1-1-2007

Multiple Particle Tracking and Two-Point Microrheology in Cells

John C. Crocker
University of Pennsylvania, jcrocker@seas.upenn.edu

Brenton D. Hoffman
University of Pennsylvania

Postprint version. Published in *Methods in Cell Biology*, Volume 83, 2007, pages 141-178.
Publisher URL: http://dx.doi.org/10.1016/S0091-679X(07)83007-X

Multiple Particle Tracking and Two-Point Microrheology in Cells

Abstract
Mechanical stress and stiffness are increasingly recognized to play important roles in numerous cell biological processes, notably cell differentiation and tissue morphogenesis. Little definite is known, however, about how stress propagates through different cell structures or how it is converted to biochemical signals via mechanotransduction, due in large part to the difficulty of interpreting many cell mechanics experiments. A newly developed technique, two-point microrheology (TPM), can provide highly interpretable, quantitative measurements of cells' frequency-dependent shear moduli and spectra of their fluctuating intracellular stresses. TPM is a non-invasive method based on measuring the Brownian motion of large numbers of intracellular particles using multiple particle tracking. While requiring only hardware available in many cell biology laboratories–a phase microscope and digital video camera, as a statistical technique, it also requires the automated analysis of many thousands of micrographs. Here we describe in detail the algorithms and software tools used for such large-scale multiple particle tracking, as well as common sources of error and the microscopy methods needed to minimize them. Moreover, we describe the physical principles behind TPM and other passive microrheology methods, their limitations, and typical results for cultured epithelial cells.

Disciplines
Cell Biology

Comments
Postprint version. Published in *Methods in Cell Biology*, Volume 83, 2007, pages 141-178.
Publisher URL: http://dx.doi.org/10.1016/S0091-679X(07)83007-X

Multiple Particle Tracking and Two-Point Microrheology in Cells

John C. Crocker and Brenton D. Hoffman
Department of Chemical and Biomolecular Engineering
Institute for Medicine and Engineering
University of Pennsylvania
Philadelphia, PA 19104

ABSTRACT

Mechanical stress and stiffness are increasingly recognized to play important roles in numerous cell biological processes, notably cell differentiation and tissue morphogenesis. Little definite is known, however, about how stress propagates through different cell structures or how it is converted to biochemical signals via mechanotransduction, due in large part to the difficulty of interpreting many cell mechanics experiments. A newly developed technique, two-point microrheology (TPM), can provide highly interpretable, quantitative measurements of cells' frequency-dependent shear moduli and spectra of their fluctuating intracellular stresses. TPM is a non-invasive method based on measuring the Brownian motion of large numbers of intracellular particles using multiple particle tracking. While requiring only hardware available in many cell biology laboratories–a phase microscope and digital video camera, as a statistical technique, it also requires the automated analysis of many thousands of micrographs. Here we describe in detail the algorithms and software tools used for such large-scale multiple particle tracking, as well as common sources of error and the microscopy methods needed to minimize them. Moreover, we describe the physical principles behind TPM and other passive microrheology methods, their limitations, and typical results for cultured epithelial cells.

I. Introduction

Cell biologists have long studied the complicated biochemical and physiological responses of cells to mechanical stress or deformation (Orr, 2006; Vogel, 2006). In addition to these responses, cells also show a purely mechanical, deformation response to applied stress, determined by their shear modulus (see the Chapter by Janmey). While mechanical and physiological responses can occur simultaneously and couple, complicating interpretation, it is usually assumed that the deformations occurring immediately after stress application or in response to small stresses are predominantly mechanical in origin.

Only in the last decade or so have techniques, termed *microrheology*, been developed that can characterize cells' dynamic shear modulus over a wide frequency range (Weihs, 2006; Waigh, 2005). Initially, many researchers hoped that these new cell rheology

measurements would display characteristic times corresponding to known molecular timescales, such as that of the myosin ATP hydrolysis cycle. Such contributions could then be dissected with genetic or pharmacological interventions to tie together cell behavior at the molecular and mesoscopic scales. Less ambitiously, it was hoped that the cell mechanical response would at least closely resemble the response of various purified biopolymer gel models (typically F-actin), whose rheology has been intensively studied since the 1980's (Kroy, 2006). In reality, however, the rheological responses of cells measured to date have not contained identifiable molecular timescales and have not been satisfactorily reproduced by any biopolymer gel model yet studied. Several recent experiments have found the (low frequency) shear modulus of cells to be well described by a weak power-law form, $G^*(\omega) \sim \omega^{\beta}$, with reported values of β varying over the range 0.1-0.3 (Yamada, 2000; Fabry, 2001; Alcaraz, 2003; Lenormand, 2004; Desprat, 2005; Hoffman, 2006). Such a power-law form has no characteristic times at all. In contrast, the low frequency behavior of reconstituted gels is either purely elastic with $\beta = 0$ (Janmey, 1990) or nearly so with a very small β exponent (Xu, 1998; Gardel, 2006).

The challenges of interpreting and modeling cell rheology measurements have several causes (Weihs, 2006). Primary among them is the obvious structural complexity within cells. Cells are composed of numerous chemically and spatially distinct sub-domains that include the cell cortex, the nuclear envelope, lamellipodia and stress fibers, not to mention the microtubule and intermediate filaments networks, endoplasmic reticula and other organelles that fill the cell interior. It seems likely that different cell microrheology methods will probe different mechanical sub-domains (or different combinations of them), and that the target and response might differ among cell types, or even among individual cells of the same type. Indeed, it is not clear *a priori* to what extent cells may be understood as a continuous viscoelastic solid, rather than a complex ensemble of discrete units. Lastly, there is also the fact that microrheology comprises a new and emerging set of methodologies, with still unresolved technical issues regarding interpretation and measurement artifacts, even with comparatively simple synthetic or reconstituted biopolymer materials.

In this chapter, we will describe a cell microrheology method developed in our laboratory, two-point microrheology (TPM) (Crocker, 2000; Lau, 2003), which computes the rheology from measurements of the statistically cross-correlated Brownian motion of pairs of embedded, intracellular tracers (for complementary discussion of particle-based microrheology see the Chapter by Wirtz). Compared to other techniques, TPM has the advantage of being more interpretable: it is essentially immune to uncertainties related to cytoskeletal heterogeneity and the tracer/network connection, and provides additional positive controls regarding whether the cell even behaves as a viscoelastic continuum. TPM has the disadvantage that, like all methods based on Brownian motion, it can be confounded by active intracellular processes, and is thus most reliable when applied to cells that have been depleted of ATP. Moreover, since TPM naturally probes a three-dimensional structure much larger than the tracers, it is best suited to measure the rheology of the thick central 'body' of the cell, rather than thin structures such as the lamellipodium, cell cortex or nuclear envelope.

The differences among the mechanical properties of different cell regions were highlighted in a recent study in our laboratory (Hoffman, 2006), which compared the results of TPM with another technique, Magnetic Twisting Cytometry (MTC; see thc Chapter by Ingber), based on rocking external, integrin-attached microparticles using a magnetic field. While qualitatively similar at first glance (Figure 1), the results of these two measurements on the same cell type are distinctly different upon close inspection. Combined with literature and control measurements, these results indicate that the mechanical properties of cells' cortical and deeper intracellular regions are different, and thus presumably so are their predominant structural elements and organization. While this difference is in line with expectations from known cell physiology, it clearly represents a potential confounding factor for all cell rheological measurements; a given rheology method may potentially probe one or the other region, or a superposition of both. The unique interpretability of TPM was essential to sorting out this confounding factor.

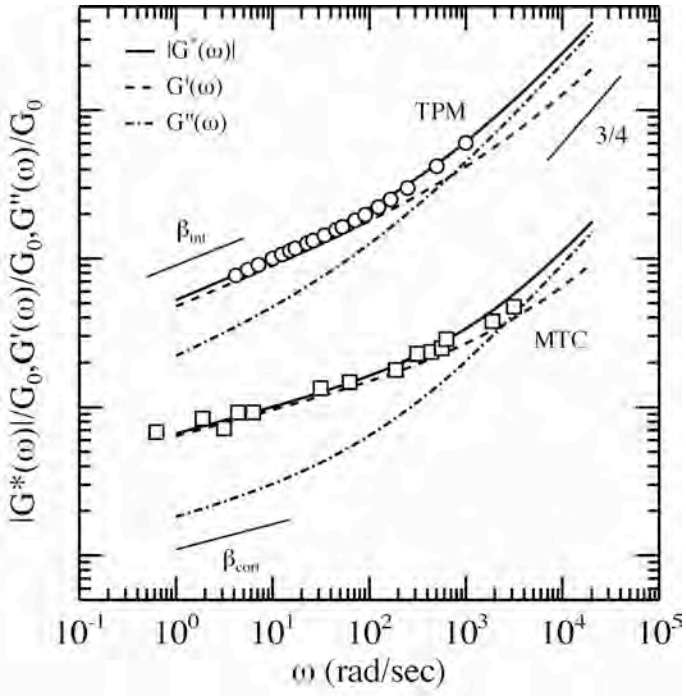

Figure 1. Comparison of the cell rheology data from two different methods applied to TC7 epithelial cells, probing the deep interior, TPM (circles), averaged over N=8 cells, or the cortex, MTC (squares), N=4 cells. Curves are best fits of a double power-law form to the data shown and other higher frequency measurements, not shown (Hoffman 2006). Curves indicate only the frequency dependence; the vertical positions of the two datasets are arbitrary. Absolute stiffness estimates are discussed in Section VII.B.

Examining the cell rheology findings in Figure 1, both the cortical and interior response can be divided into two frequency regimes. At low frequencies, the structures are predominantly elastic ($G'>G''$), with the aforementioned power-law form reported by several groups, $G^*(\omega)\sim\omega^{\beta}$, where many literature values of β fall between the two values we find ($\beta_{int}=0.26$, $\beta_{cort}=0.16$). At high frequencies, both structures are more dissipative than elastic ($G''>G'$), with a power-law exponent of about 3/4. Such high frequency behavior has a simple physical explanation—a 3/4 exponent is characteristic of the high-frequency rheology of networks formed of semi-flexible (i.e. filamentary) polymers (Deng, 2006). While this clearly suggests that deformations at high frequencies are resisted by a filamentous network, it is an open question whether low frequency deformations are resisted by the same biopolymer network or by a different structure

altogether. Indeed, the physical origin and molecular determinants of cells' power-law rheology have not been compellingly identified; clearly much research remains to be done. Given cells' structural complexity, it seems likely that no single cell microrheology method can provide a complete or compelling description of cellular responses to mechanical stress. Rather, a sensible strategy may be to apply multiple microrheology techniques to single cell types, with TPM contributing one particularly interpretable component, and to compare the results carefully for mutual consistency.

From a hardware point of view, TPM is accessible as it requires only hardware that is common in many cell biology laboratories: a high-magnification optical microscope with Phase Contrast or Differential Interference Contrast optics and a high intensity illuminator, as well as a reasonably high quality, low noise camera that can collect images at or near video rate (tens of frames per second). While we use a specialized high-speed camera to collect the above reported TPM data, this is not required to measure rheology in the weak power-law regime. Unlike laser-deflection based approaches, expertise with physical optics, analog signal processing or lasers is not required. As a statistical method, however, it does require the analysis of extremely large numbers of micrographs, at least 10,000 per single cell measurement. With the use of automated image acquisition and analysis routines and increasingly fast microcomputers, this requirement presents little real impediment. Because of the small amplitude of the Brownian motion of intracellular particles, however, care must be taken to minimize the effects of microscope vibration, and to maximize the precision of the particle tracking process. Beyond describing the underlying algorithmic and mathematical procedures required for TPM, this chapter will also discuss how to achieve high performance multiple particle tracking using an imaging system. Such an instrument is potentially useful for a number of cell biology applications other than TPM, e.g. for studies of endocytosis and intracellular trafficking.

II. Principles of passive tracer microrheology

In general, there are two approaches to measuring the rheology of soft materials, active and passive. In the active approach, a known force (or deformation) is applied to the material, and the resulting deformation (or force) is measured. In the cell rheology context, this usually relies upon atomic force microscope (AFM) (Alcaraz, 2003) or similar, force calibrated micro-cantilever instruments (Desprat, 2005). The passive approach, which is our focus in this chapter, examines the Brownian motion of tracers embedded in the soft material. No force is applied at all, only the spontaneous motion of the tracers is observed and quantified, considerably simplifying instrumentation requirements. The 'calibration' of the Brownian forces comes from simple physical principles: in general, the (squared) amplitude of the tracer motion is inversely proportional to the material's stiffness.

II.A Conventional passive microrheology

The simplest and first described example of Brownian motion is a spherical tracer

particle moving in a simple viscous fluid (like water or glycerol). Here the relationship between mechanical properties and tracer motion is the familiar Stokes-Einstein relation:

$$\eta = \frac{k_B T}{\left\langle \Delta r(\tau)^2 \right\rangle \pi a} \tau, \tag{1}$$

where η is the liquid's viscosity, a is the tracer radius and the 'driving force' is the energy of thermal fluctuations, $k_B T$, where k_B is Boltzman's constant and T the absolute temperature. The mean squared displacement (MSD) of the particle's motion, $\left\langle \Delta r(\tau)^2 \right\rangle$, is simply the square of the net distance the tracer typically moves during a given time interval, τ, in this context called a lag time. We will discuss the computation of the MSD in a later section.

To model the Brownian motion of tracers embedded in viscoelastic materials, Eq. 1 needs to be modified. Several different, but mathematically consistent versions are in use; we employ the Generalized Stokes-Einstein Relation (GSER) (Mason, 1995):

$$G^*(\omega) = \frac{k_B T}{i\omega \left\langle \Delta r(\omega)^2 \right\rangle \pi a}, \tag{2}$$

where $\left\langle \Delta r(\omega)^2 \right\rangle$ is the unilateral Fourier Transform, $f(\omega) = \int_0^\infty e^{-i\omega\tau} f(\tau) d\tau$ of the MSD and $i = \sqrt{-1}$, and the use of (Fourier) frequency, ω, rather than lag time facilitates comparison to conventional rheology models. Note that in order to determine the rheology at even a single frequency, we must formally evaluate the Fourier integral over all lag times from τ=0 to τ=∞, while it is obviously impossible to measure the MSD over that entire range. In practice, if we measure the MSD over a wide range of lag times, $\tau_1 < \tau < \tau_2$ then we can compute the rheology over a range of frequencies $\omega_2 = 1/\tau_2 < \omega < \omega_1 = 1/\tau_1$, with a little uncertainty at the frequency extrema. Numerically evaluating Fourier integrals with typical data can be quite challenging, but simple approximate methods have been developed, which we will describe in a Section IV.

Equation 2 implies several relationships between the Brownian motion and the shear modulus. Since the overall amplitude of the MSD and its Fourier transform are linearly proportional, the GSER implies an inverse relationship between mechanical properties and the tracer's MSD versus lag time. In addition, Since $\left\langle \Delta r(\omega)^2 \right\rangle$ is in general a complex function, so is $G^*(\omega)$. The real and imaginary parts of the shear modulus, $G'(\omega)$ and $G''(\omega)$, respectively called the storage modulus and loss modulus (see the Chapter by Janmey), represent the solid-like and liquid-like behavior of the material. In simple viscous liquids, the above equation reduces to Eq. 1, with MSD ~ τ^1 and $G^*(\omega) = i\omega\eta$. In simple elastic solids, the MSD is lag time independent, MSD ~ τ^0. In a viscoelastic material, the MSD will have an intermediate form, and will increase more slowly than linearly with lag time.

An important point is that the derivation of Eq. 2 assumed that the tracer is embedded in a completely homogeneous material, and has 'no slip' boundary conditions. In many complex synthetic materials, the rheology computed using the GSER has been found to

be incorrect due to a failure of one or both of these assumptions (Crocker, 2000; Valentine, 2004). In porous materials, tracers tend to diffuse inside fluid filled pores, and systematically report a softer and more fluid-like rheology. Tracers that associate strongly with the material, e.g. adhering strongly to segments of a porous material, tend to report rheology that is softer than the true bulk rheology but has the correct frequency dependence (Van Citters, 2006). Tracers in cells tend to show this latter behavior, implying that many microrheology techniques tend to yield reliable frequency dependences but may have systematically underestimated absolute stiffness values.

II.B Expected tracer motion and tracking performance

The expected amplitude of the Brownian MSD for some simple cases can be readily estimated. Consider spherical particles 1 μm in diameter, which are easily visible under a high magnification optical microscope. In water at room temperature, such particles will typically move 1 μm in each direction in a lag time of τ=1 second. In an elastic material with a shear modulus of 1000 Pascal = 1kPa, a typical value reported for cell measurements (Fabry, 2001) and roughly that of very soft agar, the rattling motion of a 1 μm tracer has an expected amplitude of only 1 nm. Such a miniscule motion is quite invisible to all but the highest-performance particle tracking instruments.
For comparison, consider the typical amplitudes of vibration and tracer position measurement error. For an optical microscope simply placed on an ordinary lab bench or table, typical sample vibration amplitudes are 10-100 nm. The same microscope on a well engineered pneumatic vibration isolation platform or table will have vibration amplitude typically about ten times smaller, but still comparable to or in excess of the expected Brownian motion signal, 1 nm, estimated above. Image based particle tracking routines that have not been carefully optimized typically locate micron-sized particles to 10-20 nm precision, (Crocker, 2000), and this can be improved to a few nanometers using the techniques described in later sections. While we find that intracellular tracers typically move somewhat more than the 1 nm estimate above, passive cell microrheology is nonetheless challenging for even a high-performance image-based particle tracking system. This is the reason why many such experiments are performed with laser-deflection particle tracking systems, which readily achieve sub-nanometer precision, or with longer lag times, where the motion is somewhat larger.

II.C Two-point microrheology

We concluded above that interpreting cell rheology measurements is often difficult due to uncertainties related to tracer boundary conditions and tracer/network association. Fortunately, two-point microrheology removes most of the interpretation uncertainties of conventional microrheology. Unlike other microrheology methods, TPM can deliver a reliable, absolute measure of stiffness, rather than just its frequency-dependence, even in heterogeneous materials and cases where the tracer boundary conditions are not known.

The basic principle of two-point microrheology is that all soft materials undergo a form of internal Brownian motion, like waves on the surface of a choppy ocean. Tracers are carried along by these random undulations of the medium, like corks bobbing on the

ocean. The Brownian motion of two separate tracers will be statistically correlated, because they will both be carried along by the Brownian motion of the segment of the material spanning between them. Two-point microrheology computes the rheology from the amplitude of the tracers' resulting cross-correlated motion. Elementary calculations show that in three dimensions, the correlation between the particles motion is inversely proportional to their separation, R. This is a consequence of larger material segments having smaller Brownian motion amplitudes (just as larger particles diffuse more slowly than small ones). This $1/R$ dependence provides a useful positive control that the material is deforming like a continuous three-dimensional object, a prerequisite for our analytical framework (Crocker 2000).

The interpretive power of two-point microrheology comes from the fact that any local motion of the tracers relative to the material generally will not be correlated with one another (Levine 2001). For example, in the troublesome case that tracers inhabit soft pores in the material, their diffusive rattling in their respective pores will be statistically independent, and will not contribute to the correlation measurement in two-particle microrheology. Details of how to compute the cross-correlation and rheology will be given later. While sample vibration does lead to correlated motion, it has a different character than that due to the material's Brownian motion, and can simply be filtered out to a large extent. It turns out that tracer position error increases the noise of the two-point measurement, and must still be minimized. Provided it is not too much larger than the tracer motion, however, it does not systematically affect the inferred rheology, facilitating the use of image based particle tracking methods.

III. Multiple particle tracking algorithms

Multiple particle tracking can be broken down into four processing stages: correcting imperfections in individual images, accurately locating particle positions, eliminating false or unwanted particles, and finally linking these positions in time to create a trajectory (Crocker, 1996). We perform these tasks with software we developed using the Interactive Data Language (IDL; ITT Visual Information Solutions, Boulder, CO), a high-level programming language used extensively in astronomy and earth sciences, which we have made available as freeware. Each sub-section below is broken into two parts: one describing the general approach and a second with specific instructions for using our IDL software.

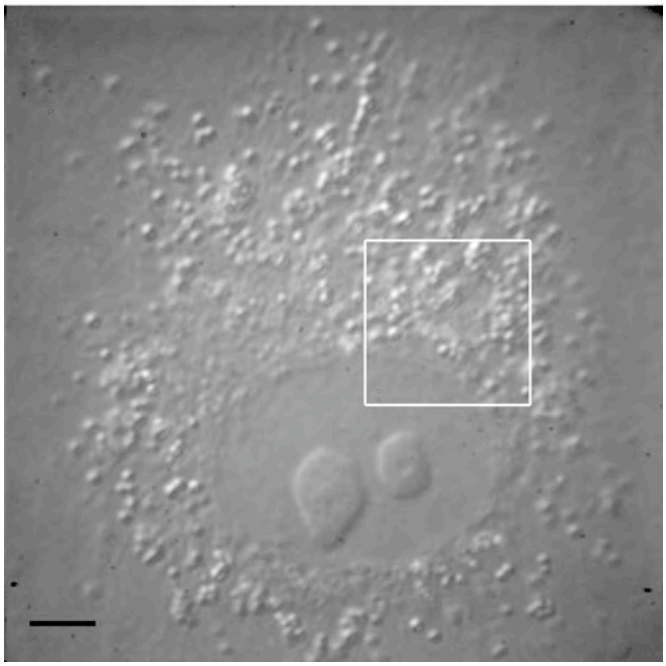

Figure 2. An image of a TC7 Epithelial Cell taken with high-performance shadow-cast DIC and illuminated by a pulsed near-IR laser. White box shows a region of interest for which further analysis will be shown. Scale bar is 5 μm.

A typical cell image is shown in Figure 2. The image shows a TC7 green monkey kidney epithelial cell illuminated with a pulsed near-IR laser, visualized using shadow-cast DIC optics (bright-field Nomarski), and acquired using a non-interlaced camera with 512x512 resolution at a magnification of 96 nm/pixel. In our example, we will track the ~400 nm diameter endogenous lipid granules that are common in these and many other cultured cells. Subsequently, we will highlight only the small region of interest marked in Figure 2, even though the algorithms process the entire image. Other organelles or probes introduced experimentally, such as polystyrene particles (see Wirtz), may be used for a similar purpose.

III.A Image restoration

Even the apparently high quality image in Figure 2 contains imperfections that can frustrate particle locating algorithms. These include both gradual variations in background brightness, 'shading', and 'snow', errors in individual pixel values. However, both effects can be sensibly removed by applying a spatial bandpass filter to the image, which reduces both low spatial frequency (shading) and high-spatial frequency (snow) contributions, while leaving intact the intermediate frequencies corresponding to our tracers.

In IDL: *bpass.pro* is the program used to filter the images the images. It uses a 'wavelet' technique corresponding to convolution with a 'Mexican hat' kernel to remove the random digitization noise and background respectively (Crocker, 1996). The

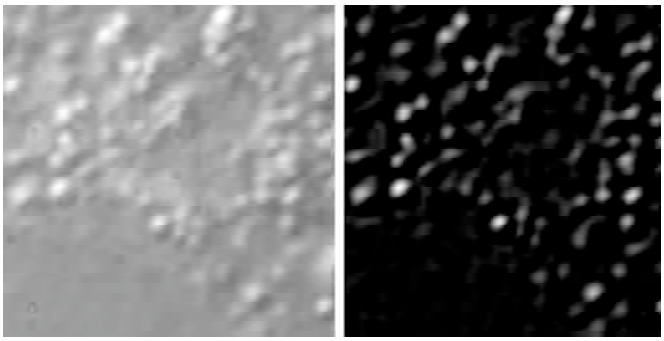

Figure 3. (Left) Close up of region of interest from Fig. 2 (Right) Image after spatial filtering. Notice subtraction of background intensity.

program requires two parameters delimiting the spatial wavelengths, in pixels. Set the first to 1 pixel, and the second to the typical particle diameter, 5 pixels in the example. Filtered images will then retain features with a linear dimension between 1 and 5 pixels. The unfiltered and filtered images are shown in Figure 3.

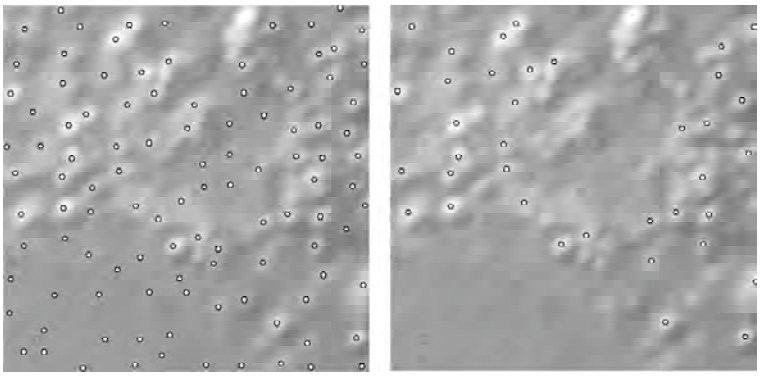

Figure 4. (Left) Region of interest with all local maxima in brightness marked with circles. (Right) Particles that will actually be tracked. Notice particles in the nucleus, particle aggregates, and dim out of focus particles have been discriminated out.

III.B Locating possible particle positions

The image now consists of bright spots on a dark background. As a first approximation for images collected at a high signal-noise ratio (SNR), the brightest pixel in each spot corresponds to its position rounded to the nearest pixel. The distribution of brightness in the rest of the spot can be used to further refine the position to sub-pixel accuracy. To further refine the particle location, we use an algorithm that computes the brightness-weighted centroid within a circular mask that is slightly larger than the particle (Crocker, 1996). While many other methods exist, all perform similarly in the limit of large SNR (high quality pictures) (Cheezum, 2001). Details of this algorithm and the limits to its performance are the subject of Section V.

In IDL: *feature.pro* in our particle tracking suite performs the centroiding operation. This program requires the user to specify the size of the mask to be used. To avoid finding multiple centers on the same particle and to minimize location error, it should be set to a value just larger than the average particle size. Since the sample has particles of different sizes, ranging from 5 to 9 pixels, we use a mask size of 11 pixels. All possible particle positions (local maxima of brightness) are shown if Figure 4(left). This algorithm also computes other properties of each particle image, including total image brightness, as well as average size (radius of gyration) and elongation of the particles. For convenience this program and *bpass.pro* have been combined into one program, *pretrack.pro,* that can run on large, multiple image datasets once the correct bpass and feature parameters have been determined interactively on single images.

III.C Eliminating spurious or unwanted particle trajectories

Because the previous stage of analysis identified all local maxima of brightness as potential particles, it readily identified a large number of maxima corresponding to very low contrast intracellular structures or camera noise. Moreover, it also located particles

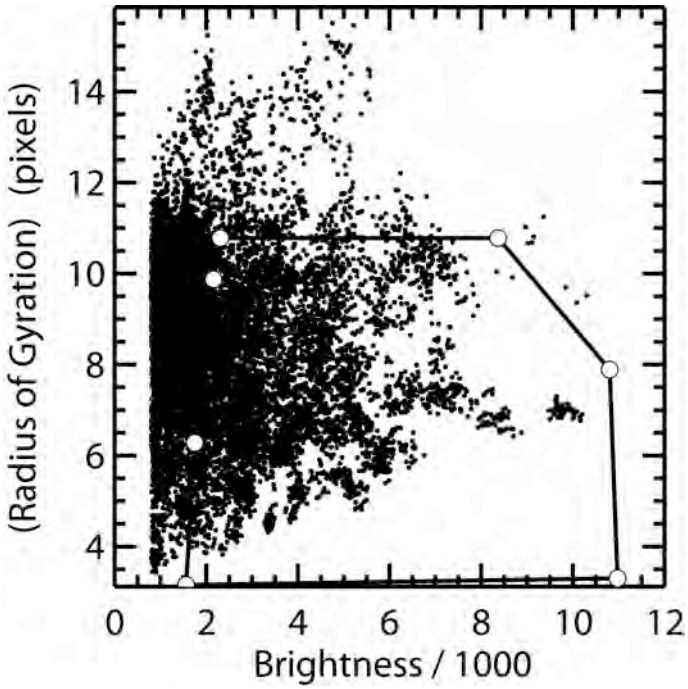

Figure 5. An example polygonal cut in the brightness, radius plane. The left edge is set to isolate the 150-250 brightest tracers. The top edge is limited to the feature size of the mask used during centroiding to eliminate the possibility of tracking multiple points on the same, oversized tracer. The upper right edge is set to avoid tracking large clumps of aggregated particles.

in the nucleus, which presumably should not be pooled together with cytoplasmic lipid granules for analysis. In general, these unwanted tracer locations must be discriminated out based on some combinations of their spatial location and morphology, e.g. total brightness or apparent size.

In IDL: The discrimination process begins by making scatter plots of two of the desired properties (e.g. x and y coordinates or particle brightness and radius). Using *polycut.pro*, the user draws a polygon over the scatter plot using the mouse. The program then removes all of the objects within the polygon. In this example, first the particles in the nucleus are located and removed, by drawing a polygon around the nucleus in the x-y scatter plot ("polygon cut"). If several discontinuous movies are being pooled together, we suggest a separate polygon cut for each, as the cell may drift between sets. Next a polygonal cut in brightness-radius plane, Figure 5, is used to further refine the particles by appearance. Particle positions due to (noise-induced) image artifacts are dim and small. Aggregates of endogenous particles are large and bright. Also points with sizes greater than the feature mask should not be used since they represent features found on large, low-contrast structures inside the cell, such as the membrane. Typically, the user experiments with different polygon placements until only the desired population of features is selected for further processing. Provided the illuminator settings are not altered, the same polygon can be used for all the images of a single cell, or even multiple cells collected in the same session. Note in Figure 4 (right), there are no particles in the background or the nucleus and each particle contains only one center.

III.D Linking positions into trajectories

Having determined a satisfactory set of particle positions, next we match locations in each image with corresponding locations in later images to produce trajectories. This involves determining which particle in a given image most likely corresponds to the particle in the next image. Since all tracer particles are fairly similar, proximity is the

only real indication that particles in successive images correspond to the same physical object. Globally, the optimum identification of particle positions should minimize the total squared distance of travel (Crocker, 1996). To make such identification computationally feasible, a maximum possible particle displacement between images must be specified. This parameter should be sufficiently larger than the maximum distance of travel in single-frame interval, to ensure that physical displacements are never rejected. The trajectories for our example are shown in Figure 6. Notice that the particle motions resemble random walks more than continuous curves. For this reason, alternative tracking algorithms that identify particles by trying to extrapolate their motion or compare their velocities are not well suited to our data. How to extract rheological information from our particles' random trajectories will be discussed in Section IV.

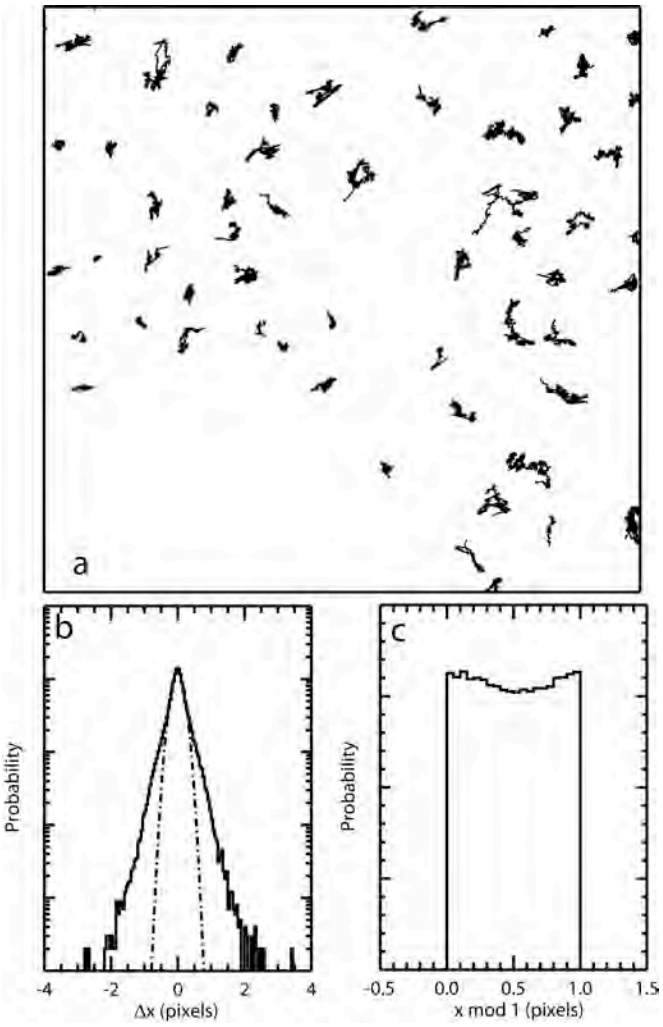

Figure 6. a) Tracer trajectories from the region of interest. b) Histogram of tracer displacements between consecutive frames, which decays well before 4 pixels, the maximum displacement set in the particle tracking software. c) Histogram of the x position of the particle modulo 1, demonstrating that there is negligible systematic error with these choice of tracking parameters.

It should be noted that particle tracking provides an additional layer of filtering for rejecting spurious, artifactual tracers. For instance, if a random fluctuation of the camera noise leads to a spurious feature being identified in a given frame (which survives discrimination), it is highly unlikely that a similar fluctuation will occur in several consecutive frames. As a result, while trajectories of physical particles often contain as many positions as movie frames, those of spurious noise particles are generally very short, and can be easily rejected based on the small number of continuous frames in which they are detected. For this reason, in the preceding discrimination stage it is usually better to adjust the settings to avoid occasional rejection of physical tracers, even if it allows a few spurious particles to be passed through.

In IDL: Particle locations from movies of many images are linked into trajectories using *track.pro*. The user needs to provide one parameter, the maximal distance that a particle may travel between consecutive frames. We find that 4 pixels is a suitable threshold for 50 frames per second data, but this is likely to vary with different cell types and magnification. In this application, setting the 'goodenough' parameter to 10 rejects spurious tracer trajectories having fewer than 10 valid locations.

Figures 6b and 6c display checks that we have found useful to verify the proper operation of the particle locating and tracking algorithms. Figure 6b shows a histogram of the particle displacements between consecutive frames (τ=1/50 second), generated using *makepdf.pro*. Note that the histogram completely decays and is not truncated by the maximum displacement setting, 4 pixels, which is larger than the largest displacement. Moreover, note that the histogram is roughly Gaussian with long tails. If particles were diffusing in water, this histogram would be completely Gaussian. Here, the long tails are an indication that different particles are trapped in heterogeneous local environments. Very long, or nearly flat tails to large displacement are an indication of the tracking of spurious particle images. In Figure 6c a histogram of the fractional part of x position of the particles, x mod 1, is shown; with proper sub-pixel accuracy in particle tracking algorithms there should be no favored sub-pixel value. That means a histogram of the x position of all particle positions modulo 1 should be completely flat. Lack of flatness indicates poor imaging or mask settings, and will be discussed in Section V.B.

III.E Available software packages and computing resources

Source code for our IDL particle tracking routines is available for free download at *http://www.physics.emory.edu/~weeks/idl/* along with a short online tutorial. These particle tracking routines have been translated into the increasingly popular Matlab environment. Both IDL and Matlab are available for a wide variety of computing platforms and many research institutions have a site license for at least one of them. Moreover, a 'stand-alone' version (with a graphical user interface) has been compiled for the Windows, Macintosh and Linux operating systems, and does not require purchasing the IDL software. Links to the Matlab and stand-alone versions may be found at the site above. Many commercially available particle tracking software packages, such as Diatrack and Metamorph, should also be able to track intracellular tracers with comparable precision and ease of use. While these solutions can be quite costly, they have the advantage of offering user support, software customization to specific imaging needs, integrated image acquisition and training courses. It should be remembered, however, that additional routines to compute rheology from the particle trajectories, described in the following section, are at this time only available in IDL form.

These software programs are somewhat computationally intensive, but can be run effectively on a high performance personal computer. The use of large stacks of images can lead to large file sizes (~256 Mbyte). We use a dual 2.4 GHz Athlon processor Linux server, with 2 Gbyte of RAM and a 500 Gbyte RAID array hard drive. In one day, it can convert ~80,000 images, each containing several hundred tracers, into trajectories. A large amount of RAM storage allows for multiple users (ideally, at least 0.5 Gbyte of

RAM per simultaneous user). Image and particle tracking data are analyzed, compressed, and then archived onto removable, external hard drives.

IV. Computing rheology from tracer trajectories

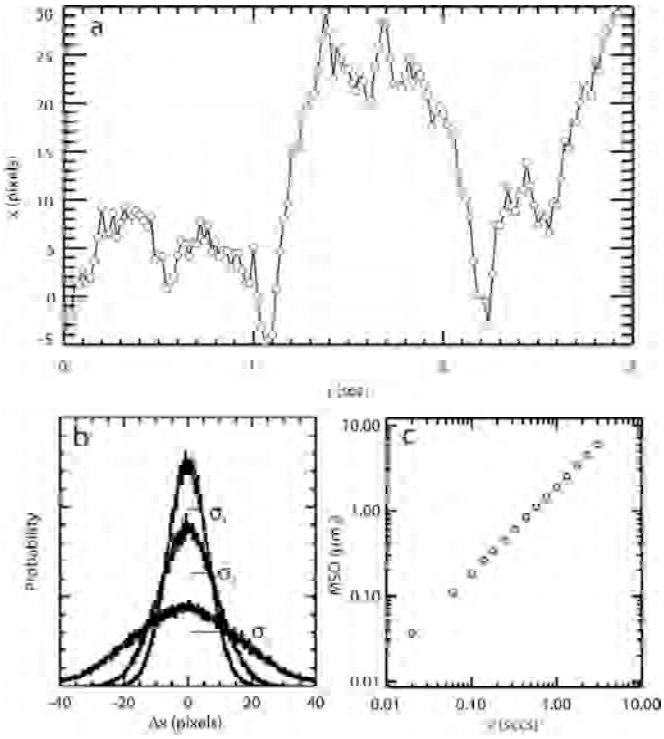

Figure 7. Simulated data for Brownian diffusion in a simple fluid a) A random walk trajectory b) Histograms of the tracer displacements at three lag times, with standard deviations labeled. c) MSD corresponding to (b).

The Central Idea behind passive microrheology is that the random, Brownian motion of small embedded tracers is determined by the stiffness of the material surrounding them; tracers in hard materials naturally move less than ones in soft materials. Here we detail how the ensembles of random particle trajectories computed in the preceding section are converted to a mean-squared displacement, MSD, or its two-point equivalent, 2P-MSD. This is followed by a description of the algorithms that convert this data to the rheology parameters of the surroundings, along with a discussion of their physical and numerical limitations. As for earlier subsections, a general description of the required procedures will be followed by specific instructions for using our IDL routines.

IV.A Computing mean-squared displacements

In Figure 7a the x-component of a numerically generated, Brownian particle is shown as a function of time. Note the random nature of this curve. Mathematically, the MSD is equal to the average, squared distance a particle travels in a given time interval, referred to as lag time, τ , expressed as:

$$\left\langle \Delta x^2(\tau) \right\rangle = \left\langle \left(x(t+\tau) - x(t) \right)^2 \right\rangle \tag{3}$$

where x indicates tracer position in one dimension, t is time, and $\langle\ \rangle$ indicates time and ensemble (over all particles) average. However, it may be best understood graphically. Figure 7b shows a histogram of the distance of travel from 200 particles diffusing in a liquid for 0.02 and 0.4 seconds respectively. The MSD at given lag time is just the square of the standard deviation of particle positions after a given lag time. To convert a

one-dimensional MSD $\langle \Delta x^2(\tau) \rangle$ to a three-dimensional MSD, $\langle \Delta r^2(\tau) \rangle$, multiply its amplitude by 3. This assumes an isotropic material, which means the material has identical properties in the x, y, and z directions.

In IDL: *msd.pro* calculates MSDs from the trajectory files. For the MSD to be in physical units the pixel size and frame rate of the camera must be well known. The magnification can be readily calibrated using microscope stage micrometers (Edmund Optics, Barrington, NJ) or diffraction gratings.

IV.B Computing two-point mean-squared displacements

The basic idea of two-point microrheology (Crocker, 2000) is that inside a viscoelastic solid, a particle cannot move without also moving its neighbors, like two people trying to jump on a trampoline without affecting each other. While in general the correlation of two vectorial displacements is a second-rank tensor, we typically correlate the motion of two particles along their line of centers, as this has the best signal to noise. This correlation may be represented mathematically as:

$$D_{rr}(R,\tau) = \langle \Delta r_1(\tau) \Delta r_2(\tau) \rangle, \tag{4}$$

where $\Delta r_1(\tau)$ and $\Delta r_2(\tau)$ are the stochastic motion of the first and second particle, respectively, along their mutual line of centers in lag time τ (Figure 8). The brackets indicate an average over all pairs of particles in the data set with separation R. In practice pairs with similar R values are pooled together to create a function of both R and τ.

Since its origin is Brownian fluctuations of the material, the τ dependence should be related to the rheology in a manner resembling the conventional MSD. Indeed, this resemblence can be made exact (Crocker 2000) by rescaling the two-point correlation tensor D_{rr} by a geometric factor, $2R/a$:

$$\langle \Delta x^2(\tau) \rangle_2 = \frac{2R}{a} D_{rr} \tag{5}$$

If the material is a) homogeneous and isotropic on length scales significantly smaller than the tracer, b) incompressible, and c) connected to the tracers by uniform no slip boundary conditions over their entire surfaces, the two MSDs will be equal $\langle \Delta x^2(\tau) \rangle_2 = \langle \Delta x^2(\tau) \rangle$. If these boundary and homogeneity conditions are not satisfied, the two MSDs will be unequal. In this case, applying the two-point MSD in the Generalized Stokes-Einstein relation will still yield the 'bulk' rheology of the material (on the long length scale '*R*'), while the conventional single particle MSD will report a rheology that is a complicated superposition of the bulk

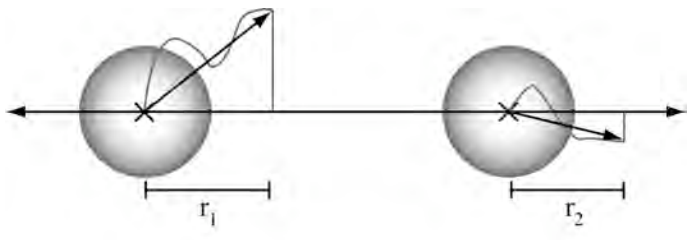

Figure 8. Schematic of displacements used to compute the two-point MSD.

rheology and the rheology of the material at the tracer boundary (Levine, 2001).

A mentioned earlier, two-point measurements will readily detect sample vibration and drift, as these effects lead to completely correlated tracer motion. When this artifactual motion is significant compared to the Brownian signal, we fit the $D_{rr}(R,\tau)$ function to a form $(A(\tau)/R)+B(\tau)$. By using the $A(\tau)$ component exclusively to compute the 2P-MSD and rheology, we can reliably remove artifacts due to sample vibration and drift.

In IDL: *msd2pnt.pro* is used to generate D_{rr} from trajectory data. Like *msd.pro,* it requires the pixel size and the frame rate of the camera to produce data in physical units. Also it requires the user to input a range of separations over which to correlate particles. The images of closely spaced particles can overlap, adding a spurious cross-correlation to their motion and confounding the TPM measurement. We find that using a two micron minimum separation is sufficient to overcome this problem (Lau, 2003). The upper limit is determined by the cell's finite thickness, which is about 4 microns for the cells we study. The fact that the cell is a thick slab rather than an infinite three-dimensional solid leads to deviations from $1/R$ decay. However, we find these deviations become negligible when R is less than about twice the cell thickness, so we typically use an upper limit of 8 microns and a smaller number for thinner cell types.

msdd.pro converts D_{rr} into $\left\langle \Delta x^2(\tau) \right\rangle_2$ and can further delimit the minimum and maximum separations used in the calculation. Since *msd2pnt.pro* can take hours to run, while *msdd.pro* takes seconds, it is more efficient to set R limits wide in *msd2pnt.pro* and then reduce them to different extents in *msdd.pro,* in order to test for the effects of changing the R limits on the final rheology. The 'lfit' keyword in *msdd.pro* causes the routine to perform the vibration/drift correction described above.

IV.C Applying automated image analysis for statistics

The statistical error in particle MSDs is readily estimated. If we approximate the distribution of tracer displacements (as in Figure 7b) as a Gaussian, the standard error for the variance is simply $2 < x^2 > /\sqrt{N_{eff}}$, where N_{eff} is the number of statistically uncorrelated measurements in the distribution. If an image series contains N_t tracers and spans a time interval T, then $N_{eff} \approx N_t T/\tau$. That is, if we image a single particle for 10 seconds at 50 frames a second, we have roughly 500 independent samples of the displacement for a lag time of 1/50 second, but only 10 independent samples for a lag time of 1 second. This τ dependence causes the statistical errors to increase dramatically at longer lag times. As an example, if we were imaging a sample containing 100 tracers at 50 frames per second, and we wanted no more than 1% statistical error in the MSD over the lag times from 1/50 second to 1 second, then we need $N_{eff}=10^4$ independent samples at τ =1 second. Since we are pooling the results of 100 tracers, then we need T=100 seconds of data. Obviously, analyzing the corresponding 5000 images, and 500,000 tracer positions for this modest example requires efficient, automated image analysis. Algorithms that require the user to manually select particles to be tracked, or to estimate their locations are not practical for this application.

In general, two-point correlation functions have much higher statistical noise, requiring the acquisition of significantly higher statistical power, higher tolerance of noisier data, or both. The origin of this is straightforward to understand. The value of D_{rr} is the mean of a distribution of numbers $\langle \Delta r_1 \Delta r_2 \rangle$, Since both Δr_1 and Δr_2 are single particle displacements, the widths of the distribution of $\langle \Delta r_1 \Delta r_2 \rangle$ is roughly $\langle \Delta x^2(\tau) \rangle$, the conventional mean squared displacement, in the limit of weak correlation. In general, the two-point correlated motion, D_{rr}, is much smaller than the single particle MSD, $\langle \Delta x^2(\tau) \rangle$. Indeed, under the most favorable case, the ratio of these two quantities according to Eq. 5 is $2R/a$, which typically has a value of 10-20. We then expect that reliable measurements of the two-point MSD would require averaging at least $(2R/a)^2$ or *several hundred times* (i.e. $(10\text{-}20)^2$) more $\Delta r_1 \Delta r_2$ measurements, relative to Δx^2 measurements to compute a conventional MSD, in order to reach a similar statistical noise. This simple estimate is consistent with our experience.

Does this mean that rather than the 5000 images in our example to measure the conventional MSD, we now need to collect 500,000 images? Fortunately, that is not the case. In a field of view containing 100 randomly located tracers, each tracer might have 10 or more neighbors within the proper distance range for computing two-point correlations. Thus, each image gives us not 100 samples of $\Delta r_1 \Delta r_2$, but more like several thousand. For this reason, the statistical noise of two-point measurements is highly sensitive to the number of tracers in the field of view. In general, if there are 100 or more tracers in a microscope field of view, then about ten times as many images are required to accurately compute a two-point MSD than a conventional MSD. Alternatively, the statistical noise of the two-point measurement will be about $\sqrt{10}$, or just a few times higher than that of the conventional MSD computed from the same data. It should be noted, however, that more statistical power is required for materials where the conventional MSD is much larger than the two-point MSD , according to the square of MSD/2P-MSD ratio. In highly porous materials, the two-point signal can be so small compared to the 'background noise' of uncorrelated tracer motion that it becomes hopelessly impractical to measure from a statistical point of view.

IV.D Converting MSDs to rheology

Earlier we defined mathematical relationships that should allow analytical calculation of rheology properties from the MSD, a process called inversion typically performed in the Fourier space. However notice that the Fourier Transform integral spans all times from zero to infinity. This means we would need data sets spanning this same time interval in order to do the conversion analytically. While this cannot be achieved in practice, many numerical transform methods have been proposed to provide an approximation (Waigh, 2005), with many stressing the importance of collecting data at a high frequency over a very wide time range.

Here we describe a very simple method (Mason, 2000). At each lag time point we estimate the logarithmic derivative (or slope on a log-log plot):

$$\alpha(\tau) = \frac{d\ln\left\langle \Delta r^2(\tau)\right\rangle}{d\ln\tau}. \tag{6}$$

One way to do this is to first take the logarithm of both the MSD and τ values and then for each τ point, fit a line to a few points of the corresponding logarithms of MSD surrounding the chosen τ. The value of this fit at the chosen τ is therefore a smoothed approximation of the MSD, and its slope is the logarithmic derivative. We then use an approximate, algebraic form of the GSER:

$$\left|G^*(\omega)\right| \approx \frac{k_B T}{\pi a\left\langle \Delta r^2(\tau = 1/\omega)\right\rangle \Gamma\left[1+\alpha(\tau = 1/\omega)\right]}, \tag{7}$$

where Γ represents the Gamma function. In using this expression, one first computes a set of ω values that are reciprocals of the measured lag times. At each frequency point, the value of $\left|G^*(\omega)\right|$ is then computed using only the values of the MSD and its logarithmic derivative at $\tau = 1/\omega$. Formally, the value of $\left|G^*(\omega)\right|$ at each frequency would require the numerical Fourier integration of the MSD over all lag times in the interval (0 to infinity), but in practice that integral is dominated by the value of the integrand at $\omega\tau \approx 1$. Eq.7 is exact in the limit that the MSD has a purely power-law form. For other more general forms, Eq. 7 is an excellent approximation at lag times where the MSD is well approximated locally by a power-law, and is seldom more than 15% in error otherwise.

To make contact with more standard rheological representations, one can compute the following from $\left|G^*(\omega)\right|$:

$$\begin{aligned} \delta(\omega) &\equiv \frac{\pi}{2}\frac{d\ln\left|G^*(\omega)\right|}{d\ln\omega} \\ G'(\omega) &= \left|G^*(\omega)\right|\cos\left(\delta(\omega)\right) \\ G''(\omega) &= \left|G^*(\omega)\right|\sin\left(\delta(\omega)\right) \end{aligned} \tag{8}$$

where $\delta(\omega)$ is the phase angle ($\delta = 0$ indicates solid-like behavior and $\delta = \pi/2$ liquid-like behavior), and *G'* and *G"* are the storage and loss moduli satisfying *G*=G'*+i*G"*. In practice, we compute $\delta(\omega)$ using a procedure identical to that for computing $\alpha(\tau)$. Note that when the phase angle is near 0 or $\pi/2$, small errors in $\delta(\omega)$ (due either to statistical noise in the MSD or systematic uncertainties in Eq. 7) can get amplified tremendously in the smaller modulus, G" or G', respectively. We have developed more accurate (but also more complicated) versions of Eq. 7-8, which also rely upon second logarithmic derivatives (Dasgupta, 2002). While these give more accurate results in general, for cell rheology data the difference is negligible.

Additionally there are some artifacts that, no matter what algorithm are used, cannot be corrected for. Because the shear modulus formally depends on the MSD value at over a

finite range of frequencies, shear moduli at the extrema of frequency range are subjected to additional 'truncation' errors. The algorithm above implicitly assumes that the power-law behavior at the extrema of the data set extends indefinitely to higher or lower frequencies. Any curvature, even the slightest ripple, at the extremal lag times can cause a disproportionate change in the shear moduli. Consequently, unless there are physical motivations for believing this extrapolation is justified, confidence in this part of the data should be low. Strictly speaking, it may be best to compute the rheology from the entire lag time range and then ignore a decade (one order of magnitude) of frequency on each end of the curve.

MSDs themselves are often subject to artifacts at their frequency extrema as well, as described in Section V.E. Any systematic deviations of the MSD due to these effects will be further amplified in the shear moduli. Therefore in some situations it might be wise to truncate the lag time range of the MSD prior to computing the rheology. In general, as the lag time range of the MSD being used is restricted, the ends of the curve will tend to 'wiggle'. This is a sign that the shear moduli being produced are subject to strong truncation errors.

One exercise we have found invaluable when utilizing this algorithm is to invert numerically generated MSD curves that resemble the actual data. Changing the dynamic range of such simulated data allows the identification of truncation effects. Furthermore, adding Gaussian distributed noise to this data is helpful for determining the artifacts associated with it. A final warning is that a "reasonable" appearance of calculated rheology (e.g. resembling something from a rheology text book) does not mean the inversion is physical. The emergence of noise floor on data from a particle diffusing in a liquid will lead to results that look almost exactly like the rheology of a Maxwell fluid, which is a common example in these texts.

In IDL: *micrheo.pro* is the program in our IDL suite that does these calculations. It has a built in smoothing to help reduce statistical noise. In addition, it uses second order formula for smaller systematic errors (Dasgupta, 2002), and provides warnings when the computed shear moduli are numerically unreliable. Its input is either a MSD or 2P-MSD.

V. Error sources in multiple particle tracking

In this section, we describe several common sources of error in particle tracking instruments. While some of these error sources can be mitigated by the use of high-quality equipment, most errors are due to irreducible physical limitations on imaging detector performance and illumination brightness. In practice, a good understanding of the origin of different errors, followed by careful adjustment of the imaging system to optimal settings, can lead to significant performance improvements. We will describe three classes of error: random error, systematic errors and dynamic error.

V.A Random Error (Camera Noise)

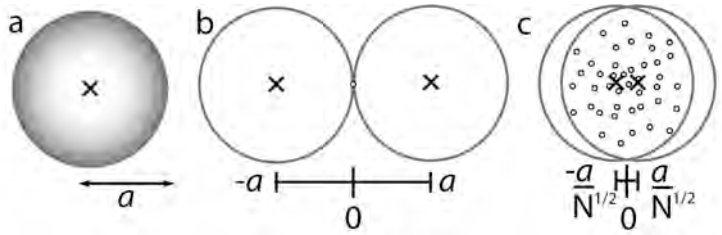

Figure 9. a) An idealized, circular particle of radius a whose center is marked by an x. b) If a single photon is detected the position of the particle (set to be at zero) can be determined to +/- a. c) If more photons are detected a cloud is detected. This limits the potential position (reduces measurement error) of the particle.

The more accurately a given particle can be located, the higher the quality of measurements of cellular rheology. In fact with optimal particles, illumination, imaging and software it is possible to measure the position of a 400 nm particle to within approximately 5 nm. While it may seem remarkable to be able to determine the location of a particle this well with optical methods, the limiting precision of particle localization is quite different from optical resolution - the limit below which structural details in complicated specimens cannot be discerned (typically about a quarter of the wavelength of light used). Instead, the situation is analogous to a familiar problem in curve fitting, finding the position of a local maximum or peak in a curve. If we have a reasonable number of evaluations of a peaked function and the values are very precise, a least squares fitter can locate the peak center to an arbitrarily high precision (relative to the width of the peak). In this analogy, the peaked function is the light intensity distribution for a tracer, with a width set by the tracer size or optical resolution.

While some noise in a camera depends on the details of its construction and electronics, there are ultimate physical limits on the performance of all cameras due to the discrete nature of light itself. Imagine that we had a 'perfect' camera that recorded the precise 2-d coordinates of all incoming photons arriving from a microscope. An 'image' of a single small tracer might resemble a round cloud of points in a 2-d scatter plot (Figure 9). From elementary statistics, we know that the standard error σ_x, σ_y when computing the mean x and y center positions of the cloud (corresponding to the particle location) is:

$$\sigma_x = \sigma_y = a/\sqrt{N} \tag{9}$$

where a is the apparent radius (technically the standard deviation) of the cloud, and N is the number of points/photons, (assuming a Gaussian distribution).
In practice, well-made cameras approach this physical limit of performance when they are operating at illumination levels near the maximum allowed by detector saturation. Of course, rather than averaging the positions of individual photons, we instead determine the tracer's position by calculating the centroid of an intensity distribution, analogous to computing the 'center of mass' for a continuous mass distribution, (Figure 10). Provided that we spread the tracer image over a sufficient number of pixels to reasonably represent the tracer's brightness distribution, the error in the centroid positions is precisely the same as that in the previous example.

In summary, whenever imaging a sample, there will be statistical fluctuations in the number of photons detected by the camera even under ideal conditions. These fluctuations are irreducible and lead to random errors in the reported brightness of individual pixels. In general, using a higher illumination will reduce this error, but is ultimately limited by detector saturation.

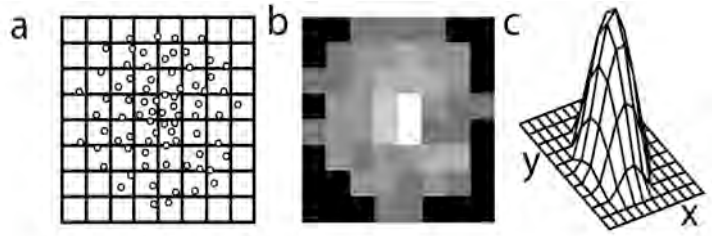

Figure 10. a) A hypothetical distribution of photons on a pixilated grid. b) A grayscale image corresponding to the photon distribution. c) The brightness distribution of a single particle from the sample imaging. Particle tracking algorithms find the center (peak) of this distribution very accurately.

V.B Systematic Errors

In general, dividing the image of a small tracer into pixels does not necessarily introduce significant errors. However, improper image pixelation and masking can introduce systematic (as opposed to random) errors into the computed particle positions. , One form of systematic error is 'pixel biasing' the tendency of the algorithm to 'round' the centroid posititon to the nearest pixel, as mentioned in section III.D. At its worst, a particle moving at a uniform speed across the field of view would appear to 'hop' from one pixel to another (Figure 11), introducing ±0.5 pixels of position error.

This problem can largely be avoided by use of the proper magnification and image processing settings. One source of pixel biasing is too low a magnification—ensure that single tracer images appear more than 3-4 pixels in diameter. This error also commonly occurs during calculation of the centroid. When computing the centroid of a single tracer image, most algorithms mask off (force into zero brightness) pixels outside some radius to avoid contributions from other nearby particles. Error is introduced if the background of the image is not zero at the edge of the mask. Physically, common image defects such as diffraction rings and out of focus background particles, or user errors such as specifying too small of a mask size, can cause this problem. For images relatively free of obvious defects, it is usually possible

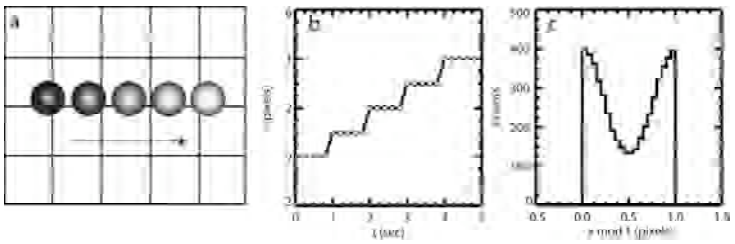

Figure 11. An example of systematic error. a) The tracer images are smaller than one pixel and slowly moving in the x direction (represented as changing from dark to light). b) The tracer's inferred position moves in 1 pixel increments. c) This leads to histograms of the fractional part of the x coordinate that has a minimum at a half pixel, instead of being flat.

to reduce this error to less than 0.1 pixels by using a suitably large mask size.

Another, rather subtle source of systematic error is due to the fact that the individual detectors corresponding to each camera pixel are not equally sensitive to light, with the typical variations between pixels being a few percent. This typically introduces a systematic error between the physical and measured position that is a few percent of a pixel, and which varies randomly with location in the field of view. This error can be largely eliminated by calibrating each pixel in the camera using a uniformly illuminated specimen (e.g. an empty bright field), and then numerically correcting all pixel intensity values by division against the calibration image prior to centroid finding. This algorithm, termed 'flat-fielding' is routinely used by astronomers, but is seldom needed for particle tracking applications.

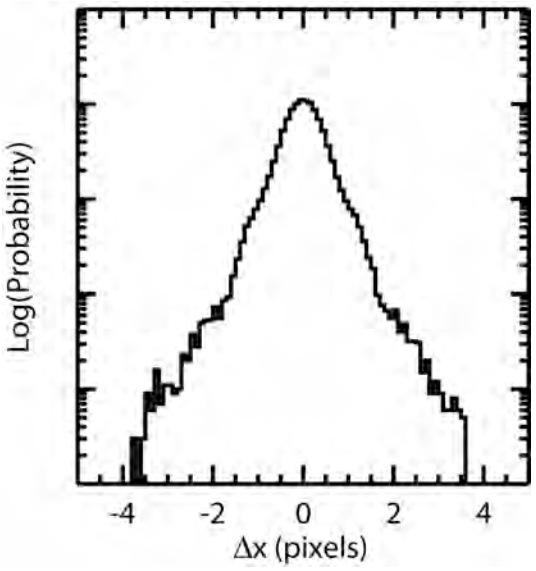

Figure 12. Another manifestation of systematic error. The ripples in this histogram of tracer displacement indicate that tracers seem likely move in increments of a pixel. This type of motion is non- physical and is most likely a artifact from choosing poor tracking parameters.

For microrheology using large ensembles of particles , the contribution of all these systematic errors can often be ignored. For example, if a tracer moves a half a pixel, it may appear to move 0.4 or 0.6 pixels due to a 0.1 pixel systematic error. When many such measurements are squared and averaged together to compute an MSD, the deviations will largely cancel. However, this may not be the case in other particle tracking situations with few particles, poor resolution, or short trajectory lengths in comparison to the lag times of interest. If one is computing a histogram of displacements, such as in Figure 12, these systematic errors will lead to an obvious 'ripple' in the measurements, with displacements of integer multiples of one pixel length being more likely than half-integer multiples.

V.C Dynamic Error

A subtle form of error that affects measurements of tracers' random motion is due to systematic underestimation of tracer motion when finite exposure times are used (Savin, 2005). This error is absent when the duration of the exposure is infinitesimal compared to the time interval between successive movie frames. If the particle moves significantly during the camera exposure, the measured centroid reports the time-averaged position during the exposure (Figure 13). Because random walks tend to loop back on

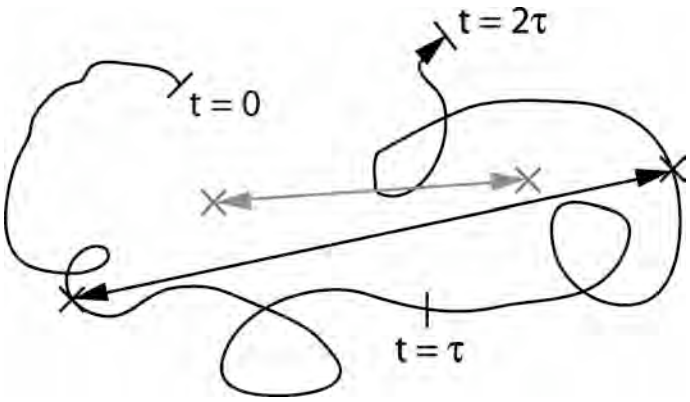

Figure 13. An example of dynamic error in a random walk trajectory. Using a short shutter time, the position at a given instant is recorded, e.g. t=0.5τ and 1.5τ (black crosses). With a long shutter time,τ, is used, the time averaged location of a tracer is found, e.g. t intervals from 0-1τ and 1-2τ (gray crosses). The distance between time-averaged locations is systematically less than for instantaneous ones.

themselves, this time averaging effect artificially reduces the amplitude of the mean squared displacement. This selective reduction of short lag time MSD value will make diffusive Brownian motion look super-diffusive (the MSD will increase more rapidly than linearly with time) at short lag times.

The requirements to reduce dynamic error and optimize the MSD precision are at odds with each other. High precision demands as much light as possible be used to form the image, which is most easily achieved by increasing the exposure integration time, while dynamic error requires minimizing the exposure time. Obviously, the solution is to use the longest exposure possible that will not introduce dynamic error. We know of no rigorous mathematical criterion in determining the optimal exposure time for an arbitrary particle tracking situation. In general, we find that if the amplitude of the mean-squared motion *during* the exposure interval is less than a quarter of the mean-squared motion *between* exposures, then the underestimation of MSD caused by dynamic error will be less than 10%. For tracers undergoing simple diffusive Brownian motion, this corresponds to an exposure time no longer than one quarter the time interval between frames, e.g. 1/200 sec for acquisition speeds of 50 frames per second. For the general case, the criterion is more difficult to apply, as it requires knowledge of the MSD at lag times shorter than those being measured. In that case, one is reduced to guessing the form of the short time MSD by extrapolating the available data, or relying on physical models or literature results. Note that in the case where the MSD is nearly time independent (as in an elastic solid) the exposure time needs to be much smaller than a quarter of the frame interval to eliminate dynamic error.

V.D Sample drift, computational detrending, and its limitations

Drift in the sample stage contributes another source of uncertainty. Thermal expansion can cause the stage or sample chambers to slowly translate across the camera field of view during long experiments, often by several pixels. Computational correction of such global error is referred to as “detrending”. Subtracting a linear fit from a single particle trajectory will remove this drift, but will often introduce artifacts into the MSD at long times because it subtract some of the tracer’s actual motion. Alternatively, if a large number of tracers are available, one can compute the average motion of the tracers’

center of mass, and subtract it from each particle position. In general, this procedure reduces the amplitude of the MSD by a factor $1/N$ (in addition to that caused by detrending), so is only recommended if $N >> 10$.

The utility of all such 'detrending' algorithms is significantly reduced by the systematic position errors described in the previous section. The non-linear mapping of actual to measured position will cause even motionless tracers to move slightly relative to one another, over the time scale of detrending. In a similar vein, when studying very slowly moving tracers, it is often tempting to time average multiple measurements (perhaps with detrending) to yield spatial resolutions that are hypothetically much smaller than a nanometer. Systematic errors tend to limit the utility of such approaches, because they are not reduced by averaging.

Lastly, all of the measurements described here are in the two dimensions x and y. While the particles also fluctuate in z, which might be detected by fluctuations in particle intensity or diffraction pattern, the centroid of any spherical particle is still well-determined as above. The distance R between any two such particles is obviously then the distance projected into the imaging plane.

V.E Effects of measurement errors on the MSD

Particle tracking errors lead to several perturbations in measured MSDs (Savin, 2005; Martin, 2002). The most common error is due to the contribution from random position measurement error, as outlined in Section V.A, which contribute additively to the value of the actual, physical MSD:

$$\left\langle \Delta x^2 \right\rangle_{measured} = \left\langle \Delta x^2 \right\rangle_{physical} + 2\sigma_x^2, \tag{10}$$

where σ_x is the standard deviation of the random position measurement error. At short lag times, where the MSD is small, the measured MSD will start to flatten and appear sub-diffusive if the physical motion approaches the precision of the MSD. Conversely, if dynamic errors are significant, the MSD can appear super-diffusive at very short times.

Another common artifact is that MSDs will tend to flatten at long times. This is due to particle loss from the imaging plane. While we image in two dimensions, particles diffuse in three dimensions and at long times will diffuse out of the imaging plane. Since large particles move less than small particles, larger ones tend to stay in the focal plane longer. This can lead to downturns in the MSD, when the typical motion exceeds the square of the depth of focus. Our MSD calculation software tabulates the number of independent measurements, N_{eff}, used as a function of lag time. As discussed in Section IV.C, this quantity should decay as $1/\tau$; deviations from this indicate significant particle loss.

Reliably estimating the precision of measured MSDs can be quite difficult. A common method is for experimenters to estimate their total tracking precision by measuring the apparent MSD of immobilized tracers in 'control' samples. This has obvious educational value, and allows the effects of different microscope adjustments and ambient sample

vibration to be rapidly determined. Still, the resulting error estimates must be taken with a grain of salt relative to the intracellular application we have described here. Obviously, ensembles of immobilized tracers will not display any dynamic error, and systematic error cannot be easily estimated. Moreover, the observed random position error will only correspond to the intracellular case if the optical parameters of that problem, such as tracer size and index of refraction are faithfully reproduced, which is often difficult in practice.

VI. Instrument requirements for high performance tracking

Even though our software tools can often yield acceptable particle tracking under less than ideal imaging conditions, acquiring high quality images facilitates image analysis and yields more reliable rheology data. The following sub-sections discuss important practices of instrumentation that minimize errors of measurements

VI.A Isolation from vibration, acoustic noise and thermal drift

Research laboratories often contain many sources of vibration (blowers for air handlers, fume hoods, cell culture hoods, as well as people walking and talking). All of these shake the microscope and the sample stage, often by a surprisingly large amount. In almost all research locations, a pneumatic isolation table is absolutely critical for dampening out these vibrations sources. Furthermore connections from the outside world (power cords, gas lines) can conduct vibrations to the floating tabletop; their number should be minimized. For example, rather than having several AC power cords hanging between the isolated table and the outside, these can be ganged into a power strip, and its single cord used to connect off the table. Moreover, many electronic devices contain small cooling fans and electrical transformers that can be significant sources of vibration. These should not be placed on the isolated table surface. For the same reason, we use an external lamp power supply in lieu of the supply built into the microscope base (which has the additional advantage or reducing thermal expansion drift in the microscope body). Additionally, many inverted microscopes have a screw to lock the illumination stalk. We find tightening this down is helpful with reducing sample vibration.

Microscopes and sample chambers will expand and contract as they change temperature. If the scope is to be heated let it reach the set-point temperature for 10-20 minutes before starting experiments. Air currents from ventilation systems can cause substantial drift and periodic motion in a microscope over time scales as short as tens of seconds. Surrounding the microscopic with a curtain to prevent drafts from blowing directly on the microscope are very helpful and will largely eliminate these effects. Lastly, stage drift can also be induced by the body heat and exhalations of an operator sitting too close to the microscope.

VI.B Microscopy – generation of high contrast tracer images.

Maximizing image quality and MSD precision requires juggling several imaging parameters. Specifically, we want to adjust the microscope to simultaneously maximize the contrast between the tracer image and its surrounding background and magnification, while using the maximal illumination intensity and minimal exposure time to avoid dynamic error.

In general, all other things being equal, the precision of MSD improves proportionally with the magnification and as the square root of the tracer brightness. Because it is difficult to predict the performance of a given microscopy condition beforehand, we take an empirical approach. Most image analysis software allows the user to acquire a single frame and to read out the brightness value of individual pixels selected with a cursor. For 8-bit cameras, this value is an integer between 0 and 255. We routinely use this function to determine the brightness difference between the center of a typical tracer and the background. Ideally, the brightness difference should be at least 100 out of 255, while for some simple specimens and high refractive index tracers it can be adjusted as high as 150-200. At the same time, it can be validated that image saturation, i.e. pixel values of 0 or 255, is not present.

Most image acquisition hardware and software allows the user to change image brightness and contrast, and to make an image collected at low illumination look like one collected at high illumination. However, it should be remembered that changing such 'gain' will not affect the final tracking precision, because both the signal and camera noise are amplified together. Even though the images may look the same to the eye, the image from the lower illumination setting will have a higher level of pixel noise and thus poorer resolution. For this reason, 'gain' should be set at the minimum setting and any automatic controls in the camera (e.g. automatic gain control) should be turned off.

Contrast Generation Methods: Care should be taken when choosing and using optical contrast generation techniques to visualize low refractive index tracers (such as the endogenous lipid granules we use here). Remember that our tracking software is not based on contrast as a ratio, but on the absolute difference in brightness. For example, extinction-mode DIC (also called dark-field Nomarski) creates a very dark background surrounding bright tracers; in other words, it produces a high brightness ratio. However, extinction mode also blocks much of the illumination, resulting in an image where the tracers have a low absolute brightness. As a result, extinction mode DIC generally yields lower quality tracking than two other contrast enhancement techniques: shadow-cast DIC (bright-field Nomarski) or phase contrast, which generate a larger absolute difference in brightness for a given specimen illumination. Epi-fluorescence imaging of fluorescently tagged microspheres can also create images with high contrast ratio, although the overall brightness of the image is significantly lower than that provided by bright-field methods. Therefore, fluorescence imaging is not well-suited for high performance particle tracking.

With any contrast generation method, the microscope's performance depends critically on operator skill and proper alignment. For example, DIC will give images that 'look

right' but have degraded contrast if the wrong DIC prism is used, the polarizers are misaligned or there are small bubbles in the immersion oil. Similarly, phase contrast will have degraded performance if the phase ring is not properly centered. Given all that can go wrong, quickly checking the actual brightness difference achieved (and verifying that it is comparable with earlier experiments) is usually a worthwhile practice prior to data collection.

Numerical aperture of condenser: For conventional bright-field microscopy, the numerical aperture (NA) of the illumination (i.e. the aperture diaphragm setting in the condenser) is critical to maximizing spatial resolution. For example, with low NA illumination, the tracer images have a very high contrast, large depth of focus and prominent diffraction rings. Because much of the light is stopped by the aperture diaphragm, however, the absolute brightness of the image is low. Conversely, at high NA, the absolute image brightness is maximized, but the contrast between the tracers and the background can be quite small. The optimum setting is clearly in between the two extremes: where tracer contrast is good, and total image brightness is not too attenuated. For our microscope operating in conventional bright field or DIC microscopy with an oil immersion condenser, a condenser aperture setting of near NA=0.4 yields the best brightness difference and particle tracking performance.

Magnification: The absolute minimum magnification is set by the size of the particles that are being tracked, particles images must be at least 3 to 4 pixels in diameter. However higher magnifications could be used to increase spatial precision, so long as total illumination can also be increased to keep the detector near saturation levels. With the higher magnification, the intensity of the particle can be spread over more pixels, allowing more accurate centroiding. However, it is important to note that, without increasing the illumination intensity, the signal at each pixel decreases in proportion to the square of magnification. Therefore, simply increasing the magnification without simultaneously increasing the illumination will not increase performance.

VI.C Using low-noise, non-interlaced camera.

The amount of random pixel noise present in a camera can be readily measured. One acquires a series of 'blank' images where the illumination is uniform across the field of view and near the saturation value. Because of camera noise, the brightness value of a given pixel will vary from image to image. One can then take the differences of intensities at a given pixel between pairs of images, and compute the standard deviation of the resulting differences. Division of this standard deviation by $\sqrt{2}$ yields the RMS (root-mean-squared) pixel noise. For a quality camera, the RMS noise amplitude should be 0.5-1% of the pixel brightness, which is calculated by subtracting out the average reading in dark images. The pixel detectors on many cameras hold about 40,000 photoelectrons at saturation (called the 'well capacity'); the corresponding statistical noise is $1/\sqrt{40000} = 0.5\%$. Some inexpensive consumer cameras can have noise figures significantly higher than this, and they are not recommended for particle tracking. Specialized camera with larger well capacities might have significantly better performance, but will likely require 12-bit digitization to take advantage of the lower

noise. 12-bit digitization is typically used to improve the precision of intensity readout, rather than to reduce camera readout noise, and thus has a limited impact on particle tracking applications. It should also be noted that the lossy compression algorithms (such as JPEG or MPEG) tend to contribute low amplitude defects to image data that can degrade particle tracking performance. If high precision is needed, uncompressed images or non-lossy compression such as LZW should be used exclusively.

Interlacing is a standard left over from the early days of video technology. To eliminate image flicker (which requires 50-60 images per second) while reducing transmission bandwidth, the display redraws half of the image at a time, alternating between the even and odd rows. This eliminates visual flicker while only transmitting 30 frames per second (for NTSC video in N. America and Japan) or 25 frames per second (for PAL or SECAM video elsewhere). While the odd and even rows can be separately analyzed to perform particle tracking at equivalent rates of 50 or 60 frames per second, interlacing degrades tracking performance because only half as many pixels are used when computing the centroid. When spatial precision is critical, a non-interlaced (also called a progressive scan) camera should be employed.

VI.D Using high-intensity, filtered illuminator.

Intensity and controlling dynamic error: As described earlier, the precision of particle tracking techniques is ultimately limited by the intensity of the illumination. In fact the ideal illumination would be a very short, very bright pulse of light, like a camera flash, which simultaneously reduces both random and dynamic error. A practical approach to illuminating samples for particle tracking is to use the microscope illuminator at its maximum setting and then vary the exposure time in the camera or the pulse length of light to obtain the required SNR for image processing. In general longer exposure time will decrease random position error, while increasing dynamic error. While random error has a fixed amplitude, dynamic error is proportional to the tracer's MSD. Thus, if the MSD at the shortest lag time is much larger than the precision of the determined MSD, one can safely reduce the exposure time to minimize the dynamic error. Conversely, if the MSD at the shortest lag time is comparable or smaller than its precision, then a longer exposure is probably warranted.

Even at the maximum setting on a 100 Watt halogen illuminator, exposure times of several milliseconds are typically required to produce optimal, near-saturation bright-field images at a magnification of 100 nm/pixel. Dynamic error is thus likely even at frame rates in excess of 50 frames per second. When we require faster acquisition rates to perform high frequency microrheology, we employ a custom-built, pulsed near-IR laser illuminator that can fully expose the detector in tens of microseconds, allowing the use of a high speed camera operating at several thousand frames per second without dynamic error. Pulsing the laser also minimizes the possibility of heating and photodamage effects.

Filtering: One downside to the use of high dose of photons is the tendency to heat or photodamage living cells. In general, halogen lamps emit copious long-wave UV light,

and their intensity actually peaks well into the near-IR. Many living cells under the continuous illumination of unfiltered, maximum setting (100 Watt) halogen illumination die in less than an hour, while cells just outside of the illuminated field appear unaffected. To minimize radiation damage we place two optical filters in the illumination light path (purchased from Edmund Scientific): one (#F54-049) that blocks UV wavelengths shorter than 400 nm and another (#F46-386) that blocks wavelengths from 700 nm to ~1800 nm. With both filters in place, cells in and out of the illumination field show indistinguishable long-term viability and proliferation on a heated 37°C microscope stage.

VI.E Using synthetic tracers to increase visibility.

One way to enhance the contrast of tracers is to use synthetic microparticles of a high refractive index n, e.g. polystyrene (n=1.6). Many such probes have been used in cellular rheology experiments, being introduced by phagocytosis (Hoffman, 2006; Caspi, 2000) or micro-injection (Tseng, 2002). Moreover, such particles could potentially be designed to target a specific structure (actin filaments or microtubules) or a specific protein (different integrins or other surface receptors) (Puig-de-Morales, 2004). In *in-vitro* studies, the surface chemistry of the bead has large effects on the resulting rheological measurements (Valentine, 2004). The effect of tracer surface chemistry on the inferred 'one-point' microrheology of cells is still an issue without a clear consensus. On the other hand, tracer chemistry is expected to have a much more limited effect on TPM measurements.

VII. Example: Cultured Epithelial Cells

VII.A Particle tracking results for TC7 epithelial cells

The results of particle tracking in TC7 cells are shown in Figure 14. The top panel shows the MSD and 2P-MSD for the endogenous tracers found in living cells, computed using 18,000 images comparable to those shown in Section III. We find that the tracers' conventional MSD looks almost purely diffusive. If we assumed the motion to be Brownian (rather than due to ATP-dependent processes), we would conclude the cell interior is a simple fluid with a viscosity roughly ten times that of water. This explains why the random motion of organelles in cells so strongly resembles the Brownian motion of tracers in water—both are diffusive motion with comparable amplitudes. The results of the two-point calculation give a 2P-MSD with a super-diffusive lag time dependence, $\text{2P-MSD} \sim \tau^{1.5}$, which is faster than a linear dependence on lag time expected for Brownian motion and clearly indicative of active intracellular processes (Lau, 2003). The functional form of D_{rr} consistently displays a $1/R$ dependence, providing a positive control that cells can be treated as three dimensional viscoelastic continua. However, the fact that single particle MSD and 2P-MSD differ in amplitude implies that the no-slip, homogeneous boundary condition assumptions needed for the conventional GSER are not valid (see Section IV.B). The difference in functional form either indicates the material in the tracers' local environment has different rheological characteristics from the bulk

material, or differential sensitivity to the (yet unknown) non-Brownian driving forces. All in all, it is difficult to definitively conclude much else from such live cell data alone.

The bottom panel of Figure 14 shows data for cells that have been depleted of intracellular ATP using a combination of sodium azide and deoxyglucose. ATP depletion inhibits motor dependent processes but may also cause structural changes, for example by inhibiting a wide range of kinase dependent processes. The amplitude of both the MSD and the 2P-MSD have dropped dramatically, as expected if the active processes have been slowed or eliminated. Both curves now display a simple power-law frequency dependence on lag time: MSD~2P-MSD~$\tau^{0.3}$. Both curves still turn upward at lag-times greater than 1 second. Measurements with varied dosages of ATP depletion agents suggest that this upturn is due to residual ATP activity, while the MSDs at shorter times are ATP independent and Brownian in origin (Hoffman, 2006). The fact that two types of MSD have the same lag time dependence suggests that they are now probing the same viscoelastic structure. The observation that they differ in amplitude by an order of magnitude suggests that individual tracers do not have the simple, homogeneous no-slip boundary conditions commonly assumed. One simple explanation would be if the tracers are strongly adhered or attached to a viscoelastic network that is porous on their length scale (Van Citters, 2006), which seems plausible given our knowledge of cytoskeletal ultrastructure. Lastly, note that the 2P-MSD has noticeably higher (but still acceptable) statistical noise versus the conventional MSD, as expected from Section VI.C.

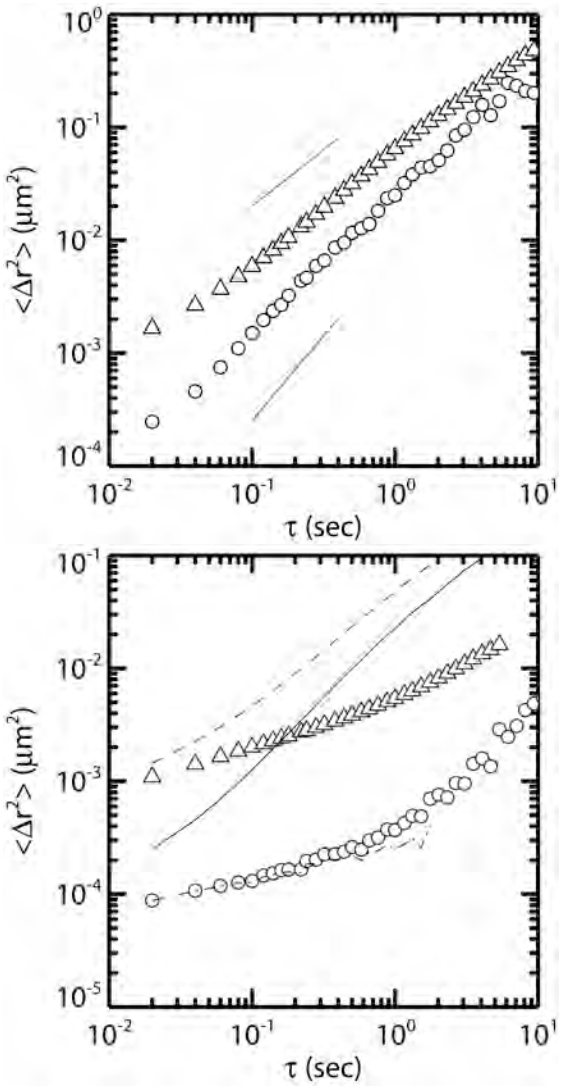

Figure 14. (top) MSD (triangles) and 2P-MSD (circles) for a single TC7 epithelial cell. Lines are eye-guides with slope 1 and 1.5 (bottom) Similar data for ATP depleted TC7 cells, averaged over N=8 cells, each with 18,000 images. MSDs are systematically larger than 2P-MSDs. Non-depleted data are reproduced as curves in bottom panel for comparison.

VII.B Two-point microrheology of TC7 epithelial cells

The storage and loss modulus computed from our example data, specifically the 2P-MSD of ATP depleted cells, is shown in Figure 15. The response has a simple power law form, $G'(\omega) \sim \omega^{0.3}$, over the range of accessible frequencies, similar to our published results (Hoffman, 2006). Notice that the statistical noise (the point to point amplitude variations) in the computed rheology appears much lower than in the corresponding 2P-MSD. This is because our algorithm averages the MSD over the time domain, to draw a 'smooth curve' through the MSD prior to estimating the logarithmic slope as described in Section IV.D. This smoothing does not eliminate the noise, but replaces it with a slow 'ripple' in the computed shear moduli, which is expected to follow a pure power-law in the frequency window, corresponding to two parallel lines for *G*' and *G*" (compare with Figure 1). The artifactual curvature shown in Figure 15 is typical of results computed from 2P-MSD with the noise shown in Figure 14. The curvature is always much more pronounced in the smaller of the two shear moduli, G" in this example. Nevertheless, this rheological measurement would provide reliable power-law fits to infer both the amplitude of the shear modulus (i.e. stiffness) and power-law exponent β.

As we noted in Section II.C, two-point microrheology can provide reliable stiffness information even in heterogeneous media, or when the connections between the tracers and the network do not satisfy the requirements of the GSER, Eq. 2. Averaging the results of TPM measurements on many TC7 cells yield a mean stiffness value of 40 Pa at ω =10 rad/sec (~1.6 Hz). Cell to cell variations are considerable, with a log-standard deviation of 1.6, meaning that ~70% of the values fall within 1/1.6 and 1.6 times the mean value, or in the range 25-60 Pa. It should be remembered that TPM reports a whole-cell average; the magnitude of the stiffness variations across a cell has not yet been determined using TPM. Moreover, probes for TPM may not uniformly sample all regions of the cell. This mean stiffness is more than an order of magnitude lower than the value (~1 kPa) inferred by several other cell rheology techniques, which we presume are probing the cell cortex. From a cell trafficking point of view, it is much easier to imagine organelles being trafficked through a soft 40 Pa intracellular space than a stiff 1 kPa network.

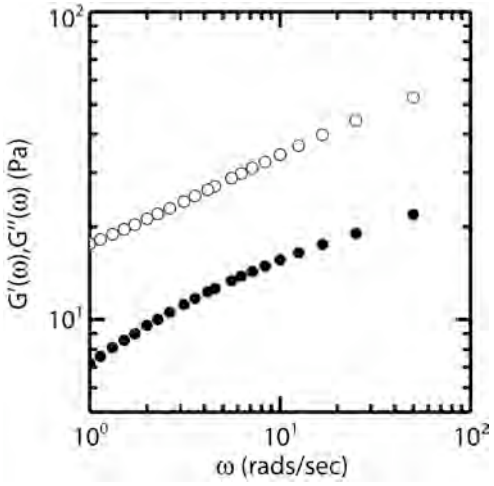

Figure 15. Inferred two-point microrheology of ATP-depleted TC7 cells, computed using the data in Fig. 13. The response is dominated by the storage modulus, G' (open circles). The curvature of the smaller loss modulus (closed circles) is artifactual.

VII.C Computing stress fluctuation spectra

In living cells where the rheology is known, it is also possible to use TPM to quantify the amplitude and frequency dependence of the cells' non-Brownian stress fluctuations (Lau, 2003), which are presumably ultimately due to forces generated by molecular motors or filament treadmilling. The basic idea is similar to that of the linear response of a spring. If we apply a force F to a spring having a spring constant k, the deflection x satisfies: $F=-kx$. If we know any two of the three variables, we can compute the third. Similarly, if we know the stiffness of a material (its shear modulus) and how much it is deforming (by measuring the 2P-MSD of embedded tracers) we can infer the magnitude of the force or stress driving the deformation (all at a given frequency, ω). In a viscoelastic material, the formula is:

$$\left\langle \Delta(\omega) \right\rangle \approx 3\pi a \left| G^*(\omega) \right|^2 \left\langle \Delta x^2(\omega) \right\rangle_2, \tag{11}$$

where $\left\langle \Delta x^2(\omega) \right\rangle_2$ is the unilateral Fourier transform of the 2P-MSD, as used in the GSER, Eq. 2, and $\left\langle \Delta(\omega) \right\rangle$ is the power spectrum of the fluctuating stresses causing the material to deform. The power spectrum of a fluctuating function is the squared amplitude of its Fourier transform. In general, such a spectral representation is used to quantify the relative contributions of low versus high frequencies to the intensity of a fluctuating quantity.

In a simple case where all three functions are power-law functions of time/frequency, e.g. $\left\langle \Delta x^2(\tau) \right\rangle_2 \sim \tau^\alpha$, $G^*(\omega) \sim \omega^\beta$, $\Delta(\omega) \sim \omega^{-\gamma}$, then this formula takes the form of a simple relation: $\alpha + 1 = 2\beta + \gamma$. In an earlier study (Lau, 2003), we applied this method to murine J774A.1 macrophage-like and F9 carcinoma cells, and found that their stress fluctuation spectrum had a nearly power-law form, with an exponent $\gamma \approx 2$. If this equality were exact, this implies the relation $\alpha = 1 + 2\beta$. For the example data in TC7 cells, a value of $\beta_{\text{int}} \approx 0.26$ would imply $\alpha \approx 1.52$, consistent with the observed power-law slope of the 2P-MSD in living TC7 cells, Figure 14(top).

A power-law stress fluctuation spectrum of the form $\Delta(\omega) \sim \omega^{-2}$ has a simple physical explanation. Such a spectrum occurs if the fluctuating stresses are due to discrete, rapid 'step' changes. For example, such steps could be the result of processive molecular motors gradually building stresses on a cytoskeletal element, which is then abruptly released when the motor detaches. Alternatively, in a pre-stressed cytoskeleton, if filaments or cross-link proteins occasionally rupture or release, this would also manifest as a $\Delta(\omega) \sim \omega^{-2}$ spectrum. All that is required is that the rate of the step change is faster relative to our fastest observed timescale (20 milliseconds in the example data), and that the repetition time is longer than the longest lag time observed (a few seconds in the example). In contrast, if we imagine the network receiving short-duration force impulses or 'kicks', the expected spectrum would be $\Delta(\omega) \sim \omega^0$, which is clearly excluded by the data. While a cell's shear modulus does not contain any readily identifiable molecular time-scales, it remains an open question whether molecular time-scales can be found in

the stress fluctuation spectrum. This question might be addressed by extending the frequency range of the measurements.

Conclusions and Future Directions

Two-point microrheology measurements on cultured cells provide a uniquely interpretable, quantitative contribution to the field of cell mechanics. They indicate that cells contain two mechanically distinct intracellular regions. The deeper intracellular interior probed by TPM is found to have a power-law frequency dependence, whose origin remains unknown. This region is remarkably soft, with a modulus that is only a few tens of Pascals. In living cells, this intracellular structure exhibits ATP dependent stress fluctuations with a 'step-like' rather than a 'kick-like' character. This appears compatible with intracellular trafficking, slow cytoskeletal remodeling or both. While these results provide novel insight into the molecular-scale mechanics of cells, they lack molecular specificity. Future studies are needed to 'dissect' the molecular determinants of both the rheology and non-Brownian fluctuations, e.g. through the use of specific cytoskeletal disruptors, motor inhibitors, and genetic manipulations, as well as externally applied stresses (Lee, 2001). Recent such studies from our lab (Van Citters, 2006) show, remarkably, that neither the intracellular rheology nor the stress fluctuations measured by TPM depend on either F-actin or myosin. Such results remind us that the field of cell mechanics is just beginning to understand the generation, propagation and relaxation of intracellular stress.

References and Suggested Reading

Orr, A. W., Helmke, B. P., Blackman, B. R. Schwartz, M. A. (2006). Mechanisms of Mechanotransduction. *Developmental Cell*. 10, 11-20.

Vogel, V. Sheetz, M. (2006). Local Force and Geometry Sensing Regulate Cell Functions. *Nature Reviews Molecular Cell Biology*. 7, 265-275.

Weihs, D., Mason, T. G. Teitell, M. A. (2006). Bio-Microrheology: A Frontier in Microrheology. *Biophys J*.

Waigh, T. (2005). Microrheology of Complex Fluids. *Reports on Progress in Physics*. 68, 685-742.

Kroy, K. (2006). Elasticity, Dynamics and Relaxation in Biopolymer Networks. *Current Opinion in Colloid & Interface Science*. 11, 56-64.

Yamada, S., Wirtz, D. Kuo, S. C. (2000). Mechanics of Living Cells Measured by Laser Tracking Microrheology. *Biophysical Journal*. 78, 1736-1747.

Fabry, B., Maksym, G., Butler, J., Glogauer, M., Navajas, D. Fredberg, J. (2001). Scaling the Microrheology of Living Cells. *Physical Review Letters*. 8714, -.

Alcaraz, J., Buscemi, L., Grabulosa, M., Trepat, X., Fabry, B., Farre, R. Navajas, D. (2003). Microrheology of Human Lung Epithelial Cells Measured by Atomic Force Microscopy. *Biophys J*. 84, 2071-9.

Lenormand, G., Millet, E., Fabry, B., Butler, J. P. Fredberg, J. J. (2004). Linearity and Time-Scale Invariance of the Creep Function in Living Cells. *Journal of the Royal Society Interface*. 1, 91-97.

Desprat, N., Richert, A., Simeon, J. Asnacios, A. (2005). Creep Function of a Single Living Cell. *Biophysical Journal*. 88, 2224-2233.

Hoffman, B., Massiera, G., Van Citters, K. Crocker, J. (2006). The Consensus Mechanics of Cultured Mammalian Cells. *Proceedings of the National Academy of Sciences of the United States of America*. 103, 10259-10264.

Janmey, P. A., Hvidt, S., Lamb, J. Stossel, T. P. (1990). Resemblance of Actin-Binding Protein Actin Gels to Covalently Cross-Linked Networks. *Nature*. 345, 89-92.

Xu, J. Y., Wirtz, D. Pollard, T. D. (1998). Dynamic Cross-Linking by Alpha-Actinin Determines the Mechanical Properties of Actin Filament Networks. *Journal of Biological Chemistry*. 273, 9570-9576.

Gardel, M. L., Nakamura, F., Hartwig, J., Crocker, J. C., Stossel, T. P. Weitz, D. A. (2006). Stress-Dependent Elasticity of Composite Actin Networks as a Model for Cell Behavior. *Physical Review Letters*. 96, -.

Crocker, J. C., Valentine, M. T., Weeks, E. R., Gisler, T., Kaplan, P. D., Yodh, A. G. Weitz, D. A. (2000). Two-Point Microrheology of Inhomogeneous Soft Materials. *Phys Rev Lett*. 85, 888-91.

Lau, A., Hoffman, B., Davies, A., Crocker, J. Lubensky, T. (2003). Microrheology, Stress Fluctuations, and Active Behavior of Living Cells. *Physical Review Letters*. 91, -.

Deng, L. H., Trepat, X., Butler, J. P., Millet, E., Morgan, K. G., Weitz, D. A. Fredberg, J. J. (2006). Fast and Slow Dynamics of the Cytoskeleton. *Nature Materials*. 5, 636-640.

Mason, T. G. Weitz, D. A. (1995). Optical Measurements of Frequency-Dependent Linear Viscoelastic Moduli of Complex Fluids. *Physical Review Letters*. 74, 1250-1253.

Valentine, M. T., Perlman, Z. E., Gardel, M. L., Shin, J. H., Matsudaira, P., Mitchison, T. J. Weitz, D. A. (2004). Colloid Surface Chemistry Critically Affects Multiple Particle Tracking Measurements of Biomaterials. *Biophysical Journal*. 86, 4004-4014.

Crocker, J. C. Grier, D. G. (1996). Methods of Digital Video Microscopy for Colloidal Studies. *Journal of Colloid and Interface Science*. 179, 298-310.

Cheezum, M. K., Walker, W. F. Guilford, W. H. (2001). Quantitative Comparison of Algorithms for Tracking Single Fluorescent Particles. *Biophysical Journal*. 81, 2378-2388.

Savin, T. Doyle, P. S. (2005). Static and Dynamic Errors in Particle Tracking Microrheology. *Biophysical Journal*. 88, 623-638.

Martin, D. S., Forstner, M. B. Kas, J. A. (2002). Apparent Subdiffusion Inherent to Single Particle Tracking. *Biophysical Journal*. 83, 2109-2117.

Caspi, A., Granek, R. Elbaum, M. (2000). Enhanced Diffusion in Active Intracellular Transport. *Physical Review Letters*. 85, 5655-5658.

Tseng, Y., Kole, T. P. Wirtz, D. (2002). Micromechanical Mapping of Live Cells by Multiple-Particle-Tracking Microrheology. *Biophysical Journal*. 83, 3162-3176.

Puig-de-Morales, M., Millet, E., Fabry, B., Navajas, D., Wang, N., Butler, J. P. Fredberg, J. J. (2004). Cytoskeletal Mechanics in Adherent Human Airway Smooth Muscle Cells: Probe Specificity and Scaling of Protein-Protein Dynamics. *American Journal of Physiology-Cell Physiology*. 287, C643-C654.

Levine, A. J. Lubensky, T. C. (2001). Response Function of a Sphere in a Viscoelastic Two-Fluid Medium. *Physical Review E*. 6304, -.

Mason, T. G. (2000). Estimating the Viscoelastic Moduli of Complex Fluids Using the Generalized Stokes-Einstein Equation. *Rheologica Acta*. 39, 371-378.

Dasgupta, B. R., Tee, S. Y., Crocker, J. C., Frisken, B. J. Weitz, D. A. (2002). Microrheology of Polyethylene Oxide Using Diffusing Wave Spectroscopy and Single Scattering. *Phys Rev E Stat Nonlin Soft Matter Phys*. 65, 051505.

Lee, J.C-M. and Discher, D.E. (2001) Deformation-enhanced fluctuations in the red cell skeleton with theoretical relations to elasticity, connectivity, and spectrin unfolding. *Biophysical Journal* 81, 3178-3192.

Van Citters, K. M., Hoffman, B. D., Massiera, G. Crocker, J. C. (2006). The Role of F-Actin and Myosin in Epithelial Cell Rheology. *Biophys J*.