A Vertical Slice of Acoustic Intensity and Velocity

2-D images of waves and turbulence

Jerome A. Smith, Robert Pinkel, and Michael A. Goldin
Scripps Institution of Oceanography
UC San Diego
La Jolla, CA, USA
jasmith@ucsd.edu

***Abstract*—Phased Array Doppler Sonars (PADS) have been developed at Scripps over the past 20 years. These provide sector-scan images of radial velocity and acoustic intensity over ranges of 0.1 m (medical ultra sound) to 1.5 km (ocean surface wave measurement). Most recently we've deployed a 200 kHz PADS from the R/V Revelle during EquatorMix (October 2012, at 140W on the Equator). The pie-shaped measurement area was oriented to probe a vertical slice under the ship parallel to the direction of travel, so that moving averages along the plane can be formed. At the surface, our speed over ground was near 1 m/s, while near 120 m depth, the core of the Equatorial Undercurrent exceeded our speed by about 0.5 m/s; i.e., the undercurrent was moving at 1.5 m/s relative to the surface, which was nearly stationary in an Earth-referenced frame. The objective was to image descending shear-layers that often occur in the early morning before sunrise, along with any associated turbulence or internal waves. Averages can be formed moving with the flow as a function of depth, yielding focused images of acoustic backscatter strength, Doppler bandwidth, and both sine- and cosine- weighted velocity estimates. Where conditions are favorable, the last pair can be recombined to yield horizontal and vertical velocities in some detail.**

Keywords—Doppler Sonar, Measurements of Waves, Turbulence, and mean flows

I. Introduction & Motivation

Imaging waves and turbulence on a segment of 2-D plane has distinct advantages over the use of a handful of "pencil-beams." For one thing, it enables us to visualize the structure of activity over the area; for another, it enables us to form averages moving with the flow (assuming the mean flow is approximately along one of the two directions imaged).

Previous deployments have focused on convection in arctic leads, Langmuir circulation (e.g., [1, 2]), surface waves (e.g., [3, 4]), and near-shore circulation (e.g., [5, 6]). An attempt was made to produce estimates of two components of flow on a fine-scale vertical grid ([7]); however the reverberations in shallow water degraded the estimates.

Most recently, as part of the "EquatorMix" experiment (Oct 8 to Nov 5, 2012, at 140°W, 0°S), a 200 kHz PADS system was deployed over the side of R/V Revelle looking downward at a vertical slice parallel to the ship's heading (see Fig. 1). For this work, a 64 element receive array resolved beamwidths of ~1.5° over a 100° aperture. Because this deployment incurs motion due to both surface waves and ship motion, two independent inertial motion units were installed (3 components of acceleration, 3 of rate-gyro rotation, and 3 of magnetic flux from each), to provide fault tolerance or, when both work, reduced noise variance in the motion estimates.

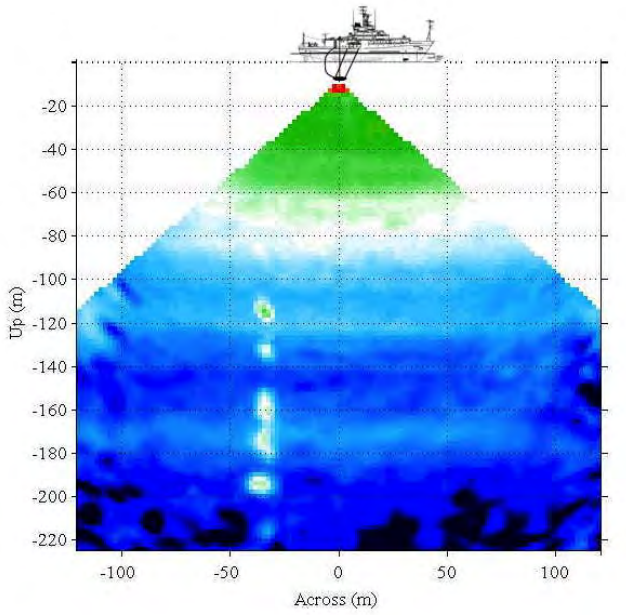

Fig. 1. Schematic illustration of PADS deployment over the side of the R/V Revelle (roughly to scale), including the path of the Fast CTD (FCTD) through the imaged sector about 40m behind the centerline.

Here, the data storage strategy and processing algorythms, including time-delay beamforming, are described. Common problems and the approaches taken to resolve them are also discussed. Other aspects of the sonar development, such as fabrication of the electronics, pressure cases, cabling, and the potting of high-quality electro-acoustic transducers, are left for another time.

This work was funded by NSF OCE 0961801.

II. Data Structure

To ensure synchronicity, time-stamps and all auxiliary data are integrated into the recorded data structure along with the sonar data. Time is encoded as "timemarks", defined as integer increments of 1/20th second starting one day before New Year's midnight (so dividing by 20*60*60*24 yields the year-day). The auxiliary data is also sampled 20 times per second, including both the inertial motion units and a CTD that was attached to the sonar frame (the CTD actually sampled at 16 per second, so some samples are repeated to pad it out to 20 per second).

The 128 sonar channels (sine and cosine homodyned data from each of the 64 staves) are sampled at 80 μs per sample (or 12.5 kHz total bandwidth). For each of 64 staves, the sampling and demodulation is done by first oversampling each amplified real input signal at 16 MHz with a sigma-delta A/D converter chip, with a built-n decimation filter that outputs at 4 MHz, followed by a digital down converter chip (DDC) that demodulates this 4 MHz digital stream using sine and cosine multiplication combined with a FIR filter into an in-phase and quadrature (cosine and sine) output stream at the output sample rate of 12.5 kHz. The final FIR filter reduces the effective output bandwidth to about 70% of the sampling bandwidth (yielding some 8.5 kHz actual usable bandwidth). Each IC handles one channel; each receiver board has 4 sets of input channels, so 16 boards are needed to do all 64 staves (i.e., 64 A/Ds and 64 DDCs are used). The published dynamic range of the FIR filter chip output is 106 dB (18 bits), but we only transmit 16 of those 18 bits, shifting the amplitude accordingly (we clip some of the samples at close range). The resulting dynamic range is sufficient to use a fixed gain, which makes the intensity simpler to interpret and eliminates a possible source of irregularities. For the unit as deployed, the maximum throughput rate of the optical link from the underwater unit to the lab was a limiting factor. In the near future much higher bandwidths will be possible.

One important lesson gleaned from the bi-static exercise [7] is the precise timing of the trasmitted signal relative to the start of a sequence. Because of pre-filtering of the transmitted signal to prevent cross-talk to other acoustic devices, and because of finite "ring-up" time of the transducers, there is a delay of some 22 sample-intervals (80 μs each) before the transmission actually starts. The transmitted signal used was a "repeat-sequence code" [8] with 6 repeats of an 11-bit code (in shorthand, 11100010010), yielding about 4 m resolution in range (but oversampled to every 2 m).

In EquatorMix, the ping-rate was set to 2 per second, permitting ranges up to ~350 m (though in practice the noise floor was reached near 200 m range). While a shorter sequence might be used to maximize the number of pings per minute (say), using these longer sequences permits a more thorough analysis of the noise characteristics, which, as we shall see, can be important. Each ping was recorded as a new record: the whole data structure, including the time, motion data, CTD, and the raw sonar data was written to a raid array (12 Terabyte capacity). Each file was closed and a new one started without interruption for every gigabyte of data, which took roughly 6 minutes to accumulate at a time. Data collection was continuous 24/7, aside from brief pauses to try various ideas for improvement of quality or to fix problems. Once a file was closed, it could be copied to a backup disk without disturbing the ongoing collection.

III. Time-Delay Beamforming

At the widest acoustic arrival angles there are 2 or more samples (or "bits" of the code) spanning the receive array; thus simple phase adjustment as in direct FFT-processing is not effective. In addition, the phase pattern associated with any given arrival angle is a function of frequency, so a wide-band signal would get smeared both in angle and code resolution. Both problems are resolved by doing time-delay beamforming, where the received signal on each stave is advanced or delayed in time relative to the center of the array. The issue is how to do this efficiently; in practice, the question is what algorythm operates the most quickly on a modern computer.

In our experience, the quickest and most accurate way to perform the time-delays is to Fourier transform the whole "ping", rotate the phases as a linear function of frequency (true, not demodulated), and inverse transform back. To implement this, we first form an N x M complex array (where N=64 is the number of staves and M is the FFT length, at least 2 times larger than the number of samples per ping to allow zero-padding). For the EquatorMix setup, the number of samples per ping is about 6200, so we used M = 2^{14} (16384). (It's probably the case that 2^{13} would have provided enough zero-padding, given the prior intensity profile normalization described below.) Then each element of the phase rotation array is of the form

$$D_{nm} = e^{i(Bd_n f_m)} \tag{1}$$

where d_n is the distance of the n^{th} stave from the center of the array, f_m is the radian frequency of the m^{th} FFT element (corrected for the complex demodulation offset), and B is a constant that determines the resulting change in angle of the beam being formed. This last is defined by the desired angle (actually sin(angle)) between the centerline beam (formed by summing stave outputs before rotation) and the next beam off center. Repeated multiplication of the result by the rotation array then moves the summed-stave beam direction further by the same sin(angle) increment, so this is repeated until all the desired beam directions to one side have been formed. The complex-conjugate of D_{nm} rotates the phases (delays) in the other direction, so the other half of the beams are formed from the (saved) unrotated data and repeated application of the conjugated phase rotation array. This array need only be computed once, after reading the sonar parameters, and can then be applied to an arbitrarily long sequence of continuously recorded data (across many files).

Because modern computer architectures are highly optimized for array multiplication and FFTs, this turns out to be very fast. That said, it was still not quite fast enough (as of 2012) to perform in real time, though it was close, and

should be do-able within another few years. In fact, it could probably be done now (2015) using modern GPU (graphics card) compilers to assist.

To get sensible behavior from this FFT analysis algorythm, the signals from all the staves are first divided by the median signal magnitude across all staves at each range for the ping being processed. This profile is also mildly smoothed in range to eliminate any especially small values. Then, before summing the phase-rotated stave data to form a beam, the data are also windowed in the stave-dimension with a 64-point 3rd-order Kaiser-Bessel window (which has an effect similar to windowing for an FFT). While a stave-by-stave magnitude correction could be combined with the window, this was not necessary for the high-quality transducers used (which were batch-matched by measured resonant frequency from a large pool of candidate staves). At this stage, after the beam is formed, the 1-code-lag covariances are formed and bin-averaged to provide Doppler estimates that are not biased by intensity. After each beam's bin-averaged intensities and Doppler covariances are formed, the ping-specific intensity profile is re-introduced, having been similarly bin-averaged. Next, to eliminate synchronous "line noise", a long-time average of a segment of noise-dominated (farthest) ranges is subracted for each beam (line noise sources can be directional). Then the beams are re-normalized by a fixed model profile, so that real changes in intensity are imaged in absolute terms, and we can see (for example) the big change between daytime (few scatterers) and night (lots). The model profile used was fit to a particularly "clear" daytime profile (see Fig. 2):

$$I_{model} = 2 \bullet 10^4 r^{-2} e^{-(r/60m)} + 0.008 \qquad (2)$$

Since this represents a minimum intensity profile, the model profile as used for the fixed normalization (2) is a factor of 4 larger than the fit shown in Fig. 2, so that such minimal intensities appear at –6 dB in the images (see Fig. 1 and below), while the brighter daytime profiles are typically up to +15 dB on this scale.

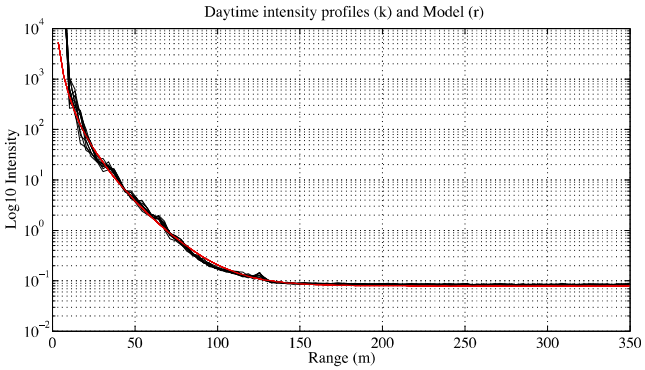

Fig. 2. A few averaged low-intensity profiles (from an especially clear daytime data segment, black), compared to the model profile (red).

Finally, a fixed beampattern formed from a long-time average over a "dependable range interval" is removed, so that intensity variations across the array reflect real intensity variations. This was determined to be quite stable, so the beampattern did not need updating.

IV. Mapping to a Cartesian Grid

The next task is to map the data from the (moving) range-angle array onto a fixed (relative to the ship) *x-z* grid (along-track and vertical directions, respectively). For this we depend on data from the two inertial motion units (IMUs; these were "YEI 3-Space Sensor™ USB" units). The most accurate angle estimates came from integrating the rate-gyro data (specifically the component of rotation orthogonal to the output grid). We found that the drift was sufficiently slow that a weak tendency toward "up" as defined from the accelerometer data kept it in check, where the accelerometer data were effectively smoothed in time with a 10 to 20 second e-folding time-constant to reduce wave-induced accelerations. The instrument typically rocked back and forth a few degrees, which corresponds to several beam-widths, and is significant at ranges over 30 m or so (see Fig. 3).

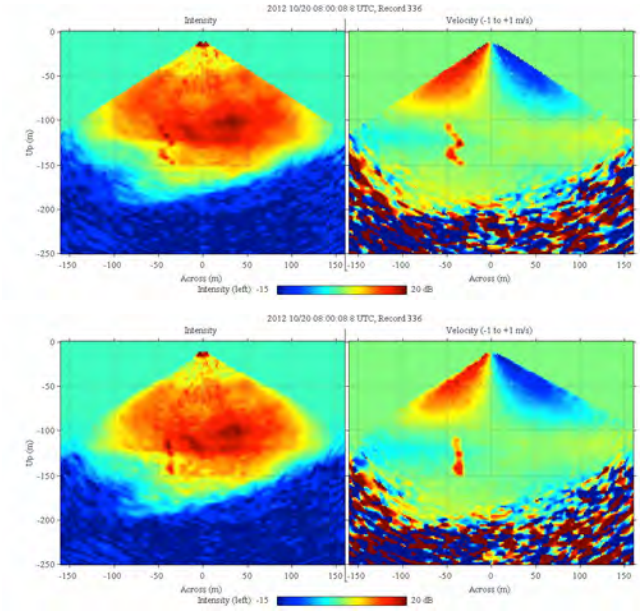

Fig. 3. The track of the Fast-CTD (along x ≈ –40m) clearly shows the effect of wave-induced rocking without motion correction (3a), both in intensity (left) and Doppler velocity (right). With correction, the track is virtually straight down, as it should be. The FCTD drops at about 3 m/s, so it provides a strong Doppler signal (in fact it's aliased, as it exceeds the alias limit of about 2.4 m/s). (See also Fig. 1.)

The vertical displacement of the PADS can be estimated from the CTD pressure, or by double-integration of accelerometer data corrected for cos(tilt) times gravity. Horizontal displacement can be inferred from double-integration of the corrected acceleration, but in this case the sin(tilt) times gravity correction can be problematic. However, typical displacements were less than a meter peak to peak, which corresponds to only 0.25 of a resolution cell, and so could be neglected. The instrument was typically lowered to a mean depth of –9 m (below the surface), and this mean location is what was used for the plots shown here.

To map the data onto a predefined *x*-*z* grid, the location and attitude of the instrument is used to compute the range and angle to each Cartesian gridpoint in sonar coordinates. The value mapped is then interpolated (bi-linearly) between the 4 closest points in the sonar's range-angle array. Grid points lying outside the range-angle limits are padded with an appropriate value (e.g., zero), and noted in a "counting array". The resulting time-series of gridded data can then be averaged or otherwise processed in the time direction. For our near-real-time displays, we block-averaged in time, keeping the previous average so we could form, for example, 1 minute averages sampled every 30 seconds. By averaging the complex covariances, the noise at farther ranges is reduced, resulting in good data over longer ranges, up to 200 m or so.

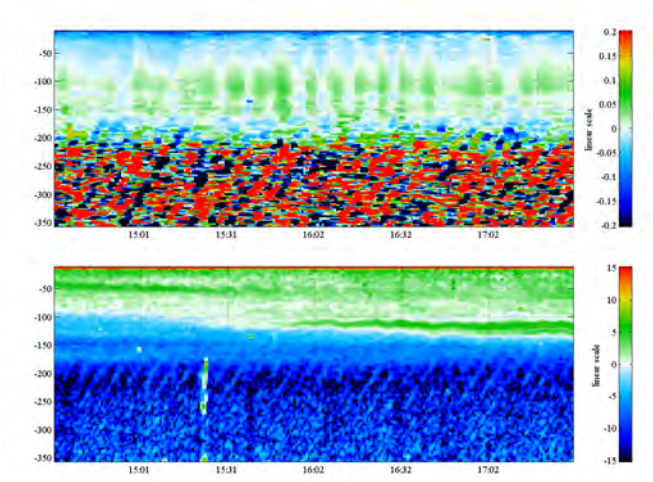

Fig. 4. A time-series of full-range velocity (upper) and intensity (lower) from the beam pointed straight down, showing an "upward propagating signal" in the noise at far ranges.

V. Other Issues

Noise from other machinery or instrumentation, and the effects of swimming scatterers, can have significant effects on data quality.

A. Asynchronous Noise

In addition to the "synchronous noise" such as the line-noise discussed above, a ship typically has numerous other noise sources to contend with. Unfortunately, a big one in EquatorMix was the ship's Doppler speed-log, which also operated at 200 kHz, so it had to be turned off during PADS data collection. While on station at the equator, we needed to maintain about 1 m/s headway relative to the surface water to keep the FCTD from tangling with the propellers, so this was inconvenient. However, the crew on watch didn't complain a bit, they just said "well, we'll have to go by the GPS. But your watch had better keep a sharp eye on the angle of the FCTD cable, since we don't really know the speed of the water." Another source of interference came from the ship's Simrad Multibeam seafloor mapper; but since we were not covering much new territory while on station, this was also tolerable (we turned on whenever in transit).

Finally, there was an unidentified noise source that gave rise to what appear to be upward propogating "stripes" of + and – velocities (see Fig. 4). At first we thought these might be internal waves (phase velocity up means group velocity down), but they are far too stable in frequency and amplitude. So we looked at the full range, and determined that they are a feature of the underlying noise. If we had shortened the sequence length, this might not have been so clear. Also, there remains the possibility of estimating and subtracting the "errant signal".

B. Swimming Scatterers

While we were steaming dead-slow to the east at about 1 m/s relative to the surface, the Equatorial Undercurrent, with a core at about 120 m depth, was flowing at over 1.5 m/s, or 0.5 m/s faster than us. Thus, at about 100 m depth, the water is moving roughly with the ship. This means that even relatively slow swimmers (e.g., squid, of which we saw many) could in theory keep a fixed orientation relative to the ship and hence PADS. Another odd feature that we noticed in both the PADS and Hydrographic Doppler Sonar System (HDSS, a more standard Janus-configuration current profiler) is a persistent bias in vertical velocity, which also often changed right at sunset and sunrise. A theory to explain this is that at least some scatterers are "cycling" in loops roughly fixed relative to the ship. When the mass of scatterers is found just ahead of the ship, the bias is downward at mid-ship; when the mass tends to cluster behind the ship, the bias is upward (see Fig. 5)

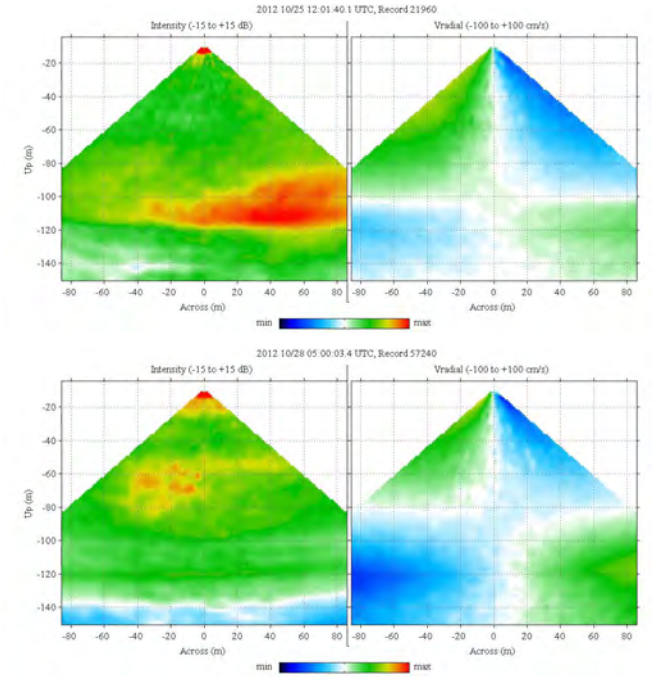

Fig. 5. Scatterer concentration location vs bias. When the scatterers stay slightly ahead (upper), the bias in the vertical beam (along x = 0) is downward (greenish Doppler shift, right panels). When they cluster behind (lower), the bias is upward (bluish).

VI. Conclusion

The instruments performed as hoped, producing nearly continuous time-series over both legs of EquatorMix (12 days on leg-1, 10 days starting 4 days later on leg-2). While some issues remain, we can look forward now to more scientific interpretations arising from this rich data set.

Acknowledgment

The authors would like to thank the scientific and technical members (including graduate students!) of the Ocean Physics Group at Scripps Institution of Oceanography for creating, programming, and operating the instruments discussed here. We also offer a hearty thanks to the crew of the R/V Revelle, and the MPL Res-techs Matt Durham and N. Ben Cohen.

References

1. Smith, J.A., *Observed growth of Langmuir circulation.* J. Geophys. Res., 1992. **97**: p. 5651-5664.
2. Smith, J.A., *Evolution of Langmuir circulation during a storm.* Journal of Geophysical Research, 1998. **103**(C6): p. 12,649-12,668.
3. Smith, J.A., *Observed Variability of Ocean Wave Stokes Drift, and the Eulerian Response to Passing Groups.* J. Phys. Oceanogr., 2006. **36**(7): p. 1381-1402.
4. Smith, J.A. and C. Brulefert, *Evolution of Persistent Wave Groups.* J. Phys. Oceanogr., 2010. **40**(1): p. 67-84.
5. Smith, J.A., *Vorticity and divergence of surface velocities near shore.* J. Phys. Oceanogr., 2008. **38**(7): p. 1450-1468.
6. Smith, J.A. and J.L. Largier, *Observations of nearshore circulation - rip currents.* J. Geophys. Res., 1995. **100**(C6): p. 10967-10975.
7. Smith, J.A., *A Bistatic Phased-Array Doppler Sonar for Wave-Current Research.* Journal of Atmospheric and Oceanic Technology, 2014. **31**(7): p. 1628-1641.
8. Pinkel, R. and J.A. Smith, *Repeat-sequence coding for improved precision of Doppler sonar and sodar.* J. Atmos. Oceanic Technol., 1992. **9**(2): p. 149-163.