Methods to Track Single-Molecule Trajectories

Stephen Anthony, Liangfang Zhang, and Steve Granick*

Departments of Chemistry, Chemical & Biomolecular Engineering, and Materials Science and Engineering, University of Illinois, Urbana, Illinois 61801

Received January 25, 2006. In Final Form: April 1, 2006

Methods are discussed to track single molecules in planar-supported phospholipid bilayers. Mainly, these methods constitute optimizations for a low signal-to-noise ratio and for the dim optical characteristics of single molecules. Algorithmic modifications to compensate and correct for misidentifications and misassignments are also described. One key advance, which exchanges the typically fragmented reconstruction of the molecules' motion for more complete trajectories, is the incorporation of information about the molecules' past and future positions into the tracking. Although the main point of these methods is to aid in extending methods of object tracking to the single-molecule regime, they may also find use in other situations where the signal-to-noise ratio is low.

I. Introduction

Tracking the motion of colloidal-sized objects through the use of video microscopy has evolved into a standard technique commonly applied to study colloidal suspensions and other objects whose fluorescence signal is bright because they contain multiple dyes.[1−3] To this end, software is freely available.[4] Recently, the rapid development of bioscience, especially cellular processes such as signaling, docking, and enzymatic dynamics, has made desirable analogous measurements concerning molecules labeled with a single fluorescent dye. Thanks to the advent of improved detection systems, the observation of single molecules is now possible, to the point that fluorescence microscopy can monitor the stochastic trajectories of single molecules in real time. For example, lipid molecules labeled with a single dye moiety are monitored to study membrane fluidity.[5−7] Single-dye-labeled proteins are used to study cellular signaling and synapse formation.[8,9] The chemical dynamics of some enzymes is investigated at the single-molecule level by monitoring fluorescence from their active sites.[10−12]

In response to this need, various methods are now standard for tracking single particles; for example, cross correlation, sum-absolute difference, centroid identification, and direct Gaussian fit, as reviewed by Cheezum et al.[13] Comparing these algorithms, the conclusion was reached that a direct Gaussian fit to the intensity distribution is the superior algorithm for tracking point sources, whose signal-to noise ratio is much less than for large particles.[13] However, these methods run into difficulty as concerns the case of single-molecule fluorophores (signal-to-noise ratio around 4), the situation that we will discuss in this article. In the case of single-molecule fluorophores, it is still problematical to discriminate relatively dim features from the upper end of noise. Inadequate choices may cause misidentification and misassignment.

Important attempts toward improving the signal-to-noise ratio of single fluorescent probes and improving the localization precision of their features were made previously. On the experimental side, using brighter fluorophores and increasing the data acquisition time of each frame to increase the fluorescence signal intensity or using proper optical filters to block the background noise of fluorescence from immersion oil and other extraneous sources will improve the situation. On the theoretical side, it is clear that precision localization inversely depends on the number of photons in a diffraction-limited spot in which background noise dominates the sources of noise.[14]

In this article, we describe algorithmic methods developed for tracking single fluorescent probes, methods apart from the advantages that are obtainable from having superior instrumentation. As will be explained below, these techniques peculiar to the study of single-molecule trajectories are used to flatten images, to identify features, to pair molecules, and to analyze the final stochastic trajectory of each single molecule.

To illustrate the implementation of these methods, the experimental portion of this article shows findings when these tracking algorithms were used to study the lateral diffusion of fluorescent-tagged lipids within supported phospholipid bilayers and compares the imputed translational diffusion coefficient of lipid to that deduced by other techniques—fluorescence correlation spectroscopy (FCS) and fluorescence recovery after photobleaching (FRAP).

This study was based on extensive modifications of the freely available software designed for tracking colloid-sized partices.[4] Because the general subject is how to deal with systems that present inherently low signal-to-noise ratios, it is hoped that these methods may be useful not only for tracking single molecules but also to the broader particle-tracking community when dealing with less than optimal experimental conditions.

* Corresponding author. E-mail: sgranick@uiuc.edu.

(1) Kasper, A.; Bartsch, E.; Sillescu, H. *Langmuir* **1998**, *14*, 5004.
(2) Maier, B.; Rädler, J. O. *Phys. Rev. Lett.* **1999**, *82*, 1911.
(3) Tseng, Y.; Kole, T. P.; Wirtz, D. *Biophys. J.* **2002**, *83*, 3162.
(4) Crocker, J. C.; Grier, D. G. *J. Colloid Interface Sci.* **1996**, *179*, 298.
(5) Schmidt, Th.; Schütz, G. J.; Baumgartner, W.; Gruber, H. J.; Schindler, H. *Proc. Natl. Acad. Sci. U.S.A.* **1996**, *93*, 2926.
(6) Fujiwara, T.; Ritchie, K.; Murakoshi, H.; Jacobson, K.; Kusumi, A. *J. Cell Biol.* **2002**, *157*, 1071.
(7) Ke, P. C.; Naumann, C. A. *Langmuir* **2001**, *17*, 3727.
(8) Kusumi, A.; Ike, H.; Nakada, C.; Murase, K.; Fujiwara, T. *Semin. Immunol.* **2005**, *17*, 3.
(9) Groc. L.; Heine, M.; Cognet, L.; Brickley, K.; Stephenson, F. A.; Lounis, B.; Choquet, D. *Nat. Neurosci.* **2004**, *7*, 695.
(10) Vale, R. D.; Funatsu, T.; Pierce, D. W.; Romberg, L.; Harada, Y.; Yanagida, T. *Nature* **1996**, *380*, 451.
(11) Ishijima, A.; Kojima, H.; Funatsu, T.; Tokunaga, M.; Higuchi, H.; Tanaka, H.; Yanagida, T. *Cell* **1998**, *92*, 161.
(12) Lu, H. P.; Xun, L.; Xie, X. S. *Science* **1998**, *282*, 1877.
(13) Cheezum, M. K.; Walker, W. F.; Guilford, W. H. *Biophys. J.* **2001**, *81*, 2378.
(14) Thompson, R. E.; Larson, D. R.; Webb, W. W. *Biophys. J.* **2002**, *82*, 2775.

10.1021/la060244i CCC: $33.50
Published on Web 05/11/2006

II. Instrumentation and Output

To demonstrate these single-molecule tracking methods and algorithms, the translational diffusion of lipids in solid-supported phospholipid bilayers was studied experimentally. The phospholipid DLPC, 1,2-dilauroyl-*sn*-glycero-3-phosphocholine, was selected because its gel-to-fluid phase transition of −1 °C was well below the experimental temperature, 23 °C. Into this, we doped at 0.01 ppm molar concentration DMPE, 1,2-dimyristoyl-*sn*-glycero-3-phosphoethanolamine, with polar headgroups labeled with rhodamine B. Both lipids were obtained from Avanti Polar Lipids, Inc. (Alabaster, AL) and were used without further purification. Using known protocols based on the fusion of single unilamellar vesicles,[15] supported bilayers of the mixture were prepared on hydrophilic quartz and then rinsed with copious amounts of PBS buffer (10 mM, pH 6.0) to remove unfused vesicles.

These experiments were carried out using a home-built single-molecule imaging apparatus consisting of a Zeiss Axiovert 200 microscope equipped with a water immersion objective lens (Zeiss Axiovert 135TV, 63× , NA = 1.2). The light from a diode-pumped Nd:YAG laser (CrystaLaser) was directed into the sample in an attenuated total reflection (ATR) geometry.[16] Fluorescence excited by the evanescent light was collected through the objective and detected by a back-illuminated electron-multiplying charge-coupled device (CCD) camera (Andor iXon DV-887 BI). The total system magnification of the CCD was 250 nm/pixel. Fluorescence images were collected at different rates but typically at a speed of 30 frames/s, and 200 frames were collected continuously. The video images were converted into digital format and analyzed using programs written in Matlab.

III. Methods

The process of particle tracking can be subdivided into four major steps: analyzing and processing the images, determining the location of the molecules within the images, pairing corresponding features between images, and analyzing the stochastic trajectories of single molecules.

i. Flattening Images. Despite the use of high-quality optics and a CCD camera with good quantum efficiency and low noise characteristics, the images are not ideal. In addition to randomly distributed noise that cannot easily be removed from the image, modulation of the background intensity and isolated bright pixels are also inevitably present. Compensation for these distortions, before any effort is made to localize molecules within the image, simplifies the computational work needed to distinguish features from noise and also reduces the number of spurious features found. The most noticeable example of such distortion is long-range spatial modulation of the background brightness, which is primarily due to nonuniform laser illumination of the sample area.

Traditionally, boxcar filtration has been used to deal with inhomogeneity of this kind: in this procedure, one subtracts the average of the values present in a square region around a given pixel from the intensity of the pixel itself.[4] Although this effectively removes variation, there is an inherent tradeoff between suppressing the spatial variation effectively and avoiding the introduction of processing artifacts.

One common artifact introduced by this filtration is a bright border around the edge of the image (Figure 1). This occurs because near the edge the boxcar extends beyond the edge of the image; therefore, when the average of all of the points within the boxcar is taken, the value is less than elsewhere because the fraction of the boxcar that extends beyond the edge of the image

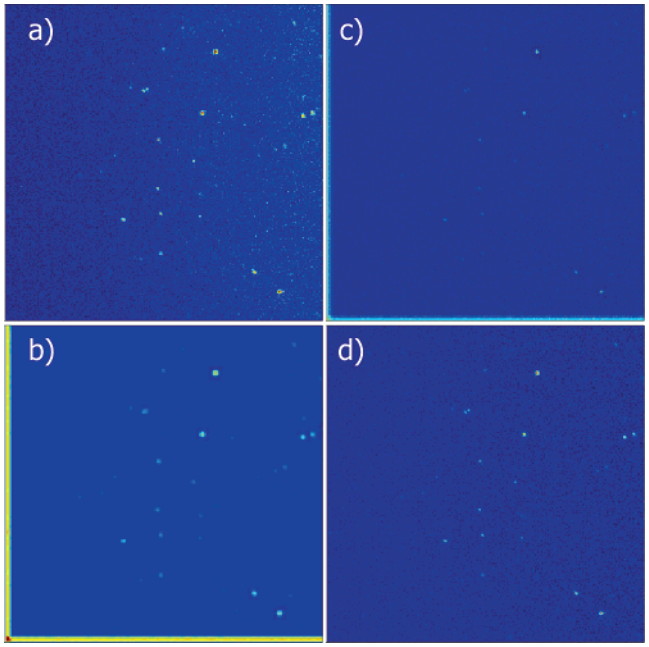

Figure 1. Comparison of various image-filtering methods, illustrated for dye-labeled lipids within DLPC-supported phospholipid bilayers as described in the text. Each image shows the bottom left corner of the same frame of the movie with an acquisition time of 30 ms. (a) The unfiltered image, showing a background gradient. (b) Traditional filtering using a Gaussian filter for the suppression of short-range noise and a boxcar filter for long-range noise.[4] (c) Boxcar filtering only. (d) Current filtering method: modified boxcar filtering with edge compensation and depression avoidance.

provides a weighting of zero. As a result, the quantity that is subtracted gradually decreases upon approaching the edge of the image, leaving a sloping raised border. Often, this region of the image is simply discarded. Alternatively, our implementation eliminates this problem by considering the number of pixels contained within the boxcar compared to the area of the boxcar and dividing by this factor. This causes the filter to subtract the average of the pixels present instead of the average of the entire boxcar region including missing pixels. As a result, although the value to be subtracted near the edge is the average of fewer pixels and hence is more variable and less precise, it does not decrease as the edge is approached, so it no longer creates a raised border around the image.

A more significant distortion introduced by the boxcar filter is localized around bright pixels or regions such as those corresponding to the molecules that one seeks to identify. This occurs because if localized regions are brighter than the average, then when the boxcar average contains that region, the value to be subtracted will incorrectly be larger than average. Consequently, each bright pixel leaves a square depression in the filtered image, the area of the depression equal to the area of the boxcar and the depth of the depression inversely proportional to the area of the boxcar. Each pixel of a bright region leaves its own depression as well, which will also appear to be roughly square, as long as the region is small compared to the size of the boxcar.

Because this depression becomes more pronounced the smaller the size of the boxcar, one may think that the simplest way to reduce the significance of this distortion is to increase the boxcar size. However, as the size of the boxcar increases, there is an accompanying increase in the length scale of the spatial inhomogeneity that can be suppressed by this method. Although the optimal balance between these opposing considerations depends on the nature of the spatial inhomogeneity, a boxcar

(15) Xie, A. F.; Yamada, R.; Gewirth, A. A.; Granick, S. *Phys. Rev. Lett.* **2002**, *89*, 24610.

(16) Lacey, A. J., Ed. *Light Microscopy in Biology: A Practical Approach*; Oxford University Press: Oxford, U.K., 1999.

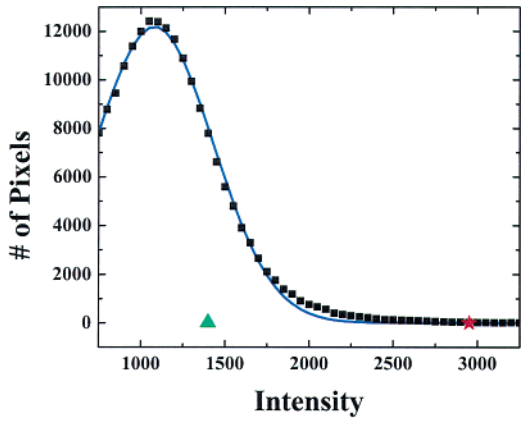

Figure 2. For dye-labeled lipids within the DLPC system described in the text, this histogram shows the number of pixels within a typical image that have a particular intensity. The blue line indicates the Gaussian fit to the histogram. The star indicates the calculated peak intensity threshold used to determine which local maxima in the image might be due to fluorescent molecules. On the basis of the Gaussian fit, random noise is unlikely to produce a pixel of this intensity more than once in this image. The triangle indicates the total intensity threshold. On the basis of the Gaussian fit, detector (CCD) noise is unlikely to produce spatial regions, the size of a typical diffraction-limited fluorescence spot, that possess this average intensity.

with each dimension roughly eight times the size of the diffraction-limited image of the single molecule was found to work best for our setup.

Even at this point, depressions were found to persist around all bright regions, especially around the extremely bright pixels occasionally caused by cosmic rays. A definitive solution to this problem consisted of setting a threshold criterion. When determining the mean value of the pixels within the boxcar, pixels above a threshold several standard deviations above the mean of the image were not used; in this situation, the threshold value was substituted instead. The difference between this value and the background was insufficient to create a noticeable artifactual depression.

ii. Locating Molecules. Before identifying a molecule in a digital image, it is necessary to decide on the parameters that distinguish a molecule from (unavoidable) background noise. Four criteria are useful to consider in making this evaluation: peak intensity, total intensity, radius of gyration, and eccentricity. The first two criteria, peak and total intensity, are the most important.

To avoid biasing the results with an arbitrary choice, it was convenient to determine threshold values automatically as follows, on an image-by-image basis. Figure 2 shows that when plotting the distribution of intensities of pixels in each image, an approximately Gaussian distribution was observed. This stemmed mainly from the noise distribution of the CCD camera. However, its asymmetry stemmed from the desired signal from the molecules. How should with this situation be handled? First, the profile of the CCD noise was determined by fitting the histogram to a Gaussian distribution. Having determined in this way the noise profile, it was possible to determine the probability that a pixel had a true intensity larger than could be attributed to CCD noise. The peak intensity threshold depicted in Figure 2 was then selected such that the probability was low that more than one pixel in the entire image possessed this intensity solely because of statistical noise. In the analysis of images, preliminary screening of each pixel consisted of determining whether the pixel was brighter than the minimum peak intensity calculated above, signifying that it was too bright to be CCD noise.

However, owing to diffraction, the Airy spot produced by a fluorescent molecule spans several pixels. With this in mind, we next set a criterion dictating that a region (adjoining pixels) of the image must possess a total intensity larger than a given threshold. Because calculating this probabilistically is more difficult than for the single-pixel criterion described in the previous paragraph, generally it was not feasible to calculate explicitly. We set a criterion for the total intensity by using judgment, with the intention of reducing the chance that an erroneous feature sneaks past the first criterion.

It was necessary not only to determine which regions of the image were brighter than the noise but also to localize their center. For this purpose, a corollary image was constructed whose local maxima represent the best matches for fluorescent signals in a given region of the image. This was accomplished relatively easily by convoluting the original, filtered image with a convolution kernel roughly approximating the expected shape of the feature being sought. Applied to the problem of locating a fluorescent dye, a Gaussian convolution kernel was employed as an appropriate approximation of the pixilated form of the expected Airy spot. Having localized the center of the molecule to within one pixel, the entire feature was considered. The size of the expected Airy spot was easily calculated. This region was examined. Its total brightness, radius of gyration, and eccentricity were determined.

The total brightness was used as a secondary screening criterion to prevent noise being falsely classified as a molecule. At present, we make little use of the final two criteria (radius of gyration and eccentricity) but include them for completeness.

As part of this procedure, features where the majority of the intensity comes from a single pixel are removed because these cannot correspond to molecules when one considers that the diffraction-limited spot size of a molecule ensures that it occupies several pixels. Finally, the location of the particle is further refined using a brightness-weighted centroid. Because the use of Gaussian convolution has already localized the center to within one pixel, iterated centroid determination was not employed. The signal-to-noise ratio heavily affects the accuracy of this calculation, as discussed in a related context by Cheezum et al.[13] Depending upon the system studied, the noise level may be too high, in which case accuracy is limited to the nearest pixel.

iii. Matching Features. What should be done when the paths of two molecules cross? Determining which features in images taken close together in time correspond to the same molecule is a significant challenge that presents inherent limitations regarding what kind of trajectories can be tracked. Although the exact implementation that we used differs from the reference that follows, an explanation of how to link locations into trajectories is covered well in a paper by Crocker and Grier.[4] Reasoning from the necessary condition that trajectories rarely intersect within a single time step, it is possible to reduce this problem to many smaller subnetworks, most of them with a unique solution. For the remaining subnetworks, a Gaussian probability distribution is assumed for the diffusion of the particles, and the most probable solution is chosen. This biases the solution toward minimizing the motion of the particles. Although usually accurate, in rare cases it will result in a misassignment.

At first glance, the assumption of a specific value for the Gaussian distribution of the particle motion may appear problematical when the objective is to determine the presence of Brownian motion and the diffusion constant. However, from repeated simulations with a wide range of both diffusion

coefficients and distribution types, we concluded that this is a weak assumption with almost negligible impact on the final conclusion.

Relative to traditional prior practice in this field[1−3] is the following: a significant modification was made to implementing the connection between brightness and feature identity. Traditionally, brightness was scaled appropriately and treated as an additional physical dimension. Although appropriate to 3-D tracking where the brightness is an indicator of *z*-axis height, this is less useful for tracking single molecules in two dimensions. Instead, the reasons that identical single molecules on a flat surface might differ in fluorescence intensity must be considered. One reason could be different alignment of the transition dipoles. Another consideration is that for many fluorescent dyes the rate of fluorescence is not constant, the molecule switching instead ("blinking") between various emissive and nonemissive modes.[17] The 30 ms time resolution of our data acquisition apparatus is shorter than the lifetime of some of these modes, resulting in variations in molecule intensity between different modes. Therefore, a partial correlation existed between the intensity of a single molecule in one frame and the next.

Before analyzing brightness data in this context, it was corrected for the intensity of the background, as described above. This criterion was used not only as a weak distinguishing characteristic between molecules but also to suppress noise further. By considering the relative brightness when evaluating possible matches, the partial correlation of a particle's brightness between frames improves the tracking. This frequently allowed the correct trajectory to be determined when two otherwise indistinguishable molecules move past each other. When implementing this in a computer program, the weighting of the ratio of brightness was made to be an adjustable parameter to allow for varying degrees of correlation depending upon the relevant experimental parameters.

The absolute brightness of each feature was also considered; this both further suppressed noise and further improved trajectory assignments. Although weakly emissive and nonemissive fluorescent states are not uncommon, going from being one of the brightest molecules in one image to zero emission in the next would require an improbable transition between states at exactly the wrong moment. It is reasonable to conclude that the probability of observing the brightest features of an image also in the next temporal image is greater than the probability of not observing them. Similarly, the probability is low that one of the brightest features is in fact noise. These differences were incorporated into the matching probability discussed earlier, again with an adjustable coefficient to vary the weighting depending upon the strength of the correlation.

One consequence of the frequency with which single fluorescent molecules blink is that their overall trajectories are fragmented into shorter visible segments. To overcome this, the algorithm was modified to allow the reconstruction of trajectories even in the presence of an intervening hole, as illustrated in Figure 3. The main idea is to supplement the trajectory information determined from the actual time sequence of images with trajectory information from the same movie downsampled to half the time resolution.

Specifically, after tracking all adjacent frames of the movie, alternating frames of the movie were tracked similarly; this gave two additional sets of trajectories—one containing the odd frames and the other containing the even frames. When a sequence in the ordinary tracking terminated, the additional trajectories were consulted. Frequently, these trajectories connected to another trajectory from the ordinary tracking, in which case the possibility that the two trajectories from the sequential tracking represented the same molecule was considered. When attempting to join two sequences together, two points of overlap on either side of the break to be bridged are required.

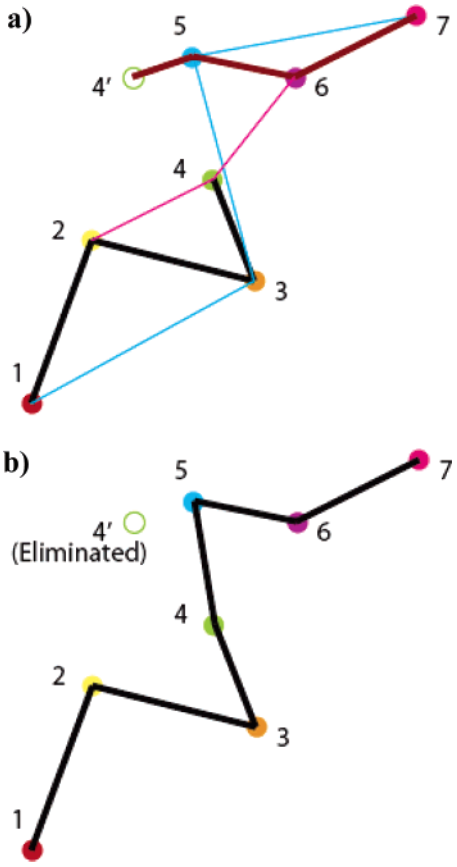

Figure 3. (a) Although a molecule may move sequentially from point 1 to point 7, reconstructing that motion from a movie can be complicated. Complications can include failure to detect the molecule in a particular frame, the trajectories of two molecules crossing each other, or interference from noise, as simulated here. Traditional tracking would find two segments, one following the molecule from point 1 to point 4 (black) and a second following the molecule from point 4′ (mistakenly) to point 7 (brown). This can easily occur if the noise present at 4′ provides a better match with point 5 than the real molecule at point 4 does. Tracking alternating frames, as described in the text, results in the trajectories containing the odd points (blue) and the even points (pink). (b) These trajectories connect the two single trajectories sufficiently to allow reconstruction of the proper time sequence.

In this analysis, the terminating steps of any trajectory were less reliable than the middle steps. The inability to continue tracking was a good indicator that some kind of an abnormality was present, either a misidentification or a vanished molecule. Therefore, if the trajectories were joined and there was a point in dispute, then the dispute was resolved in favor of the trajectory that contained the point in the middle. This represents a significant improvement over the practice of retaining the location of lost trajectories and attempting to find a match for them at a later point in time. By using multiple points of overlap, mismatches were minimized, and accuracy was maintained because (as already mentioned) the alternating frame tracking is equivalent to the regular tracking of a movie with half the time resolution. As a result, the only drawback is that this cut the effective time resolution in half for the purposes of determining which movies could be successfully tracked. Generally, this was not an issue.

(17) Göhde, W.; Fischer, U. C.; Fuchs, H.; Tittel, J.; Basché, Th.; Bräuchle, Ch.; Herrmann, A.; Müllen, K. *J. Phys. Chem. A* **1998**, *102*, 9109.

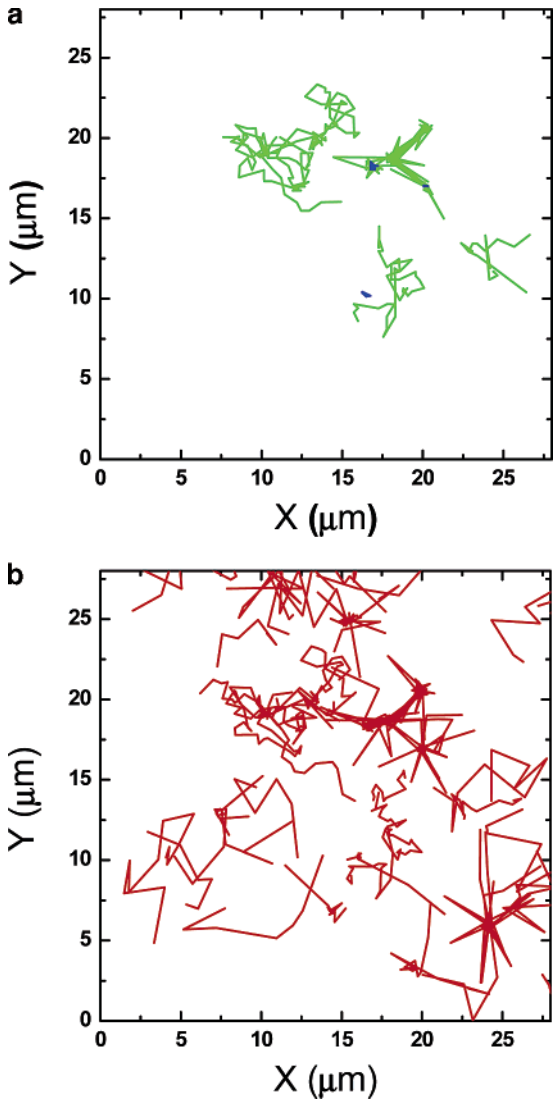

Figure 4. For dye-labeled lipids within the DLPC system described in the text, panel a shows trajectories at least five frames long found using algorithms described in the text. Among these, the trajectories in blue (up to 30 frames in length) are examples of trajectories that exhibited no motion beyond that which could be attributed to uncertainty in locating the particles and therefore were discarded. The green lines illustrate trajectories used to determine the computed diffusion constant, but it was not very simple. Although most of these trajectories appeared to exhibit random walk behavior, some were similar to the trajectories in blue in the sense that they were stuck to a fixed position or else they were stuck to a few fixed positions, occasionally jumping between them. Experience showed that although these would have profoundly reduced the inferred mean squared displacement if included in the final data analysis, when the distribution of displacements was plotted, their motion was distinctly separate from the real motion. (b) The same data analyzed using a standard particle-tracking algorithm.[4] The trajectories in panel a are partially reproduced, but close examination reveals that the equivalent trajectories were frequently either shorter than in panel a or else were composed of multiple fragments and also that they exhibited a greater uncertainty in position. Most important in panel b is the appearance of a large number of new trajectories in which the standard particle-tracking routine mistakes noise for signal (i.e., actual molecules). The main point of panel b is to show that the real single-molecule trajectories were overwhelmed by trajectories that in fact reflected noise in the data, thus confounding the determination of the real translational diffusion coefficient.

Although the main advantage of this procedure was to extend the length of the time sequences that could be tracked successfully,

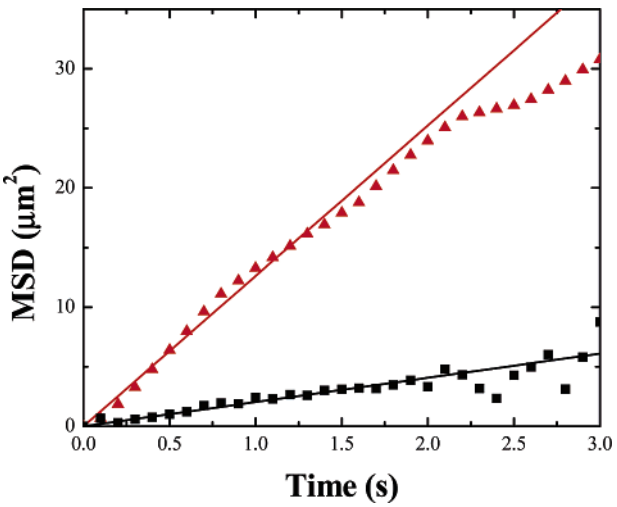

Figure 5. For dye-labeled lipids within the DLPC system described in the text, the mean squared displacement is plotted vs time for all trajectories at least 15 segments long. (Lower curve) Linearity shows Brownian motion. A vertical offset, 0.62 μm^2, attributable to the random error in locating the particles, was subtracted. The number of trajectories contributing to this calculation was 35 for the first 15 data points and decreased for longer times. The diffusion constant calculated from the slope of this plot is artificially low. Figure 6 explains why: both the mobile and immobile populations were mixed together indiscriminately. Using an alternative method described in the text, for one of the longest trajectories that we succeeded in tracking (76 successive frames, 100 ms per frame), the mean squared displacement was plotted against elapsed time. (Upper curve) A vertical offset, 0.20 μm^2, attributable to the random error in locating the particles, was subtracted. Note that because frames in this trajectory were better localized and brighter than for the average trajectory the random error in determining the particle position was less for this plot. One observes linearity for the first 3 s. A reasonable fit (shown) implies $D = 3.2\ \mu m^2/s$, which agrees well with the FCS finding that $D = 2.6\ \mu m^2/s$.[19]

it was also useful in error correction and in dealing with the inaccuracy introduced by having edges to the image.

Discrepancies occurred occasionally. The single frame and one of the alternating frame sequences agreed about earlier and later points but disagreed about the middle point in a sequence. Upon examination, often the point from the alternating frame sequence was found to be more probable because it minimized the motion of the particle. Therefore, it was substituted back in. Additionally, it is self-evident that tracking accuracy is inherently poorer near the edge of an image because observing molecules that diffuse into or out of this image is technically impossible. Therefore, we find it convenient to evaluate all molecules within one diffusion step of the image edge. If the preceding and following points were tracked successfully using alternating step tracking and lay far enough from the edge of the image, then (bearing in mind that the presence of the edge did not influence the trajectory) these trajectories were maintained. If not, then the trajectory was fragmented into two segments.

Although this would occasionally split legitimate trajectories, it eliminated the artifactual difference in accuracy between trajectories that were near the edge of an image and those that are not. In turn this made it easier to determine the fundamental characteristics of the motion and to improve the tracking accuracy.

iv. Analyzing Trajectories. At this point, the trajectories of the molecules were known, but the list of trajectories included trajectories or segments caused artifactually by noise, in addition to the true trajectories. To eliminate these, the first step was recognition that seeming trajectories owing to random noise generally contained fewer segments than true trajectories that

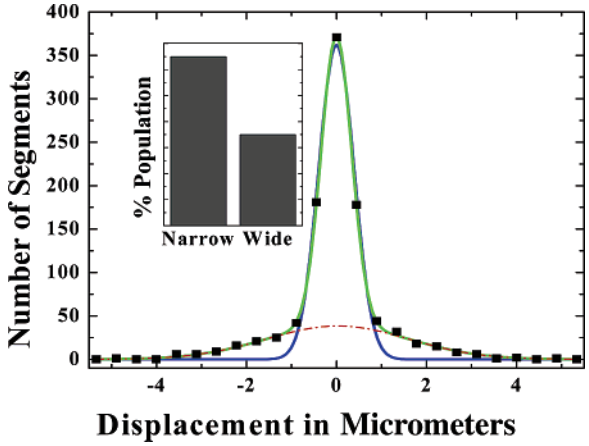

Figure 6. Histogram of the translational displacement of the molecules for a four-frame interval. Attempts to fit this histogram with a single Gaussian (blue) and two Gaussians (green) reveal the presence of two subpopulations. (The red line shows just the wide component of the two-Gaussian fit.) The width of the narrow Gaussian is approximately equal to the uncertainty in position. The width of the wider Gaussian, when adjusted for the time interval, corresponds to a diffusion rate of $D = 2.5\ \mu m^2/s$ using these parameters. By varying computational parameters, as shown in Figure 7, the Fickian diffusion coefficient was determined to be $D = 2.6 \pm 0.4\ \mu m^2/s$. This agrees with the diffusion constant of the mobile molecules determined by fluorescence correlation spectroscopy, FCS ($2.6\ \mu m^2/s$).[19] (Inset) A heuristic histogram showing that 65% of the statistics summarized here reflected unphysically small steps attributable to various artifacts and that the physically meaningful (larger) displacements comprised only 35% of the statistical sample.

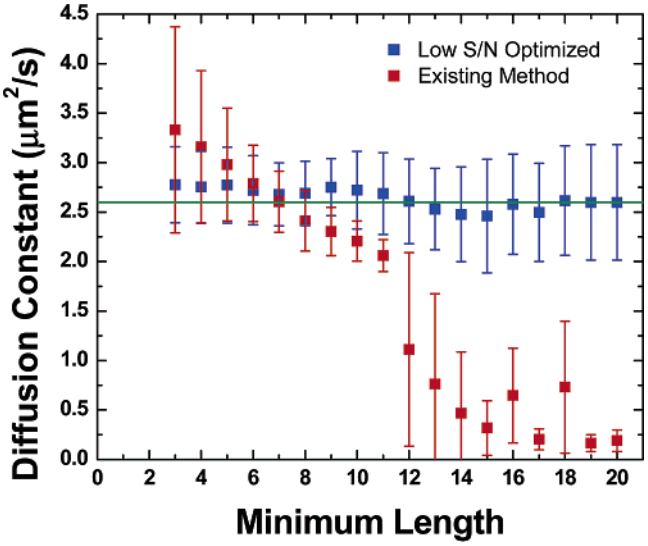

Figure 7. Comparison of the translational diffusion constant determined using algorithms described in this article (blue) and algorithms in the literature[4] (red) while varying two computational parameters: the threshold length below which trajectories were discarded (*x* axis) and the number of successive frames used to compute the translational displacement trajectories. Each data point represents the mean value for frame intervals from three to six frames, and the error bars represent the standard deviation. This range was found to be the optimal range for these data based upon the quality of the fitting. When we sought to fit shorter intervals, the combination of noise and smaller difference between components in the two-Gaussian fit complicated matters to the point that unequivocal fits could not be obtained. However, though longer trajectories would in principle be desirable, too few longer trajectories could be acquired to generate reliable statistics. In this plot, one notices that the choice of minimum trajectory length was not critical for the algorithms described in this article, which agreed well with the FCS value (green line) regardless of how this parameter was chosen. However, the application of literature algorithms[4] to this situation of low signal-to-noise ratio gave problematical results. When the shortest trajectories were included, the distribution of displacements was dominated by random noise, yielding an erroneously large computed diffusion coefficient. When only the longer trajectories were included, standard algorithms[4] did not succeed in following enough of them, resulting in so few long trajectories that the statistics of their analysis was not robust.

followed true molecules. This was easily understood when one considered that if the trajectory was true and the dye did not blink then there should be a feature to match up to in the next frame. For noise, it is not so. Therefore, the majority of these artifactual, seeming trajectories could be excluded simply by discarding the shortest trajectories.

Testing this idea, one finds that the statistical result depends strongly on the trajectory length. Removing the very shortest time trajectories from consideration significantly suppressed the number of inferred long spatial trajectories, but no significant further change was found when medium-length trajectories were discarded. The residual differences were obvious by considering the enhanced chance of a match owing to artifactual noise when the displacement amplitude is high. In our implementation of this idea, all trajectories of fewer than 10 time segments were discarded. This arbitrary value produced a stable histogram of findings.

Noise might be manifested as seemingly complete trajectories and might also degrade accuracy in determining the true trajectories of the tracked molecules. For the reasons already discussed, the ends of trajectories are the least accurate. Therefore, after the trajectories were otherwise finalized, the first and last segments of each trajectory were discarded. Although this reduced somewhat the number of largest-displacement segments, because further reduction was not observed when this algorithm was repeated this was a strong indication that this was due to a reduction of noise. Because the beginnings and ends of trajectories were distributed evenly through all the frames of a time sequence and other segments of the trajectory should not depend on the discarded segments, discarding those segments should not introduce bias.

At this point, the motion of the molecules was recorded. It was used in a variety of ways depending upon the purpose of the experiment. Frequently, the mean squared displacement of the system was of interest, and the diffusion constant was determined from the slope of the relevant plot. For a system exhibiting a random error in determining the position of the molecules, the equation of the plot for Brownian motion in two dimensions is[18]

$$\mathrm{MSD}(\Delta t) = 4D\Delta t + 2\sigma^2 \qquad (1)$$

and a linear dependence of the mean squared displacement (MSD) on time elapsed (Δt) indicates Brownian diffusion. Here, D is the translational diffusion coefficient, and σ^2 is the variance of the experimental noise.

This analysis is insensitive to the possible presence of subpopulations whose D differs from the majority component. Rather than plot the mean squared displacement according to eq 1, the distribution of motion in single steps can be plotted. This method of analysis is insensitive to the correlation of motion between time steps but readily reveals the presence of multiple diffusion rates or mechanisms. Therefore, these two methods of analysis are complementary.

(18) Dietrich, C.; Yang, B.; Fujiwara, T.; Kusumi, A.; Jacobson, K. *Biophys. J.* **2002**, *82*, 27.

IV. Data and Results

To test the efficacy of these algorithms and to illustrate their operation in practice, they were applied to analyze images taken for supported phospholipid bilayers of DLPC containing fluorescent-labeled lipids as described above. This research group had already studied this same system using fluorescence correlation spectroscopy (FCS).[19]

First, we analyzed a single (relatively long) trajectory (7.6 s long, 100 ms per frame). Because the analysis of a single random walk trajectory is not standard, it is appropriate to explain the method. The main idea is to consider the trajectory to be an assembly of steps: 76 segments (in this case) composed of 1 unit step length, 75 segments composed of 2 step lengths, 74 segments composed of 3 step lengths, and so on. Each step length corresponds to the time interval of a single data acquisition frame (100 ms in this example). In the top curve of Figure 5, the mean-squared displacement is plotted against elapsed time. One sees that the data are linear for the initial 2 s (i.e., before the statistics became too few to be reliable because the number of segments with >20 step lengths was too few to be reliable). The slope corresponds to $D = 3.2\ \mu m^2/s$, which agrees well with the FCS finding that $D = 2.6\ \mu m^2/s$.[19]

However, a blind implementation of single-molecule tracking proved to be misleading. Considering all data collectively, the mean squared displacement is plotted against time in the bottom curve of Figure 5. Upon first observing this, we were shocked that the slope implied a diffusion coefficient a factor of 5 less than that deduced from the top curve of Figure 5 and the FCS data.[19] Why? The explanation emerged from considering a histogram of the displacement steps, plotted in Figure 6. Here one observes two distinct subpopulations. One subpopulation consisted of relatively large displacements; another consisted of miniscule displacement from the original position, not larger than the statistical uncertainty. However, because of some trajectories having sections from each subpopulation, a separation of trajectories based upon the subpopulation was not feasible. The second subpopulation probably stems in part from fluorescent dyes that were nearly immobilized by the surface, in part from surface defects and finally in part from fixed noise sources (scattering points, fluorescence impurities in the substrate, and defects in the optics or the detector). We tested this interpretation extensively. Detailed analysis (not shown) revealed that displaying displacement over a single-frame interval, instead of the four-frame interval plotted in Figure 6, approximately halved the width of the wider Gaussian, just as should be expected if the width corresponded to Brownian diffusion. At the same time, it did not change the narrow center subpopulation, just as expected if this subpopulation were artifactual for the reasons just discussed.

This illustrates a distinct advantage of single-molecule tracking: the capacity to distinguish distinct subpopulations whose mobilities differ. In implementing this technique, one must be fully aware of the disproportionately large contribution taken by trajectories that are exceptionally bright and for this reason contribute exceptionally strongly to the raw statistics, but for the trivial reason that they are nearly stationary. This near-static population should be ignored and was ignored for our determination of the diffusion constant.

Varying the computational parameters used for analysis highlighted the need to modify the standard particle-tracking algorithms[4] before working with systems that possess lower signal-to-noise ratios. As seen in Figure 7, our algorithms uniformly gave results in close agreement with the FCS finding of $D = 2.6\ \mu m^2/s$[19] for a wide range of parameters. The overall result obtained using our algorithms was $D = 2.6 \pm 0.4\ \mu m^2/s$.

V. Summary

Modifications to standard particle-tracking methods[4] allow their extension to the lower signal-to-noise regime that holds when one seeks to track the trajectories of single molecules. Although some of the optimizations proposed above are based upon the unique properties of single-molecule fluorescence, most dwell on reducing the impact of noise upon the final set of trajectories and thus may also be relevant more broadly.

Acknowledgment. This work was supported by the U.S. Department of Energy, Division of Materials Science, under award no. DEFG02-02ER46019. S.A. acknowledges the NSF for financial support in the form of a graduate research fellowship.

LA060244I

(19) Zhang, L.; Granick, S. *Proc. Natl. Acad. Sci. U.S.A.* **2005**, *102*, 9118.