWR context, and show how it can be extended to handle the inverse problem of designing networks to meet specified performance levels. In order to illustrate the steps involved in formulating and applying the methodology, the discussion is framed in the context of a specific scenario: the design of an HFSWR network for providing surveillance and remote sensing of the South China Sea.

The spatial resolution and ultimate sensitivity of HFSWR is primarily a function of radar design, but performance in its various candidate roles is also dependent on a wide variety of geophysical factors, lithospheric, oceanic, atmospheric and ionospheric. Further, the relative priority of different missions reflects economic, geopolitical and strategic considerations. As all these aspects would (or should) be taken into account by network designers, it is appropriate to examine ways in which they can be incorporated in the objective functions employed for optimisation. In the following section we set the scene for the subsequent analysis by reviewing the physical environment and the associated human activities which an HFSWR radar network might be expected to monitor. Next we outline the capabilities and limitations of HFSWR in this context, based on the nominal performance of four existing radar systems. Once the radar capabilities have been established, we turn to the central issue, namely, that of formulating the network design problem in mathematical terms, which leads us to focus on evolutionary algorithms of nonlinear optimisation. Here the genetic algorithm approach of Anderson (2013) is emphasised, as it lends itself naturally to multi-objective optimisation, though in order to handle the enormous computational burden in the present case, a recently-reported convergence acceleration technique (Anderson et al, 2013) is introduced. We proceed to describe practical methods for constructing chromosomes and objective functions for a number of missions, illustrating these by relating them to the South China Sea context.

2. The South China Sea

2.1. Physical geography

Formally the South China Sea extends from Bangka Belitung, between Sumatera and Borneo, to the northern extremity of Taiwan, and from the Gulf of Thailand to the Philippines, as shown in Figure 1.

Within its area of some 3,500,000 square kilometres lie several hundred islands, of which most are grouped into two clusters, the Paracel and Spratly Island chains. A great many of the islands are little more than exposed reefs and even the important Spratly Island group has a total land area of less than 5 square kilometres and a maximum elevation above sea level of only 4 metres. Some important features are entirely submerged, as is the case with Macclesfield Bank – actually an atoll-which is on average about 10 m below sea level, yet has an area of some 6500 square kilometres. Scarborough Shoal (*aka* Panatag Shoal), has reefs and small islets above water amounting to only a few hectares, surmounting an area of some 150 square kilometres of about 15 metres depth. beyond which the sea floor drops away rapidly to a depth of several kilometres. Only a handful of islands are large enough to be home to an airstrip and some of these facilities may not survive even a modest rise in sea level.

The large-scale bathymetry of the South China Sea is particularly striking. South of a line joining Brunei to the southern tip of Vietnam, the depth is less than 100 metres, but north of that line the sea floor descends rapidly to 1000 – 5000 metres, except around the island chains and atolls. With the exception of a few narrow but deep channels between Luzon and Taiwan, connecting to the East China Sea, the South China Sea is essentially a basin.

Ocean surface conditions are influenced by the orography of adjacent land masses which helps steer the prevailing winds. In the case of the South China Sea the principal land feature that is relevant to HF radar system performance is the mountain range along almost the entire coast of Vietnam.

2.2. Meteorology and oceanography

The wind regime over the South China Sea is dominated by the monsoon winds, punctuated by mesoscale systems such as tropical cyclones. During the boreal winter, the northeasterly winter monsoon winds impose a fairly uniform stress over most of the South China Sea, whereas in summer, June, July and August, the southwesterly monsoon winds show somewhat greater spatial variability, especially south of about 6°N. Average wind speeds in winter tend to fall in the range 8 – 12 m/s whereas the southwesterly summer monsoon winds are typically approximate 6 – 8 m/s in the Southern SCS and somewhat less in the northern SCS. Highly variable winds and surface currents are observed during the transitional periods. Moreover, synoptic systems often pass by the SCS and causes temporally and spatially varying wind fields. Severe weather most frequently takes the form of an increase in the strength of the prevailing monsoon winds or as meso-scale disturbances concentrated in either of two regions: a localised area east of the southern part of Vietnam, centred on 10°N, 110°N, and the band between Luzon and southern China. The mean wind regimes for summer and winter are shown in Figure 2.

Figure 1. The bathymetry of the South China Sea (adapted from the World Data System for Marine Environmental Sciences, <http://www.wdc-mare.org/>)

Tropical cyclones form in the waters between 12°N and 24°N , usually making landfall over Hong Kong and southern China, the north and central coasts of Vietnam or the northern Philippines. The most severe cyclones occur to the east of the Philippines and Taiwan, as shown in Figure 3 but, even so, the South China Sea north of 15°N is occasionally subjected to category 3 and 4 events.

Wave and current distributions due to the wind forcing are less uniform than the wind fields. Significant waveheight (SWH) distributions are higher in the northern and central SCS (north of 10°N) than in the southern SCS (south of 10°N) with upper quartile values exceeding 2.25 m. (The Wavewatch III model has been found to yield fairly accurate results (Chu et al, 2004), so serves as a useful adjunct in modelling radar performance.) As shown in Figure 4, the orientation of the high SWH region coincides with the orientation of the monsoon winds.

Figure 2. Synoptic-scale wind patterns during the summer and winter monsoon seasons (from Chu et al, 2003)

Figure 3. The distribution of tropical cyclones over the period 1945 – 2006.

Figure 4. Seasonal variation of mean wave height for the period 1979-2009, WaveWatch III hindcast (Mirzaei et al, 2013). Note that the dominant wave direction is aligned with the monsoon winds, which is to say southward-propagating in DJF and northward-propagating in JJA.

The prevailing winds have a direct effect on the surface water currents of the shelf region. The current speeds are about 0.6 knots to the SW during the winter monsoon. They change to 0.2 to 0.4 knots to the NE during the summer monsoon. Stronger currents flow adjacent to the Vietnamese coast in particular, attaining speeds in excess of 1 m/s, while the islands of the Spratly archipelago can induce fairly complex local current variations. Primary or climatological current patterns in summer and winter are shown in Figure 5, but these convey an incomplete picture of the flow field. To gain a better understanding of the complexity of the current distribution, consider Figure 6, which shows the outputs of a detailed hydrodynamic model of the current field in the northern and central South China Sea. A key feature of this model is the inclusion of the wind-induced current, which was found to dominate the

geostrophic current in many places. Moreover, note the appearance of the mesoscale eddies. These have been validated by observation. In a similar vein, Marghany (2009, 2011, 2012) has shown how local current patterns can be extracted from spaceborne SAR using observations off the east coast of peninsular Malaysia. The lesson to be drawn from this kind of modelling is that the flow field has significant structure on length scales of 50 km or less; given that the cross-range dimension of a typical HF radar resolution cell at long range may approach this magnitude, it is evident that HFSWR could provide unique validation data, though conventionally measured Doppler spectra will not always have discrete Doppler shifts and hence current velocity estimation will be compromised on those occasions.

One particular form of current perturbation which has received a lot of attention by the HF radar community is that associated with a tsunami (Lipa et al, 2012). It has been demonstrated that HFSWR is an effective tool for early warning of tsunamis provided that the bathymetry is favourable, which is to say reasonably shallow so that the speed of the tsunami is much reduced from its high deep water value. As Figure 1 shows, the southern part of the sea occupied by the Sunda Shelf and the north-western margins certainly satisfy this requirement.

Figure 5. Primary currents during the summer and winter monsoon periods (Chen et al, 1985)

The SW winds blowing along the SW-to-NE part of continental shelf may induce upwelling during the summer, bring nutrients to the eutrophic zone on the outer portion of the shelf and, enhance primary production of the waters (Wang and Kester, 1988). The seasonal stratification stimulates the seasonal changes in primary production and nutrient cycling, with a strong signature evident in the high chlorophyll distributions in two coastal upwelling regions: the northwestern Luzon in winter and the eastern coast of Vietnam in summer. Mesoscale eddies provide another mechanism responsible for seasonal and interannual variability of the surface chlorophyll distribution.

from 33.2 to 34.5 PSU. In summer, due to the influence of southwest monsoon, the surface temperature is generally 28-29°C and the salinity is low – near 32 PSU – in the north and south, and high – about 33.6 PSU – in the central region. In summer, the SW monsoon results in the large increase in rainfall and river discharges. This results in the reduction of salinity in the coastal waters and the production of seasonal pycnoclines. Particularly low salinity occurs off the east coast of peninsular Malaysia.

The diurnal and semi-diurnal tides are of about equal magnitude in the South China Sea, though the latter is more effective at generating internal waves. These are exceptionally strong in the northern region, where they are generated in the Luzon Strait before propagating westward. Amplitudes reach 200 m with horizontal scales upwards of 200 km. These internal waves take about 4 days to cross the South China Sea, modulating the surface gravity wave field as they progress and hence influencing the radar scattering properties of the sea surface.

2.3. Shipping

The volume of shipping activity in the South China Sea can be illustrated by a few key statistics:

- i. nearly half the world's annual merchant fleet tonnage moves through its waters, carrying commodities valued at over \$5 trillion
- ii. one third of global oil tanker traffic and over half of global LNG traffic crosses the South China Sea, most from the Strait of Malacca but the very largest supertankers via the Sulu Sea
- iii. ore carriers, predominantly from Australia, transport roughly half a billion tonnes of iron ore and a similar amount of coal through the South China Sea annually
- iv. six of the world's ten largest ports lie on the coastlines of the South China Sea
- v. the annual growth rate for liquid petroleum fuels consumption in recipient countries – mainly China and Japan and South Korea – is presently 2.6 %, while that for natural gas is 3.9 %. For Australian minerals the figure is 4.6%
- vi. over half a billion people live within 100 miles of its margins
- vii. perhaps as many as 18,000 small fishing boats ply its waters

The major shipping routes are shown on Figure 7, using data derived from Wang et al, 2013).

Given the density of traffic, it is perhaps not surprising that shipping hazards in the South China Sea continue to take a toll on vessels in transit, as exemplified by several recent incidents: the sinking of the Bright Ruby (severe storm, November 2011), Royal Prime (hit reef, and sank, December 2012), Harita Bauxite (sank after engine failure, February 2013), Jung Soon (sank after hull failure, September 2013). Another form of hazard is piracy, for which the South China Sea was once notorious. While less frequent than a few years ago, hijacking and armed robbery remain a significant threat in some waters. Mimicking the 'mothership' refuelling station tactic used by pirates off the coast of Somalia, pirates in Indonesia and Malaysia tend to camp on a small island near to narrow shipping lanes and launch their strikes from there. Pirates in South

Figure 7. Principal shipping lanes through the South China Sea, and islands selected for use in radar mission definitions.

East Asia also tend to launch their attacks at night, which makes it much harder for ship captains to spot them coming. Between 2008 and 2010, 57 incidents of ‘cluster piracy’ took place around the Abambas / Natuna/ Tambalan corridor. In the first six months of 2013, attacks involving pirates boarding vessels and assaulting the crew were recorded in the Singapore Straits, in Malaysian waters, in the Straits of Malacca and in the Philippines. Within the main body of the South China Sea, 2013 escaped serious incident.

Another consequence of heavy ship traffic is oil pollution, both accidental and deliberate, such as that caused by tankers flushing their tanks on the voyage back to the Middle East. Offshore facilities and undersea pipelines are other man-made sources.

2.4. Economic activity

It is evident from the preceding discussion that ‘through traffic’ is critically important to the destination countries of China, Japan and South Korea, but, for the littoral states around the South China Sea, fishing is the most vital maritime activity, as it has been for centuries. Fish protein constitutes nearly a quarter of the average Asian diet and demand continues to grow

strongly. Thus, whereas the extent of oil and gas reserves beneath the South China Sea remains questionable, the value and importance of its fisheries and aquaculture is not in doubt. It is therefore of great concern that the relative stability of traditional fishing practices is now threatened by over-fishing, together with rising water temperatures which appear to be resulting in migration of fish populations, primarily further north. These developments are stoking tensions between the countries whose populations depend on accessible and reliable stocks.

It is widely reported that the South China Sea holds immense untapped natural reserves of oil and gas, and that the contested ownership of the Spratly Islands and other parts of the sea is primarily a fight for these resources. It is certainly the case that confrontation and armed skirmishes have taken place where exploration has been pursued in disputed waters. Yet a considered analysis does not support the more extreme assertions regarding the magnitude of the reserves in the contested regions. The most recent assessment by the US Energy Information Administration estimates that the total of proven and probable reserves in South China Sea amounts to approximately only 11 billion barrels of oil and 190 trillion cubic feet of natural gas. Another US expert source places the figure for oil at 2.5 billion barrels. Allowing for additional reservoirs in under-explored areas, the EIA says, could add between 5 and 22 billion barrels of oil and 70 to 290 trillion cubic feet of gas. These figures contrast with those of the Chinese National Offshore Oil Company which estimates undiscovered reserves amount to 125 billion barrels of oil and 500 trillion cubic feet of natural gas. In the absence of detailed prospecting, the actual quantities cannot be known with any certainty. What is undeniable is that the preponderance of known resources resides in the uncontested areas close to the coasts of the surrounding countries, especially Vietnam, Malaysia and Brunei. Thus the fierce competition for control, if not ownership, of the islands, reefs and shoals of the South China Sea, is probably driven by a combination of factors, economic, political and strategic.

We note too that Malaysian researchers have identified potentially valuable elements in seabed sediment, including manganese, zinc, chromium, lead, copper and aluminium.

2.5. Strategic and geopolitical issues

The South China Sea has a long history of tension and conflict. Much of this derives from overlapping territorial claims and disputed ownership of maritime features, as indicated in Figures 7 and 8, exacerbated by a race to exploit the maritime zone's natural resources. In addition, there is also a growing element of overt strategic rivalry and nationalism, which poses a substantial risk to regional security and prosperity. Specific areas of dispute include:

- the Spratly Islands, disputed between the People's Republic of China, the Republic of China, and Vietnam, with Malaysia, Brunei, and the Philippines claiming part of the archipelago
- the Paracel Islands, disputed between the People's Republic of China, the Republic of China, and Vietnam
- the Pratas Islands, disputed between the People's Republic of China and the Republic of China

- the Macclesfield Bank, disputed between the People's Republic of China, the Republic of China, the Philippines, and Vietnam
- the Scarborough Shoal, disputed between the People's Republic of China, the Philippines, and the Republic of China.

Figure 8. Territorial claims over the South China Sea, together with occupied islands.

Many detailed discussions of these issues can be found in the open literature (International Crisis Group 2012a, 2012b); here it suffices to make the point that timely, comprehensive, robust and persistent surveillance can be a useful means of establishing trust and defusing incidents which could spiral out of control.

3. Capabilities and limitations of HF surface wave radar

3.1. The suitability of HFSWR for maritime remote sensing and surveillance

The physical quantities which impact directly on HF radar capability in both its remote sensing and surveillance roles are (i) water electrical conductivity, (ii) surface currents, and (iii) the geometry and dynamics of the sea surface, usually represented as a spectrum of surface gravity waves. It is no exaggeration to state that radar performance in any of its missions is highly dependent on these primary quantities. Moreover, the primary quantities are coupled with other geophysical variables and processes, as illustrated in Figure 10. Therefore, as part of the

network design procedure, it is absolutely essential to take into account the kind of information presented in Section 2. Figure 10 shows the linkages between the primary quantities and other geophysical variables and processes. From this figure it is apparent that, by appropriate analysis and interpretation of the radar echoes, it may be possible in some circumstances to use the primary measurements of currents and wave spectra to infer secondary phenomena, surface winds being the prime example.

Figure 9. The relationships between the geophysical parameters and phenomena which impact on HF radar performance

It is evident from the discussion in Section 2 that real-time monitoring of these environmental conditions and ship traffic over the South China Sea could have substantial value for a wide range of users. Yet the practical application of any remote sensing technology requires that we first establish whether the coverage, resolution and accuracy of the measurements are commensurate with the needs of the users. To illustrate this step, we shall consider the nominal performance of four well-known HFSWR products.

The commercial marketplace for so-called ‘oceanographic’ HFSWR systems is dominated by two manufacturers: CODAR, with its Seasonde radars (Barrick, 1998), and Helzel Messtechnik, with its WERA systems (Helzel et al, 2010). These radars each cost in the vicinity of 0.5M\$ per system consisting of one transmitting station and one receiving station, and have excellent track records for delivering ocean current information at ranges out to 200 km, with sea state measurements available for significantly shorter ranges. While some extravagant claims are made about the ability of these low power radars to detect ship targets at ranges of several hundreds of kilometres, experience has tended to show that reliable detection is confined to 50 – 150 km, depending on target type, time of day, and other factors. These radars have quite

different characteristics so it is appropriate to include them both in the list of options for a heterogeneous network.

The South China Sea is roughly 1000 km east-to-west at 10° N so it is clear that oceanographic radars based on undisputed territories are not able to provide comprehensive surveillance. The possibility of deploying radar systems on small islands may change this assessment somewhat, as discussed later in this paper, but a priori it would seem that radars with far superior long range performance are required if comprehensive surveillance is to be achieved. For the study reported here, two commercial HFSWR radars were chosen to represent such ‘military-class’ systems: Raytheon’s SWR-503 (Ponsford, 2012) and Daronmont Technology’s SECAR radar (Anderson et al, 2003). Each of these systems has demonstrated ship detection at ranges well in excess of 400 km. It is important to point out that the diverse observations and opinions from which these figures were inferred correspond to a variety of environmental conditions, so the estimates are really just indicative. Still, the table serves to provide numerical values for the purpose of exercising the network optimisation suite.

Observable	Typical performance			
	low-cost civilian radar		military radar	
	max. range (km)	accuracy	max. range (km)	accuracy
surface current	60 - 200	± 0.02 – 0.20 m/s	350 - 450	± 0.02 – 0.10 m/s
wave height	30 - 100	± 10 – 25 %	150 - 350	± 10 – 20 %
wind direction	50 - 180	± 30° - 60°	320 - 400	± 20° - 30°
wind speed	30 - 150	± 20 %	150 - 350	± 20 %
large ship	50 - 180	± 0.5 - 3 km	300 - 450	± 0.5 – 3 km
fishing boat	20 - 65	± 0.5 – 2 km	120 - 280	± 0.5 – 2 km
small boat	10 - 45	± 0.5 – 1 km	70 - 150	± 0.5 – 1 km

Table 1. Capabilities of selected HFSWR systems, expressed in terms of typical maximum ranges at which measurements can be made reliably

A natural first test is to see whether even the most potent (and expensive) network could deliver the desired coverage. In order to obtain a large and realistic set of possible sites in the present case, we visually searched the coastlines around the South China Sea as presented in Google Earth™, selecting as candidate locations all those places characterised by reasonably flat, low-lying ground with linear sea frontage in excess of 300 m. These criteria were applied to ensure a choice of radar type at every location; only the CODAR Seasonde is able to be deployed on almost any topography. Some 141 sites emerged from this procedure; they are marked on Figure 10.

Figure 10. Maps of the South China Sea showing the major shipping lanes, candidate radar sites (blue dots), the associated potential radar coverage (yellow sectors), the discrete points in the sea area at which objective functions can be evaluated (magenta dots), and selected islands of interest (dots, various colours). Figure 10a shows radar coverage to 200 km, Figure 10b to 400 km.

Figure 10a shows the nominal current measurement coverage for oceanographic-class radars deployed at each of these sites, while Figure 10b shows the corresponding information for a military-class radar. These figures reveal that no fully-compliant solution exists, even with the maximal deployment of radars, but they suggest that a solution employing a combination of radar types, at a suitable subset of sites, might achieve an acceptable outcome leaving relatively few areas unsurveyed for this particular mission.

Regarding spatial resolution, while the cross-range dimension of a cell at nominal maximum range exceeds the along-range dimension by up to an order of magnitude, the broad features of oceanographic fields remain distinguishable, and discrete ship echoes can be finely resolved in the Doppler domain, so HFSWR is certainly able to provide the required detail for most objectives. The accuracy of measurements is limited not so much by radar design as by the intrinsic spatial and temporal variability of natural phenomena; the widespread acceptance of HFSWR remote sensing products confirms that the information is of adequate fidelity.

3.2. Performance limitations and constraints

It is helpful to be aware of the factors which limit HFSWR performance, as radar and network design can be adapted to minimise the deleterious effects of some of them. First there is the nature of surface wave propagation, which results in increasingly rapid signal decay as one moves beyond about 50 km range, with higher frequencies decaying much more quickly. Second, there is the frequency dependence of the radar signatures of both ship targets and the ocean surface, which have complicated forms that jointly have a strong influence on radar performance. Third, there is the external HF noise from lightning and man-made emissions,

which almost always defines the noise floor against which radar echoes of interest must compete. External noise is highly dependent on time of day. Fourth, there are Doppler-spread echoes from the ionosphere, which have a complex spatial and temporal pattern of occurrence and can mask echoes of interest.

Measures which can be taken to mitigate these factors include antenna design, advanced signal processing, frequency agility and, of special relevance to the present study, siting relative to the locations and velocities of the phenomena under observation. As a simple example, Figure 11 shows the obscuration of ship echoes due to sea clutter, plotted in Doppler space.

Figure 11. Blind speeds for various ship types, against a specific sea clutter spectrum, for a 32 second integration time

4. Radar siting and configuration design as a multi-objective optimisation problem

4.1. Elements of the formulation

At the outset we need to identify the data structures, procedures and supporting information that need to be integrated into the problem formulation. These are:

- i. the parameter space P in which the solutions must lie. A particular solution $x^{mn} \in X^{mn}$ will have a fixed number m of transmitting systems, each specified by location, orientation and design, together with a fixed number n of receiving systems, each similarly described by its location, orientation and design. The ability of receiving system p to acquire and process signals from transmitting system q is represented by a coupling matrix C^{mn} :

$$C_{pq}^{mn} = \begin{cases} 1 & \text{if receiver } p \text{ can process signals from transmitter } q \\ 0 & \text{otherwise} \end{cases}$$

Usually compatibility demands that the systems belong to the same product family. Thus, for example, a Seasonde receiver can process signals from a Seasonde transmitter, but not signals from a WERA transmitter. In the problem under consideration, we do not know a priori what the network membership numbers m and n should be; accordingly we define P as the disjoint union of the X^{mn} ,

$$P = \bigcup_{n=1}^N \bigcup_{m=1}^M X^{mn}$$

where M and N are upper bounds on the numbers of radar transmit and receive sites.

A solution x thus can be written in the form of a pair of two-dimensional arrays,

$$x \triangleq [\{R_j^{lat}, R_j^{lon}, \Psi_j, \varphi_j\}_{j=1,N}, \{T_k^{lat}, T_k^{lon}, \Psi_k, \varphi_k\}_{k=1,M}]$$

where the dimensions correspond to parameter type (latitude, longitude, radar class, orientation) and parameter index (labelling the set of Tx/Rx sites which make up the configuration)

- ii. the specific coordinates of feasible sites. These constitute a subset of the set of points C which comprise the coastlines which border the South China Sea or, in the case of radar sites on very small islands, the nominal location at which the installation is most feasible
- iii. the amenity of each location to the installation of a radar, taking into account factors such as accessibility, power supply, field of view and environmental impact
- iv. the range and azimuthal coverage of the individual radars to be used
- v. the wind, wave and current climatology of the waters of the South China Sea
- vi. the recognised shipping lanes
- vii. the types of vessels of interest and their typical speeds and radar cross sections
- viii. the surveillance and remote sensing missions assigned to the radar system and the associated performance thresholds which must be exceeded (at least in a statistical sense)
- ix. algorithms which compute the radar network response for any given combination of ship type, course, speed and environmental conditions

- x. the objective function space Y , that is, the k -dimensional space whose coordinates measure the radar performance against the k tasks assigned to the radar
- xi. a search algorithm which finds the extrema of a scalar function over a specified domain
- xii. a criterion for ranking solutions which achieve extrema in one or more coordinates of Y

It is common practice to formulate optimisation problems in terms of minimising the objective functions rather than maximising them, which is trivially achieved by redefining the coordinates of Y ; we shall follow this practice.

4.2. Criteria for network optimality

With this palette of ingredients, various radar siting problems with cost and performance constraints can be formulated. Three of the most important are:

- find the solution $x \equiv x^{min} \in X$ which maximises performance against a specific task for a given cost
- find the minimum cost solution $x \equiv x^{min} \in X$ which exceeds a specified threshold of performance
- given an existing deployment x_{old} of a number of radars illuminating parts of an area of interest, find the solution x_{new} such that the augmented network $x_{aug} \equiv x_{old} \cup x_{new}$ meets some specified performance/cost criterion

but there are many other possibilities. We observe that some of these can be expressed as inverse problems with the threshold vector taking the role of the data vector.

4.3. Multi-objective optimisation via Pareto dominance

The definition of the problem given above is in one sense incomplete – it does not specify the choice of norm for the space Y . In a single objective optimisation problem, the objective space is usually a subset of the real numbers and a solution $x_1 \in P$ is better than another solution $x_2 \in P$ if $y_1 < y_2$ where $y_1 = \mu(x_1)$ and $y_2 = \mu(x_2)$. In the case of a vector-valued objective function mapping, comparing solutions is more complex and one must endeavour to capture the essential priorities of the problem in the choice of norm. Herein lies the crucial distinction between single objective and multi-objective problems-whereas the former afford simple scalar measures of fitness that can be used to rank individual members of the design space, the latter are characterised by conflicts of interest among the competing objectives as measured by μ_i , $i=1, m$.

There are several ways to deal with this complication. Perhaps the simplest is to create a scalar figure of merit as a weighted sum of the separate objective measures,

- i. minimize $\mu^{(1)} = \sum_{i=1}^m \alpha_i \mu_i$

Another approach is to convert all but one of the objectives into constraints,

- ii. minimise μ_j subject to $\mu_i \leq z_i \quad \forall i=1, m; i \neq j$

While convenient, these methods shed little light on the nature of the trade-offs made. As there may be subtle, non-quantifiable considerations involved in site selection, such as risks to personnel or to equipment, a better approach is to map the trade-off surface so that the decision maker can execute judgment in making a final selection. To perform this mapping, it is not necessary to run (i) or (ii) above for a large number of parameter selections α_i, z_i and to inspect the outcomes. Instead, we can use an evolutionary stochastic optimisation algorithm to reveal the Pareto front, as described below.

Pareto optimality is based on the binary relation of dominance. A solution $x_1 \in X$ is said to be dominated by another solution $x_2 \in X$, written $x_2 \prec x_1$, if x_2 is at least as good on all counts (objectives) and better on at least one, that is,

$$\mu_i(x_2) \leq \mu_i(x_1) \quad \forall i=1, m \text{ and } \mu_j(x_2) < \mu_j(x_1) \text{ for some } j.$$

With this relation, the Pareto set of optimal (non-dominated) solutions P^* will usually have multiple entries, associated with different trade-offs between the objectives. The image $Y^* \subset Y$ of the Pareto set $P^* \subset P$ is referred to as the Pareto front and knowledge of its shape greatly assists in choosing the best compromise solution.

4.4. Implementation via genetic algorithms

Classical techniques for finding extrema of functions defined on prescribed domains rely, in most cases, on gradient search methodologies. Such techniques are vulnerable to being trapped on local extrema, rather than the global extremum of main interest. In addition, the convergence may be slow, especially near the extrema, necessitating the invocation of higher-order derivatives. While there are ways to alleviate these weaknesses, they come at considerable cost. An alternative approach, now in widespread use, is to emulate evolutionary mechanisms which we observe in action in the natural world. The best known of these evolutionary optimisation techniques are genetic algorithms.

Genetic algorithms encode the parameter values associated with each candidate solution as a string, usually in binary format. For each parameter, the number of bits provided must be sufficient to encode the full range of possible values associated with that parameter. The string representing a solution is simply the concatenation of the sub-strings corresponding to the individual parameters; by analogy with biology, this string is referred to as a chromosome. Starting with an initial population of candidate solutions (ie, chromosomes) constructed by means of a random number generator, a genetic algorithm iteratively applies three basic steps: (i) rank the members of the current population according to fitness, (ii) select superior members which will be used to breed the next generation, and (iii) apply operators on randomly-selected pairs of these members to mimic the transfer of genetic material to offspring that occurs during biological reproduction, thereby producing a new generation with statistically superior characteristics.

A common mechanism for the transfer of information from one generation to the next is variable length cross-over. For each pair of chromosomes selected to breed together, the start and end indices of a sub-string are selected by a random number generator and the corresponding sub-strings are exchanged. The excisions are not forced to align with the parameter sub-string boundaries. The offspring of this coupling have parts in common with each parent, and in general will represent new solutions. A small fraction of this new set of chromosomes is then subjected to mutation, that is, one or two bits may be flipped to produce a different string, which of course maps onto a different candidate solution. This completes the process of constructing a new generation.

With single objective optimisation, it is a simple matter to rank the members of the resulting population so that selection of candidates for constructing the next generation can proceed. Chromosomes representing the best solutions are carried over unchanged to the next generation, as well as participating in the breeding cycle, while the least fit are discarded. The resulting population is then allowed to breed in its turn, via cross-over and mutation. After passing through a large number of generations, the population tends to converge towards a uniform composition whose members share the most desirable parameter values. Importantly, by virtue of the randomness of the cross-over and mutation operations, candidate solutions from all over the solution domain are potentially represented, and mutation ensures that this property is maintained, so that the population is unlikely to be trapped on a local extremum if a superior solution exists.

With multi-objective optimisation, the key objective is to find the Pareto front, but experience has shown that coverage and convergence can be improved by relying on more than just Pareto dominance for selection. In our approach, each chromosome was tested against its contemporaries and those which were Pareto dominant were automatically selected, while those which had only one or two dominators were also short-listed. In addition, members that performed particularly well against just one objective function were retained. Supplementing these criteria, a scalar figure of merit was defined by taking the product of the individual objective functions; this provided another metric for selection. The total size of the population was maintained at the initial value by allowing each of these different selection mechanisms to contribute a fraction of the membership, with the relative proportions changing with the generation index. We modified the single objective genetic algorithm developed by Anderson (2013) to embody these ideas and hence to compute an estimate of the Pareto front.

4.5. Acceleration techniques

Genetic algorithms tend to be computationally expensive, so special techniques continue to be developed to accelerate convergence. Some methods which have proven efficacious are:

- eugenics – a recent hybrid scheme which combines the virtues of GA with a very efficient gradient search
- class identifiers – partitioning chromosome space into dissimilar clusters and constraining cross-over to avoid in-breeding, thereby increasing and maintaining diversity

- smart seeds – using intuition, experience and common sense to insert some chromosomes with high potential

4.6. Methods for handling variable solution space dimensionality

One of the challenges of the general network optimisation problem is that, unlike the case in Anderson (2013), the number of radars is itself a variable. Given that the gene length for representing an individual radar is fixed, it follows that the minimum chromosome length will change. This introduces some very fundamental modifications to the elements of the GA algebra, so a number of approaches have been explored:

- loop through dimension index; a straight-forward extension of conventional GA structure
- set the chromosome length to the maximum number of radar sites considered feasible and work within this space; likely to be computationally expensive
- adopt a hierarchical scheme, with fixed length chromosomes containing genes serving as pointers to subspaces of different dimensionality; potentially effective but complicated
- employ variable length chromosomes; this requires a whole new class of genetic operators able to work with strings of different lengths

In the implementation we have used for the South China Sea example, the first of these options has been adopted.

4.7. Constructing the chromosomes

The chromosomes do not need to encode all the detailed information about site properties, radar characteristics, and so on. It is more efficient to use the genes as pointers to data files in which the numerical specifications are stored. In our illustrative example, we allow for four different radar types, so 2 bits are required for that purpose. Our survey of the coastlines of the South China Sea identified 141 candidate sites, so it might seem logical to allocate $\log_2 141$ bits to represent them. This causes a problem, as not all 8-bit strings correspond to radar sites, and the extra algorithmic structure that would be required to deal with this issue would arise in a section of the code which is run intensively. In the present case it is better to prune the set back to 128 sites, with minimal impact on the outcome, though conceivably another problem could justify increasing the state space to 256 sites. Thus our basic gene has $2 + \log_2 128 = 9$ bits. As it necessary to extract the separate radar-type and radar-site parameters, an efficient ‘gene scissors’ is required, easily implemented in Matlab.

The specific context imposes other constraints that need to be carefully considered. For instance, in the present network design study, we found candidate sites located on the mainlands of Vietnam, Malaysia, Brunei and the Philippines, as well as on several islands some of which are of disputed sovereignty. It may be that network operations embracing radars in all ASEAN nations could be negotiated, but such arrangements area never simple. The situation becomes even more complicated when we contemplate radars on those islands which presently are home to airstrips, ideal for basing array-type HFSWR systems, since islands

meeting that description are owned or occupied by China, Indonesia, Malaysia, the Philippines, Taiwan and Vietnam. In addition, a number of mostly submerged reefs and seamounts in the Spratly Islands bear constructions on which CODAR Seasonde radars could easily be fitted. All these possibilities need to be taken into account when proposing the extent of the solution space in which the set of optimal solutions is to be sought.

5. Constructing the objective function space

5.1. Objective functions for priority missions

While it is certainly possible to conceive of many useful missions which could be addressed by a network of HFSWR systems, it is generally the case that one focusses on those which have a high level of economic or geo-strategic relevance. As an example, the palette of tasks which one might wish to address could take the form:

- i. maintain surveillance around most, preferably all, of the important islands with the targets of interest being vessels of at least patrol boat size, typically 50 – 80 m in length and long endurance research ships of around 100-120 m length
- ii. provide full ocean current vector information over those parts of the South China Sea which are traversed by large vessels such as tankers and container ships; with emphasis on the major shipping routes
- iii. provide sea state information for the areas in which fishing fleets operate

It is readily seen that the spatial domains over which the performance of these three tasks is of interest are of different dimensionality. As the objective function used to define fitness for a given task involves integration over the corresponding domain, there is a strong link between task domain and computational load. The cases of most concern to the network optimisation problem under consideration are as follows:

Domain dimensionality	Examples
0	islands, shoals, offshore oil platforms,
1	shipping lanes, transects, sovereignty and EEZ boundaries
2	fishing grounds, oil exploration leases, wave energy surveys

For each of the designated tasks, it is necessary to define some criterion that quantifies performance and which can thus be used to govern the search for the Pareto optimum configurations. To illustrate, we shall outline the construction of objective functions – also known as fitness functions or figures of merit – for the first two tasks mentioned above.

5.2. Ship detection

Suppose an HF surface wave radar operating at a fixed centre frequency f is deployed with the goal of detecting ships whose radar cross section (RCS) exceeds some specified threshold. For detection we require that a ship echo exceed the clutter and noise power in the same Doppler bin by some margin ε , that is, there exists $\omega \in [-\Omega, \Omega]$ such that $s(\omega) > c(\omega) + n(\omega) + \varepsilon$ where $s(\omega)$, $c(\omega)$ and $n(\omega)$ are the target, clutter and noise power spectral densities respectively and $[-\Omega, \Omega]$ is the extent of the Doppler domain. At modest ranges, dependent on the radar type, the clutter power spectral density exceeds that of external noise, but at longer ranges external noise dominates and sets the detection limit. This we need to have a database which provides these distributions. From the description in Section 2 we know that the wind stress and hence the sea state is relatively constant over the South China Sea during each of the two monsoon periods, comprising some 80% of the year, so a reasonable approach is to compute clutter Doppler spectra for just these two sea states. In the context of large ships on the major lanes in the South China Sea, proceeding along known shipping lanes at fairly uniform speeds, $v \in [v_{min}, v_{max}]$, the Doppler perceived by a radar from a given ship is a function of a single coordinate, representing the ship's position along its chosen lane, since that determines the viewing geometry. Accordingly, for these targets it makes good sense to define a figure of merit which measures the fraction of the time (equivalently distance along the route under surveillance) for which such ships are detectable. In the case of small or medium size ships near particular islands or facilities, the direction of travel and the speed cannot be assumed, so the figure of merit should reflect the need to maximise detectability against all eventualities. Assuming a maximum speed v_{max} ,

$$OF_1 = \frac{1}{2\omega_m} \sum_{k=1}^n \int_{-\omega_m}^{\omega_m} H(s(\omega; r_k) - c(\omega; r_k) - n(\omega) - \varepsilon) d\omega \quad (1)$$

where $H(x)$ is the Heaviside function,

$$H(x) = \begin{cases} 0 & x < 0 \\ 1 & x \geq 0 \end{cases}$$

and

$$\omega_m = \frac{2v_{max}f}{c}$$

where f is the radar frequency and c the speed of light. The r_k are the coordinates of the discrete islands, offshore oil platforms or other discrete features of interest.

It is a computationally trivial but operationally useful generalisation to apply a priority weighting to the individual islands,

$$OF_2 = \frac{1}{2\omega_m} \left(\frac{1}{\sum_k w_k} \right) \sum_{k=1}^n w_k \int_{-\omega_m}^{\omega_m} H(s(\omega; r_k) - c(\omega; r_k) - n(\omega) - \varepsilon) d\omega \quad (2)$$

which could reflect the distribution of navigation hazards, risk of piracy, cross-Strait traffic density and so on. To evaluate these integrals, we need expressions for $s(\omega; r)$ and $c(\omega; r)$, as well as noise data. The first of these can be written

$$s(\omega; r) = R(\psi_{Rx}) \left(\frac{c^2}{4\pi f^2} \right) G(r_{Rx}, r) \sigma(\omega; \varphi_{scat}, \varphi_{inc}, r) \delta(\omega - \omega_D) G(r, r_{Tx}) T(\psi_{Tx}) P_{Tx} \quad (3)$$

with P_{Tx} the transmitted power, $T(\psi_{Tx})$ and $R(\psi_{Rx})$ denoting the azimuthal gain patterns of the transmit and receive antennas, $G(r_2, r_1)$ representing the propagation loss factor between positions r_1 and r_2 , and $\sigma(\varphi_{scat}, \varphi_{inc})$ the bistatic radar cross section for an incident angle φ_{inc} and scattered angle φ_{scat} as defined at r , and ω_D the Doppler shift associated with the target echo,

$$\omega_D = \frac{-f}{c} \times \frac{d}{dt} (|r - r_{Tx}| + |r - r_{Rx}|) \quad (4)$$

For target-specific criteria, the RCS must be calculated using a computational electromagnetics code such as NEC4 or FEKO™.

The corresponding expression for $c(\omega; r)$ takes the form

$$c(\omega; r) = R(\psi_{Rx}) \left(\frac{c^2}{4\pi f^2} \right) G(r_{Rx}, r) \sigma(\omega; \varphi_{scat}, \varphi_{inc}, r) G(r, r_{Tx}) T(\psi_{Tx}) P_{Tx} A \quad (5)$$

Here A denotes the area of the resolution cell, whose cross-range dimension increases with range from the receiver. The cell's range extent is determined, in general, by the bandwidth B of the transmitted waveform and, for a phased array system of aperture L_{Rx} we can write

$$A \approx \frac{c^2 |r - r_{Rx}|}{2BL_{Rx} F \cos \psi_{Rx}} \quad (6)$$

The sea surface scattering coefficient $\sigma(\omega; \varphi_{scat}, \varphi_{inc}, r)$ has a continuum of spectral content and, being dependent on sea state, will normally vary with position.

$$\begin{aligned} \sigma(\omega; \varphi_{scat}, \varphi_{inc}, r) = & 2^6 \pi k_0^4 \left[\sum_{m=\pm 1} \int S(m\kappa; r) \delta(\omega - m\sqrt{g\kappa}) \delta(\kappa + k_{inc} - k_{scat}) d\kappa \right. \\ & + \sum_{m_1=\pm 1} \sum_{m_2=\pm 1} \iint \Gamma^2(m_1\kappa_1, m_2\kappa_2) S(m_1\kappa_1; r) S(m_2\kappa_2; r) \\ & \left. \times \delta(m_1\kappa_1 + m_2\kappa_2 + k_{inc} - k_{scat}) \delta(\omega - m_1\sqrt{g\kappa_1} - m_2\sqrt{g\kappa_2}) d\kappa_1 d\kappa_2 \right] \end{aligned} \quad (7)$$

where $S(\kappa; r)$ is the directional wave spectrum at location r and $\Gamma^2(m_1\kappa_1, m_2\kappa_2)$ is a kernel which contains, inter alia, the polarisation dependence, though that does not play a role here. An segment of the resulting database of Doppler spectra, for a given frequency, evaluated for all bistatic angles and wind directions, is shown in Figure 12. The sea parameters were those from Section 2.

Figure 12. A small subset of the database of Doppler spectra evaluated for a particular wind speed and radar frequency, but with all combinations of wind direction and bistatic scattering angle

For an operational deployment, one would compute figures of merit averaged over time of day and the seasons, for which we would need wind, wave and current climatologies. If appropriate, a weighting factor could be applied to effect diurnal or seasonal priorities.

For each of these figures of merit, the value lies in the interval $[0, 1]$, increasing with the merit of the solution. Two simple options for the function to be minimised are $(1 - OF)$ and OF^{-1} .

The figures of merit developed above apply to individual radars but the essence of the problem under consideration is optimisation of a network. The extension to the network case begins

from the observation that, at any given moment, the target will be detected if at least one radar is able to achieve detection. For a set of radars operating in monostatic mode – what has been termed ‘stereoscopic radar’ (Anderson, 1990) – this can be encapsulated in the following expression:

$$OF_3 = \frac{1}{n} \sum_{k=1}^n \left[1 - \prod_{j=1}^n \left[1 - \max_{\omega \in Z} H(s_j(\omega; r) - c_j(\omega; r) - \varepsilon_j) \right] \right] \quad (8)$$

However, we need also to allow for bistatic detection, which has been shown (Anderson, 1990) to increase the probability of detection by circumventing the possibility of double blind speeds in stereoscopic configurations. This leads to

$$OF_4 = \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \left[1 - \prod_{i=1}^n \prod_{j=1}^n \left[1 - \max_{\omega \in Z} H(s_{ij}(\omega; r) - c_{ij}(\omega; r) - \varepsilon_i) \right] \right] \quad (9)$$

While this formulation seems reasonable as far as detection is concerned, it does not take into account the advantage of detecting a target with two radars simultaneously, from different directions. Not only is the probability of detection increased but detection-to-track association is improved; this is an important consideration in the dense traffic environment of the South China Sea where ships are on average only ~ 10 km apart, not much more than the radar range resolution and less than the azimuthal resolution of the smaller radars. Accordingly, when two radars can view a region, we could take dual detectability into account via a performance enhancement factor which is a function of the angle subtended at the target by the two radars.

5.3. Current mapping

In principle, surface current mapping is a relatively simple operation, relying as it does on two very strong peaks in the Doppler spectrum. A fairly rudimentary objective function is the predicted clutter-to-noise ratio, which can be defined for both monostatic and bistatic measurements. Given the length scales of fine structure in the current field in the South China Sea, as shown in Figure 6, some refinements are needed if the function is to serve its purpose effectively.

The most important is consideration of the phenomenon of geometric dilution of precision. The parameters which govern the GDOP for current measurement are (i) the bistatic angle 2β subtended by the two radar axes, and (ii) the crossing angle χ , that is the angle between the nominal current direction and the bisector of the two radar axes. These are indicated on Fig. 4. The theory of GDOP is widely reported (see for example Chapman et al, 1997, Emery et al, 2004) and will not be repeated here. If, instead of the total current vector, one is interested in the components along and perpendicular to a given direction, a slightly different function emerges; we are then interested in the component shown as $u_{\perp}$, so the relevant crossing angle is φ and the error associated with GDOP must be computed using this angle.

Figure 13. Geometry of bistatic illumination, with current vector and shipping lane

Let the radial component of current velocity derived from a measurement by radar 1 be designated $u_1 + \varepsilon_1$, where ε_1 represents random measurement error, and similarly, that of radar 2 by $u_2 + \varepsilon_2$. The estimates of velocity parallel to and normal to the bisector axis are then given by

$$u_p = \frac{(u_1 + u_2 + \varepsilon_1 + \varepsilon_2)}{2 \cos \beta} \quad (10)$$

$$u_n = \frac{(u_1 - u_2 + \varepsilon_1 - \varepsilon_2)}{2 \sin \beta} \quad (11)$$

We transform this vector measurement into the coordinate system defined by the lane axis and its normal by means of a rotation:

$$\begin{bmatrix} u_{\parallel} \\ u_{\perp} \end{bmatrix} = \begin{bmatrix} \cos \varphi & \sin \varphi \\ -\sin \varphi & \cos \varphi \end{bmatrix} \begin{bmatrix} u_p \\ u_n \end{bmatrix} \quad (12)$$

Solving for $u_{\perp}$,

$$\begin{aligned} u_{\perp} = & \left(\frac{\cos \varphi}{2 \sin \beta} - \frac{\sin \varphi}{2 \cos \beta} \right) u_1 - \left(\frac{\cos \varphi}{2 \sin \beta} + \frac{\sin \varphi}{2 \cos \beta} \right) u_2 \\ & \pm \left[\left(\frac{-\sin \varphi}{2 \cos \beta} + \frac{\cos \varphi}{2 \sin \beta} \right) \varepsilon_1 - \left(\frac{\cos \varphi}{2 \sin \beta} + \frac{\sin \varphi}{2 \cos \beta} \right) \varepsilon_2 \right] \end{aligned} \quad (13)$$

Assuming ε_1 and ε_2 are independent and identically distributed, the rms error ε is found by squaring and averaging the error term,

$$\varepsilon = \frac{\left[\cos^2(\varphi + \beta) + \cos^2(\varphi - \beta) \right]^{1/2}}{\sin \beta} \varepsilon_1 \quad (14)$$

The GDOP is defined as the ratio of the rms error ε to the error ε_1 associated with an individual radar. Numerical evaluation (Anderson, 2013) shows that the geometry has a major bearing on the accuracy of HFSWR current estimates, more than doubling the errors once the radars depart from orthogonal viewing geometry by more than 50° .

5.4. Visibility and topographic constraints

The figures of merit and associated objective functions developed in the preceding sections have made one assumption which demands explicit representation – the spatial integrations have made no allowance for blocking of the signal path from radar to patch of interest by an intervening land mass, either an island or part of the mainland. As it happens, HF surface waves can propagate across land, though with much greater attenuation than across sea, and there is an unusual effect (the Millington effect) through which a considerable fraction of signal strength is restored once the signal reaches the sea beyond the intervening land mass. Nevertheless, unless it cannot be avoided, it is better not to entertain the possibility of exploiting signals which have propagated across one or more islands. We can formalise this constraint on single site acceptability as follows.

Suppose there are K landmasses $\{D_j\}_{j=1,K}$ with coastlines $\{\partial D_j\}_{j=1,K}$ adjoining a sea or ocean of which a region W is to be monitored. From the k -th coastline, ∂D_k , construct ∂D_k^+ as follows:

$$\partial \tilde{D}_k^+ = \left\{ r \in \partial D_k \left| \begin{array}{l} \{ \alpha r + (1 - \alpha) r' \} \cap D_m = \{0\} \\ \forall r' \in W, \\ \forall m = 1, K; \alpha \in [0, 1] \end{array} \right. \right\} \quad (15)$$

Then $\partial D_k^+ \subset \partial D_k$ is the subset of the coastline of the k -th landmass which has an unobstructed view of the region W . If a radar is to be placed on D_k , then it must lie on ∂D_k^+ .

In the present study, where each radar can hope to survey at most a part of the area of concern, it is more appropriate to assign the coverage arc at each candidate radar site and measure the effectiveness of that coverage according to the metrics defined earlier.

Accordingly we have chosen to define the coverage arcs by the requirement that they exclude any directions which meet regions W for which the site does not belong to ∂D_k^+ .

For the present illustrative purposes we shall not impose other site-specific constraints such as conditions on local topography or coastline orientation and curvature, though these too could be added if desired.

6. Radar networks for the South China Sea

The tools and procedures described in the preceding sections are of general applicability, but the success of the network optimisation relies on making best possible use of site-specific environmental information, not only oceanographic and meteorological but also the levels and patterns of HF noise. Often this information is unavailable or incomplete, but in most cases one can find climatological data which will serve adequately. We have used the South China Sea mainly as a context to illustrate the ways in which the geophysical information can be exploited, as well as the general issues that could drive network deployment. Needless to say, the South China Sea is of particular interest, so a number of network design experiments have been undertaken. They confirm the applicability of the genetic algorithm methodology to this context, reinforcing the conclusions of Anderson (2013) for the Strait of Malacca.

Figure 14. Part of an optimum solution, focussing on missions over the Spratly islands, showing the receive beam structure for a military-class phased array system.

An example of a candidate solution to a particular network optimisation problem is presented in Figure 14. It shows the coverage of the Spratly Islands afforded by a network of three military-class radars, addressing a combination of mission types. In this example, the optimality criterion of maximising performance for a fixed cost, discussed in Section 4.2, was used to constrain the solution space. A more comprehensive set of solutions, invoking a variety of optimality constraints, is in preparation (Anderson, 2014).

Note again that performance is a statistical quantity, with somewhat greater or perhaps far lesser coverage achieved on any given occasion.

7. Conclusion

The optimum deployment of a network of HFSWR systems is a highly complex task with many factors to be considered, especially when the radars are expected to perform multiple roles. Failure to treat the design problem with appropriate care could seriously degrade performance in one or more radar missions.

In this paper we have described a practical technique for HFSWR network design, based on a genetic algorithm adapted to multi-objective optimisation, and illustrated the method by placing it in the context of designing a multi-radar configuration system for remote sensing and surveillance of the South China Sea. The treatment pays particular attention to the construction of chromosomes and objective functions and extends previous measures of performance to allow for bistatic radar operations. In addition, we emphasise the importance of exploiting a priori knowledge about the regional geography, meteorology and oceanography in the design procedures.

A key advantage of the Pareto dominance formulation developed by Anderson (2013) and incorporated here is that it efficiently identifies those solutions which are superior to their fellows according to every criterion tested, and hence greatly reduces the range of design options which need to be considered. The designer is presented with a range of candidate optimal solutions which can be assessed according to additional considerations which may not be meaningfully quantifiable. When we consider domains such as the South China Sea, where complex geopolitical issues may arise, the virtues of this approach are self-evident.

Author details

S. J. Anderson

Address all correspondence to: stuart.anderson@dsto.defence.gov.au

Defence Science and Technology Organisation, Edinburgh SA, Australia

References

- [1] Anderson, S.J. (1990). Stereoscopic and Bistatic Skywave Radars : Assessment of Capabilities and Limitations, Proceedings of Radarcon-90, Adelaide, Australia, 305-313.
- [2] Anderson, S.J., Edwards, P.J., Marrone, P., and Abramovich, Y.I. (2003). Investigations with SECAR - a bistatic HF surface wave radar, Proceedings of IEEE International Conference on Radar, RADAR 2003, Adelaide.
- [3] Anderson, S.J. (2013). Optimizing HF Radar Siting for Surveillance and Remote Sensing in the Strait of Malacca, IEEE Transactions on Geoscience and Remote Sensing, Vol.51, No.3, 1805-1816.
- [4] Anderson, S.J., Darces, M., Helier, M., and Payet, N. (2013). Accelerated convergence of genetic algorithms for application to real-time inverse problems, Proceedings of the 4th Inverse Problems, Design and Optimization Symposium, IPDO-2013, Albi, France, 149-152.
- [5] Anderson, S.J. (2014). HF radar network optimisation : Case studies for the South China Sea. In preparation.
- [6] Barrick D E. (1998). EEZ Surveillance-The Compact HF Radar Alternative, EEZ Technology, Edition 3. London: ICG Publishing Ltd,. 125-129. See also the CODAR website <http://www.codar.com/> and references therein for detailed information about the SeaSonde family of HFSWR systems.
- [7] Chapman, R.D., Shay, L.K., Graber, H.C., Edson, J.B., Karachintsev, A., Trump, C.L., and Ross, D.B. (1997). On the accuracy of HF radar surface current measurements : Inter-comparisons with ship-base sensors, Journal of Geophysical Research, vol. 102, C8, 118737-18748.
- [8] Chu, P.C., Qi, Y., Chen, Y., Shi, P., and Mao, Q. (2003). Validation of Wavewatch-III Using the TOPEX/POSEIDON Data, Proceedings of SPIE Conference on Remote Sensing of the Ocean and Sea Ice, Barcelona, Spain.
- [9] Chu, P.C., Qi, Y., Chen, Y., and Mao, Q. (2004). South China Sea Wind-Wave Characteristics. Part I: Validation of Wavewatch-III Using the TOPEX/POSEIDON Data, Journal of Oceanic Technology, Vol.21, No.11, 1718-1733.
- [10] Emery, B.M., Washburn, L., and Harlan, J.A. (2004). Evaluating radial current measurements from CODAR high-frequency radars with moored current meters, Journal of Atmospheric and Oceanic Technology, vol. 21, no. 8, pp. 1259-1271.
- [11] Helzel, T., Kniephoff, M., and Pettersen, L. (2010). Oceanography radar system WERA : features, accuracy, reliability and limitations, Turkish Journal of Electrical Engineering and Computer Science, vol. 18, no. 3, 389-397. See also the Helzel Messtechnik website <http://www.helzel.com/> and references therein for detailed information about the WERA family of HFSWR systems.

- [12] International Crisis Group. (2012a). Stirring up the South China Sea (I), International Crisis Group Asia Report No. 223.
- [13] International Crisis Group. (2012b). Stirring up the South China Sea (II): Regional Responses, International Crisis Group Asia Report No. 229.
- [14] Lipa, B., Isaacson, J., Nyden, B., and Barrick, D. (2012), Tsunami arrival detection with high frequency radar, *Remote Sensing*, Vol.4, No.11, 1448-1461.
- [15] Marghany, M. (2009). Volterra - Lax-wendroff algorithm for modelling sea surface flow pattern from Jason-1 satellite altimeter data. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* Volume 5730 LNCS, 1-18 .
- [16] Marghany, M. (2011). Developing robust model for retrieving sea surface current from RADARSAT-1 SAR satellite data, *International Journal of the Physical Sciences*. Vol. 6 (29), 6630-6637.
- [17] Marghany, M. (2012). Three-Dimensional Coastal Front Visualization from RADARSAT-1 SAR Satellite Data. In Murgante B. et al. (eds.): *Lecture Notes in Computer Science (ICCSA 2012)*, Part III, LNCS 7335, 447-456.
- [18] Mirzaei, A., Tangang, F., Liew, J., Mustapha, M.A., Husain, M.L., and Akhir, F.A. (2013). Wave climate simulation for southern region of the South China Sea, *Ocean Dynamics*, Vol.63, 961-977.
- [19] Ninh, P.V., Quynh, D.N., Lanh, V.V., and Lien, T.V. (2000). Geostrophic and drift current in the South China Sea, Area IV: Vietnamese waters, *Proceedings of the SEA-DEC Seminar on Fishery Resources in the South China Sea, Area IV: Vietnamese Waters*, 365-373.
- [20] O'Rourke, R. (2013). Maritime Territorial and Exclusive Economic Zone (EEZ) Disputes Involving China: Issues for Congress, Congressional Research Service Report for Congress 7-5700 R42784.
- [21] Ponsford, A.M. (2012). Persistent surveillance of the 200 nautical mile Exclusive Economic Zone using Raytheon's land-based high frequency surface wave radar, *Raytheon Technology Today*, 2012, issue 2, 25-27.
- [22] Wang, J., Li, M., Liu, Y., Zhang, H., Zou, W., and Cheng, L. (2014). Safety assessment of shipping routes in the South China Sea based on the fuzzy analytic hierarchy process, *Safety Science*, Vol.62, 46-57.