Chapter 6

Experimental Methods

This chapter describes the two different techniques utilized in this dissertation for observing the behavior of diffusing photons in turbid media. In the first, I measure the average intensity and phase of light generated by an intensity modulated source and investigate the properties of the photon density as it diffuses through highly scattering media. Media with and without spatially varying optical properties are considered. In the second, I indirectly measure the temporal fluctuations of diffuse electric fields emanating from highly scattering media. From the time scale of the temporal fluctuations it is possible to derive information about the dynamical properties of the scattering media. Media with spatially varying dynamical and optical properties are considered.

6.1   Measuring Diffuse Photon Density Waves

When photons are introduced into a highly scattering medium, they undergo a random walk. Some photons will scatter only a few times before exiting the system while others will experience a thousand or more scattering events. The average number of scatter events for an infinite homogeneous medium is estimated as $(r/l^*)^2$ where r is the source-detector separation and l^* is the photon random walk step. When the source of photons is intensity modulated at a frequency f, then the photon fluence in the medium will also modulate with a frequency f. This modulation in the photon fluence can be thought as a traveling wave propagating outwards from the source and

is referred to as a diffuse photon density wave (DPDW) [3, 21, 22, 23, 68].

6.1.1 Experimental

The experimental apparatus for generating and detecting these DPDW's is depicted in fig. 6.1. The random medium used to multiply scatter the photons is usually a liquid called Intralipid. Intralipid is an emulsion of water and soy bean fat used for intravenous feeding of hospitalized patients. It is a polydisperse suspension of particles with an average diameter of $\sim$ 0.4 μm, but a relatively wide range of sizes (i.e. from $\sim$ 0.1 μm to $\sim$ 1.1 μm). The optical properties of Intralipid have been described by others [72, 73, 74]. We purchase our Intralipid from the supply room at the University of Pennsylvania. Other researchers obtain their Intralipid for free by accepting the spoils that are discarded by hospitals. By changing the solution concentration, it is possible to vary the light diffusion coefficient of the medium. The photon random walk step l^*, often referred to as the photon transport mean free path, is about 0.2 cm at 800 nm for a solution of 0.5% Intralipid by volume.

Polystyrene microspheres in suspension are often used as a turbid medium. They can be purchased from Seradyne and Bang Labs. These microspheres are monodisperse and well characterized so that the scattering lengths can be calculated. We prefer Intralipid over polystyrene because of the price; $20 per pint compared to $10 per milliliter. We generally measure the optical properties using our optical technique and therefore don't require the well characterized microspheres. The disadvantage with Intralipid is that it spoils after a few days, much like milk, and it is a breeding ground for bacteria. Also since the scatterer is fat, the scattering properties are temperature dependent.

For the experiments described in chapters 2 and 3, a large tank (30 cm x 30 cm x 60 cm) is typically filled with an Intralipid solution. Source and detector optical fibers ($\sim$4 mm in diameter) are immersed in the solution at the same height above the tank floor. The fibers are pointed in orthogonal directions to minimize gradient systematics; specifically the detector fiber is perpendicular to the radial vector from
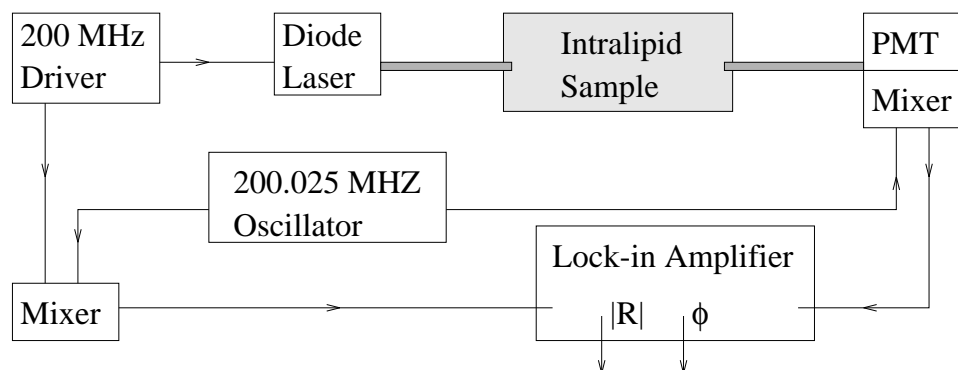
Figure 6.1: A known concentration of Intralipid solution fills a glass tank. The source is a diode laser that is amplitude modulated by a 200 MHz driver. Light is delivered into the sample through a source fiber, and picked up by a movable detector fiber. The fibers are pointed in orthogonal directions to minimize gradient systematics. The signal at the photomultiplier tube (PMT) is heterodyned down to 25 kHz (by modulating the second PMT dynode at 200.025 MHz), and then fed into a lock-in amplifier (Stanford, model #SRS-530). The 25 kHz lock-in reference signal is derived from the 200 MHz driver by standard mixing techniques. The two-phase lock-in amplifier provides amplitude ($|\mathbf{R}|$) and phase (ϕ) output signals.

the source fiber.

The source light is derived from a 3 mW diode laser operating at 816 nm. The light is coupled into the source fiber by butting the fiber bundle up to the surface of the laser diode. The source fiber position in the tank is fixed.

Diffuse light is collected with the detector fiber which can be positioned anywhere in a plane parallel to the bottom of the tank by using a home-built translation stage driven by stepper motors (described below). The fiber delivers the collected light to a photo-multiplier tube (PMT). We use an R928 PMT from Hamamatsu. The PMT is modified to facilitate the phase and amplitude measurements. Basically the voltage at the second dynode is modulated at 200.025 MHz. Modified this way, the PMT acts both as a detector and a mixer. The PMT signal thus has a 400.025 MHz modulation and a 25 kHz modulation. We use a notch filter to pick out the 25 kHz signal. A 25 kHz reference is generated from the 200 MHz and 200.025 MHz oscillators by a mixer and low pass filter. The low frequency signals are then measured using
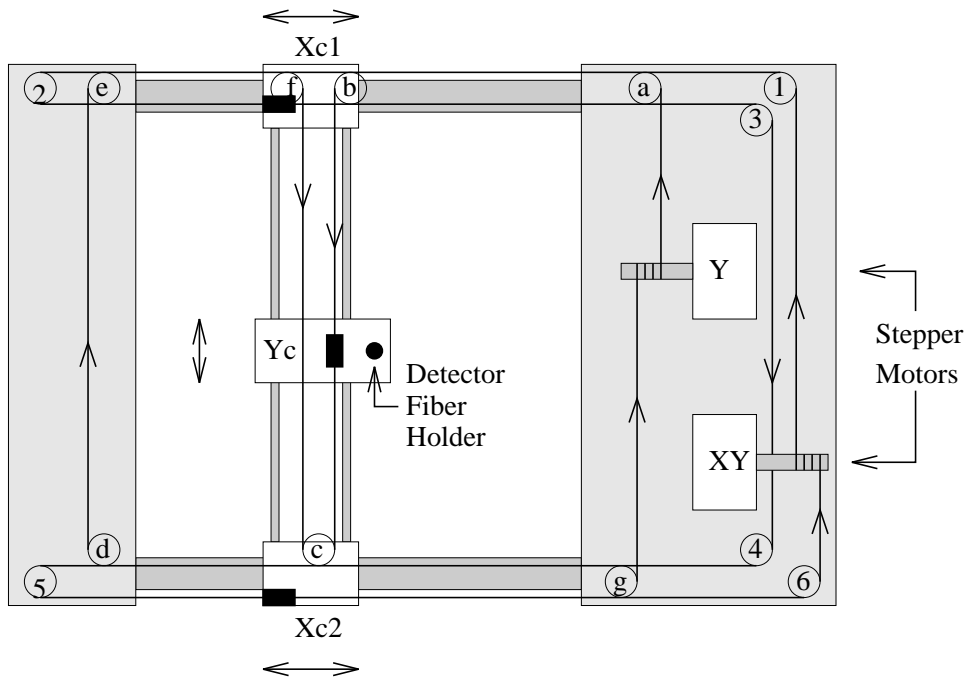
Figure 6.2: Home-made stepper motor-driven translation stage.

a lock-in amplifier (Stanford, model #SRS-530). The phase shift and AC amplitude of the detected light relative to the source are determined by the lock-in amplifier. The dynamic range of this apparatus is about 10^5. Since the signal amplitude decays by more than $\exp(-2\pi)$ in one DPDW wavelength, the range of the experiments is limited to slightly more than one DPDW wavelength. Nevertheless it is possible to clearly distinguish the essential physical phenomena of diffuse photon density waves.

The stepper motor-driven translation stage is diagrammed in fig. 6.2. The x-axis carriages (Xc1 and Xc2) and the y-axis carriage (Yc) are moved by a wire and pulley system. The wire attached to the y-axis carriage (at the solid block) is drawn by the Y stepper motor. The wire attached to the x-axis carriages (at the solid blocks) is drawn by the XY stepper motor. Because of the wire configuration, the XY motor actually moves the Xc1, Xc2, and Yc carriages. The Y motor only moves the Yc carriage. The Y wire must have the designed configuration in order to permit full motion with a wire of constant length. From the XY motor the wire wraps the pulleys in the following order: 1, 2, 3, 4, 5, and 6. From the Y motor the wire wraps the
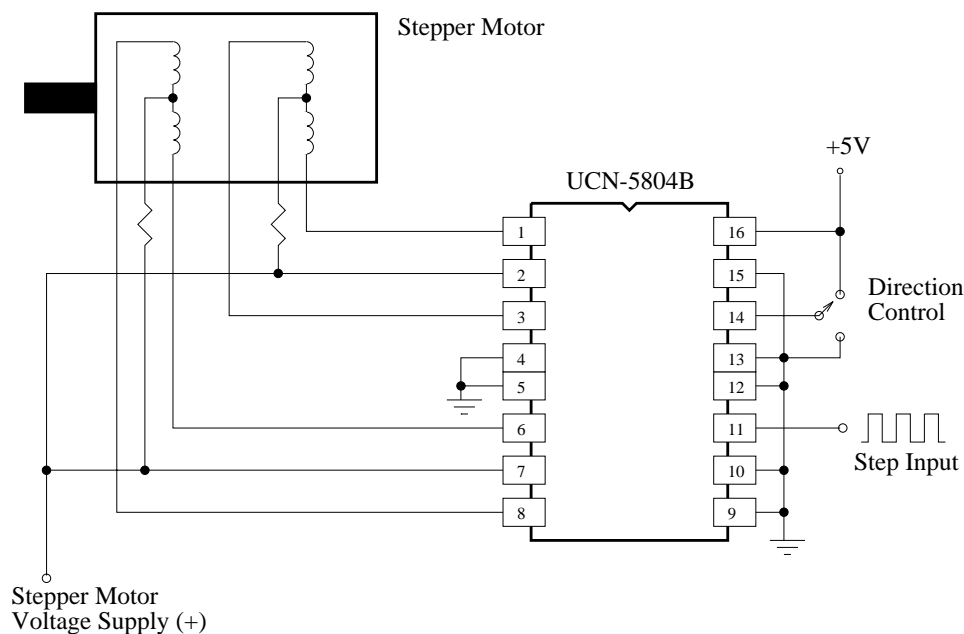
Figure 6.3: Circuit diagram for stepper motor controller. Pin 16 is the logic supply. Pin 15 is the output enable (when set high the chip turns off). Pin 14 is for direction control. Pin 11 is for step control. The chip number is UCN-5804B.

pulleys in the order: a, b, c, d, e, f, c, g. There are two pulleys in position c. All pulleys are on appropriate levels to keep wires from intersecting. The detector fiber is held in place by the Yc carriage.

We use 4-phase stepper motors with 200 steps per revolution. The motors are driven using a specially designed, inexpensive integrated circuit (Herbach and Rademan 1-800-848-8001, H&R # Q5029 $9.95). A schematic of a simple circuit using this chip for controlling the motor is given in fig. 6.3. The chip requires TTL pulses to advance the motor one step and to control the step direction. These pulses are provided by the parallel port (LPT1) of an IBM clone PC. We use Q-basic to directly assigning values to the parallel port. The base address for LPT1 is 888 (decimal) and the Q-basic command for assigning a value to address 888 is `out 888,n` where `n` is a value between 0 and 255 (see table 6.1) and is the decimal value of the binary word represented at the base address by pins 2 through 9. For example, if you want pins 2 and 3 to be high and 4 through 9 to be low then `n` would be 00000011 in binary which

Table 6.1: Configuration of the parallel port on an IBM-clone PC. The numbers correspond to pins on the 25 pin D-shell parallel port connector. Logic is inverted for numbers with a bar. x means that the bit is not used. The base address for port LPT1 is 888 (decimal). Pins 18-25 are logic ground.

Address	MSB			$\leftarrow$				LSB	function
base	9	8	7	6	5	4	3	2	output
base + 1	$\overline{11}$	10	12	13	14	x	x	x	input
base + 2	x	x	x	0	$\overline{17}$	16	$\overline{14}$	$\overline{1}$	in/out

is 3 in decimal.

Experiments are performed in several different geometries and are discussed in chapter 2.

In section 2.2 measurements of the propagation of DPDW's in a homogeneous, approximately infinite medium are described. For these measurements, the phase shift and AC amplitude of the detected diffusive wave is measured with respect to the source at each point on a 0.5 cm square planar grid. The measurements are made sufficiently far from the edges of the tank such that the medium could be approximated as being infinite.

In section 2.4 measurements of the refraction and diffraction of DPDW's are described. Experiments include measurements of the refraction by a planar interface between different concentrations of Intralipid and the refraction and diffraction by spherical objects with different optical properties than the surrounding medium. The planar interface is created by a 2 mm thick sheet of acrylic which separated the two different solutions of Intralipid. The spherical objects are created by containing different concentrations of Intralipid and black ink in a 4.0 cm diameter glass bulb. The wall of the glass bulb is approximately 3 mm.

In section 2.5 measurements of the scattering of DPDW's from spherical objects is described. Perfectly absorbing spheres and spheres with different scattering coef-

ficients were measured. The perfectly absorbing spheres are wooden spheres from a hobby shop painted matte black. The scattering spheres are made by suspending titanium dioxide in casting resin and letting it cure in hemispherical molds. Titanium dioxide/resin phantoms have been described in detail by Firbank *et al.* [137, 138] and are briefly discussed in section 6.3.

6.1.2 Monte Carlo

In addition to experimental measurements of DPDW's, I use Monte Carlo simulations of migrating photons to analyze certain properties with a signal-to-noise ratio that is difficult to achieve in the laboratory. Simulations are performed for a point source in an infinite, homogeneous medium. Results are obtained for different source modulation frequencies and medium absorption coefficients. These simulations permit a detailed examination of the breakdown of the photon diffusion approximation and the validity of higher order transport approximations for large photon absorption coefficient and modulation frequencies (see section 2.7).

The Monte Carlo code that I developed comprises only 500 lines of C code, this includes components for spherical objects (see appendix C). The basic idea is to launch N photons into the medium at time $t = 0$ and histogram partial photon flux in radial and temporal channels. For my simulations N is typically 1 million to 10 million and the code takes $\sim$2 hours to run on a Sun Sparc 10 Model 512 50MHz processor or a 75 MHz Pentium. Once the temporal response to a pulse is found, a fast Fourier transform of the result into the frequency domain is performed for analyzing the properties of diffuse photon density waves.

There are many techniques for propagating and histogramming photons within Monte Carlo calculations [67, 139, 140]. To keep the code simple I mimic the physical process as closely as possible instead of relying on reduction techniques that purport to increase statistics while reducing computation (see [67, 139, 140] for discussions of different reduction techniques). The algorithm is charted in fig. 6.4. To propagate a photon from one interaction event to the next, the program calculates a scattering
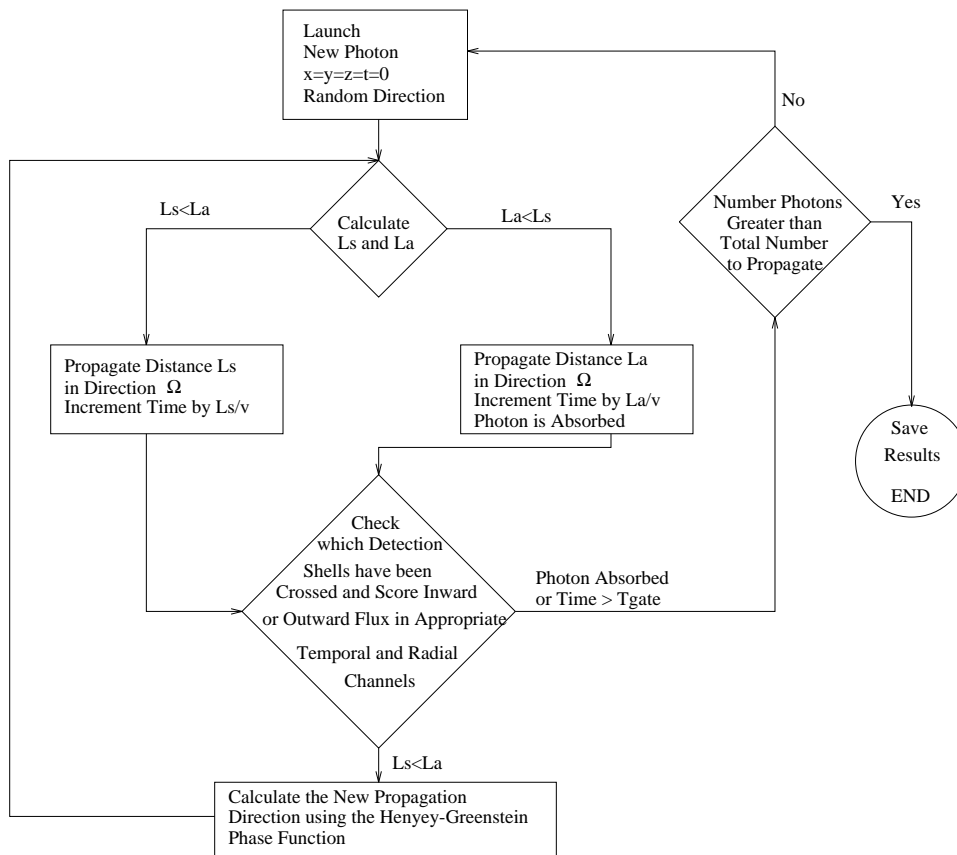
Figure 6.4: Flow chart for Monte Carlo simulations for a point source in infinite homogeneous medium.

length and an absorption length based on the exponential distributions derived from the scattering and absorption coefficients respectively. If the absorption length is shorter than the scattering length, then the photon is propagated the absorption length, scored if necessary, terminated, and a new photon is launched at time $t = 0$ from the source position. If the scattering length is shorter, then the photon is propagated the scattering length, scored if necessary, the scattering angle is calculated based on the Henyey-Greenstein phase function [139, 140, 141], and then a new scattering length and absorption length are calculated. Photon propagation continues until the photon is absorbed, escapes, or the time exceeds a maximum set by T_{gate}.

To exploit the spherical symmetry of the problem (the source is isotropic), spherical shell detectors are centered on the point source and the crossing of photons across each shell is scored in the appropriate temporal and radial channel. The width of each temporal channel is 20 ps. Inward and outward crossings are scored separately so that the Monte Carlo simulation can report the radial components of the photon flux. In this way we can obtain the photon fluence, $\Phi(r)$, and the net photon flux, $-D\nabla\Phi(r)$, from the data. For my analyses I use the photon fluence.

6.2 Photon Correlation Spectroscopy Techniques

Photons scattered by moving particles have their frequency Doppler-shifted by an amount proportional to the particle's speed, the photon's wavenumber, and the scattering angle. The frequency shifts are generally a very small fraction of the absolute frequency, typically ranging from 10^{-9} to 10^{-12}. These relatively small shifts are difficult to measure directly. Instead they are usually determined indirectly by measuring the beating of different frequencies as revealed in the fluctuating intensities of a single coherence area (i.e. speckle) of the scattering light. These fluctuations can be analyzed by looking at the power spectrum or temporal autocorrelation function of the fluctuations. I measure the temporal intensity autocorrelation function of the fluctuating speckles. This method is prefered over measuring the power spectrum for

two reasons. First, the correlation function is the quantity which is derived in the correlation diffusion equation (see chapter 4). Second, using digital correlators and photon counting techniques, it is possible to analyze smaller signals.

There is one major advantage of direct measurements. Indirect measurements of the Doppler broadening of the laser linewidth requires that single (or only a few) coherence areas of scatter light be detected. For systems which multiply scatter light, these coherence areas are on the order of ~ 1 μm^2. Small aperture light collectors are thus necessary, resulting in the collection of small numbers of photons. This is not the case for systems which scatter light no more than once since then the coherence area is then given by the laser beam size and coherence length. For direct measurements of Doppler broadening it is not necessary to collect light from a single coherence area and thus the number of photons collected can be increased by at least a factor of 10^3. Low signal-to-noise ratio applications would benefit from the development of better filters for directly measuring these relatively small Dopper shifts.

6.2.1   Experimental

A schematic of the experimental apparatus for measuring the temporal intensity auto-correlation function (discussed in chapters 4 and 5) of a speckle's intensity fluctuations is presented in fig. 6.5. For the initial experiments, the 514 nm line of an Argon-ion laser (operating with an etalon) is used because of its long coherence length and strong power. In later human subject experiments, the 800 nm line of a Ti-Sapphire laser is used to coincide with the isosbestic point in the absorption spectra of oxy- and deoxy-hemoglobin. For the animal experiments, the 632 nm line of a HeNe laser is used because of its portability. In the future, it will be desirable to utilize laser diodes with stable, single longitudinal mode operation.

The laser beam is coupled into a multimode fiber using a fiber launch system purchased from Thor Labs (part # MDT612). Fibers with core diameters ranging from 50 to 200 μm and various numerical aperatures (NA) are used. Generally, for large source-detector separations the diameter and NA are not critical, although
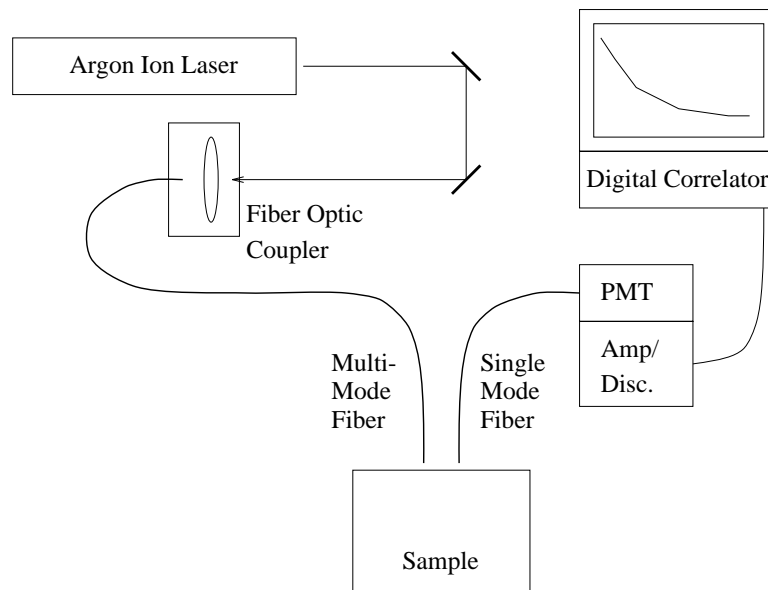
Figure 6.5: Experimental setup for photon correlation spectroscopy.

coupling with the laser beam is easier with large diameters and NA's. For the smaller source-detector separations used on the human subject and animal trials, it is best to use a small NA. This minimizes the divergence of the beam from the output of the fiber to the sample, thus reducing the possibility of detecting light that has reflected at the surface of the sample. For these measurements I use 200 μm core diameter fiber with an NA of 0.16 (Thor Labs FG-200-LAT).

Measurements are made on various samples with many source-detector separations. The source-detector positions are controlled by repositioning the source and detector fibers. Diffuse back reflection is measured from turbid, homogeneous, and solid slabs with either a spherical cavity containing a turbid colloid or a cylindrical vein through which a colloid is pumped. Transmission and reflection measurements are made on solid cylinders with spherical cavities. Human and animal subjects are also employed. These various samples are described in more detail in the appropriate chapters.

Single speckles of the diffuse light emanating from the sample are collected with a single mode fiber. I use different single mode fibers, all from Thor Labs. Basically it is best to use a single mode fiber designed for the wavelength of light to be detected.

Single mode fibers designed for short wavelengths will not propagate the desired longer wavelength thus killing the signal. Single mode fibers designed for longer wavelengths will pass many more modes (speckles) increasing the photon count rate but decreasing the relative magnitude of the fluctuations resulting in a decreased signal-to-noise ratio. This trade-off can be exploited if the photon count rate with the appropriate single mode fiber is near the dark current. The dark current is a constant noise source independent of the photon count rate and thus the signal-to-noise ratio decreases more than expected as the photon count rate approaches the dark current. In this case a better signal-to-noise ratio can be achieved by using a fiber for a longer photon wavelength which increases the photon count rate to a level such that the dark current is insignificant. To be more quantitative we must look at the relative trade-off between a significant dark current associated with the correct single mode fiber and the reduced relative magnitude of the fluctuations associated with using a single mode fiber for a longer wavelength.

A single mode fiber does not actually collect light from a single speckle but from a single spatial mode. Actually the single mode fiber propagates both polarizations so that it is really a dual mode fiber. The difference between observing a speckle and a mode is best demonstrated by comparing the measured correlation function for different distances between the single mode fiber and the sample. Because of the non-zero NA of the single mode fibers, the area from which light is collected increases with the distance between the sample and collecting fiber. The number of observed speckles increases with the collection area and thus the observed intensity fluctuations should decrease relative to the average intensity because of the averaging of uncorrelated, fluctuating speckles. In addition the average intensity should increase. However, the fluctuations are observed not to decrease and the average intensity remains constant. This result, perhaps surprising, arises because a single mode fiber projects the collected light onto a single spatial mode which by definition is coherent. The use of single mode fibers for observing speckles has been discussed by Ricka [142] and Van Keuren *et al.* [143].

The collected light is delivered to a photon counting photo-multiplier tube (Hamamatsu HC120). The dark count is less than 10 counts per second at room temperature. Smaller dark counts are achieved by cooling the PMT to 0°C. After-pulsing can occur for up to 100 ns and the response of the PMT is linear up to 200,000 counts per second. The signal from the PMT is amplified and then discriminated using an amplifier/discriminator package supplied by Brookhaven Instruments Corporation. The TTL signal from the discriminator goes to the digital correlator card housed in a 486 computer. The correlator is from Brookhaven Instruments Corporation (BI9000). To measure correlation functions faster than 200 ns it is necessary to cross-correlate the signal from two PMT's. This is best achieved by using a 50-50 fiber optic beam divider to split the beam between two PMT's and operate the digital correlator in cross-correlation mode.

Since the decay of the correlation function depends on the position of the source and detector relative to the dynamical regions of the sample, it is necessary to optimize the range of correlation times for which data is collected. It is important to choose the minimum correlation time to capture the early time decay. However it is not good to set this minimum arbitrarily low since smaller minimum correlation times reduce the signal-to-noise ratio. The maximum correlation time must be chosen so that the full decay, and thus the baseline, of the correlation function is recorded.

The samples used in my experiments are non-ergodic (a time average is not equal to an ensemble average). As described in the theory section on photon correlation spectroscopy and ergodicity (section 4.2.4), care must be taken with non-ergodic samples in order to obtain repeatable measurements. Basically, for non-ergodic samples each speckle has a constant component and a fluctuating component. The constant component is comprised of photons which have not scattered from any moving particles. The fluctuating component arises from the photons which have scattered from at least one moving particle. If we fix the source and compare speckles in a localized region some distance from the source, then we will see that the fluctuating component is the same from speckle to speckle (the ensemble average equals the time average) but
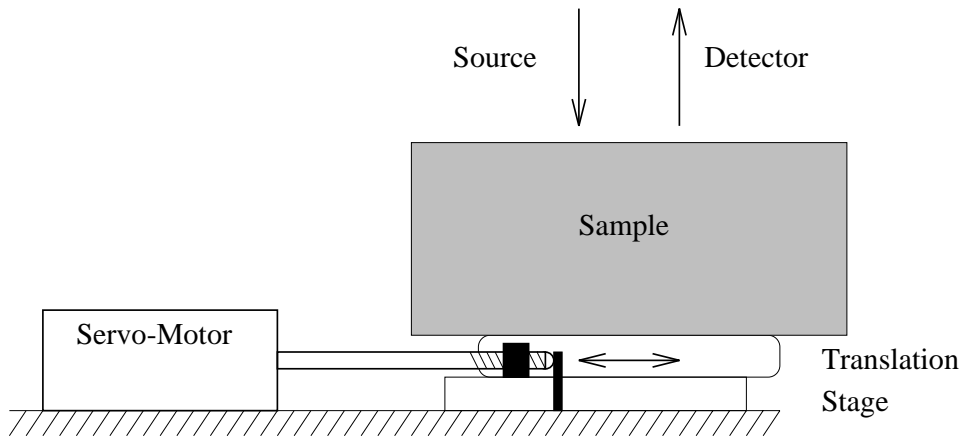
Figure 6.6: Setup for doing the ensemble average.

that the constant component differs (the ensemble average does not equal the time average). This variation in the constant component causes the measured temporal intensity correlation function to vary from speckle to speckle.

In order to measure the proper correlation function it is necessary to ensemble average the signal. I do this by moving the sample (or sometimes the source and detector) during the integration of a correlation function. In this way I measure an ensemble of constant components and thus obtain the desired ensemble average of the speckles' intensity. Moving of the sample affects the correlation function by increasing the observed intensity fluctuation. If the motion is slow enough then the decay of the correlation function due to moving the sample occurs on a time scale that is long compared to the time scale of interest. However, if we move the sample too fast then the ensemble averaging decay overlaps with the decay due to the internal dynamics.

A schematic of the system I use for moving the sample is shown in fig. 6.6. The sample is placed on a translation stage which is driven by a servo-motor. It is important to use a motor which moves smoothly. Stepper motors do not work well because their motion is jerky. This jerky motion results in fast intensity fluctuations which obscure the intensity fluctuations due to the internal dynamics. The speed of the servo-motor is set to move the sample approximately 50 μm s^{-1}. It is important to ensemble average the signal in a localized area. If we average over too large of an area

then source-detector positions relative to objects in the sample are not well-defined and accurate comparisons with theory cannot be made. I use limit switches on the servo-motor to make it oscillate along a 400 μm path.

6.2.2 Monte Carlo

When the accuracy of my experimental results was questionable or I did not have the experimental data, I used the results from Monte Carlo simulations to check the accuracy of the correlation diffusion equation (see section 4.4). In many cases these simulations provide a signal-to-noise ratio that is difficult to achieve in the laboratory and therefore permit a more accurate test of the validity of the correlation diffusion theory. I run simulations for a point source in an infinite medium. Media with different dynamical properties are considered. I first gather data for a system with spatially uniform Brownian motion. The correlation diffusion equation is known to be valid for this system at short correlation times. Therefore, these first simulations worked as a test for the Monte Carlo code at early correlation times and to demonstrate the breakdown of the diffusion equation at long correlation times.

Next, I ran simulations for a homogeneous solid system containing a spherical region with scatterers undergoing Brownian motion. The optical properties of the spherical region are varied with respect to the background optical properties. These simulations are necessary to unravel systematic discrepancies between experimental data and correlation diffusion theory.

Finally, simulations are executed for a homogeneous system with different volume fractions of random flow. This is a model of tissue blood flow. All simulations are performed assuming isotropic scattering (i.e. $g = 0$).

The theoretical details pertinent to the Monte Carlo simulation are reviewed here. A complete discussion of deriving temporal electric field correlation functions $(g_1(\tau))$ ie given in section 4.2. The correlation function of light that scatters once from a

dilute suspension of noninteracting uncorrelated particles is

$$g_1^s(\tau) = \exp\left(-\frac{1}{6}q^2\left\langle\Delta r^2(\tau)\right\rangle\right) \ , \tag{6.1}$$

where $\mathbf{q} = \mathbf{k_{out}} - \mathbf{k_{in}}$ is the momentum transfer imparted by the scattering event (see fig. 4.2) and $\left\langle\Delta r^2(\tau)\right\rangle$ is the mean square displacement of the scattering particles in time τ. The magnitude of the momentum transfer is $q = 2k_o\sin(\theta/2)$. When photons are multiply scattered by non-interacting uncorrelated particles, the correlation function is computed for a given photon path α with n uncorrelated scattering events as

$$g_1^{(\alpha)}(\tau) = \exp\left(-\frac{1}{6}\sum_{j=1}^n q_{\alpha,j}^2\left\langle\Delta r^2(\tau)\right\rangle\right) \ . \tag{6.2}$$

$q_{\alpha,j}$ is the momentum transfer experienced along path α at scattering site j.

The general procedure for considering multiple paths is to first relate the total dimensionless momentum transfer $Y = \sum_{j=1}^n q_{\alpha,j}^2/2k_o^2 = \sum_{j=1}^n(1 - \cos\theta_{\alpha,j})$ to the dimensionless path length $S = s/l^*$. Here s is the length of the photon trajectory through the sample and l^* is the photon random walk step length. For large n, Y is accurately approximated by the average over the scattering form factor and thus

$$Y \approx n\left\langle 1 - \cos\theta\right\rangle = \frac{nl}{l^*} = \frac{s}{l^*} = S \ . \tag{6.3}$$

Here $\left\langle ...\right\rangle$ denotes the average over the scattering form factor and l is the photon scattering length which equals the photon random walk step length when the scattering is isotropic. Next the total correlation function is obtained from the weighted average of eq. (6.2) over the distribution of path lengths, i.e.

$$g_1(\tau) = \int_0^\infty P(S)\exp\left(k_o^2\left\langle\Delta r^2(\tau)\right\rangle S/3\right) \ . \tag{6.4}$$

Although $P(S)$ can be determined using Monte Carlo simulations, it is usually found with the help of the photon diffusion equation.

This procedure has built into it two assumptions, the relation between Y and S and that $P(S)$ is accurately given by the photon diffusion equation. For the purposes of the Monte Carlo simulations it is desirable to take a different approach that doesn't

make these two assumptions. As suggested by Middleton and Fisher [120] and Durian [119], the total correlation function can be obtained from a weighted average over the total dimensionless momentum transfer experienced by all photon trajectories, i.e.

$$g_1(\tau) = \int_0^\infty P(Y) \exp\left(\frac{1}{3} Y k_o^2 \left\langle \Delta r^2(\tau) \right\rangle\right) \ . \tag{6.5}$$

There are no assumptions in this formulation other than the standard noninteracting uncorrelated particles assumption. The drawback is that $P(Y)$ cannot be analytically calculated. However, Monte Carlo simulations provide a simple numerical approach to finding $P(Y)$ for different geometries. Such Monte Carlo simulations are described by Middleton and Fisher [120], Durian [119], and Koelink *et al.* [144].

The Monte Carlo simulation follows the trajectory of a photon using the algorithm described in section 6.1.2 with the addition that the dimensionless momentum transfer Y is incremented in dynamic regions. When the photon reaches a detector, the Y associated with that photon is scored in a $P(Y)$ histogram. After sufficient statistics have been accumulated for $P(Y)$ (typically 1 million to 10 million photons) $g_1(\tau)$ can then be calculated using eq. (6.5).

6.3   Making Resin Phantoms

Intralipid and India ink phantoms are often insufficient. Either a more permanent phantom is required or a system with spatially varying optical properties is desired. These properties are not available with Intralipid phantoms since Intralipid spoils after a few days (very much like milk) and, since it is a liquid, it is necessary to use a physical barrier that perturbs the light field to separate regions with different optical properties. Both of these limitations are overcome with polyester resin phantoms.

Polyester resin phantoms are solid and inert. The scattering and absorption coefficients of the solid can be controlled by mixing different amounts of scatterer (we use titanium-oxide) and dye to the liquid resin before the curing process. The resin can be formed into any shape using an appropriate mold. Optical heterogeneities can be

introduced before or after the curing process. The recipe for making a resin phantom follows.

6.3.1   Recipe for Resin Phantoms

Ingredients

- Polyester Resin (a.k.a. Castin Resin from Creative Wholesale (404)474-2110 $25 per gallon)

- Catalyst (Creative Wholesale (404)474-2110 $2)

- Titanium Oxide Powder (TiO_2 from Sigma Chemical Co. (800)325-3010 product # T8141 $44.20 per kg)

- Dye (ICI plc, product name: Pro jet 900 NP)

- 100% Ethanol

Equipment

- Scale with 1 mg accuracy

- Test tube with a cap

- Sonicator

- Ventilation hood

- Disposable container for mixing resin, titanium oxide, and catalyst

- Molds

Procedure

Warning: all work with polyester resin should be done in a ventilation hood.

- Find an appropriate mold (molds are discussed below).

- Determine the desired value of μ_s' and μ_a for the phantom.

- Weigh 1 mg of TiO_2 powder for every 1 ml of polyester resin to obtain a μ_s' of ~ 8 cm^{-1}. Put the weighed quantity into a test tube that can be capped.

- Weigh 20 mg of dye for every liter of resin to obtain a μ_a of ~ 0.1 cm^{-1}. Put the weighed quantity with the TiO$_2$.

- Add 100% ethanol to the TiO$_2$ and dye. Use a volume of ethanol that is 1% to 2% of the volume of resin to be used.

- Stir and then sonicate the mixture. Stir occasionally to re-suspend any clumps that settle to bottom. The mixture needs to be sonicated until the TiO$_2$ is completely suspended as well as the dye. This takes ~ 15 minutes.

- Pour the desired amount of resin into the mold or a mixing container if the mold is not appropriate for mixing.

- Pour the ethanol, TiO$_2$, and dye mixture into the resin.

- Stir until uniform.

- Add catalyst to the resin mixture and stir well. For resin volumes greater than 500 ml use ~ 60 drops of catalyst per liter of resin. For smaller volumes you can use more catalyst, e.g. use 10-20 drops for 100 ml of resin.

- If the resin mixture is not yet in the mold then pour it into the mold.

- Let the sample cure for 12-24 hours in a ventilation hood.

- Remove solid resin sample from mold.

- Machine if desired.

6.3.2   Molds

It is important to choose the right mold for forming the resin sample. The mold has to be made of a material that will not react with resin as it is catalyzed. Furthermore, the mold has to have smooth surfaces so as to release the resin once it has hardened. In our lab, we have found that Tupperware containers, PVC, and teflon make good molds.

We use Tupperware when we want a semi-infinite medium with a flat interface and the shape of the sides and bottom is not important. The flat interface comes from the surface of the resin that is exposed to air. We use sheets of PVC to make rectangular molds when the shape of the edges is important. To make cylindrical molds we use PVC tubing. When making large samples, i.e. volumes greater than 1.5 liters, the resin tends to crack during the curing process. We have found that the large diameter PVC tubing reduces the risk of cracking over rectangular molds. Finally, we use teflon blocks for hemispherical molds to make spheres. The hemispherical molds are drilled in the teflon using ball end mills.

Sometimes it is desirable to make resin phantoms with spatially varying optical properties. There are a couple of techniques for doing so. One approach is to place a previously made solid resin object with different optical properties in the mold for the background medium before curing the background medium. This works well as long as care is taken to stabilize the inserted object and to remove any air bubbles that remain at the liquid-solid interface. Phantoms with capillary networks are easy to make by using teflon tubing in the mold to define the network. The tubes are then easy to remove after the resin has cured provided they don't form too many loops that greatly increase the resistance to being removed. Another approach is to machine the homogeneous phantom after the curing and then introduce media with different optical properties.

6.3.3   Machining

Hardened resin is a soft material that machines well in the lathe, mill, and drill press. Care must be taken with the band saw to prevent binding of the material with the blade. The age-old advice of cutting away small amounts of material, backing up, and repeating is sufficient for successfully using the band saw on hardened resin. I often use the lathe to flatten the face of a resin block or cylinder. I recommend removing no more than 0.020 inches of material per cut. Removing more material will result in divots. Some people prefer to use the mill to make a flat surface, but I find that the

lathe gives a smoother finish. In some instances, however, the sample is too large for the lathe and the mill must be used.

I next describe the technique for making resin samples with spherical cavities (see section 4.3.1). A block with cavity is made by first squaring off the faces of the block. Then I cut 2.2 cm of material off the top of the block, flatten the newly cut surfaces so that the smaller piece is 2.0 cm thick, use a 1 inch ball end mill to cut a hemisphere in the center of each of the two newly flattened surfaces, and drill a 3 mm filling hole through the thicker block to the hemisphere. Next I glue the two pieces together with resin and catalyst. It is necessary to use resin as the glue so as to match the indices of refraction between the two pieces. With other glues the interface is clearly visible indicating that light propagation is severly perturbed by the interface, whereas when using resin as the glue the interface disappears. Note that some glue will seep into the cavity. For this reason I let the glue cure so that gravity will cause seeping resin to flow into the filling hole. I can then remove this material by re-drilling the filling hole after the glue has cured.

A cylinder with a cavity is made in a similar way. First the cylinder is made smooth in the lathe. Then it is cut in half and the two faces are flattened in the lathe. A $\frac{1}{2}$ inch ball end mill is used to cut hemispheres into each piece, a 3 mm filling hole is drilled parallel to the axis of the cylinder, and the two pieces are glued together with resin.

6.3.4   Warnings

There are several mistakes that one can make that will ruin the sample. Here I list the ones that I know personally.

Use 100% ethanol to suspend the TiO_2 and dye. Water does not mix with resin, and TiO_2 has a high affinity for water. Therefore if the ethanol contains water, then the TiO_2 will preferentially suspend in the water and separate from the resin while the ethanol easily mixes with the resin. This suggests an interesting demonstration. Prepare identical samples except in one case use 95% ethanol and in the other use 100%.

Another interesting demonstration is to mix just water and resin. This produces a highly scattering emulsion that is stable for days if the volume fraction of water is small enough. Unfortunately this emulsion cannot be hardened.

Incomplete sonication of the TiO_2 and dye will result in an inhomogeneous sample. This warning holds especially for the TiO_2. The sonication breaks large aggregates of TiO_2 into submicron crystallites. Incomplete sonication will leave behind large aggregates which, when mixed with the resin, will quickly settle. In addition, after complete sonication, it is important not to shake or stir the ethanol/TiO_2 mixture since this often leads to aggregation. If the aggregation happens before you have added the TiO_2 to the resin then simply re-sonicate. Note that even under the best of circumstances that some settling of TiO_2 will occur during the curing process. I have not seen deviations in the scattering coefficient from the top of the sample to the bottom of the sample any greater than 10%. This deviation can be minimized by stirring the resin/TiO_2 and catalyst approximately 1 hour after the initial mixing of the resin and catalyst. The quoted optical properties of TiO_2 and dye in resin are only approximate. Therefore the properties should be measured optically once the sample has hardened.

Note that suspending the TiO_2 in the alcohol can be difficult. Add more alcohol if necessary but not more than 5% the volume of resin to be used. If too much ethanol is used then the resin will not harden but will instead become spongy.

We have had difficulty making polyester resin samples that were larger than 1.5 liters because of sample cracking during the curing process. This cracking occurs because the polyester resin has a large exothermic reaction with the catalyst which can result in significant volume changes. Firbank *et al.* have suggested using an epoxy resin which produces less heat during the curing process [138].

Finally, we have observed that x-rays will change the absorption coefficient of the resin. This was found accidentally when an x-ray was taken of a sample in the hospital.