Accuracy and precision in photonic Doppler velocimetry

D. H. Dolan[a)]
Sandia National Laboratories, Albuquerque, New Mexico 87185-1195, USA

(Received 8 March 2010; accepted 22 April 2010; published online 27 May 2010)

While photonic Doppler velocimetry (PDV) is becoming a common diagnostic in dynamic compression research, its limiting accuracy and precision are not well understood. Velocity resolution is known to be inversely proportional to the time resolution, but resolution estimates differ by one to two orders of magnitude. Furthermore, resolution varies with the number of recorded signals and how these signals are analyzed. Numerical simulations reveal factors that affect accuracy and precision in PDV, and the results may be extended to a broad class of measurements. After systematic effects are eliminated, the limiting velocity uncertainty in a PDV measurement is governed by the sampling rate, the signal noise fraction, and the analysis time duration. © *2010 American Institute of Physics.* [doi:10.1063/1.3429257]

I. INTRODUCTION

Photonic Doppler velocimetry (PDV), also known as heterodyne velocimetry,[1] is a powerful diagnostic for tracking velocity (m/s to km/s) on the short time scales (nanoseconds to microseconds) of a dynamic compression experiment. Essentially a fiber-based Michelson interferometer, PDV combines Doppler shifted light from a moving reflector with a reference source ($\lambda_0 \approx 1550$ nm), yielding a complete fringe when the reflector moves a distance $\lambda_0/2$. Unlike traditional displacement interferometry, which is impractical in most shock-wave experiments,[2] PDV can be applied wherever velocity interferometer sensor for any reflector (VISAR)[3] is used. Being simpler to build than other optical velocimeters,[3,4] yet robust to nonideal conditions (such as multiple Doppler shifts), PDV is rapidly finding uses throughout dynamic compression research.

Direct comparison tests indicate that PDV measurements are consistent with VISAR and shorting pins,[5] with the velocity resolution of PDV seemingly better by an order of magnitude. However, the limiting resolution of PDV is quite vague. Resolution estimates exist for PDV,[5] but common experience suggests that these estimates significantly exaggerate uncertainty. That there is a competition between time and velocity resolution in PDV analysis is well known, but the actual limits have not been investigated.

The goal of this work is to develop a rigorous uncertainty estimate for PDV. In doing so, both the accuracy and precision of various types of PDV measurement/analysis are considered. First, a brief overview of PDV analysis is given with extreme estimates of the limiting uncertainty. Next, Monte Carlo simulations are used to investigate the accuracy and precision of PDV in various settings. The results of these simulations are combined with theory to identify broader trends, and these trends are validated experimentally. Direct comparison of PDV and VISAR performance indicates that the former is competitive with (and possibly superior to) the latter in dynamic velocity measurements.

[a)]Electronic mail: dhdolan@sandia.gov.

II. BACKGROUND

As a displacement interferometer, individual points in a PDV measurement contain no velocity information. Converting the measured signal(s) to velocity requires an analysis time scale τ. This time scale is related to the minimum response time for a velocity step, analogous to the interferometer delay in VISAR (Appendix). It is assumed throughout this discussion that the detector(s) and digitizer(s) used in a PDV measurement are much faster than τ, and that all components are linear over the frequency range of interest.

When a reflector in a PDV measurement moves at fixed velocity, there is a simple relationship between the apparent[6] velocity v^* and the signal frequency $\bar{f}$,

$$v^* = \frac{\lambda_0}{2}\bar{f}. \tag{1}$$

The frequency content of PDV signals is typically determined by a short-time Fourier transform (STFT),[5] generating power spectra $P(f)$ from signal segments of duration τ; $\bar{f}$ is determined from the peak location in each spectrum. Velocity resolution is defined by how well frequency can be resolved during τ. The uncertainty principle[5] describes the relationship between τ and characteristic peak width $\Delta\bar{f}$,

$$(\Delta\bar{f})\tau \geq \frac{1}{4\pi}. \tag{2}$$

For example, PDV analysis over 1 ns has a minimum spectral width of 0.08 GHz, which corresponds to a velocity uncertainty of 62 m/s.

Uncertainty estimates from Eq. (2) may be overly conservative, particularly when the measurement contains a single Doppler shift. To demonstrate this, consider a fixed amplitude sinusoid with a 1.880 GHz frequency sampled 25 times over 1 ns. Figure 1 shows an example of the sampled signal and its power spectrum. Although the power spectrum is quite wide ($\sigma = 1$ GHz), the peak position can be much more narrowly defined. A Gaussian curve fit locates the peak location more precisely than the uncertainty principle: the

 © 2010 American Institute of Physics

Downloaded 18 Jan 2011 to 134.253.26.11. Redistribution subject to AIP license or copyright; see http://rsi.aip.org/about/rights_and_permissions

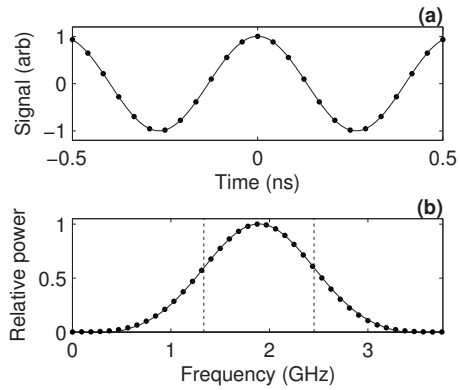

FIG. 1. (a) Hypothetical measurement of a 1.880 GHz sinusoid with 25 samples over 1 ns. The solid line shows the ideal signal, while dots indicate the sampled signal. (b) Calculated power spectrum (dots) using a Hamming window (zero padded to 256 points). The solid line shows a Gaussian fit to data ±0.5 GHz from the peak. Dashed lines indicate the frequency bounds determined from the uncertainty principle.

peak fit in Fig. 1(b) is precise to about 0.003 GHz. If such resolution enhancement could be obtained in an actual PDV measurement, it would correspond to a 2 m/s uncertainty with a 1 ns time window!

In practice, resolution enhancement via peak fitting is limited by systematic bias. In the present example, Gaussian fitting places the peak at 1.894 GHz instead of 1.880 GHz, a systematic difference nearly five times larger than the estimated resolution. Such discrepancies are quite common and vary with both the analysis method and signal properties. Monte Carlo simulations described in the next section illustrate how these factors limit frequency resolution in a PDV measurement.

III. MONTE CARLO SIMULATIONS

Consider signal s_k containing frequency f_0 sampled $N=2M+1$ times over a $\tau=2MT$ duration,

$$s_k = \cos(2\pi f_0 Tk + \delta) + \sigma R_k \quad (k=-M,\ldots,M). \tag{3}$$

As a basic model of real PDV measurements, the signal is contaminated with Gaussian noise that is a fraction σ of the total amplitude ($\langle R_k\rangle=0$ over many samples). The frequency $\bar{f}$ inferred from s_k depends on the duration, the number of samples, the signal phase (δ), and the noise fraction as well as the actual frequency and the analysis method. Given the many measurement variations and analysis parameter values, it is not obvious that an analytic expression can be derived for the limiting performance of PDV.

Numerical simulations can be used to investigate general resolution trends. Two key restrictions—fixed duration ($\tau=1$ ns) and fixed number of samples ($N=25$)—are considered here. The first restriction focuses on analysis with 1 ns time window, a characteristic goal for dynamic compression research. The second restriction is representative of digitizers used in PDV measurements (25 Gsamples/s, 6 GHz bandwidth). With these restrictions, the accuracy and precision can be investigated as a function of frequency, noise fraction, and analysis technique. Since the signal phase cannot be controlled during an actual measurement, a Monte Carlo approach is used to determine the distribution of frequency estimates for a particular combination of frequency and noise fraction. Both single-signal and multiple-signal measurements are considered.

The following procedure is used to simulate the resolution of single-signal PDV measurements.

(1) A random phase δ is selected from the range $(0,2\pi)$.
(2) The sampled signal s_k is calculated from Eq. (3) for the current value of f_0 and σ (drawing R_k from a normal distribution centered at zero with a standard deviation of unity).
(3) The power spectrum $P(f)$ is determined from a fast Fourier transform (FFT) (Ref. 7) of s_k using boxcar/rectangle, Hamming, and Hann window functions.[8]
(4) The peak location $\bar{f}$ is determined by fitting $P(f)$ with a Gaussian curve.
(5) Steps 1–4 are repeated 20 000 times to create a distribution of $\bar{f}$ values for each window function at the current value of f and σ.
(6) Accuracy is calculated from the difference between $\langle\bar{f}\rangle$ and f_0.
(7) Precision is calculated from the standard deviation of $\bar{f}$.

This procedure is applied over a range of frequencies (0.1–6 GHz, 0.02 GHz steps) and noise fractions (1%, 5%, 10%, 20%, and 50%) to create the curves shown in the next section.

Multiple-signal simulations follow the above procedure with a few modifications. Three versions of s_k are generated at each Monte Carlo iteration with linked phases: $\delta_n=\delta+2\pi n/3$ $(n=-1,0,+1)$. The three signals are analyzed with a quadrature reduction method[9] to determine displacement, which is differentiated with a local polynomial fit[10] to determine velocity; the results are scaled by $2/\lambda_0$ for comparison with single-signal simulations. A separate estimate of $\bar{f}$ is also made from the average of independent Gaussian fits to the three power spectra (boxcar window only).

IV. RESULTS AND DISCUSSION

A. Single-signal calculations

Figure 2 shows the calculated resolution of single-signal PDV measurements containing 10% noise. Each plot displays the absolute accuracy and precision estimates for a particular window function. In all three cases, PDV resolution is accuracy limited at low frequencies (about 300 MHz), while precision limitations dominate above 1 GHz (for $\tau=1$ ns). Both the accuracy and precision are quite poor in a low-frequency "shoulder" below 1 GHz, where frequency estimates are reliable to no better than 0.1 GHz. Above 2 GHz, PDV measurements would be accurate to about 10^{-4}–10^{-3} GHz (0.08–0.8 m/s) but only precise to 0.015–0.043 GHz (12–33 m/s). The Hamming window provides the best overall resolution (0.020–0.023 GHz) for 10% signal noise, while the Hann window has slightly worse resolution (0.024–0.026 GHz) and the boxcar window resolution is variable (0.015–0.043 GHz).

Qualitatively, the low-frequency shoulder in Fig. 2 arises from spectral analysis of a partial interferometer fringe.

Downloaded 18 Jan 2011 to 134.253.26.11. Redistribution subject to AIP license or copyright; see http://rsi.aip.org/about/rights_and_permissions

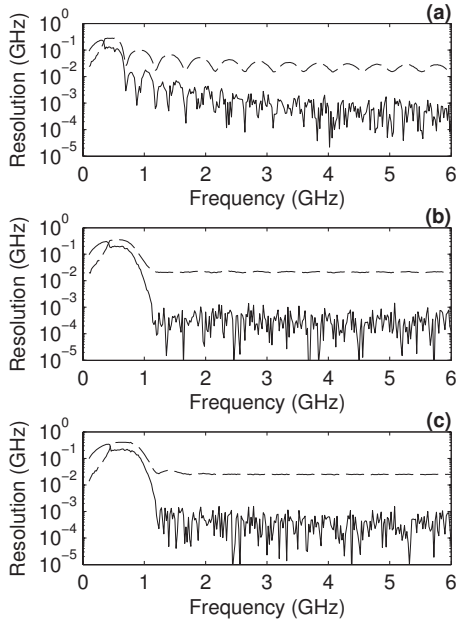

FIG. 2. Single-signal PDV resolution for measurements with 10% noise. Solid lines indicate the accuracy and dashed lines show the absolute resolution for Gaussian fits of the FFT power spectrum using boxcar (a), Hamming (b), and Hann (c) windows.

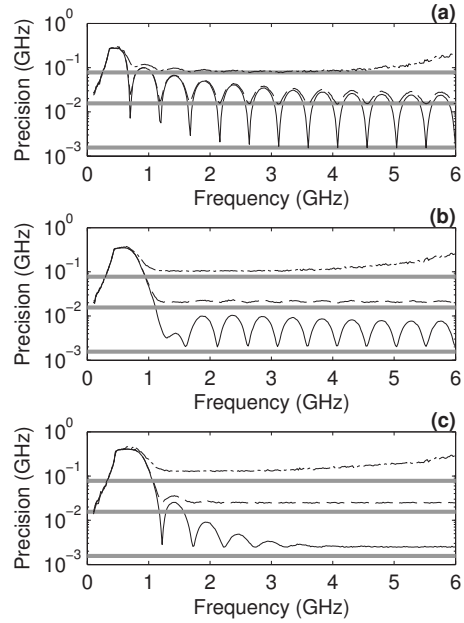

FIG. 3. Single-signal PDV precision using boxcar (a), Hamming (b), and Hann (c) windows. Solid lines correspond to measurements with 1% noise, dashed lines correspond to 10% noise, and dash-dot lines correspond to 50% noise. Gray horizontal lines indicate the estimated resolution limit [Eq. (9)].

More rigorously, this effect is due to practical limitations in discrete Fourier transforms (DFT) of real signals. The power spectrum of Eq. (3) is

$$|\tilde{S}(f)|^2 \propto S_+^2(f) + S_-^2(f) + 2\cos(2\delta)S_+(f)S_-(f), \tag{4}$$

$$S_\pm(f) = \int_{-\tau/2}^{+\tau/2} w(t)e^{-2\pi i(f\mp f_0)t}dt \to T\sum_{k=-M}^{M} w_k e^{-2\pi i(f\mp f_0)kT}, \tag{5}$$

where $w(t)$ is the window function (assumed to be real). The power spectrum contains positive and negative frequency components $[S_\pm(f)]$ that depend on the window function. The simplest example is the boxcar window

$$S_\pm(f) = T\frac{\sin[\pi(\tau+T)(f\mp f_0)]}{\sin[\pi T(f\mp f_0)]}, \tag{6}$$

which has a $1/|f|$ dependence in the limit $T\to 0$. At low frequencies, the negative component [second term in Eq. (4)] plays a significant role in the power spectrum, biasing the peak toward $f=0$. PDV analysis becomes accurate at high frequency because this bias diminishes rapidly $(1/|f|^2)$.

Numerical simulations indicate that PDV accuracy is quite similar for noise fractions of 1–20%, but the precision is strongly tied to noise fraction and frequency (Fig. 3). In all cases, PDV becomes less precise with increasing noise, spanning a range of 0.001–0.01 GHz (for $\tau=1$ ns). Precision can be highly variable in low noise measurements, particularly when using a boxcar window, though this variability diminishes with increasing noise. The Hann window provides the most consistent precision (especially at higher frequencies), but the Hamming windows yields better precision at higher noise levels. In all three cases, precision diminishes at high frequencies for 50% signal noise, as does the accuracy (not shown).

Signal phase and DFT limitations are responsible for variable precision above the low-frequency shoulder. Signal phase varies uncontrollably during STFT analysis of a PDV signal, causing the power spectrum peak to oscillate about the mean value. The oscillation arises from the third term in Eq. (4), which has a noticeable impact on peak location (even when the second term is negligible) because it decreases as $1/|f|$ instead of $1/|f|^2$. This term adds a negative frequency side lobe to the positive frequency contribution, altering peak location in a phase-dependent fashion; the effect disappears for certain phases, such as $\delta=45°$. Peak position is minimally altered when the side lobe is aligned with the $S_+(f)$ peak, a condition satisfied by

$$\tau\sin[2\pi f_0(\tau+2T)] = (\tau+2T)\sin[2\pi f_0\tau] \tag{7}$$

for a boxcar FFT. Under the present conditions, boxcar resonances occur at 1.18, 1.67, 2.15, 2.64, 3.12, 3.60, 4.08, 4.56, 5.05, and 5.53 GHz, a perfect match to the optimal resolution frequencies in Fig. 3(a).

By strongly attenuating points at the domain boundaries, the Hamming and Hann windows enforce FFT periodicity, making the power spectrum less sensitive to signal phase than the boxcar window. However, this benefit is accompanied by reduced noise rejection. In the low noise domain (1%), phase sensitivity is the dominant factor in PDV precision, and the Hann window is preferable. When noise (5–10%) begins to dominate the frequency precision, precision becomes less phase dependent and the Hamming window provides consistently low precision. Phase sensitivity is negligible at higher noise levels (50%), and the boxcar window

Downloaded 18 Jan 2011 to 134.253.26.11. Redistribution subject to AIP license or copyright; see http://rsi.aip.org/about/rights_and_permissions

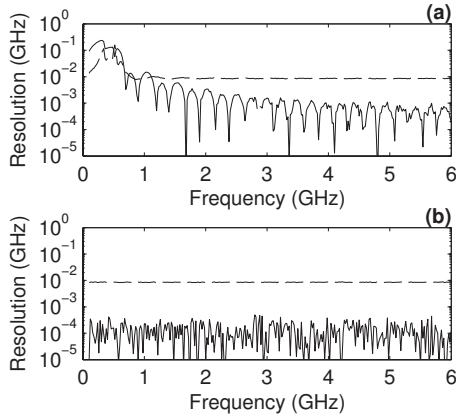

FIG. 4. Multiple-signal PDV resolution for measurements with 10% noise. Solid lines show the absolute accuracy and dashed lines shown the absolute precision for averaged boxcar FFT (a) and quadrature analysis (b).

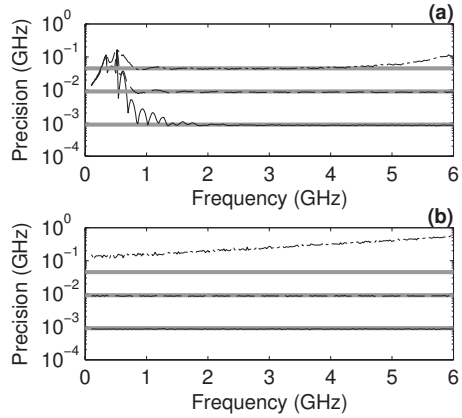

FIG. 5. Multiple-signal PDV precision for averaged boxcar FFT (a) and quadrature analysis (b). Solid lines correspond to signals with 1% noise, dashed lines correspond to 10% noise, and dash-dot lines correspond to 50% noise. Heavy gray lines indicate the estimated resolution limit [Eq. (9) scaled by $\sqrt{3}$].

provides better results due to its higher noise rejection. Thus, the optimal FFT window is determined by the noise fraction in a PDV measurement.

Apart from the low-frequency shoulder and phase variability, Fig. 3 indicates that the limiting precision of a PDV measurement scales with the signal noise fraction, regardless of the window function. To quantify this effect, consider the signal variation Δs at a single point due to a frequency variation Δf. The value of Δs averaged over all possible phase values is

$$\left(\frac{\Delta s}{\Delta f}\right)^2 = \frac{1}{2M+1}\sum_{k=-M}^{M}\left\langle\left(\frac{\partial s_k}{\partial f}\right)^2\right\rangle_\delta = (2\pi T)^2\frac{M(M+1)}{6} \approx \frac{(\pi\tau)^2}{6}, \tag{8}$$

where M is generally much larger than unity. Inverting Eq. (8) and equating Δs with the noise fraction yield an expression for the frequency uncertainty due to signal variation at one sample point; this result must be scaled by $\sqrt{N}$ to account for all N samples (collected at sampling rate f_S),

$$\Delta f = \sqrt{\frac{6}{N}}\frac{\sigma}{\pi\tau} = \left(\sqrt{\frac{6}{f_S}}\frac{\sigma}{\pi}\right)\tau^{-3/2}. \tag{9}$$

This expression (which can be derived[11] rigorously for sinusoidal fitting) links the limiting frequency resolution to fixed measurement quantities (sampling rate and noise fraction) and the adjustable analysis duration. The plots in Fig. 3 indicate that Eq. (9) is a reasonable estimate for the best possible frequency resolution.

B. Multiple-signal calculations

Figure 4 shows the calculated resolution for multiple-signal PDV measurements containing 10% noise. Accuracy in a three-signal boxcar FFT measurement is generally similar to single-signal measurements: at low frequencies, analysis is inaccurate and averaging does nothing to change this; beyond roughly 1 GHz, the averaged boxcar FFT method becomes precision limited. Unlike the FFT approach, quadrature analysis is always precision limited, and there is no apparent frequency dependence in either the accuracy or the precision (at this noise level).

Figure 5 shows the calculated precision for multiple-signal PDV measurements containing 1%, 10%, and 50% noise. Precision resonances are absent from the boxcar FFT approach (though the low-frequency shoulder remains) because of the 120° phase shifts between the three signals. This behavior can be demonstrated analytically by constructing an average power spectrum from the three signals, where the $S_+(f)S_-(f)$ term scales with

$$\sum_{n=-1}^{1}\cos\left[2\left(\delta+\frac{2\pi n}{3}\right)\right] = 0. \tag{10}$$

In quadrature analysis, precision is frequency independent at 1% and 10% noise levels because fringe shift is calculated directly. At 50% noise fraction, quadrature analysis shows worsening precision above 1 GHz, an effect that does not appear at lower noise levels ($\leq$20%). Like the single-signal calculations, accuracy is frequency dependent at 50% noise (not shown), but the average boxcar FFT method is much less affected than the quadrature approach.

The limiting precision in PDV measurement using L signals can be estimated by scaling Eq. (9) by $\sqrt{L}$. Precision estimates for 1%, 10%, and 50% noise levels shown in Fig. 5 indicate that averaged boxcar FFT and quadrature analysis can match the limiting resolution for 1–10% noise (similar performance is observed at 20% noise).

C. Implications

The single-signal and multiple-signal simulations reveal two critical trends in PDV.

- Measurements that rely on spectral analysis (particularly of a single signal) have poor accuracy and resolution at low frequencies.
- The limiting resolution is governed by the analysis time window, the sampling rate, and the signal noise fraction.

Downloaded 18 Jan 2011 to 134.253.26.11. Redistribution subject to AIP license or copyright; see http://rsi.aip.org/about/rights_and_permissions

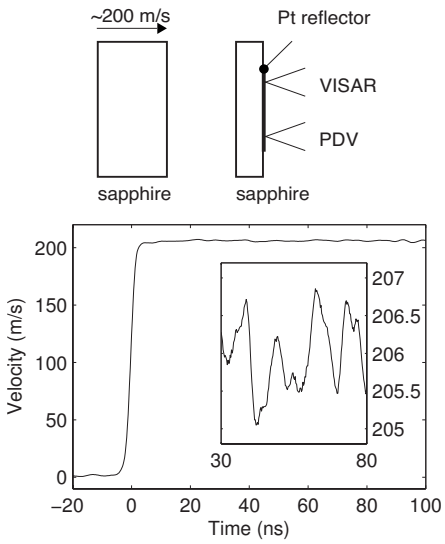

FIG. 6. (Top) Validation experiment layout, not to scale. (Bottom) Free surface velocity measurement using frequency-conversion PDV (τ=10 ns applied every 0.2 ns, Hamming window with 16 384 frequency points). The inset reveals measurement variations after shock arrival.

The first trend may be overcome through quadrature analysis and/or frequency conversion (Sec. V) to move away from the low-frequency shoulder. The second trend places a fundamental limit on velocity resolution. Under reasonable operating conditions (τ=1 ns, N=25, and σ=10%), the limiting velocity resolution of a single-signal PDV is about 12 m/s, roughly five times smaller than the bounds estimated from the uncertainty principle. Velocity resolution can be improved by expanding the analysis window (at the expense of time resolution) and/or faster data sampling at a fixed noise level. Precision can be also improved by using L *independent* measurements, but the improvement scales with $\sqrt{L}$ instead of L.

V. EXPERIMENTAL VALIDATION

Figure 6 shows an experimental test for the limiting performance of PDV. A c-axis sapphire impactor (63.5 mm diameter, 12.7 mm thick) was launched at a c-axis sapphire target (69.9 mm diameter, 6.4 mm thick) using a single-stage gas gun. Shorting pins indicate that the impact velocity was 206.6 ± 1.1 m/s, generating an elastic shock wave (4.6 GPa) in the target. After the shock wave reflected from the rear of the target, this surface moved for several microseconds at the original impact velocity (to a very good approximation[12]).

Simultaneous PDV and VISAR measurements of the free surface velocity were performed using a deposited platinum reflector (with a chromium adhesion undercoat); this configuration prevents reflector detachment due to tensile spall. Two NP Photonics "Rock" lasers (operating near 1550.000 nm) were used to create a frequency-conversion PDV: the target was illuminated with a tunable laser, and the Doppler shifted reflection was mixed with a fixed wavelength laser. The PDV signal frequency started at roughly 4 GHz and increased as the free surface moved to the right. Figure 6 shows a velocity history extracted from the PDV signal. A second measurement of the steady state velocity was obtained with a custom air-delay VISAR (National Security Technologies) set to its highest sensitivity (532 nm wavelength, 19 ns delay).

TABLE I. Free surface velocity measured in the validation experiment. VISAR uncertainty is calculated as 2% of the fringe constant (14 m/s), while PDV uncertainty is estimated from velocity variations in the steady state (Fig. 6).

Diagnostic	Time scale (ns)	Velocity change (m/s)
Pins	>1000	206.6 ± 1.1
VISAR	19	206.11 ± 0.28
PDV	19	205.95 ± 0.32
PDV	10	205.93 ± 0.64
PDV	1	206 ± 21

Table I shows the free surface velocity determined by shorting pins, VISAR, and PDV in the validation experiment. The shorting pin measurement is shown for comparison, but this value is not a strong constraint on the free surface velocity; the VISAR measurement (206.11 ± 0.28 m/s) provides a much more rigorous test for PDV performance. When PDV analysis is performed with a common time scale (19 ns), the result (205.95 ± 0.32 m/s) differs from the VISAR result by 0.16 m/s, which is less than the uncertainty in either diagnostic. VISAR has slightly less uncertainty than PDV in this example, though the difference can be attributed to the number of detectors in each measurement.

Unlike VISAR, different analysis time scales may be applied to PDV measurements after the experiment. Analysis at different time scales (Table I) indicates that τ determines the uncertainty, but not the mean value, of the PDV measurement. This behavior is consistent with the Monte Carlo simulations: PDV measurements are precision, not accuracy, limited above the low-frequency shoulder. Furthermore, the measured uncertainties vary with the $-3/2$ power of analysis time scale (Fig. 7) as expected from Eq. (9), though the variations are systematically larger than the limiting resolution. Much of this bias comes from the use of a Hamming window, which leads to a steeper rise time (relative to the boxcar window) at the expense of reduced frequency precision (Appendix). To compensate for this effect, Fig. 7 shows a scaled version of the data in Table I that are close to, but always above, the theoretical prediction.

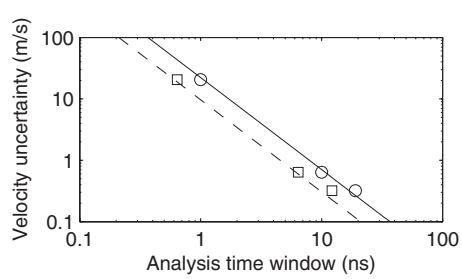

FIG. 7. PDV velocity uncertainty as a function of analysis time. Circles indicate the measured values (Table I) with a best fit line (slope set to $-3/2$); squares show the same data scaled by the relative rise time of the Hamming window (64%). The dashed line indicates the limiting uncertainty for a 25 Gsamples/s sampling rate and the ≈8% noise fraction of the frequency-conversion PDV.

Downloaded 18 Jan 2011 to 134.253.26.11. Redistribution subject to AIP license or copyright; see http://rsi.aip.org/about/rights_and_permissions

For a broader understanding of the relative performance of PDV and VISAR, consider measurements having a common τ (analysis duration and interferometer delay, respectively). The limiting velocity resolution in VISAR (Refs. 3 and 13) is approximately 1–2% of the fringe constant $K = \lambda_0/2\tau$. The velocity resolution ratio $\rho \equiv \Delta v_1/\Delta v_2$ (PDV/VISAR) is as follows:

$$\rho = \frac{\lambda_1}{\lambda_2}\sqrt{\frac{6}{f_S \tau}}\frac{\sigma}{\pi\epsilon}, \quad \epsilon \approx 1-2\%. \tag{11}$$

While PDV typically operates at a longer wavelength than VISAR ($\lambda_1 = 1550$ nm and $\lambda_2 = 532$ nm, respectively), the value of ρ is often close to unity, implying similar performance between diagnostics. A single PDV measurement with 25 samples/ns and 8% signal noise (comparable to the validation experiment) has a performance ratio of 3.6–1.8 when $\tau = 1$ ns, but several factors may lower this ratio in favor of PDV.

- The value of ρ decreases with $\sqrt{\tau}$ ($\rho = 1.1-0.57$ for $\tau = 10$ ns).
- PDV rise times are faster than VISAR using a common τ (Appendix). Incorporating this effect changes ρ to 2.0–1.0 for $\tau = 1$ ns.
- Multiple-signal PDV measurements decrease ρ as $\sqrt{L}$: averaging two independent signals brings the performance ratio to 1.4–0.7 ($\tau = 1$ ns, rise time corrected). For comparison, push-pull VISAR (Ref. 14) measurements use four signals in quadrature sensing, and eight signals are typically needed to resolve fringe ambiguity.
- Obtaining 1–2% fringe accuracy in VISAR requires precise system characterization,[15] without which the value of ϵ increases considerably (5–10% is not uncommon). Nonquadrature PDV analysis requires little or no system characterization.

Although displacement interferometry is less sensitive at 1550 nm than at 532 nm, the ability to make accurate frequency measurements allows PDV to be competitive with (and sometimes superior to) standard VISAR measurements. By the same token, there is no resolution advantage for a 1550 nm VISAR relative to a 1550 nm PDV.

VI. SUMMARY

While the uncertainty principle describes the general scaling between time and velocity resolution, PDV performance is usually much better than predicted by these criteria. Extremely accurate PDV measurements can be made with the appropriate configuration, leaving precision as the dominant uncertainty. Some analysis methods are advantageous in specific settings, but the limiting velocity resolution in PDV is controlled by the sampling rate, the noise fraction, and the analysis duration; multiple-signal measurements can further reduce uncertainty. With little effort, PDV time/velocity resolution can rival or exceed VISAR under typical operating conditions.

ACKNOWLEDGMENTS

David Holtkamp, Scott Jones, Tom Ao, Michael Furnish, Gerald Stevens, and Michael Furlanetto contributed to this paper through many useful discussions. The validation experiment was performed with the assistance of Sheri Payne, Mike Willis, Andy Shay, Jess Lynch, and Randy Hickman. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Co., for the United States Department of Energy's National Nuclear Security Administration under Contract No. DE-AC04-94AL85000.

APPENDIX: TIME RESOLUTION

All PDV measurements smear velocity changes over a time scale related to τ. The corresponding time resolution Δt depends on the velocity history and convention. For this discussion, Δt is defined as the 10–90% rise time in the analysis of a velocity step. The purpose of this Appendix is to investigate the relative values of Δt in PDV and VISAR, not necessarily the absolute magnitudes, for a consistent comparison.

Consider a velocity step with magnitude v_m at time $t = 0$. For convenience, let v_m match a complete VISAR fringe (266 m/s) for $\lambda_0 = 532$ nm and $\tau = 1$ ns. The velocity step may be converted to ideal (noise-free) PDV and VISAR signals and analyzed to reveal the limiting time resolution of each method. Following Sec. III, the PDV signal is sampled 25 times every nanosecond and the signal is frequency shifted by 3 GHz to avoid the low-frequency shoulder. Figure 8(a) shows an example of the time-frequency analysis ($\tau = 1$ ns) for a particular PDV signal (phase $\delta = 0$). The velocity profile (obtained by Gaussian fits at each time point) is clearly smeared during the analysis, though $\Delta t < \tau$.

Figure 8(b) shows phase-averaged step profiles for boxcar, Hamming, and Hann STFT analysis over 1 ns. The PDV time resolutions in this example are 0.58, 0.37, and 0.34 ns using boxcar, Hamming, and Hann windows (respectively);

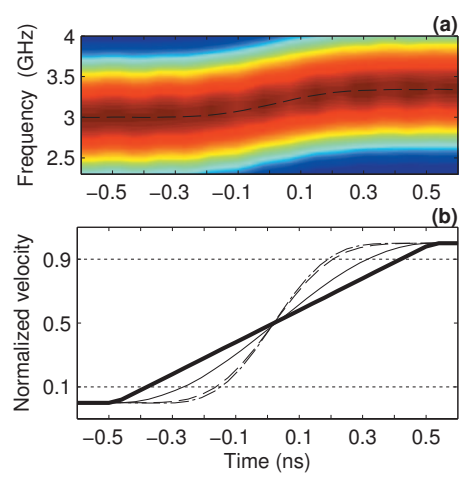

FIG. 8. (Color) (a) Time-frequency analysis (Hamming window, 1024 points) for a PDV measurement of a 266 m/s velocity step (frequency shifted to 3 GHz). The colors indicate relative intensities of 0.5 (blue) to 1.0 (red), and the dashed line shows the extracted profile. (b) Step responses for boxcar (solid line), Hamming (dashed line), and Hann (dash-dot line) FFT windows using a 1 ns analysis window. The heavy line indicates an ideal VISAR measurement with a 1 ns interferometer delay.

Downloaded 18 Jan 2011 to 134.253.26.11. Redistribution subject to AIP license or copyright; see http://rsi.aip.org/about/rights_and_permissions

quadrature analysis with a first-order Savitzky–Golay derivative has a similar step profile (not shown) as the boxcar FFT. The step response of a 532 nm VISAR measurement is also shown in Fig. 8(b); this profile is essentially linear, corresponding to a time resolution of 0.80 ns.

Since Eq. (11) compares PDV and VISAR measurements with the same τ, the performance ratio is biased toward VISAR, which has poorer time resolution. To achieve a common rise time $(0.80 \times \tau_2)$ with a VISAR having delay τ_2, PDV analysis should be performed over a $\tau_1 \approx 1.38 \times \tau_2$ time scale using a boxcar FFT window. Hamming and Hann windows yield higher time resolution that the boxcar window, though at the expense of frequency resolution (as shown in Fig. 3).

[1] O. T. Strand, D. R. Goosman, C. Martinez, and T. L. Whitworth, Rev. Sci. Instrum. **77**, 083108 (2006).
[2] L. M. Barker, Exp. Mech. **12**, 209 (1972).
[3] L. M. Barker and R. E. Hollenbach, J. Appl. Phys. **43**, 4669 (1972).
[4] C. F. McMillan, D. R. Goosman, N. L. Parker, L. L. Steinmetz, H. H. Chau, T. Huen, R. K. Whipkey, and S. J. Perry, Rev. Sci. Instrum. **59**, 1 (1988).
[5] B. J. Jensen, D. B. Holtkamp, P. A. Rigg, and D. H. Dolan, J. Appl. Phys. **101**, 013523 (2007).
[6] Apparent velocity includes the effects of optical windows and non-normal measurement. The distinction between apparent and actual velocity is thoroughly described in Ref. 15.
[7] E. O. Brigham, *The Fast Fourier Transform and Its Applications* (Prentice Hall, Upper Saddle River, NJ, 1988).
[8] The choice of window functions used here is representative of many possibilities (Ref. 16). The boxcar and Hann windows are extreme choices, the former using all sample points equally while the latter fully attenuates samples at the boundaries; the Hamming is an intermediate choice with partial attenuation at the boundaries. Other windows can be used in PDV analysis, but the general trends discussed here do not change.
[9] D. H. Dolan and S. C. Jones, Rev. Sci. Instrum. **78**, 076102 (2007).
[10] A. Savitzky and M. J. E. Golay, Anal. Chem. **36**, 1627 (1964).
[11] M. H. Montgomery and D. O'Donoghue, Delta Scuti Star Newsletter, NASA Astrophysics Data System **13**, 28 (1999).
[12] J. Walsh and R. Christian, Phys. Rev. **97**, 1544 (1955).
[13] L. M. Barker, AIP Conf. Proc. **429**, 833 (1998).
[14] W. F. Hemsing, Rev. Sci. Instrum. **50**, 73 (1979).
[15] D. H. Dolan, "Foundations of VISAR analysis," Technical Report No. SAND2006–1950, 2006.
[16] F. J. Harris, Proc. IEEE **66**, 51 (1978).

Downloaded 18 Jan 2011 to 134.253.26.11. Redistribution subject to AIP license or copyright; see http://rsi.aip.org/about/rights_and_permissions