SYNTHETIC APERTURE SONAR CONCEPTS IN OUTLINE

What Can SAS Do for Me?

Synthetic Aperture Sonar (SAS) makes centimetric-resolution side-scan survey with high coverage rates a now affordable reality. The new generation of commercial systems deliver 5cm by 5cm resolution, maintained to the edge of a 300m swath, using a 2.5m towfish surveying at over 6 knots.

This article aims to help the interested surveyor gain an understanding of the advantages and limitations of commercial SAS technology. The examples given here are from the GeoAcoustics GeoSAS, developed jointly by GeoAcoustics Ltd and QinetiQ.

What Does SAS Do?
Synthetic aperture sonar uses the forward motion of the survey towfish and many consecutive pings to create a large effective sonar aperture. This large synthetic aperture can be many times the length of the real sonar array, and has the ability to resolve details on the seafloor that would otherwise require a real array of twice that size. The GeoSAS has a 2.1m receive array made up of 67 elements 3.1cm apart, and would typically be used to create a 15m synthetic aperture to give along-track resolution of 5cm. The main advantage of SAS is that a real 30m array does not need to be mobilised to get this kind of performance. The problem with SAS is that the shape of the huge synthetic aperture depends on the path the real aperture takes when it moves through the water. This shape must be known to much greater accuracy than any sensible Inertial Navigation System (INS) could achieve. Before going into how this has been solved, let us first look at some of the concepts behind a practical SAS.

Basics and Concepts
Our example SAS system uses a 60kHz-bandwidth chirp centred at 150kHz; in the 15ms sonar pulse the frequency rises from 120kHz to 180kHz. The resolution of sonar using chirp-pulse compression can be very high, matching the shortest single-frequency ping. The advantage is that a lot of energy can be put into the water in a long, chirped pulse. This gives the chirp processing gain of the system. SAS processing combines many pings from a moving real aperture to create the image. This gives the SAS processing gain. As well as improving resolution, it enhances the signal to noise level and hence the image dynamic range giving more levels of dark grey to black in the shadows). One of the core concepts in the SAS process is the phase centre: the centre point between the separate transmitter and one of the elements in the receive array. The array of phase centres has half the length of the real array, see Figure 1. The separation of the phase centres relates to the best resolution theoretically achievable by the SAS; the examples here illustrate SAS performance achieved in real survey situations.
The sonar pressure at each receive element can be plotted on a line going out from the appropriate phase centre. In the seismic community this is known as a â€˜wiggle traceâ€™. This line is, of course, not truly where the signal comes from, since the field of view of the receive elements is quite wide; the target needs to be in every elementâ€™s real field of view over the whole of the synthetic aperture. In Figure 1 several of these wiggle traces have been laid out, as they would occur if there were only one hard target in the area. For each phase centre in Figure 1 the largest return occurs at the range to the target. We plot this return on the wiggle trace for the phase centre at the correct range, and the shape made by joining up the large returns from each of the wiggle traces is approximately a hyperbola.
The core technique of SAS is to add up the returns from the hyperbola and plot them at its apex. This is known as migration: the returns from the â€˜armsâ€™ of the hyperbola are â€˜migratedâ€™ to the apex. This migration must be done coherently (maintaining phase information) to achieve high resolution. SAS processing does this for every image point in the field of view: only the hyperbola centred on a target will coherently sum all the returns from that target. This gives the high resolution and dynamic range gain of SAS processing and leads to one other important result: the resolution along-track does not change with range. A SAS truly does give range-invariant resolution.
We have seen that a SAS has a very high across-track resolution due to the chirp processing, high range-invariant along track resolution due to the SAS processing and high dynamic range. So why arenâ€™t these sonars everywhere? The answer lies in the challenges of achieving a big enough synthetic aperture and processing the large amounts of data generated.

Difficulties and Answers
The first obstacle to a real-world SAS is that no practical INS can give the accuracy required. For coherent migration to work, all the wiggle traces in the synthetic aperture must line up properly. But if INS is not useable for this, what is? The answer is to use a micro-navigation algorithm. This uses correlation between overlaps from consecutive pings to work out how the sonar has moved between the pings. But this means there must be enough overlap to do the correlation. For reliable micro-navigation, about half the phase centres must overlap from one ping to the next. This means that the real aperture must have advanced by L/4, where L is the real aperture length. This kind of figure is a key buzzword in discussing a real SAS: the type of question heard is "Can your micro-navigation get to L/3 without going wrong?"
With micro-navigation limiting the advance rate between pings, the speed at which the fish can be towed through the water depends only on the time between pings. For maximum 150m SAS range (~165m real range due to the chirp length) the ping rate will be 4.5 pings per second, and a 2.1m real array operated in the L/3 regime will be able to go at over 6kts. But to achieve robust performance in the L/3

regime on all types of seafloor the micro-navigation needs some help. The most critical aspect of this is fish yaw. A good yaw aiding sensor combined with a fish designed to minimise unwanted motion can allow comfortable performance at L/3, with some margin of error.
Another limitation of SAS lies in the region close to the fish. In a side-scan image the nadir is not so useful. A SAS system is similar, and will only start generating good data from about 30m out (at 10m fly height). This leaves a central gap that is hard to fill at the same resolution without technology just as advanced as the SAS itself. A more practical approach is to adjust the survey pattern to cover the gaps â€“ this also provides two looks at most targets, which helps greatly with classification schemes.
As mentioned earlier, the amount of data collected is massive, and a lot of processing power is needed to generate a SAS image within a reasonable time. Until recently this was a serious limitation for SAS technology. Faster processors and advances in SAS algorithms have now enabled on-line SAS processing. But 5cm resolution over 300m fills a lot of screens: how do you make use of all this data? Automated tools are part of the answer (especially for small targets), and automated target candidate identification is a key part of a successful SAS package.
For immediate inspection of the bottom, lower-resolution data can be used. A SAS fish is a very advanced piece of sonar engineering combined with very capable data processing hardware, so beam-forming the data is straightforward. This means a SAS can produce on-line displays of high-resolution sector-scan side-scan to help the operator identify areas of interest for more detailed SAS inspection.

How Does This Compare?

The closest rival to SAS is the high-frequency beam-forming side-scan. A typical top end system might operate at 450kHz with a 1.5m array. To achieve reasonable swath widths this will need a long pulse length (>100us), so will have a resolution of about 10cm across track and 35cm along track and operate out to 150m range. Note that this along-track resolution will vary with range.
Sonar images are hard to interpret, especially for unfamiliar targets. To illustrate the effect of different resolutions, Figure 5 shows a GeoSAS towfish on a stand, together with a series of simulated images adjusted to the resolutions of a SAS and a beam-formed side-scan, in vertical (corresponding to across track), horizontal (corresponding to along track), and dynamic range. The simulated SAS image is a lot clearer than the side-scan, and would pose no problem for computer-aided classification (CAC) or a surveyor.

Conclusions

SAS works and is now commercially affordable. There is a market for high- end side-scan surveys, and survey practitioners need to understand what SAS technology means and how it can help them achieve survey objectives. This article has attempted to explain some of the basic concepts behind the commercial realisation of SAS technology. How well a real SAS performs depends on how well all the high-technology pieces have been put together into a complete system. "What can SAS do for me?" is a question much like "What can a car do for me?" It will get you from A to B - but would you rather travel for a couple of days in a 1970s Yugo or a few hours in a new Mercedes?

Acknowledgements

The authors were purely the reporters for this article. The real work was done by the SAS team at QinetiQ Bincleaves and GeoAcoustics Ltd. Special mention is due to Richard Orme and Nigel Glover for their help in explaining the GeoSAS technology, although we take full responsibility for any failures in their message being correctly interpreted. We recognise (and apologise for) the shortcuts taken in our attempts at clarity, and hope that the SAS gurus (you know who you are) will one day forgive us.