

Modeling the Performance of an Acoustic Doppler Current Profiler

T. K. CHERESKIN AND A. J. HARDING

Scripps Institution of Oceanography, La Jolla, California

(Manuscript received 13 January 1992, in final form 15 July 1992)

ABSTRACT

A systematic examination of measurement error in acoustic Doppler current profiler (ADCP) velocity estimates, in the limit of large signal-to-noise ratio, is made using a system model and sonar signal simulations coupled into an ADCP. The model is extremely successful in predicting ADCP performance. The signal simulations provide model validation. Three main sources of error are examined: frequency tracking, measurement variance (inherent variance of pulse-to-pulse incoherent volume reverberation), and measurement bias.

A theoretical lower bound on measurement variance is directly tested by coupling simulated signals into an ADCP. The observed measurement variance is approximately twice the theoretical value and varies as the inverse of the product of the pulse and averaging period (bin). Model predictions of velocity errors for back-to-back beam pairs measuring a sequence of increasing velocity-shear profiles in a medium of randomly distributed scatterers are in excellent agreement with errors measured from simulated signals coupled into an ADCP. Trade-offs between velocity error, vertical and temporal resolution, and maximum range are discussed, with specific focus on optimizing parameters available to users of commercial instruments.

For reasonable parameter choices in low velocity-shear ocean conditions, the predicted error in horizontal velocity from effects considered in this study is $1\text{--}2\text{ cm s}^{-1}$. In large-shear conditions, the predicted error using the same parameters as in low shear is much worse, approximately 10 cm s^{-1} . Optimal parameter choices, however, can reduce the error in large-shear conditions to $1\text{--}4\text{ cm s}^{-1}$.

1. Introduction

As acoustic Doppler techniques for measuring ocean velocities gain widespread acceptance, a need arises to evaluate the measurement technique and to estimate measurement uncertainty. To date, a common evaluation of acoustic Doppler current profilers (ADCPs) has been looking at the measured velocity profiles and deciding how "reasonable" they are (e.g., Firing 1987; Cutchin 1991). Obviously, this method fails to give absolute accuracy estimates. A second form of evaluation has been comparisons with other types of current measurements (Chereskin et al. 1987; Kosro 1985; Pettigrew and Irish 1984). Such verification requires availability of other measurements at the same time and place, and usually requires some assumptions about the homogeneity and stationarity of the flow since the instruments used for comparison often do not sample the same volume of water. Because of changes in the velocity and backscatter characteristics from place to place and over time, it is difficult to make

general accuracy estimates without having many comparisons over which to average. One is also limited by the accuracy of the instrument used for comparison. A third type of evaluation has been made from tow-tank studies (Appell et al. 1988). These studies have provided much information, particularly with respect to beam calibration. Such facilities, however, are expensive and can only be used to test some aspects of system performance. Two important aspects that cannot be easily tested in a tow tank are measurements at far vertical range and in a vertically sheared current.

An alternative approach to evaluating instrument performance is through a system model. Although more complex than a mechanical current meter, the ADCP can be modeled (e.g., Hansen 1985a). A model of one part of the RD Instruments (RDI) current profiler was successful in identifying a source of measurement bias due to the disabling of the feedback control loop (Chereskin et al. 1989). In this paper, the model is extended to include a functional feedback control system, sometimes referred to as a frequency tracker. Hansen's model neglects tracker effects. In addition to the numerical model, a sonar signal synthesizer was developed. The signal synthesizer is a useful tool in its own right for evaluating end-to-end system performance; it allows one to drive a known signal through

Corresponding author address: T. K. Chereskin, Scripps Institution of Oceanography, Physical Oceanography Research Division, 0230, 9500 Gilman Drive, La Jolla, CA 92093-0230.

the system and to test differences between systems. The numerical model, however, has distinct advantages with respect to ease in isolating parts of the system and experimenting with different system parameters. In this paper, results from the model and from signal simulation experiments are used 1) to demonstrate that ADCP performance can be adequately modeled; 2) to present estimates of errors from controlled simulation experiments in the spirit of setting lower bounds on the measurement accuracy that can be achieved in ocean measurements; 3) to identify and quantify predictable biases that may be correctable, a posteriori, in ocean data; and 4) to facilitate the choice, a priori, of system parameters that optimize the trade-off between higher sampling resolution (temporal and spatial) and larger errors.

Sections 2 and 3 present details of the model and the sonar signal synthesizer. In section 4, the model is benchmarked against an ADCP, using the signal simulator to drive the same signal through the ADCP as specified numerically in the model calculations. The benchmark comparisons isolate three aspects of performance: tracking characteristics, measurement variance, and measurement bias. In section 5, model and ADCP performance for one sampling scheme of back-to-back beam pairs measuring a sequence of velocity-shear profiles in a medium of randomly distributed scatterers are compared; the model is then used to compare errors between different sampling schemes for one of the sequence of shear profiles. The discussion of section 6 focuses on improving instrument performance via those parameters that are typically accessible to users of commercial instruments. The particular in-

strument used in this study was an RD Instruments 153.6-kHz pulse-to-pulse incoherent ADCP, an instrument in widespread use on oceanographic research vessels and current-meter moorings. The parameters that can be easily varied (i.e., user friendly) are the pulse length, the bin length, and the ensemble length. Limited choice (possible, although not user friendly) can be made for the low-pass-filter bandwidth and the tracker time constant. Conclusions are presented in section 7.

2. System model

An ADCP sends out a short pulse of high-frequency sound along narrow acoustic beams. The component of velocity along the beam is determined by measuring the Doppler shift in frequency of the echo return, reflected from scatterers assumed to be passively advected by the water motion. Figure 1 is a simplified block diagram of the signal processing in an RDI ADCP. A detailed discussion of the measurement technique and the processing scheme is presented in Chereskin et al. (1989). In this study we are concerned with modeling the frequency estimator and feedback control loop, corresponding to that part of the block diagram enclosed by a dashed line. Prior to entering this loop, the signal is processed by an automatic gain control (AGC) loop. For the purpose of this study, the AGC loop is assumed ideal when the signal-to-noise ratio (SNR) is high; that is, we assume that the AGC loop maintains the signal amplitude at a constant level.

In Fig. 1, the input signal is heterodyned during the mixer-low-pass-filter stage by modulation with in-

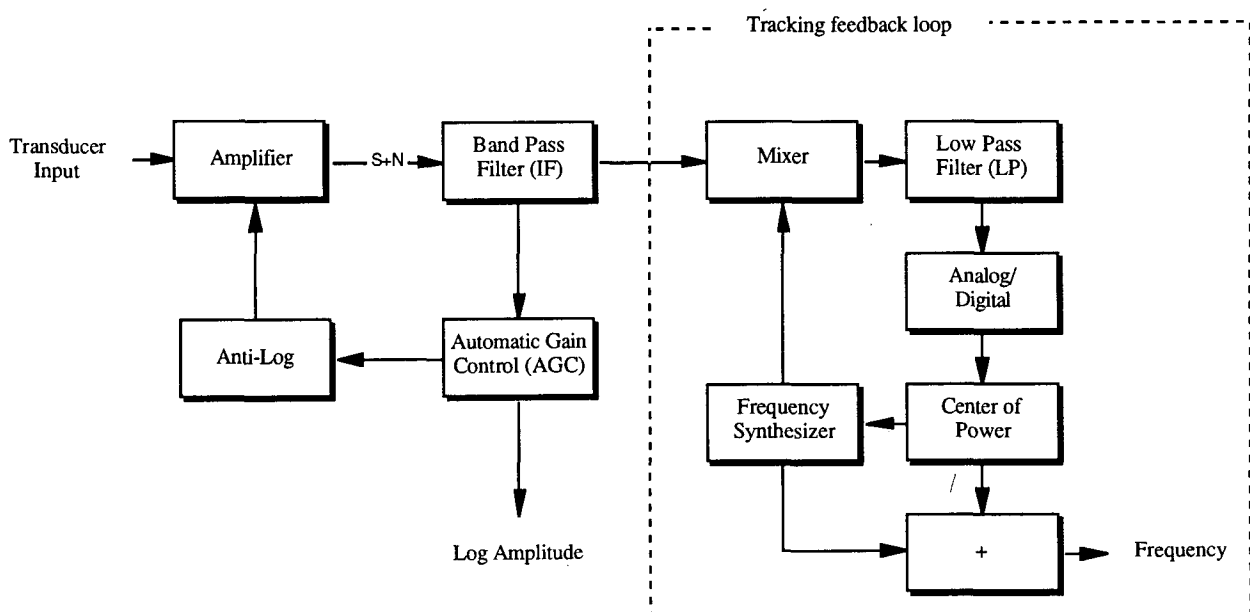


FIG. 1. Block diagram of signal processing in an RDI ADCP. Dashed line encloses the blocks modeled.

phase and quadrature components of a reference signal whose frequency is chosen to maintain the difference (error) frequency near zero. The dynamic adjustment of the reference signal frequency is made by the tracking loop. For the purpose of discussion, it is assumed that the input signal frequency f_i has been mixed to base band, that is, that the transmit frequency has been removed and only the Doppler shift of the transmit frequency remains, $f_D = f_i - f_t$. The frequency of the heterodyned signal (error frequency) is then the Doppler shift minus the reference frequency, $f_e = f_D - f_{osc}$. The purpose of the signal-processing loop is twofold: to determine the appropriate reference frequency and to estimate the error frequency; the sum of these two frequencies forms the estimate of the Doppler shift and hence velocity. In order to calculate the error frequency, the in-phase $I(t)$ and quadrature $Q(t)$ components of the heterodyned signal are combined to form the complex envelope

$$Z(t) = I(t) + iQ(t).$$

The envelope is sampled M times with sampling period t_s , and the discrete, complex autocovariance function $C(t, kt_s)$ at lag k is formed,

$$C(t, kt_s) = \frac{1}{M} \sum_{m=1}^M Z(t + mt_s)Z^*[t + (m + k)t_s],$$

where t is the time relative to the pulse transmission and Mt_s is the averaging time (range bin).

The zero-lag autocovariance ($k = 0$) is the low-pass-filtered signal power, referred to as the energy level. The energy level is used to determine the SNR in the low-pass stage; it must exceed a threshold before a velocity estimate is computed. The one sample-period lag autocovariance ($k = 1$) is used to estimate the first spectral moment of the low-pass-filtered signal (Miller and Rochwarger 1982). The first spectral moment is the mean error frequency,

$$\hat{f}_e = \frac{1}{2\pi t_s} \arctan \left\{ \frac{\text{Im}[C(t, t_s)]}{\text{Re}[C(t, t_s)]} \right\} = \frac{1}{2\pi t_s} \arg[C(t, t_s)].$$

In the limit of large SNR and short sample period, $t_s \ll \sigma_s^{-1}$, where σ_s^{-1} is the bandwidth of the low-pass-filtered signal spectrum, \hat{f}_e is an unbiased estimator of f_e (Plueddemann and Pinkel 1991). Note that the directly measured quantity is the phase change over the sample period, $\arg[C(t, t_s)]$.

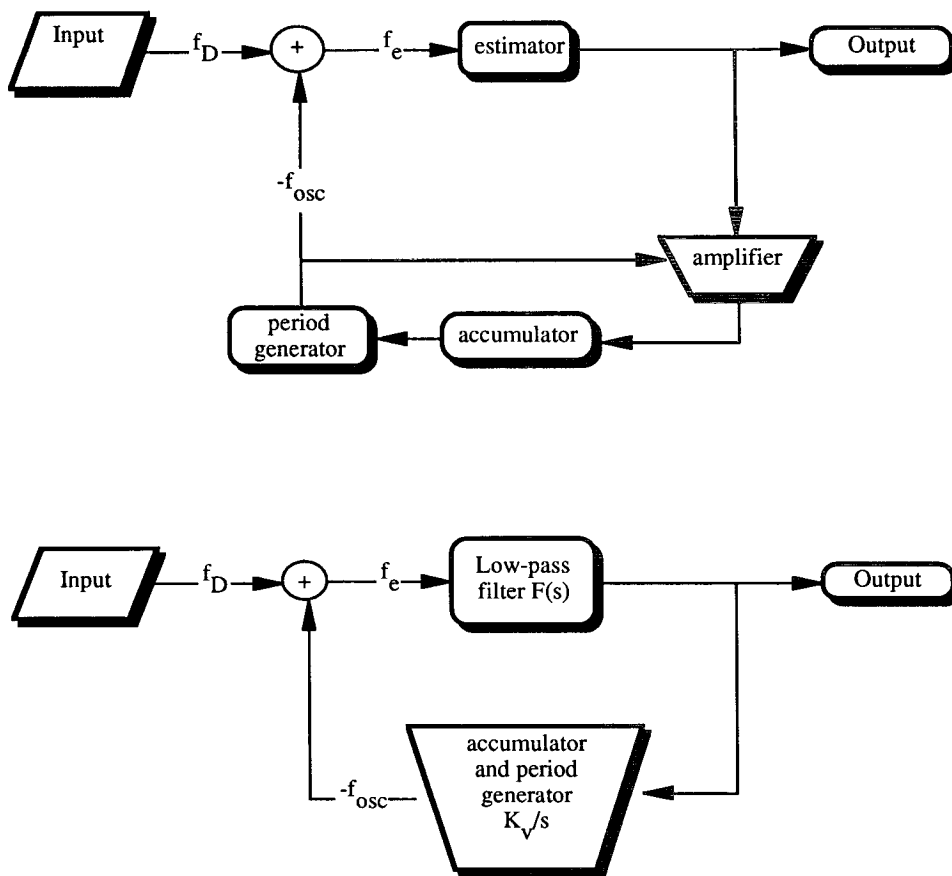


FIG. 2. (a) Model of signal demodulation and frequency-tracking feedback loop corresponding to blocks enclosed in dashed line in Fig. 1. (b) Simple implementation of (a).

The reference frequency is determined by the tracking feedback loop (Fig. 2a). The feedback loop uses the same basic control mechanism as a phase-locked loop. Using Laplace transform notation, a simple implementation of such a loop is shown in Fig. 2b, where s is a complex frequency (the Laplace transform variable), $F(s)$ represents a low-pass filter, K_v is a gain, and $1/s$ denotes integration. The transfer function of the system shown in Fig. 2b is

$$\frac{f_{\text{osc}}(s)}{f_D(s)} = \frac{F(s)K_v}{s + F(s)K_v}.$$

This system can be solved analytically for some simple, illustrative cases. An ideal low-pass filter equals unity in the pass band and zero in the stop band. The simplest case assumes that the filter is absent; that is, $F(s) = 1$, and the gain $K_v(s) = K_0$, a constant. For a linear ramp input frequency $f_D(t) = t$, that is, a constant velocity shear, $f_D(s) = 1/s^2$. Substituting and computing the inverse Laplace transform yields

$$f_{\text{osc}}(t) = t - \frac{1}{K_0} + \frac{1}{K_0} \exp(-K_0 t).$$

After the initial transient decays (for large t),

$$f_{\text{osc}}(t) \approx t - \frac{1}{K_0}.$$

The reference frequency is equal to the Doppler shift, lagged by a fixed amount that is the inverse of the gain K_0 and is termed the time constant for the loop. For this case, increasing the gain improves the response time of the loop by reducing the time constant.

The RDI feedback loop is a discrete version of such a loop that uses a Chebyshev low-pass filter and a time-varying gain. In the model of the RDI feedback loop, the gain K_v has the form of a constant K_0 times factors of 2,

$$K_v = K_0 2^p 2^{q2^n},$$

to reflect the fact that three parameters control the gain and that the gain is always varied by factors of 2. The three parameters are the energy level, the reference frequency, and the basic gain setting (for high energy). In this model, which assumes high SNR (energy), the basic gain and the reference frequency determine K_v . As before, increasing K_v decreases the response time of the loop. The loop response time, however, is not constant for this more complicated case. The tracker does not respond continuously but in discrete steps that are of variable size, and the response time of the loop depends on the varying step size of the tracker. The error phase is allowed to accumulate until it reaches a threshold value. At threshold, the generated period and, therefore, the reference frequency are changed. The generated period changes linearly, but the reference frequency changes nonlinearly due to the inverse relation between frequency and period N ,

$$f_{\text{osc}} = f_{\text{max}} - \frac{c_n}{N}; \quad df_{\text{osc}} = \frac{c_n}{N^2} dN,$$

where f_{max} is the (asymptotic) maximum reference frequency, c_n is a constant, and df_{osc} is the local tracker step size. Therefore, the tracker step size decreases with increasing reference frequency. If K_v were constant, the tracking-loop response time would increase with increasing reference frequency (smaller tracker step sizes). To prevent the response time from varying by more than a factor of 2 over the range of frequencies synthesized, there is a feedback between the period generator and K_v , indicated in Fig. 2a. As the reference frequency reaches switch points defined by a decrease in step size by a factor of 2 from the previous switch point, K_v is increased by a factor of 2 to compensate. We will refer to the response time as the time "constant"; in fact, it lies within a factor of 2 of a fixed value.

A second analytic case that incorporates a realistic filter is required to investigate the stability of the tracker feedback loop. The actual low-pass filter in the ADCP is a second-order Chebyshev filter of the form

$$F(s) = \frac{k}{(s-p)(s-p^*)}, \quad \frac{k}{pp^*} \approx 1,$$

where k is a constant and (p, p^*) are complex conjugate poles in the complex s plane. The transfer function of this analytic system is

$$\frac{f_{\text{osc}}(s)}{f_D(s)} = \frac{kK_v}{s(s-p)(s-p^*) + kK_v}.$$

Asymptotically, the f_{osc} response behaves as in the previous case, following an input ramp with a fixed lag that is inversely proportional to the gain. In this case, however, increasing the gain, in addition to decreasing the time constant, causes two of the poles of the transfer function to migrate into the right half s -plane, producing an unstable system. Roughly speaking, when the time constant is less than the inverse bandwidth of the filter, nondecaying transients dominate. The ADCP's default gain value ensures that its loop is stable and behaves like an overdamped oscillator. It is possible to explicitly change the basic gain setting from its default value by factors of 2. This was done in order to compare the stability of the model and the ADCP; an example of the instability is given at the end of section 5. The gain switching described earlier, based on the reference frequency, is relative to the initial prescribed gain.

The error frequency, $f_e = f_D - f_{\text{osc}}$, is estimated in the feedback control loop. The estimate of the Doppler shift is $\hat{f}_D = f_{\text{osc}} + f_e$. The function of the feedback control loop is to minimize f_e , so that the error signal has a much narrower bandwidth than the input signal, and a narrower processing filter can be used to estimate f_e than would be required to estimate f_D . Use of a nar-

rower processing filter increases the SNR and allows the ADCP to profile to longer ranges.

The error frequency is used to estimate the component of water velocity along the acoustic beam, \hat{v}_r , through the Doppler relation,

$$\hat{v}_r = \frac{\hat{f}_D c}{2f_i} = \frac{(f_{osc} + \hat{f}_e)c}{2f_i},$$

where c is the speed of sound in seawater, nominally 1500 m s^{-1} . The errors in the estimate of velocity, directly related to frequency through the Doppler equation, are the chief interest, and one of the main purposes of the model is to predict instrument performance (velocity error) in different flow regimes. There are many possible sources of error, and due to complexity, only those deemed most significant are included in the model. The incoherent Doppler technique assumes that the scatterers of the sound, primarily plankton, act as passive tracers of the water motion. Active motion (e.g., swimming) will introduce error into the water-velocity estimate. The method of estimation, employing intrapulse covariance calculations, requires correlated reverberation over the averaging interval employed. Uncorrelated reverberation introduces random error into the velocity estimate. Background noise, such as from the instrument's receivers and other sources such as from a ship (for a hull-mounted instrument), degrades the estimate. These error sources apply to a broad class of pulse-to-pulse incoherent Doppler sonars (Hansen 1985a; Chereskin et al. 1989; Plueddemann and Pinkel 1991). In addition, feedback-loop characteristics, such as the time constant of the loop and the frequency resolution of the period generator, together with characteristics of the signal-processing filter, also affect the accuracy of the velocity estimate.

There are really two models used in this paper. The results of the first model were used to create a parameterized tracking model that runs efficiently on a personal computer. The first model, which does not track, is a direct simulation of the signal processing of the ADCP. It uses a statistical model of the backscattering to produce base-band (modulation) data that is mixed with (constant) Doppler shifts, filtered, sampled, and used to calculate the autocovariance and hence velocity and velocity error. This model was used to verify our understanding of the ADCP signal processing and to parameterize single-ping skew error as a function of error frequency, transmit frequency, pulse, and bin length (results of section 4c). The second model tracks, but is not a direct simulation of the ADCP because it does not mix base-band data with Doppler-shifted frequencies. Rather, the tracking model parameterizes the effect of the base-band data on the true Doppler shift as a random variable, resulting in random error in the velocity estimate, and it predicts the skew error from the direct simulation results. The estimate of the error frequency, the true error frequency plus random and skew errors, is used in the tracking feedback loop. The

variance of the random variable is the velocity variance observed in measurements made by an ADCP. The ADCP was driven by signal simulations employing the same statistical scattering model as in the nontracking direct simulation of ADCP processing. The statistical scattering model is described in section 3; the variance determination is described in section 4b. All results except those in section 4c employ the parameterized tracking model. Comparison of output from the two numerical models and from an ADCP verified the validity of the parameterized model results.

The numerical models include a third source of error, noise bias, parameterized in terms of SNR, as described in Chereskin et al. (1989). This error, however, is insignificant for the high SNR results of this study.

3. Sonar signal simulator

A programmable sonar signal simulator was developed in order to test ADCPs under controlled velocity-shear and backscattering conditions. The simulator developed for this study generates an electrical waveform that corresponds to the echoes of the acoustic backscattered field that would be observed by an ADCP after it transmits a given waveform. The echoes are computed from a numerical model that simulates variations in backscatter strength and velocity shear. A hydrophone converts the electrical waveform to sound that is coupled into an ADCP transducer.

The numerical model follows the single-beam model of Theriault (1986) and is presented in the complex envelope notation of Van Trees (1971). In this notation, a narrowband signal $s(t)$ is expressed as

$$s(t) = 2 \text{Re}[\tilde{s}(t)e^{i\omega t}],$$

where ω is the center angular frequency and $\tilde{s}(t)$ is the complex envelope of $s(t)$. The ADCP transmit signal is modeled as $d(t)$,

$$d(t) = 2 \text{Re}[\tilde{d}(t)e^{i\omega_i t}],$$

with center frequency $\omega_i = 2\pi f_i$ and complex envelope $\tilde{d}(t)$:

$$\tilde{d}(t) = \sqrt{E}\Pi(t/T),$$

where Π is a boxcar function of unit amplitude and duration T seconds; that is,

$$\Pi(x) = \begin{cases} 1, & \text{if } 0 \leq x \leq 1; \\ 0, & \text{otherwise.} \end{cases}$$

The energy of the transmitted waveform is thus ET and is proportional to the pulse length T .

The properties of the medium—scattering response and radial velocity (Doppler shift)—are expressed as functions of distance along the beam (range) in units of the two-way travel time. A feature at range r will occur at time $\zeta = 2r/c$. The reflection response of a volume of scatterers to a narrowband pulse is repre-

sented as a superposition of the response of the individual scatterers; the superposed reflection response can be characterized as a complex zero-mean Gaussian random variable (Theriault 1986; Van Trees 1971). It is assumed that within a range increment all the scatterers move at the same radial velocity (constant Doppler shift over a range increment). The medium response to a scattering increment at ζ_0 is written

$$\sqrt{E} \tilde{b}(\zeta_0) \Pi\left(\frac{t - \zeta_0}{T}\right) \exp[-i\omega_D(\zeta_0)t] \Delta\zeta,$$

where $\tilde{b}(\zeta_0)$ is a zero-mean complex Gaussian random variable that specifies the reflection response at range ζ_0 , and ω_D is the Doppler shift associated with the increment. At any instant in time, the backscattered signal is received from a range of $cT/2$, centered at $r = ct/2 - cT/4$. The response of the entire medium is expressed by superposition of the increments of the ensonified volume

$$\begin{aligned} \tilde{s}(t) &= \sqrt{E} \int_{-\infty}^{\infty} \Pi\left(\frac{t - \zeta}{T}\right) \tilde{b}(\zeta) \exp[-i\omega_D(\zeta)t] d\zeta \\ &= \sqrt{E} \int_{t-T}^t \tilde{b}(\zeta) \exp[-i\omega_D(\zeta)t] d\zeta. \end{aligned}$$

Writing $\tilde{b}(\zeta) = \tilde{B}(\zeta) \exp\{-i[\phi(\zeta)]\}$ where \tilde{B} and ϕ are Rayleigh and uniform random variables, respectively, and splitting ω_D into the sum of a volume-averaged Doppler shift ω_D^0 and the spatially varying shift over the volume Ω , yields

$$\begin{aligned} \tilde{s}(t) &= \sqrt{E} \int_{t-T}^t \tilde{B}(\zeta) \\ &\quad \times \exp\{-i\{[\omega_D^0 + \Omega(\zeta)]t - \phi(\zeta)\}\} d\zeta. \end{aligned}$$

The echo-simulator software models the envelope of the echo return as the following convolution sum:

$$\tilde{s}(t) = \alpha(t) \sqrt{E} \exp(-i\omega_D^0 t) \int_{t-T}^t \tilde{b}(\zeta) d\zeta,$$

where $\alpha(t)$ is an (ad hoc) representation of transmission loss due to spherical spreading and absorption. (For the high SNR assumption of this study, $\alpha = 1$.) This model neglects the contribution

$$\sqrt{E} \int_{t-T}^t \tilde{B}(\zeta) \exp[-i\Omega(\zeta)t] d\zeta,$$

which represents the amplitude modulation of the spatially varying Doppler shift over the integration interval. The model could easily be modified to include this effect by retaining the Doppler-shift term within the convolution integral. A significant advantage of the simplification, however, is the separation of variables: amplitude, frequency, modulation. Controlling any two of these while varying the third isolates effects due to the variable of interest. Varying amplitude allows ex-

amination of SNR dependence; varying frequency produces different ocean-velocity regimes; and modulation produces velocity variance due to the assumed statistics of the scatterers. Hence, the model does not claim to be an accurate representation of the ocean-scattering environment but rather a useful conceptualization of some of the effects. The Gaussian model creates a narrowband signal modulated by an envelope with random amplitude and phase. The modulation results in the Doppler spreading of the resultant spectrum (Van Trees 1971). The spectrum has a center frequency equal to the sum of the transmit and the Doppler frequency, and a finite bandwidth due to the modulation. This property of Doppler spreading is fundamental to incoherent pulsed Doppler systems. The central estimation problem is that of determining the mean frequency of such a spectrum. Therefore, simulation of the Doppler spread was considered a key aspect of the model. Although the design criteria proposed by Belliveau and Whitman (1989) included the ability to generate signals of prescribed bandwidth, their prototype simulator did not include this feature. Since the simulator is software driven, other models of interest can be developed and implemented in the future.

Four sets of data are required to generate the acoustic echo. Referring to the random access memory (RAM) arrays indicated in Fig. 3, the required datasets are the real and imaginary components of the modulation signal (cosine and sine RAM), the frequency data of the backscattered echo (Doppler RAM), and the attenuation data. The frequency of the backscattered signal is the transmit plus the Doppler frequency. The transmit frequency is a constant. The Doppler frequency results from the component of water velocity along the beam, converted to frequency via the Doppler relation. The frequency is integrated to form a 16-bit phase; the phase is input to a 22-bit numerically controlled oscillator, implemented by the digital signal processor. Table 1 lists the specifications of the oscillator. The base-band data are computed as the convolution sum of the transmit pulse with a zero-mean complex Gaussian random variable. By varying the start values of the sine and the cosine data in random fashion, a background scattering field with the assumed statistics can be generated without loading new base-band data into RAM for each ping.

The simulator is a sampled data system with sample frequency $f_s = 4f_t = f_c/N$, where f_c is the simulator master-clock frequency. A master-clock divider ratio $N = 16$ yields a transmit frequency of 153.603 kHz. The desired transmit frequency was about 3 Hz less and was achieved by a (constant) phase-rate correction in the signals calculated by the numerical model. To reduce data storage requirements, the input base band and frequency data are specified at the coarser sample interval of $2/f_t$ and then interpolated to the desired sample interval of $1/4f_t$. The interpolation of the base-band data is performed by zero padding and filtering

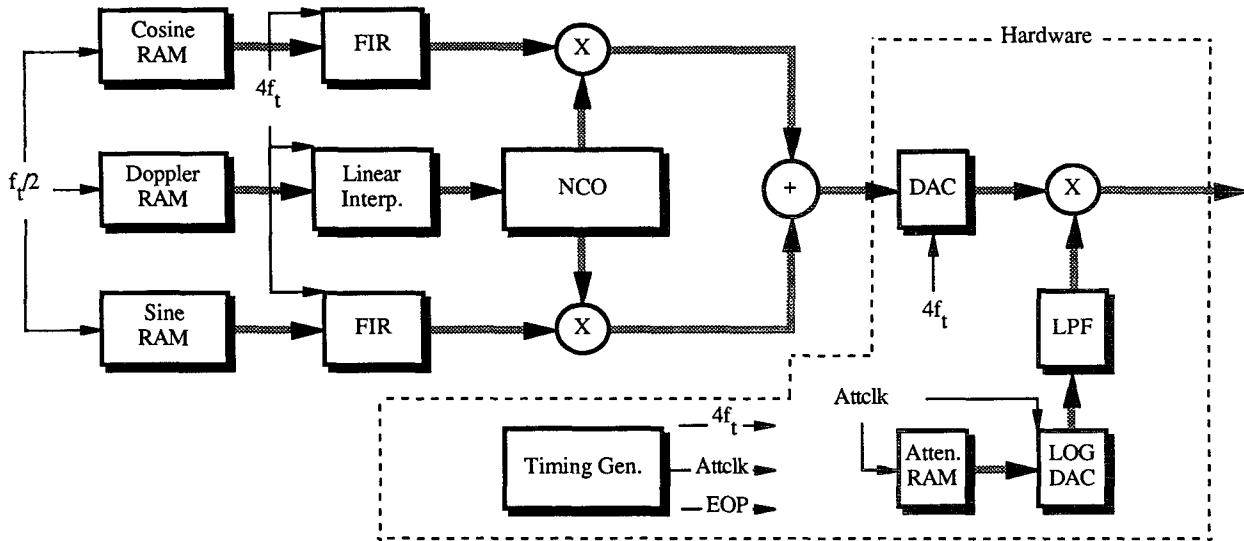


FIG. 3. Block diagram of echo signal simulator.

with a 40-tap low-pass finite-impulse response (FIR) filter. The frequency data are linearly interpolated. The maximum number of data samples (after interpolation) is 262 144. For a transmit frequency of 153.6 kHz, this corresponds to a maximum ping length of 0.4267 s or, using two-way travel time, a maximum range of 320 m.

The interpolated sine and cosine base-band data are multiplied by the interpolated sine and cosine frequency data and added to form the echo signal. The system is triggered by detecting an outgoing pulse from an ADCP. When triggered, the echo data stored in RAM are converted to an analog waveform by a digital-to-analog converter (DAC). The attenuation data are input to a logDAC that generates an analog waveform corresponding to transmission loss. The echo waveform is multiplied by the transmission loss waveform, and the product is converted to sound by a small hydrophone. The hydrophone is coupled to an ADCP transducer face with silicon grease.

4. Model benchmarks

The model was benchmarked against an actual system in three areas: tracking feedback loop performance, velocity measurement variance, and velocity measurement bias. The system used for comparison was an

RDI 153.6 kHz ADCP. Frequency was converted to the component of velocity along the beam (radial velocity, v_r) using the Doppler relation. Range was converted to depth using an assumed beam angle θ of 30° with respect to vertical. The maximum range of 320 m corresponds to 277 m in depth. For the assumed beam angle, horizontal velocity v_h is twice the magnitude of the radial velocity v_r , $v_h = v_r/\sin\theta$. Unless stated otherwise, pulse and bin lengths of 8 m (12.32 ms) were used. The bin length is the time interval over which the covariance is averaged before estimating velocity; it is frequently chosen equal to the pulse length. In this study, pulse and bin lengths are specified in seconds or vertical meters, depending on the context. Two-way travel time and the beam angle are used to convert from τ seconds to z vertical meters, $z = c\tau \cos\theta/2$. These lengths represent the vertical dimension of the sampled volume.

a. Tracking feedback-loop performance

The synthesized input signal to the ADCP was a narrowband signal, corresponding to $\alpha = 1$ (no attenuation) and $\hat{b} = 1$ (no velocity variance from modulation) in the previous discussion, that is, the convolution of a sine wave with a boxcar pulse. Similarly, the model assumed large SNR and no velocity variance. The explicitly modeled errors (random, skew, noise) were zero for this case, and the principal interest was to verify the model tracking algorithm. An error due to the group delay of the processing filter was apparent in both model and ADCP estimates, however, and will be discussed at the end of this section.

Figure 4 compares the model and ADCP output velocities with the specified input velocities; the model and ADCP tracker positions are also shown. Extreme

TABLE 1. Numerically controlled oscillator specifications.

Frequency resolution	$9.537 \times 10^{-7} f_i$ ($=0.1465 \text{ Hz}, f_i = 153.6 \text{ kHz}$)
Velocity resolution	$7.153 \times 10^{-4} \text{ m s}^{-1}$
Frequency range	$\pm 3.125 \times 10^{-2} f_i$ ($= \pm 4.8 \text{ kHz}, f_i = 153.6 \text{ kHz}$)
Velocity range	$\pm 23.438 \text{ m s}^{-1}$

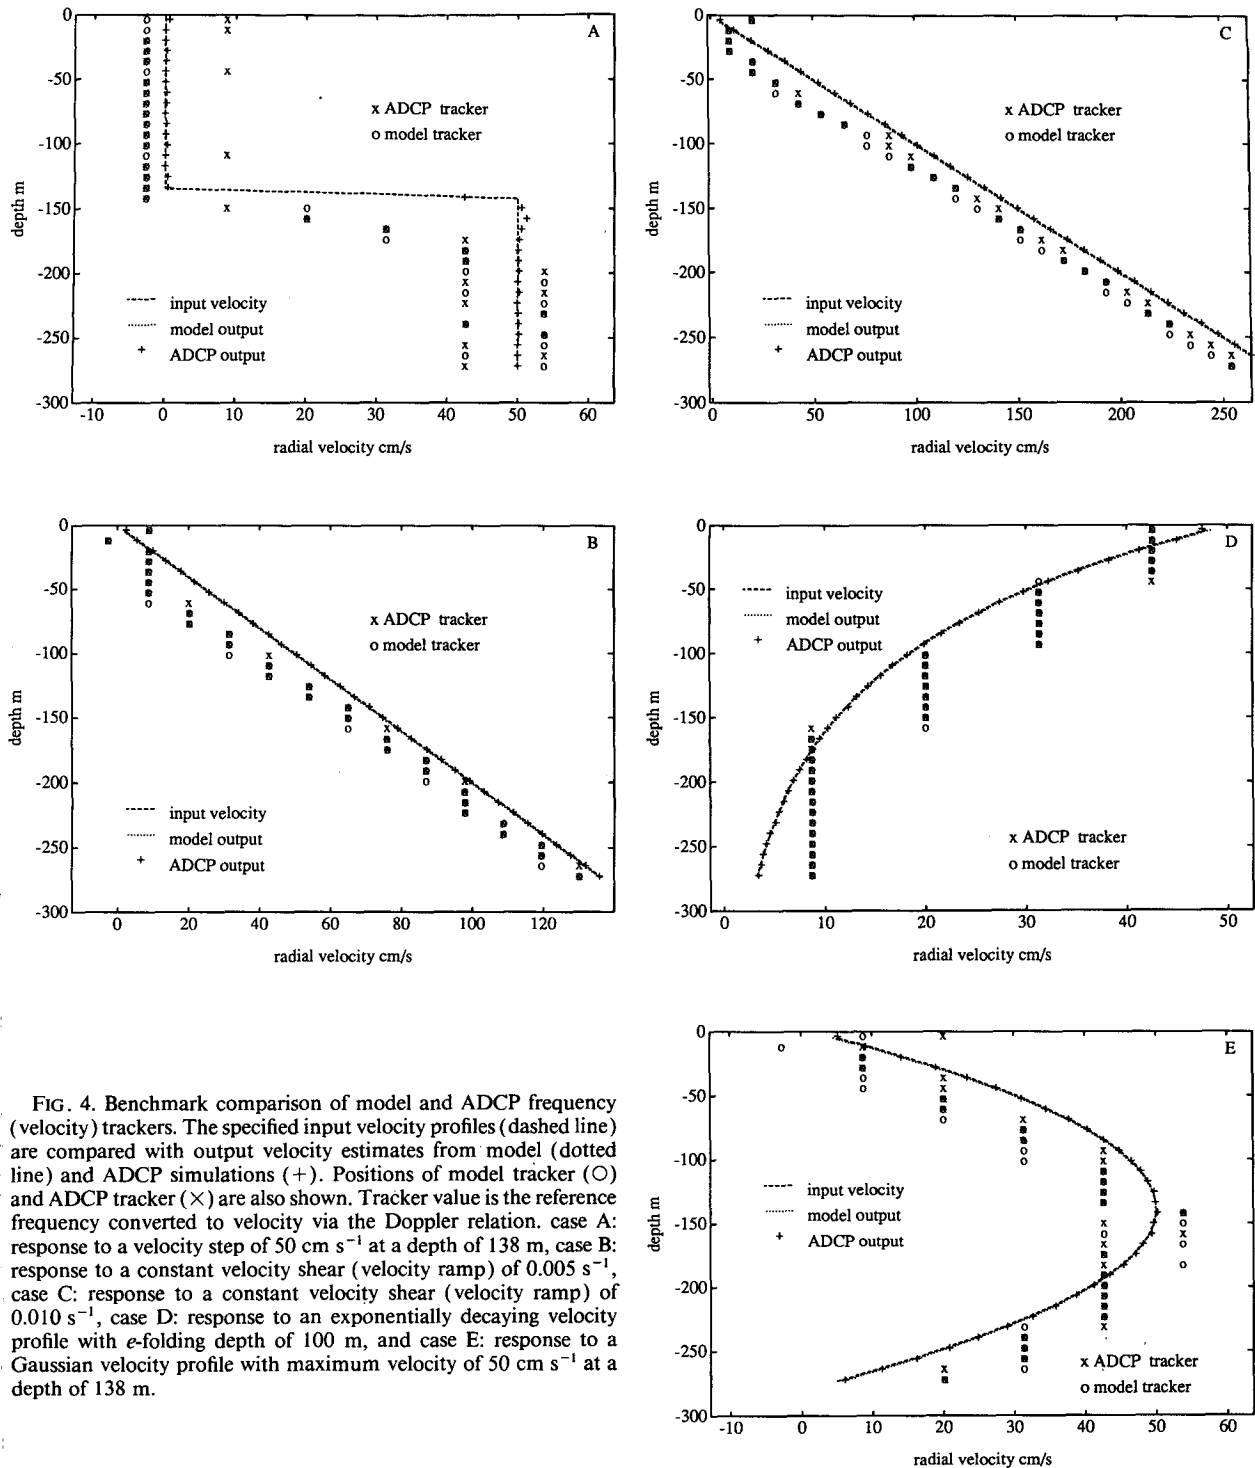


FIG. 4. Benchmark comparison of model and ADCP frequency velocity trackers. The specified input velocity profiles (dashed line) are compared with output velocity estimates from model (dotted line) and ADCP simulations (+). Positions of model tracker (O) and ADCP tracker (X) are also shown. Tracker value is the reference frequency converted to velocity via the Doppler relation. case A: response to a velocity step of 50 cm s^{-1} at a depth of 138 m, case B: response to a constant velocity shear (velocity ramp) of 0.005 s^{-1} , case C: response to a constant velocity shear (velocity ramp) of 0.010 s^{-1} , case D: response to an exponentially decaying velocity profile with e -folding depth of 100 m, and case E: response to a Gaussian velocity profile with maximum velocity of 50 cm s^{-1} at a depth of 138 m.

changes in input velocity were chosen for two reasons: to compare the trackers over a wide range of Doppler shifts, and to verify that the largest expected physical signal would be tracked. Four different idealized velocity profiles were used: a step, a ramp, an exponential, and a Gaussian. These are fairly standard test inputs;

in the ocean they also correspond to reasonable model velocity profiles. A slab mixed-layer model assumes a step velocity profile; the surface Ekman-layer velocity predicts an exponential decay of velocity with depth; the shallow (between the surface and 250 m) equatorial currents (South Equatorial Current overlying the

Equatorial Undercurrent) have a Gaussian profile shape; regions where velocity changes rapidly with depth are often approximated by a constant velocity shear (ramp).

Although the shapes of these profiles are reasonable approximations to ocean profiles, it must be emphasized that the magnitude of the velocity changes are extreme values of what might be observed in the ocean. In particular, the size of the velocity step, 0.5 m s^{-1} (1 m s^{-1} in horizontal velocity), represents a nonphysical, discontinuous change in velocity (Fig. 4a). The shear (velocity change with depth) in the maximum rate ramp profile (Fig. 4c) is observed in highly sheared ocean currents such as the Gulf Stream, but is a factor of 4 larger than typical ocean shear. The velocity changes for the Gaussian profile (Fig. 4e) are typical of the equatorial currents, but like the Gulf Stream, these currents are among the most highly sheared flows that exist in the ocean.

The model output velocity profile was computed from a single ping. (Since there was no velocity variance, multiple pings yielded the same result.) For the ADCP, the output velocity profile was an average of ten pings in order to reduce any (nonmodeled) system noise that may have been present. The ADCP and model used the same algorithm for bin-to-bin tracking. Due to an oversight, they did not use identical algorithms for first-bin positioning (ping-to-ping tracking). Although both used the average velocity over the initial 10% of the profile (32 m), the ADCP also used a geometric weighting of previous pings to choose the initial tracker position (the n th previous ping was weighted by 2^{-n}). Nonetheless, their initial positions agreed most of the time. The ADCP was configured with a dummy power amplifier so that it could not actually transmit a waveform; synchronization between the synthesizer and ADCP was achieved by using the outgoing pulse signal of the ADCP electronics to trigger the simulator card to transmit. The ADCP began processing the signal immediately after transmit (no blanking period).

Overall, the positions of the trackers agreed remarkably well and indicated a similar response time to velocity changes (time constant). For the RDI feedback control loop, there is a feedback between the period generator and the amplifier (Fig. 2) that occurs when the tracker value reaches 43 cm s^{-1} . The feedback consists of increasing the gain in the loop by a factor of 2, resulting in a halving of the time constant and thus faster tracker response. This change in tracker response is evident in the ramp input cases (Figs. 4b,c). The tracker positions are noticeably farther from the profile when the tracker is less than 43 cm s^{-1} . After the loop gain is doubled, the error phase accumulates more quickly, resulting in a shorter response time of the trackers. By about 80 cm s^{-1} , the average tracker misposition is less than observed in the initial (shallow) part of the profile. As previously discussed, the tracker time constant depends upon the loop gain and the

tracker step size, which in turn depends on the reference frequency. The change in gain at predetermined switch points returns the tracking loop time constant to its value at the previous switch point; the time constant increases by a factor of 2 between switch points.

The signed error profiles, (ADCP – input) and (model – input), are shown in Fig. 5. The standard error of the ADCP sample mean (standard deviation divided by square root of the number of samples, 10) is used as the uncertainty in the ADCP error. The model error for the velocity step, case A, is exactly zero. The ADCP error profile shows a large error in the bin just below the depth of the step, 138 m, and an overshoot error in the next bin that the model does not duplicate. The errors over the ADCP profile away from the step appear small and randomly distributed. Although not identically zero, they are negligibly small, and the model and ADCP agree reasonably well. For the error profiles of cases B–E, a bias is evident that is proportional to the derivative of the velocity profile, that is, constant for the ramp profiles, exponential for the exponential profile, and linear for the quadratic profile. This apparent bias in velocity is due to the group delay of the low-pass filter; the system output is a delayed version of the input, and this delay needs to be accounted for in the depth algorithm. For the ADCP, the group delay is inversely proportional to the bandwidth of the filter and also depends on the size of the error frequency; at zero error frequency and 300-Hz bandwidth, the lag is approximately 1.1 m. In the model of the RDI processing filter, the lag is exactly 1.1 m. When the filter width is doubled (600 Hz), the lag is halved (0.55 m).

Figure 6 shows the difference profiles for the quadratic case (E, Figs. 4e and 5e) using a 1.1-m-depth lag correction. This correction improved the comparison for cases B–E, and the appropriate lag correction will be used in all subsequent comparisons in this study. Error averaged over 34 depth bins for cases A–E is presented in Table 2. The error in velocity is defined as the difference between the modeled or measured velocity and the true input velocity, lagged by the indicated amount. For cases B–E the average error is less than 0.25 cm s^{-1} . The root-mean-square difference is less than 0.33 cm s^{-1} . The errors are negligibly small when the lag correction is used.

b. Estimate variance

The complex modulation of the echo return by the scattering volume produces variance in the velocity measurement. In this section, an estimate of velocity variance from simulations ($\alpha = 1$, $\omega_b^2 = 0$, $\tilde{b} = \text{Gaussian}$) is compared to theoretical predictions. The observed variance is least-squares fit to a function that varies as the inverse of the pulse–bin product. This fit specifies the velocity variance used in the model. The functional form corresponds to the Cramer–Rao (CR) lower bound in the asymptotic limit of large SNR.

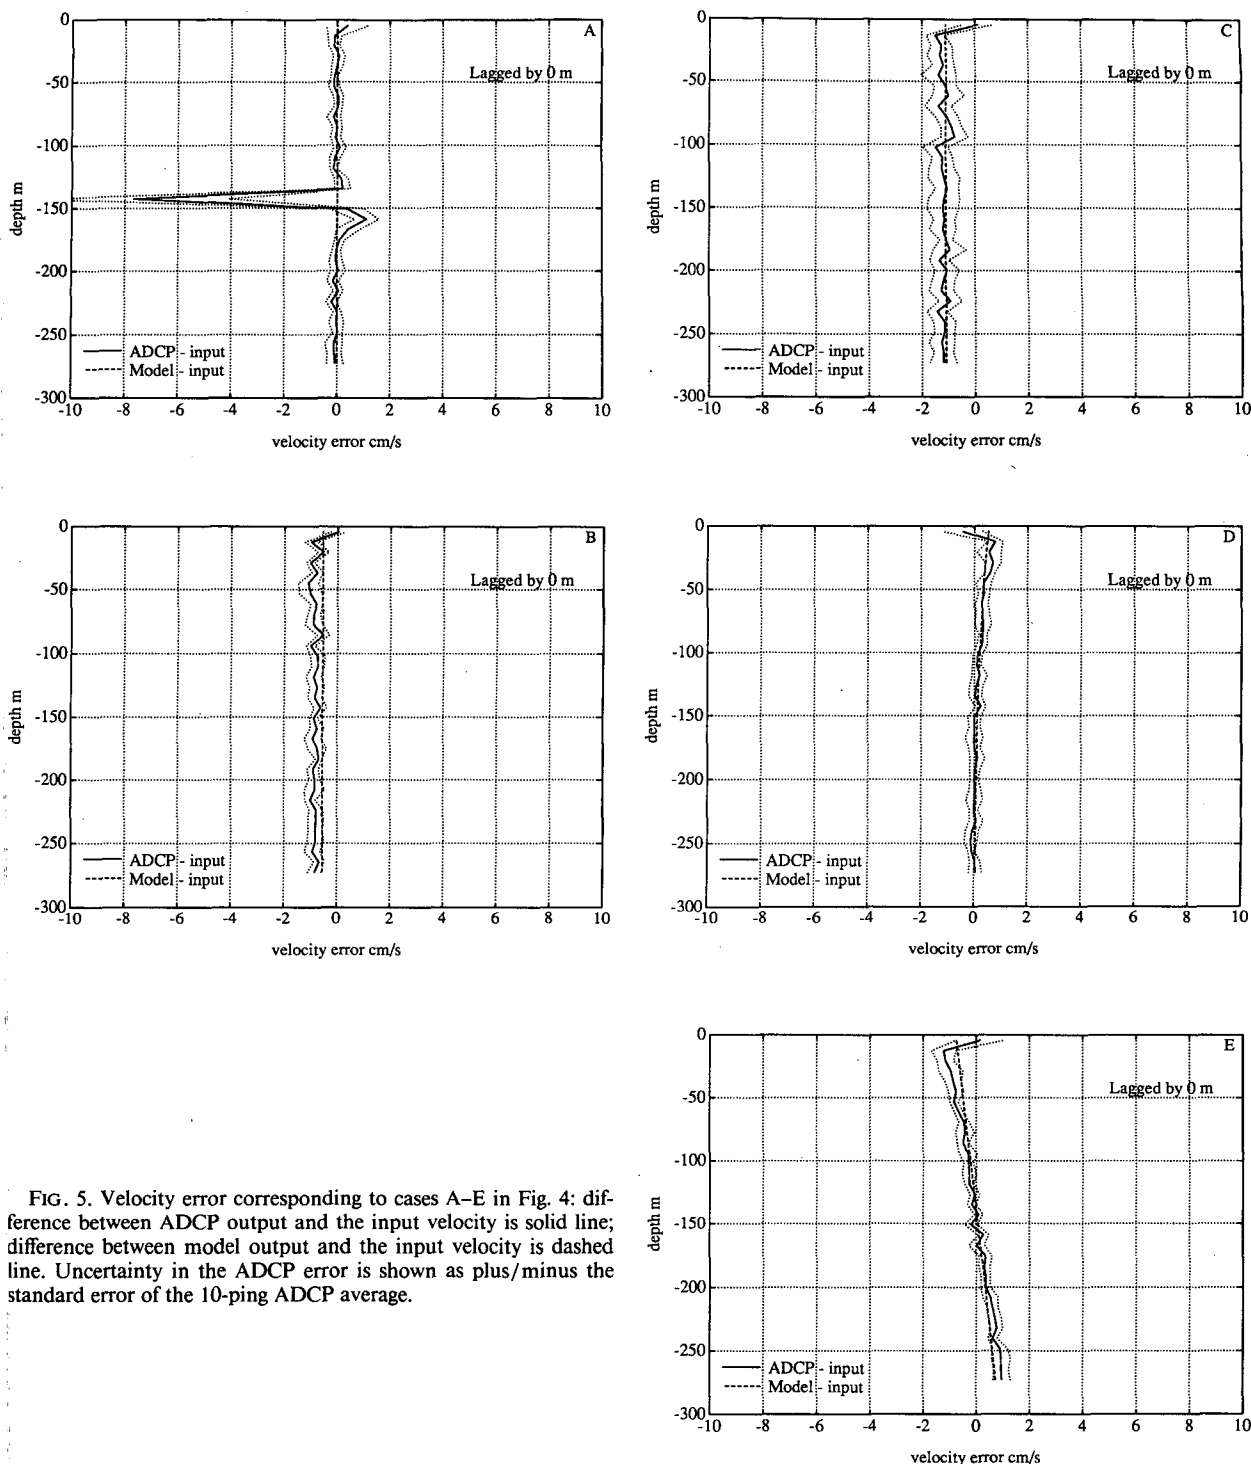


FIG. 5. Velocity error corresponding to cases A–E in Fig. 4: difference between ADCP output and the input velocity is solid line; difference between model output and the input velocity is dashed line. Uncertainty in the ADCP error is shown as plus/minus the standard error of the 10-ping ADCP average.

Theriault (1986) derives the CR bound for the case where the pulse T_p and bin T_a time intervals are equal:

$$\sigma_v^2 = \text{var}^T = \left(\frac{c}{2\omega_t T_p} \right)^2.$$

Following Theriault's analysis, Smith (1989) derives a CR bound that considers discrete sampling and unequal bin and pulse lengths. If the sampling interval is short compared to the pulse interval, Smith's predicted velocity variance is

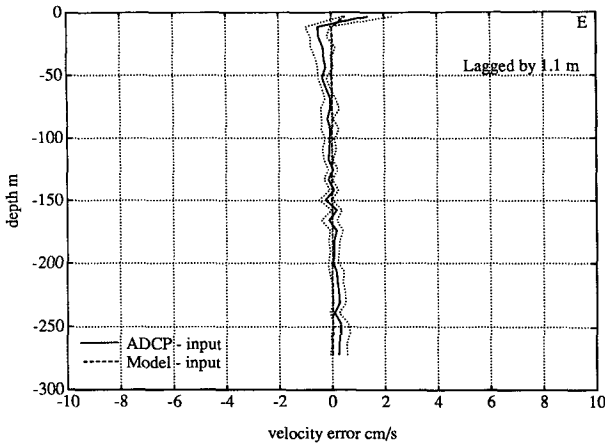


FIG. 6. Velocity error as in Fig. 5e except depth of output velocities has been corrected by the group delay of the filter (1.1 m) prior to subtracting the input velocity.

$$\sigma_v^2 = \text{var}^S = \left(\frac{c}{2\omega_l T_p} \right)^2 \left(\frac{T_p}{T_a} - \frac{T_p - T_a}{2T_p - T_a} \right)^{-1}$$

In the limit of bin = pulse length, Smith's estimate recovers that of Theriault.

Since the CR bound is a lower bound for the variance of an unbiased estimator, the actual variance is expected to be larger than the CR prediction. Estimates of velocity variance from ocean data indicate that the observed variance is about a factor of 2 larger than the CR prediction (RD Instruments 1989; Hansen 1985b). Smith and Pinkel (1991) derive an alternative variance estimate for incoherent Doppler sonar data that is based on the error inherent in the covariance estimation technique (Miller and Rochwarger 1972). In the limit

TABLE 3. Estimate variance (based on 100 ping averages, 153.6 kHz, number of independent range bins $M = 272$ m/pulse).

Pulse (m)	Pulse (ms)	Bin (m)	Number of independent estimates	Mean (cm s ⁻¹)	Standard deviation (cm s ⁻¹)
pulse = bin					
4	6.16	4	6793	-0.13	18.02
8	12.32	8	3400	0.02	10.37
16	24.63	16	1700	0.04	5.37
32	49.27	32	800	-0.24	3.32
pulse > bin					
8	12.32	4	3377	-0.04	14.24
16	24.63	4	1699	0.35	12.59
32	49.27	4	799	-0.13	7.66

of uncoded pulses, such as used in this study, their predicted variance is

$$\sigma_v^2 = \text{var}^{SP} = \left(\frac{c}{2\omega_l T_p} \right)^2 \left(\frac{T_p + T_a/2}{T_a} \right)$$

The predicted error of the covariance estimator falls between the CR bound and the observed standard deviation. In the limit of bin = pulse length, the covariance estimator error is $\sqrt{1.5}$ times the CR bound.

Table 3 presents the random error (standard deviation) observed in single-ping velocity estimates made by an RDI 153.6-kHz ADCP, driven by a Gaussian-distributed random modulation of a zero-Doppler-shift input signal, synthesized by the echo simulator. The signal was not attenuated ($\alpha = 1$), and since zero Doppler shift was close to the default tracker position, there was virtually no tracking. (The tracker does not move; however, due to the discrete frequencies that can be synthesized, there is a small misposition of the tracker for zero Doppler shift.) Pulse and bin lengths were varied. The pulse length determined the vertical resolution; for example, for a 16-m pulse and a 4-m bin, every fourth velocity estimate is independent. Only independent estimates were used to estimate the mean and standard deviation. A hundred pings were collected for each case; the upper bound on the number of independent estimates is therefore 100 times the number of independent bins. The ADCP, however, excluded some velocity estimates on a per ping and per bin basis, based on its energy criterion.

For these simulations, skew and noise bias errors should be negligible, and this appears to be the case. The means lie within ± 0.35 cm s⁻¹ of the true mean and indicate no particular bias. A comparison of the measured standard deviations with that predicted by Theriault (1986), Smith (1989), and Smith and Pinkel (1991) is presented in Fig. 7. Theriault's curve (dashed) assumes the abscissa is the pulse length, that is, that the bin = pulse. This curve is also plotted offset by 3.1

TABLE 2. Velocity error estimates: Unmodulated Doppler shifts (based on 10-ping averages: 153.6 kHz, 8-m bin and pulse, 34 bins).

Case	Mean error		rms error	
	ADCP (cm s ⁻¹)	Model (cm s ⁻¹)	ADCP (cm s ⁻¹)	Model (cm s ⁻¹)
0-m lag				
A (step)	-0.18	0.00	1.34	0.00
B (ramp 1)	-0.80	-0.55	0.83	0.55
C (ramp 2)	-1.15	-1.10	1.18	1.10
D (exp)	0.16	0.19	0.29	0.23
E (quad)	-0.01	0.00	0.63	0.44
1.1-m lag				
A (step)	-0.18	0.00	1.34	0.00
B (ramp 1)	-0.25	0.00	0.32	0.00
C (ramp 2)	-0.05	0.00	0.27	0.00
D (exp)	-0.04	-0.01	0.20	0.01
E (quad)	0.02	0.03	0.33	0.09

cm s^{-1} , the root-mean-square difference between the observations and the CR bound. The solid curve is the least-squares fit to the measurements; it was chosen to have the same shape as the CR prediction. It is a first-order fit to $1/\sqrt{T_p T_a}$ and has a larger slope and offset than the CR bound. At small abscissa values, less than about 15, the observed standard deviation is about 40% larger than the CR bound. At large values, greater than about 30, the standard deviation reaches a higher noise floor than the CR bound prediction.

The variance estimates presented in Table 3 are applicable to the component of velocity along the beam. The calculation of horizontal velocity depends on the number and geometry of the beams. For a four-beam (Janus) configuration, two horizontal velocity components are calculated from two pairs of back-to-back beams using

$$\hat{v}_h = \frac{\hat{v}_r^1 - \hat{v}_r^2}{2 \sin \theta},$$

where 1 and 2 superscripts denote two opposite facing beams and θ is the beam angle with respect to vertical. Assuming that the variance in radial velocity is Gaussian distributed and that the errors are uncorrelated between beams, the random error in horizontal velocity is the standard deviation from Table 3, multiplied by $1/\sin\theta\sqrt{2N}$, where N is the number of independent estimates (pings) averaged together. For a beam angle of 30° , the standard deviation in Table 3 is scaled by $\sqrt{2}/N$. This is the random variability in the estimate. A major source of bias in the estimate is discussed in the next section.

c. Estimate bias

Chereskin et al. (1989) presented numerical predictions of skew bias made from a model of the signal and the RDI processing filter, but direct confirmation using signal simulations and an ADCP was not possible due to an inability to generate signals of specified bandwidth. In this section, errors in ADCP measurements of simulated signals ($\alpha = 1$, $\omega_b^0 = \text{constant}$) are compared to numerical model predictions. The tracking feedback loop in the ADCP was disabled for these measurements. Since ideally the tracking loop keeps the filter centered on the signal, when the loop is disabled the filter is mispositioned by an amount equal to the error frequency. The advantages of disabling the loop are that the exact misposition (error frequency) is known a priori, large mispositions are possible, and no changes in loop gain occur during the measurements.

When the processing filter is not flat or the signal is not centered in the passband, the estimate of the first moment of the Doppler spectrum will be biased if the processing filter is not symmetrically positioned on the spectrum. The resulting bias, termed "skew bias," is a function of the filter misposition, the shape of the pro-




FIG. 7. Comparison of observed and predicted velocity standard deviation. Open symbols (\circ) were determined from simulations; the values are given in Table 3. The solid line is the least-squares fit to $1/\sqrt{T_p T_a}$. The CR lower bound derived by Theriault is shown by the dashed line, offset by zero and 3.1 cm s^{-1} . The CR bound derived by Smith (+) is calculated for the same cases as Table 3. The error predicted for the covariance estimator by Smith and Pinkel (*) is also shown for comparison.

cessing filter, and the width of the Doppler spectrum. Because of the discrete step size of the tracker and the loop time constant, some filter misposition is always present (e.g., Fig. 4). The Doppler spectrum has a finite bandwidth due to the modulation envelope, as discussed in the previous section, and to the width of the acoustic beams. The beamwidth effect is discussed in Chereskin et al. (1989) but was not included in the model used to create the synthesized signals.

A series of signal simulations were calculated, employing a Gaussian random variable scattering function convolved with the transmit pulse and constant Doppler shifts, that is, sinc functions. Two different low-pass-filter (LPF) bandwidths were used. Narrowband LPF will refer to 300-Hz bandwidth, and broadband LPF will refer to 600-Hz bandwidth. (Narrowband LPF is the default for the ADCP that was used.) A maximum of 4000 independent estimates were collected for nine different mispositions for an 8-m bin and pulse. Screening by energy eliminated some estimates; the minimum number of usable estimates (3756) occurred for the largest misposition, 85 cm s^{-1} measured in the narrowband LPF. The single-ping random error in Table 3 was estimated at 10.4 cm s^{-1} . Averaging over 3756 independent estimates reduces the random error to 0.17 cm s^{-1} . This error is small relative to the observed difference of -21.31 cm s^{-1} , which is attributed to skew bias. The measured values for the narrowband LPF width are shown in Fig. 8a (solid line, + symbols), together with two numerical model predictions. Corresponding results for the broadband LPF width are shown in Fig. 8b.

One numerical model prediction was computed from the model of Chereskin et al. (1989), modified

for radial-velocity error. (A misposition of 40 cm s^{-1} in radial velocity corresponds to 80 cm s^{-1} in horizontal velocity, and the error doubles as well.) This model predicted skew bias from the average misposition and an assumed signal spectrum. Since skew bias is a nonlinear function of misposition, the skew of the average misposition underestimates the average skew based on single-ping misposition. The second numerical model prediction is based on direct numerical simulations of the ADCP, using the same statistical scattering function that was used to drive the signal simulator and the same low-pass filter as used in both the ADCP and Chereskin et al. (1989). The numerical simulations (dash-dot line, \times symbols) and the Chereskin et al. (1989) results (dash line, \circ symbols) are shown in Fig. 8a for the narrowband filter and in Fig. 8b for the broadband filter. The curves are least-squares cubic polynomial fits to the discrete model estimates. The coefficients of the polynomials that were fit from the direct numerical simulations of the ADCP are presented in Table 4 and correspond to the dash-dot curves shown in Fig. 8.

Some insight into the skew bias can be gained by examining the distribution of velocity estimates for different mispositions. The distributions of the ADCP velocity estimates for the minimum and maximum misposition (0 and 84.29 cm s^{-1} , respectively) are shown in Fig. 9. In each case, 4000 pings were collected, but screening by the ADCP eliminated some estimates, primarily at large misposition. The distribution of estimates at zero misposition is approximately Gaussian (Fig. 9a), but becomes skewed toward lower estimates at large misposition (Fig. 9b). The mean and standard error of the skew were computed using a bootstrap (Efron and Gong 1983), where the 4000 independent estimates formed an empirical probability distribution of velocity for each misposition. One thousand estimates of 1000-ping sample means were formed and used to determine the mean and standard error of the mean. The standard errors ranged from 0.3 to 0.45 cm s^{-1} , with the largest standard error corresponding to the largest misposition.

Figure 10a shows the distribution of the estimates of the numerical simulations. A total of 3930 estimates was used. The model velocity distribution shows a skewed shape similar to that of the ADCP measurements. Velocity versus energy is shown in Fig. 10b. No energy screening was used for these velocity estimates. The largest variability in velocity occurs at low energy values, and screening based on energy level would reduce the variance of the estimate of the mean. The highest energy estimates, however, do not correspond to the misposition velocity but rather to the peak of the distribution.

Skew bias results in an underestimate of the error frequency and therefore the Doppler shift, with negative bias for positive misposition (error frequency) and positive bias for negative misposition. Skew bias de-

pends strongly on the filter bandwidth, and the errors more than halved when the filter bandwidth was doubled for both models and ADCP (Fig. 8b). The single-ping direct simulation estimates provide better agreement with measured skew than the average skew prediction of Chereskin et al. (1989), primarily because they include the nonlinear dependence on misposition. The ADCP and the single-ping direct simulation estimates agree to within about 1.0 cm s^{-1} for mispositions less than 60 cm s^{-1} . When tracking is enabled, mispositions are typically less than 30 cm s^{-1} .

Skew bias is also strongly dependent on the spectral width of the signal. For pulses less than about 30 m , the pulse bandwidth determines the signal bandwidth. Since shorter pulses have larger spectral width, they are associated with larger skew bias error in the velocity estimate. Figures 8c and 8d show numerical model predictions for a 4-m pulse for the narrowband and broadband LPF widths, respectively. Numerical predictions for a 16-m pulse are shown in Figs. 8e and 8f. The smallest errors occur for the longest pulse and the widest filter.

The skew estimates presented in Figs. 8–10 are for radial velocity. The magnitude of the horizontal velocity, as determined by a beam pair measuring the same velocity profile, is equal to the magnitude of the radial velocity measured by either beam, divided by $\sin\theta$. Assuming the magnitude of the skew bias velocity error is identical for each beam and $\theta = 30^\circ$, the skew error in horizontal velocity is twice the skew error of the radial velocity, and its sign causes the horizontal velocity estimate to be biased toward zero. The validity of the assumption of symmetry in skew bias in back-to-back beam pairs is examined in the next section.

5. Janus configuration errors in velocity shear

In the Janus configuration, four transducers equally spaced in azimuth (90°) are used to determine the three-dimensional velocity vector. Back-to-back beam pairs yield the velocity vector in the beam plane. The beam pairs yield the orthogonal components of horizontal velocity and redundant estimates of vertical velocity. The preceding interpretation assumes that the velocity field is the same over the volume spanned by all four beams. This assumption is valid for the simulations and model calculations. In the ocean, velocity variation within the volume spanned by the beams is present due to small-scale fluctuations such as turbulence and internal waves. In actual data, the horizontal velocity determined from the Janus-configured beams is interpreted as the average over the volume spanned by the four beams.

A series of signal simulations and model computations were conducted to estimate the errors observed by a beam pair in the Janus configuration due to measurement uncertainties. The simulations employed a Gaussian random-variable scattering function con-

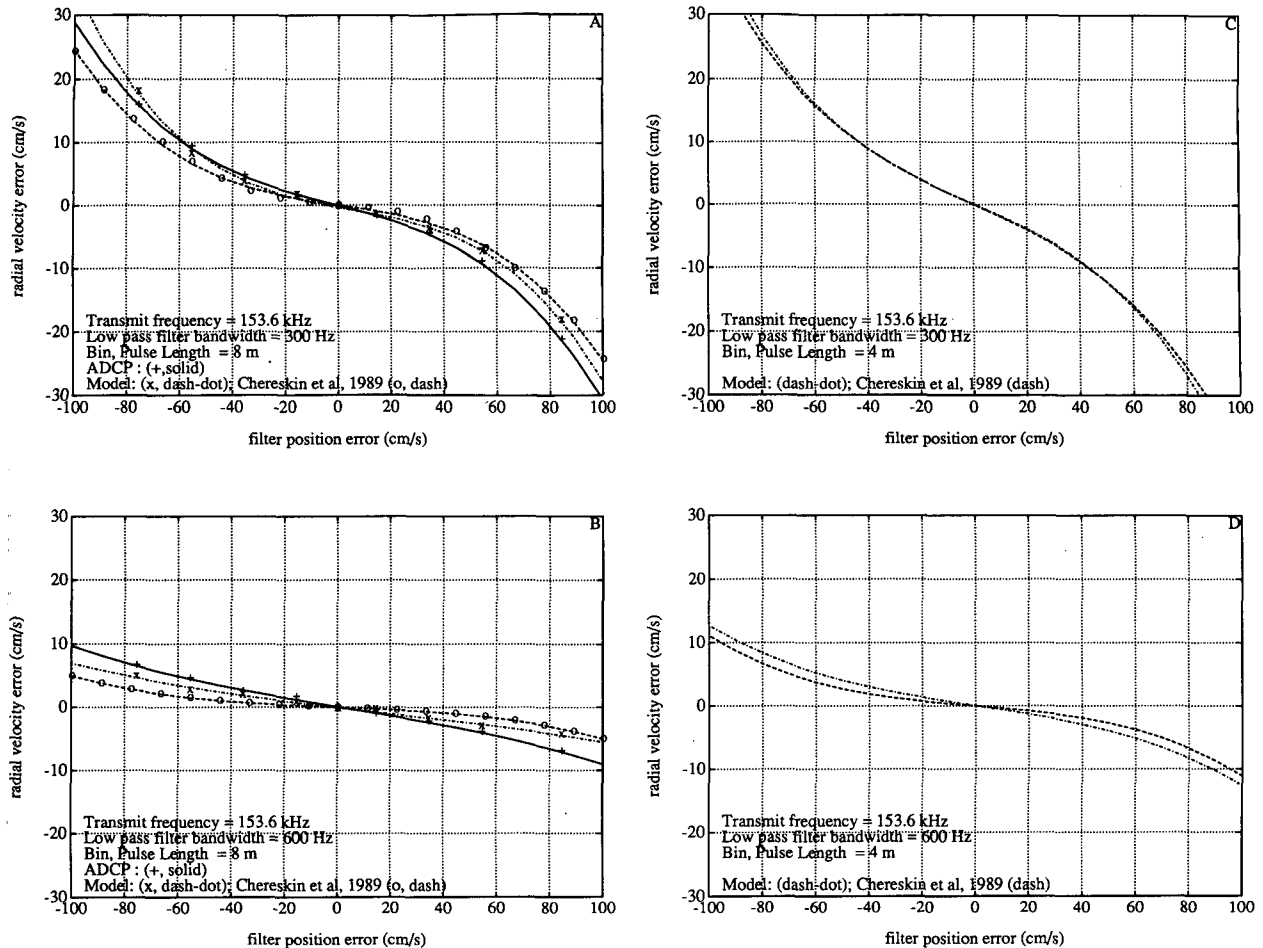


FIG. 8. Benchmark comparison of model and ADCP skew error. The (+) symbols are the bias in ADCP measurements of simulated signals; the solid curve is the least-squares fit of a cubic polynomial to the ADCP bias measurements. The (x) symbols and dash-dot curve are the model skew prediction based on simulated signals employing the same scattering function as used to drive the simulated signals coupled into the ADCP. The (o) symbols and dashed curve are the model results from Chereskin et al. (1989), which predicted the skew based on the average filter misposition rather than single-ping misposition. (a) Low-pass-filter bandwidth 300 Hz, 8-m bin and pulse; (b) low-pass-filter bandwidth 600 Hz, 8-m bin and pulse; (c) low-pass-filter bandwidth 300 Hz, 4-m bin and pulse, model only; (d) low-pass-filter bandwidth 600 Hz, 4-m bin and pulse, model only; (e) low-pass-filter bandwidth 300 Hz, 16-m bin and pulse, model only; and (f) low-pass-filter bandwidth 600 Hz, 16-m bin and pulse, model only.

involved with constant-shear velocity profiles. The velocity profiles are listed in Table 5, characterized by the velocity V_0 at $z = 0$ and the velocity shear V_z . The appellations “forward” and “backward” beams are used to differentiate between positive and negative Doppler shifts, since opposing beams see Doppler shifts of opposite signs. Nonzero values of V_0 can be interpreted as the ship speed for vessel-mounted ADCPs. A total of 300 pings were collected for each case. For a ping rate of once per second, these averages correspond to 5-min averages. ADCP and model output velocities were lagged by 1.1 m or 0.55 m^{-1} , depending on the LPF bandwidth, as discussed in section 4.

The forward-backward beam input and output velocity profiles, tracker positions, and errors for three of the shear cases in Table 5 are shown in Figs. 11–16.

ADCP error uncertainty [dotted lines in part (b) of Figs. 11–16] is the standard or random error in the ADCP mean, based on 300 pings. The tracker position plotted in part (a) of these figures is averaged over 300 pings, and then rounded to the nearest realizable tracker position; that is, it reflects the discrete step size of the actual tracker. Part (b) of these figures lists the average misposition of the tracker over 300 pings and 20 bins. This position is not rounded to a realizable tracker position, and so the ADCP and model mean tracker positions differ very slightly in Fig. 11b, even though they are seen to occupy the same positions in Fig. 11a.

Figure 11 presents the case where there is no velocity shear and the initial tracker position lies close to the true velocity. As a result, the trackers seldom move,




FIG. 8. (Continued)

but due to random velocity variation and the accumulated misposition, albeit small, they do occasionally step to the other side of the profile. The back beam for this case is shown in Fig. 12. For the back beam, the

TABLE 4. Coefficients of polynomial fit of radial velocity error y as a function of filter position error x for different pulse lengths and low-pass-filter bandwidths. System frequency is 153.6 kHz. Here $y = c_3x^3 + c_2x^2 + c_1x + c_0$, (x, y) in centimeters per second.

Pulse (m)	Bin (m)	c_3 ($\times 10^5$)	c_2 ($\times 10^5$)	c_1	c_0
Low-pass-filter bandwidth 300 Hz					
4	4	-2.3327	-1.4689	-0.1856	0.07264
8	8	-2.3253	36.0510	-0.0778	-0.4032
16	16	-1.3063	-0.8514	-0.0560	0.0561
Low-pass-filter bandwidth 600 Hz					
4	4	-0.6223	-0.2328	-0.0644	0.0638
8	8	-0.1354	8.0800	-0.0492	-0.0609
16	16	-0.1626	1.1968	-0.0321	-0.0820

input velocity lies about midway between adjacent tracker positions; the trackers spend about equal amounts of time on one side, although the model and ADCP trackers choose opposite sides. Hence, the tracker positions in Fig. 12 agree in character, if not in actual positions. The model and ADCP errors agree well, and the skew is not significantly different from zero, using the bootstrap upper bound of 0.45 cm s^{-1} as the skew uncertainty. There is random error in the profile that would be reduced by averaging over a larger ensemble.

The nonzero shear cases (Figs. 13–16) indicate a significant bias toward the tracker position (and zero velocity). As in the zero-shear case, there is obvious random variability that remains. The bias magnitude increases for increasing shear (Table 5). It remains essentially the same for cases C and D, which have identical shear but a different V_0 . For shears larger than 0.005 s^{-1} , the back-beam errors become significantly larger than the front beams. This is due to the larger

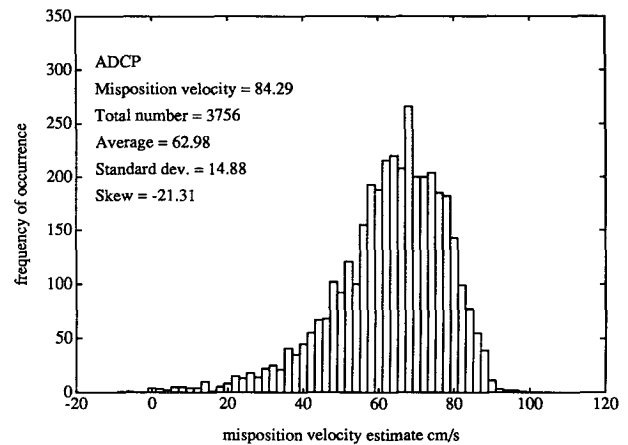
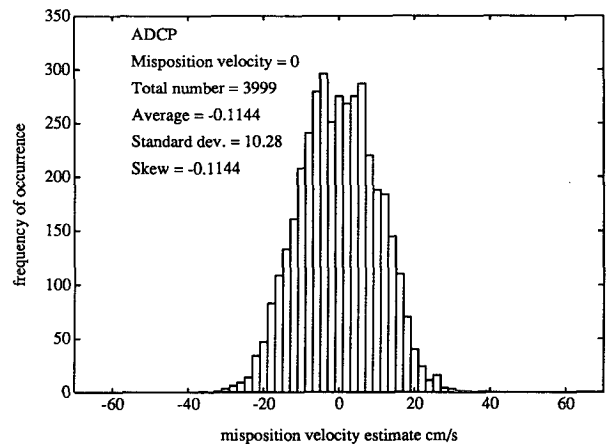


FIG. 9. Histograms of misposition velocity estimates made by the ADCP: (a) True misposition of 0 cm s^{-1} ; (b) true misposition of 84.29 cm s^{-1} .

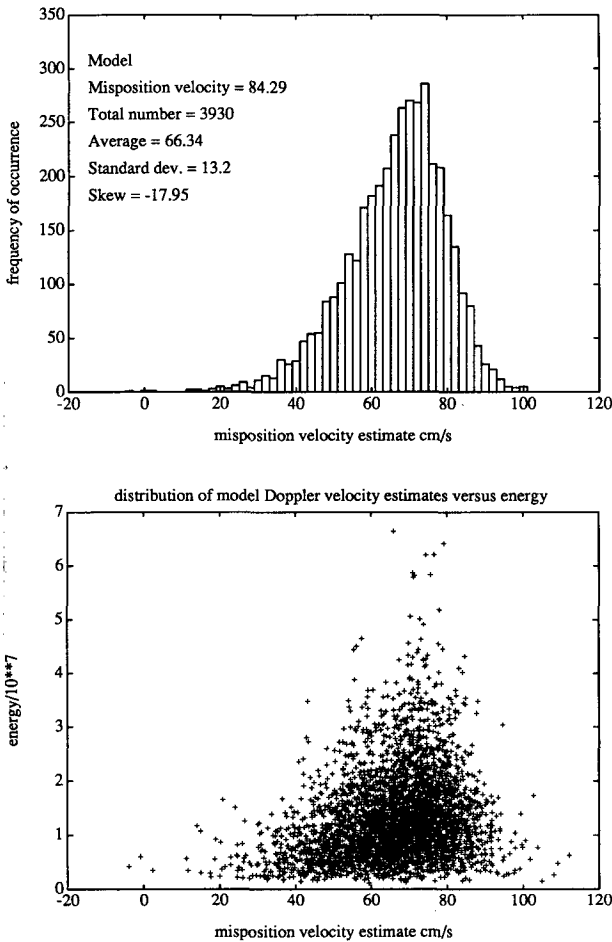


FIG. 10. (a) Histogram of misposition velocity estimates from the model where the true misposition was 84.29 cm s^{-1} . (b) Distribution of misposition velocity estimates versus energy (zero-lag autocorrelation) corresponding to (a).

mispositions for the back beam relative to the front beam (Fig. 17). The model overpredicts the misposition for this largest-shear case, but both model and ADCP indicate that on average the misposition is at least 10 cm s^{-1} greater on the back beam.

The initial velocity of the profile for cases A–C and F, 43 cm s^{-1} , corresponds to a gain switch point. This (deliberate) choice yields the largest discrepancy between front- and back-beam tracking time constants because the tracking time constant halves at the switch point (from 56 ms to 28 ms for the cases discussed here). In addition, the time constant is a function of the tracker step size, which varies with Doppler shift. For the largest shear, case F, the Doppler shift varies considerably over the depth range of the profile. As the Doppler shift increases linearly with depth, the time constant of the front beam increases linearly from 28 to 34 ms. The Doppler shift measured by the back beam decreases linearly with depth; correspondingly,

the time constant decreases from 53 to 48 ms. The average front-beam time constant is 31 ms, corresponding to a lag of 20 vertical meters; for a shear of 0.015 s^{-1} , a lag of 20 m yields an average misposition velocity of 30 cm s^{-1} . The corresponding calculation for the back beam is an average time constant of 51 ms, corresponding to a lag of 33 m or a misposition velocity of -49 cm s^{-1} . The profiles of average misposition velocity shown in Fig. 17 are in good agreement with these estimates over the depth range 60–160 m. The initial part of the profile corresponds to the transient phase in the tracking loop. Misposition velocity (or more correctly, error phase) must accumulate before the tracker changes the reference fre-

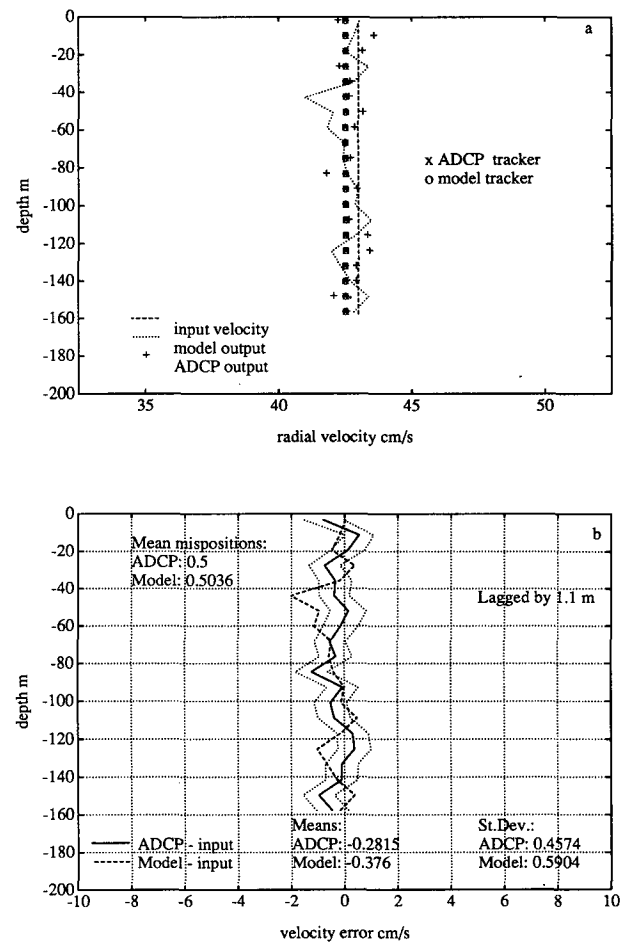


FIG. 11. (a) Forward-beam velocity comparison of model and ADCP, where single-ping velocity variance is a Gaussian-distributed random variable. Symbols as in Fig. 4. Output velocities lagged by 1.1 m. Constant velocity (43 cm s^{-1}), zero-shear profile, corresponding to Table 5, case A. (b) Velocity error corresponding to (a); symbols as in Fig. 6. Uncertainty in the ADCP error is shown as plus/minus the standard error of the 300-ping ADCP average (dotted lines). The mean and standard deviation of the error profiles and the mean tracker misposition over 20 depth bins are also shown. Units are centimeters per second.

TABLE 5. Velocity error estimates for constant shear (based on 300-ping averages: 153.6 kHz, 8-m bin and pulse, 20 bins).

Case	V_0 (cm s ⁻¹)	V_z (s ⁻¹)	Mean error		Standard deviation	
			ADCP (cm s ⁻¹)	Model (cm s ⁻¹)	ADCP (cm s ⁻¹)	Model (cm s ⁻¹)
Forward beam						
A	43	0.000	-0.28	-0.38	0.46	0.59
B	43	0.002	-0.54	-0.88	0.65	0.63
C	43	0.005	-1.30	-1.07	0.93	0.62
D	350	0.005	-1.31	-1.49	0.69	0.82
E	350	0.010	-2.50	-2.48	1.32	0.95
F	43	0.015	-3.38	-3.60	1.16	1.18
Backward beam						
A	-43	0.000	0.45	-0.39	0.78	0.76
B	-43	-0.002	0.29	0.23	0.46	0.88
C	-43	-0.005	1.35	1.06	0.55	0.93
D	-350	-0.005	1.92	1.11	1.04	1.21
E	-350	-0.010	3.54	2.30	1.10	1.77
F	-43	-0.015	6.15	5.16	2.57	2.75
Backward beam, double filter width						
F	-43	-0.015	1.83	1.65	1.24	1.25
Backward beam, decrease time constant successively by factors of 2						
F	-43	-0.015	2.09	1.97	1.16	1.01
F	-43	-0.015	1.79	0.67	0.96	0.69

quency; thus, misposition velocity increases over the first 60 m, then maintains a constant lag. The rather subtle effect that the front tracker time constant is increasing while the back-beam tracker time constant is decreasing is not readily apparent.

For smaller shear, the effect of the time constant discrepancy is not as dramatic, but the average misposition velocity is smaller for the front beam (Fig. 13b) than for the back (Fig. 14b). The transient phase of the tracking loop is also less noticeable. In the smaller-shear case (Figs. 13 and 14), the bias is approximately constant over the profile, and the error profile indicates random fluctuations about a nonzero mean. In contrast, the large-shear case shown in Figs. 15 and 16 indicates an increasing bias over the upper part of the depth range, then approximately constant over the remainder of the profile. The depth range over which the bias increases is larger for the back beam, and the maximum value that the bias attains is also larger.

The most extreme case (case F, back beam) was treated as a hypothetical worst-case error, and attempts to decrease the bias by increasing the bandwidth of the LPF and by decreasing the tracker time constant are shown in Table 5 and Figs. 18–20. A remarkable improvement results from increasing the LPF from 300 to 600 Hz. Although the mean misposition (Fig. 18)

remains nearly identical to Fig. 16, the bias is greatly reduced because the passband of the filter is wider. Changing the default time constant by factors of 2 was also examined. Figure 19 shows the result of decreasing the default time constant by a factor of 2. About four bins of error accumulate; then the tracker maintains a constant misposition by moving every bin, in contrast to Fig. 18, where about six bins of error accumulate before the tracker needs to move routinely. The mean misposition, and therefore the bias error, is reduced. Model and ADCP mispositions are nearly identical, whereas in Fig. 18 the model noticeably lagged the ADCP after it skipped one step at about 75 m. The improvement in bias is comparable to that achieved by using the wider filter. Figure 20 shows the result of decreasing the default tracker time constant by a factor of 4. Although misposition is reduced, the errors are not significantly different from Fig. 19.

Caution must be exercised if the time constant is decreased. Figure 21 shows the result of decreasing the time constant by a factor of 8, resulting in tracker instability. For clarity, the input velocity was chosen to be a constant (zero), and no modulation (velocity variance) was specified, corresponding to the narrow-band signals that were discussed in section 4a. Only a single ping from the ADCP is shown. The tracker diverges from the signal, with large oscillations, and there

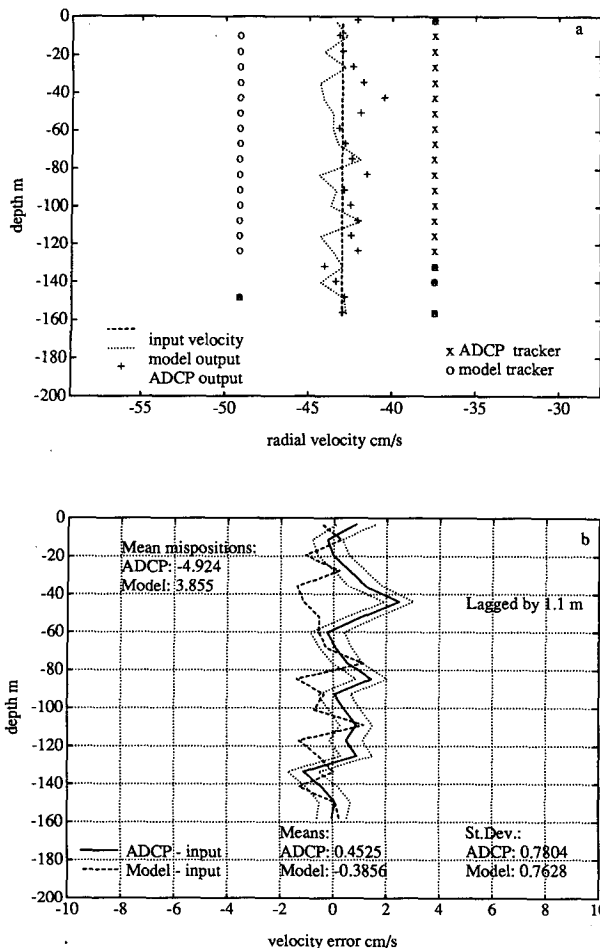


FIG. 12. (a) Comparison for the backward-beam velocity of a Janus pair; forward beam shown in Fig. 11. Constant velocity (-43 cm s^{-1}), zero-shear profile, corresponding to Table 5, case A. (b) Velocity error corresponding to (a).

are large errors in the ADCP velocity estimate when the misposition is large. When the misposition is so large that the low-pass energy level is below threshold, no velocity estimate is computed. This particular case, constant velocity and no modulation, could be measured with negligible error for a single ping if the time constant were longer. The main point of these last two comparisons is that reducing the misposition beyond 20 cm s^{-1} results in negligible improvement in the bias error, and in fact, the tracker will be unstable if the time constant is too short. Because the gain (time constant) varies by a factor of 2 over the range of frequencies synthesized, the instability could theoretically occur for some input velocity profiles when the time constant was reduced by a factor of 4.

The results in Table 5 can be used to estimate the bias in horizontal velocity where a linear shear is a good approximation. Horizontal velocity shear must be scaled by $\sin\theta$ to get the beam shear, and the bias as a function of shear for the forward and backward

beams taken from Table 5. The magnitude of the bias in horizontal velocity will equal the sum of the absolute value of the forward- and backward-beam biases, and the sense of the bias will be toward zero (underestimate). For small shears where the beam errors are symmetric, the error is approximately twice the single-beam error.

Table 6 compares the worst-case error (case F, back beam) as a function of bin and pulse lengths and filter bandwidth for the default time constant. For each filter width, the errors are smallest for the longest pulse combined with the shortest bin. Shorter bins can improve the time constant, and longer pulses decrease the spectral width. Doubling the filter bandwidth decreases the bias by more than a factor of 2.

6. Discussion

The results presented in this study assume a statistical scattering model. The assumptions of the statistical model appear reasonable at high signal to noise, but it

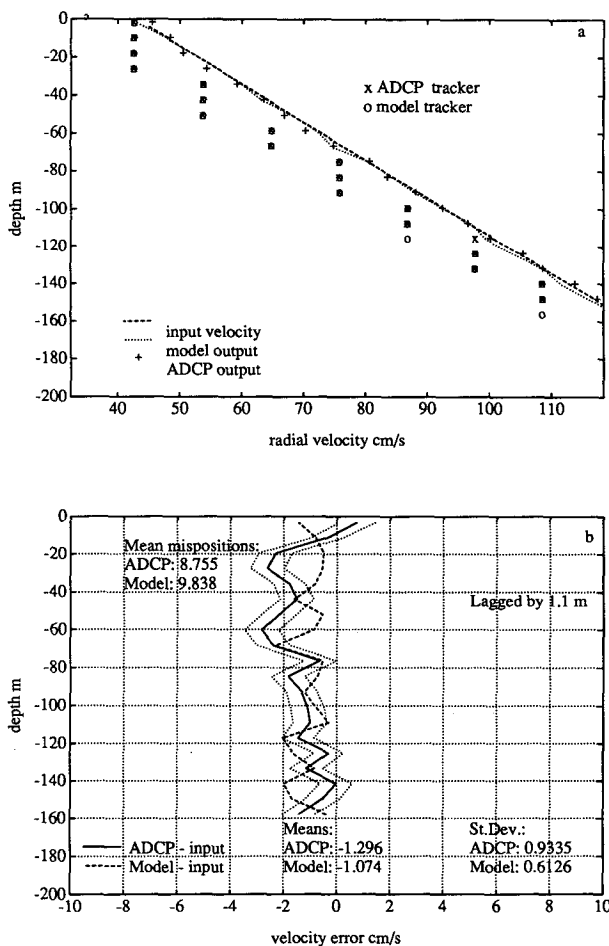


FIG. 13. (a) Forward-beam velocity comparison of model and ADCP. Velocity shear is 0.005 s^{-1} , corresponding to Table 5, case C. (b) Velocity error corresponding to (a).




FIG. 14. (a) Comparison for the backward-beam velocity of a Janus pair; forward beam shown in Fig. 13. Velocity shear is -0.005 s^{-1} , corresponding to Table 5, case C. (b) Velocity error corresponding to (a).

should be emphasized that these results are restricted to a high-SNR regime. Applied to ocean data, where the backscattered signal strength decays due to spreading and attenuation, approximately 90% of the velocity profile obtained by the ADCP used in this study is expected to lie within a high-SNR regime (above 10 dB).

The variance results (Table 3, Fig. 7) indicate that at short pulse and bin lengths the measurement variance is dominated by the bandwidth of the pulse-bin convolution. The variance can be decreased by increasing the pulse and bin up to lengths of 30 m, after which the variance drops off very slowly, essentially reaching a noise floor. The variance can also be decreased by increasing the length of the averaging ensemble (number of pings averaged together). For N pings and a matched bin and pulse width P , the standard error of the N -ping average is inversely proportional to $P\sqrt{N}$. Increasing P is thus a more efficient way to decrease the error; however, no practical gain

is made for increasing P beyond 30 m. Increasing P decreases the vertical resolution; increasing N decreases the temporal resolution. Since 30-m vertical resolution is typically too coarse for many oceanographic studies, the trade-off is to choose a smaller P and decrease the variance through loss of temporal resolution.

Skew bias can be decreased by choosing a longer pulse, a wider filter, or a faster tracker. Instability of the tracker will result if the time constant is too short; therefore, caution must be exercised if the time constant is changed from its default (stable) value. If scientific requirements set P , then only one option remains. Choice of a wider filter results in a loss of approximately 5%–10% of the maximum range (about 25 m for 150 kHz) through decreased SNR at the end of the profile. It is possible to decrease skew and minimize range loss by switching the filter bandwidth from wide to narrow in midprofile.

An encouraging result of this study is that in typical (low-shear) ocean conditions, velocity errors will typ-




FIG. 15. (a) Forward-beam velocity comparison of model and ADCP. Velocity shear is 0.015 s^{-1} , corresponding to Table 5, case F. (b) Velocity error corresponding to (a).

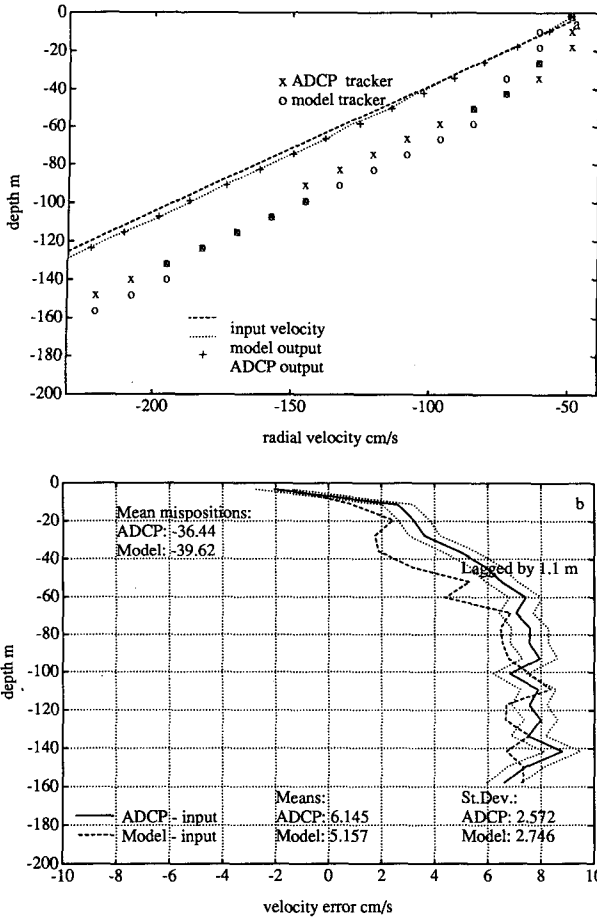


FIG. 16. (a) Comparison for the backward-beam velocity of a Janus pair; forward beam shown in Fig. 13. Velocity shear is -0.015 s^{-1} , corresponding to Table 5, case F. (b) Velocity error corresponding to (a).

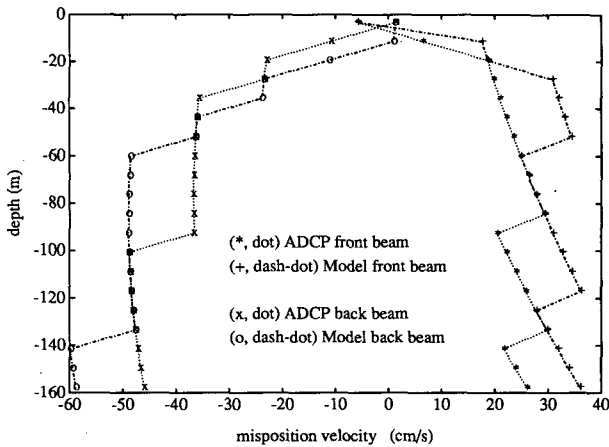


FIG. 17. Average misposition velocity (true input velocity minus tracker velocity) for case F, Figs. 15 and 16. The front-beam tracker misposition velocity is denoted by (*, dot) for the ADCP average and (+, dash-dot) for the model average. The back-beam tracker misposition velocity is denoted by (x, dot) for the ADCP average and (o, dash-dot) for the model average. Averages were calculated using 300 pings.

ically lie within $1\text{--}2 \text{ cm s}^{-1}$. For example, if the horizontal velocity shear is less than 0.004 s^{-1} , the corresponding beam shear is less than 0.002 s^{-1} . For a 15-min ensemble (900 pings), employing a 153.6-kHz four-beam Janus system with 30° beam angles and an 8-m bin and pulse, the random error is predicted to be approximately 0.5 cm s^{-1} . For a narrowband (300-Hz) LPF width, the skew is predicted to be about 1 cm s^{-1} .

In contrast, strongly sheared currents such as the Gulf Stream can have shears approaching the worst-case shear considered (case F, back beam). An important distinction, however, is that extremely large ocean shear seldom persists over a large depth range. Both the magnitude of the shear and its persistence (the tracker never catches up) contribute to the worst-case errors. Assuming the same operating parameters as discussed earlier, the random error is unchanged, but the skew error is about 10 cm s^{-1} . This error is a small percentage of the total velocity, but unfortunately, since it is a bias, it cannot be reduced by further averaging.

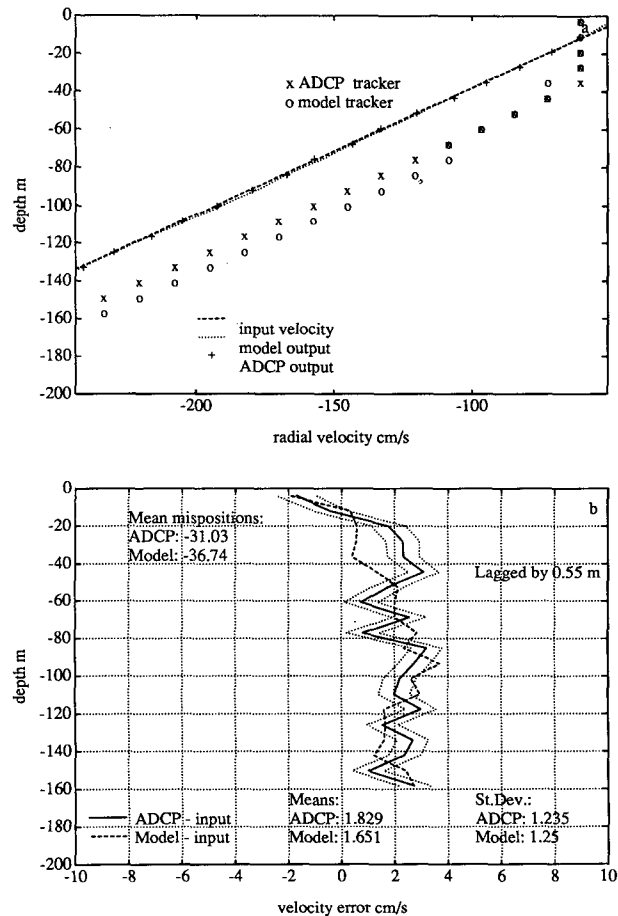


FIG. 18. (a) Backward-beam velocity of a Janus pair, with low-pass-filter bandwidth of 600 Hz, double that used in Fig. 16. Velocity shear is -0.015 s^{-1} , corresponding to Table 5, case F. (b) Velocity error corresponding to (a).




FIG. 19. (a) Backward-beam velocity of a Janus pair, with tracker time constant one-half that used in Fig. 16. Velocity shear is -0.015 s^{-1} , corresponding to Table 5, case F. (b) Velocity error corresponding to (a).

Doubling the filter bandwidth results in a total predicted skew of 4 cm s^{-1} . If 16-m vertical resolution suffices, however, the skew can be reduced to $1\text{--}2 \text{ cm s}^{-1}$. The random error for a 16-m pulse, 900 pings, is also less, about 0.25 cm s^{-1} .

Skew bias depends on the tracker misposition, which is large when the shear is large and constant, since the tracker allows some error to accumulate and then maintains the maximum lag. It may be possible to reduce the skew error in large shear by rotating the beams so that the beam shear is minimized. For example, in shipboard surveys of ocean currents, the sampling scheme that is frequently chosen is to cross the current at right angles. If the beams are aligned fore-aft and port-starboard, then one pair primarily measures the ship speed and the other pair measures the current structure and, more significantly, most of the shear. By rotating the beams 45° off-axis, all four beams measure the current shear, and the maximum shear measured by any single beam is reduced by $\sqrt{2}$. Obviously,

this is only sensible when the orientation of the current with respect to the beams is known a priori.

Velocity shear is the main source of tracker misposition discussed in this study. Another source of tracker misposition is from high-frequency ADCP platform motion, such as the effect of the surface wave field on a ship or buoy. Unless the ping-to-ping tracker samples fast enough to resolve the surface waves, large misposition can occur in the first bins, resulting in loss of the initial profiling range. Use of a wider filter can greatly reduce the error in these bins.

7. Conclusions

A study of bias and random error in ADCP velocity measurements was conducted using a system model and signal simulations. The model was extremely successful in predicting actual ADCP performance. One goal of the present study was to remove some of the "black box" mystique of commercial ADCPs. Although the tracking feedback loop in the model fol-

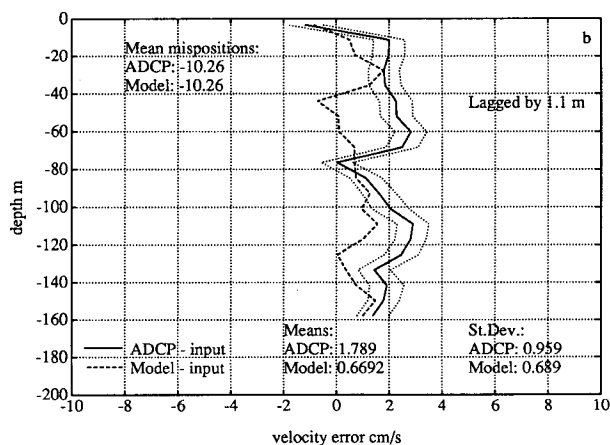
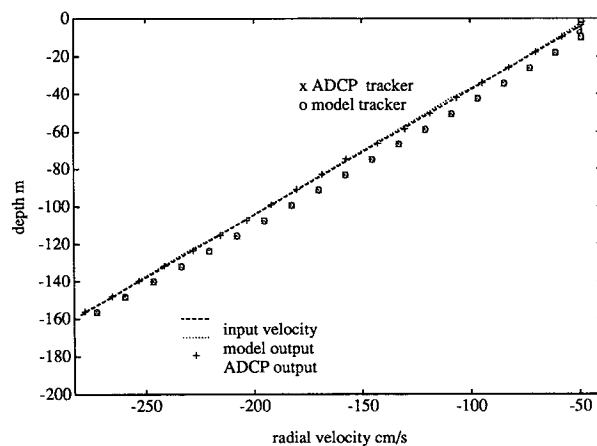


FIG. 20. (a) Backward-beam velocity of a Janus pair, with tracker time constant one-fourth that used in Fig. 16. Velocity shear is -0.015 s^{-1} , corresponding to Table 5, case F. (b) Velocity error corresponding to (a).

lowed general principles of phase-locked loops and feedback control, the loop time constant and processing filter were chosen to model the RDI ADCP, a commercial ADCP in widespread use. The signal simulator provided the capability of driving the same signal through an ADCP as specified in the model. Although the signal simulations were necessary to validate the model, the model alone is sufficient to predict system performance under different parameter settings and different velocity profiles. A parameterized tracking model will run on a personal computer.

The model benchmark studies, designed to isolate various aspects of system performance (positioning of tracker and skew bias), showed good agreement with simulation experiments. A third benchmark compared the estimate variance of simulations to theoretical predictions and determined the variance specified in the parameterized tracking model. The variance was found to have the functional form of the Cramer–Rao lower bound, with amplitude roughly double the predicted lower-bound value for pulses less than 30 m.

A Gaussian statistical model of reverberation was employed. At high SNR, the reverberation model appeared a reasonable choice, but appropriate modeling of the reverberation is a topic that deserves further investigation. The only modeled source of bias was filter skew error, since noise bias was negligible at high SNR. Skew error was observed to vary with the pulse length, filter bandwidth, and filter misposition. Smaller errors were observed for longer pulses, wider filter bandwidths, and smaller misposition. In general, the predicted skew errors are larger than those presented in Chereskin et al. (1989) and are in better agreement with skew errors measured by an ADCP driven by signal simulations. The larger skew errors result from the non-Gaussian distribution of the velocity estimates at large misposition. These skew error results supersede those of Chereskin et al. (1989).




FIG. 21. Example of tracker instability when default time constant was reduced by a factor of 8. Input velocity profile is solid line (narrowband, zero Doppler shift). A single ping measured by the ADCP is shown.

TABLE 6. Model back beam, case F, error as a function of pulse, bin, and filter width (based on 300-ping averages, 153.6 kHz, 20 bins).

Pulse (m)	Bin (m)	300-Hz LPF BW		600-Hz LPF BW	
		Mean error (cm s ⁻¹)	Standard deviation (cm s ⁻¹)	Mean error (cm s ⁻¹)	Standard deviation (cm s ⁻¹)
4	4	7.94	5.18	1.63	1.83
8	4	3.38	2.94	0.96	1.29
8	8	5.16	2.75	1.65	1.25
16	4	2.19	1.61	0.46	1.08
16	8	3.08	1.76	0.90	0.81

Errors observed in Janus beam-pair measurements of a sequence of increasing shear velocity profiles indicated that bias increased with increasing shear due to increased misposition of the frequency tracker. For the largest shears considered, an asymmetry was observed in the skew error of the forward (positive Doppler shift) and backward (negative Doppler shift) beams. Larger errors were observed for the back beam, due to larger mispositions due to the variation of the tracker time constant with Doppler shift. Overall, the model predictions and ADCP measurements showed remarkable agreement.

The discussion focused on the trade-offs between temporal and vertical resolution, range, and bias error through the choice of system parameters typically available to users of commercial instruments. These parameters include the pulse, bin, ensemble length, and filter bandwidth. Bias can be minimized by decreasing the vertical resolution and the range, that is, by using long pulses and large filter bandwidths. Random error can be minimized by increasing the averaging ensemble and the pulse and bin lengths, thereby decreasing the temporal and vertical resolution. For reasonable parameter choices in typical ocean conditions, the predicted error in horizontal velocity from effects considered in this study is 1–2 cm s⁻¹. In extreme ocean conditions, the predicted error using the same parameters as discussed earlier is much worse, about 10 cm s⁻¹. Optimal parameter choices, however, can reduce the worst-case error to 1–4 cm s⁻¹. While minimizing errors is always desirable, these trade-offs need to be determined in the context of the scientific requirements of the measurements, that is, the vertical and temporal resolution required for the processes being observed.

Acknowledgments. Numerous discussions with J. Gast and K. Deines of RD Instruments are gratefully acknowledged. S. Roy (InterOcean Systems) designed the signal simulator. This work was supported by the National Science Foundation under Contract OCE-8917566.

REFERENCES

- Appell, G. F., J. A. Gast, R. G. Williams, and P. D. Bass, 1988: Calibration of acoustic Doppler current profilers. *Proc. Oceans '88*, 346–352.

- Belliveau, D. J., and J. W. Whitman, 1989: Design of an acoustic Doppler current profiler test unit. *Proc. Oceans '89*, 1308–1313.
- Chereskin, T. K., D. Halpern, and L. A. Regier, 1987: Comparison of shipboard acoustic Doppler current profiler and moored current measurements in the equatorial Pacific. *J. Atmos. Oceanic Technol.*, **4**, 742–747.
- , E. Firing, and J. A. Gast, 1989: Identifying and screening filter skew and noise bias in ADCP current profiler measurements. *J. Atmos. Oceanic Technol.*, **6**, 1040–1054.
- Cutchin, D., Clifford, K., and J. Christensen, 1991: Towards a simpler acoustic Doppler current profiler. *Proc. Oceans '91*, 1258–1261.
- Efron, B., and G. Gong, 1983: A leisurely look at the bootstrap, the jackknife, and cross-validation. *Am. Stat.*, **37**, 36–48.
- Firing, E., 1987: Differences among nearby ADCP profiles as related to the speed of the ship. *Trans. Amer. Geophys. Union*, **68**, 1305.
- Hansen, D. S., 1985a: Receiver and analog homodyning effects on incoherent Doppler velocity estimates. *J. Atmos. Oceanic Technol.*, **2**, 644–655.
- , 1985b: Asymptotic performance of a pulse-to-pulse incoherent Doppler sonar in an oceanic environment. *IEEE J. Ocean Eng.*, **10**, 144–157.
- Kosro, P. M., 1985: Shipboard acoustic current profiling during the Coastal Ocean Dynamics Experiment. Ph.D. thesis, Scripps Institution of Oceanography, 119 pp.
- Miller, K. S., and M. M. Rochwarger, 1972: A covariance approach to spectral moment estimation. *IEEE Trans. Inf. Theory*, **IT-18**, 588–596.
- Pettigrew, N. R., and J. D. Irish, 1983: An evaluation of a bottom-mounted Doppler acoustic profiling current meter. *Proc. Oceans '83*, **83**, 182–186.
- Plueddemann, A. J., and R. Pinkel 1991: Biasing of the covariance-based spectral mean estimator in the presence of band-limited noise. *J. Atmos. Oceanic Technol.*, **8**, 172–178.
- RD Instruments, 1989: *ADCP Principles of Operation: A Practical Primer*. RD Instruments, 38 pp.
- Smith, J. A., 1989: Doppler sonar and surface waves: Range and resolution. *J. Atmos. Oceanic Technol.*, **6**, 680–696.
- , and Pinkel, R., 1991: Improvement of Doppler estimation through repeat-sequence coding. *Proc. Oceans '91*, 977–984.
- Theriault, K., 1986: Incoherent multibeam Doppler current profiler performance. Part I: Estimate variance. *IEEE J. Ocean Eng.*, **11**, 7–15.
- Urick, R., 1975: *Principles of Underwater Sound*. McGraw-Hill, 384 pp.
- Van Trees, H. L., 1971: *Detection, Estimation, and Modulation Theory. Part III: Radar–Sonar Processing and Gaussian Signals in Noise*. John Wiley and Sons, Inc., 626 pp.