

HF INTERFERENCE TO IMPROVE COVERAGE PERFORMANCE

Proliferation SeaSonde Coastal Current-Mapping Radars

For several decades while coastal HF current-mapping radars were under development, they radiated rarely and only for a few weeks at a time. Now 120 commercial SeaSonde units have been sold and seventy operate continuously in permanent locations. Their numbers are presently increasing at a rate of about 35 per year and they are used for a variety of maritime purposes. Although emitting only 50 watts of power they can map currents to 200km on the lower frequencies; these signals may also travel half way around the world, thanks to overhead reflecting ionospheric layers.

After all, these same layers have provided paths for radio signals used by short-wave broadcasters and amateur radio enthusiasts for the past eighty years. The band is crowded and the days are gone when each of the few HF radar users could ask his national authorities for separate frequencies - there just aren't enough to go round! And if several radars operate on the same channel they may interfere with each other, even hundreds of kilometres apart.

Solving the Challenge

CODAR Ocean Sensors, Ltd. (COS) solved this challenge by using a universally available resource: precisely timed signals beamed down from the 24 satellites of the Global Positioning System (GPS). Normally used for geo-location, we harness the signals here to synchronise the modulation phasing of multiple SeaSonde units so that they operate simultaneously and without mutual interference.

This precision marching-in-step by multiple stations opens another door: that to enhanced coverage performance. One receiver can hear sea echoes from several transmitters. Thus two receivers served by three transmitters may become six radars, but requires only the two computers and real-time processors that originally constituted two conventional radars. When transmitter and receiver are separated, the radar is referred to as being "bistatic", in contrast to the common backscatter, or "monostatic" radar. Combinations in a network are referred to as "multistatic". Until our GPS augmentation, bistatic radar operation has been extremely rare because of the high cost of solving the synchronisation challenge. Multistatic operation pushes the coverage area outward, as well as increasing measurement accuracy and robustness, for little extra cost. The present paper focuses on the GPS addition and resulting enhancements.

Background to SeaSondes

CODAR is a coastal HF radar the primary use of which has been for mapping surface currents. Other output products are wave information and ship tracks. Operating at frequencies a thousand times lower than their numerous microwave cousins, these radars have a major advantage. When their signals' electric fields are vertically polarised they follow the spherical earth well beyond the normal horizon line-of-sight cut-off. Having nothing to do with the atmosphere or ionosphere, this diffraction is made possible by the high conductivity of sea water at HF; it is essentially a perfect conductor.

Currents are mapped as a result of Bragg echoes: signals scattered from surface gravity waves half the radar wavelength for backscatter. These echoes are well defined from their spectra. The extra Doppler shift beyond that induced by the gravity-wave dispersion relation is due to currents near the surface. These Bragg waves (from 50 to 5 metres corresponding to HF frequencies from 3 to 30 MHz)

"feel" the currents to a depth approximately the Bragg wavelength divided by 4 (i.e., the upper meter or two).

The generic "CODAR"-type of radar departs from older, conventional HF radars in its antenna technology. Rather than forming and scanning beams to determine echo bearing, which requires arrays hundreds of meters in length, CODAR uses a compact three-element antenna that is mounted on a post or rooftop. Direction-finding principles replace beam-forming. Invented by the present author and his staff thirty years ago at the NOAA Environmental Research Laboratories, this advance made the system both affordable and easy to deploy and operate. Combined with its low input and radiated powers, this technology has replaced the cumbersome phased arrays as a viable commercial oceanographic product. The family of CODAR products sold today are called SeaSondes.

The maximum coverage distance of the systems depends critically on the radar frequency. Our latest, most popular product is the Long-Range SeaSonde that operates near 4.5MHz. In the backscatter mode, its 50-watt signal reaches to 200km. This allows coverage across the continental shelves, and their real-time current maps are being eyed as assimilative inputs to numerical models. These may soon offer valuable forecasts of water circulation, temperature and salinity within the 200nmi continental EEZs (exclusive economic zones). Real-time current-map now-casts on various web-sites are being used ever more frequently for search and rescue, environmental/pollution management, fisheries, navigation and recreation.

A recent example illustrates a powerful cause-effect relation between surface winds and currents. As a powerful winter "nor' easter" storm moved up the east coast of the US, the four GPS-synchronised Long-Range SeaSondes operated as part of the Rutgers University observatory caught its impact on surface flows. The figures below show (left) NOAA surface-pressure and wind analyses, and (right) the underlying surface currents observed by the radars. The first, upper set occurred when the low pressure centre lay directly off the coast of New Jersey. In the second set, the storm centre has passed over Long Island to the North. Observe how the strong cyclonic winds drive

the currents on the surface in a similar pattern, diverting them from their normal, weak background tidal and geostrophic-dominated circulation. After the fury of the storm with its onshore surges has passed six hours later, the circulation is nearly all offshore as high water levels at the coast relax and ebb back to normal conditions.

How GPS Allows Sharing

The four Rutgers Long-Range SeaSondes that produced the maps above are not the only systems in the world whose 4-5 MHz signals propagate thousands of kilometres at night. Just in the past two years, more than sixteen of these permanent stations have been or are presently being installed on US coasts alone. But in addition, many radio stations world-wide share these bands. A single channel used by all radars would minimise interference to and from these radio stations. Essentially, there are three ways to share: time, frequency, or modulation multiplexing. Time multiplexing means several radars operate sequentially in time, so that no two are "on the air" simultaneously. This results in a big drop in performance, because as the radar-signal data collection period increases, accuracy improves and better coverage is obtained. Frequency multiplexing simply means that all radars are on at the same time but occupy different frequency bands. This is the dilemma we are trying to escape, because of its excessive use of our valuable spectral resources. The CODAR invention that allows them all to operate on the same frequency and not interfere employs modulation multiplexing. Piggybacking on another CODAR invention - our unique digitally generated and processed FMCW waveform - it uses universally available GPS for synchronisation purposes, adding only incrementally to system cost. Exaggerated sketches of waveforms at the reference radar and second GPS-timed radar are shown in the upper two panels of Figure 2. The duration of the sweep, t_s , in SeaSonde systems is typically one second. The reference frequency versus time is shown as the solid red ramp in the third panel, while the solid blue ramp is the second radar's frequency versus time. The dashed ramps represent target echoes from the two signals (red and blue) shifted to the right by their time delays.

In the third panel these time delays to the right show up as instantaneous vertical frequency offsets, all measured with respect to the solid red reference radar ramp. The key to measuring these signal frequency offsets is demodulation in the receiver. A process known as "mixing" in radio and radar receivers produces a frequency that is the difference between two signals: a reference (the red ramp generated at the reference Radar 1) and its received signals (the second blue Radar 2 transmitter ramp, along with dashed echo signal ramps). The output of the mixer (after low-pass filtering) are horizontal lines rather than ramps, as shown in the lower panel. These represent the frequency offsets of the three signals entering the receiver: (i) the dashed echo from Radar 1, (ii) the blue direct solid transmit signal from Radar 2 and (iii) the blue dashed echo signal from Radar 2.

A key advantage realised here is the vast compression of signal bandwidth inherent in our demodulation scheme from what was transmitted. Signal bandwidth or sweep width, f_s , must be 100kHz for a range resolution of 1.5km, but demodulation reduces the frequency span over which one needs to digitise the receiver output at less than 100Hz to contain all of the inherent range information. We achieve a thousand-fold data-rate reduction, but all of the original information is preserved. This allows SeaSonde radars to employ laptop computers and audio digitisers, providing real-time output at an affordable cost.

In addition, we have separated the signal from the blue Radar 2 so it does not interfere. But even more important, we can process and use its echoes with the same, single Radar 1 receiver. Echo information from Radar 2 appears as the upper region on the lower panel, occupying another 100Hz of bandwidth. One gets the idea how a third and fourth radar's transmit signal and echoes may be stacked in an orderly fashion above this region, all to be processed by the Radar 1 receiver and laptop computer. By adding simple GPS-synchronised transmitters, we turn one radar into many, a kind of "loaves and fishes" miracle of modern radar!

One must be careful the other way, too. A signal from Radar 1 must not interfere with the Radar 2 receiver. This seems complicated, but it's a solvable problem. Part of the CODAR invention are algorithms that derive mutually non-interfering timings for multiple radars, based on their locations and separation geometry. Up to fifty radar transmitters grouped within several hundred kilometres can share a common frequency without mutual interference. Already, more than ten SeaSondes in the US and ten in Japan share the common frequencies using GPS. Groups of fifty radars can repeat the same modulation timing, because now their group separations become sufficiently great that any inter-group signals seen are reduced to acceptable levels.

Multistatic Advantages for Current Mapping

The ubiquitous backscatter radars make their measurements in a polar coordinate system: range cells are circular, while bearing grids and radial velocities are constant along radiating azimuth spokes. If you "pull" the transmitter and receiver apart, you stretch the polar system into an elliptical system within which the transmitter and receiver become the focal points of this coordinate family. As any high-school student knows, constant distance (time delay) defines the ellipses: stretch a given length string fixed at the focal points to trace out an ellipse. Thus, the same receiver with two transmitters - one at the same location (backscatter) and one separated (bistatic) - see the same point on the sea at different angles. Said another way, they measure different vector components at this same point. This is shown in the figure below, where a single receiver at Santa Cruz, CA serves a bistatic transmitter to the Southeast (left), as well as its own transmitter (right).

How does all of this improve performance? It happens in two ways. Putting a transmitter on a buoy or offshore oil rig at nearly its maximum backscatter range will extend coverage another 50 per cent, at the same transmitted power. This was predicted theoretically and proven in tests with Rutgers, both at 25MHz and 5MHz (see www.marine.rutgers.edu). The low-power buoy transmitter was solar powered.

The second boost to performance has to do with better robustness and accuracy of current measurements. Theoretically, only two different components of horizontal surface velocity, from two radar observations, are needed to create a total vector. Due to noise in each scalar radar look, and/or to a bad viewing geometry - as when the individual components are nearly parallel - a total velocity estimate may be quite erroneous. Having multiple data values at the same point on the sea at different directions from a multi-station network, along with uncertainties in each radar's scalar observation as outputted by present SeaSondes, will stabilise the measurements and greatly increase their accuracy.

To study these improvements, we use a coverage and quality factor display. This employs all of the measurements that can reach a given point, along with their uncertainties in terms of geometry considerations. The latter include GDOP (geometrical dilution of precision), familiar from accuracy readings on GPS screens, and dilution of spatial resolution in going from radar to the Cartesian coordinates preferred by users. As an example, we began with the four present

Long-Range backscatter Seasonodes operated by Rutgers on the New Jersey coast that produced Figure 1. The coverage and quality are shown as the top figure below. Darkest colours indicate best quality. As shading tends to white, quality fades to poor or no radar coverage at all.

Now, to these four we add a fifth backscatter radar at Fire Island in New York (green dot), and a bistatic transmitter offshore on a spar buoy shown as the red dot, both coming online in Spring 2003. They are GPS synchronised, as are the original four shown as blue dots. But we add a new wrinkle: the Fire Island radar signal and its sea echoes are seen and used bistatically by the original four. The spar buoy transmitter signal and sea echo are seen by all five coastal receivers. Thus, by the addition of a fifth radar and a simple transmitter, we have turned the original four radars into fourteen!

The effect on coverage and performance is striking. Coverage into the Atlantic is extended. To the north, it now reaches well beyond the

tip of Long Island. But the area of highest quality, lying within the innermost green contour, is also dramatically expanded. This shows how networks of SeaSondes, planned for the US as part of coastal observatory programmes, can work together. Thanks to GPS synchronisation they avoid mutual interference, as well as interference to and from other users of the radio spectrum. Coverage and quality of current mapping are improved over those of standard backscatter radars by capitalising on coherent multistatic echoes at the coastal receivers. Finally, “dual use”™ of these systems for wave monitoring and beyond-the-horizon vessel tracking promises to increase their utility immensely amongst whole new classes of users.

<https://www.hydro-international.com/content/article/proliferation-seasonde-coastal-current-mapping-radars>
