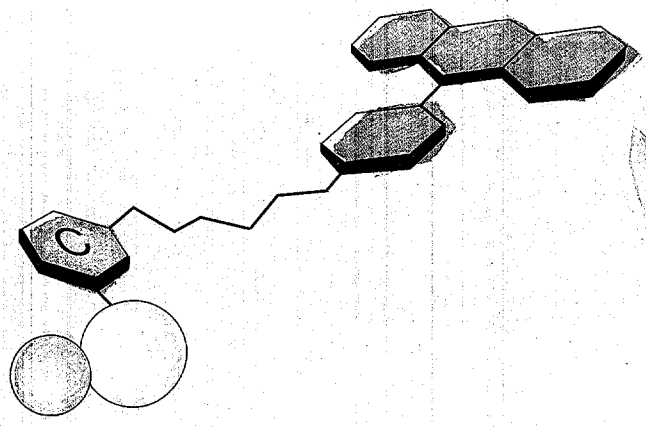

Rapid DNA Sequencing Based On Single-Molecule Detection

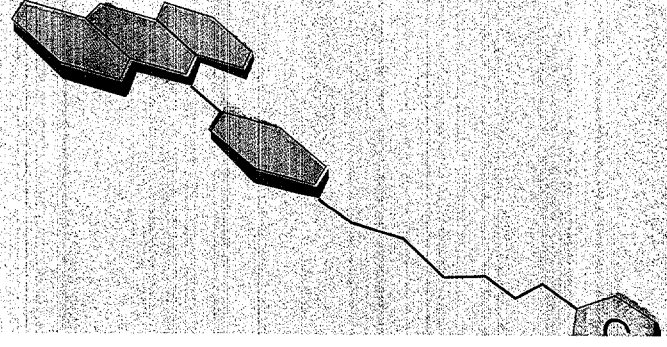

by Lloyd M. Davis, Frederick R. Fairfield, Mark L. Hammond, Carol A. Harger, James H. Jett, Richard A. Keller, Jong Hoong Hahn, Letitia A. Krakowski, Babetta Marrone, John C. Martin, Harvey L. Nutter, Robert R. Ratliff, E. Brooks Shera, Daniel J. Simpson, Steven A. Soper, and Charles W. Wilkerson

It is no secret that faster and cheaper DNA-sequencing techniques must be developed to make large-scale sequencing projects possible and to provide a routine tool for diagnostic medicine. Current methods are probably a hundred to a thousand times too slow and at least a hundred times too expensive for routine use (see pp. in "Mapping the Genome"). Several years ago, a group of us set out to find a better way.

Concluding that gel-based techniques could never achieve the required speed or sequence long stretches of DNA in one pass, we considered a variety of radically new potential methods, some more realistic than others. Ultimately we settled on a scheme that we believed might provide the needed high speed and low cost.

The method involves a combination of technologies from several fields, including molecular biology, enzymology, chemistry, and physics. In concept the method is simple. We would take a single strand of DNA, and starting from one end, remove and identify one nucleotide at a time. Building on Los Alamos expertise in flow cytometry, we intended to suspend the selected DNA strand in a gently flowing stream of water, which would carry the nucleotides to a detector after they were cleaved from the strand. The nucleotides would be removed in order, and if we avoided scrambling that order during the detection process, the sequence would be obtained immediately.

We knew that enzymes (called exonucleases) exist that sequentially remove one nucleotide at a time. Those enzymes can cleave several hundred or more nucleotides from DNA each second. Enzymology was a key element in our method—if nature had not provided those enzymes, our approach would have been unworkable.

Identification of the cleaved nucleotides posed a difficult problem since nature had no ready solution. Laser-excited fluorescence of a dye molecule attached to a nucleotide seemed to be the most likely method for detecting one nucleotide at a time. Such sensitivity had never before been achieved for a molecule in solution, surrounded by trillions of others. But a long history of ultrasensitive detection experiments by our group and others, demonstrating ever increasing sensitivity, encouraged us to believe that new insight and a few tricks might bring us to the ultimate sensitivity limit. After some considerable effort, single-molecule detection was achieved.

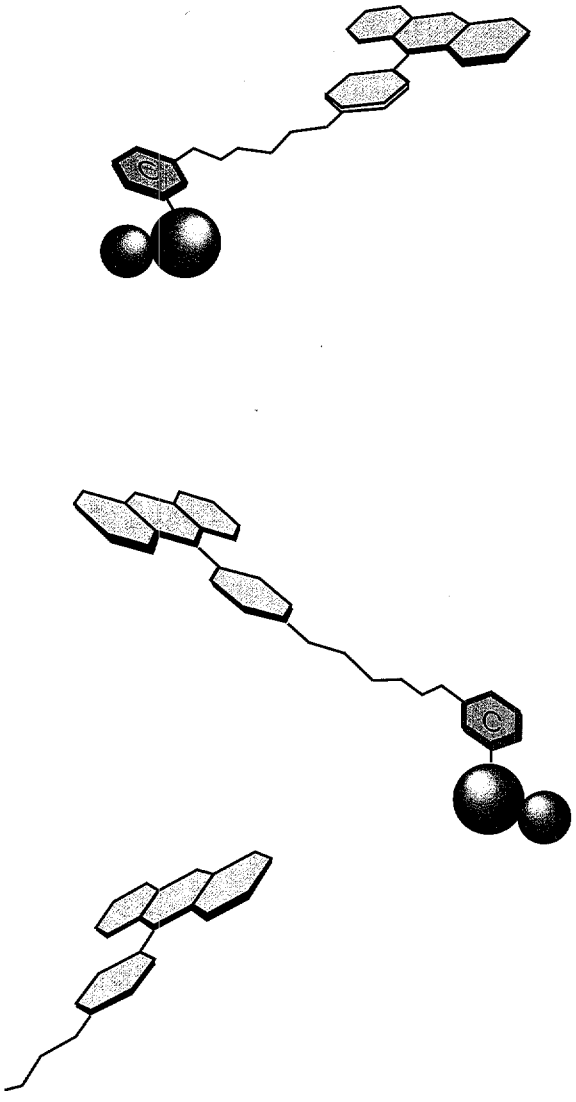

But detection of a nucleotide is not enough. We must identify which of the four possible types of nucleotides, A, C, G, or T, is passing through the detector. For identification we planned to attach a distinctive fluorescent dye molecule to each type of nucleotide. The color of the attached dye's fluorescence would provide the identification. We needed to work out how to chemically attach the labels to each base of the DNA, making sure that the dyes did not impede the cleavage reaction. Our approach was to synthesize pieces of DNA using as raw materials nucleoside triphosphates (dNTPs) bound to dye molecules; the template would be a strand of DNA that we want to sequence (see "DNA Replication" in "Understanding Inheritance"). If no errors occur during synthesis, the dye-labeled strand we created would have a sequence exactly complementary to the template DNA. Thus the sequence of the template would be revealed by sequencing the labeled strand.

A separate technical problem was how to select only a single dye-labeled DNA molecule and suspend it in the flowing water. Using optical traps, in which counter-propagating laser beams hold small objects, seemed like a possibility. Then perhaps the DNA strand could be attached to a tiny bead that would serve as a handle by which the laser beams could hold it.

We estimated that, once we put all these components together, we would be able to obtain rough sequence at the rate of hundreds of bases per second and to sequence tens of thousands of bases consecutively, surpassing by orders of magnitude the performance of current technology.

The Progress and the Problems

The past two years have brought progress in developing some of the techniques we need to make our sequencing scheme a reality. We have chemically attached highly fluorescent dyes to nucleoside triphosphates and used them to synthesize 7000-base DNA strands in which all the C and U nucleotides on one strand were labeled. (Uracil substitutes for the very similar thymine in these experiments.) We are working on replicating longer strands because one of the great potential strengths of our method is the ability to provide long stretches of sequence data in a single pass. (We hope to sequence entire cosmid clones, about 40,000 bases long; gel-based methods can sequence fewer than 1000 bases at a time.) In addition, we have demonstrated that various exonucleases cleave DNA in which all the nucleotides of one type on one strand are labeled, albeit the cleavage is slower than normal. Much of the biochemistry and enzymology is being developed with the help of an industrial partner, Life Technologies, Inc. (LTI), under a Cooperative Research and Development Agreement.

Our plan calls for synthesizing strands in which all the nucleotides are labeled, but if synthesis with all four species of labeled nucleotides continues to prove difficult we could fall back on using only two labeled nucleotide species at a time. By making DNA with each of the six combinations of two labeled nucleotides (that is, a strand with labeled A and C only, one with labeled A and G only, and so forth) and carrying out the rest of our procedure for each strand, we would determine the complete sequence.

To suspend DNA in flowing water we may first attach it to a microscopic bead. Therefore we need to determine automatically whether a bead is attached to exactly one fluorescently labeled DNA molecule. We have identified single DNA molecules by the visual appearance of their fluorescence under a microscope, but we have not yet automated the process. We have also used flow sorting to select beads attached to DNA molecules based on their fluorescent brightness; we are working on increasing the sensitivity so that we can reliably identify single DNA molecules. All our identification experiments use stained DNA molecules that fluoresce roughly as brightly as labeled DNA molecules

of the size we plan to sequence. Construction of an adequate optical trap or development of another way to manipulate the beads remains to be done. Among the standard alternatives to optical trapping are making the bead out of a magnetic material and manipulating it with magnetic fields as well as using mechanical means such as micropipettes.

The detection of single dye-labeled nucleotides is the last step in our scheme, and there we have achieved an important success. We have been able to detect single dye molecules and dye-labeled nucleotides dissolved in water. Further, we have now developed techniques that allow us to distinguish single molecules of one dye from molecules of another by the difference in the color of their fluorescence. As discussed in the following sidebar, "Single-Molecule Spectroscopy in Solution," detection of individual molecules in solution may also find applications in fields other than sequencing.

The nucleotides must be detected in the order in which they were cleaved. However, in the distance between the suspended DNA molecule and the detector, molecular diffusion may scramble the order in which the nucleotides pass through the beam. We are trying to devise designs for the flow and the attachment of the DNA that minimize this problem.

Developing our rapid-sequencing method involves several technological advances. Although none of the parts is ready yet for inclusion in the overall scheme, we have made great progress and continue to be optimistic about overcoming the remaining obstacles because no fundamental principles stand in the way, only difficulties that can be overcome with ingenuity.

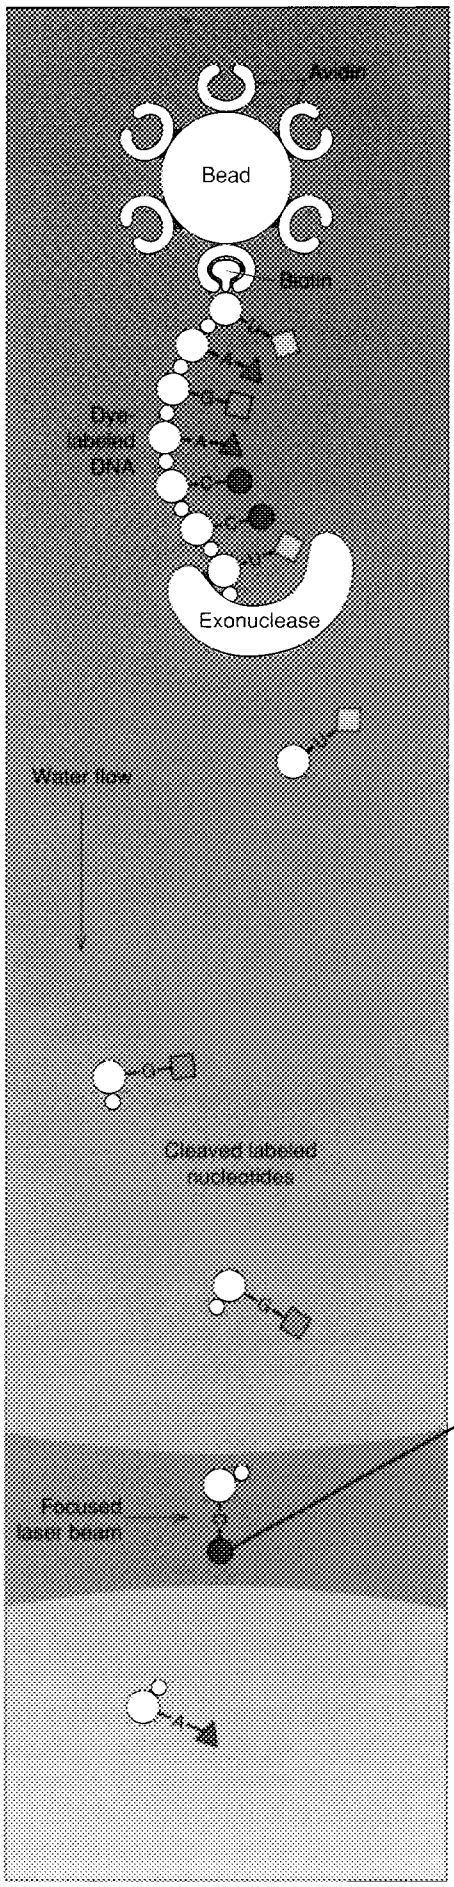

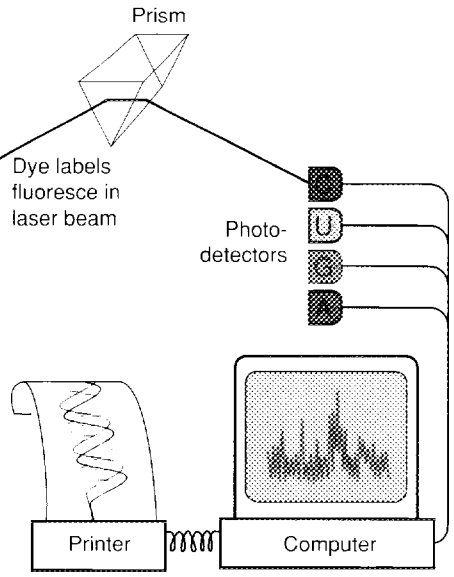

In our projected sequencing method, a bead supports a single strand of DNA in which each base is labeled with a fluorescent dye molecule. (We could also use a double-stranded molecule in which the bases on one strand are labeled.) Each species of base is labeled with a different dye, depicted here with the colors of their fluorescence and with different shapes. An exonuclease enzyme cleaves one nucleotide at a time from the end of the labeled DNA strand. Once a labeled nucleotide is cleaved, flowing water carries it into the detection volume, which is illuminated by a laser beam. The laser light causes the dye label to fluoresce and the fluorescence signal is registered by one of four detectors, depending on the color of the fluorescence emission. Detectors measure the frequency of the fluorescent light, identifying each base as it passes through the beam. The record of the bases going through the beam gives the sequence of the DNA.

1. Dye-labeled Nucleotides

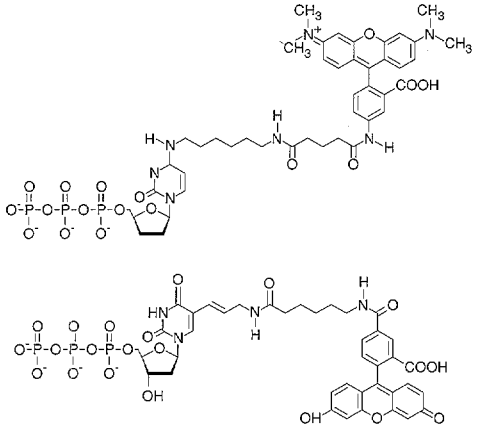

2. Synthesis of Dye-labeled DNA

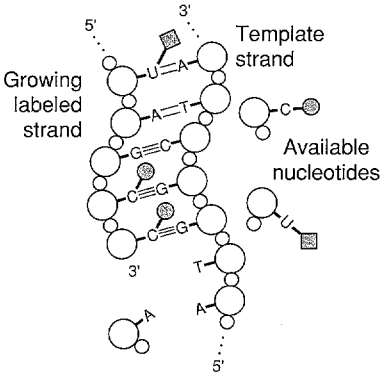

3. Biotin-Avidin Attachment of DNA to a Bead

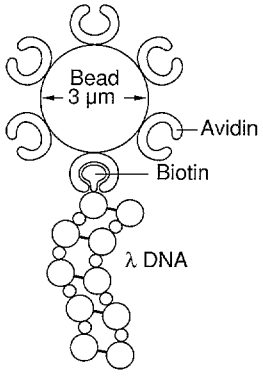

Current Status

We are working on the various components of the project individually. The figures and captions on the next two pages outline both the current status of each component and some of the challenges that remain.

1. To make labeled DNA, we need four species of nucleoside triphosphates (dNTPs), each species labeled with a different fluorescent dye. We have experimented with several commercially available dye-labeled nucleotides, two of which are shown. The first is dCTP linked to tetramethyl rhodamine (commercially available); the second is dUTP linked to fluorescein (synthesized by Life Technologies, Inc.). Since our replication and cleavage experiments work with only some of the labeled nucleotides we have tried, we expect to construct many new ones to find four with which our full scheme can succeed. Those labeled nucleotides will be developed by LTI.

2. To synthesize labeled DNA strands, we will replicate a template strand of DNA using the labeled nucleotides. In collaboration with LTI, we have made DNA molecules about 7000 bases long in which all the C's and U's on one strand are dye-labeled. (The template is the genome of the M13 phage.) The figure is a schematic "snapshot" of DNA replication in which the C's and U's available in the reaction have dye labels, indicated by circles and squares. We determined the completeness of the replication by gel electrophoresis of the product; we are in the process of testing the fidelity.

3. Next we must attach a labeled strand of DNA to a solid support in order to hold it in flowing water. Working with unlabeled DNA strands, we have studied two standard ways of attaching DNA strands to beads about 3 microns in diameter. In one method the beads are glass and so the DNA binds electrostatically. In the other (shown schematically) a polystyrene bead is coated with avidin, a biotin molecule is attached to the 5′ end of the labeled DNA strand, and the tight binding between avidin and biotin attaches the DNA to the bead. The electrostatic attachment method does not interfere with digestion of DNA by an exonuclease. Since that method involves contact between the DNA and the bead at many points, whereas the biotin-avidin method only involves contact at the last base of the DNA, we expect that the biotin-avidin method will not interfere with exonuclease cleavage either.

4. Since the attachment step just described may be performed by mixing beads into batches of labeled DNA, we need to be able to identify beads carrying exactly one DNA molecule, and then move one of those beads into a stream of water and hold it in position there. We have learned to identify under a microscope single fluorescently stained λ DNA molecules (about 48.5 kbp long) not attached to beads. The objects we identified are confirmed to be λ DNA by chemical tests and by length measurements with gel electrophoresis. The method involves some computer assistance; we have begun work on automating it fully. With an alternative method, flow sorting, we have distinguished beads carrying one to three fluorescently stained

DNA molecules from other fluorescent objects. The sensitivity of that method needs to be improved so that we can be confident of identifying single molecules. We have not yet tested ways of manipulating and suspending single beads, but hope to use one of the techniques mentioned in the main text.

5. Once a DNA molecule is suspended in the stream of water, an exonuclease will cleave bases one by one from the labeled DNA. Working with LTI at their laboratories, we tested the exonuclease cleavage of DNA in which one species of base is labeled. (The DNA was produced in our replication experiments.) All six exonucleases we studied digested the DNA. Some of the exonucleases act on single-stranded DNA, others on double-stranded. Four of the enzymes are polymerases or components of polymerases. (Most polymerases include a part that acts as an exonuclease, apparently to proofread replication.) We studied one enzyme, the polymerase made by the T7 phage, in more detail. The enzyme completely digested 300-base DNA strands made with rhodamine-labeled C or rhodamine-labeled U (but degradation of strands made with fluorescein-labeled U was incomplete). The completeness of the digestion was determined by gel electrophoresis of the reaction products, which showed no fragments longer than one base. The cleavage rates under the conditions we used were around 10 to 20 nucleotides per second, two to five times slower than rates for unlabeled DNA. (To measure cleavage rates, we stop the reaction at various times by adding chemicals that "poison" it, and determine the size of the undigested portion by gel electrophoresis or other methods.) We are continuing to search for enzymes and reaction conditions that give rapid and complete cleavage.

6. In the final step a stream of water will carry the labeled bases one by one through a laser beam, where each base will be identified by the laser-induced fluorescence of its dye label. We have not yet constructed a prototype flow system. The flow must not change the order of the cleaved bases and must ensure that each base passes through the detection volume of about 10^{-12} liters. As described in "Single-Molecule Spectroscopy in Solution", we have succeeded in detecting single dye molecules and dye-labeled nucleotides, and in distinguishing single molecules of two species of dyes (shown as circles and triangles). We need to make the detector capable of distinguishing four species, and to improve the accuracy of detection and identification from about 65 percent to at least 99 percent to be adequate for DNA sequencing. ■

Further Reading

J. H. Jett, R. A. Keller, J. C. Martin, B. L. Marrone, R. K. Moyzis, R. L. Ratliff, N. K. Seitzinger, E. B. Shera, and D. J. Simpson. 1989. High-speed DNA sequencing: An approach based upon fluorescence detection of single molecules. *Journal of Biomolecular Structure and Dynamics* 7: 301–309.

S. A. Soper, L. M. Davis, F. R. Fairfield, M. L. Hammond, C. A. Harger, J. H. Jett, R. A. Keller, B. L. Marrone, J. C. Martin, H. L. Nutter, E. B. Shera, and D. J. Simpson. 1991. Rapid DNA sequencing based on single molecule detection. *Proceedings. SPIE: The International Society for Optical Engineering* 1435: 168.

J. D. Harding and R. A. Keller. 1992. Single-molecule detection as an approach to rapid DNA sequencing. *Trends in Biotechnology* 10: 55.

5. Cleavage of labeled DNA

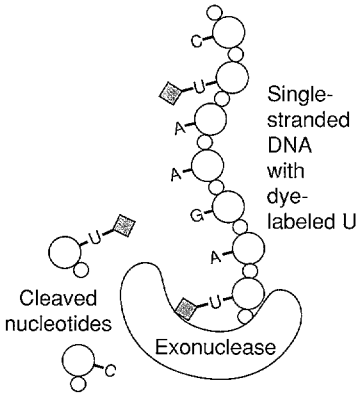

6. Single-molecule Detection

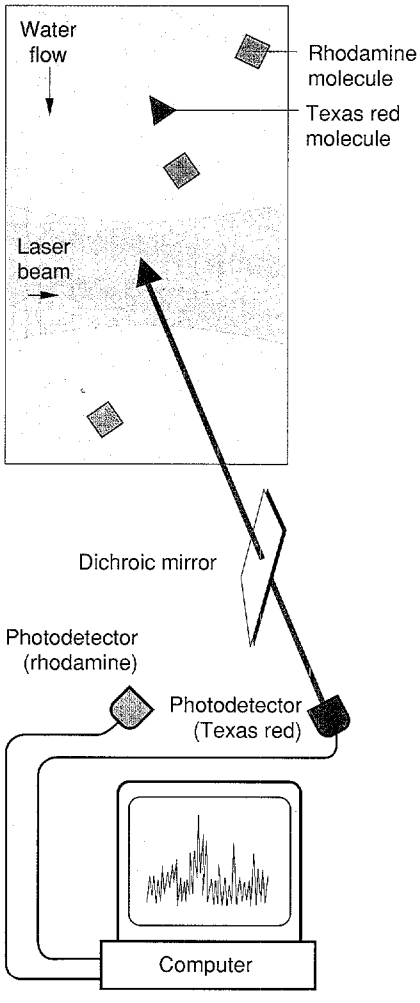

Single-Molecule Spectroscopy in Solution

Steven A. Soper, Lloyd M. Davis, and E. Brooks Shera

Detecting minute concentrations of chemicals in liquid environments has many applications. Consequently for over a decade scientists have been pushing the limits of detection sensitivity to lower and lower chemical concentrations. The Los Alamos plan for rapid DNA sequencing, described in the preceding pages, calls for the ultimate sensitivity—detection of single molecules. This article reviews the first technique used to accomplish that feat.

The technique involves measuring the fluorescence emission of a molecule as it passes through a laser beam. In addition to identifying the molecule, such spectroscopic measurements can reveal information about that molecule's chemical or physical surroundings, and thus single molecules may be used as probes to explore biological processes and structures at a microscopic level.

Physical Basis for Detection

Our detection method is applicable to fluorescent dye molecules, or to any molecules, such as nucleotides, that have been labeled with fluorescent dyes. The molecules, in solution, pass through a rapidly pulsed laser beam of such a wavelength that the dye can absorb the light. The solution must be so dilute that only one fluorescent molecule at a time passes through the beam. As each molecule passes through, many extremely brief pulses of light illuminate it; each pulse may cause the molecule to emit a fluorescence photon. Thus while in the beam each molecule produces a short burst of photons. We ascertain when a fluorescent molecule is in the beam by detecting some of the individual fluorescence photons that make up the burst and discriminating them from other photons that reach the detector.

Figure 1a illustrates both the process of fluorescence and other processes that compete with it. The processes begin when a molecule in its ground electronic state absorbs a photon and is thereby excited to a higher electronic state. (Our experiments use photons of visible light that excite the molecule to the first excited singlet state.) Fluorescence occurs when the molecule then quickly relaxes to a slightly lower energy through changes in its rotational and vibrational motion, and finally returns to the ground electronic state by emitting a photon. The photon emitted in the transition from the first excited singlet state to the ground state is called the fluorescence photon. A molecule is said to be fluorescent if it has a high probability of returning to the ground state by that path rather than by the other paths shown in Figure 1a. That probability is called the fluorescence quantum yield. As Figure 1b shows, the frequency (and energy) of the fluorescence photon is lower than that of the absorbed photon. The average difference, called the Stokes shift, is roughly the same for most organic dyes. For example, a dye we often use, rhodamine 6G, is excited by green light and emits yellow fluorescence.

Once the molecule returns to the ground state, by any path, it is again available for excitation. However, absorption of a photon does not always bring about a reversible process. Sometimes absorption causes the molecule to undergo photobleaching, an irreversible change into a different chemical species, after which the molecule can no longer fluoresce and often cannot absorb light of the frequency of the absorbed photon. (The same process causes the fading of dyed materials exposed to sunlight.) If photobleaching occurs, the production of fluorescence photons stops even if the molecule is still illuminated. The probability that an excited molecule will bleach instead of returning to its original state is called the photobleaching efficiency (though in our work photobleaching is a source of inefficiency).

For single-molecule detection, an important feature of fluorescence is that the time between the excitation of the molecule and the emission of a

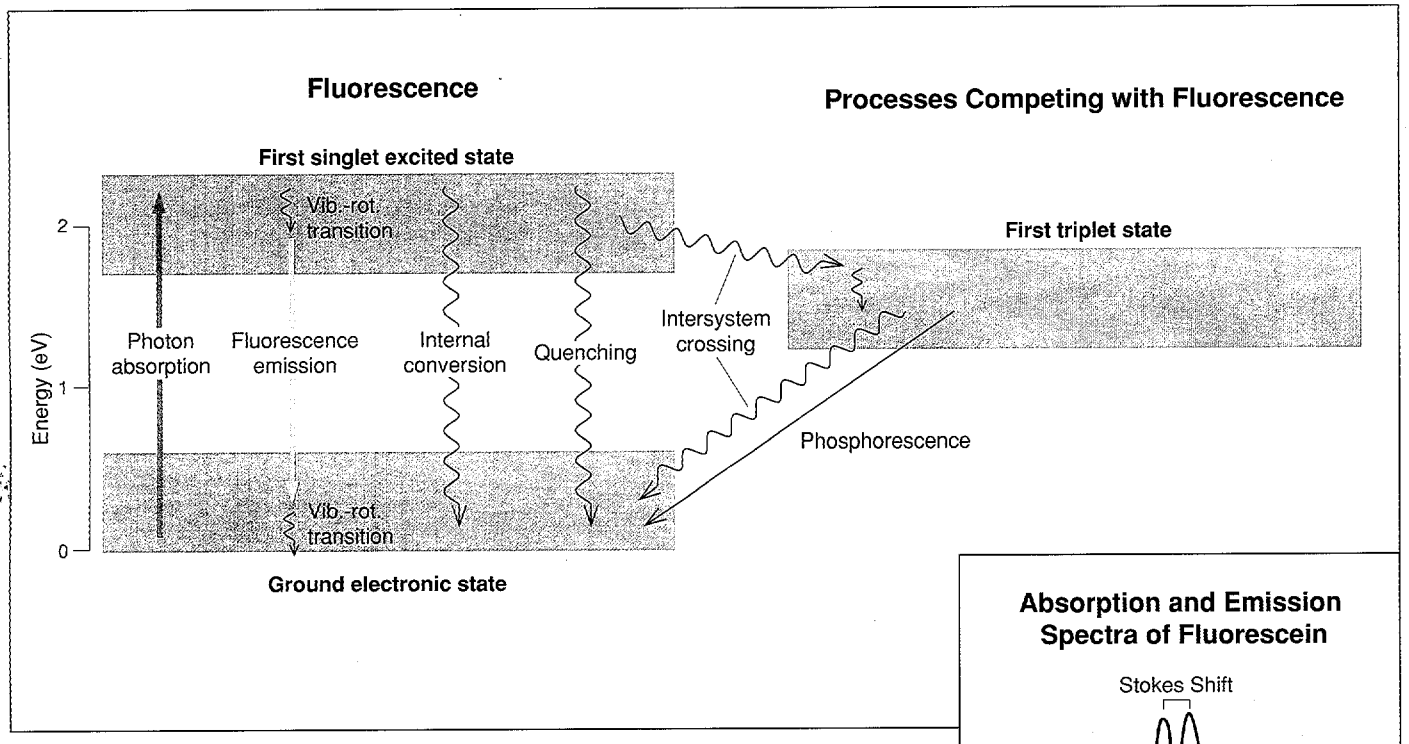

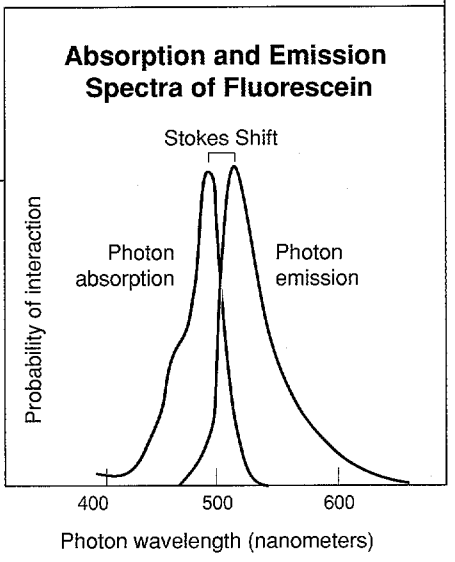

Figure 1. (a) Fluorescence of a Dye Molecule

An energy-level diagram of a dye molecule showing the processes involved in fluorescence, as well as those that compete with fluorescence. In all cases the first step is absorption of a photon (green arrow), which causes the molecule to make a transition from its ground electronic state (a singlet state for most organic molecules) to the first excited singlet state. (Each electronic state is shown as a continuous band because within each such state are many closely spaced rotational and vibrational levels, which overlap one another when the molecule is in solution.) The excited molecule may return to the ground state in several ways, one of which is fluorescence. When undergoing fluorescence, the molecule first loses energy by a rapid series of rotational and vibrational transitions (wavy arrow), remaining in the same electronic excited state. The molecule then makes an electronic transition to some level within the ground electronic state (yellow arrow), emitting a photon—the fluorescence photon. Finally the molecule relaxes to a low-lying level within the ground electronic state by further vibrational and rotational transitions (wavy arrow). The average time from photon absorption to fluorescence emission is called the fluorescence lifetime. Most organic fluorescent dyes have lifetimes of a few nanoseconds. Processes that compete with fluorescence are indicated by lighter arrows. A molecule in the first singlet excited state can decay to the ground state without fluorescing. The energy is converted either into heat (internal conversion) or into the excitation of a molecule of another species (quenching). Another possibility is that the molecule can decay (without emitting a photon) to an excited triplet state (intersystem crossing). After a time that for organic dyes is much longer than the fluorescence lifetime, the molecule decays to the ground electronic state, with or without photon emission. (Photon emission in a transition from the first triplet state is called phosphoresce, and has a much longer lifetime than fluorescence.)

(b) The Stokes Shift

The probabilities of photon absorption and fluorescence emission as functions of wavelength for a typical fluorescent dye (fluorescein). Note that the emission curve peaks at a longer wavelength than the absorption curve. The difference between the two peaks is the Stokes shift. The fluorescence photon has a lower energy than the absorbed photon because some of the molecule's excitation energy is converted into heat though rotational and vibrational transitions.

Figure 2. The Time Scales of Single-Molecule Detection

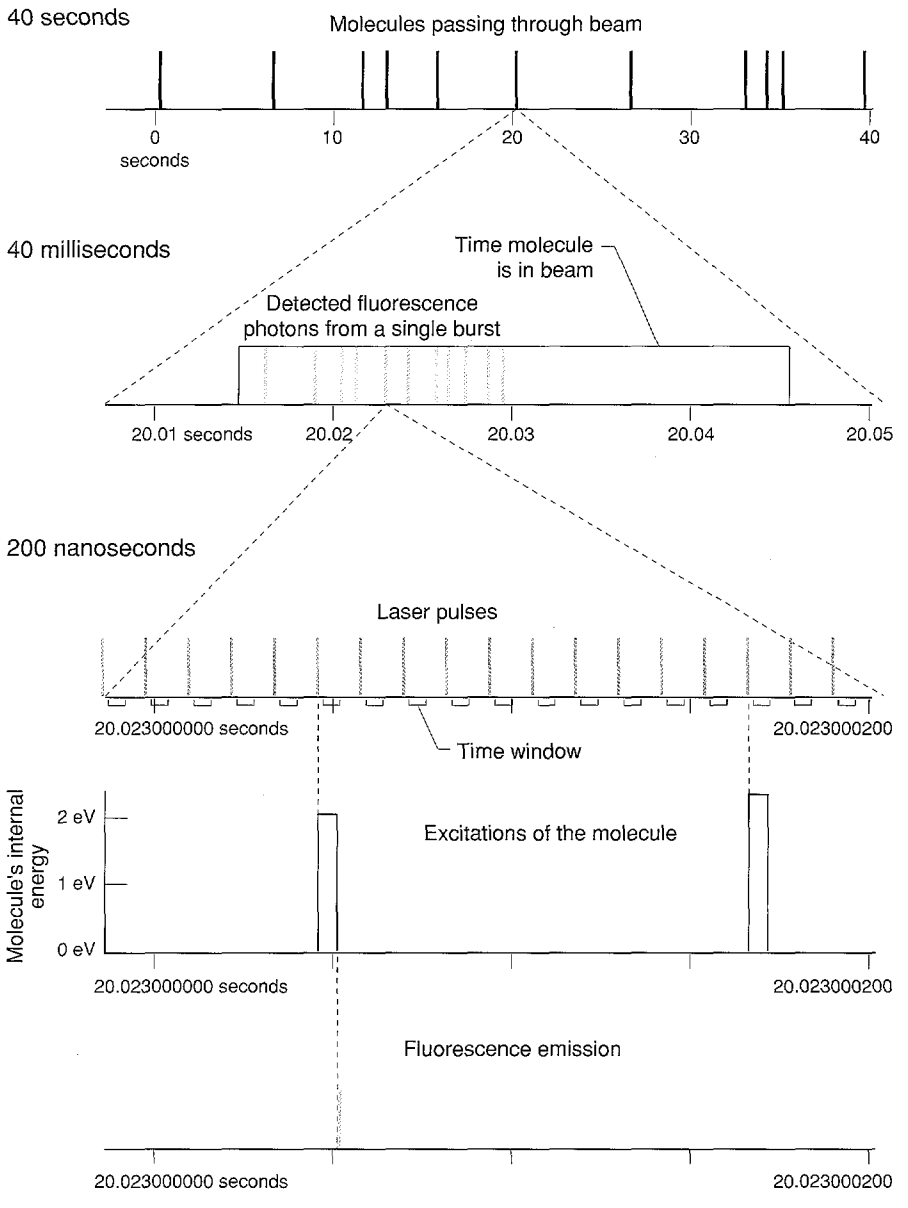

Molecules pass through the laser beam at random, 4 seconds apart on average.

The average time a molecule spends in the beam is about 30 milliseconds. During that time, the molecule emits a burst of photons, which ends when the molecule photobleaches.

The laser pulses every 12.2 nanoseconds. The time periods from 1 to 5 nanoseconds after each pulse (small brackets) are the time windows during which we detect photons.

On average one laser pulse in ten excites the molecule. The molecule remains excited for an average of a few nanoseconds (the fluorescence lifetime), most of that time near the bottom of the first excited singlet state.

For a molecule whose quantum yield is 0.5, half the excitations lead to the emission of a fluorescence photon.

fluorescence photon is usually finite—and long enough to measure. That time difference follows an exponential probability distribution whose average is called the fluorescence lifetime. The dyes we use have fluorescence lifetimes of a few nanoseconds. Timing is crucial to the design of our experiments on single-molecule detection. Figure 2 illustrates the relevant time scales. The flashes of light from our laser last about 0.07 nanoseconds and repeat every 12.2 nanoseconds.

Since the duration of the laser pulses is much shorter than the fluorescence lifetime, each pulse can bring about fluorescence at most once, thus producing at most one photon. The time between pulses is much longer than the fluorescence lifetime, so a molecule that absorbs a photon is practically certain to return to the ground state and be ready for another excitation by the time the next pulse arrives. Therefore, in principle, every pulse could cause the molecule to emit one fluorescence photon. Since our apparatus detects individual photons, we can take advantage of the interval between the arrival of the laser pulse and fluorescence to distinguish between fluorescence photons and photons produced by scattering of light from the laser pulse, as will be described below in the discussion of background light.

Signal Strength from Single Molecules

For the fluorescence signal to be detectable, its strength (the number of photons the signal comprises) must be large enough to be distinguished from photons produced by background

sources. To make the simplest estimate of the signal strength, we recall that the molecule can produce one photon, but no more, for every pulse that illuminates it. In our apparatus the molecules' transit time through the beam is about 30 milliseconds, so with the laser pulsing at 12-nanosecond intervals we might expect a molecule with a large quantum yield to emit a burst of 2,400,000 photons—provided that the laser is powerful enough to excite the molecule with each flash and that photobleaching does not occur. Actually, the laser intensity we use is so low that typically a molecule will be excited by only about one laser pulse in ten. Even so, we might expect about 240,000 photons. However, photobleaching practically always stops the emission of photons before the molecule leaves the beam, thereby greatly reducing the size of the signal. The average number of excitations a molecule endures before bleaching is the reciprocal of the photobleaching efficiency, which for rhodamine 6G dissolved in water is 1.8×10^{-5}. Accordingly each rhodamine 6G molecule is excited, on average, 56,000 times before it photobleaches (so a more powerful laser would not increase the signal). Of those excitations, the fraction that induces fluorescence is by definition the fluorescence quantum yield. The quantum yield of rhodamine 6G is 0.45, so on average, a rhodamine 6G molecule is expected to produce about 25,000 fluorescence photons.

Our photon detector records about one photon of every thousand emitted, so we might hope to see 25 photons from one molecule. (Our group has recently obtained a detector that should do ten times better than our present equipment.) For more precise estimates of the signal size, we constructed an elaborate Monte Carlo computer simulation of our experiments that includes all the physics and photochemistry that significantly affect our detection ability. Of the processes not yet mentioned but included in the simulation, the most important is diffusion, or random motion, of dye molecules, through which they can move into or out of the beam. The results from the simulation suggested that a typical signal would contain 10 to 15 photons. The simulation also proved to be very helpful in designing our experiment and quickly optimizing such experimental conditions as the flow rate and the size of the detection volume.

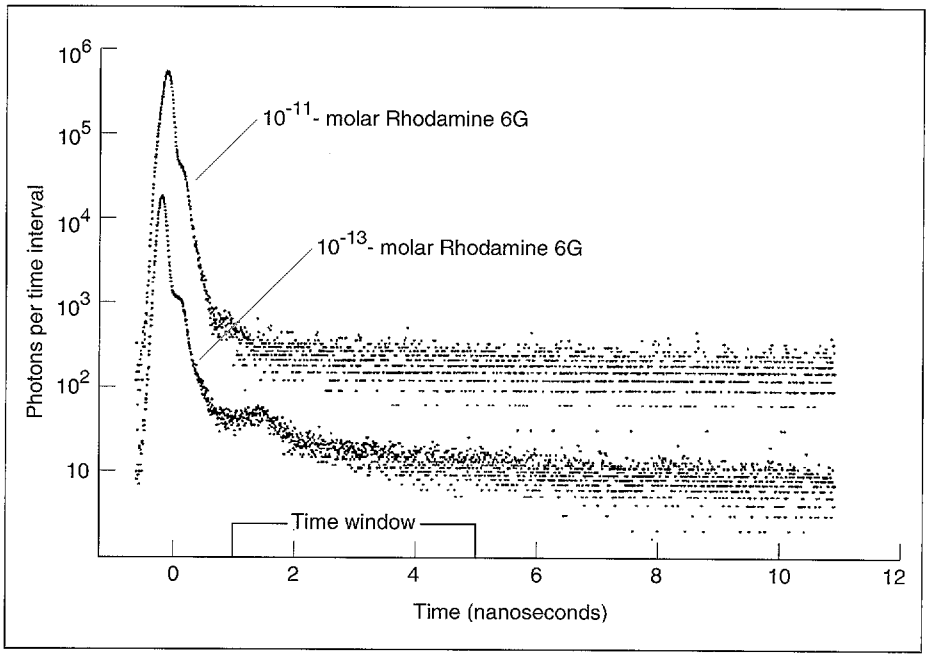

Figure 3. The Basis for Time-gated Detection of the Fluorescence Signal

The plotted data were accumulated during experiments in which laser pulses illuminated two dye solutions with different concentrations of R6G: 10^{-11} molar and 10^{-13} molar. The graph shows the number of detected photons as a function of the time interval (measured from the time of excitation) in which they were emitted. During the first nanosecond, the light intensity is independent of the dye concentration and therefore must be due mainly to light scattered from the laser pulse. After 1 nanosecond, the light intensity decreases exponentially and is proportional to the dye concentration, as one would expect from fluorescence emission. After about 5 nanoseconds, minor background sources with long lifetimes (perhaps from the glass in our apparatus) start to compete with the more quickly decaying fluorescence signal. Therefore to minimize the contribution from the background, we count only photons emitted in the time window shown, between 1 and 5 nanoseconds after each pulse.

Reducing the Background

Detection of light from single molecules in solution is difficult because the number of background photons is far larger than the number of photons in the fluorescence burst. Most background photons are laser photons that have undergone either Raman scattering from water molecules, Rayleigh scattering primarily from small-scale fluctuations in the density of the water, or reflection from surfaces of water and glass in

the equipment. A smaller source of background is laser-induced fluorescence of impurities in the water. Light from Rayleigh scattering and reflection has the same color as the green laser light, so we can block most of it by placing a color filter in front of our detector. For example, we can use a filter that reflects green light but transmits yellow fluorescent light from rhodamine 6G. The color filters are made of thin layers of various colorless materials deposited on a clear glass support. By means of optical interference, the filters reflect all but one photon in a million at the frequency of the laser light while transmitting about 60 percent of the fluorescence photons.

The light from Raman scattering can be from a hundred to a thousand times more intense than a typical fluorescence signal. Some of that light has the same color as the fluorescence emission of the molecule, so it can not be eliminated by filtering. However, photons scattered from a laser pulse are nearly simultaneous with the pulse, whereas fluorescence photons are likely to be emitted well after the pulse. The time dependences of light from those two sources are shown in Figure 3 (previous page), a plot of the light intensity (versus time) obtained when a pulsed laser beam illuminated solutions of rhodamine 6G containing two different concentrations of the dye. The height of the initial light-intensity peak is independent of the dye concentration, so most of the photons constituting that peak are background photons. On the other hand, the light intensity at later times depends strongly on the concentration; therefore, much of the light at later times must come from the dye molecules, presumably from fluorescence. With these considerations in mind, we reduce the background by counting only those photons produced after the bright laser pulse is over but while the probability of emission of a fluorescence photon is still relatively high, a procedure called time gating. The time window during which we record photons is shown in Figures 2 and 3. Ignoring photons emitted outside that window causes a small loss of signal, but decreases the number of accepted background photons by a factor on the order of a thousand.

Together, time gating and color filtering reduce the background intensity reaching the detectors by a factor of roughly a billion. Consequently, we can see the faint fluorescence signal from a single dye molecule even though tens of trillions of surrounding water molecules are illuminated by the laser. By the same token, we detect fluorescence from impurities in the water, so they are an important source of background. We minimize their effect, first, by purifying the water as much as possible and, second, by minimizing the volume of water illuminated by the laser and monitored by our detectors.

Identification by Color

For many applications, including high-speed DNA sequencing, merely detecting each dye molecule that passes through the laser beam is not sufficient; we also need to distinguish different types of molecules. Because the molecules are detected by their fluorescence emission, it is natural to distinguish molecules of different chemical species by observing some spectroscopic property such as the color of their fluorescence emission. The emitted photons of different colors can be separated by using a glass prism to bend the light in different directions according to color. However, in our set-up, it is more efficient to use color filters and color-selective (dichroic) mirrors, which, like the color filters described earlier, work by thin-film interference. Each type of molecule we want to detect is assigned its own photodetector, which is shielded by a color filter that transmits light of that molecule's emission wavelength but reflects light from molecules of other types in the solution. In principle, every dye molecule that passes through the beam can be identified by noting which detector signals its presence.

A fluorescent dye molecule absorbs and emits light over a range of wavelengths (as shown in Figure 1). Consequently, to be readily distinguishable by fluorescence color, different dyes must have emission spectra whose peaks are well separated in wavelength, by at least some tens of nanometers. Since the Stokes shift is roughly the same for most organic molecules (about 20 nanometers), it is difficult to find two dyes that absorb efficiently at the same wavelength but fluoresce at wavelengths far enough apart to be distinguishable. Therefore we need a different laser to excite each species of dye—a complication in designing the experiment.

Apparatus for Spectroscopy

Figure 4 shows the apparatus we have developed for a demonstration of spectroscopy on single molecules. Two lasers, one activated by the other, provide two colors of excitation light in synchronized pulses so that we can observe two types of molecules in one solution. Since dyes and dye-labeled nucleotides behave the same way in the experiment and labeled nucleotides must be made specially, we have used the dyes rhodamine 6G and Texas red in our experiments to date. Water containing an extremely low concentration of dye molecules flows through a glass tube, called the flow cell. Laser light is focused into a very narrow beam that

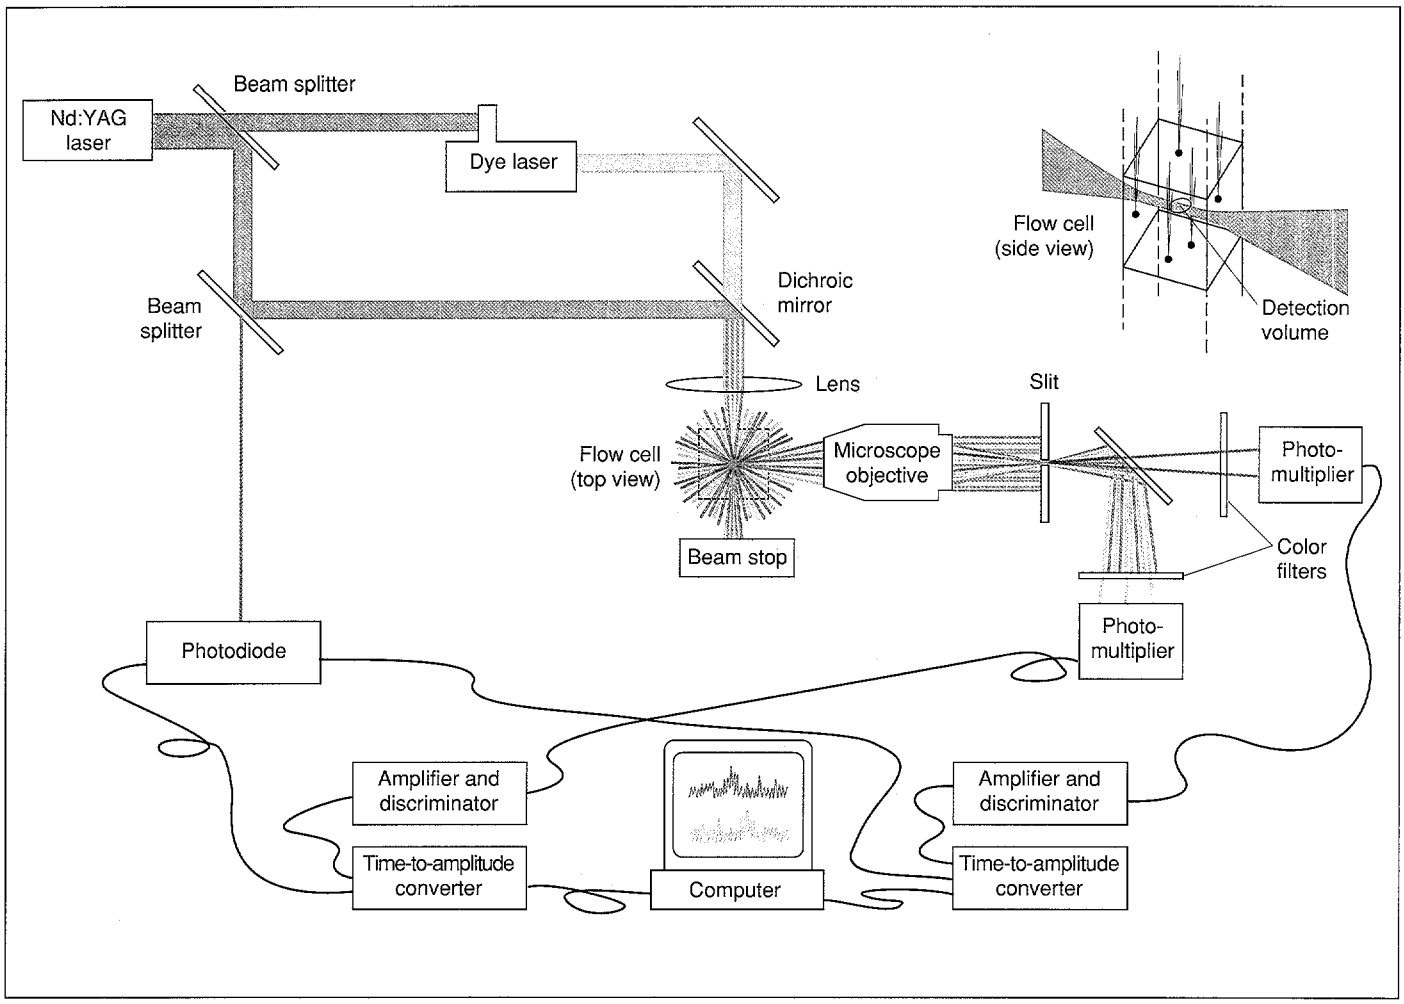

Figure 4. Apparatus for Single-Molecule Spectroscopy

The figure schematically illustrates the optics for delivering two pulsed laser beams to a flow cell containing fluorescent dye molecules of two different kinds as well as the apparatus for detecting fluorescence signals emitted by the two dyes, rhodamine 6G and Texas red. A mode-locked neodymium:YAG laser produces pulsed green light (532 nanometers). Part of the light is deflected by a beam splitter toward a dye laser and causes that laser to produce yellow light (585 nanometers, a wavelength that can excite the dye Texas red) pulsed with the same frequency as the green light that stimulates it. To generate a start signal for time-gated detection of the fluorescence signal, another beam splitter sends a small fraction of the original green light to a photodiode. Various lenses and mirrors (most of them not shown) direct the rest of the green laser light and all the yellow laser light to a dichroic mirror. A delay in the path of the green light synchronizes the green and yellow light pulses. The dichroic mirror combines the two beams into one by transmitting yellow light and reflecting green light. Then a lens focuses the combined beam into the detection volume within the glass flow cell (inset). As an extremely dilute solution of the dyes flows through the cell, the laser light induces the emission of fluorescence light from dye molecules. Some of the laser light is also scattered from the water in the flow cell. To detect the fluorescence signal, a microscope objective (which subtends about 2 percent of the total solid angle) collects light from the flow cell and focuses it on an opaque plate with a slit. Only light that comes from the detection volume, a small fraction of the water illuminated by the beam (about 10^{-12} liters), passes through the slit. This arrangement minimizes background light from outside the detection volume, including fluorescent impurities. Likewise if light comes from other directions (meaning that it is background) the edges of the slit block it. Light emerging from the slit continues to another dichroic mirror that reflects the long-wavelength light, including the orange fluorescence from Texas red, toward one detector, while transmitting the yellow fluorescence from rhodamine 6G to the other detector. Color filters in front of each detector transmit only photons in the expected frequency range for fluorescence of the appropriate dye. As explained in the text, the photomultiplier tubes produce current pulses for about 5 percent of the incident photons. The pulses are then amplified, and small pulses, which probably result from random instrumental noise, are filtered out by discriminators. Next pulses from photons that were not emitted during the time-gating window are rejected. The necessary time-keeping is performed by time-to-amplitude converters (TACs). A TAC begins measuring time at an electrical start signal produced by the photodiode each time a flash of green laser light reaches it. Then every time the TAC receives an electrical pulse from the photomultiplier, which indicates the arrival of a photon from the detection volume, it generates an output pulse whose peak voltage is proportional to the time since the start signal. Other electronic components accomplish the time gating by transmitting only those pulses whose voltage indicates that the corresponding photons were emitted inside the time-gating window.

Figure 5. Computer Processing of the Data

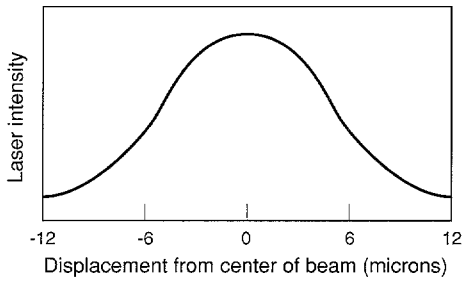

The laser intensity has a Gaussian dependence on distance from the center of the beam.

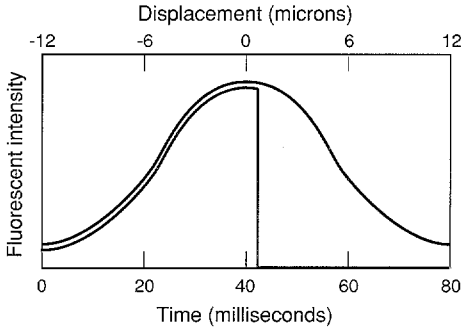

With a fairly low-intensity beam the average fluorescence emission rate is proportional to the laser intensity at the molecule's position. If the molecule moves with constant velocity through the center of the beam, the average emission rate has the same dependence on time as on space. After photobleaching (which can occur at any time), fluorescence stops, so the average emission rate might look like the red curve.

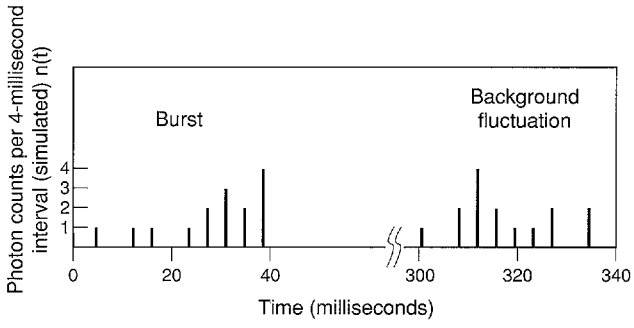

These fictitious data include two 40-millisecond time periods during which photons arrive much faster than the average background rate. Such groups of photons may be fluorescence bursts. The first group has the characteristic profile of a burst (red curve above); the second looks more like an extremely unusual fluctuation of the background.

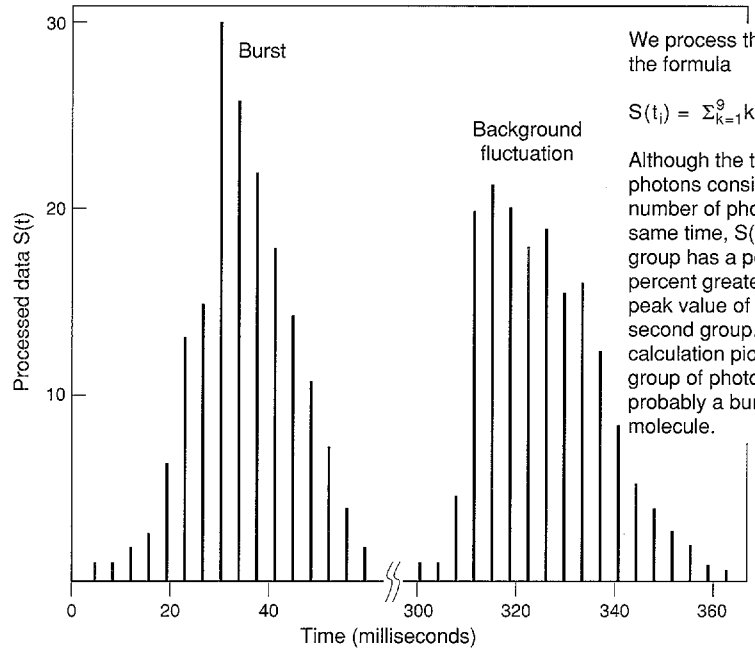

We process the data with the formula

$$S(t_i) = \Sigma_{k=1}^{9} k[n(t_{i-9+k})]^2/9.$$

Although the two groups of photons consist of the same number of photons in the same time, S(t) for the first group has a peak value 50 percent greater than the peak value of S(t) for the second group. Thus the calculation picks out the group of photons that is probably a burst from a molecule.

passes through the flow cell, where it causes the dye molecules to fluoresce, and produces background light through scattering processes. A fraction of the light from the flow cell falls on lenses, which focus that light onto a plate with a slit. Light emanating from the detection volume, a very small volume around the focal point of the laser beam, passes through the slit. A dichroic mirror directs light of the color produced by rhodamine 6G to one photodetector while sending light of the color produced by Texas red to the other detector. At the detector, each photon of the selected wavelength passes through a color filter and strikes the cathode of a photomultiplier tube where it can produce a free electron by the photoelectric effect. Electric fields accelerate the electron toward the anode, causing it to jar other electrons loose from solid structures of the tube, which in turn are accelerated and jar still more electrons loose. Thus about one photon in twenty that strikes the cathode gives rise to a current pulse large enough for the rest of the electronics to discriminate it from noise in the detector. As indicated in Figure 4, electronic components measure the time, relative to the most recent laser flash, at which each photon arrives at the detector, and reject photons that do not arrive during the time-gating window.

Figure 5 outlines the computer data processing that distinguishes fluorescence bursts emitted by molecules of each type from the background photons that reach the detector. The distinguishing feature of a fluorescence burst is the time dependence of the rate at which photons are detected. Typically the rate increases as the molecule moves toward the center of the laser beam, where the laser light is most intense, then drops abruptly when photobleaching occurs. Though random fluctuations in the background

may occasionally produce many photons in a short time, they are unlikely to duplicate the characteristic time profile of a fluorescence burst. Every few milliseconds, the computer program calculates a function $S(t)$ that depends on the detection rate of the photons that arrived at the detector during the previous few dozen milliseconds. The function's value is large when the temporal pattern of incoming photons is typical of a fluorescence burst, but smaller when photons arrive in other patterns. If many photons arrive in rapid succession but their rate of arrival does not increase with time as in a burst, the value of $S(t)$ calculated from those data will be smaller than that for a typical burst. We record the presence of a molecule when the value of $S(t)$ exceeds a set threshold.

Experimental Results for Single-Molecule Detection

We used the apparatus shown in Figure 4 with solutions of single dyes, and more recently with a solution of rhodamine 6G and Texas red, both diluted to 10^{-14} molar. In the latter experiment the flow speed through the detection volume was about 290 microns per second. Figure 6 shows raw and processed data for both dyes. Peaks in the processed data above the thresholds (dotted lines) are interpreted as signals from dye molecules. Thus we are able to detect and distinguish individual dye molecules of two types in a mixed solution.

Computer processing does not eliminate all errors. We set the thresholds so that the false-positive rate in experiments with no dye present is no more than 0.01 per second. In an experiment with only rhodamine 6G we saw 87 percent of the dye molecules (calculated by comparison with the

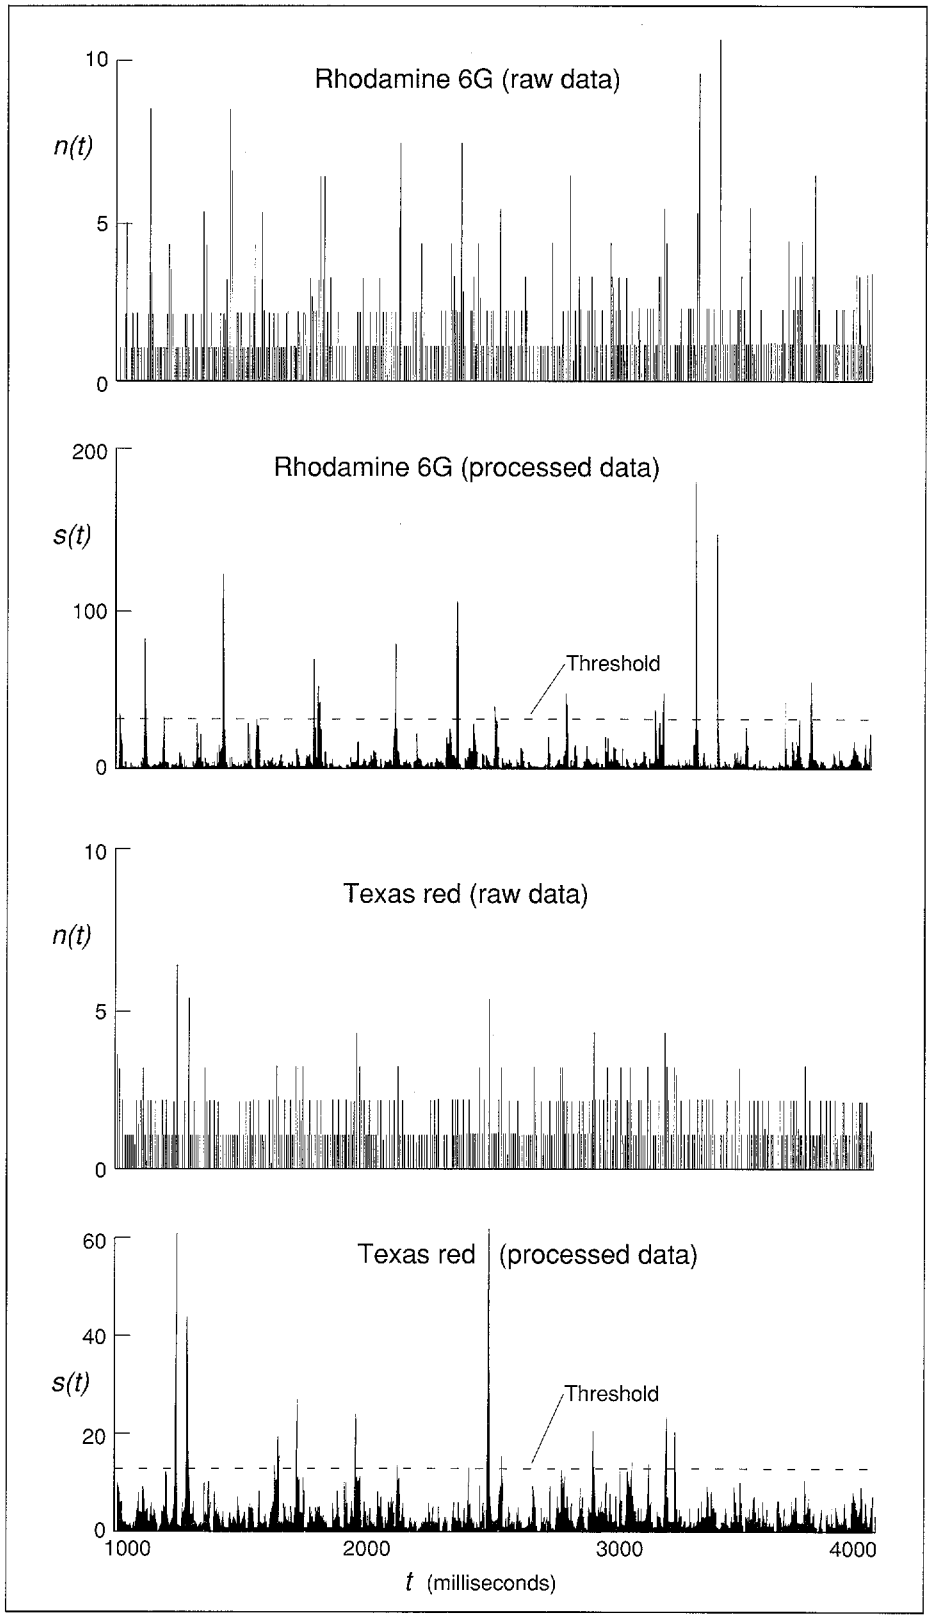

Figure 6. Single-Molecule Discrimination of Two Dyes

The apparatus depicted in Figure 4 was used to detect and distinguish two different dye molecules, R6G and Texas red. The upper pair of graphs shows raw data *n(t)* and processed data *S(t)* from R6G. All peaks in *S(t)* higher than a threshold marked by the dotted line were interpreted as indications of fluorescence bursts from R6G molecules. Use of the threshold restricted the false-positive rate to 0.01 per second. The lower graphs show the analogous data for the same time period from Texas red.

estimated rate at which molecules pass through the detection volume). Since the optimum chemical conditions for reducing photobleaching of rhodamine 6G are incompatible with those for Texas red, when we ran experiments with both dyes we detected 79 percent of the rhodamine 6G molecules and 54 percent of the Texas red molecules that flowed through the detection volume. (Improvements in the apparatus, such as our new photomultipliers, should soon allow much better efficiencies.) These experimental results agree approximately with our Monte Carlo prediction of the rate at which we should detect single molecules. The agreement gives us confidence that we understand the photophysics of single molecules in solution.

Identification by Lifetime

Spectroscopic properties other than emission wavelength can be used to distinguish different types of molecules. Fluorescence lifetime is convenient for us to measure. It is particularly useful because molecules of different types usually have different lifetimes, as do molecules of any one type in different chemical environments. As stated above, fluorescence lifetime is the average amount of time that a molecule remains excited before returning to the ground electronic state through the emission of a fluorescence photon, and the individual times from excitation to emission are random and follow an exponential probability distribution.

The standard way to measure a fluorescence lifetime is to excite a concentrated solution of a dye with a pulse of light and observe the exponential decay in intensity of the light that the many dye molecules produce. (An exponential-decay curve in data from essentially the same experiment appears in Figure 3.) On the other hand, to determine the average lifetime of a single dye molecule, we must re-excite that molecule many times and measure the time to fluorescence following each excitation. Our apparatus is already set up to observe the individual times between excitation and fluorescence, denoted ΔT, since that measurement is required to implement the time gating described earlier. Because the resulting sample of individual ΔT values is small, it is more efficient to calculate the lifetime from the data by taking the mean of the time differences between excitation and fluorescence than to fit the data to an exponential distribution (as shown by Peierls in 1935).

The main purpose of our experiments was to demonstrate the feasibility of measuring the fluorescence lifetimes of single molecules with enough precision to discriminate between chemical species. Therefore we took steps that are incompatible with certain applications of lifetime measurements, such as high-speed sequencing. In particular, to maximize the number of fluorescence photons from each molecule, we reduced photobleaching by using methanol as the solvent and extended the time window almost to the next laser pulse. We also used a low flow speed so that each molecule would remain in the beam longer. The solvent moved so slowly, in fact, that the length of time the

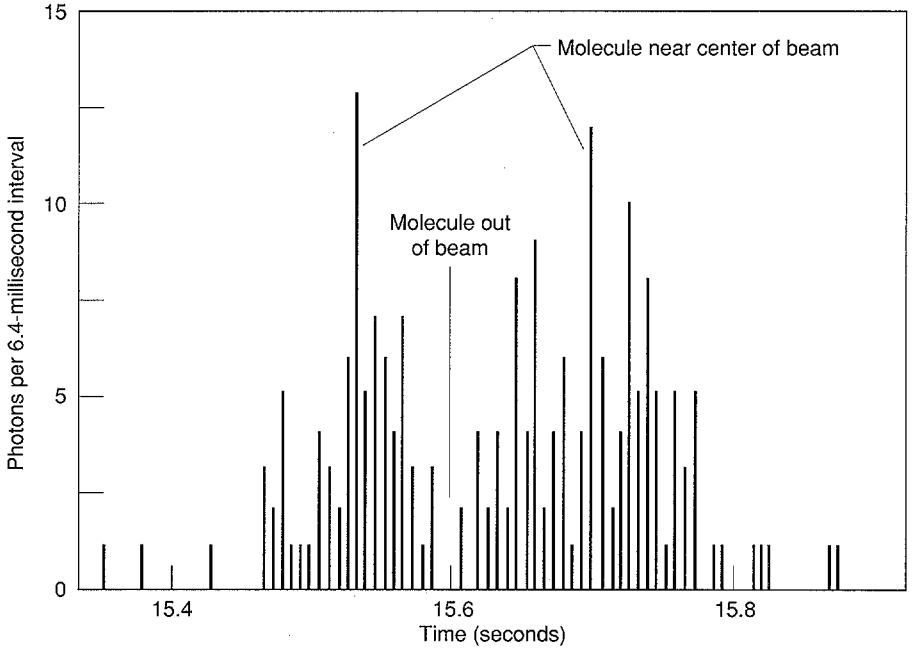

Figure 7. Photons from a Single Molecule

The bars represent the number of photons detected in 6.4-millisecond time intervals during a measurement of the fluorescence lifetime of a single molecule. The profile of the burst differs from the typical profile of fluorescence bursts shown in Figure 5 because the lifetime experiment differed from our color-discrimination experiments in having a much slower flow and reduced photobleaching. Those conditions allow molecules to wander randomly into and out of the beam while still fluorescing; the data shown suggest that a molecule entered and left the beam twice.

molecules stayed in the detection volume was determined more by diffusion than by the flow. Bursts often exhibited multiple peaks as a molecule wandered into, out of, and back into the beam. To ensure that every burst, even those with multiple peaks, came from a single molecule, we made the dye so dilute that during experiments lasting several minutes, only a few dye molecules passed through the detection volume. Between bursts, the photon detection rate was low and approximately equal to the background rate.

Figure 7 shows a typical burst consisting of about 200 photons from a Texas red molecule (compared to 10 to 15 photons in our color-discrimination experiments). From those data we can determine the fluorescence lifetime of the molecule with an accuracy of $200^{-1/2}$ or 7 percent, sufficient to distinguish many species of dyes.

Figure 8 gives the ΔT values for each of the photons that made up the burst shown in Figure 7. Since time gating is still necessary to reduce the background, we must ignore photons emitted shortly after the laser pulse. We can still determine the average lifetime of each molecule by measuring ΔT from the beginning of the time window rather than from the time of the laser pulse, because the lifetimes have an exponential distribution. Accordingly we ignore the ΔT values less than 0.7 nanoseconds and subtract 0.7 nanoseconds from all the larger values. Then the average of the values in Figure 8 is a reasonably accurate measurement of the fluorescent lifetime of that molecule. (We make a small correction for our inability to record ΔT values greater than 11 nanoseconds, when the next laser pulse begins to interfere. In principle the background makes another correction necessary, but in this case the expected background in the short duration of a burst is negligible.)

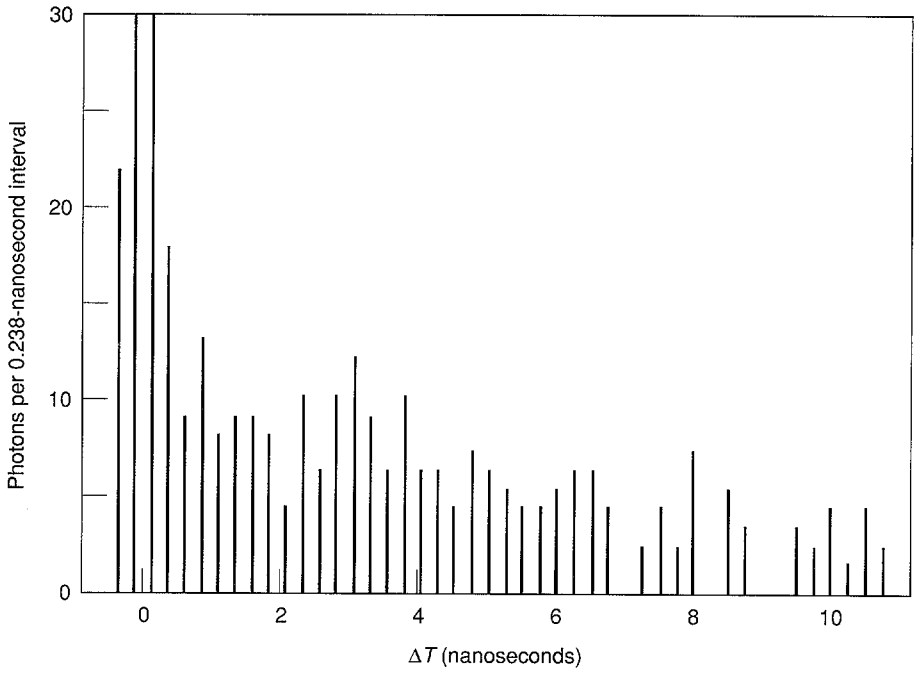

Figure 8. Lifetime of a Single Molecule

The data are from the same photon burst shown in the previous figure, but here the vertical bars give the number of photons whose emission times (measured from the previous excitation time) fell in each 0.238-nanosecond interval of ΔT. (We have not used time gating so that we can show times less than 1 nanosecond after the laser pulses, where a contribution from Raman scattering, too large for the scale of this graph, is present.) With only about 200 photons, the data follow the expected exponential distribution very roughly, but still give a reasonably accurate value for the average lifetime.

From the data plotted in Figure 8, we computed the lifetime of the molecule that produced that burst to be 4.5±0.3 nanoseconds, in agreement with the value 4.17±0.01 nanoseconds measured on bulk solutions of Texas red. Lifetimes computed for all the individual bursts clustered near the known value. Our measurements are the first determinations of fluorescent lifetimes for single molecules in solution. In future experiments such measurements could be used to identify the molecular species that produced the burst by comparison to values previously measured from bulk samples of the dye. In rapid sequencing, identification by lifetime has the advantage that one might use related dyes with similar spectra but different lifetimes, and thus one would need only one laser and one photomultiplier.

With proper experimental design, two or more independent spectroscopic properties, such as lifetimes and emission spectra, could be measured simultaneously on each passing molecule. It might also be possible eventually to measure other quantities such as photobleaching efficiency and the molecule's effectiveness at absorbing photons. Knowledge of two or more parameters would be useful in cases where the value of a single parameter is insufficient to make a definitive identification.

Other Applications of Spectroscopy

We have discussed single-molecule spectroscopy primarily from the point of view of identification of molecular species, as when a dye molecule is attached as an identification tag to a nonfluorescent molecule, for example a DNA base. However, data collected by observing individual molecules can reveal features that are not evident in the average behavior of a group. For example, in a group of dye molecules, one in fifty might be bound to another molecule that diminishes its quantum yield, causing the average yield to be slightly less than the true value. A bulk experiment would reveal only the average value, whereas the true yield as well as the statistics underlying the decreased average would be readily accessible using single-molecule spectroscopy.

Moreover, to the extent that spectroscopic properties are modified by the immediate environment of a molecule, measuring those properties can supply information about that environment on a microscopic scale (provided the environment changes little in the time needed for the measurement). For example, fluorescence lifetime can be used to measure distances on the atomic scale. An excited fluorescent molecule (donor) can lose its excitation energy to a nearby acceptor molecule if the donor's range of emission energies overlaps the acceptor's range of absorption energies. (See the depiction of quenching in Figure 1.) Such energy transfer reduces the fluorescence lifetime of the donor. Because the probability of energy transfer depends strongly on the distance r between donor and acceptor (as r^{-6}), measurement of the donor's fluorescence lifetime provides a measurement of its distance from the acceptor. We hope to use this molecular-level "yardstick" to determine distances that are inaccessible by other means, particularly in biological systems. Such potential applications, as well as the promise of rapid DNA sequencing, will maintain our interest in developing single-molecule spectroscopy. ■

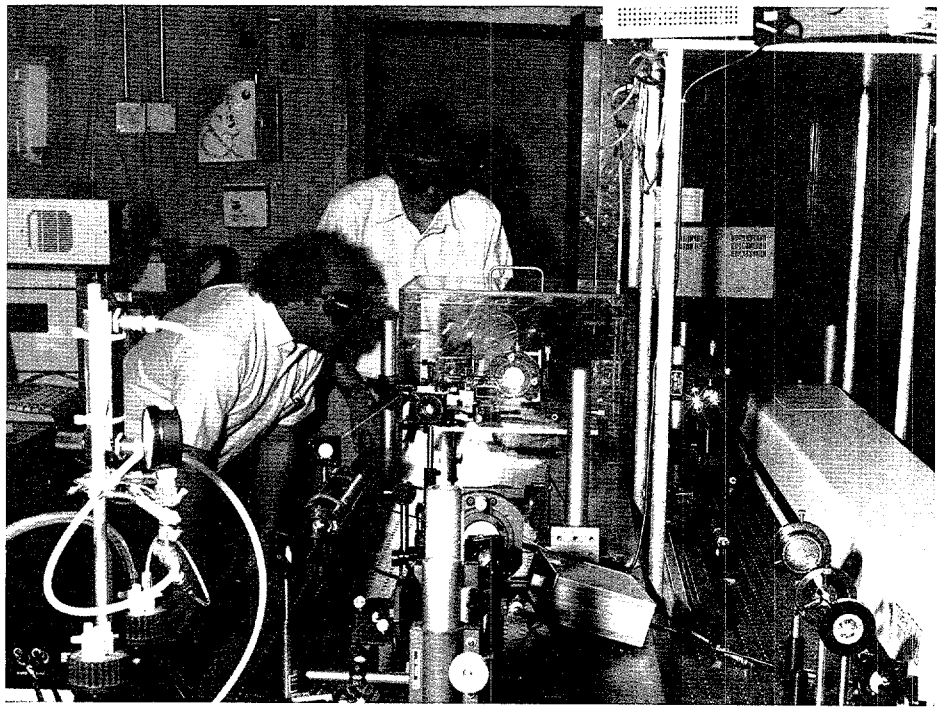

Further Reading

R. Peierls. 1935. Statistical errors in counting experiments. *Proceedings of the Royal Society (London)* A149:467.

E. B. Shera, N. K. Seitzinger, L. M. Davis, R. A. Keller, and S. A. Soper. 1990. Detection of single fluorescent molecules. Chemical Physics Letters 175:553–557.

S. A. Soper, L. M. Davis, and E. B. Shera. 1992. Detection and identification of single molecules in solution. Accepted for publication in *Journal of the Optical Society of America* B.