Continuous Time-Space Sampling of Near-Surface Velocities Using Sound

JEROME A. SMITH

Scripps Institution of Oceanography, University of California, San Diego, La Jolla, California

(Manuscript received 6 December 2000, in final form 19 April 2002)

ABSTRACT

A phased-array Doppler sonar (PADS) system is described that uses sound at frequencies near 200 kHz to probe an area several hundred meters on a side with 7–20-m spatial resolution. The area can be sampled every second or less with under 2 cm s⁻¹ rms velocity error per sample. Radial velocity estimates from two or more systems can be combined to produce time series of horizontal velocity vector maps over the irregularly shaped overlapping region. Such extensive and continuous sampling in time and space permits analysis via direct 3D Fourier transformation, for example, producing complete wavenumber–frequency spectra. Free waves, bound harmonics, finite-amplitude effects, Doppler shifting by currents, etc., can be studied. Extended temporal sampling permits investigations into lower-frequency vortical and internal wave modes as well as surface waves, and of the modulation of these by tides. A pair of PADS was deployed as part of SandyDuck, a large collaborative field experiment held in 1997 near Duck, North Carolina. An example drawn from SandyDuck data illustrates use of the technique, demonstrating that both mean flow and oscillatory (wave) motions can be detected.

1. Introduction

Many aspects of research into surface waves and wave/current interactions call for comprehensive field measurements of near-surface horizontal velocities (or an equivalent). For example, the combined time-space coverage required to address the evolution of wave groups as they propagate in the field is demanding. In addition, very high resolution directional spectra are needed to resolve details that are important in the calculation of nonlinear transfers (e.g., Young et al. 1995). Studies of the kinematics and dynamics of wind-driven mixing in the deep ocean are also facilitated by continuous time-space coverage of surface currents as the mixed layer develops and evolves (Smith 1998, 1999), and further benefit will result from the simultaneous resolution of surface wave propagation. Finally, studies of waves and wave-driven flows near shore have motivated ever larger and more dense sampling of the currents and surface waves (e.g., Elgar et al. 1994; Feddersen et al. 1998). Ideally, measurements should resolve the spatial distributions and gradients of wave energy, momentum, and "radiation stress" (Longuet-Higgins and Stewart 1964) as well as the underlying currents and shears (vorticity), since these link directly to the dynamics. To fully resolve the 2D propagation and evolution of surface waves along with the underlying currents requires a dense array of measurements at hundreds or even thousands of locations. It is this need that has motivated development of the approach described here: the prospect of resolving surface waves and currents continuously over an area several hundred meters on a side is compelling.

Here two systems are described that use sound to probe the velocity field over finite areas, phased-array Doppler sonars (PADS). The general approach is similar to high-frequency (HF) radars used to map surface currents [Coastal Ocean Dynamics Application Radar (CO-DAR), and Ocean Surface Current Radar (OSCR); e.g., see Paduan and Rosenfeld 1996], but using underwater sound and covering a smaller area with higher spatial, temporal, and velocity resolution. Acoustic signals at frequencies near 200 kHz are used to produce radial velocity estimates over pie-shaped areas roughly 400m radius by 90° in bearing. The measurements are resolved to 7.6 m (range) by 7° (bearing), with new estimates produced every 0.75 s. The bearing is resolved via phased-array beamforming, and the range by time since transmission. By combining the radial velocities from two such devices located some 300 m apart (Fig. 1), both horizontal components of velocity can be estimated on a grid several hundred meters on a side, at thousands of locations.

The purpose of this paper is to describe the technique, its limitations, and the associated error bounds. An example illustrating the technique is drawn from a recent Office of Naval Research (ONR)-funded deployment of two PADS off Duck, North Carolina as part of

Corresponding author address: Dr. Jerome A. Smith, Scripps Institution of Oceanography, University of California, San Diego, 354 OAR Building, Mail Code 0213, La Jolla, CA 92093-0213. E-mail: jasmith@ucsd.edu



FIG. 1. Plan view of the SandyDuck experimental site, showing the area covered by two PADS. North is about 20° clockwise of left. The arrows indicate velocity estimates from a single snapshot, dominated by swell from the upper right (SE); the longer arrows correspond to velocities approaching 1 m s⁻¹. Both horizontal components are estimated in the overlap region, but only one component is in the corners covered by a single sonar. The circles show locations of instrumented frames deployed by various collaborating investigators. The heavy circles denote the four locations from which data comparisons are shown here (the tide increases the depths by 0.5 m at the time of comparison). The location is the Field Research Facility of the U.S. Army Corps of Engineers, near Duck, North Carolina.

"SandyDuck" (Fall 1997), a large, multi-institution multiagency experiment to investigate nearshore dynamics. While that deployment was in shallow water, this paper focuses on aspects of the technique that apply in deep water as well. A companion paper (Smith 2002) investigates effects specific to using the method in shallow water, the most significant of which is competition between the desired volume backscatter signal and bottom backscatter.

2. PADS design and operational parameters

The basic approach is to use high-frequency acoustic repeat-sequence codes (Pinkel and Smith 1992) to estimate radial velocities as a function of range and angle from each of two PADS. The signal is emitted in a wide horizontal fan, radiating outward in the water from the instrument package and intersecting the surface. The sound scatters off particles in the water, especially the near-surface bubble layer, and also (perhaps) off the bottom. Some backscattered sound returns to the sonar, where the signal is received on a 16-element array and analyzed for frequency shift versus angle and elapsed time since transmission. The frequency shift of the back-

scattered signal is proportional to the radial component of the velocity of the scatterers. Horizontal angle (bearing) is resolved by spatially Fourier transforming the signal received on the regularly spaced array of 16 transducers, synthesizing a set of beams spanning about 90°. For direct-path transmission and return, the time delay after transmission translates to distance from the sonar via knowledge of the soundspeed in the water. For example, in SandyDuck the soundspeed was evaluated from daily conductivity-temperature-depth probe/profiler (CTD) casts taken at the end of the FRF pier, using middepth values of temperature and salinity. A typical value was 1515 m s⁻¹. The vertical location is not directly resolved; the effective depth of measurement depends on the vertical distribution of scatterers. Typically, bubbles are the most efficient scatterers, and these tend to be concentrated toward the surface with about a 1-2-m e-folding depth (e.g., Crawford and Farmer 1987).

To permit simultaneous operation in SandyDuck, two PADS were designed to operate at slightly different center frequencies: 226 kHz (PADS 1, closer to the FRF pier) and 195 kHz (PADS 2, at the north end of the surveyed area). In practice, the latter was shifted to 190



FIG. 2. Layout of the curved transmitter arrays as used in the 226-kHz sonar at SandyDuck. Transmit sections were laid out in four groups, each split symmetrically about the center axis (denoted by different shades of gray). This permits up to four independent power amplifiers to be used, providing up to 4 times the total transmitted power. While it also permits limited array shading, this was found to be unnecessary in practice: the curvature of the array face is adequate (see Fig. 3).

kHz to further reduce cross-talk between the systems. The transducer heads (transmit and receive) and frontend electronics need to be tailored closely to the operational frequency to maintain beampatterns and efficiency. To facilitate beamforming over 90° ($\pm 45^{\circ}$), the widths of the individual receiving staves are 0.404 cm (226 kHz) and 0.447 cm (190 kHz). The receive arrays consist of 66 elements or staves that can be reconfigured in the lab to hook up any combination of 16 channels: 16 in a row (as used in SandyDuck), or 16 pairs, triplets, or quadruplets (permitting higher angular resolutions



FIG. 3. Computed transmit patterns from the curved transducer head: black for an ideal array, and gray) for an array with random 0.04-cm rms positional error (in both x and y) for each element. Even with realistic positional errors, the predicted beampattern is excellent.

over a smaller angular spans). Experiments were carried out with sparse arrays as well as linear (dense) arrays; however, the increase in angular resolution possible with a sparse array (using the same number of channels) comes at the expense of phase uncertainty, translating directly to larger errors in the Doppler estimates. For the present purpose this is a bad trade; a regular linear array provides much more stable estimates of Doppler shift versus angle, and hence less velocity error. To preserve performance characteristics, the curved transmit heads are scaled with operating frequency: each is formed of 48 staves laid out on an 88° arc, with inner radii of 12.28 cm for 226 kHz (Fig. 2), and 13.7 cm for 190 kHz. This transmitter design produces a beampattern that rolls off to 1/4 power at $\pm 45^{\circ}$, with excellent suppression of energy outside the intended field of view even with modest positional errors or amplitude variations (Fig. 3). It also provides enough surface area to transmit sufficient power to reach 400 m (~300 W electrical input power) without cavitation or saturation effects (Smith 1989, and references therein).

The electronics to synthesize and transmit the repeatsequence code, and to receive, mix, and digitize the backscattered signal, are contained in an underwater housing. Only AC power and fiber-optic communications links to shore are needed for continuous operation. Timing information is sent to each system to synchronize within a microsecond. The underwater controller and associated hardware can be reprogrammed, providing flexibility in the operating frequencies (transmit and mix), coding scheme, and filtering characteristics of the



FIG. 4. Spectral diagram of the off-center mixing scheme. The mix frequency is set at one edge of the desired signal band, and the mixed signal is filtered and sampled with sufficient bandwidth to span the signal band. "Extra" noise is admitted from an equal size band on the other side of the mix frequency from the signal band, and folded over onto the signal (note that the noise level in the postmix section is twice as high). However, the additional noise is deemed acceptable to reduce the electronics requirement by half, and to eliminate problems associated with balancing two channels in quadrature. The mix frequency can be at either the low end (as shown) or the high end of the signal band. In the latter case the postmix frequencies are a mirror image of the signal.

transmitted and received signals. In SandyDuck, the systems were operated with a 32-µs sample rate, corresponding to 15,625-Hz net usable bandwidth. To fill the available bandwidth, the repeat-sequence code bits are two samples long ($\tau = 64 \ \mu s$). For SandyDuck, 13 repeats of 13-bit codes were transmitted (M = 13, L =13). The code used for the 226-kHz system can be represented as $\{1010110011111\}$, where the ones are 0° phase (positive) and zeros are 180° phase (negative) segments of the carrier wave. The code used for the 190-kHz system is the time-reversed version of this: {1111100110101}. These codes individually satisfy the condition of having minimal autocovariance magnitudes at all lags other than multiples of L = 13. While these are probably not optimal orthogonal codes, in the sense of satisfying both autocovariance and cross-covariance minimization criteria, it was found that with one code reversed and 36-kHz frequency offset, the cross-talk between systems was reduced to an almost undetectable level even for the direct path transmissions from one system to the other. The systems use an off-center mixing scheme: after mixing and filtering, the signal carrier (center) frequency is sampled at four samples per cycle, and the usable bandwidth extends from 0 to 2 times this intermediate frequency (Fig. 4). Although this admits some additional noise (from the negative post-mix frequencies), it reduces the hardware requirement by half (a crucial trade-off with mid-1990s technology). The synthesized transmit codes are prefiltered to reduce the transmitted energy outside the desired frequency band; it is particularly important to filter out energy that would cross into the negative frequency region with the offcenter mixing. The 226-kHz received data are mixed high, with a 233,812 Hz waveform; the 190-kHz data are mixed low with 182,187 Hz. These mix frequencies differ by 0.5 Hz from those needed to exactly center the signals; this results in deterministic biases of 0.167 and 0.20 cm s⁻¹ for the 190- and 226-kHz systems, respectively (these are easily corrected). Mixing with a frequency farther away from the other sonar's frequency band reduces cross-talk between systems.

The mixed, filtered, and digitized data are transmitted back to shore (about 1 MB s^{-1} for each system), where it is processed (beamforming, Doppler estimation, range averaging, and ping averaging) and stored. Shore-based processing facilitates debugging and development of improved analysis schemes. An in situ calibration process was developed to equalize and phase-correct the acoustic receiver transducers. Covariance pairs between receive elements are formed between all samples at one time and samples one code length later, producing a 16 \times 16 matrix. For the SandyDuck setup, the appropriate time lag τ is 26 samples or 0.832 ms (one code length), corresponding to alias velocities V_{max} of 2.0 m s⁻¹ at 226 kHz and 2.4 m s⁻¹ at 190 kHz. The matrix of covariances is averaged in range over a user-selected range interval (a middle segment where the data are both in the far field and reliable, say, from 50 to 170 m), and is further averaged in time according to the following scheme:

$$C_{ij}(t) = (1 - p)C_{ij}(t - \tau) + pC_{\text{new}},$$
 (1)

where $C_{ii}(t)$ represents the accumulating average at time t, τ is a fixed time lag (either 0 or 1 code length = 0.832 ms), C_{new} is the new array of products, and p sets the time constant for the average (e.g., a value p = 0.01corresponds to a time constant of 100 pings, or 75 s). The diagonal elements of C_{ii} are just the autocovariances of each receive element at time lag τ , from which the relative magnitude responses can be calibrated. Using the finite time lag τ helps eliminate noise variance (which may vary for each receive element) from the calibration estimates. Column averages of this matrix correspond to the net correlation between each element and the signal from the beam aimed straight ahead (the center beam). From these column averages the phase of each receive element relative to the mean can be derived. The mean phase across the array is the estimated Doppler shift for the center beam, which must not be included in the calibrations. The linear trend in phase across the receive array must also be removed; otherwise, this scheme will steer the center beam to point at the brightest scatterer. An optimal solution would be to adjust the phases to maximize the mean-square deviations in intensity (Attia and Steinberg 1989); however, the present scheme is much simpler and can be computed quickly. If there are only a few bright, compact scatterers within the field of view of the center beam (such as instrumented frames that pierce the surface), this method converges quickly on appropriate phase corrections.

In practice, it was found that 1) the beamforming is acceptable without calibrations (the transducers were hand-picked for uniform response); 2) amplitude corrections alone provide most of the improvement in

beamforming; and 3) the phase corrections (once properly implemented) do produce some additional sharpening of the image, though barely detectable (even though some of the phase corrections land in the neighborhood of 10° – 20°). The complex calibrations resulting from this procedure are stable over times of at least 1 week. After an average is accumulated over a few minutes, the values can be locked in. Since noise and interference from other sonars or from reflections off the Coastal Research Amphibious Buggy (CRAB) are occasionally loud enough to disrupt the calibration estimates, switching off the updating scheme after they stabilize is desirable. This procedure can also be applied to non-time-lagged (lag 0) covariances, if the individual noise levels are small or uniform across the array. In practice this yields comparable results, which increases confidence that the scheme is properly implemented. The use of detrended lag-1 arrays permits us to skip computing and storing lag-0 values altogether.

Beamforming can be performed either from covariance matrices as described above or from direct (windowed) FFTs of the 16-element data. Range-averaging covariances over (M - 1)L/2 = 78 samples (3.78 m) and block averaging over two pings reduces the data by a factor of 156, more than enough to compensate for an initial expansion of data from 16 elements to 256 covariances. To provide windowing (sidelobe suppression), a 16-element window is crossed with itself to produce a 16×16 window matrix (formed a priori), which is applied element by element to the covariance array prior to averaging all similar spatial lags. The result is equivalent to windowing prior to forming the covariances, and produces no negative sidelobes (which can disrupt phase estimates). For the conditions at SandyDuck, for example, a second-order Kaiser-Bessel window was found to provide a good trade-off between sidelobe suppression and angular resolution (Harris 1978). Note that the (-) spatial lags are different from the (+) lags for nonzero time lag, so there are 31 distinct spatial lags. The reduced array of spatial lags can then be Fourier transformed, resulting in equivalent timelagged covariances versus angle and range. To provide interpolation, a 64-element FFT was employed.

For long time averages it is useful to store the full set of 256 covariances, since this permits arbitrary reworking of the beamforming. For 1.5-s sampled data or faster, however, this produces a significant data storage burden. To reduce data storage needs, only the 31 similar spatial lags (like lags) were retained for each range and ping pair (a reduction by a factor of just over 8 relative to keeping the complete covariance matrix). However, once the covariance matrix has been collapsed to 31 similar spatial lags, the windowing cannot be redone effectively, so window selection is important.

Before estimating Doppler shifts, the data are further pair averaged in range, producing range bins with length equivalent to 12 code repeats $[T_a = (M - 1)L_{\tau}]$. This corresponds to range resolution (at 1515 m s⁻¹) of 7.56 m, sampled every 3.78 m. As noted, two-ping averages were formed to reduce data storage by 1/2, resulting in samples everywhere (32 angles by 125 range bins) every 1.5 s. Some runs were made storing both lag-1 and lag-0 covariances, and some storing just lag 1 (time lagged by one code length, or 0.832 ms; cf. Pinkel and Smith 1992; Trevorrow and Farmer 1992). The latter are the minimum data needed to estimate velocity. Inclusion of lag-0 covariances (i.e., total intensities) permits estimation of coherence, providing useful diagnostics and/ or bandwidth estimates (Miller and Rochwarger 1972). The resulting data storage rate is 80 MB (just lag 1) or 120 MB (both lags) per hour per sonar. About 5 GB of data were collected per day at SandyDuck. With two magneto-optical drives per sonar, using media that hold 1.3 GB per side, the operational schedule consisted of a quadruple disk flip or swap occurring once per day. To minimize the adverse effect of a disk failure, 1-h files were stored with odd hours on one disk and even hours on the other; thus with a disk failure we would still have every other hour for semicontinuous coverage. While there were several disk failures (to be expected when over 100 disks are employed), these were caught and fixed quickly, and this was not a major cause of data loss.

3. Doppler estimation error

Doppler estimation error with repeat-sequence coding was addressed by Pinkel and Smith [1992, Eq. (11)]. For the systems as specified above, the estimated error variance of the derived mean frequency is

$$\sigma_{\omega}^{2}(\text{ideal}) = \frac{(1 + T_{a}/2T_{o})}{PLT_{a}T_{o}} = T_{o}^{-2} \left[\frac{(T_{o}/T_{a} + 1/2)}{PL} \right]$$
$$= (24.06 \text{ rad s}^{-1})^{2}, \qquad (2)$$

where we have modified the equation to include ping averaging (P = 2), and inserted the value for the overlap time $T_o = (M - 1)L\tau$, and where (as noted above) M= 13, L = 13, $\tau = 64 \ \mu$ s. To convert this to rms velocity error, take the square root and multiply by ($c/4\pi f$):

$$\Delta V = \left(\frac{c}{4\pi f}\right) (\sigma_{\omega}^2)^{1/2} \approx f^{-1} (2902 \text{ m s}^{-2}).$$
(3)

This yields rms error estimates of $\Delta V = 1.28$ cm s⁻¹ (226-k system) and 1.53 cm s⁻¹ (190-k system). In practice, the measured rms errors tend to be larger, from 1.1 to 1.3 times these values, depending on environmental conditions and other (incompletely understood) factors (Pinkel and Smith 1992).

While this is adequate for the high signal-to-noise ratio (SNR) areas, we desire error estimates to the farthest usable ranges, as the signal fades out. For this we need finite SNR estimates. Theriault (1986) derives a Cramer–Rao lower bound on the error variance of a single-ping uncoded pulse for finite SNR, with equal For example, with just two radial estimates this simaveraging and overlap times $(T = T_a = T_o)$:

$$\sigma_{\omega}^{2} \geq \frac{\text{SNR}^{2} + 36\text{SNR} + 30}{T^{2}\text{SNR}^{2}}$$
$$= T^{-2} \left(1 + \frac{36}{\text{SNR}} + \frac{30}{\text{SNR}^{2}} \right).$$
(4)

Here we use the empirical finding that the actual error variance is typically about twice this value, and modify the equation to include the effects of an ideal code (divide by L) and averaging pings (P = 2):

$$\sigma_{\omega}^{2} = \frac{2}{\mathrm{PLT}^{2}} \left(1 + \frac{36}{\mathrm{SNR}} + \frac{30}{\mathrm{SNR}^{2}} \right).$$
(5)

At the high SNR limit, we recover the value (27.78 rad $(s^{-1})^2$, or about 1.15 times the value in (2), in line with the real-world performance reported by Pinkel and Smith (1992). Thus we adopt (5) to describe the error associated with each radial velocity estimate. The SNR value is estimated from the ratio of the observed acoustic intensity at each location to the mean signal strength at the farthest ranges in the same direction (where the signal has attenuated beyond detection).

4. Combining data from two or more systems

a. An optimal solution

At each location x, suppose there are two or more radial velocity estimates S_i from different directions φ_i , with associated error estimates ϵ_i :

$$S_i = u \cos \varphi_i + v \sin \varphi_i + \varepsilon_i, \qquad i = 1, 2, \dots$$
 (6)

We want an estimate

$$\tilde{u} = a_1 S_1 + a_2 S_2 + \cdots \tag{7}$$

such that $\langle (u - \tilde{u})^2 \rangle = \text{minimum}$.

The best possible estimate would incorporate as much dynamics as we know: the data would be assimilated into a model. Where the dynamics are not well established, or where we wish to independently test modeling assumptions, we might instead opt for a purely statistical optimization. Following Bretherton et al. (1976), set

$$0 = \frac{\partial}{\partial a_i} \left\langle \frac{1}{2} (\tilde{u} - u)^2 \right\rangle = \left\langle (\tilde{u} - u) \frac{\partial \tilde{u}}{\partial a_i} \right\rangle$$
$$= \left\langle (\tilde{u} - u) S_i \right\rangle = \left\langle (a_1 S_1 + a_2 S_2 + \dots - u) S_i \right\rangle \quad \text{or}$$

$$\langle uS_i \rangle = a_1 \langle S_1 S_i \rangle + a_2 \langle S_2 S_i \rangle + \cdots$$
 (8)

Defining the data covariance matrix and signal-data covariance vector,

$$C_{ij} \equiv \langle S_i S_j \rangle$$
 and $D_i \equiv \langle u S_i \rangle$, (9)

we can write the solution in vector and matrix form (summing over *j*):

$$a_i = C_{ij}^{-1} D_j. (10)$$

plifies to

$$a_{1} = \frac{\langle S_{2}^{2} \rangle \langle u S_{1} \rangle - \langle S_{1} S_{2} \rangle \langle u S_{2} \rangle}{\langle S_{1}^{2} \rangle \langle S_{2}^{2} \rangle - \langle S_{1} S_{2} \rangle^{2}} \quad \text{and} \qquad (11a)$$

$$a_{2} = \frac{\langle S_{2}^{2} \rangle \langle uS_{2} \rangle - \langle S_{1}S_{2} \rangle \langle uS_{1} \rangle}{\langle S_{1}^{2} \rangle \langle S_{2}^{2} \rangle - \langle S_{1}S_{2} \rangle^{2}}.$$
 (11b)

Implementation requires evaluation of the covariances, the quantities in braces $\langle \rangle$. For a general-purpose solution. set

$$\langle u^2 \rangle = \langle v^2 \rangle = 1, \quad \langle \varepsilon_i \varepsilon_j \rangle = N_i \delta_{ij}, \text{ and}$$

 $\langle uv \rangle = \langle u\varepsilon \rangle = \langle v\varepsilon \rangle = 0.$ (12)

This corresponds to normalizing the error variances N_i (etc.) by the true velocity variances. The results (after a little math) can be written

$$a_{1} = \frac{(1+N_{2})\cos\varphi_{1} - \cos\Delta\varphi\cos\varphi_{2}}{(1+N_{1})(1+N_{2}) - \cos^{2}\Delta\varphi} \text{ and } (13a)$$

$$a_{2} = \frac{(1+N_{1})\cos\varphi_{2} - \cos\Delta\varphi\cos\varphi_{1}}{(1+N_{1})(1+N_{2}) - \cos^{2}\Delta\varphi}$$
(13b)

where $\Delta \varphi \equiv \varphi_2 - \varphi_1$.

Likewise, the other component of velocity is estimated from

$$\tilde{V} = b_1 S_1 + b_2 S_2 + \cdots,$$
 (14)

where, for the case of just two signals,

$$b_1 = \frac{(1+N_2)\sin\varphi_1 - \cos\Delta\varphi\sin\varphi_2}{(1+N_1)(1+N_2) - \cos^2\Delta\varphi} \quad \text{and} \quad (15a)$$

$$b_2 = \frac{(1+N_1)\sin\varphi_2 - \cos\Delta\varphi\sin\varphi_1}{(1+N_1)(1+N_2) - \cos^2\Delta\varphi}.$$
 (15b)

b. Comparison with Gaussian elimination

Given radial velocities from two directions [i = 1, 2, ..., 2]in Eq. (6), with $\epsilon_i = 0$], we can eliminate v:

$$u = \frac{V_1 \sin\varphi_2 - V_2 \sin\varphi_1}{\cos\varphi_1 \sin\varphi_2 - \sin\varphi_1 \cos\varphi_2}$$
$$= \frac{V_1 \sin\varphi_2 - V_2 \sin\varphi_1}{\sin\Delta\varphi}.$$
 (16)

Hence, the simple elimination solution for u takes the form

$$a_1 = \frac{\sin\varphi_2}{\sin\Delta\varphi}$$
 and $a_2 = \frac{\sin\varphi_1}{\sin\Delta\varphi}$. (17)

This solution has the serious defects that 1) it blows up as $\Delta \varphi \rightarrow 0^{\circ}$ or 180°, and 2) it weights measurements equally regardless of error content. Reasonable results can be obtained by setting thresholds for the SNR of each beam and for how close to 0° or 180° is to be



FIG. 5. A snapshot of horizontal velocities from the dual-PADS (thin black arrows) and from in situ current meters (thicker gray arrows based within circles) near the start of the comparison segment. An arrow long enough to reach from one grid point to another corresponds to 50 cm s⁻¹. The gray shading compares the velocity signal strength to the estimated error variance for the PADS (SNR; darkest means error dominates, lightest means the signal is reliable). Methodology for estimating error variances is described in Smith (2002). Where only one component is estimated from the PADS, the maximum SNR possible is 1 (half the variance of the signal is resolved), or 0.0 dB. At the farthest ranges, where the acoustic signal is weak, the error variance increases. Comparisons are made near the heavier circles; because of acoustic interference from the frames, the PADS estimates are taken about 15 m offshore of the in situ sensors.

allowed. However, this is both arduous and unreliable in the long run: if the thresholds are set high, there are never spikes, but many good data are discarded; if they are set low, many spikes occur.

For comparison, rewrite the optimal solution [e.g., Eq. (13a)] in a similar form: expanding $\cos\Delta\varphi$ and $\sin\Delta\varphi$ and rearranging, the optimal a_1 (for example) can be written

$$a_{1} = \frac{N_{2}\cos\varphi_{1} + \sin\Delta\varphi\,\sin\varphi_{2}}{N_{1} + N_{2} + N_{1}N_{2} + \sin^{2}\Delta\varphi}.$$
 (18)

This goes asymptotically to the simple elimination solution as the noise values go to zero. Moreover, it is well behaved in the presence of noise: 1) as N_2 becomes large, it recovers just the portion of signal captured in S_1 ; 2) as N_1 becomes large, it goes to zero; and 3) as the two signals approach parallel or antiparallel with finite noise, the unresolved component goes to zero rather than infinity. The method also provides objective estimates of the error variance (see Bretherton et al. 1976).

5. Example: 3D spectra of horizontal velocity

The three-dimensional (two space and time) views of horizontal velocity sampled by PADS systems can be analyzed via 3D Fourier transformation. Similar analyses have been performed on radar intensity image sequences (Young et al. 1985) and with Doppler estimates from a deployment of the focused phased-array imaging radar (FOPAIR) at Duck, North Carolina, in 1994 (Frasier and McIntosh 1996). Here the dual-Doppler arrangement permits examination of both horizontal velocity components over an overlap area a few hundred meters on a side (Fig. 1).

Data for the example shown here were gathered as part of SandyDuck, a multiagency experiment conducted in the fall of 1997 at the Field Research Facility of the Army Corps of Engineers, near Duck, North Carolina. The sample is an hour's worth of data taken starting at 1400 local time (1900 UTC), 10 September 1997. An example snapshot of the horizontal velocities estimated from PADS data is shown in Fig. 5, along with currents from in situ sensors [sonic altimeter, pressure,



FIG. 6. Spectral comparisons of measurements of the cross-shore component of velocity from the PADS and in situ current meters (SPUV). (a) Frequency spectra at the 5.8-m-deep site: (solid) PADS current, (dotted) SPUV. (b) Spectra at 4.5-m depth, averaged over three sites (see text and Fig. 5): (solid) PADS, (dotted) SPUV. Note especially enhancement of the peak near 0.054 Hz (19.4-s period) from the deeper to the shallower sites. (c) Spectral coherence between PADS and SPUV data at (solid) 5.8-m-deep site and (dotted) averaged over the 4.5-m-deep sites. Note that the coherences are especially high across the surface wave band (0.05–0.17 Hz). At lower frequencies, stratification may facilitate larger differences between near-surface (PADS) and near-bottom (SPUV) currents. At higher frequencies, the finite separation between the PADS and SPUV measurements (~15 m) becomes more important.

and U and V currents (SPUV); see above] and estimates of PADS data quality [data quality assessment is addressed in some detail in Smith (2002); for the present purpose note that lighter means better]. The sample area spans a modest depth range (4–6.5-m depth; see Fig. 5). Since surface wave dispersion (for example) depends functionally on depth as the square root or weaker, significant insight can be gained from the 3D spectra estimated without corrections for these changes in depth. Also, in shallow water there is competition between bottom and volume backscatter, introducing a systematic attenuation of the velocity estimates. This is considered in a companion paper (Smith 2002). For now we note that these variations in response are large scale and



FIG. 7. Spatial maps of shore-normal velocity variance at three frequencies: (left) 0.018 Hz (54-s period), (center) 0.029 Hz (34-s period), and (right) 0.052 Hz (19.4-s period). Each panel is scaled relative to the maximum observed value at the given frequency. Nodes and antinodes parallel to shore indicate partial reflection of wave energy. Note that node locations change with frequency, and that the spacing is consistent with shallow-water dispersion as expected for gravity waves at these frequencies (this would hold true for shore-trapped edge waves as well as leaky-mode or unbound waves). The patterns may vary weakly in the alongshore direction due to slight variations in bathymetry. The sizes of the Xs in the figure represent the relative magnitudes of the spectral estimates from the in situ data at each location where this was available. It is fairly clear that the Xs are small in the lighter zones and larger in the darker zones, indicating qualitative agreement between the in situ (bottom current) estimates and the PADS data. Quantitative comparisons are made in the companion paper (Smith 2002).



FIG. 8. Alongshore-averaged variance vs frequency and distance offshore. The variances have been multiplied by $h^{3/2}$ (where *h* is depth; this scaling would equalize variances for nondissipating waves propagating nearly orthogonal to shore) and are scaled by the maximum at each frequency. At frequencies below 0.045 Hz, the pattern is consistent with nondissipative waves reflecting from the shore. Between 0.045 and 0.07 Hz, the pattern is consistent with partial dissipation and partial reflection. For all frequencies below about 0.07 Hz, the curving stripes of light and dark (corresponding to antinodes and nodes of cross-shore current variance, respectively) are consistent with nearly shore-normal waves reflecting off the beach, with the spacing between nodes increasing as the wavelength increases. Above 0.07 Hz, it appears that the waves dissipate consistently shoreward of the 400-m FRF coordinate (within 300 m from the shoreline, as the water shoals to less than 4-m depth).

slowly varying with respect to the surface wave field. Thus, while the estimated velocity amplitudes may be low, the location of nodes and antinodes in space are faithfully reproduced, as are the overall spectral form and peak locations. The large-scale variations in response act to modify slightly the effective spatial window applied before Fourier transformation.

Figure 6 shows 1D frequency spectra (square root power density) and the squared spectral coherence between PADS data and in situ current meter data at two water depths, using just the cross-shore component of velocity {in situ data provided by S. Elgar [Woods Hole Oceanographic Institution (WHOI)], R. Guza [Scripps Institution of Oceanography (SIO)], T. H. C. Herbers [Naval Postgraduate School (NPGS)], W. O'Reilly (NPGS), and B. Raubenheimer (WHOI)}. All spectra were computed from identically sampled 30.72-min time series (1536 samples at 1.2-s interval; current meter data were smoothed and resampled to match the PADS data). Power densities and squared coherences were averaged with two passes of a five-frequency running mean, yielding 10 degrees of freedom. A prominent peak appears near 19.4-s period in all data for this day, representing long swell from the distant Hurricane Erica. Peaks appear near 11 and 14 s, also, incident from nearly the same direction (from the east-southeast) with a broader spectrum at higher frequencies. The offshore spectra (top panel) are from a single site (designated 'spuvt72," at x = 500 m, y = 829 m, in the local coordinates as shown in Fig. 1; note the heavier circles). The shallower-water spectra (middle panel) are averaged over three data sites, increasing the degrees of freedom by a factor up to 3 (depending on the true spatial correlation): "spuvt62" (x = 385 m, y = 828 m), "puvt63" (x = 385 m, y = 815 m), and "spuvt25" (x = 385 m, y = 100 m)y = 778 m). The PADS estimates are multiplied by 1.4 to match spectral densities at the shallower site; this is probably due to biasing of PADS estimates toward zero



FIG. 9. Spectral density of the cross-shore current component vs (vector) wavenumber at (left) 0.018 Hz (54-s period), (center) 0.029 Hz (34-s period), and (right) 0.052 Hz (19.4-s period). Each panel is scaled relative to the maximum value at the given frequency. The orientation of these maps is such that a dark area to the left of the origin indicates energy propagating toward the left. The circles represent linear surface gravity wave dispersion in finite depth (for a mean depth of 5.25 m). The lines indicate the mean direction of the 19.4-s period waves and their reflections; these are repeated unmodified on the plots for the other two frequencies. The variance at all three frequencies appears very well aligned in direction. The lowest-frequency case (left) indicates roughly equal incoming and reflected energy, with both components propagating slowly to the left, consistent with edge waves (either adiabatic or with generation roughly equal to dissipation). The swell case (right) indicates about 20% reflection, consistent with previous analyses for waves of similar frequency at this location (Elgar et al. 1994).

by bottom interference, as discussed in Smith (2002). The PADS sample data were taken about 15 m offshore of the corresponding current meters to avoid the effects of acoustic interference due to strong reflections from the frames holding the current meters. For subsequent illustrations of PADS data, contaminated estimates from all the known frame sites were blanked out, and values interpolated in range to fill these voids (the interpolation



FIG. 10. Linear finite-depth phase and group speeds at 5.25-m depth vs frequency. The wavelength differs from the shallow-water limit by less than 5% for all wave frequencies below 0.12 Hz, and so the dispersion circles (as in Fig. 9) for finite-depth and the shallow-water limit are indistinguishable there. Note, however, that where group velocity is key (as for adiabatic variations in amplitude) finite depth must be considered more carefully.

distance is shorter in range than in angle). The crossspectral coherences between the PADS and in situ data (Fig. 6, bottom panel) are extremely high wherever there is nonnegligible variance density in the spectrum. The coherence levels drop slightly at frequencies below the incident waves. This may be due to stratification and baroclinicity, since the PADS estimates correspond to averages over the top 2-3 m while the spuvt currents are measured about 0.5 m above the bottom. The coherences drop at the highest frequencies as well; this is likely caused by spatial filtering of the PADS data due to the finite sample area (7.6-m range by 7° azimuth).

Insight can be gained from spatial maps of the variance at each frequency (or wave period). For example, maps of variance near 54-, 34-, and 19-s periods show distinct nodes and antinodes, indicating reflection of these long waves off the beach (Fig. 7). The locations of the nodes (as functions of frequency) depend on the bottom topography shoreward of the viewed area. This is true for both incident surf (e.g., for 19-s waves) and for edge-wave (shore trapped) or leaky modes (e.g., 54or 34-s periods). The alongshore uniformity in nodal position seen in Fig. 7 encourages examination of the alongshore-averaged positions versus frequency (Fig. 8). A pattern of nodes and antinodes (dark and light areas) is seen for all frequencies below about 0.07 Hz. The offshore spacing between nodes increases as the wavelength increases (frequency decreases). The curving stripes formed by the nodeantinode pattern are consistent with nearly shore-normal propagation.

The 3D spectra are unconstrained with respect to dis-



FIG. 11. Spectral density vs 2D wavenumber at 0.217 Hz (4.6-s period) for (left) the shorenormal component of velocity and (right) the shore-parallel component, each scaled by the maximum. At this frequency the free-wave dispersion (outer circle) is distinguishable from the shallowwater limit (inner circle) with the present measurement technique. Variance lying on the shallowwater limit curve is probably associated with harmonics of lower-frequency surface waves. Note that the spectral peak centered on the inner circle indicates a propagation direction roughly coincident with the lower-frequency peaks (see Fig. 9). The peaks on the free-wave dispersion curve may be associated with two separate generation systems (from the NE and the SE), or might be generated by localized breaking of longer waves and swell.

persion and can be used to examine whether variance lies near linear or weakly nonlinear surface wave dispersion surfaces, or near dispersion surfaces appropriate to harmonics or other forms of motion. Figure 9 shows three slices from the full 3D spectral variance-density estimates (vs f, k_x , k_y) for the cross-shore component of velocity. Each slice is a map of power density versus vector wavenumber at a fixed frequency: 1) 0.0185 Hz and 2) 0.0294, where edge waves and leaky modes dominate, and 3) 0.0515 Hz, focusing on long swell incident from the distant Hurricane Erica.

At frequencies higher than about 0.1 Hz, the finitedepth dispersion relation for free surface waves diverges from the shallow-water limit (Fig. 10). It is illuminating to examine the relative power density associated with each of these two dispersion surfaces. Figure 11 shows the wavenumber spectra at 0.217 Hz for the cross-shore and alongshore components, with three peaks discernable in each. For either velocity component, one peak is found that lies on the shallow-water limit curve, while the other two lie on that for free waves. A physical interpretation is that the peak on the shallow-water dispersion circle is associated with waves bound to the lower-frequency motions: for example, harmonics of the 12- and/or 19-s period swell peaks. The direction is consistent with this interpretation: the swell peaks are directed from southeast to east-southeast (right panel of Fig. 9), as is the peak on the inner circle of Fig. 11. Wave nonlinearity and harmonics have been investigated previously, in particular via bispectra (Elgar and Guza 1985; Elgar et al. 1995). The bispectral approach permits explicit identification of interacting triads of waves (more than just harmonics), but requires more degrees of freedom for stable estimates. The bispectra can be summed to estimate the net variance due to all such interactions. In contrast, the 3D spectra provide direct estimates of the total variance at each wavenumber and frequency regardless of origin. The peaks on the free-wave dispersion circle in Fig. 11 may be the result of local or nearby wind generation (one set from the northeast, consistent with the weak local wind, the other from the southeast, perhaps linked to one of the two sets of swell from that direction). Alternatively, they could be generated like a localized breaking event as the longer waves constructively interfere (creating a boatlike wake of higher-frequency waves).

The ability to make fine distinctions in the dispersion characteristics, to separate a propagating wave from its 180° reflection, and to simultaneously monitor the underlying mean flow makes this technique appealing for work on waves, currents, and their interactions. Complete 3D spectra for this and other example time periods may be viewed as movies, showing the power densities versus wavenumber at one frequency per frame (available online at http://jerry.ucsd.edu/NSWspectra.html).

6. Discussion

The patterns of nodes and antinodes revealed in the spatial maps of variance at each frequency provide a compelling tool for the identification of modal structure versus frequency. The simple, nearly plane-beach case shown here illustrates an apparently faithful reproduction of partially reflected waves (including edge waves).



FIG. 12. Illustration of time-averaged (\sim 90 s) velocities and the associated vorticity (vertical component). The upper feature 1) resembles a "vortex pair" as is sometimes seen in models. Over a sequence of several minutes duration, this moves through the domain from left to right, along a trajectory indicated by the gray arrow. The strength of this feature is fairly constant as it moves across. The lower feature 2) is a rip current, originating near the gap in the sandbar. The leading edge of this feature propagates to near (350, 900) and then the whole "rip" fades (dissipates). Stronger acoustic backscatter is associated with both features, suggesting that bubblerich fluid is retained and transported. For the weak winds and stratification at the time of the data, baroclinic (3D) flows are expected. It is plausible that the offshore feature is detached from the bottom, and hence from the main source of friction, helping to explain its longevity. (Data interval ending 14 Oct 1997, 0123:31 UTC, at Duck, NC).

Further benefit would come in more complex systems, where the modal structure may not be well known beforehand. An example is where a submarine canyon intersects the shoreline. How effectively coupled are alongshore-propagating and canyon-propagating modes? How does the overall distribution of variance at a given frequency change as the tides alter the sill depths? In addition to spatial variance maps, the phase structure at each frequency can be analyzed, assisting in understanding of the dynamic modes. In the simplest cases, like that discussed above, this provides information about propagation speeds, directions, and reflections. In more complex cases, the 2D structure of the motion can be investigated as a function of frequency; for example, in ways analogous to locating amphidromic points and creating phase maps of tidal motion in ocean basins. This information would be helpful in cataloging and understanding the forms of motion seen in such a combined system.

Densely sampled wave data can also be used to assess wave propagation and evolution. Small-area windows of data may be used to form a mean direction and directivity index (cf. Smith and Bullard 1995) as a function of location (the directivity index divided by phase speed is the appropriate factor to convert mean wave energy and direction into the net wave-momentum vector). The resulting maps of direction and directivity provide strong constraints for models of wave propagation, and sufficient information to evaluate radiation stress and mass transport as functions of location. In shallow water, a few additional independent measurements (e.g., a few pressure sensors) may be needed to calibrate the absolute magnitude (Smith 2002).

Finally, although the examples discussed so far focus mainly on surface gravity waves, the data permit investigations of lower-frequency motions as well. Continuous data collection over weeks at a time, as implemented at SandyDuck, captures motions such as nearshore eddies, freshwater outflows trapped to the shore, and tides. For example, Fig. 12 shows a snapshot of lower-frequency motion (time-averaged with a moving window about 90 s in duration).

The averaging time is long enough to suppress most surface wave motion (including edge waves), but short enough to permit examination of vorticity features embedded in the flow. The offshore feature appears stable, and propagates across the field of view with little change; in contrast, the inshore feature is ephemeral, propagating only a short distance offshore (to about x = 350 m, y = 900 m) before fading beyond detection. At such lower frequencies, baroclinicity can be important, emphasizing the need for auxiliary measurements of the vertical structure; for example, velocity and density profiles at several locations in or near the sampled area (on the day of this example, the density profile at the end of the pier shows a distinct pycnocline). Further analysis of the lower-frequency motions observed at SandyDuck is under way.

7. Conclusions

High-frequency sound (e.g., near 200 kHz) can be used to probe an area several hundred meters on a side with better than 8-m resolution. Receiving the backscattered signal on an array, with digital beamforming, permits the entire area to be sampled every second or less, with under 2 cm s⁻¹ rms velocity error per sample. Radial velocity estimates from two or more systems can be combined to produce time series of horizontal velocity vector maps. Modal structures (e.g., nodes and antinodes) are revealed in spatial maps of horizontal velocity variance as a function of wave frequency. The extensive and dense sampling in time and space invites analysis via 3D Fourier transformation, producing complete directional wavenumber-frequency spectral estimates. Attractions of the technique include unambiguous estimates of propagation speed and direction, objective evaluation of the reflection of wave energy off a beach, ability to distinguish spectral density maxima lying near free-wave versus bound-wave dispersion curves, and spectral estimates free of assumptions concerning the appropriate dispersion relation or physics. For simple interpretation of the results, there must be scatters such as microbubbles in the water, embedded in the flow. This implies that the measurements may be less useful during dead calm conditions, when such bubbles are not generated except within the surf zone.

Acknowledgments. This work was supported by ONR Grant N00014-90-J-1285. I wish to thank R. Pinkel for his role in the conception and development of the phased-array sonar, and E. Slater, M. Goldin, L. Green, C. Neely, C. Bonner, D. Scarla, and A. Aja for their roles in the design, construction, and deployment of the sonar systems. I also wish to thank S. Elgar (WHOI), R. Guza (SIO), T. H. C. Herbers (NPGS), W. O'Reilly (NPGS), and B. Raubenheimer (WHOI) for providing data for comparison.

REFERENCES

- Attia, E. H., and B. D. Steinberg, 1989: Self-cohering large antenna arrays using the spatial correlation properties of radar clutter. *IEEE Trans. Antennas Propag.*, **37**, 30–8.
- Bretherton, F. P., R. E. Davis, and C. B. Fandry, 1976: A technique for objective analysis and design of oceanographic experiments applied to Mode-73. *Deep-Sea Res.*, 23, 559–582.
- Crawford, C. B., and D. M. Farmer, 1987: On the spatial distribution of ocean bubbles. J. Geophys. Res., 92, 8231–8243.
- Elgar, S., and R. T. Guza, 1985: Observations of bispectra of shoaling surface gravity waves. J. Fluid Mech., 161, 425–448.
- —, T. H. C. Herbers, and R. T. Guza, 1994: Reflection of ocean surface gravity waves from a natural beach. J. Phys. Oceanogr., 24, 1503–1511.
- —, —, V. Chandran, and R. T. Guza, 1995: Higher-order spectral analysis of nonlinear ocean surface gravity waves. J. Geophys. Res., 100C, 4977–4983.
- Feddersen, F., R. T. Guza, S. Elgar, and T. H. C. Herbers, 1998: Alongshore momentum balances in the nearshore. J. Geophys. Res., 103, 15 667–15 676.
- Frasier, S. J., and R. E. McIntosh, 1996: Observed wavenumberfrequency properties of microwave backscatter from the ocean surface at near-grazing angles. J. Geophys. Res., 101C, 18 391– 18 407.
- Harris, F. J., 1978: On the use of windows for harmonic analysis with the discrete Fourier transform. *Proc. IEEE*, **66**, 51–83.
- Longuet-Higgins, M. S., and R. W. Stewart, 1964: Radiation stress in water waves; A physical discussion, with applications. *Deep-Sea Res.*, 11, 529–562.
- Miller, K. S., and M. M. Rochwarger, 1972: A covariance approach to spectral moment estimation. *IEEE Trans. Info. Theory*, IT-18, 588–596.
- Paduan, J. D., and L. K. Rosenfeld, 1996: Remotely sensed surface currents in Monterey Bay from shore-based Hf radar (Coastal Ocean Dynamics Application Radar). J. Geophys. Res., 101C, 20 669–20 686.
- Pinkel, R., and J. A. Smith, 1992: Repeat-sequence coding for improved precision of Doppler sonar and sodar. J. Atmos. Oceanic Technol., 9, 149–163.
- Smith, J. A., 1989: Doppler sonar and surface waves: Range and resolution. J. Atmos. Oceanic Technol., 6, 680–696.
- —, 1998: Evolution of Langmuir circulation during a storm. J. Geophys. Res., 103, 12 649–12 668.
- —, 1999: Observations of wind, waves, and the mixed layer: The scaling of surface motion. *The Wind-Driven Air–Sea Interface*, M. L. Banner, Ed., University of New South Wales, 231–238.
- —, 2002: The use of phased-array Doppler sonars near shore. J. Atmos. Oceanic Technol., 19, 725–737.
- —, and G. T. Bullard, 1995: Directional surface wave estimates from Doppler sonar data. J. Atmos. Oceanic Technol., 12, 617– 632.
- Theriault, K. B., 1986: Incoherent multibeam Doppler current profiler performance. Part I: Estimate variance. J. Ocean Eng., OE-11, 7–15.
- Trevorrow, M. V., and D. M. Farmer, 1992: The use of Barker codes in Doppler sonar measurements. J. Atmos. Oceanic Technol., 9, 699–704.
- Young, I. R., W. Rosenthal, and R. Ziemer, 1985: A three-dimensional analysis of marine radar images for the determination of ocean wave directionality and surface currents. J. Geophys. Res., 90C, 1049–1059.
- —, L. A. Verhagen, and M. L. Banner, 1995: A note on the bimodal directional spreading of fetch-limited wind waves. J. Geophys. Res., 100C, 773–778.