

Power spectrum and Allan variance methods for calibrating single-molecule video-tracking instruments

Bob M. Lansdorp and Omar A. Saleh^{*}

Citation: *Rev. Sci. Instrum.* **83**, 025115 (2012); doi: 10.1063/1.3687431

View online: <http://dx.doi.org/10.1063/1.3687431>

View Table of Contents: <http://aip.scitation.org/toc/rsi/83/2>

Published by the [American Institute of Physics](#)

Power spectrum and Allan variance methods for calibrating single-molecule video-tracking instruments

Bob M. Lansdorp¹ and Omar A. Saleh^{2,a)}

¹Materials Department, University of California Santa Barbara, Santa Barbara, California 93106, USA

²Materials Department and Biomolecular Science and Engineering Program, University of California Santa Barbara, Santa Barbara, California 93106, USA

(Received 14 October 2011; accepted 2 February 2012; published online 29 February 2012)

Single-molecule manipulation instruments, such as optical traps and magnetic tweezers, frequently use video tracking to measure the position of a force-generating probe. The instruments are calibrated by comparing the measured probe motion to a model of Brownian motion in a harmonic potential well; the results of calibration are estimates of the probe drag, α , and spring constant, κ . Here, we present both time- and frequency-domain methods to accurately and precisely extract α and κ from the probe trajectory. In the frequency domain, we discuss methods to estimate the power spectral density (PSD) from data (including windowing and blocking), and we derive an analytical formula for the PSD which accounts both for aliasing and the filtering intrinsic to video tracking. In the time domain, we focus on the Allan variance (AV): we present a theoretical equation for the AV relevant to typical single-molecule setups and discuss the optimal manner for computing the AV from experimental data using octave-sampled overlapping bins. We show that, when using maximum-likelihood methods to fit to the data, both the PSD and AV approaches can extract α and κ in an unbiased and low-error manner, though the AV approach is simpler and more robust. © 2012 American Institute of Physics. [<http://dx.doi.org/10.1063/1.3687431>]

I. INTRODUCTION

Single-molecule manipulation (SMM) instruments, such as the magnetic tweezer (MT), optical trap (OT), or atomic force microscope (AFM), measure properties of single molecules by tethering the molecule to a probe (i.e., a magnetic bead in the MT, a dielectric sphere in the OT, or a cantilever in the AFM), applying force to the probe, and evaluating the molecule's response by measuring the probe's position (Fig. 1). Performing a SMM experiment with accuracy and precision requires calibration of the applied force.

Force calibration is frequently achieved by comparing the noisy probe trajectory to an analytical model derived from the Langevin equation. When probe mass is negligible, the Langevin equation depends on two parameters: the hydrodynamic drag on the probe, α , and the spring constant of the system, κ . Typically, estimates of α and κ are made by comparing $P(f)$, the theoretical power spectral density (PSD) of probe fluctuations found from the Langevin equation, with $\tilde{P}(f)$, the discrete PSD calculated from the measured probe position.

In practice, accurately and precisely comparing $P(f)$ and $\tilde{P}(f)$ is difficult for three reasons:

1. *Instrument response.* The instrument used to measure probe position invariably has a non-uniform frequency response, distorting $\tilde{P}(f)$ away from $P(f)$.
2. *Finite sampling distortions.* $\tilde{P}(f)$ is distorted by the effect of aliasing, in which the estimated power at a given frequency f_k contains contributions from all frequencies $f_k + nf_s$, where $n = \dots, -1, 0, 1, \dots$, and f_s is the sam-

pling frequency. Further, $\tilde{P}(f)$ is also distorted by the effects of the finite duration of the measured probe trajectory, which result in spectral leakage in the resulting estimate of the PSD.

3. *Biased fitting.* The model/data comparison method itself can bias the estimated parameters, as the typical strategy of least-squares fitting is correct only when the error in the data is normally distributed, which is not the case for the estimated PSD.¹

In this article, we present and compare two approaches to low-error, bias-free α and κ estimation that remove or sidestep each of the three difficulties. Both strategies are formulated for a particular instrumental response function corresponding to zero-dead-time video-based tracking of probe position; that is, tracking by imaging the probe at a frame rate f_s with the shutter open for a time $\tau = 1/f_s$. Our approaches remove all three sources of bias, and do so using either time-domain or frequency-domain analysis procedures; this contrasts with prior work that focused on only some of the sources of bias, and used only frequency-domain methods.¹⁻⁴

In the frequency domain, we derive a closed-form expression for the measured probe PSD, $P_{A,B}(f)$, that accounts for both aliasing and the video-tracking instrument response. While approximations to this expression have been made before,⁵ our expression is exact and, to the best of our knowledge, novel. We show that finite-duration (spectral leakage) effects can be removed by calculating the windowed, blocked PSD, $\tilde{P}_{b,w}(f)$, using established discrete Fourier transform (DFT) algorithms.^{4,6} Finally, following Nørrelykke and Flyvbjerg,¹ we show that bias-free fitting of $P_{A,B}$ to $\tilde{P}_{b,w}(f)$ can be achieved using a maximum-likelihood estimate (MLE) algorithm.

^{a)}Electronic mail: saleh@engineering.ucsb.edu.




FIG. 1. (Color online) A schematic representation of a video-tracking SMM instrument. A single molecule of DNA is attached to a magnetic bead on one end and a substrate on the other, and magnets are used to apply a vertical force to the bead. Light from a collimated LED strikes the probe particle and the resulting diffraction pattern can be used to track the bead in three dimensions at 60 Hz using a CCD camera. Although the magnetic tweezer is a commonly used video-tracking instrument, the methods we present in this paper are generally applicable to any video-tracking instrument.

Our second strategy is based on a time-domain analysis, the Allan variance (AV). The AV is a two-sample difference measurement⁷ that is well-known in the frequency-stability literature^{8,9} and has recently been introduced to the SMM community.^{10–13} Since the AV is a time-domain analysis, neither aliasing nor spectral leakage occur. Further, the calculation of Allan variance intrinsically assumes a finite averaging time that exactly matches the video-tracking instrument function, so no additional correction for the instrument needs to be applied. Thus, AV is well-suited to the problem of determining α and κ , since no corrections must be made for the instrument response, aliasing or spectral leakage. However, using the AV is not completely free of subtleties: we show that obtaining low-error, bias-free parameter estimates requires that the AV is estimated from data using bins that fully overlap, and that range in length by factors of 2. Finally, we point out that biased fitting is also an issue with the AV due to the non-normally distributed noise in AV estimates, and estimating the noise distribution is more involved for the AV than the PSD. To overcome this, we extend the Nørrelykke–Flyvbjerg strategies¹ to the Allan variance, and show that unbiased parameter estimates can be found from an MLE approach that accounts for the actual distribution of Allan variance estimates.

Upon comparing the two strategies, we find that the AV and the PSD, when applied in an optimal manner, can both estimate α and κ with zero-bias and with nearly-identical error. However, use of the AV is conceptually simpler, since fewer corrections need to be applied, and more practical, since the AV compresses the data more efficiently than the PSD. Furthermore, AV does not require the choice of a data compression (blocking) parameter; thus, the AV calculation has less indeterminacy.

II. POWER SPECTRAL DENSITY

A. Langevin dynamics and the physical PSD

For micron-scale probe beads trapped by optical or magnetic forces in water, inertial effects can be ignored up to roughly megahertz frequencies,¹⁴ and thus are insignificant at typical video-tracking frequencies of approximately 100 Hz. Thus, the bead position $x(t)$ vs. time t obeys the overdamped Langevin equation of motion:

$$\kappa x + \alpha \dot{x} = F_L, \quad (1)$$

where κ is the spring constant of the system and $\alpha = 6\pi\eta r$ is the dissipation of the spherical bead of radius r in a solution of viscosity η , and F_L is the Langevin force, a time-varying stochastic force that, in thermal equilibrium, obeys the fluctuation-dissipation relation $\langle F_L(t + t_0)F_L(t) \rangle = 2\alpha k_B T \delta(t_0)$ given thermal energy $k_B T$. The power spectral density of bead motion, $P(f)$, where f is the frequency in hertz, can be calculated by taking the magnitude of the Fourier transform of Eq. (1) to obtain

$$P(f) = \frac{k_B T}{2\pi^2 \alpha (f_c^2 + f^2)}, \quad (2)$$

where $f_c = \kappa/2\pi\alpha$. We use two-sided power spectra throughout so that integrating $P(f)$ on the range $(-\infty, \infty)$ results in the equipartition result $\langle x^2 \rangle = k_B T/\kappa$.

B. An analytical expression for the PSD function, accounting for aliasing and the instrument function

The PSD of Eq. (2) corresponds to the true physical motion of the bead; any actual measurement $\tilde{P}(f)$ will deviate from $P(f)$ due to the distorting effects of aliasing and the instrumentation. In video-tracking experiments in which the bead position is measured by a camera with a shutter time τ_s , any motion of the bead that occurs on timescales faster than τ_s is diminished by the measurement;¹⁵ i.e. video tracking introduces a low-pass filter. The low-pass filtering effect of the camera can be quantitatively accounted for: when the illumination is constant over τ_s , the measured position x_i from frame i at time t_i corresponds to the average over the interval $(t_i - \tau_s/2, t_i + \tau_s/2)$:

$$\begin{aligned} x_i &= \frac{1}{\tau_s} \int_{t_i - \tau_s/2}^{t_i + \tau_s/2} x(t) dt \\ &= [x * \Pi_{\tau_s}](t_i), \end{aligned} \quad (3)$$

which we have written in the second line as the convolution of $x(t)$ with the boxcar function $\Pi_{\tau_s}(t)$. We define $\Pi_{\tau_s}(t')$ to have value $1/\tau_s$ for $|t'| < \tau_s/2$, and to be zero otherwise. The frequency-space equivalent of convolution is simply multiplication of the Fourier transforms of the two convolved functions. Thus, since the Fourier transform of $\Pi_{\tau_s}(t)$ is $\sin(\pi f \tau_s)/(\pi f \tau_s)$, the PSD of bead motion including the boxcar response, $P_B(f)$, is

$$P_B(f) = P(f)H(f), \quad (4)$$

where the instrument response $H(f)$ is

$$H(f) = \frac{\sin^2(\pi f \tau_s)}{(\pi f \tau_s)^2}. \quad (5)$$

$P_B(f)$ only accounts for the effect of instrumental low-pass filtering in distorting the measured PSD; it does not account for the further distortions created by aliasing. The aliasing effect causes the PSD at a particular positive frequency f , with $0 < f < f_s/2$, to contain the summed power at all frequencies in the series $|f + n f_s|$, where $n \in \{\dots, -1, 0, 1, \dots\}$.⁵ Thus, the measured PSD accounting for both aliasing and boxcar filtering is given by

$$P_{A,B}(f) = \sum_{n=-\infty}^{\infty} P_B(|f + n f_s|) \quad (6)$$

with P_B given by Eqs. (4) and (5).

In the special case that the sampling frequency is the inverse of the shutter time, $f_s = 1/\tau_s$, the sum in Eq. (6) can be performed analytically:

$$P_{A,B}(f) = \frac{2k_B T \alpha}{\kappa^3} \left(\kappa + \frac{2\alpha f_s \sin^2\left(\frac{\pi f}{f_s}\right) \sinh\left(\frac{\kappa}{\alpha f_s}\right)}{\cos\left(\frac{2\pi f}{f_s}\right) - \cosh\left(\frac{\kappa}{\alpha f_s}\right)} \right). \quad (7)$$

Equation (7) can be verified in multiple ways. First, we can use $P_{A,B}$ and the Wiener–Khinchin relation to find the variance of the measured signal in the time domain; the result matches previously predicted results for boxcar-averaged probe positions.⁵ Second, we can compare $P_{A,B}$ with $P_A(f)$, the PSD expression that accounts only for aliasing (and not instrumental low-pass filtering)¹⁶:

$$P_A(f) = \frac{k_B T}{\kappa f_s} \left(\frac{\sinh\left(\frac{\kappa}{\alpha f_s}\right)}{\cosh\left(\frac{\kappa}{\alpha f_s}\right) - \cos\left(\frac{2\pi f}{f_s}\right)} \right). \quad (8)$$

P_A corresponds to the PSD measured by an instrument with sampling frequency f_s , but with an infinitely fast measurement interval (i.e., the boxcar measurement function, Π_τ , is replaced by a delta function). We compare $P_A(f)$ and $P_{A,B}(f)$ to experimental data in Fig. 2. The data are clearly best-described by $P_{A,B}$, lending confidence to the expression in Eq. (7).

We further verified the prediction of Eq. (7) through simulation of probe trajectories in the time domain, computation of the PSDs of the simulated trajectories, and comparison to $P_{A,B}$. We simulated discrete particle trajectories using an algorithm based on the solution to the Smoluchowski equation for a Brownian particle in a harmonic well.¹⁷ To account for the averaging effect of the boxcar filter, we simulated discrete points with a short time step of $\delta < \tau$, then averaged together consecutive clusters of τ/δ points, thus arriving at a set of discrete positions x_i spaced by time τ . We did this for multiple time steps $\delta = \tau/m$, with $m = 1, 2, 4, 8, 16$. Figure 3 shows comparisons of the simulation with both $P_{A,B}$ and P_A for both $f_c \ll f_s$ and $f_c \approx f_s$. For $m = 1$, there is no averaging of the time-domain signal, so the simulated PSDs match P_A , as expected. As m increases, the simulated PSDs converge on $P_{A,B}$.




FIG. 2. (Color online) Experimental PSD (gray) compared to MLE fits to $P_{A,B}$ (Eq. (7), blue) and P_A (Eq. (8), purple). The data trajectory contained 4096 points, and was acquired by 60 Hz video-tracking of a $1 \mu\text{m}$ diameter bead tethered to a DNA molecule in a magnetic tweezer (see Ribbeck (Ref. 28) for experimental procedure). The trace was split into half-overlapping blocks of 512 points and Hann-windowed before PSD calculation. The best-fit parameters from the $P_{A,B}$ fit are $\kappa = 2.46 \times 10^{-4}$ pN/nm and $\alpha = 1.39 \times 10^{-5}$ pN s/nm.




FIG. 3. (Color online) Plots of PSDs measured from time-domain simulations compared to the analytical PSDs, P_A (gray dashed line; Eq. (8)) and $P_{A,B}$ (solid black line; Eq. (7)). The measured PSDs (colored lines) correspond to averages of 10^4 simulations. Each simulation contained 1024 points with $\alpha = 10^{-5}$ pN s/nm and $f_s = 100$ Hz, with each point representing the average of $m = 1, 2, 4, 8, 16$ simulation time steps, as labeled. The spring constant was set to give (top) $f_c = 5$ Hz, or (bottom) $f_c = 50$ Hz. The insets highlight the convergence of the measured PSD to $P_{A,B}$ with increasing m . We note that the bottom figure has an apparent corner frequency which is due to the low-pass filtering effect of the camera (see Eq. (5)).

This directly indicates that Eq. (7) correctly accounts for the camera's averaging effect on the measured probe position.

C. Estimating the PSD from a time series of data

Spectral leakage and windowing: Given a set of N discrete measurements of probe position vs. time, an estimate, $\tilde{P}(f)$, of the PSD can be made from the discrete Fourier transform. Directly comparing $\tilde{P}(f)$ to $P_{A,B}(f)$, (Eq. (7)) leads to difficulties due to the finite duration of the measurement: the frequency-domain analysis of the Langevin equation implicitly assumes that the trajectory is infinitely long, while any measurement will naturally be of finite duration. This truncation leads to the phenomenon of spectral leakage, in which some power present in the signal at one frequency shows up at other frequencies. For white noise, spectral leakage has no effect, as neighboring frequencies leak equal and opposite power to each other, resulting in zero net change. However, for limited bandwidth signals (such as the Lorentzian, where the power is focused in the range $f < f_c$), spectral leakage leads to PSD distortions.

Spectral leakage can be minimized by “windowing” the input data.⁴ In particular, given N measurements of probe position, $x_j, j \in \{0, 1, 2, \dots, N-1\}$, at sampling rate f_s , the PSD can be estimated from:

$$\tilde{P}_{w,N} = \frac{1}{f_s N} \left| \sum_{j=0}^{N-1} w_j x_j e^{-2\pi i j k / N} \right|^2, \quad (9)$$

where w_j is the windowing function, which can be chosen from a range of standard functions, all of which taper to zero at the edges of the data trace.¹⁸ Here, we use a Hann window, defined by

$$w_j = \sqrt{\frac{8}{3}} \sin^2\left(\frac{\pi j}{N}\right). \quad (10)$$

Application of the window reduces the total power in the PSD in a frequency-independent manner; this is corrected by the leading factor of $\sqrt{8/3}$.

Blocking: In the absence of stochastic noise, the windowed PSD $\tilde{P}_{w,N}$ would converge on the analytical aliased filtered Lorentzian $P_{A,B}$. However, a feature of the discrete Fourier transform is that increasing N does not decrease the stochastic noise; instead, it increases the number and density of estimates in frequency space. Moreover, while the noise in x_j is normally distributed, the stochastic noise in a single power spectrum is distributed exponentially, complicating least-squares fitting routines that assume normally distributed noise.

The stochastic noise can be decreased and the noise distribution made more normal by compressing the data. Typically, this is done by blocking:¹ the data trace is separated into b blocks of length m , the PSD of each block, \tilde{P}_m , is calculated, and averaged with all others to give the blocked PSD $\tilde{P}_b = \sum \tilde{P}_m / b$. Blocking should be adjusted when also windowing, as application of the window minimizes the contribution of data near the termini of each block. All of the data

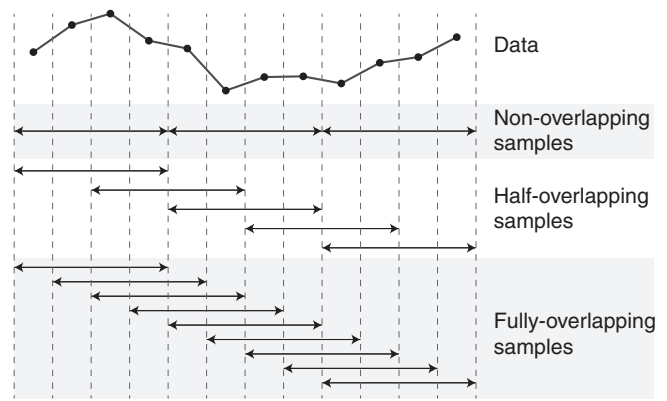


FIG. 4. Illustration of different manners of binning (“blocking”) data, shown here for an example trace containing $N = 12$ points split into bins of length $m = 4$. Non-overlapping sampling corresponds to no sharing of points between bins, and gives 3 bins in this example. This is used in calculation of the “standard” AV, or in non-windowed estimates of the PSD; neither method is optimal, as described in the main text. Half-overlapping sampling uses bins that are offset by half the binning interval from each neighbor, giving 5 bins offset by two points in this example. This insures maximal data usage when calculating windowed PSDs, and is known as Welch’s method (Ref. 6). Fully-overlapping sampling uses bins offset by a single point from each neighbor, giving 9 bins here. This method is used to calculate the OAV.

can be efficiently utilized by following Welch’s method⁶: instead of blocking into non-overlapping bins, $b = N/m$, the data trace is blocked into half-overlapping bins (Fig. 4), giving $b = 2N/m - 1$. The windowed PSD of each block, $\tilde{P}_{w,m}$, is calculated, and averaged with all others to give the final blocked, windowed PSD estimate $\tilde{P}_{w,b} = \sum \tilde{P}_{w,m} / b$. This insures efficient exploitation of all of the available data while keeping each bin statistically independent of its neighbors.

The noise distribution in a PSD: Recently, Nørrelykke and Flyvbjerg¹ have pointed out that, even after blocking, the distribution of noise in the estimated PSD is typically still far from normally distributed, but that fitting can proceed in a precise manner by accounting for the actual distribution: averaging b spectra together causes the noise distribution to be the convolution of b exponential distributions, which is a Gamma distribution. In a Gamma distribution, the probability f of observing a certain value x depends on parameters η, θ as

$$f(x; \eta, \theta) = x^{\eta-1} \frac{e^{-x/\theta}}{\theta^\eta \Gamma(\eta)}, \quad (11)$$

where $\Gamma(\eta)$ is the gamma function. The two free parameters are the shape parameter η and the scale parameter θ . For a PSD calculated from b blocks, and whose true value is given by $P_{A,B}(f_k)$, the probability of observing $\tilde{P}_{b,w}$ is proportional to $f(\tilde{P}_{b,w}; b, P_{A,B}(f_k)/b)$; that is, the shape is $\eta = b$ and the scale is $\theta = P_{A,B}(f_k)/b$. Explicitly using the Gamma distribution allows for the application of bias-free maximum-likelihood estimation.¹ We discuss this below (Sec. IV), in conjunction with a similar discussion of fitting to the Allan variance.

III. ALLAN VARIANCE

A. Definition of the Allan variance

An alternative to frequency-domain analysis of probe fluctuations is the AV, denoted by $\sigma^2(\tau)$. The AV is a time-domain measure: it is half of the ensemble-averaged variance of the difference between two consecutive samples of position,^{19,20} where each sample is itself a local average of probe position (Eq. (12)). The timescale τ denotes both the time between consecutive samples, and the time over which each sample is averaged. Thus, for a probe position $x(t)$ and integer j , the AV is defined by

$$\sigma^2(\tau) = \frac{1}{2} \langle (\bar{x}_{\tau, j+1} - \bar{x}_{\tau, j})^2 \rangle, \quad (12)$$

where $\bar{x}_{\tau, j}$ is given by

$$\bar{x}_{\tau, j} = \frac{1}{\tau} \int_{\tau(j-\frac{1}{2})}^{\tau(j+\frac{1}{2})} x(t) dt \quad (13)$$

$$= [x * \Pi_{\tau}](j\tau). \quad (14)$$

The identical forms of Eqs. (3) and (14) emphasize that the camera measures an average probe position, exactly as called for by the AV. Thus, while the PSD requires special corrections for a boxcar filtering instrument function, the AV does not, simplifying its use in this situation. In situations where data contains dead-time, for example, when the frame-read time of a CCD approaches the frame acquisition time, the AV can be compensated by bias correction when the noise sources have known power laws.^{21,22} In our experiments using a JAI CV-A10CL camera, the ratio of frame-read time to frame-acquisition time is very small, $3.5 \mu\text{s} \times 60 \text{ Hz} \ll 1$, so the instrument function is well-approximated by the boxcar function introduced in Eq. (3).

B. An analytical expression for the AV of a damped harmonic oscillator

In analogy to Sec. II B, we derive an analytical expression for the AV for a Brownian probe in a harmonic well. By expanding the product in Eq. (12), we find that the AV is related to the variance and autocorrelation of \bar{x}_{τ} :

$$\sigma^2(\tau) = \langle \bar{x}_{\tau}^2 \rangle - \langle \bar{x}_{\tau, j+1} \bar{x}_{\tau, j} \rangle, \quad (15)$$

where we have made use of the stationary nature of the process, so $\langle \bar{x}_{\tau, j+1}^2 \rangle = \langle \bar{x}_{\tau, j}^2 \rangle \equiv \langle \bar{x}_{\tau}^2 \rangle$. The Wiener-Khinchin theorem can be used to relate the variance and autocorrelation of \bar{x}_{τ} to its PSD; in turn, given the properties of convolution, the PSD of \bar{x}_{τ} depends on the PSD of $x(t)$ and of Π_{τ} . This allows the AV to be related to $P(f)$, giving:⁸

$$\sigma^2(\tau) = \int_{-\infty}^{\infty} \frac{4 \sin^4(\pi f \tau) P(f)}{(\pi f \tau)^2} df. \quad (16)$$

Using Eq. (2), this integral can be performed to obtain an analytical expression for the AV of an overdamped bead in a




FIG. 5. Experimentally derived OAV calculating using bins of octaved lengths (points), or all possible lengths (gray line), compared to the best-MLE-fit to the theoretical prediction (Eq. (17), dashed line). The data were acquired as described in Fig. 2. The best-fit parameters are $\kappa = 2.40 \times 10^{-4}$ pN/nm and $\alpha = 1.38 \times 10^{-5}$ pN s/nm.

harmonic well:

$$\sigma_{SMM}^2(\tau) = \frac{2k_B T \alpha}{\kappa^2 \tau} \left(1 + \frac{2\alpha}{\kappa \tau} e^{-\frac{\kappa \tau}{\alpha}} - \frac{\alpha}{2\kappa \tau} e^{-\frac{2\kappa \tau}{\alpha}} - \frac{3\alpha}{2\kappa \tau} \right). \quad (17)$$

This expression has long been known in the frequency-stability literature,²³ but we have rewritten it here in terms of α and κ with special application to SMM experiments. For long times, $\tau \gg \alpha/\kappa$, the AV reduces to $\sigma_{SMM}^2(\tau) \approx 2k_B T \alpha / \kappa^2 \tau$, consistent with previously published approximations.¹⁰⁻¹² The large time limit can be understood by noting that neighboring \bar{x}_{τ} values become uncorrelated for large τ , so the variance arises from the standard error in determining the mean of a normally distributed value, i.e., variance $\sim 1/(\# \text{ of points}) \sim 1/\tau$. In the opposite limit, $\tau \ll \alpha/\kappa$, we find $\sigma^2(\tau) \approx 2k_B T \tau / 3\alpha$; that the variance increases linearly with τ arises from the purely diffusive nature of the probe motion at short times. Thus, a feature of the AV of a damped, Brownian harmonic oscillator is that it increases at short times and decreases at long times (Fig. 5); the location of the intervening peak can be found numerically to be $\tau_{max} \approx 1.89 \alpha/\kappa$, i.e., the peak scales with the system's correlation time, α/κ .

C. Estimating the AV from a time series of data

Overlapping AV: For a discrete set of measurements of probe positions x_j at times $t_j = j\tau_s$ for $j = 1, \dots, N$, the “standard” AV can be calculated for timescale $\tau = m\tau_s$ by computing the mean of successive bins of m points, then taking the mean-squared difference of the resulting $N/m - 1$ neighboring pairs of bins. So, in particular, the first pair of bins corresponds to the sets of points $(1, \dots, m)$ and $(m+1, \dots, 2m)$, the second pair are points $(m+1, \dots, 2m)$ and $(2m+1, \dots, 3m)$, and so on, through the N th point. However, the standard AV is known to not fully utilize the data.²⁴ A superior option is the “overlapping” AV (OAV), which uses every possible m -bin (Fig. 4). In particular, the first OAV pair of bins is $(1, \dots, m)$ and $(m+1, \dots, 2m)$, the second pair is $(2, \dots, m+1)$ and $(m+2, \dots, 2m+1)$, and so on, through the N th point. The

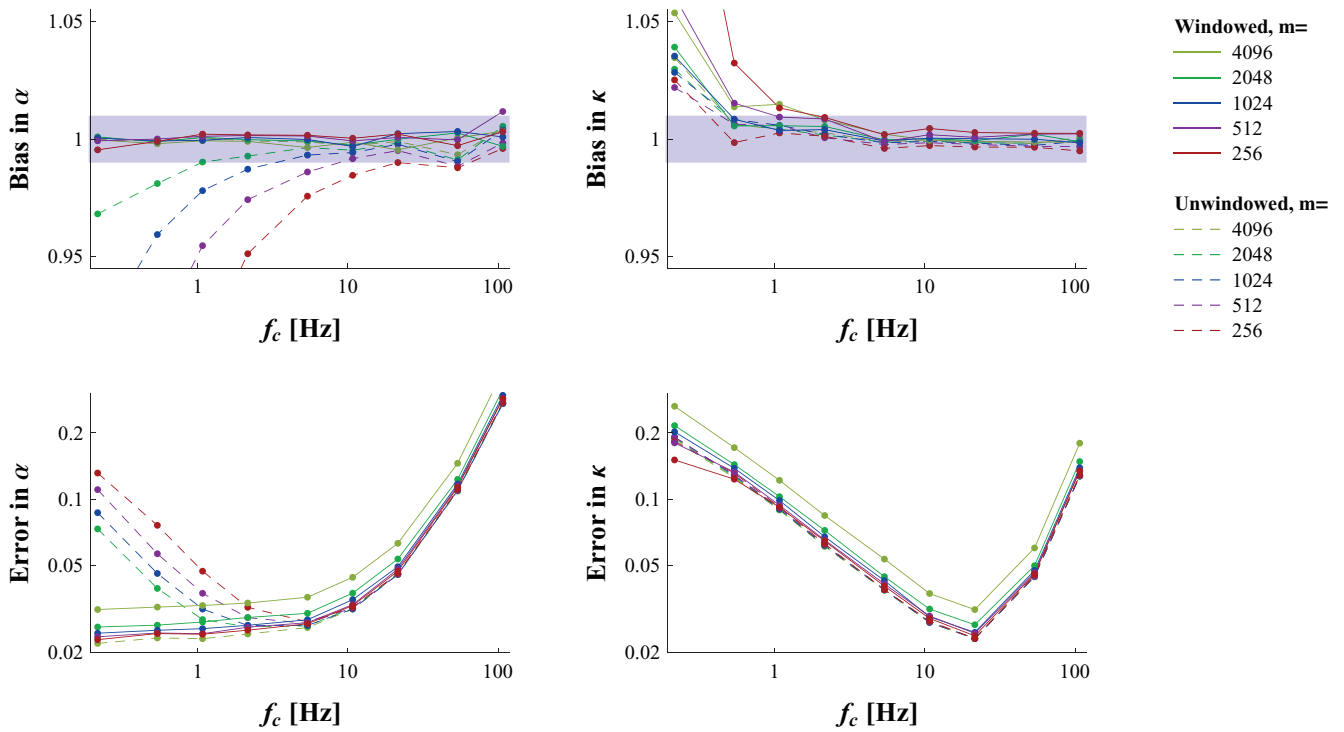


FIG. 6. (Color online) Bias and error in the best-fit parameters α and κ found from MLE fitting of Eq. (7) to simulated PSDs. A total of 1000 simulations of a video-tracking experiment were performed; each simulation contained 4096 points and used $\alpha = 10^{-5}$ pN s/nm and $f_s = 100$ Hz. The spring constant κ was set to give a range of f_c between 0.2 and 100 Hz. Each trajectory was blocked into subsequences of length m , as indicated, and the PSD was calculated either using a Hann window (solid lines) or not (dashed lines). We used half-overlapping blocks (Fig. 4) for the windowed calculation, and non-overlapping blocks for the unwindowed calculation. The bias (top row) is reported as the median best-fit value across all simulations relative to the true value, while the error (bottom row) is reported as the standard deviation in best-fit results relative to the median value. In the bias plots, the gray bar indicates values within $\pm 1\%$ of being unbiased.

standard and overlapping AVs are identical for $m = 1$ and $m = N/2$, but elsewhere the OAV more efficiently exploits the available data, and so provides a better estimate of the true AV. The experimentally derived OAV can be computed from²⁵

$$\tilde{\sigma}_m^2 = \frac{1}{2(N-2m)(m\tau_c)^2} \sum_{k=1}^{N-2m} (x_{k+2m} - 2x_{k+m} + x_k)^2. \quad (18)$$

We use only the OAV in calculations throughout this article.

Octave sampling: The OAV can be determined for any integer value of the bin size, $m = 1, 2, \dots, N/2$ (see Fig. 5). However, upon estimating and plotting the OAV, it is clear that neighboring values of m have a correlated noise, which stems from the nearly-identical sets of points used to calculate the OAV for similar values of m . While calculating all possible m values might be useful for visualizing the data, the noise correlations preclude most fitting algorithms, which assume independent error in the fitted points. Allan pointed out that this issue can be solved by octave sampling,²⁶ i.e., only using bin sizes in powers of 2: $m = 1, 2, 4, 8, \dots$ up to $m \leq N/2$. Using bins of 2^n ensures a minimal correlation between successive estimates, and improves the results of fitting, as discussed in Sec. V and shown in Fig. 7.

Error in the AV estimate: At the root of the AV is the difference $\bar{x}_{j+1} - \bar{x}_j$, which is normally distributed with mean zero. Thus, an estimate for the AV (which is a sum of squares of such differences) is Gamma distributed. The estimated blocked and windowed PSD, $\tilde{P}_{b,w}$ (Sec. II C), is also

Gamma distributed, and $\tilde{P}_{b,w}$ has the same shape factor, $\eta = b$, for all frequencies f_k . However, the estimated AV does *not* have a constant shape factor because the number of available differences depends on the bin length m ; this is the major complicating factor in working with the AV relative to the PSD.

Generally, the shape factor is given by the number of exponential distributions that are convolved to form the metric (this is why $\eta = b$ for the PSD). The standard AV for a given m is calculated from $N/m - 1$ differences; since each difference has a Chi-Square distribution, two of which must be convolved to get an exponential distribution, the shape factor is

$$\eta_{AV}(m) = \frac{1}{2} \left(\frac{N}{m} - 1 \right). \quad (19)$$

As discussed, the OAV more efficiently exploits the available data, meaning that there are more available degrees of freedom per m . Thus, for the OAV, the shape factor of Eq. (19), while nearly correct, is an underestimate. The exact calculation of the available degrees of freedom for the OAV is quite involved, and beyond the scope of this paper; it has been the subject of much discussion in the literature.^{20,27} However, we have found that using η from the relatively simple Eq. (19) gives near-identical results to more exact, involved estimates, so we use Eq. (19) in all computations discussed below.

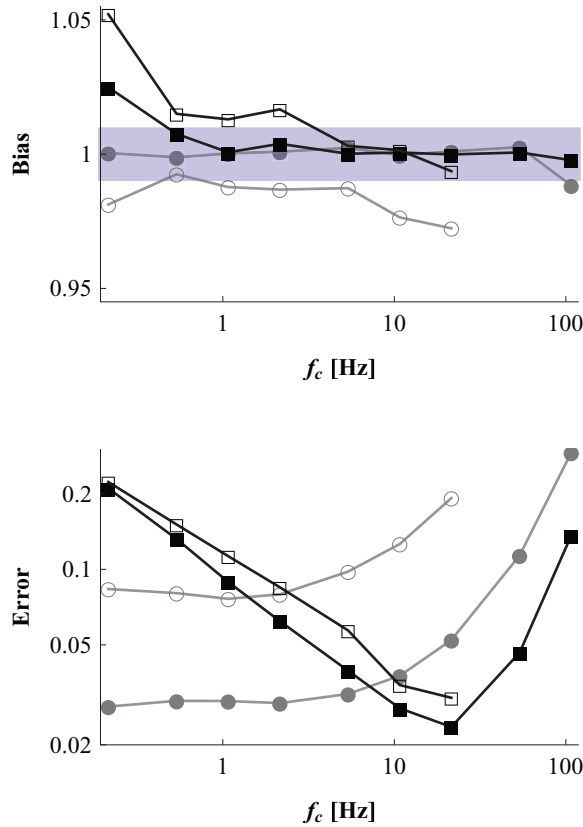


FIG. 7. (Color online) Bias and error in the best-fit parameters α and κ found from MLE fitting of Eq. (17) to simulated AVs. Trajectories were simulated with parameters as described in Fig. 6. From each trajectory, we calculated and fit to both the octave-sampled, overlapping AV (α : \bullet and κ : \blacksquare), and the all- τ , overlapping AV, (α : \circ and κ : \square). The bias (top) is reported as the median best-fit value across all simulations relative to the true value, while the error (bottom) is reported as the standard deviation in best-fit results relative to the median value. In the bias plots, the gray bar indicates values within $\pm 1\%$ of being unbiased.

IV. MAXIMUM LIKELIHOOD ESTIMATION

As discussed by Nørrelykke and Flyvbjerg,¹ least-squares fitting to the PSD always leads to bias in the best-fit parameters. The bias occurs because of the least-squares assumption that the noise is normally distributed, whereas the noise in the PSD is Gamma-distributed (Sec. II C). Similarly, the noise in the AV is Gamma distributed (Sec. III C), so least-squares fitting of AVs will also lead to biases. Although these biases can be analytically corrected,¹ an unbiased and more accurate alternate to least-squares fitting is to perform a maximum likelihood estimate (MLE) that accounts for the correct noise distribution in the data.¹ The MLE approach stems from the assumption that the optimal parameters will maximize the joint probability $p_J = \prod_k p_k(\tilde{y}_k; y_k(\alpha, \kappa))$, where $p_k(\tilde{y}_k; y_k(\alpha, \kappa))$ is the probability of measuring a value \tilde{y}_k given a true value y_k that depends on parameters (α, κ) . By substituting the Gamma distribution, Eq. (11), for p_k , it can be shown that maximizing p_J is equivalent to minimizing the cost function:¹

$$F' = \sum_k \eta_k \left(\frac{\tilde{y}_k}{y_k(\alpha, \kappa)} + \ln y_k(\alpha, \kappa) \right), \quad (20)$$

where η_k is the shape parameter for point k . Eq. (20) applies to both the PSD and AV. For the PSD, the shape factor is constant for all points, and equal to the number of blocks, $\eta_k = b$, while the true value is found from Eq. (7) as $P_{A,B}(f_k, \alpha, \kappa)$. For the AV, the shape factor is point-dependent, $\eta_k = \eta_{AV}$, where η_{AV} is defined in Eq. (19). The true value is found from Eq. (17) as $\sigma_{SMM}^2(\tau = k\tau_c, \alpha, \kappa)$. In either case, minimization of Eq. (20) with respect to (α, κ) can be carried out relatively efficiently using most numerical analysis software; we use the MLE approach to generate the results described below.

V. BIAS AND ERRORS IN AV AND PSD FITS

To test and compare the AV and PSD algorithms, we performed a set of simulations of probe trajectories using a variety of input parameters and tested the ability of the algorithms to correctly extract the known parameters. In particular, we simulated video-tracking trajectories using a frequency-space algorithm, as discussed in the Appendix. All trajectories contained $N = 4096$ points sampled at $f_s = 100$ Hz. The drag was $\alpha = 10^{-5}$ pN s/nm, as appropriate for a spherical bead of diameter $1 \mu\text{m}$ in water. We varied the spring constant, κ , between $\approx 1.4 \times 10^{-5}$ and $\approx 6.8 \times 10^{-3}$ pN/nm; this gave a set of corner frequencies f_c that ranged from ≈ 0.2 to ≈ 100 Hz. For each spring constant, we simulated 1000 trajectories. Each trajectory was processed according to Secs. II C and III C to give estimates of the PSD and AV, and fit using the MLE approach to extract the optimal parameters (α, κ) . All numerical work was carried out in *Mathematica*, with the MLE minimization carried out using the *FindMinimum* function.

To emphasize the need for windowing in the PSD calculation, we compared the best-fit parameters from PSDs that did or did not use a Hann Window (Fig. 6) for a variety of block lengths m . When windowing, we used half-overlapping samples (blocks), while the unwindowed analysis used non-overlapping samples; so, for example, for blocks of length $m = N/4 = 1024$, the windowed PSD was calculated from $b = 7$ blocks, while the unwindowed PSD used $b = 4$ blocks. Both methods performed similarly in estimating the spring constant, giving unbiased estimates with relatively low noise over a broad range of f_c . In contrast, the use of windowing was clearly superior in the estimate of drag: the unwindowed method gave significant bias and increased stochastic error, in the estimate of α , particularly at lower f_c . For both parameters, the stochastic error in the windowed estimates decreased as m decreased, with $m = N/8 = 512$ points or fewer needed to obtain the minimal error. We attribute this to end-effects: using a single block ($m = N$) and applying a window cause a loss of information near each terminus of the trace, which increases the error. Instead, using many overlapped blocks decreases the information lost at the termini (because the blocks are shorter), while overlapping avoids information loss in the interior of the trace.

The results of the AV analysis (Fig. 7) demonstrate the importance of octave-sampling: we compared the best-fit parameters from AVs estimated using bins of all possible lengths (all- τ), $m = 1, 2, 3, \dots, N/2$, and those estimated from octave-bins, $m = 1, 2, 4, \dots, N/2$. The octave-sampled method

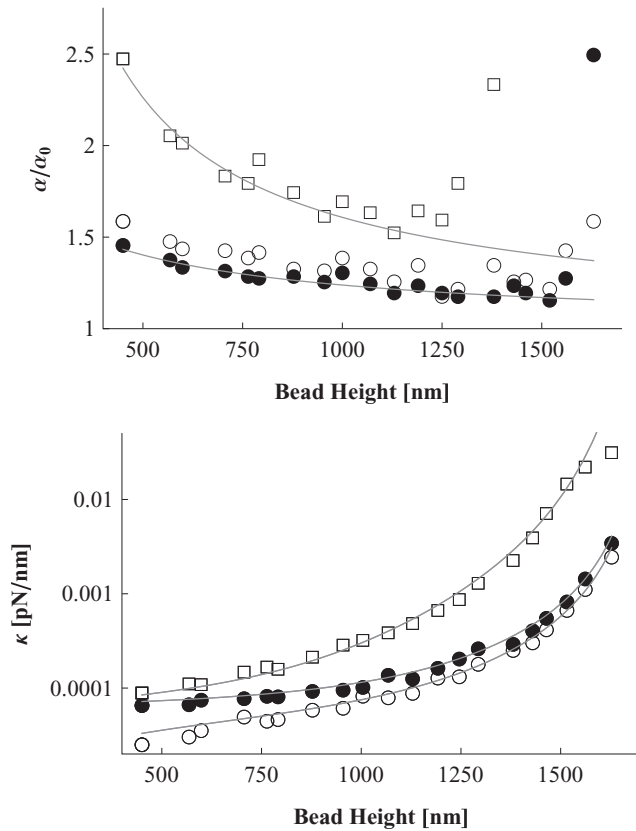


FIG. 8. Experimental results showing the relationship between the measured drag coefficients (α_x/α_0 : \bullet), (α_y/α_0 : \circ), (α_z/α_0 : \square), and the bead height above the surface in comparison to the theoretical Faxén' correction (Ref. 29). For all theoretical curves, we use radius $R = 530$ nm, viscosity $\eta = 0.00115$ Ns/m², and $\alpha_0 = 6\pi\eta R$. The bottom panel demonstrates the relationship between the extension of dsDNA and the measured stiffness coefficients (κ_x : \bullet), (κ_y : \circ), (κ_z : \square) in comparison to the predictions of a model that combines dsDNA's worm-like-chain elasticity (Ref. 30) with the inverted-pendulum model of a magnetic tweezer (Ref. 15). For all analytical worm-like-chain curves, we use persistence length $l_p = 53$ nm, contour length $L_0 = 1720$ nm, and $R = 530$ nm.

gives bias-free estimates of α and κ across nearly the entire range of f_c investigated, while the all- τ approach has a small, but significant, bias in both parameters. More significantly, octave-sampling leads to lower error in the estimated parameters, particularly in α where the noise in the all- τ estimate is roughly three-fold worse. Finally, we were unable to get the all- τ method to consistently converge for the largest two f_c values investigated, and so, Fig. 7 shows no all- τ point at $f_c = 50$ or 100 Hz. Generally, numerical minimization routines must be carefully handled to insure convergence. Here, we attribute the difference in the robustness of the octave and all- τ algorithms to the different numbers of points that are fit: whereas the all- τ method uses $N/2$ points, the octave-sampled AV fits to only $\log_2 N$ points, a vast decrease that increases both the robustness and speed of the minimization procedure.

VI. EXPERIMENTAL RESULTS

To demonstrate the applicability of our AV fitting method to experimental data, we use a magnetic tweezer instrument to stretch double-stranded DNA molecules over a range of

forces.²⁸ From each particle trajectory taken at a constant force, we calculate the octave-sampled, overlapping AV, and use MLE fitting of Eq. (17) to estimate α_x , α_y , α_z , κ_x , κ_y , and κ_z . We then compare these parameters to those independently estimated by theories of bead drag or polymer elasticity.

In Figure 8 we plot the drag coefficients α_x , α_y , α_z as a function of the average bead height. As the magnetic force is decreased, the bead height decreases and the drag coefficient of the bead increases due to surface effects, which can be theoretically predicted using the Faxén correction.²⁹ Both α_x and α_z closely match the expected behavior. We consistently found that $\alpha_y > \alpha_x$, which we attribute to the rotational constraints imposed on the magnetic beads by the applied magnetic field. At high force, the fitting routine was unable to accurately determine α_z , since the corner frequency of the system approached and exceeded the Nyquist frequency of our camera.

In Figure 8 we also present the experimentally derived κ_x , κ_y , κ_z values as a function of extension. At each measured tether extension L , we can generate three independent predictions for κ_x , κ_y , κ_z by, first, using the worm-like chain model for double-stranded DNA (Ref. 30) to convert L into an estimate of force, F . Then, we use the inverted-pendulum model for magnetic tweezers¹⁵ to convert F and L into expected stiffness through the relations: $\kappa_z = \partial F/\partial z$, $\kappa_x = F/L$, and $\kappa_y = F/(L + R)$. For all predictions, we use a persistence length of 53 nm, consistent with previous measurements,³⁰ and a contour length of 1720 nm. Although we do not correct for bead mistethering,³¹ we nonetheless find a remarkably good agreement between theory and experiment over a wide range of forces in all six measured parameters.

VII. CONCLUSION

In this article, we have presented two comprehensive and novel strategies for calibrating the spring constant κ and measuring the drag α in a SMM instrument that uses full-frame video tracking to measure probe position. These strategies are based on the AV and PSD. In particular, given N discrete measurements of position $x_j(t)$, with $j \in \{0, 1, 2, \dots, N-1\}$, we suggest the following fitting procedure for the PSD:

1. Divide the trace into half-overlapping bins of length $m \leq N/8$, resulting in $b = 2N/m - 1$ bins.
2. For each bin, multiply x_j by a Hann window and calculate a windowed PSD $\tilde{P}_{w,m}$ for each bin (Eq. (9)).
3. Average the Power spectra over all bins to obtain a single blocked and windowed $\tilde{P}_{b,w}$.
4. Minimize the cost function in Eq. (20) using $y_k = P_{A,B}(\alpha, \kappa, f_k)$ and $\tilde{y}_k = \tilde{P}_{b,w}(f_k)$ to obtain α and κ .

The corresponding recipe for calculating α and κ using an AV fit to the same data is

1. Calculate the octave-sampled OAV $\tilde{\sigma}_m^2(\tau)$ using Eq. (18).
2. Calculate the shape factor $\eta_{AV}(m)$ from Eq. (19).
3. Minimize the cost function in Eq. (20) using $\eta_{AV}(m)$, $y_k = \sigma_{SMM}^2(\tau, \alpha, \kappa)$, and $\tilde{y}_k = \tilde{\sigma}_m^2(\tau)$ to obtain α and κ .

We conclude by comparing the AV and PSD approaches: as can be seen from the results (Figs. 6 and 7), both the AV and

PSD algorithms are capable of unbiased, low-error parameter estimates of equal quality. Thus to differentiate between the two we focus on ease-of-use, both in terms of general complexity and, in particular, for the numerical minimization required by the MLE algorithm. For both criteria, we find that the AV algorithm outperforms the PSD: In terms of complexity, while both the AV and PSD require fitting with a MLE algorithm rather than least-squares, the PSD approach requires applying multiple corrections, particularly for instrument response, discrete sampling (aliasing), and finite duration of the data (spectral leakage). The AV is not subject to any of these three issues, as the instrument response intrinsic to video tracking exactly matches the assumptions made in calculating the AV, and the PSD finite-sampling effects that manifest in the DFT sum are simply not present in the time-domain AV calculation. Further algorithmic simplicity is achieved by the AV in that no choice of blocking number, b , needs to be made; this is particularly important since we show that the choice of b has a minor, but significant, effect on the fitting results of the PSD.

Numerical minimization of the cost function (Eq. (20)) is also more straightforward with the AV because it achieves superior data compression without sacrificing accuracy or precision. The octave-sampled OAV for a data trace of N points contains $\log_2 N$ points, while the estimated PSD (using the maximum bin length that still reduces stochastic error) contains $N/8$ points. Therefore, the total number of points that must be fit grows linearly with N for the PSD, but only logarithmically with N for the AV. For example, for the $N = 4096$ traces simulated here, the AV contains 12 points, while the PSD contains 513 points. This means that the MLE cost function (Eq. (20)) contains only 12 terms when working with the AV, but contains 513 terms for the PSD. Because of the advantages in simplicity, speed, and robustness of fitting with the AV as compared to the PSD, we suggest that the Allan variance is the ideal calibration method for many SMM instruments.

ACKNOWLEDGMENTS

This work was supported by the NIH under Grant No. 1R21GM079581-01A1, and the NSF under Grant No. PHY-0748564. B.L. acknowledges support from an NSERC PGS-D fellowship. We thank T. M. Squires, D. B. McIntosh, and A. A. M. Saleh for helpful discussions.

APPENDIX: EFFICIENT SIMULATION OF VIDEO-TRACKING TRAJECTORIES IN FREQUENCY SPACE

The efficient simulation of probe trajectories is important as a tool to test calibration algorithms. A typical strategy to simulate particle trajectories is to recast the Langevin equation, Eq. (1), as an equation of probability flux (a Smoluchowski equation), then apply an iterative scheme to generate successive discrete time points in the trajectory.¹⁷ The discrete points generated in this way are not a realistic representation of data measured by a video-tracking instrument, as the method implies that the instrument has an infinitely-fast time

response. A more correct approach is to generate points with a small time step, then to average groups of points together, thus approximating the averaging effect of the camera. However, this can be quite onerous as it requires the simulation of many more points than are finally desired, particularly when f_c approaches f_s (see Fig. 3).

There is an alternative approach: since the Fourier transform of a trajectory contains all the information needed to recapitulate the trajectory, it is possible to use $P_{A,B}$ (Eq. (7)) as an efficient simulation tool. In particular, we can get an N -point trajectory by simulating N points in the frequency domain that are distributed according to $P_{A,B}$, then applying an inverse Fourier transform. Here, we briefly describe this approach.

Consider a set of points x_j that represent a probe trajectory at N time points $t_j = j\tau_c = j/f_s$; we assume N is an even number. We define the Fourier transform of x_j to be \hat{x}_k , which is a complex number and is defined for a set of N frequencies $f_k = kf_s/N$, with $k = -(N/2) + 1, \dots, -1, 0, 1, \dots, N/2$. All x_j are real numbers, which constrains the allowed values of \hat{x}_k :

$$\begin{aligned} \text{Re}(\hat{x}_k) &= \text{Re}(\hat{x}_{-k}), \\ \text{Im}(\hat{x}_k) &= -\text{Im}(\hat{x}_{-k}), \\ \text{Im}(\hat{x}_0) &= \text{Im}(\hat{x}_{N/2}) = 0. \end{aligned} \quad (\text{A1})$$

Thus, it suffices to choose N random values, one each for $k = 0$ and $k = N/2$, and two each (one imaginary and one real) for $k = 1, \dots, (N/2) - 1$, suggesting the following algorithm:

1. Use a random number generator to generate a set of N points, q_j for $j = 0, \dots, N - 1$, that are normally distributed with mean zero and unit variance.
2. Set the real-only values at zero frequency and the Nyquist frequency by multiplying a random number by the expected value at those frequencies:

$$\begin{aligned} \hat{x}_0 &= P_{A,B}(0)q_0, \\ \hat{x}_{N/2} &= P_{A,B}(f_s/2)q_{N/2}. \end{aligned} \quad (\text{A2})$$

3. Set the remaining complex values, \hat{x}_k for $k = 1, \dots, (N/2) - 1$, by using different random numbers for the real and imaginary components. Each is scaled by $P_{A,B}$, but must be divided by $\sqrt{2}$ to account for ratio between the expected value of the magnitude and each complex component:

$$\hat{x}_k = \frac{1}{\sqrt{2}}(P_{A,B}(f_k)q_k + iP_{A,B}(f_k)q_{k+N/2}). \quad (\text{A3})$$

4. Apply the constraints, Eq. (A1), to find \hat{x}_k for all k .
5. Calculate x_j by applying a discrete inverse Fourier transform to \hat{x}_k :

$$x_j = \sqrt{\frac{f_s}{N}} \sum_k \hat{x}_k e^{2\pi i j k / N}. \quad (\text{A4})$$

¹S. F. Nørrelykke and H. Flyvbjerg, *Rev. Sci. Instrum.* **81**, 075103 (2010).

²W. P. Wong and K. Halvorsen, *Opt. Lett.* **34**, 277 (2009).

³C. Gosse and V. Croquette, *Biophys. J.* **82**, 3314 (2002).

⁴F. J. Harris, *Proc. IEEE* **66**, 51 (1978).

- ⁵W. P. Wong and K. Halvorsen, *Opt. Express* **14**, 12517 (2006).
- ⁶P. Welch, *IEEE Trans. Audio Electroacoust.* **15**, 70 (1967).
- ⁷J. von Neumann, *Ann. Math. Stat.* **12**, 367 (1941).
- ⁸J. A. Barnes, A. R. Chi, L. S. Cutler, D. J. Healey, D. B. Leeson, T. E. McGunigal, J. A. Mullen, W. L. Smith, R. L. Sydnor, R. F. C. Vessot, and G. M. R. Winkler, *IEEE Trans. Instrum. Meas.* **IM-20**, 105 (1971).
- ⁹T. Walter, in *Proceedings of the 24th Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting* (National Aeronautics and Space Administration, Goddard Space Flight Center, 1993), Vol. 1, pp. 413–426.
- ¹⁰F. Czerwinski, A. C. Richardson, and L. B. Oddershede, *Opt. Express* **17**, 13255 (2009).
- ¹¹M. Andersson, F. Czerwinski, and L. B. Oddershede, *J. Opt.* **13**, 044020 (2011).
- ¹²O. Otto, F. Czerwinski, J. L. Gornall, G. Stober, L. B. Oddershede, R. Seidel, and U. F. Keyser, *Opt. Express* **18**, 22722 (2010).
- ¹³G. M. Gibson, J. Leach, S. Keen, A. J. Wright, and M. J. Padgett, *Opt. Express* **16**, 14561 (2008).
- ¹⁴B. Lukić, S. Jeney, C. Tischer, A. J. Kulik, L. Forró, and E.-L. Florin, *Phys. Rev. Lett.* **95**, 160601 (2005).
- ¹⁵A. J. W. te Velthuis, J. W. J. Kerssemakers, J. Lipfert, and N. H. Dekker, *Biophys. J.* **99**, 1292 (2010).
- ¹⁶K. Berg-Sørensen and H. Flyvbjerg, *Rev. Sci. Instrum.* **75**, 594 (2004).
- ¹⁷N. G. van Kampen, *Stochastic Processes in Physics and Chemistry* (North Holland, 2007).
- ¹⁸R. C. Castberg, Master's thesis, University of Oslo, 2008.
- ¹⁹D. W. Allan, *Proc. IEEE* **54**, 221 (1966).
- ²⁰W. J. Riley, *Handbook of Frequency Stability Analysis* (Hamilton Technical Services, 2007).
- ²¹P. Lesage and C. Audoin, *IEEE Trans. Instrum. Meas.* **28**, 6 (1979).
- ²²J. A. Barnes and D. W. Allan, in *Proceedings of the 19th Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting* (National Bureau of Standards, Boulder CO, 1987), pp. 227–234.
- ²³C. M. van Vliet and P. H. Handel, *Physica A: Stat. Mech. Appl.* **113**, 261 (1982).
- ²⁴D. A. Howe, D. W. Allan, and J. A. Barnes, in *Proceedings of the 35th Annual Frequency Control Symposium, Philadelphia, PA*, 1981.
- ²⁵T. Walter, *IEEE Trans. Instrum. Meas.* **43**, 69 (1994).
- ²⁶D. W. Allan, M. A. Weiss, and J. L. Jespersen, in *Proceedings of the 45th Annual Symposium on Frequency Control* (Institute of Electrical and Electronics Engineers, 1991), pp. 667–678.
- ²⁷A. Makdissi, F. Vernotte, and E. Clercq, *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **57**, 1011 (2010).
- ²⁸N. Ribeck and O. A. Saleh, *Rev. Sci. Instrum.* **79**, 094301 (2008).
- ²⁹E. Schäffer, S. F. Nørrelykke, and J. Howard, *Langmuir* **23**, 3654 (2007).
- ³⁰C. Bouchiat, M. D. Wang, J.-F. Allemand, T. Strick, S. M. Block, and V. Croquette, *Biophys. J.* **76**, 409 (1999).
- ³¹D. Klaue and R. Seidel, *Phys. Rev. Lett.* **102**, 28302 (2009).