1. OVERTURE

This is a book about **cytometry**, in general, emphasizing **flow cytometry**, in particular. In it, I hope to tell you what cytometry is, how it works, why and how to use it, when you should favor one type of cytometry or another, and when cytometry won't solve your problem. This chapter, like the overture to an opera or a musical, presents important themes from the body of the work, but may also stand alone.

1.1 WHAT (AND WHAT GOOD) IS CYTOMETRY?

Cytometry is a process in which physical and/or chemical characteristics of **single cells**, or by extension, of other biological or nonbiological particles in roughly the same size range, are **measured**. In **flow cytometry**, the measurements are made as the cells or particles pass through the measuring apparatus, a **flow cytometer**, in a fluid stream. A **cell sorter**, or **flow sorter**, is a flow cytometer that uses electrical and/or mechanical means to divert and collect cells (or other small particles) with measured characteristics that fall within a user-selected range of values.

Neither the cells nor the apparatus are capable of putting the process of cytometry in motion; the required critical element for that is a human interested in obtaining information about a cell sample and, in the case of sorting, extracting cells of interest from the sample. At the most basic level, a cytometer might be considered to be a "black box" with cells as "inputs" and numbers as "outputs"; the outputs of a cell sorter would include both numbers and cells. However, while some modern cytometers (and some modern users) can obtain the desired results while running unattended in "black box" mode, it is fair to say that most of the applications, and all of the interesting applications, of cytometry call for some understanding and some intellectual effort on the part of the user.

Tasks and Techniques of Cytometry

From the time of van Leeuwenhoek and Hooke until the mid-20th century, determining:

1) whether cells were present in a specimen,
2) how many were there,
3) what kinds of cells were represented, and
4) what their functional characteristics might be

required that a human observer interpret a microscope image. The same tasks remain for modern cytometry.

Although electrical and acoustic properties of, and nuclear radiation emission from, single cells can be measured, it is fair to say that **optical** measurements are by far the most common in cytometry. A typical cytometer is thus a specialized **microscope**; the degree of physical resemblance is dictated by the requirements of the measurement(s) to be made, which in turn are dictated by what the user needs to know about the cell sample. In successful applications of cytometry, electro-optics, electronics, and computers are employed to improve on what could be obtained "by eye," although interpretation is required more often than not. The successful applications are many, increasing in number, and commonplace in locales as diverse as clinical laboratories and breweries.

Some Notable Applications

Cytometry is currently used to obtain the helper T lymphocyte counts needed to monitor the course and treatment of HIV infection, and to determine tumor cell DNA content

and proliferative activity, which may aid in assessing prognosis and determining treatment for patients with breast cancer and other malignant diseases. The technology has also been used to crossmatch organs for transplantation, to isolate human chromosomes for the construction of genetic libraries, to separate X- and Y-chromosome bearing sperm for sex selection in animal breeding and *in vitro* fertilization in humans, to identify the elusive hematopoietic stem cell and an expanding family of other stem cell types, and to reveal several widely distributed but previously unknown genera of marine microorganisms.

Biological particles that have been subjected to cytometric analysis range, in order of decreasing size, from multicellular organisms (e.g., *Drosophila* embryos and adult *Caenorhabditis elegans* nematodes) through cell aggregates (e.g., pancreatic islets and tumor cell spheroids), eukaryotic cells, cellular organelles (e.g., mitochondria), bacteria, liposomes, individual virus particles and immune complexes, down to the level of single molecules of proteins, nucleic acids, and organic dyes. Cytometers can also be used for sensitive chemical analyses involving the binding of suitably labeled ligands to solid substrates or to particles such as polystyrene beads.

The first practical applications of flow cytometry, beginning in the 1940's, were to counting blood cells in liquid suspension, on the one hand, and bacteria and other small particles in aerosols, on the other, based on measurements of **light scattering** or **electrical impedance**; these signals were also used to provide estimates of cell size.

In the early 1960's, **light absorption** measurements were used for quantitative flow cytometric analyses of cellular nucleic acid and protein. Flow cytometers in modern clinical hematology laboratories perform counts of red cells (erythrocytes), white cells (leukocytes), and platelets (thrombocytes) in blood, as well as differential leukocyte counts, using combinations of electrical impedance, light scattering, and light absorption measurements.

However, many people who know the term "flow cytometer" tend to use it – incorrectly – to describe only instruments that measure **fluorescence** as well as light scattering. The first fluorescence flow cytometers were built in the late 1960's; although there are now well over 10,000 in use in clinical and research laboratories worldwide, they are still outnumbered by impedance and scattering-based hematology analyzers. So much for fluorescence chauvinism.

What is Measured: Parameters and Probes

The novice should not be intimidated by the jargon of cytometry; there are no native speakers, and he or she can soon enough become as fluent in it as the rest of us. The term **parameter** is, unfortunately, used in several different senses in our jargon. It can refer to a **physical or chemical characteristic of a cell** (e.g., cytoplasmic granularity or nuclear DNA content) that is measurable by cytometry; it can also describe a **physical property, measured by a sensor**, defined broadly (e.g., light scattering or fluorescence), or more narrowly (e.g., orthogonal light scattering or red fluorescence), or a **physical property of a cell-associated reagent** (e.g., propidium fluorescence). A fairly comprehensive list of measurable cellular parameters appears as Table 1-1 on the facing page.

I have characterized cellular parameters as **intrinsic** or **extrinsic**, depending upon whether they can or cannot be measured without the use of **reagents**, which are often referred to in cytometric jargon as **probes**. Some parameters can, at least in principle, be measured either with or without probes; cellular DNA content, for example, can be estimated from ultraviolet (UV) absorption at 260 nm in unstained cells, but it's much more practical to use a fluorescent dye probe such as propidium iodide. A deeper philosophical dilemma arises when considering fluorescence from *Aequorea* green fluorescent protein (GFP) or one of its genetically engineered offshoots, introduced by cloning into cells of other species to report gene expression; one could characterize this as intrinsic or extrinsic, but I lean toward the latter.

Parameters can also be defined as **structural** or **functional**, again with some ambiguity. For example, the glycoprotein efflux pump responsible for multidrug resistance in tumor cells can be detected, and the amount present in a cell quantified, using fluorescent antibodies, but such antibodies might also bind to an inactive mutant protein, and thus provide a measurement (in this case, inaccurate) based on structure. The function of the glycoprotein pump can be demonstrated by measurement of uptake or loss of fluorescent drugs or dyes by cells over periods of time.

In a **kinetic measurement** such as that just described, **time** itself can be used as a parameter. When such analyses are done by flow cytometry, the dynamic behavior of a cell population must be inferred from observations of different cells at different times, because conventional flow cytometers cannot make successive measurements of a single cell over time periods exceeding a few microseconds.

Both the novice and the expert in flow cytometry should be aware that almost every parameter that can be measured by flow cytometry can also be measured by alternative cytometric methods such as **microspectrophotometry**, **confocal microscopy**, **image analysis**, and **scanning cytometry**. These methods are often applicable where flow cytometric methods are not, e.g., for true kinetic analyses involving repeated examination of the same cell or cells over a period of time, or for *in situ* analyses of cells growing in aggregates attached to solid substrates. In general, the fluorescent probes used for flow cytometry can be used with alternative measurement techniques. However, most dyes and other reagents that are commonly employed in absorption microspectrophotometry are not readily usable in fluorescence flow cytometers.

1.2 BEGINNINGS: MICROSCOPY AND CYTOMETRY

It recently (i.e., since the last time I wrote an introduction to cytometry) occurred to me that the best way in

PARAMETER	MEASUREMENT METHOD AND PROBE IF USED
Intrinsic Structural Parameters (no probe)	
Cell Size	Electronic (DC) impedance, extinction, small angle light scattering; image analysis
Cell shape	Pulse shape analysis (flow); image analysis
Cytoplasmic granularity	Large angle light scattering, Electronic (AC) impedance
Birefringence (e.g., of blood eosinophil granules)	Polarized light scattering, absorption
Hemoglobin, photosynthetic pigments, porphyrins	Absorption, fluorescence, multiangle light scattering
Intrinsic Functional Parameter (no probe)	
Redox state	Fluorescence (endogenous pyridine and flavin nucleotides)
Extrinsic Structural Parameters (probe required)	
DNA content	Fluorescence (propidium, DAPI, Hoechst dyes)
DNA base ratio	Fluorescence (A-T and G-C preference dyes, e.g., Hoechst33258 and chromomycin A_3)
Nucleic acid sequence	Fluorescence (labeled oligonucleotides)
Chromatin structure	Fluorescence (fluorochromes after DNA denaturation)
RNA content (single and double-stranded)	Fluorescence (acridine orange, pyronin Y)
Total protein	Fluorescence (covalent- or ionic-bonded acid dyes)
Basic protein	Fluorescence (acid dyes at high pH)
Surface/Intracellular antigens	Fluorescence; scattering (labeled antibodies)
Surface sugars (lectin binding sites)	Fluorescence (labeled lectins)
Lipids	Fluorescence (Nile red)
Extrinsic Functional Parameters (probe required)	
Surface/intracellular receptors	Fluorescence (labeled ligands)
Surface charge	Fluorescence (labeled polyionic molecules)
Membrane integrity (not always a sign of "viability")	Fluorescence (propidium, fluorescein diacetate [FDA]); absorption or scattering (Trypan blue)
Membrane fusion/turnover	Fluorescence (labeled long chain fatty acid derivatives)
Membrane organization (phospholipids, etc.)	Fluorescence (annexin V, merocyanine 540)
Membrane fluidity or microviscosity	Fluorescence polarization (diphenylhexatriene)
Membrane permeability (dye/drug uptake/efflux)	Fluorescence (anthracyclines, rhodamine 123, cyanines)
Endocytosis	Fluorescence (labeled microbeads or bacteria)
Generation number	Fluorescence (lipophilic or covalent-bonded tracking dyes)
Cytoskeletal organization	Fluorescence (NBD-phallacidin)
Enzyme activity	Fluorescence; absorption (fluorogenic/chromogenic substrates)
Oxidative metabolism	Fluorescence (dichlorofluorescein)
Sulfhydryl groups/glutathione	Fluorescence (bimanes)
DNA synthesis	Fluorescence (anti-BrUdR antibodies, labeled nucleotides)
DNA degradation (as in apoptosis)	Fluorescence (labeled nucleotides)
"Structuredness of cytoplasmic matrix"	Fluorescence (fluorescein diacetate [FDA])
Cytoplasmic/mitochondrial membrane potential	Fluorescence (cyanines, rhodamine 123, oxonols)
"Membrane-bound" Ca^{++}	Fluorescence (chlortetracycline)
Cytoplasmic [Ca^{++}]	Fluorescence ratio (indo-1), fluorescence (fluo-3)
Intracellular pH	Fluorescence ratio (BCECF, SNARF-1)
Gene expression	Fluorescence (reporter proteins)

Table 1-1. Some parameters measurable by cytometry.

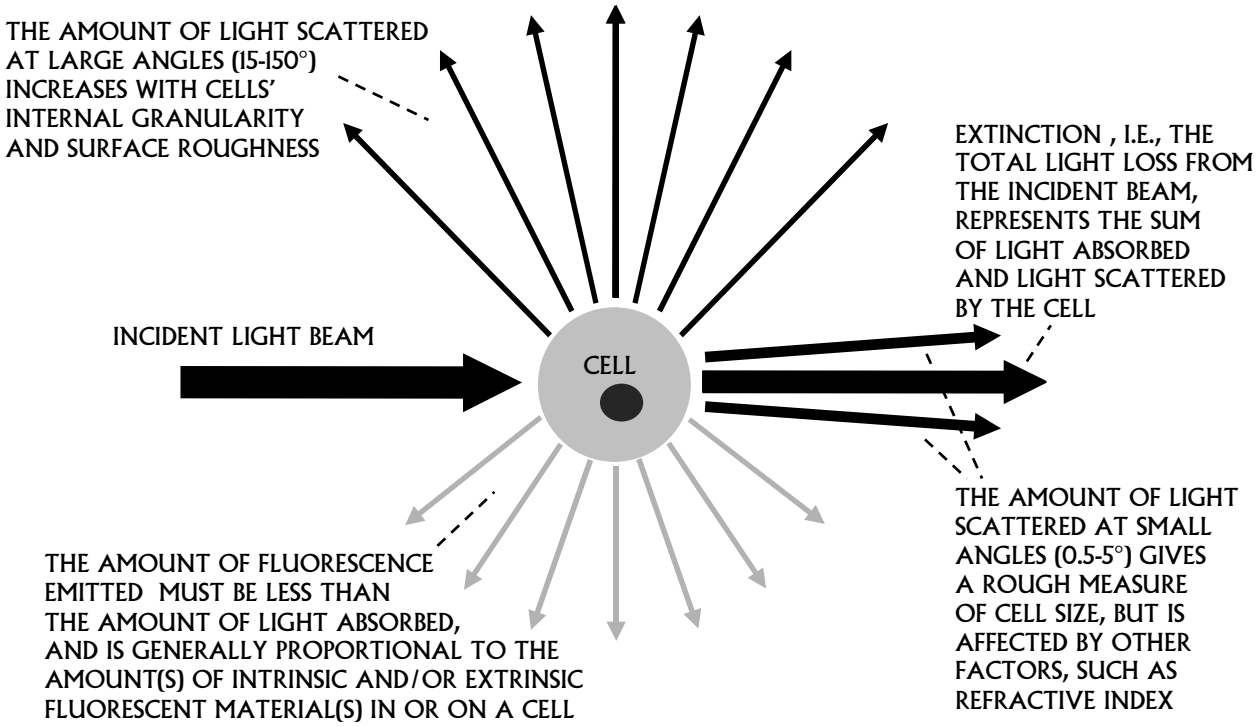

Figure 1-1. Interaction of light with a cell.

which to introduce the subject might be to consider how cytometry developed from microscopy, emphasizing both the similarities and the differences between the two, and stressing how the information gets from the cells to the user. That is what I will try to do in the remainder of this chapter. I hope this will be helpful for the uninitiated reader, but, also, that it will be equally thought-provoking, informative, and at least moderately amusing to those who have been over the terrain one or many times before.

The first order of business in both microscopy and cytometry is discriminating between the cells and whatever else is in the sample; the next is often discriminating among a number of different cell types that may be present. Optical microscopes first allowed cells to be discovered and described in the seventeenth century, and were refined in design in the eighteenth and early nineteenth, but the capacity of microscopy to discriminate among different cell types remained limited by the relative difficulty of obtaining **contrast** between cells and the background in microscope images.

A Little Light Music

While all the senses can provide us with pleasure and discomfort, it is predominantly vision that shapes our perception of the world around us, and, without light, our visual imagery is restricted to memories, dreams, and hallucinations. According to the Book of Genesis, the discrimination of light from darkness is the divine achievement of the first day of creation, and we humans, despite taming fire and inventing light bulbs and lasers, remain aware of and profoundly affected by the daily difference, not least during power outages.

What most of us know as **light** is defined by physicists as **electromagnetic radiation** with wavelengths ranging between about 400 and about 700 nanometers (nm). Other species can detect shorter and longer wavelengths, but most lack our ability to discriminate among wavelengths, i.e., color vision, and some of us have genetic deficiencies that restrict this capacity.

When we look at the macroscopic world, most of our retinal images are formed by light that we say is **reflected** from objects around us, and an early concept of light was that of **rays** traveling in straight lines, and reflecting from a surface at the same angle at which they strike it. If we look at an object under water and attempt to grab it, we find that it is not exactly where it appears to be; this is explained by the concept of **refraction**, according to which light passing from one material medium into another is bent at an angle depending on a macroscopic property of the medium known as the **refractive index**, and on the wavelength of the light. The "white" light emitted by the sun and by incandescent and fluorescent bulbs comprises a range of visible wavelengths; objects and materials that **absorb** some, but not all, wavelengths reflect others, and thus appear colored.

As we turn our attention to smaller and smaller objects, the concepts of reflection and refraction become less and less useful, and we instead make use of the concept of **light scattering**. Figure 1–1 describes the interaction of light with a

cell in terms of **scattering**, **absorption**, and **fluorescence**. The last of these phenomena is not readily explicable in terms of either ray (geometrical) or wave optics, and can only be dealt with properly by the theory of **quantum electrodynamics**, which considers light as particles, or **photons**, which interact with **electrons** in atoms and molecules. The energy of a photon is inversely proportional to the corresponding wavelength; i.e., photons of short-wavelength, 400 nm violet light have a higher energy content than photons of long-wavelength, 700 nm red light.

Scattering, which explains both reflection and refraction, typically involves a brief interaction between a photon and an electron, in which the photon is annihilated, transferring its energy to the electron, which almost immediately releases all of the energy in the form of a new photon. Thus, light scattered by an object has the same (or almost exactly the same) wavelength, or color, as the incident light. However, the new photon does not necessarily travel in the same direction as the old one, so scattered light usually appears to be at an angle to the incident beam.

In empty space, there are, by definition, no atoms or molecules, and there are thus no electrons available to interact with photons. Although, according to quantum electrodynamics, a photon has a finite probability of going in any direction, when we actually calculate the probabilities that apply in the case of photons in empty space, we come up with what look like rays of light traveling in straight lines.

As a general rule, the density of atoms and molecules in atmospheric air is fairly low, meaning that there are few opportunities for light to be scattered as it appears to traverse distances of a few meters or tens of meters. However, we note the blue appearance of a cloudless sky, resulting from light scattering throughout the atmosphere; the color results from the fact that shorter wavelengths of light are more likely to be scattered than longer ones, with the intensity of scattering inversely proportional to the fourth power of the wavelength.

The well-known laws of reflection and refraction emerge from quantum electrodynamics applied to objects substantially bigger than the wavelength of light. Materials that appear transparent to the human eye, e.g., glass and water, still contain relatively high densities of atoms and molecules, and thus provide numerous opportunities for scattering.

Some light appears to be reflected at the interfaces between layers of different materials, with the angle of reflection equal to the angle of incidence. The total amount of light reflected is found to be a function of the thickness of the layers and the wavelength of the incident light; that is, layers of different thicknesses reflect different colors of light to different extents. This **interference** effect, explained by the theory of wave optics, accounts for the patterns of color seen in peacock feathers, butterfly wings, diffraction gratings in spectrophotometers, on credit cards, and in cheap jewelry, and in opals in somewhat more expensive jewelry. It is exploited in optical design, notably in the production of **interference filters** used to select ranges of wavelengths to be observed and/or detected in microscopes and other optical instruments. Quantum electrodynamics comes up with the same results for interference and reflection as wave optics, even while taking into account that the phenomena are due to scattering throughout objects, not just from front and back surfaces.

The apparent bending of light striking an interface between two materials is described in classical optics with the aid of invented quantities, called **refractive indices**, which are characteristic of the materials involved. Light appears to travel more slowly through a material of higher refractive index than through a material of lower index, and a "ray" appears to "bend" toward the normal (i.e., toward a line perpendicular to the interface) when passing from a lower-index medium to a higher one, and away from the normal when passing from a higher-index medium to a lower one. The apparent velocity of light in a material is less than in empty space; the higher the refractive index, the lower the apparent velocity. Light of a shorter wavelength is "bent" more than light of a longer one, allowing a transparent object with surfaces that are not parallel (i.e., a **prism**) to **disperse** light of different wavelengths in different directions.

Armed with ray optics and the classical law of refraction, we can calculate how an object with appropriately curved surfaces, i.e., a **lens**, will "bend" light originating from two points separated in space. If the surfaces are convex, divergent "rays" coming through the lens from two points a given distance apart on the "input" side can be made to converge at two points a greater distance apart on the "output" side; this provides us with a **magnified image**. A magnifying lens is, of course, the fundamental ingredient of a microscope.

Not surprisingly, everything useful that classical optics tells us about refraction can be obtained using quantum electrodynamics. Although actually doing this usually involves a great deal of advanced mathematics, Richard Feynman, who received his Nobel Prize for work in the field, wrote a small book called *QED*[641], in which he used simple diagrams and concepts to make the subject accessible to a lay audience (which, in this context, includes me). What I am writing here paraphrases the master.

The light scattering behavior of objects of dimensions near the wavelength of light is not predictable from ray optics. For spherical particles ranging in diameter from one or two wavelengths to a few tens of wavelengths, most of the light scattering occurs at small angles (0.5° to 5°) to the incident beam; the intensity of this "**small angle**," or "**forward**," light scattering is dependent on the refractive index difference between the particle and the medium, and on particle size. However, the relationship between particle size and small angle scattering intensity is not monotonic, meaning that, although a particle 10 μm in diameter will probably produce a bigger signal than one of the same composition 5 μm in diameter, a particle 5.5 μm in diameter might produce a smaller signal than one 5 μm in diameter. It is thus wise to avoid thinking of the small angle scatter signal as an accurate measure of cell size.

Smaller particles scatter proportionally more light at larger angles (15° to about 150°) to the incident beam; the amplitude of such signals, variously described as "**side**," "**orthogonal**," "**large angle**," "**wide angle**," or "**90°**" light scattering, is, all other things being equal, larger for cells with internal granular structure, such as blood granulocytes, than for cells without it, such as blood lymphocytes.

Ray optics and wave optics break down when we consider the process of **light absorption**. This comes down to photons and electrons, period. Quantum theory tells us that the electrons in a given atom or molecule can exist only in discrete energy states. The lowest of these is referred to as the **ground state**, and the absorption of a photon by an electron in the ground state raises it to a higher energy **excited state**. An electron in an excited state can absorb another photon, ending up in a still higher energy excited state.

Like scattering, and all other quantum phenomena, absorption is probabilistic. We cannot say that a particular electron will absorb a particular photon; the best we can do is calculate the probability that an electron in a particular energy state will absorb a photon of a particular energy, or wavelength. This probability increases as the difference in energy between the current energy state of the electron and the next higher energy state gets closer to the energy of the photon involved.

In many molecules, the energy difference between states is greater than the energy in a photon of visible light. Such molecules may exhibit substantial absorption of higher energy, shorter wavelength photons, e.g., those with wavelengths in the **ultraviolet (UV)** region between about 200 and 400 nm. Substances made up of such molecules appear transparent to the human eye; smearing them on exposed skin decreases the likelihood that ultraviolet photons will interact with electrons in DNA and other macromolecules of dermal cells, and reduces the likelihood of sunburn (yay!) and tanning (boo!). We're not sure yet about skin cancer.

For a molecule to absorb light in the visible region, the energy differences between electronic energy states have to be rather small. This condition is satisfied in some inorganic atoms and crystals, which have unpaired electrons in *d* and *f* orbitals, in metals, which have large numbers of "free" electrons with an almost continuous range of energy states, resulting in high absorption (and high reflectance) across a wide spectral range, and in organic molecules with large systems of conjugated π orbitals, including natural products such as porphyrins and bile pigments, and synthetic dyes such as those used to stain cells.

The interaction of light with matter must obey the law of conservation of energy; the amount of light transmitted should therefore be equal to the amount of incident light minus the amount scattered and the amount absorbed. But what happens to the absorbed light? One would not expect the electrons involved in absorption to remain in the excited state indefinitely, and, indeed, they do not. In some cases, all of the absorbed electronic energy is converted to vibrational or rotational energy, and lost as heat. In others, some energy is lost as heat, but the remainder is emitted in the form of photons of lower energy (and, therefore, longer wavelength) than those absorbed. Depending on the details of the electronic energy transitions involved, this emission can occur as **fluorescence** or as **phosphorescence**. Fluorescence emission usually occurs within a few tens of nanoseconds of absorption; phosphorescence is delayed, and may continue for seconds or longer. As is the case with absorption, fluorescence and phosphorescence are inexplicable by ray and wave optics; they can only be understood in terms of quantum mechanics.

Making Mountains out of Molehills: Microscopy

When we are not looking at luminous displays such as the one I face as I write this, most of our picture of the world around us comes from reflected light. Contrast between objects comes from differences in their reflectivities at the same and/or different wavelengths. When ambient light levels are high, we utilize our retinal cones, which give us color vision capable of prodigious feats of spectral discrimination (humans with normal vision can discriminate millions of colors), at the expense of relatively low sensitivity to incident light. The high light levels bleach the visual pigments in our more sensitive retinal rods; if the light level is decreased abruptly, it takes some time for the rod pigment to be replenished, after which we can detect small numbers of photons, sacrificing color vision in the process. Thus, while we can perceive large numbers of 450 nm photons, 550 nm photons, and 650 nm photons, respectively, as red, green, and blue light, using our cones, we cannot distinguish individual photons with different energy levels as different colors. Night vision equipment typically utilizes monochromatic green luminous displays because the rods are most sensitive to green light, but the cone system also exhibits maximum sensitivity in the green region, making the spot from a green laser pointer much more noticeable than that from a red one emitting the same amount of power.

While the spectral discrimination capabilities of the unaided human visual system are remarkable, its spatial discrimination power is somewhat limited. The largest biological cells, e.g., ova and large protists, are just barely visible, and neither the discovery of cells nor the appreciation of their central role in biology would have occurred had the light microscope not been invented and exploited.

When unstained, unpigmented cells are examined in a traditional **transmitted light**, or **bright field**, microscope, light absorption is negligible; contrast between cells and the background is due solely to scattering of light by cells and subcellular components, and the only information we can get about the cells is thus, in essence, contained in the scattered light. Some of this is scattered out of the field of view; we must therefore rely on slight differences in transmission between different regions of the image to detect and characterize cells. We are working against ourselves by presenting our eyes (or the detector(s) in a cytometer) with a large amount of light that has been transmitted by the specimen.

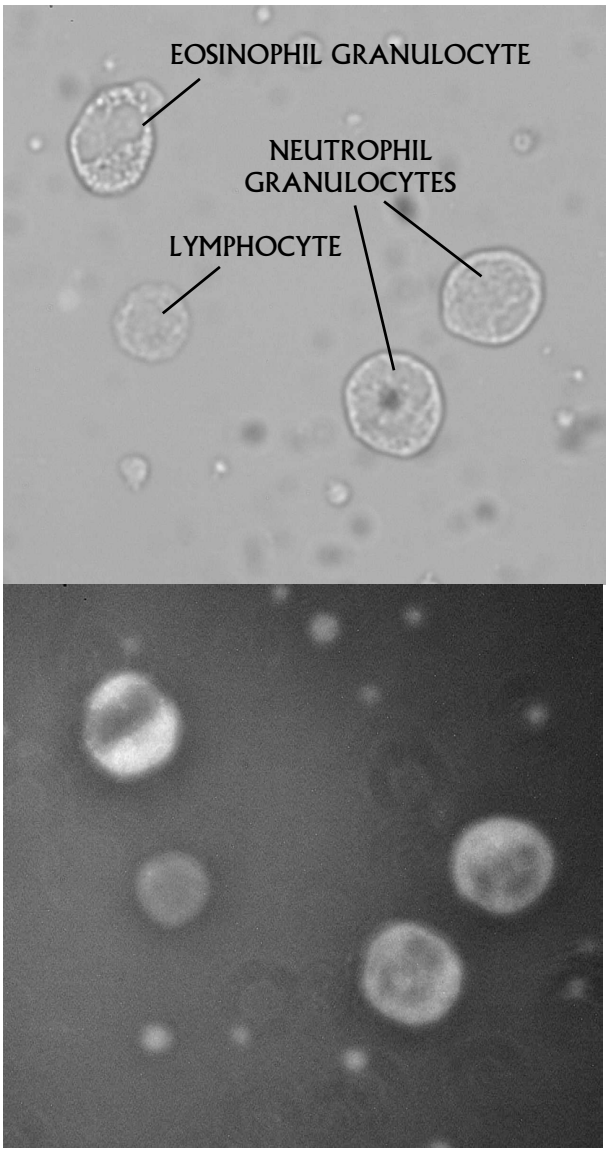

Figure 1-2. Transmitted light (bright field) (top panel) and dark field (bottom panel) images of an unstained suspension of human peripheral blood leukocytes. The objective magnification was 40 ×.

As it happens, the maximum spatial resolution of a microscope is achieved, i.e., the distance at which two separate objects can be distinguished as separate is minimized, when illuminating light reaches, and is collected from, the specimen at the largest possible angle. The **numerical aperture (N.A.)** of microscope condensers and objectives is a measure of the largest angle at which they can deliver or collect light. However, when the illumination and collection angles in a transmitted light microscope are large, much of the light scattered by objects in the specimen finds its way back into the microscope image, increasing resolution, but decreasing contrast. The top panel of Figure 1-2 shows a bright field microscope image of a suspension of human peripheral blood leukocytes; the condenser was stopped down to increase contrast between the cells and background. The cytoplasmic granules in the eosinophil and neutrophil granulocytes are not particularly well resolved, nor is it easy to distinguish the nuclei from the cytoplasm. Increasing the level and angle of illumination might, as just mentioned, increase resolution, but this would not be useful, as contrast would not be increased.

Modern microscopy exploits both differences in **phase** and **polarization** of **transmitted light** and the phenomenon of **interference** to produce increased contrast in bright field images. However, **staining**, which came into widespread use in the late 1800's, largely due to the emergence of synthetic organic dyes, was the first generally applicable practical bright field technique for producing contrast between cells and the medium, and between different components of cells in microscope images. Paul Ehrlich, known for his later researches on chemotherapy of infectious disease, stained blood cells with mixtures of acidic and basic dyes of different colors, and identified the three major classes of blood granulocytes, the basophils, eosinophils (which he termed acidophils), and neutrophils, based on the staining properties of their cytoplasmic granules.

Stained elements of cells are visually distinguishable because of their **absorption** of incident light, even when the refractive index of the medium is adjusted to be equal or nearly equal to that of the cell. The dyed areas transmit only those wavelengths they do not absorb, resulting in a difference in spectrum, or color, between them and undyed areas or areas that take up different dyes. Absorption by pigments within cells, such as the hemoglobin in erythrocytes, also makes the cells more distinguishable from the background.

Microscopy of opaque specimens, such as samples of minerals, obviously cannot use transmitted light bright field techniques. Instead, specimens are illuminated from above, and the image is formed by light reflected (i.e., scattered) from the specimen. In **incident light bright field microscopy**, illumination comes through the objective lens, using a partially silvered mirror, or **beam splitter**, to permit light to pass between source and specimen and between specimen and eyepiece at the same time. In **dark field microscopy**, illumination is delivered at an oblique angle to the axis of the objective by a separate set of optics. The bottom panel of Figure 1-2 is a dark field image of the same cells as are shown in the top panel. In the dark field microscope, none of the illuminating light can reach the objective unless it is scattered into its field of view by objects in the specimen. The illumination geometry used in this instance ensured that the only light contributing to the dark field image was light scattered at relatively large angles to the illuminating beam. It has already been noted that this is the light represented in the side scatter signal, and it can be seen that the lymphocyte, which would have the smallest side scatter signal, appears dimmer than the neutrophil granulocytes and the eosinophil, which would have higher side scatter signals. Although the cytoplasmic granules within the granulocytes are not well resolved, the intensity of light coming from the

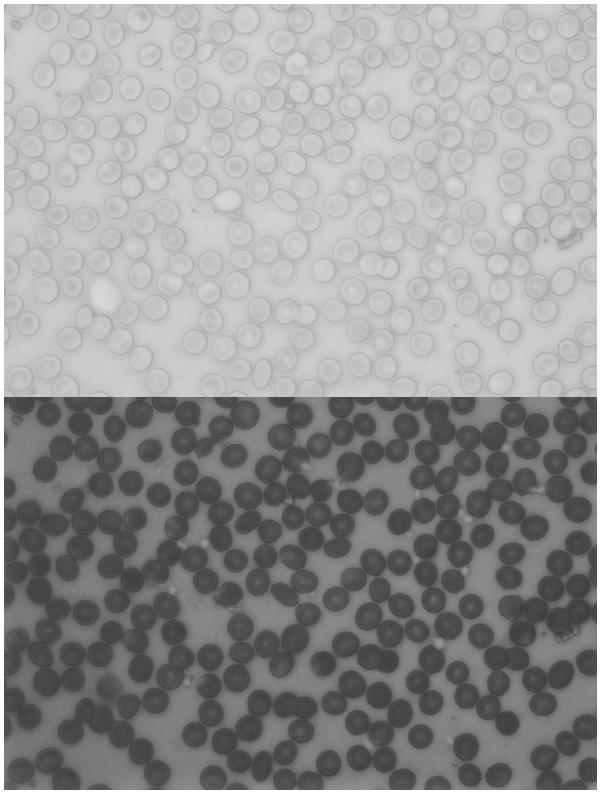

Figure 1-3. Transmitted light microscope images of an unstained smear of human peripheral blood. The picture in the top panel was taken with "white light" illumination; that in the bottom panel was taken with a violet (405 nm, 15 nm bandwidth) band pass optical filter, and demonstrates the strong absorption of intracellular hemoglobin in this wavelength region. Objective: 40 ×.

cytoplasm provides an indication of their presence; indeed, it is much easier to resolve nucleus from cytoplasm in the granulocytes in the dark field image than in the bright field image. Thus, we can surmise that it may be possible to get information about subcellular structures from a cytometer operating at an optical resolution that would be too low to allow them to be directly observed as discrete objects. In fact, using dark field microscopy, one can observe light scattered by, and fluorescence emitted from, particles well below the limit of resolution of an optimally aligned, high-quality optical microscope; the dark field "ultramicroscope" of the 1920's allowed researchers to see and count viruses, although it was obviously impossible to discern any structural detail.

Absorption measurements are bright field measurements, and they work best, especially for quantification, when the absorption signal is strong. The material being looked for should have a high likelihood of absorbing incident light, as indicated by a high **molar extinction coefficient**, and there should be a lot of it in the cell. Figure 1-3 shows the absorption of hemoglobin in the cytoplasm of unstained red blood cells. Note that the "white light" image in the top panel gives little hint of strong absorption, which is restricted to the violet region known as the Soret band; the "white" light used here, which came from a quartz-halogen lamp, contains very little violet, and the exposure time used for the picture in the bottom panel was about 100 times as long as that for the picture in the top panel.

Fluorescence microscopy is inherently a dark field technique; even in a "transmitted light" fluorescence microscope, **optical filters** are employed to restrict the spectrum of the illuminating beam to the shorter wavelengths used for fluorescence **excitation**, and also to allow only the longer-wavelength fluorescence **emission** from the specimen to reach the observer. As is the case in dark field microscopy, fluorescent cells (ideally) appear as bright objects against a dark background.

Most modern fluorescence microscopes employ the optical geometry shown in Figure 1-4. Excitation light is usually supplied by a mercury or xenon arc lamp or a quartz-halogen lamp, equipped with a lamp condenser that **collimates** the light, i.e., produces parallel "rays." These components are not shown, but would be to the left of the excitation filter in the figure. The excitation filter passes light at the excitation wavelength, and reflects or absorbs light at other wavelengths. The excitation light is then reflected by a **dichroic mirror**, familiarly known simply as a **dichroic**, which transmits light at the emission wavelength. The microscope objective is used for both illumination of the specimen and collection of fluorescence emission, which is transmitted through both the dichroic and the emission filter.

In any microscope, a real image of the specimen is formed by the objective lens; the eyepiece and the lens of the observer's eye then project an image of this image onto the retina of the observer. Light falling on sensitive cells in the retina produces electrical impulses that are transmitted along the optic nerves. What happens next is the province of neurology, psychology, and, possibly, psychiatry.

It has already been noted that humans are very good at color discrimination, and we also know that humans, with some training, can get pretty good at discriminating cells from other things. With more training, we can become proficient at telling at least some kinds of cells from others, usually on the basis of the size, shape, color, and texture of cells and their components in microscope images; it is not always easy to program computers to make the same distinctions on the same basis.

The human visual system can detect light intensities that vary over an intensity range of more than nine decades; in other words, the weakest light we can perceive is on the order of one-billionth the intensity of the strongest perceptible light. However, we can't cover the entire range at once; as previously mentioned, we need dark-adapted rods to see the least intense signals, and do so only with monochromatic vision. And we aren't very good at detecting small changes in light intensity. This has forced us to invent instruments to make precise light intensity measurements to meet the needs

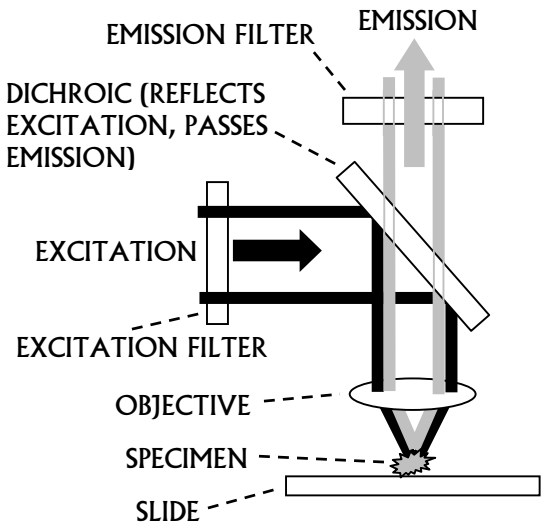

Figure 1-4. Schematic of a fluorescence microscope.

of science, technology, medicine, and/or art (remember when the light meter was not built into the camera?). It was this process that eventually got us from microscopy to cytometry.

Why Cytometry? Motivation and Machinery

In the 1930's, by which time the conventional histologic staining techniques of light microscopy had already suggested that tumors might have abnormalities in DNA and RNA content, Torbjörn Caspersson[34], working at the Karolinska Institute in Stockholm, began to study cellular nucleic acids and their relation to cell growth and function. He developed a series of progressively more sophisticated **microspectrophotometers**, which could make fairly precise measurements of DNA and RNA content based on the strong intrinsic UV absorption of these substances near 260 nm, and also found that UV absorption near 280 nm, due to aromatic amino acids, could be used to estimate cellular protein content. When Caspersson began working, it had not yet been established that DNA was the genetic material; he helped move others toward that conclusion by establishing, through precise measurement, that the DNA content of chromosomes doubled during cellular reproduction[7].

A conventional optical microscope incorporates a **light source** and associated optics that are used to illuminate the specimen under observation, and an **objective** lens, which collects light transmitted through and/or scattered, reflected and/or emitted from the specimen. Some means are provided for moving the specimen and adjusting the optics so that the specimen is both properly illuminated and properly placed in the field of view of the objective. In a microscope, a mechanical stage is used to position the specimen and to bring the region of interest into focus.

A microspectrophotometer was first made by putting a small "pinhole" aperture, or **field stop**, in the image plane of a microscope, restricting the field of view to the area of a single cell, and placing a photodetector behind the field stop. The diameter of the field stop could be calculated as the product of the magnification of the objective lens and the diameter of the area from which measurements were to be taken. If a 40× objective lens were used, measuring the transmission through, or the absorption of, a cell 10 μm in diameter would require a 400 μm diameter field stop.

Using a substantially smaller field stop, it would be possible to measure the transmission through a correspondingly smaller area of the specimen; for example, a 40 μm field stop would permit measurement of a 1 μm diameter area of the specimen. By moving the specimen in the x and y directions (i.e., in the plane of the slide) in the raster pattern now so familiar to us from television and computer displays, and recording and adding the measurements appropriately, it was possible to measure the integrated absorption of a cell, and/or to make an image of the cell with each pixel corresponding in intensity to the transmission or absorption value. This was the first, and, at the time, the only feasible approach to **scanning cytometry**.

The use of stage motion for scanning made operation extremely slow; it could take many minutes to produce a high-resolution scanned image of a single cell, and there were no computers available to capture the data. Somewhat higher speed could be achieved by using moving mirrors, driven by galvanometers, for image scanning, and limiting the tasks of the motorized stage to bringing a new field of the specimen into view and into focus; this required some primitive electronic storage capability, and made measurements susceptible to errors due to uneven illumination across the field, although this could be compensated for.

Since the late 1940's and early 1950's had already given us Howdy Doody, Milton Berle, and the Ricardos, it might be expected that somewhere around that time, someone would have tried to automate the process of looking down the microscope and counting cells using video technology. In fact, image analyzing cytometers were developed; most of them were not based on video cameras, for a number of reasons, not the least of which was the variable light sensitivity of different regions of a camera tube, which would make quantitative measurements difficult. There was also the primitive state of the computers available; multimillion dollar mainframes had a processor speed measured in tens of kilohertz, if that, and memory of only a few thousand kilobytes, and this made it difficult to acquire, store, and process the large amount of data contained even in a digitized image of a single cell.

By the 1960's, a commercial version of Caspersson's microspectrophotometer had been produced by Zeiss, and several groups of investigators were using this instrument and a variety of laboratory-built scanning systems in attempts to automate analysis of the Papanicolaou smear for cervical cancer screening, on the one hand, and the differential white blood cell count, on the other[42,43,52,53,,57-60]. It was felt that both of these tasks would require analysis of cell images with reso-

lution of 1 μm or better, to derive measures of such characteristics as cell and nuclear size and shape, cytoplasmic texture or granularity, etc., which could then be used to develop the cell classification algorithms needed to do the job. Although it was widely recognized that practical instruments for clinical use would have to be substantially faster than what was then available, this was not of immediate concern in the early stages of algorithm development, and few people even bothered to calculate the order of magnitude of improvement that might be necessary.

Flow Cytometry and Sorting: Why and How

Somewhat simpler tasks of cell or particle identification, characterization, and counting than those involved in Papanicolaou smear analysis and differential white cell counting had attracted the attention of other groups of researchers at least since the 1930's. During World War II, the United States Army became interested in developing devices that could rapidly detect bacterial biowarfare agents in aerosols; this would require processing a relatively large volume of sample in substantially less time than would have been possible using even a low-resolution scanning system. The apparatus that was built in support of this project[29-31] achieved the necessary rapid specimen transport by injecting the air stream containing the sample into the center of a larger (**sheath**) stream of flowing air, confining the particles of interest to a small region in the center, or **core**, of the stream, which passed through the focal point of what was essentially a dark-field microscope. Particles passing through the system would scatter light into a collection lens, eventually producing electrical signals at the output of a photodetector. The instrument could detect at least some *Bacillus* spores, objects on the order of 0.5 μm in diameter, in specimens, and is generally recognized as having been the first flow cytometer used for observation of biological cells; similar apparatus had been used previously for studies of dust particles in air and of colloidal solutions.

By the late 1940's and early 1950's, the same principles, including the use of **sheath flow**, as just described, for keeping cells in the center of a larger flowing stream of fluid, were applied to the detection and counting of red blood cells in saline solutions[48]. This paved the way for automation of a diagnostic test notorious for its imprecision when performed by a human observer using a counting chamber, or **hemocytometer**, and a microscope.

Neither the bacterial counter nor the early red cell counters had any significant capacity either for discriminating different types of cells or for making quantitative measurements. Both types of instrument were measuring what we would now recognize as side scatter signals; although larger particles would, in general, produce larger signals than smaller ones composed of the same material, the correlations between sizes and signal amplitudes were not particularly strong. In the case of the bacterial counter, a substantial fraction of the spores of interest would not produce signals detectable above background; the blood cell counters had a similar lack of sensitivity to small signals, which was advantageous in that blood platelets, which are typically much smaller than red cells, would generally not be detected. White cells, which are larger than red cells, would be counted as red cells; however, since blood normally contains only about 1/1000 as many white cells as red cells, inclusion of white cells in the red cell count would not usually introduce any significant error.

An alternative flow-based method for cell counting was developed in the 1950's by Wallace Coulter[49]. Recognizing that cells, which are surrounded by a lipid membrane, are relatively poor conductors of electricity as compared to the saline solutions in which they are suspended, he devised an apparatus in which cells passed one by one through a small (< 100 μm) orifice between two chambers filled with saline. A constant electric current was maintained across the orifice; when a cell passed through, the **electrical impedance** (similar to **resistance**, which is the inverse of **conductance**) increased in proportion to the volume of the cell, causing a proportional increase in the measured voltage across the orifice. The Coulter counter was widely adopted in clinical laboratories for blood cell counting; it was soon established that it could provide more accurate measurements of cell size than had previously been available[50-1].

In the early 1960's, investigators working with Leitz[61] proposed development of a hematology counter in which a fluorescence measurement would be added to the light scattering measurement used in red cell counting. If a fluorescent dye such as acridine orange were added to the blood sample, white cells would be stained much more brightly than red cells; the white cell count could then be derived from the fluorescence signal, and the raw red cell count from the scatter signal, which included white cells, could, in theory, be corrected using the white cell count. It was also noted that acridine orange fluorescence could be used to discriminate mononuclear cells from granulocytes. However, it does not appear that the device, which would have represented a new level of sophistication in flow cytometry, was ever actually built.

A hardwired image analysis system developed in an attempt to automate reading of Papanicolaou smears had been tested in the late 1950's; although it was nowhere near accurate enough, let alone fast enough, for clinical use, it showed enough promise to encourage executives at the International Business Machines Corporation to look into producing an improved instrument.

Assuming this would be some kind of image analyzer, IBM gave technical responsibility for the program to Louis Kamentsky, who had recently developed a successful optical character reader. He did some calculations of what would be required in the way of light sources, scanning rates, and computer storage and processing speeds to solve the problem using image analysis, and concluded it couldn't be done that way.

Having learned from pathologists in New York that cell size and nucleic acid content should provide a good indica-

tor of whether cervical cells were normal or abnormal, Kamentsky traveled to Caspersson's laboratory in Stockholm and learned the principles of microspectrophotometry. He then built a flow cytometer that used a transmission measurement at visible wavelengths to estimate cell size and a 260 nm UV absorption measurement to estimate nucleic acid content[1, 65].

Subsequent versions of this instrument, which incorporated a dedicated computer system, could measure as many as four cellular parameters[78]. A brief trial on cervical cytology specimens indicated the system had some ability to discriminate normal from abnormal cells[77]; it could also produce distinguishable signals from different types of cells in blood samples stained with a combination of acidic and basic dyes, suggesting that flow cytometry might be usable for differential leukocyte counting.

Although impedance (Coulter) counters and optical flow cytometers could analyze hundreds of cells/second, providing a high enough data acquisition rate to be useful for clinical use, scanning cytometers offered a significant advantage. A scanning system with computer-controlled stage motion could be programmed to reposition a cell on a slide within the field of view of the objective, allowing the cell to be identified or otherwise characterized by visual observation; it was, initially, not possible to extract cells with known measured characteristics from a flow cytometer. Until this could be done, it would be difficult to verify any cell classification arrived at using a flow cytometer, especially where the diagnosis of cervical cancer or leukemia might be involved.

This problem was solved in the mid 1960's, when both Mack Fulwyler[67], working at the Los Alamos National Laboratory, and Kamentsky, at IBM[66], demonstrated **cell sorters** built as adjuncts to their flow cytometers. Kamentsky's system used a syringe pump to extract selected cells from its relatively slow-flowing sample stream. Fulwyler's was based on ink jet printer technology then recently developed by Richard Sweet[68] at Stanford; following passage through the cytometer's measurement system (originally a Coulter orifice), the saline sample stream was broken into droplets, and those droplets that contained cells with selected measurement values were electrically charged at the droplet breakoff point. The selected charged droplets were then deflected into a collection vessel by an electric field, while uncharged droplets went, as it were, down the drain.

Fluorescence and Flow: Love at First Light

Fluorescence measurement was introduced to flow cytometry in the late 1960's as a means of improving both quantitative and qualitative analyses. By that time, Van Dilla et al[79] at Los Alamos and Dittrich and Göhde[83] in Germany had built fluorescence flow cytometers to measure cellular DNA content, facilitating analysis of abnormalities in tumor cells and of cell cycle kinetics in both neoplastic and normal cells. Kamentsky had left IBM to found Bio/Physics Systems, which produced a fluorescence flow cytometer that was the first commercial product to incorporate an argon ion laser; Göhde's instrument, built around a fluorescence microscope with arc lamp illumination, was distributed commercially by Phywe.

Leonard Herzenberg and his colleagues[2], at Stanford, realizing that fluorescence flow cytometry and subsequent cell sorting could provide a useful and novel method for purifying living cells for further study, developed a series of instruments. Although their original apparatus[82], with arc lamp illumination, was not sufficiently sensitive to permit them to achieve their objective of sorting cells from the immune system, based on the presence and intensity of staining by fluorescently labeled antibodies, the second version[86], which used a water-cooled argon laser, was more than adequate. This was commercialized as the FACS in 1974 by a group at Becton-Dickinson (B-D), led by Bernard Shoor.

By 1979, B-D, Coulter, and Ortho (a division of Johnson & Johnson that bought Bio/Physics Systems) were producing flow cytometers that could measure small- and large-angle light scattering and fluorescence in at least two wavelength regions, analyzing several thousand cells per second, and with droplet deflection cell sorting capability. DNA content analysis was receiving considerable attention as a means of characterizing the aggressiveness of breast cancer and other malignancies, and monoclonal antibodies had begun to emerge as reagents for dissecting the stages of development of cells of the blood and immune system. Instruments with two lasers were used to detect staining of cells by different monoclonal antibodies conjugated with spectrally distinguishable dyes.

Image cytometers existed; they were much slower and even less user-friendly than the early flow cytometers, weren't easily adapted for immunofluorescence analysis, and couldn't sort. Meanwhile, the early publications and presentations based on flow cytometry and sorting created a large demand for cell sorters among immunologists and tumor biologists. By the early 1980's, when a mysterious new disease appeared, best characterized – using flow cytometry and monoclonal antibodies – by a precipitous drop in the numbers of circulating T-helper lymphocytes, clinicians, as well as researchers, had become anxious to obtain and use fluorescence flow cytometers – and, often, to avoid sorting!

In the decades since, confocal microscopes, scanning laser cytometers, and image analysis systems have come into use. They can do things flow cytometers cannot do; they typically have better spatial resolution and can be used to examine cells repeatedly over time, but they cannot analyze cells as rapidly, and there are many fewer of them than there are flow cytometers. They are also, unlike flow cytometers, not subject to:

Shapiro's First Law of Flow Cytometry:
A 51 μm Particle CLOGS a 51 μm Orifice!

That notwithstanding, in these first years of the 21st century, most cytometry is flow cytometry, and, for almost all

applications except clinical hematology analysis, flow cytometry involves fluorescence measurement.

Fluorescence and flow are made for each other for several reasons, but primarily because fluorescence, at least from organic materials, is a somewhat ephemeral measurement. Recall that fluorescence occurs when a photon is absorbed by an atom or molecule, raising the energy level of an associated electron to an excited state, after which a small amount of the energy is lost as heat, and the remainder is emitted in the form of a longer wavelength photon, as fluorescence. However, there is a substantial chance that a photon at the excitation wavelength will not excite fluorescence but will, instead, **photobleach** a fluorescent molecule, producing a nonfluorescent product by breaking a chemical bond. In general, you can expect to get only a finite number of cycles of excitation and emission out of each fluorescent molecule (**fluorophore**) before photobleaching occurs.

If you look at a slide of cells stained with a fluorescent dye under a fluorescence microscope, you are likely to notice that, each time you move to a new field of view, the fluorescence from the cells in the new field is more intense than the fluorescence from the field that you had been looking at immediately before, which has undergone some photobleaching. This effect makes it difficult to get precise quantitative measurements of fluorescence intensity from cells in a static or scanning cytometer if you have to find the cells by visual observation before making the measurement, because the extent of photobleaching prior to the measurement will differ from cell to cell. In a flow cytometer, each cell is exposed to excitation light only for the brief period during which it passes through the illuminating beam, usually a few microseconds, and the flow velocity is typically nearly constant for all the cells examined. These uniform conditions of measurement make it relatively easy to attain high **precision**, meaning that one can expect nearly equal measurement values for cells containing equal amounts of fluorescent material; this is especially desirable for such applications as DNA content analysis of tumors, in which the abnormal cells' DNA content may differ by only a few percent from that of normal stromal cells.

A basis for the compatibility between fluorescence measurements and cytometry in general is found in the dark field nature of fluorescence measurements. It has already been noted that precise absorption measurements are best made when the concentration of the relevant absorbing material is relatively high. When one is trying to detect a small number of molecules of some substance in or on a cell, this condition is not always easy to satisfy. In the 1930's, unsuccessful attempts were made to detect antibody binding to cellular structures by bright field microscopy of the absorption of various organic dyes bound to antibodies. In 1941, Albert Coons, Hugh Creech, and Norman Jones successfully labeled cells with an antibody containing a fluorescent organic molecule[44], enabling structures binding the antibody to be visualized clearly against a dark background. In general, fluorescence measurements, when compared to absorption measurements, offer higher **sensitivity**, meaning that they can be used to detect smaller amounts or concentrations of a relevant analyte; this is of importance in attempting to detect many cellular antigens, and also in identifying genetic sequences and/or fluorescent protein products of transfected genes present in small copy numbers.

It is also usually easier to make simultaneous measurements of a number of different substances in cells, a process referred to as **multiparameter** cytometry, by fluorescence than by absorption, and the trend in recent years in both flow and static cytometry has been toward measurement of an increasingly large number of characteristics of each cell subjected to analysis, as can be appreciated from Table 1-1, way back on page 3.

Conflict: Resolution

When I first got into cytometry in the late 1960's, and for the next twenty years or so, there was a "farmer vs. rancher" feud going on between the people who did image analysis and the people who did flow, especially in the areas of development of differential white cell counters and Pap smear analyzers.

The first automated differential counters to hit the market were, in fact, image analyzers that scanned blood smears stained with the conventional Giemsa's or Wright's stains. Most of them are gone, now; modern hematology counters, which produce total red cell, reticulocyte (immature red cell), white cell, and platelet counts and red cell and platelet size (and, in at least one case, red cell hemoglobin) distributions, in addition to the differential white cell count, are typically flow based. Various instruments may measure electrical impedance (AC as well as DC), light absorption, scattering (polarized or depolarized), extinction, and/or fluorescence. None of them uses Giemsa's or Wright's stain.

Of course, with hundreds of monoclonal antibodies available that react with cells of the blood and immune system in various stages of development, we can use fluorescence flow cytometry to count and/or classify stem cells and other normal and abnormal cells in bone marrow, peripheral blood, and specimens from patients with leukemias and lymphomas, taking on tasks in hematology that few of the pioneers seriously believed could be approached using instruments. However, while the hematology counters run in a highly automated mode and produce numbers that can go directly into a hospital chart, most of the more sophisticated fluorescence-based analyses require considerable human intervention at stages ranging from the selection of a panel of antibodies to be used to the performance of the flow cytometric analysis and the interpretation of the results. This may facilitate reimbursement for the tests, but it leaves some of us unfulfilled, although perhaps better paid.

Cytometric apparatus that facilitated the performance and interpretation of the Papanicolaou (Pap) smear reached the market much later than did differential white cell counters. The first improvements were limited to automation of sample preparation and staining; there are now several image

analysis based systems approved for clinical use in aiding screening (locating cells and displaying images of them to a human observer), and at least one approved for performing screening itself. All use the traditional Papanicolaou stain, a witches' brew of highly nonspecific acidic and basic dyes known since the 19th century and blended for its present purpose before the middle of the 20th.

Why the difference? What made the Pap smear survive the smear campaign and the Wright's stained blood smear go with the flow? The answer is simple. Both Pap smear analysis and blood smear analysis on slides depend heavily on morphologic information about the internal structure of cells. Criteria for cell identification in these tasks may include cell and nuclear size and shape, cytoplasmic granularity or texture, and, especially in dealing with abnormal blood cells, finer details such as whether nucleoli or intracellular inclusions are present.

Some of these characteristics, e.g., cytoplasmic granularity (which, as has already been noted, is a major contributor to a side scatter signal), can be determined using flow cytometers. While the fluorescent antibodies used for such tasks as leukemia and lymphoma classification using flow are highly specific (although not, in general, specific to a single cell type), most of the instruments that perform the differential leukocyte count do not need to use particularly specific reagents. In fact, it is possible, using only a combination of polarized and depolarized light scattering measurements, to do a differential white cell count with no reagents other than a diluent containing a lysing agent for red cells.

In the case of differential leukocyte counting, we have learned to substitute measurements that can be made of whole cells in flow, requiring only low-resolution optics, for those that would, if we were dealing with a stained smear, require that we make and analyze a somewhat higher-resolution image of each cell. Flow is faster, simpler, and cheaper, and, although morphologic hematologists still look at stained smears of blood and bone marrow from patients in whom abnormal cells have been found, we no longer need to look at a stained smear by eye or by machine to perform a routine white cell differential count. Although there may be combinations of low-resolution flow-based measurements that could provide a cervical cancer screening test comparable in performance to Pap smear analysis, none have yet been clinically validated; we therefore still rely on image analysis in approaches to automation of cervical cancer cytology and on visual observation where automation is not available.

Researchers face problems similar to those faced by clinicians. If you want to select and sort the 2,000 cells out of 10,000,000 cells in a transfected population that express the most green fluorescent protein (GFP), you will probably use a flow cytometer with high-speed sorting capability and settle for a low-resolution optical measurement that detects all of the GFP in or on the cell without regard to its precise location. If you have arranged for the GFP to be coexpressed with a particular structural protein involved, say, in the formation of the septum in dividing bacterial cells, you will very likely want to look at images of those cells at as high a **resolution** as you can achieve in order to get the information you need from the cells. There are, of course, tradeoffs.

It is January 1, 2002, as I write this, and therefore particularly appropriate to continue this New Year's resolution discussion. In the age of the personal computer and the digital camera and camcorder, there is little need to introduce the concepts of **digital images** and their component **pixels** (the term originally came from "picture elements"); most of us are exposed to at least 1024 × 768 almost 24/7. In this instance, the familiar 1024 × 768 figure describes the **pixel resolution** of an image acquisition or display device, with the image made up of 768 rows, each containing 1,024 pixels (or of 1,024 columns, each containing 768 pixels). However, the pixel resolution of the device doesn't, in itself, tell us anything about the **image resolution**, i.e., the area in the specimen represented by each pixel.

This depends to a great extent on what's in the image. In an image from the Hubble Space Telescope, each pixel could be light-years across; in an image from an atomic force microscope, each pixel might only be a few tenths of a nanometer (Ångstroms) across. But the image resolution also depends on the combination of hardware and software used to acquire and process the image. We are free to collect a transmitted light microscope image of a 10.24 μm by 7.68 μm rectangle (close quarters for a single lymphocyte) somewhere on a slide containing a stained smear of peripheral blood, using a digital camera chip with 1024 × 768 resolution, but we are not free to assume that each of the pixels in the image represents an area of approximately 0.01 by 0.01 μm. In this instance, the optics of the light microscope will limit our effective resolution to somewhere between 0.25 and 0.5 μm, and using a camera with a high pixel resolution won't help resolve smaller structures any more than would projecting the microscope image on the wall. Either strategy provides what microscopists have long known as **empty magnification**; the digital implementation, by allowing us to collect many more bits worth of information than we need or can use, slows down the rate at which we can process samples by a factor of at least several hundred, and is best avoided.

So what do we do when we really need high-resolution images? As it turns out, one of the physical factors that limits resolution in a conventional fluorescence microscope, in which the entire thickness of the specimen is illuminated, is fluorescence emission from out-of-focus regions of the specimen above and below the plane of what we are trying to look at. In a **confocal microscope**, the illumination and light collection optics are configured to minimize the contributions from out-of-focus regions; this provides a high-resolution image of a very thin slice of the specimen. Resolution is improved further in **multiphoton confocal microscopy**, in which fluorescence is excited by the nearly simultaneous absorption of two or more photons of lower energy than would normally be needed for excitation. The illumina-

tion in a multiphoton instrument comes from a tightly focused high-energy pulsed laser, and it is only in a very small region near the focal spot of the laser that the density of low energy photons is sufficient for multiphoton excitation to occur. This produces an extremely high-resolution image (pixel dimensions of less than 0.1 μm are fairly readily achieved), and also minimizes bleaching of fluorescent probe molecules and photodamage to cells.

As usual, we pay a price for the higher-resolution images. We are now looking at slices of the specimen so thin that we need to construct a three-dimensional image from serial slices of the specimen to fully visualize many cellular structures. Instead of two-dimensional pixels, we must now think in terms of three-dimensional **voxels**, or volume elements. Let's go back to the single lymphocyte which, on the blood smear discussed on the previous page, was confined in a two-dimensional, 10.24 μm by 7.68 μm, rectangular area. For three-dimensional imaging, we would prefer that the cell not be flattened out, especially if we want to look at it while it is alive, so we will assume it to be roughly spherical, and imprison it in a cube 10 μm on a side. If we used a multiphoton microscope with each voxel representing a cube 0.1 μm on a side, building a 3-D image of that single cell would require us to collect data from 100 × 100 × 100 voxels, or 10^6 voxels, and, even if it only took one microsecond to get data from each voxel, it would take a second just to collect the data.

This is a perfectly acceptable time frame for an investigator who needs information about subcellular structures; even with the computer time required for image processing, one can examine hundreds, if not thousands, of cells in a working day. However, even this is feasible only if the experimenter and/or the hardware and software in the instrument first scan the specimen at low resolution to find the cells of interest.

1.3 PROBLEM NUMBER ONE: FINDING THE CELL(S)

Continuing with the scenario just described, suppose we have cells at a concentration of 10^6/mL, dispersed on a slide in a layer 10 μm thick. A 1 × 1 cm area of the slide will contain 10,000 of the 10 μm cubicles in which we could cache a lymphocyte. Recalling that 10 μm is 1/1,000 cm, and that 1 cm^3 is 1 mL, we can calculate the aggregate volume of these 10,000 little boxes as 1/1,000 mL. If the cell concentration is 10^6/mL, we can only expect to find about 1,000 cells in 1/1,000 mL, and it would take us 16 minutes, 40 seconds to scan all of them at high resolution. However, if we adopted the brute force approach and did 3-D scans over the entire 1 × 1 cm area, instead of finding the locations of the cells and restricting the high-resolution scanning to those regions, we would waste 9,000 seconds, or 2 hours and 15 minutes, scanning unoccupied cubicles.

There's another problem; although we may arbitrarily divide the 1 × 1 cm × 10 μm volume into 10 μm cubicles, we have not created actual physical boundaries on the slide, and we can expect the cells to be randomly distributed over the surface, which means that parts of the same cell could lie in more than one cubicle. If we deal with a specimen thicker than 10 μm or so, the positional uncertainty extends to a third dimension, further compounding the problem of finding the cells, which gets even more difficult if we are trying to get high-resolution images of specific cell types in a tissue section, or in a small living organism such as a *Drosophila* embryo or a *C. elegans* worm.

When I first got into the cytometry game, in the late 1960's, my colleagues and I at the National Bureau of Standards and the National Institutes of Health built a state-of-the art computerized microscope, with stage position and focus, among other things, under computer control[73]. The instrument could be operated in an interactive mode, which allowed an experimenter to move the stage and focus the microscope using a small console that included a keypad and a relatively primitive joystick; the actual motion remained under computer control at all times. This made it possible to scan a slide visually, find cells of interest, store their locations in the computer, and have the instrument come back and do the high-resolution scans (resolution, in this instance, was better than 0.25 μm) needed for an experiment.

We didn't have a computer algorithm for finding cells automatically; since scanning the area immediately surrounding a cell took us not one second, but two minutes, there would have been little point to automating cell finding. The actual scanning time required to collect integrated absorption measurements of the DNA content of 100 cells, stained by the Feulgen method, was 3 hours, 20 minutes. We could find the cells that interested us by eye in a few minutes; scanning the slide looking for them might have taken days.

We were able to make life a little easier for ourselves by developing an algorithm to remove objects from the periphery of an image. A typical microscope field would contain a cell of interest, which we had positioned in the center of the field, surrounded by other cells, parts of cells, or dirt and/or other junk. Since the algorithm was relatively simple-minded, our visual selection process required us to exclude cells that touched or were overlapped by other cells. Figure 1-5, on the next page, shows the results of applying the algorithm.

The figure also shows how difficult it might be to develop algorithms to find cells. Even among the few cells present in the image shown, there are substantial differences in size and shape, and there are marked inhomogeneities in staining intensity within cells. Humans get very good very fast at finding cells and at discriminating cells from junk, even when cell size, shape, and texture vary. If staining (or whatever else produces contrast between the cell and the background) were relatively uniform, recognizing a cell by computer would be fairly easy; one would only have to find an appropriately sized area of the image in which all the pixel values were above a certain threshold level. This simple approach clearly won't work with cells such as those shown in the figure.

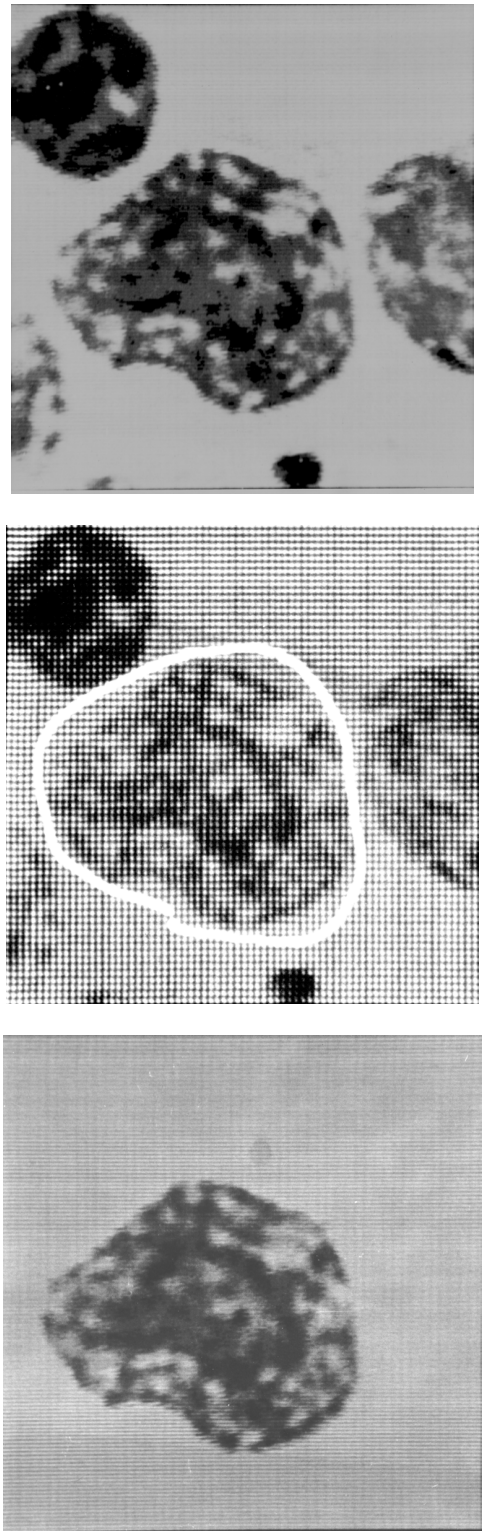

Figure 1-5. Top panel: scanned image of Feulgen-stained lymphoblastoid cells. In the middle panel, a boundary drawn around the cell of interest is shown; the bottom panel shows results of applying an algorithm to remove all objects except the cell of interest.

Before we developed the procedure for removing unwanted material from images, we had the option of picking out the cell of interest by drawing a boundary around it using a light pen, as shown in the middle panel of Figure 1-5. Many researchers working with cell images still find it convenient to locate cells and define boundaries for analysis in this fashion, and essentially the same procedure is used to draw the boundaries of regions of interest in two-parameter data displays from flow cytometers. This sidesteps the issue of automated cell finding (or of automated cluster finding, in the case of data displays). The boundary drawing is now commonly done using a personal computer and a mouse; in 1970, there were no mice, at least not the computer kind, and the interactive display and light pen we used cost tens of thousands of dollars, and had to be attached to the mainframe computer we needed to do the image processing. Very few laboratories could have afforded to duplicate our apparatus; today, you can introduce your children and grandchildren to the wonders of the microscopic world using a digital video microscope that costs less than $100 and attaches to your computer's USB port. But, although your computer is probably hundreds of times faster than the one we used and has thousands of times the storage capacity, which could allow it to be used to implement cell finding algorithms of which we could only dream, it still takes a long time to capture high-resolution cell images, and the detail in those images makes it more difficult for those algorithms to define the boundaries of a cell or a nucleus than it would be if the images used for cell finding were of lower resolution.

A cell 10 μm in diameter occupies thousands of contiguous pixels in a high-resolution image with 0.1 × 0.1 μm pixels, such as might be obtained from a multiphoton confocal microscope, but fewer than 100 contiguous pixels in a lower-resolution image with 1 × 1 μm pixels, such as might be obtained from a scanning laser cytometer. The high-resolution image may contain many pixels with intensity near that of the background (as is the case with the image shown in the top panel of Figure 1-5), making it necessary to do fairly convoluted analyses of each pixel in the context of its neighbor pixels to precisely define the area of a cell or an internal organelle. However, each of the 1 × 1 μm pixels of the lower-resolution image can be thought of as representing contributions from a hundred 0.1 × 0.1 μm areas of the cell, and, since it is unlikely that all of these are at background intensity, it is apt to be easier to define an area as composed of contiguous pixels above a certain intensity level if one uses larger pixels.

When one is working with isolated cells, it becomes attractive to attempt to confine them to defined areas of a slide rather than to have to scan the entire surface to find cells distributed at random. By the 1960's, it had occurred to more than one group of investigators that depositing cells in a thin line on a glass or plastic tape would allow an automated cytology instrument to restrict stage motion to one dimension instead of two, potentially speeding up processing. The concept is illustrated in Figure 1-6 (next page).

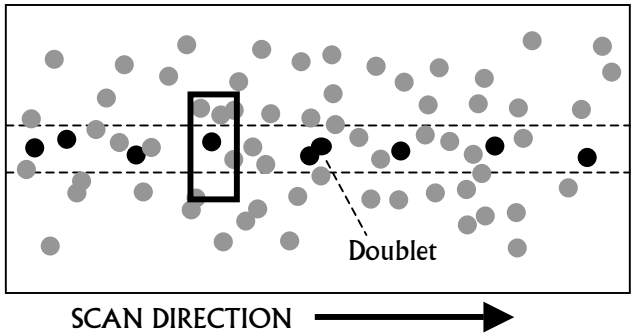

Figure 1-6. One-dimensional scanning of cells deposited in a narrow line (between dotted lines) on a slide or tape simplifies finding cells in a specimen. Black dots represent cells deposited in the line, gray dots represent cells deposited at random, and the small rectangle shows the field of view.

You can actually try this trick at home, if you happen to keep a microscope there, or in the lab, if you don't. Simulate the "cells" with dots in different colors made by a permanent marker with a fine or extra fine point; make dots in one color, corresponding to the black dots in the figure, along a straight edge placed parallel to the long edge of a slide, and make dots of another color (or enlist a [much] younger associate to do so), corresponding to the gray dots in the figure, all over the slide. Put the slide under the microscope, using a low- (10× or lower) power objective; place one of the "black" dots in the center of the field of view. Stop down the substage iris diaphragm until you get a field a few times the diameter of the "cell." Then move only the horizontal stage motion control. You should note that, although the "black" cells you encounter as you scan along the slide in one dimension remain entirely in the field of view (up to a point; if the line along which you scribed wasn't exactly parallel to the edge of the slide, there will be some drift), you will almost certainly find "gray" cells cut off at the edges of the field of view. Now, looking at Figure 1-7, we can consider what a photodetector "looking" at the field of view would "see" if the slide in Figure 1-6 were scanned. We can regard this signal as a series of images, each made up of a single pixel that is considerably larger than the cells of interest.

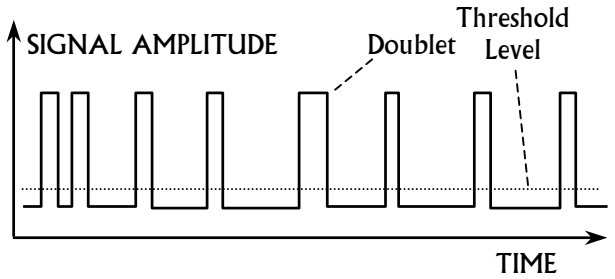

Figure 1-7. Idealized plot of signal amplitude vs. time representing a scan at constant speed along the cell deposition line of Figure 1-6; only signals corresponding to the "black" cells in that figure are shown.

For the moment, we can make believe the "gray" cells in Figure 1-6 aren't there; the simulated detector signal shown in Figure 1-7 only goes above threshold when it scans a "black" cell. This would actually happen if, for example, the slide were illuminated with blue light, and the "gray" and "black" cells were, respectively, unstained and stained with a green fluorescent dye. Figure 1-7 could then represent the electrical signal from a photodetector with a green filter in front of it.

When we look at the slide by eye, we don't scan very rapidly, and we almost never scan at uniform speed, so we don't instinctively relate what we see to the exact time at which we see it. When we scan with a cytometer, it is at least an advantage, and often an imperative, to scan at a constant speed, putting the times at which signals from objects appear at the detector output(s) in a fixed and precise relation to the positions of the objects in space.

In constructing Figure 1-7, the assumption was made that both the illumination intensity in the field of view and the scanning rate remained constant. If we look at the signal amplitude in the figure, it remains at a relatively low **baseline** level most of the time, and there are eight **pulses** during which the amplitude rises to a higher level and returns to the baseline value after a brief interval. If we glance up from Figure 1-7 to Figure 1-6, we notice that the positions of the pulses in time correspond to the positions of the black cells on the slide.

Flow Cytometry: Quick on the Trigger

The signal(s) used to detect cells' presence in the field of view (also called the **measurement point**, **region**, **station**, or **zone**, or the **analysis point**, **interrogation zone or point**, or **observation point**) of a cytometer is (are) called **trigger signal(s)**. The amplitude of a trigger signal must be substantially different in the cases in which a cell is and is not present at the observation point; in other words, it must be possible to define a **threshold level** above which the amplitude will invariably rise when a cell is present. If we pick a threshold level indicated by the dotted horizontal line in Figure 1-7, we see that the signal shown in the figure can serve as a trigger signal; its amplitude is well above the threshold level whenever a cell or cells are present in the field of view, and comfortably below that level when the field of view contains no cells.

Now, suppose that, instead of scanning cells deposited in a line on a slide or tape, we confine cells to the center of a flowing stream, and look at that through a microscope. We'll get rid of the gray cells this time, and only consider the black ones. And, if we want to draw a schematic picture of this, what we get is Figure 1-6, except that the gray cells aren't there, and the arrow indicates "Flow Direction" instead of "Scan Direction." Instead of defining the boundaries of the cell deposition area, the dotted lines define the diameter of the core stream containing the cells. We have sneakily built ourselves a **flow cytometer**.

Of course, if we were actually looking at the stream of cells in a flow cytometer, it would probably be flowing fast enough so that we couldn't distinguish the individual cells as they went by; remember that the visual system makes a "movie" out of images displayed at rates of 25-30/second (/s). Most photodetectors don't have this problem; they can respond to changes in light intensity that occur in nanoseconds (ns). So we could get a signal pretty much like the signal in Figure 1-7 out of a photodetector in a flow cytometer; the major difference would be in the time scale.

When scanning a slide by eye, we are apt to take at least 100 milliseconds (ms) to examine each cell; slide-scanning apparatus is substantially faster, producing pulse durations of hundreds of microseconds (μs) or less. Flow cytometers are faster still; most current commercial instruments produce pulses with durations in the range between 0.5 and 12 μs. Thus, the hardware and software responsible for detecting the presence of a cell need to do their job in a relatively short time, particularly in cell sorters, where the cell must be detected and analyzed, and the decision to sort it or not made and implemented, in the space of a few microseconds. If the signal in Figure 1-7 were coming from a detector in a flow cytometer, we could use it as a trigger signal.

Many of the signals of most interest to users of flow cytometers are of very low amplitude. Routine immunofluorescence measurements often require detection of only a few thousand fluorescently labeled antibody molecules bound to a cell surface. In such cases, the signal from the fluorescence detectors may be only slightly above background or baseline levels, and their use as trigger signals is likely to result in an unacceptably high level of **false triggering**, resulting in accumulation of spurious data values, due to the influences of stray light and electronic noise fluctuations. Even in cases when relatively weak fluorescence signals can be used as trigger signals to indicate the presence of stained cells, they will be of no help in detecting unstained cells. It has thus become customary to use a small-angle (forward) light scattering signal as the trigger signal when measuring immunofluorescence; all cells scatter light.

When none of the pulses from cells of interest are expected to be of high amplitude, requiring that a threshold level be set close to the baseline, discrimination of cells from background noise may be improved by using multiple triggers, requiring that two or more signals go above threshold at the same time to indicate a cell's presence. I almost always use forward light scattering and fluorescence as dual trigger signals when working with bacteria.

The Main Event

Looking back at Figures 1-6 and 1-7, though, we can see that there is another catch to triggering; it is not Catch-22, but Catch-2. Two of the black cells in Figure 1-6 are stuck together, and delineated as a "doublet" in that figure; the corresponding pulse, similarly delineated in Figure 1-7, is, though wider than the other pulses, still only a single pulse. Since cells going through a flow cytometer (or cells deposited on a slide) arrive (or appear) at more or less random intervals, there is always the chance that two or more cells will be close enough in space, and their corresponding output signals close enough in time, so that they produce only a single pulse at the detector output. Note that the cells do not have to be physically stuck together for such **coincidences** to occur, they must simply be close enough so that the detector signal does not fall below the threshold value between the time the first cell enters the measurement region and the time the second (or last, if there are more than two) cell leaves it.

When we get technical about what we are really measuring in a flow cytometer (and now is one of those times), rather than saying that a pulse above threshold level represents a **cell**, we say that it represents an **event**, which might correspond to the passage of one cell, or multiple cells, or one or more pieces of noncellular junk capable of generating an equivalent optical/electronic signal, through the system, or which might result from stray light and/or electronic noise or some other glitch in the apparatus.

The Pulse Quickens; The Plot Thickens

There are ways of identifying pulses that result from coincidences; the height, width, and/or area of such pulses is/are typically different from those resulting from the transit of single cells, and, with the aid of appropriate hardware and/or software, it is possible to identify coincidences and correct counts. And now is probably an opportune time for me to confess that the pulses of Figure 1-7 are highly idealized, in that all of the pulses from single cells look pretty much the same; that definitely isn't the way things really are.

In fact, **all** of the information about a cell that can be gotten from flow cytometers is contained in, and must be extracted from, the **height**, or **amplitude**, the **area**, or **integral**, and the **width** and **shape** of the pulses produced at the detector(s) as the cell passes through the measurement region(s). Generally speaking, there isn't much point to doing flow cytometry if you expect all of the cells you analyze to look alike; the usual purpose of an experiment is the characterization of **heterogeneity** within a cell population, and the rest of this book is intended to help you make sure that the differences in pulses you see from cell to cell represent biological differences you are looking for, rather than reflecting vagaries of apparatus, reagents, and technique.

And now, at last, we have gotten our fingers on the pulse of flow cytometry. For the fact is that, while the information in scanning and imaging cytometers ultimately makes its way into the processing electronics in the form of a series of pulses, often referred to as a **pulse train**, it is only in flow cytometers and in the lowest resolution scanning devices that all of the information a detector gets about a cell (or, more accurately, an event) is contained in a single pulse. This was recognized early on as an important and distinctive characteristic of flow cytometry; before the term "flow cytometry" itself was coined in the 1970's, many workers in the field referred to it as **pulse cytophotometry**.

1.4 FLOW CYTOMETRY: PROBLEMS, PARAMETERS, PROBES, AND PRINCIPLES

Since the 1970's, it has become possible for users blissfully unconcerned with the nuts and bolts (or the atoms and bits) of instrumentation to buy flow cytometers capable of extracting more and more pulses from an increasingly diverse variety of objects, ranging downward from eukaryotic cells and microorganisms to organelles and large molecules, and upward to pancreatic islets, *C. elegans*, *Drosophila* embryos, and multicellular plankton organisms.

From reading the manufacturers' brochures and visiting their Web sites, interested researchers and clinicians can learn that it is possible to analyze and sort over a hundred thousand cells per second, to identify rare cells that represent only one of every ten million cells in mixed populations, to simultaneously measure light scattering at two or three angles and fluorescence in twelve or more spectral regions, to measure fluorescence with a precision better than one percent, and to detect and quantify a few hundred molecules of fluorescent antibody bound to a cell surface. It is somewhat harder to discern that it may be difficult or impossible to accomplish two or three of these amazing feats at once. If you're contemplating pushing the envelope, you definitely need to look at the **problem(s)** you're trying to solve, the measurement **parameters** and **probes** with which you can extract the necessary information from the cells, and the **principles** that may allow you to get your answers – or prevent you from getting them. I will take this approach in considering how the technology has gotten to its present state, starting with relatively simple problems and the relatively simple systems for solving them.

Since flow cytometers are designed to analyze single cells in suspension, it is not surprising that their development and evolution have been directed in large part by workers in the fields of hematology and immunology, who deal primarily with cells that are either in suspension, as is the case in blood samples, or relatively easy to get into suspension, as is the case when it is necessary to examine cells from bone marrow or lymphoid tissues or tumors.

In addition to being conveniently packaged, cells from the blood and immune system provide us with a number of models for fundamental biological processes. With the analysis of the genome behind us, we still need the details of differentiation that allow politically sensitive fertilized ova to develop through the politically sensitive embryonic stem cell stage into multicellular organisms who, after some years, can be dropped from the welfare rolls with the blessings of the same legislators who so staunchly defended them at smaller cell numbers. Cells in the blood and immune system develop from a single class of stem cells, which were hypothesized about and sought for years, and were finally identified with the aid of flow cytometry, and we now traffic in blood stem cells for patients' benefit as well as studying the cells' development in the interest of science. Differentiation gone wrong, with the aid of somatic mutation, produces leukemias and lymphomas, and we use flow cytometry both to clarify the biology of neoplasia and to determine the prognosis and treatment in individual instances. The processes of clonal selection underlying both cellular and humoral immune responses provide a picture of evolution at work, as well as examples of a wide variety of mechanisms of inter- and intracellular signaling.

Counting Cells: Precision I (Mean, S.D., CV)

The simplest flow cytometers, and the first to be widely used, solved the **problem** of providing precise counts of the number of cells per unit volume of a sample, without explicitly characterizing the cells otherwise. Such instruments have only a single detector, and, because they measure an **intrinsic parameter**, typically **light scattering** or **extinction** or **electrical impedance**, do not require that the cells be treated with any reagent, or **probe**.

The **sample** used for cell counting may be taken directly from the **specimen** containing the cells, or may be an aliquot of that specimen **diluted** by a known amount, or **dilution factor**. If, for example, the specimen is diluted 1:20 to produce the sample, the dilution factor is 20.

The **principle** of operation of a cell counter is almost embarrassingly simple. An electronic counter is set to zero at the beginning of each run. Next, sample is passed through the system at a known, constant flow rate. As cells go through the measurement system, they produce pulses at the detector output; the count is increased by one whenever the output from the detector goes above the threshold level. Those cells that produce pulses with amplitudes above threshold are counted; those that do not are not. Any particle other than a cell that produces a signal above threshold is counted as a cell; any transient electrical disturbance or noise that causes the sensor output to go above threshold is also counted as a cell.

Although this sounds like a very simple-minded approach, it usually works, can be implemented using relatively primitive electronics, and can deal with thousands of cells per second. And, as will be amply illustrated later in this section, it is relatively easy to get from this point to a flow cytometer that makes one or several additional measurements of cells. The principal requirement is that, in addition to (or instead of) being used to increase the number in the counter, the trigger signal(s) initiate(s) the capture and recording of information about the height, area, and/or width of pulses from one or more detectors.

In the late 1950's and 1960's, the first optical and electronic (Coulter) cell counters reached the market. They were designed to count blood cells; I have already noted that red cell counts were done by setting a threshold high enough to prevent platelets from triggering, and that white cells were counted with red cells, but did not normally introduce significant inaccuracy into the red cell count because of their relatively low numbers. White cell counts were done on samples in which the red cells had been lysed by addition of a chemical such as saponin or one of a number of detergents to the diluent.

Before counters became available, people did cell counts by examining diluted blood (or another cell sample) in a **hemocytometer** under a microscope. A hemocytometer is a specially designed microscope slide with a ruled grid that defines square or rectangular areas, each fractions of a millimeter on a side, and with ridges on either side of the ruled area that insure that the thickness of the layer of diluted blood under the cover slip will be constant (usually 0.1 mm). For a white cell count, blood is typically diluted 1:20 with a solution that lyses red cells and stains white cells; the number of cells in four 1 × 1 mm squares is counted. The total volume of diluted blood counted is therefore 0.4 mm^3, or 0.4 μL. To obtain the count of white cells/mm^3 (the old-fashioned unit used when I was a medical student), one divides this number by 0.4 (the volume counted) and multiplies the result by 20 (the dilution factor). Because red cells are so much more numerous than white cells, blood is diluted 1:200 for red cell counts (without lysis, obviously), and a smaller area of the slide is used for counting.

Poisson Statistics and Precision in Counting

So what's wrong with hemocytometer counts, apart from the fact that they used to be done by slave labor (for which read medical students, or at least those of my generation)? The problem is with the **precision** of the counts. Precision, as was noted on p. 12, refers to the degree to which replicate measurements agree with one another. The precision of a measurement is often characterized by a statistic called the **coefficient of variation (CV)**, which, expressed as a percentage, is 100 times the **standard deviation (S.D.)** divided by the **mean** (and by **mean** I mean the **arithmetic mean**, or **average**, i.e., the sum of the individual measurements divided by the number of measurements). Well, you might say, "What mean and standard deviation? The count is only done once; how much time do you think those overworked medical students can spare?"

Enter another Student; not a 1960's medical student, this time, but a man of an earlier generation named William Sealy Gossett, who published his basic statistical works as "Student" because his employers at the Guinness Brewery worried that their competitors might improve their positions by using statistics if they discovered his identity. He showed in 1907[2317] that, if one actually counted n cells in a hemocytometer (that's before the division and multiplication steps), one should expect the standard deviation of the measurement to be the square root of n (I will use the notation $n^{1/2}$ rather than $\sqrt{n}$ for this quantity for typographic reasons), meaning that the coefficient of variation, in percent, would be $100/n^{1/2}$. We would now say that the statistics of counts conform to the **Poisson distribution**, which was described by Siméon Poisson in 1837[2318], but "Student" was apparently unaware of Poisson's work, and reached his conclusions independently. In fact, the Poisson distribution was only given that name seven years after Gossett's paper appeared[2319]. We will encounter the Poisson distribution in several other contexts related to cytometry, flow and otherwise.

Now, if we consider looking at a sample with a white cell count of 5,000/mm^3, which is in the normal range, the number of cells you would actually have counted in the hemocytometer to obtain that value would be 5,000 divided by the dilution factor (20) and multiplied by the volume (in mm^3) counted (0.4), which works out to 100 cells. The standard deviation would therefore be the square root of 100, or 10; the CV would be 10 percent. If you were dealing with an abnormally low white cell count, say one that you read as 1,250/mm^3, you would only have counted 25 cells; the standard deviation would be 5, and the CV would be 20 percent. And all of this assumes that the counting process is perfect; we know that it isn't, and we also know that other factors, such as dilution and pipetting errors, will further decrease precision. So the precision of a hemocytometer white cell count in the normal range is barely acceptable. Getting a CV of 1 percent, which is more than respectable, would require that you count 10,000 objects, which would be 100 hemocytometers' worth if you were dealing with our original white cell count of 5,000/mm^3. Nobody is going to sit there and do that by eye, but it's a piece of cake for an electronic or optical counter.

A typical hematology counter uses a **constant volume pump**, such as a syringe pump, to deliver sample at a constant flow rate. The flow rate is the volume of sample analyzed per unit time; dividing the number of cells counted per unit time by the flow rate gives the number of particles per unit volume. Blood specimens are usually diluted before being run in a counter, so the raw value must be multiplied by the dilution factor to get the particle count per unit volume of blood. For example, if the counter's sample flow rate is 1 μL/s, and a blood sample is diluted 1:20 (with a solution that lyses red cells) to count white cells, and running the counter for 40 seconds yields a raw count of 10,000 cells, the white cell count in the blood is:

$$\frac{10{,}000\ (\#\ \text{of cells counted in 40 s}) \times 20\ (\text{dilution})}{40\ (\#\ \text{of}\ \mu\text{L counted in 40 s})}$$

or 5,000/μL. Since the raw count is 10,000, the standard deviation is 100, and the CV is 1 percent.

Rare Event Analysis: The Fundamental Things Apply as Cells Go By

Many of the tasks in modern cytometry are examples of **rare event analysis**. Examples are looking for primitive stem cells, leukemic cells or cancer cells in blood or bone marrow, for fetal cells in maternal blood, or for transfected cells present at low frequency in a culture. In comparing different samples, it is frequently necessary to determine the statistical significance of small differences between large numbers. Some people seem to think that counting hundreds of thousands or millions of cells lets them beat the Poisson statistics; what's important, however, is the number of cells of interest you count, not the total. Suppose, for example, that you find your cells of interest present at a frequency of 0.04% posi-

tives in one sample of 200,000 cells and 0.15% in another. Simple arithmetic tells you that 0.01% of 200,000 is 20 cells, so the first sample has 80 cells of interest and the second has 300. The Poisson standard deviations for the numbers of cells of interest counted would be about 9 for the 80 cells in the first sample and about 18 for the 300 cells in the second. The two values are thus separated by several standard deviations, which is to say that there is a statistically significant difference between them. However, the statistics provide no information as to the source of the difference. If the cells came from the same pot, one would suspect instrumental factors related to data collection and/or analysis, unless there is reason to believe that a process such as differential settling of the rare cell type would change the composition of a sample aliquot with time. A mild degree of paranoia is probably an asset when dealing with rare event analysis.

Poisson statistics apply to counting anything, from cells to photons and photoelectrons, and even to votes. Digressing briefly from rare event analysis to not-so-current event analysis, if 3,000,000 votes are counted, one expects a Poisson standard deviation of 1,732 votes, or roughly 6 parts in 10,000, meaning that if the vote counting process is supposedly even less reliable or accurate than Poisson statistics would predict (Florida's was said to be 99.9% reliable, or accurate to 10 parts in 10,000), neither candidate had a strong claim to having won the state's Presidential vote.

We have a little more control over cell counting than over vote counting. If you count enough cells, you can accurately discriminate between, say, .01% and .02%. If you only count 10,000 cells total, you'd expect to find one cell (and a CV of 100%) in the sample with .01% and 2 cells (CV of 70.7%) in the sample with .02%; so 10,000 cells total is too small a sample to let you discriminate. If you count 1,000,000 cells total, you end up with 100 cells in the .01% sample (10% CV) and 200 cells (7.1% CV) in the .02% sample, and this difference will be statistically significant.

Count Constant Numbers for Constant Precision

The best way to do counts, although almost nobody does them this way, is to always count the same number of cells of interest, which gives you equal precision no matter what the value is. Normally, we do absolute counts by analyzing a fixed volume of blood (or other sample) and percentage counts by analyzing a fixed number of cells. The alternative is to decide on the level of precision you want – suppose it is 5%. Then you have to count 400 cells (the square root of 400 is 20, and 100/20 = 5). What you do is measure the volume of sample (in the case of absolute counts), or the total number of cells (in the case of percentage counts), which has to be analyzed to yield 400 of the cells of interest. If the cells of interest are at .01%, you'll have to count 4,000,000 cells total to find your 400 cells of interest; if they are at 1%, you'll only have to count 40,000 cells, but, instead of the .01% value being much less precise than the 1% value, both will have the same 5% precision. The down side of doing things this way is that it may require some reprogramming of the apparatus, and probably uses more reagent, but, if you want good numbers, there is simply no better way to get them.

Alternative Counting Aids: The Venerable Bead

As it happens, most fluorescence flow cytometers do not use constant volume pumps for sample delivery, nor do they provide an alternative means of measuring the sample volume flow rate with sufficient precision to allow calculation of cell counts per unit volume by the method described above. Carl Stewart, being a leukocyte biologist, must have felt deprived of one of the major tools of his trade when he arrived at Los Alamos National Laboratory many years ago and discovered that the very fancy fluorescence flow cytometers built there did not provide a cell count. He and John Steinkamp solved that problem by adding fluorescent beads at known concentrations to cell samples[1539]. If you have a bottle full of beads that contains a known number of beads per unit volume, adding a known volume of bead suspension (and it had better be well-mixed bead suspension) to a known volume of cell sample allows you to calculate the number of beads per unit volume in the sample. You can then run the sample for an arbitrary length of time, tallying the total numbers of beads and cells counted. The cell count per unit volume is then given by:

$$\frac{\#\text{ of cells counted} \times \#\text{ of beads per unit volume}}{\#\text{ of beads counted}}$$

and the number of cells per unit volume in the original material from which the cells were taken can be obtained by multiplying by the dilution factor, as in previous examples.

There are a few caveats here. If the determination of the concentration of beads per unit volume is done by a relatively imprecise method (Stewart and Steinkamp used a hemocytometer), the precision of the cell count cannot be improved by counting large numbers of cells and beads. One must also take into account the frequency of clumps and coincidences among both cells and beads, which affect the **accuracy** of the count, i.e., the degree to which the measured value agrees with the "true" value. And, of course, the cytometer must be capable of accurately identifying and counting both cells and beads.

Addition of beads to the sample is now widely practiced in the context of counting CD4 antigen-bearing (**CD4-positive**, or **CD4+**) T lymphocytes in HIV-infected individuals. The identification of these cells is most often done by staining with fluorescently labeled monoclonal anti-CD4 antibody (and, usually, at least one other monoclonal antibody labeled with a different fluorescent label). Before counting beads became available, the standard procedure was the so-called "two-platform" method, in which a hematology counter with a constant volume sample feed is used to obtain both the total white cell count per unit volume of

blood, and the differential white cell count, which includes the percentage of lymphocytes among the white cells; the number of lymphocytes per unit volume is then calculated. The fluorescence flow cytometer is used to define the lymphocyte population and the fraction of that population represented by CD4+ T-cells, allowing calculation of the number of these cells per unit volume. Using counting beads, the procedure can be done on a single platform, i.e., the fluorescence flow cytometer, and this appears to improve accuracy.

And Now to See with Eye Serene the Very Pulse of the Machine: Display, Digitization, and Distributions

In general, people who use flow cytometry want to know more about their samples than how many cells are contained in each milliliter, and that translates into getting more information about the signal pulses than whether their amplitudes exceed the threshold level. In a single-parameter electronic (Coulter) counter, the heights of pulses are proportional to the volumes of the cells passing through. However, whereas only relatively simple circuitry, triggered by the rise above threshold in the signal, is required to increment and store the cell count, more complex hardware and software are needed to capture and store measured values of the volumes of cells. Information about the measured particle may be extracted from the peak amplitude (height), the integral (area), the duration (width), and the shape of signal pulses.

The earliest electronic counters did not come equipped with the means to collect and display **distributions**, i.e., **histograms**, or **bar graphs**, of cell volumes; investigators interested in such information acquired it by feeding the pulse train from a counter into a gadget called a **pulse height analyzer**, a hardwired digital computer originally used by nuclear physicists to measure and discriminate among gamma ray energies.

The prerequisite to pulse height analysis, and to just about anything else that one might want to do in the way of data analysis in cytometry, is the conversion of information from an **analog** form, usually a **voltage** representing one of the pulse characteristics mentioned above, to a **digital** form, i.e., a **number**, using a device appropriately named an **analog-to-digital converter (ADC)**. Digital processing in the flow cytometers of the early 1970's was pretty much restricted to the use of pulse height analyzers, which had the disadvantage that their single ADCs (ADCs were expensive in those days) could only provide information on one measured quantity, or parameter, at a time. It was, however, possible to use live display and storage **oscilloscopes**, without benefit of digitization, to provide simple **dot plots** showing the interrelation of two parameters.

The pulses produced during a cell's passage through the measurement system typically last for only a few microseconds at most (making them veritable "phantoms of de light"), and, until recently, the only ADCs that could practically be used in flow cytometers required more time than this to digitize signals. As a result, it was necessary to use hybrid circuits, which combine analog and digital electronics, to store the appropriate analog values for long enough to permit analog-to-digital conversion. These **peak detector**, **integrator**, and **pulse width measurement** circuits must be **reset** as each particle passes through the illuminating beam, allowing new analog signal levels to be acquired; it is then necessary to **hold** their outputs at a constant level until digitization is complete. The "reset" and "hold" signals must be delivered to the analog storage circuits at the proper times by additional hybrid "**front end**" electronics, which compare one or more trigger signal levels with preset threshold values to determine when a cell is present.

Luckily, a flow cytometer is an example of what is known as a **low duty cycle** device. Even when a sample is being run, cells pass the sensors rather infrequently; what goes by the sensors, most of the time, is the water or saline suspending medium, meaning that a certain amount of **dead time**, during which the pulse measurement circuits are occupied with data from one cell and cannot respond to signals from a second, is tolerable. Because cells arrive at random times, rather than at fixed intervals, **coincidences**, when a second cell arrives before processing of signals from the first is complete, are inevitable. The probability of coincidences can be calculated from – guess what – the **Poisson distribution**, and, while they cannot be eliminated entirely, it is possible to reduce them to acceptable levels by limiting the number of cells analyzed per unit time in accordance with the instrument's dead time.

Once held signals have been digitized, further analysis is accomplished with a digital computer, which, in modern instruments, is typically either an Intel/Microsoft-based or Apple Macintosh personal computer. The necessary software is now available from both flow cytometer manufacturers and third parties, in some cases at no cost. In recent years, as inexpensive, fast, high-resolution ADCs have become available (due largely to the needs of the consumer electronics and telecommunications markets), **digital signal processing (DSP)** hardware and software have replaced analog and hybrid circuits for peak detection, integration, and pulse width measurement, and for some other common tasks in flow cytometry, such as **fluorescence compensation** and **logarithmic conversion** of data. There will be a great deal more said about this further on in the book; for now, however, we will go back to another old problem, its old and newer solutions, and their implications for science, medicine, and society.

DNA Content Analysis: Precision II (Variance)

Most users of flow cytometers and sorters have at least a passing acquaintance with measurements of the **DNA content** of cells and chromosomes, which can be done rapidly and precisely by flow cytometry using a variety of fluorescent stains.

As a rule, all normal diploid cells (nonreplicating or G_0 cells and those in the G_1 phase of the cell cycle) in the same eukaryotic organism should have the same DNA content; this quantity is usually expressed as 2C. DNA syn-

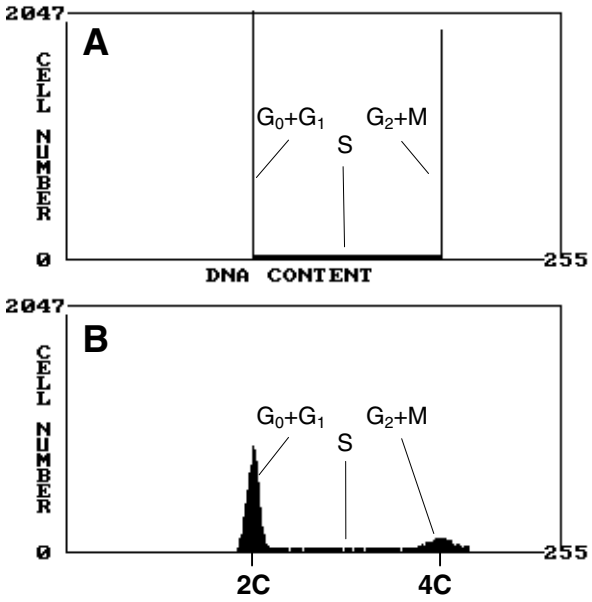

Figure 1-8. Ideal (A) and "real" (B) DNA content distributions, with the same ratios of $(G_0+G_1)/S/(G_2+M)$ cells represented in both.

thesis during the S phase of the cell cycle results in an increase in cellular DNA content, which reaches 4C at the end of S phase and remains at this value during the G_2 phase and during mitosis (M phase), at the completion of which the original cell has been replaced by two daughter cells, each of which has a DNA content of 2C. The haploid germ cells have a DNA content of approximately C; there are approximately equal populations of sperm with DNA content slightly greater than and slightly less than C due to differences in the DNA content of male and female sex chromosomes.

An idealized DNA content **distribution**, that is, a **bar graph** or **histogram** of values of DNA content that would be expected to be observed in a population of cells, some or all of which were progressing through the cell cycle, is shown in panel A of Figure 1-8, above. A "real" distribution, actually synthesized by a mathematical model, but more like those actually obtained from flow cytometry, appears in Panel B. Real (really real) DNA content distributions always exhibit some variance in the G_0/G_1 peak, which may be due to staining procedures, to instrumental errors, and/or to cell-to-cell differences in DNA content. The belief in the constancy of DNA content in diploid cells has been strengthened by the observation that the variances have diminished in magnitude with improvements in preparative and staining techniques and in instrumentation since the first DNA content distributions were published in the 1960's.

When I used the word "variance" in the paragraph above, I meant it, and you probably took it, to denote variability from measurement to measurement. However, the term also has a defined (and related) meaning in statistics; the **variance** of a set of measurements is the sum of squares of the differences between the individual measurements and the **arithmetic mean**, or **average**, divided by one less than the number of measurements. In fact, the statistical **variance** is the square of the **standard deviation**, or, to put it more accurately, the standard deviation is the square root of the variance, and is calculated from it instead of the other way round. For purposes of this discussion, and in most of the rest of the book, I will try to use "variance" to mean the statistical entity unless I tell you otherwise. I may slip; word processors have spelling checkers and grammar checkers, but not intention checkers.

The Normal Distribution: Does the Word "Gaussian" Ring a Bell?

Although the number of cells counted does have some effect on the observed variance of a set of measurements, we are not dealing with Poisson statistics here; the variance of a Poisson distribution is not independent of the mean, but is always equal to it. The peak representing the G_0 and G_1 phase cells of a real DNA content distribution is generally considered to be best approximated by what statisticians define as a **normal** or **Gaussian** distribution, sometimes popularly known as a **bell curve**. The normal distribution is **symmetric**; the arithmetic **mean**, the **median** (the value separating the upper and lower halves of the distribution), and the **mode** (the highest point, or most common value) coincide. The coefficient of variation (CV) (which, you may recall, is expressed in percentage terms as 100 times the S.D. divided by the mean) remains a valid measure of precision, but there is an obvious problem in calculating the CV for a G_0/G_1 peak in a DNA content distribution. The peak falls off as one would expect on the left (low) side, but, on the right (high) side, it merges into the part of the distribution made up of S phase cells, and there isn't a convenient way to decide where the G_0/G_1 cells leave off and the S cells begin.

Because the anatomy of the normal distribution is well known and predictable, we have a statistical trick available to us. The width of the distribution between the two points on the curve at half the maximum (modal, mean, median) value, often referred to as the **full width at half maximum (FWHM)**, is 2.36 standard deviations, and the width between the two points at 0.6 times the maximum value is very nearly two standard deviations.

Binned Data: Navigating the Channels

The process of analog-to-digital conversion that occurs in a pulse height analyzer or in a modern flow cytometer's computer-based data acquisition and analysis system puts data into **bins**, to which we frequently refer as **channels**. These **binned data** are used to compile distributions of measured values of cellular parameters. The distributions in Figure 1-8 are broken into 256 channels, which, by convention, are numbered from 0 to 255. That is the number of bins, or channels, into which an 8-bit ADC distributes its output; an ADC with m bits resolution will have 2^m possible outputs, which, by convention, would be described as channels 0 to 2^{m-1}. Although the outputs of ADCs are often the

same unsigned binary numbers between 0 and 2^{m-1} that denote the channel numbers, ADCs with outputs in different binary formats are not uncommon. For our purposes, it is safest to think in terms of channel numbers, and leave the raw binary formats to the engineers and computer people.

Suppose that the maximum value, i.e., the largest number of cells, in the G_0/G_1 peak of such a distribution is 500 cells, occurring at channel 100, and that channels 97 and 103 each contain 300 cells (that is, 0.6 times the maximum number, 500). It is assumed here that each of the channels between 98 and102 contains 300 or more cells, and that each of the channels below 97 and above 103 contains fewer than 300 cells. The width of the distribution at 0.6 times the maximum value, representing two standard deviations, is then 7 channels, one standard deviation is 3.5 channels, and the CV, expressed as a percentage, is 100 × (3.5/100), or 3.5 percent. It is obviously easier for most people to calculate a CV in their heads using the width at 0.6 times maximum than it is using the width at half maximum, and a real piece of cake if you set the gain so that the maximum value ends up at channel 100, but we've all got calculators, anyway.

So what's the big deal about precision in DNA content measurement? To appreciate this, we go back to the 1960's again. The first cell counters had become available, and they were being used for counting and sizing blood cells. The 1960's also saw a great deal of progress in the field of tissue culture, resulting in substantial numbers of investigators having ready supplies of cells other than blood cells that were either in suspension or could conveniently be put into suspension. People became interested in the details of the cell cycle in cells derived from healthy tissues and from tumors, and in the effect of drugs on the cell cycle.

Once it became convenient to culture cells, it was possible to observe enough mitotic figures to establish that humans had 46 chromosomes, and not 48, as had once been believed, and to establish that cells from many tumors had more or fewer chromosomes, whereas cells from others had what appeared to be chromosomal deletions and translocations. This would mean that the amount of DNA in G_0/G_1 cells from a tumor could be different from the amount in G_0/G_1 cells from the normal stromal elements found in the tumor, potentially providing at least a means of identifying the tumor cells, and, possibly, an objective measurement with prognostic implications.

The catch here is that, as the difference you are trying to detect between two populations becomes smaller, you need better and better precision (lower CVs) in the measurement process. Generally speaking, two populations are resolvable if their means are a few standard deviations apart. If a tumor cell has one or two small chromosomes duplicated, adding, say, two percent to its G_0/G_1 DNA content, you would need a measurement process with a CV well under one percent to resolve separate G_0/G_1 peaks, although you might get a hint of the existence of two populations in a tumor specimen by observing broadening and/or **skewness (asymmetry)** in the peak of a distribution measured with a less precise process. A triploid tumor population, with 50% more DNA than was found in stromal cells, could, of course, be resolved using very imprecise measurements.

DNA Content: Problem, Parameter, Probes

So, the **problem** became one of measuring DNA content with reasonably high precision. It was then necessary to find a suitable measurement **parameter** to solve it. Although Caspersson, in his microspectrophotometers, and Kamentsky, in his early flow cytometers, had used absorption at 260 nm for nucleic acid content measurement, the absorption measurements were difficult to make (among other things, they required special, very expensive quartz optics, because the UV wavelength used is strongly absorbed by glass), and not precise enough to detect small differences.

In the 1920's, Feulgen[35] developed a staining method that coupled a dye to the backbone of the DNA molecule, allowing DNA content in cells to be quantified by measuring absorption of visible light, but some fundamental problems with absorption measurements still limited the precision of DNA analysis. However, in one of the first publications describing fluorescence flow cytometry, in 1969, Van Dilla et al, at Los Alamos National Laboratory, reported the use of a modified **Feulgen procedure**, with fluorescent stains (acriflavine and auramine O) and an argon laser source flow cytometer, to produce DNA content distributions with a coefficient of variation of 6% for the G_0/G_1 cell peak[79]. The Feulgen staining procedure was relatively technically intensive, due to its requirement for fairly elaborate chemical treatment of the cells, and the search for dyes that were easier to use began almost immediately. The first step in this direction was taken in 1969, when Dittrich and Göhde published a relatively sharp DNA content distribution obtained using their arc source flow cytometer to measure the fluorescence of fixed cells stained with **ethidium bromide**[83]. Thus, **fluorescence** became the **parameter** of choice for DNA content measurement.

Ethidium, which increases its fluorescence about thirty-fold when intercalated into double-stranded DNA or RNA, quickly replaced the fluorescent Feulgen stains as the **probe** of choice, and was then largely supplanted by a close chemical relative, **propidium**[217], which remains widely used as a DNA stain. Both dyes require that the cell be fixed, or that its membrane be permeabilized, in order to achieve good stoichiometric staining; they are frequently used to stain nuclei released from cells by treatment with one of a variety of **nonionic detergents**, such as Nonidet P-40 or Triton X-100[223]. Precise measurement of DNA in whole cells, and the best precision measurements in nuclei, require treatment of the sample with RNAse to remove any residual double-stranded RNA.

Once cell sorters became available, in the 1970's, it was realized that a dye that could enter living cells and stain DNA stoichiometrically would make it possible to sort cells in different phases of the cell cycle and analyze their subsequent biological behavior and/or their chemical composi-

tion. Several *bis*-benzimidazole compounds originally synthesized as antiparasitic drugs by Hoechst AG turned out to meet these requirements[238]; the one that has been most widely used, by far, is **Hoechst 33342**[239]. This dye, like the other Hoechst dyes, has two characteristics that limit its use in some situations. Ultraviolet light is required to excite its blue fluorescence, preventing its use in the majority of fluorescence flow cytometers, which are equipped only with a 488 nm (blue-green) argon ion laser as a light source. And, although Hoechst dye staining is highly specific for DNA, the dyes, which do not intercalate but instead bind to the minor groove of the macromolecule, are selective for sequences of three **adenine-thymine (A-T)** base pairs[268]. The latter characteristic is disadvantageous for such applications as DNA content determination in plants, which is widely used as an aid in classification of species, because the Hoechst dyes would yield different results for two species having the same amount of DNA but different **base compositions**, i.e., different ratios of A-T and **G-C (guanine-cytosine)** base pairs. However, the base specificity of the Hoechst dye is an advantage in other circumstances; the combination of the A-T-selective Hoechst 33258 and G-C selective, DNA-specific dyes such as **chromomycin A_3** and **mithramycin**[230], has been used to stain chromosomes from humans and other species, enabling chromosomes with similar total DNA content but different base composition to be distinguished and sorted separately[278]. **High-speed sorting** of dual-stained human chromosomes[904] provided a valuable set of DNA libraries in the early phases of the Human Genome Project, but I'm getting ahead of myself. We can't get into that until we take at least a first look at one- and two-parameter **data displays**.

One-Parameter Displays: Pulse Height Distributions

The cells represented in Figures 1-9 are from the CCRF-CEM T-lymphoblastoid line. They were incubated with **Hoechst 33342**, which, as has already been mentioned, stains DNA stoichiometrically (neglecting, for the moment, differences in base composition). The cells were also exposed to **fluorescein diacetate (FDA)**, a nonfluorescent ester of fluorescein, which should more properly be called diacetylfluorescein but which almost never is. Both compounds are taken up by living cells; once inside cells, FDA is hydrolyzed by nonspecific esterases to **fluorescein**, which exhibits intense green fluorescence when excited with blue or blue-green light, and which, because of its anionic character, is retained in cells for minutes to hours. The cells were measured in a flow cytometer with two separated laser beams; they were first illuminated by a UV laser beam, and the blue fluorescence of Hoechst 33342 (panel A of Figure 1-9) excited by this beam was used as the trigger signal. The cells then passed through the 488 nm beam of a second laser, which provided excitation for the fluorescein fluorescence signal (panel B of Figure 1-9). The histograms of the distributions were collected at different times during a single sample run, using a multichannel pulse height analyzer. The horizontal axis of each histogram indicates fluorescence intensity, on a 512-channel scale; the vertical axis of each histogram represents the number of cells with the corresponding fluorescence intensity. This, by the way, is not a historically informed modern performance on period instruments; the histograms are from around 1980, when one pulse height analyzer and a storage oscilloscope (see Figure 1-11) were all I had to work with for data capture and analysis.

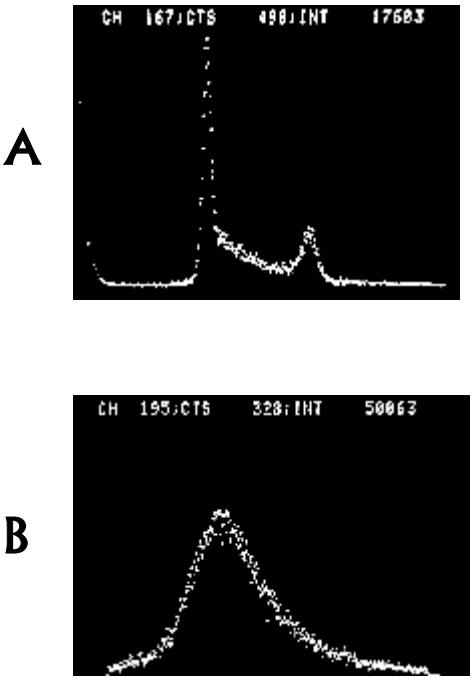

Figure 1-9. Two single parameter histogram displays from the oscilloscope screen of a multichannel pulse height analyzer. A: Fluorescence of the stoichiometric DNA stain Hoechst 33342. B: Fluorescence of intracellular fluorescein. Cells from the same sample are represented in the two histograms.

That said, the data are pretty respectable; their quality is determined primarily by the design and the state of alignment of the flow cytometer optics and fluidics. The CV of the G_0/G_1 peak of the histogram is about 3%, which is excellent for live cells stained with Hoechst 33342.

I have often described sharp peaks, such as the G_0/G_1 peak of a DNA content distribution, as being shaped like a **needle**. Such distributions are not common in flow cytometry data, unless one happens to be analyzing populations of objects that have been intentionally designed to be highly homogeneous, such as the fluorescent plastic microspheres used for instrument alignment and calibration. Although nuclei stained for DNA content, which exemplify one of the best of nature's own quality control processes, yield needles, the shapes of the distributions of most cellular parameters are closer to that of the fluorescein fluorescence distribution

in panel B of Figure 1-9, which resembles a **haystack**, in which it will be unlikely to find a needle.

The pulse height analyzer used to accumulate and display the histograms shown in Figure 1-9 is a specialized computer system that also incorporates some of the features of a flow cytometer's front end electronics and a peak detector. It can accept as input signal a train of pulses ranging in height from 0 to 10 V, using the input signal or another pulse train as a trigger signal, with a threshold set by the operator. Once a signal above threshold is encountered, the peak height is captured by the peak detector, and the signal is digitized by an ADC that, in this instance, produces a **9-bit** output, i.e., a number between 0 and 511.

The pulse height analyzer stores its histograms in 512 memory locations. The **program**, or, more accurately, the procedure, or **algorithm**, for calculating a histogram is fairly simple. First, set the contents of all memory locations to zero. Then, every time a new numerical value emerges from the A-D converter, add one to the contents of the memory location corresponding to that numerical value. Stop when the total number of cells reaches a preset value.

This particular analyzer actually had several options on when to stop: at a preset value for the total number of cells, or for the number of cells in a single channel or memory location, or for the number of cells in a **region of interest**, a range of contiguous channels settable by the operator. It also had some refinements in its display; it would show the channel location of a cursor (CH) and the number of counts in that channel (CTS), as well as the total number of counts in the region of interest (INT). The histogram, sans numbers, could also be drawn on an X-Y plotter; several could be compared by eye in overlays using different color pens.

Pretty much the same algorithm is used for histogram computation today as was used in the analyzer. The difference is that in 1973, when the pulse height analyzer was built, a small startup company called Intel had just begun to ship samples of the first 4-bit microprocessor, and computer memory costs were on the order of 10 cents a byte. The smallest minicomputers available cost around $10,000. The pulse height analyzer didn't have a central processing unit, couldn't process alphanumeric data, couldn't calculate a sine or a logarithm; it used special-purpose hardware to implement the algorithm, and, even at that and even then, it sold for about $5,000. I'm not sure you can even buy a standalone pulse height analyzer today; instead, there are boards containing the necessary front end electronics that plug into standard personal computers. But, even if I could have afforded a second pulse height analyzer in 1980, it wouldn't have helped me do correlated analyses of two parameters.

Mathematical Analysis of DNA Histograms: If It's Worth Doing, It's Worth Doing Well

It was noted on p. 22 that, when one looks at a DNA content histogram, there isn't a convenient way to decide where the G_0/G_1 cells leave off and the S cells begin; there also isn't a convenient way to decide where the S cells leave off and the G_2/M cells begin, or identify debris and cell aggregates in a sample, and things get worse in tumor samples.

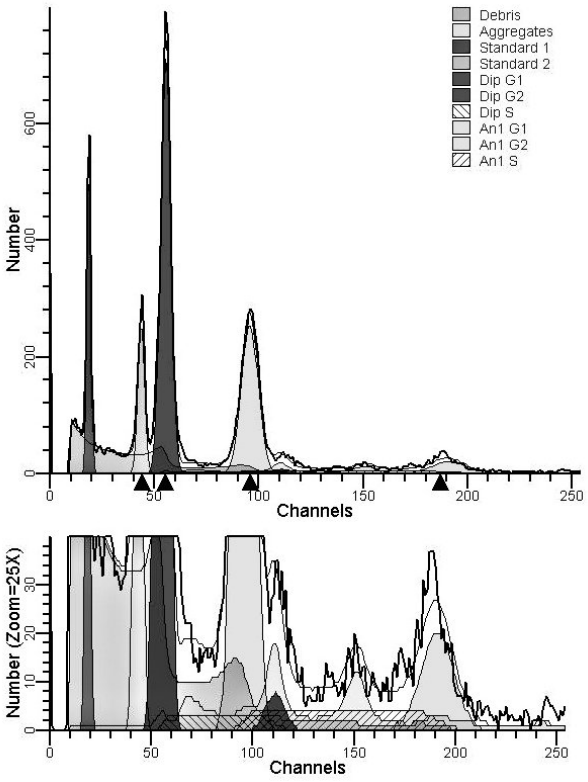

Figure 1-10. Use of a mathematical model to determine fractions of DNA-aneuploid breast cancer cells and normal stromal cells in different cell cycle phases in a sample from a tumor. Chicken and trout erythrocytes are added to the sample to provide standards with known DNA content. Contributed by Verity Software House.

Although tumor cells with abnormal numbers of chromosomes are correctly described as **aneuploid**, a tumor in which the neoplastic and stromal G_0/G_1 cells have different DNA contents is, by convention[741], referred to as **DNA aneuploid**. Mathematical models for DNA histogram analysis have been developed over the years, first, to estimate the fractions of cells in different cell cycle phases in an otherwise homogeneous population, and, later, to determine cell cycle distributions of both stromal cells and DNA aneuploid tumor cells. Further refinements allow for modeling of cellular debris and cell aggregates, enabling them to be largely excluded from analysis. An example of the application of one of the more sophisticated such models (ModFit LT™, from Verity Software House) appears in Figure 1-10.

The earliest publications on fluorescence flow cytometry[79,83] dealt with DNA analysis, and cancer researchers and clinicians began to use the technique almost immediately to attempt to establish the prognostic significance of both DNA aneuploidy and the fraction of cells in S phase in tu-

mors. The development of a method for extracting nuclei from paraffin-embedded tissue for flow cytometric analysis of DNA content[610,891] allowed these issues to be approached by retrospective as well as by prospective studies. By the early 1990's, DNA analysis of breast cancer had come into reasonably widespread clinical use as a prognostic tool. However, in 1996, the American Society of Clinical Oncology recommended against the routine use of flow cytometry in breast cancer[2320], and the volume of specimens analyzed has declined substantially since then. Bagwell et al[2321] have recently demonstrated, based on reanalysis of data from several large studies of node-negative breast cancer, that, after application of a consistent method of analysis and adjustment of some previously used criteria, DNA ploidy and S phase fraction again become strong prognostic indicators. This is not the only publication that shows that how and how well a laboratory test is done can profoundly affect its clinical significance, and that message is important whether or not flow cytometric DNA analysis comes back into vogue.

Linear Thinking

Noncycling cells with known DNA content, such as chicken and trout erythrocytes, can be added to a sample to serve as standards, as was done in the sample shown in Figure 1-10. Such standards are useful in establishing the **linearity** of the instrument and data acquisition system. A system is said to be linear when a proportional change in its input changes its output by the same proportion. In a simple DNA histogram, if the system is linear, and the mean or mode of the G_0/G_1 peak, representing cells with 2C DNA content, is at channel n, the mean or mode of the G_2/M peak, representing cells with 4C DNA content, will be at channel $2n$, or, because of the inherent error of ADCs, within one channel of channel $2n$. In practice, somewhat larger degrees of nonlinearity can be tolerated and corrected for, provided the nonlinearities are stable over time.

Lineage Thinking: Sperm Sorting

Since X- and Y-chromosomes in most species do not contain the same amount of DNA, one would expect a highly precise fluorescence flow cytometer to be able to distinguish them. The necessary precision has been achieved in high-speed sorters by modifications to flow chamber geometry and light collection optics, and sperm vitally stained with Hoechst 33342 have been successfully sorted by sex chromosome type and used for artificial insemination and/or *in vitro* fertilization in animals and, more recently and with a great deal more attention from the media, in humans[2322-3]. Gender selection in humans using sorted sperm, while still under attack from some quarters, is now deemed preferable to other methods that involve determination of the sex of preimplantation embryos. Gender selection in animals using the same methodology appears not to have generated as much controversy as has introducing a foreign gene or two into tomatoes, and may yet become big business[2324].

Two-Parameter Displays: Dot Plots and Histograms

Histograms of the individual parameters do not provide any indication of correlations between Hoechst 33342 and fluorescein fluorescence values on a cell-by-cell basis. In modern flow cytometers, computer-based data acquisition and analysis systems make it trivial to capture, display, and analyze correlated multiparameter data from cells, but, until the 1980's, many instruments could only obtain correlated data on two parameters in the form of a display on an oscilloscope. Such a display was called a **cytogram** by Kamentsky and is now more commonly known as a **dot plot**. One showing both Hoechst 33342 and fluorescein fluorescence values for the cells from the same sample analyzed to produce Figure 1-9, appears in Figure 1-11, below.

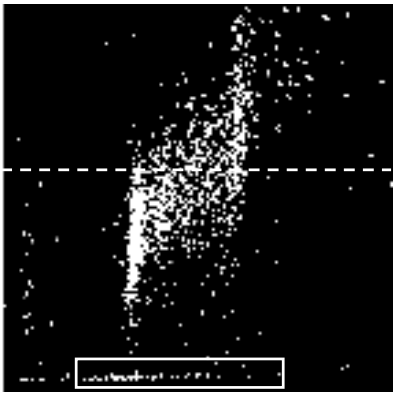

Figure 1-11. Dot plot (cytogram) of Hoechst 33342 fluorescence (x-axis) vs. fluorescein fluorescence (y-axis) for CCRF-CEM cells from the same sample shown in Figure 1-9. Cells in the box are dead; the dotted line is explained below.

Dot plots were the first, and remain the simplest, multiparameter displays in cytometry, and, as we shall presently see, tell us more than we could find out simply by looking at single-parameter histograms. In order to demonstrate this point, we should keep the histograms of Figure 1-9 in mind as we proceed.

In order to appreciate why two parameters are better than one, we need only look at the dot plot in Figure 1-11. One of the first things we notice is that cells with higher Hoechst dye fluorescence intensities, i.e., cells containing more DNA, show higher fluorescein fluorescence intensities. This shouldn't be surprising; if cells didn't get bigger during the process of reproduction, they'd eventually vanish, and it would seem logical that the amounts of FDA cells would take up, and the amounts of fluorescein they would produce and retain, would be at least roughly proportional to cell size. The horizontal dotted line across the dot plot defines two ranges of fluorescein fluorescence values that almost completely separate the diploid and tetraploid populations.

Even more significant, but less obvious to the untrained eye, are the cells represented in the box near the bottom of

the cytogram. These exhibit Hoechst 33342 fluorescence, but not fluorescein fluorescence; they are **dead cells**, or would be so defined by the criteria of a **dye exclusion test**. Such tests actually detect a breach in the cell membrane, which allows dyes such as propidium iodide and Trypan blue, which normally do not enter intact cells, to get in. In this case, the hole in the membrane allows the fluorescein produced intracellularly to leak out very rapidly. As a result, the dead cells exhibit little or no fluorescein fluorescence; their Hoechst dye fluorescence intensities remain indicative of their DNA content.

Dot plots, then, could readily generate an appetite for multiparameter data analysis capability which, given the state of instrumentation and computers in the early days of flow cytometry, was not readily satisfied. A few people could afford what were called **two-parameter pulse height analyzers**. These devices could produce distributions tabulating the number of events (cells, in this case) corresponding to each possible pair of values for two variables. They were about ten times the price of single-parameter pulse height analyzers; they also didn't have great resolution, due to the high cost of memory. Even if the two variables were digitized to only 6 bits' precision, with each yielding a number between 0 and 63, storage of the two-parameter, or **bivariate**, distribution would require 64 × 64, or 4,096, memory locations. However, much of the information contained in bivariate distributions could be obtained, at much lower cost, by adding relatively simple **gating electronics** to the circuitry used to generate dot plots.

Multiparameter Analysis Without Computers: Gates Before Gates

Multiparameter analysis and **gating** may be the most important concepts in flow cytometry. Overall progress in the field was undoubtedly slowed during the 1970's and early 1980's because many of the people studying the really interesting biological problems didn't have either information about or access to the tools needed to implement even relatively simple multiparameter analysis and gating, let alone the sophisticated schemes that are now commonplace.

A dot plot, made using an oscilloscope, and demonstrating simple electronic gating, is shown in Figure 1-12. In order to understand how the gating works, we need first to consider how the dot plot is generated. An oscilloscope, like a television set, is built around a **cathode ray tube**. Electrons are accelerated toward a screen coated with a **phosphor** by the electric field generated by a high voltage applied between the cathode and the screen. The electrons are focused into a beam by a magnetic field. The trajectory of the beam, i.e., the horizontal and vertical locations at which it will hit the screen, is determined by voltages applied to pairs of **deflection plates** inside the tube. A **modulation voltage** may be applied to control how much of the beam reaches the screen. Electrons that do reach the screen are absorbed by the phosphor, which subsequently emits some of the absorbed energy as light, by the process of **phosphorescence**.

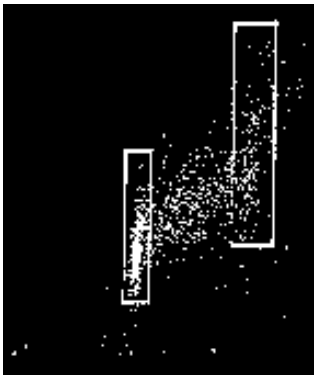

Figure 1-12. Gating regions for counting or sorting set electronically and drawn electronically on an oscilloscope display of a dot plot of DNA content (Hoechst 33342 fluorescence, shown on the x-axis) vs. RNA content (pyronin Y fluorescence, shown on the y-axis) in CCRF-CEM cells.

The dot plot above displays Hoechst 33342 fluorescence on the horizontal or x-axis, and the fluorescence of **pyronin Y**, which stains RNA, on the vertical or y-axis. To generate it, the output from the Hoechst dye fluorescence peak detector was connected to the horizontal deflection plate drive electronics, and the output from the pyronin fluorescence peak detector was connected to the vertical deflection plate drive electronics. The peak detector outputs are both **analog** signals; when applied to the deflection plates, they determine the x- and y- coordinates of the point at which the electron beam will hit the oscilloscope screen. Whether or not an intensified spot, representing the Hoechst 33342 and pyronin fluorescence values associated by the cell, is produced on the screen is determined by the oscilloscope's modulation voltage, which, in this instance, is controlled by what is called a **strobe** signal, generated by the same front end electronics that send the reset and hold signals to the peak detectors.

The strobe signal is a **digital** signal, or **logic pulse**, meaning that its output voltage values are in one of two narrow ranges, or **states**. In this case, voltages at or near about 5 volts (V) represent a "**(logical) 1**," or "**on**," or "**true**" output state, and voltages at or near 0 V, or **ground**, represent a "**(logical) 0**," or "**off**," or "**false**" output condition. The transitions between those two voltage ranges are made rapidly, which, in this instance, means within a small fraction of a microsecond; the interval required is known as the **rise time**.

Some systems use a **positive going** or **positive true** strobe signal, i.e., the strobe output is at ground when the strobe is "off" or "false" and at 5 V when the strobe is "on" or "true"; others use a **negative going** or **negative true** strobe signal, with the output at 5 V when the strobe is off and at ground when the strobe is on. The strobe signal described above is positive true.

The front end electronics are designed so that the strobe signal does not start until the analog signal value in the peak detectors, which can vary continuously between ground and 10 V, has stabilized, and the hold signal is applied to the peak detectors to keep this value from changing during the time the strobe is "on."

When a computer is used for data acquisition and analysis, the beginning, or **leading edge**, of the strobe signal is used to start analog-to-digital (A-D) conversions of the data in the peak detectors; when a dot plot is generated on an oscilloscope, no computer is used, and no digitization is done. The modulation electronics are set so that the beam will reach the screen when the strobe is on and not reach the screen when the strobe is off. Thus, every strobe signal received by the modulation electronics causes a dot to appear on the screen in a position corresponding to the values of the parameters represented on the x- and y-axes.

Early flow cytometers often used **analog storage oscilloscopes**, which incorporated special tubes, with long-persistence phosphors, and associated circuitry that could keep any region of the screen already intensified by the beam "on" until the screen was cleared, or erased, by the user. When an oscilloscope without such storage capacity was used, the dot plot could be recorded by taking a time exposure photograph of the screen.

A dot plot, whether it is recorded on an oscilloscope or using a digital computer (and today's oscilloscopes are increasingly likely to be special-purpose digital computers), does not contain as much information as a **bivariate distribution**. When you see a dot at a given position on the display, you know only that at least one cell in the sample had values of the two measured parameters corresponding to the position of the dot; and you can't get a better estimate of the actual number of cells that shared those values. That's where **gating** comes in. The strobe signal itself can be connected to a digital electronic **counter**, which will store a count and increase the count by one each time a strobe pulse is received. If the value in the counter is set to zero before analysis of a sample begins, the counter will maintain a tally of the total number of cells counted during the analysis.

Now, suppose we were interested in finding out how many of the cells in our dot plot had Hoechst dye fluorescence signals in the range between 3.5 and 4.75 V and pyronin fluorescence signals between 2.5 and 7 V. We could do this if we connected the relevant peak detector signals to an electronic circuit called a **window comparator**.

A **comparator** is a circuit element with two analog inputs, termed **positive** and **negative**, and a **digital**, or **logic level** (e.g., ground for "0" or "off"; 5V for "1" or "on") output. The digital output is on when the voltage at the positive input is higher than the voltage at the negative input, and off otherwise. Comparators are used in the analog front end circuitry of a flow cytometer to determine when the trigger signal (positive input) rises above the threshold level (negative input); one comparator is required for each trigger signal used.

A window comparator is made by connecting the logical outputs of four comparators together in a logical "AND" configuration. The inputs to the the individual comparators are appropriate combinations of the two input signals and two sets of upper and lower limits such that the combined output is "on" only when both signals fall within the limits. The limits would typically be set by turning the knobs of **variable resistors**, or **potentiometers**, which are best known in their roles as volume controls in relatively unsophisticated and older radios and television sets.

Gating is accomplished by connecting the digital output of the window comparator to one input, and the digital strobe signal to the other input, of a purely digital circuit called an **AND gate**. The output of an AND gate is on only when both inputs are on; in this case, the output of the AND gate will be a pulse train containing only the strobe pulses from those cells with parameter values falling between the set limits.

While one counter, working off the strobe signal, is counting all the cells in the sample, another counter, connected to the output of the AND gate, accumulates a count of the cells falling within the gating region. The output of the AND gate can also be used as an input to the electronics that control cell sorting, allowing the cells with values within the set limits to be physically separated from the rest of the sample.

By incorporating a few other bits of analog and digital circuitry into the window comparator modules of my earliest "Cytomutt" flow cytometers, I could, at the press of a button, draw the boundaries of rectangular gating regions on the oscilloscope, as is shown in Figure 1-12; this greatly facilitated setting the upper and lower boundaries of the gating regions. Early commercial instruments had similar features. Of course, they were still limited to rectangular gating regions, and there were clearly situations when one could not separate the cells one wanted to count or sort from the unwanted cells using rectangular gates.

It was possible, by adding still more analog electronics to generate sums and differences of signals from two parameters, and feeding the sums and differences, rather than the original signals, into a window comparator, to define a gating region that corresponded to a parallelogram or other quadrilateral, rather than a rectangle, in the two-dimensional measurement space. This feature was incorporated in the instruments Kamentsky built at Bio/Physics Systems in the early 1970's.

Kamentsky also described, but did not put into production, a clever alternative counting/sort control circuit made by placing opaque black tape over all of an oscilloscope screen except the area corresponding to the gating region, and mounting a photodetector in front of the screen. The gating region defined in this manner could be any arbitrary shape, or even a set of disconnected arbitrary shapes, limited in size and scope only by the user's dexterity with scissors or a knife blade and black tape. Every time a cell lying within the region was encountered, the uncovered portion of the

screen was intensified, generating an output pulse at the photodetector that could be sent to a counter and/or used to initiate sorting. In the era of Bill Gates, we describe freeform gating regions of this type, implemented with mice and computers rather than blades and tape, as one type of **bit-mapped (or bitmap) gates**.

Well, most of the above is all ancient history, right? You must be wondering why I've devoted so much time to searching the souls of old machines when we do everything with computers now.

There are two reasons. The first is that the computers, in most cases, are doing the same things we did with hardwired electronics years ago, and if you understand how things worked then, you'll understand how they work now. The second is that there were, and still are, a few advantages to the old-fashioned electronics, especially for time-critical tasks.

I should mention that, then and now, nothing precluded or precludes us from defining a one-dimensional gating region, using either a simpler window comparator or a computer, and I did note that one-dimensional gating capability, allowing definition of a "region of interest," was typically built into pulse height analyzers. One-dimensional gating was widely used to control cell sorting in the earliest cell sorters, a logical choice when one considers that they typically measured only one relevant parameter.

Two-Parameter Histograms: Enter the Computer

As I wrote in 1994 for the 3rd Edition of this book, "Digital computers are extremely versatile. The same notebook computer on which I am writing this book with the aid of word processing software can be, and has been, used to acquire and analyze data from my flow cytometer. All I have to do is load and start a different program; I can even continue writing while I wait for the cytometer to get data from a new sample. Using additional telecommunications hardware and software, I can, and have, set gates on the cytometer, which is in Massachusetts, from a conference room in Maryland. However, while the computer's overall speed and its ability to switch rapidly between tasks make it appear as if it's doing many things at once, this is an illusion. About the only thing a computer can really do while it is running whatever program is occupying its attention is read or write data from or to a single source. Otherwise, digital computers do one thing at a time, even if they do that one thing really fast."

It's all still pretty much true. Of course, the notebook computer on which the 3rd Edition was written cost nearly $5,000, weighed about seven pounds, had a 50 MHz 80486 processor, at most a couple of MB of RAM, a 500 MB disk drive, and a 640 by 480 pixel screen, and the one I use now cost about $2,000, weighs three pounds, has a 750 MHz Pentium III processor, 256 MB of RAM, a 30 GB disk drive, and a 1024 by 768 pixel screen. The operating system and word processing software have also supposedly been improved. Last time around, my telecommunications were limited to what I could do over standard telephone lines using a 9,600 baud modem; now, I gripe when my cable modem or DSL connections slow to even fifty times that speed. So, what I or you can do with a single computer can be done faster than what could be done eight years ago. But there's more; cytometry today can take advantage of both digital signal processing and multiprocessor systems in ways that, while obvious, were simply infeasible then.

A window comparator implemented in electronics is really making four comparisons at the same time, and they are accomplished in well under a microsecond. If you build a sorter using two window comparators to control deflection into left and right droplet streams, the two comparators work simultaneously. If you want to use a digital computer for sort control in a brute force kind of way, the computer has to fetch the value of the x-axis parameter for the left gating region, check it against the lower and upper bounds for that region, fetch the value of the y-axis parameter, check it against both bounds, and repeat the same steps for the right gating region. Obviously, the computation for a particular gating region can be stopped as soon as a parameter value is found to be out of bounds, but, if you think about it, the full four comparisons for one gating region or another have to be done for any cell that falls in either region, and, until they get done, no signal can be sent to initiate droplet deflection.

In droplet cell sorting, a sort decision has to be made within the few dozen microseconds it takes at most for a cell to get from the observation point to the point at which droplets break off from the cell stream. Up to 10 μs may be required for the signals from the peak detectors (or integrators) to become stable. When hardwired electronics, e.g., window comparators, are used to control sorting, the sort decision signal is sent within a microsecond or so after this time. When a digital computer is used to control sorting, another time interval of at least a few μs is required for A-D conversion before the computer can process the data. And, although the computers have gotten faster, the emergence of high-speed sorting has made it necessary for them to respond within even shorter time intervals.

Until the late 1970's, even minicomputers weren't really fast enough to be competitive with hardwired electronics for sort control. Today's much faster personal computers can easily accomplish the computations required for the window comparison described above well within the time period in which a sort decision must be made. The same computers, however, might not be able to get through a more complex computation, which, say, involved calculating four logarithms and solving quadratic equations to determine whether a cell falls in an elliptical gating region, in time to issue a sort signal, largely because while modern computer hardware is extremely fast, often requiring less than 1 ns to execute a machine instruction, the real time response of the hardware is literally slowed to a crawl by the design of the graphical user interface (GUI) based operating systems (various versions of Microsoft's Windows™, Linux with GUI

extensions, and Apple's Macintosh™ OS) now in most widespread use. The sorting problems are now solved by using some combination of external analog and digital electronics, frequently including one or more digital signal processors, or DSP chips, to implement time-critical processes such as sorting decisions, taking the load off the personal computer's central processing unit (CPU), leaving it free to do what it does best, namely, display the data informatively and attractively.

For plain old flow cytometric data analysis, in which there is no need to initiate action within a few microseconds after a cell actually goes through the beam, computers have always been better than hardwired electronics. That's why Kamentsky used one in his original instrument at IBM. Computers for the rest of us only came along as we could afford them. A few lucky souls, myself included, had computers on their cytometers in the mid-1970's; they were minicomputers, and they were expensive. Now, it's virtually impossible to buy a flow cytometer that doesn't have at least one computer external to the box; most have one or more inside, as well.

Figure 1-13 is a histogram, collected, displayed, and annotated using my own competent, if ancient, MS-DOS-based **4Cyte™** data acquisition software, showing 90° (side) scatter values from a human leukocyte population. The data are plotted on a linear scale. The sample was prepared by incubating whole blood with fluorescently labeled antibodies to the CD3, CD4, and CD8 antigens, and lysing the erythrocytes by addition of an ammonium chloride solution. The "Cytomutt" cytometer used 488 nm excitation from an air-cooled argon ion laser, and measured forward and side scatter and fluorescence in 30 nm bands centered at 520 nm (green; principally fluorescence from **anti-CD4** antibody labeled with **fluorescein**), 580 nm (yellow, principally fluorescence from **anti-CD8** antibody labeled with the **phycobiliprotein**, **phycoerythrin**), and 670 nm (red, principally fluorescence from **anti-CD3** antibody labeled with a **tandem conjugate** of **phycoerythrin** and the cyanine dye **Cy5**). The forward scatter signal was used as the trigger signal.

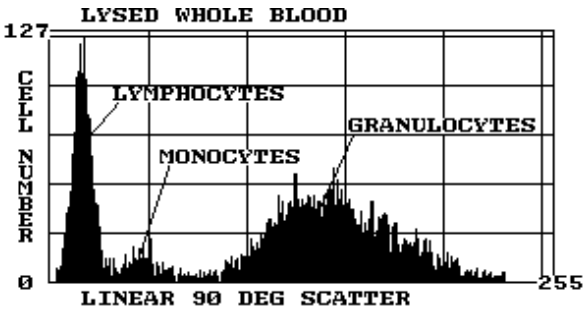

Figure 1-13. Histogram of 90° (side) scatter from leukocytes in lysed whole blood stained with fluorescent antibodies to lymphocyte antigens.

Modern Multiparameter Analysis: List Mode

The histograms and dot plots appearing in Figures 1-9, 1-11, and 1-12 are preserved for posterity only in the form of photographs. Figure 1-9A was photographed after the Hoechst dye fluorescence signal was connected to the pulse height analyzer and some 17,000 cells were run through the cytometer. The analyzer's memory was then cleared, the input was connected to the fluorescein fluorescence signal, and another 50,000 cells from the same sample were run through the instrument to generate the histogram of Figure 1-9B. The dot plots are taken from photos of the screen of an analog storage oscilloscope. I don't suppose the fact that we and a lot of other people stopped buying all of that Polaroid black-and-white film for our oscilloscope cameras loomed large in the company's eventually going bankrupt, but you never know. In the context in which we were using it, the film was a highly unsatisfactory archival medium.

The data represented in the histogram of Figure 1-13 were acquired in **list mode**, meaning that values of all parameters from all cells were stored in the computer's memory and, subsequently, on disk. List mode data acquisition doesn't preclude generating histograms, dot plots, or multivariate distributions while a sample is being run, and it does offer the user the considerable advantage of being able to reanalyze data well after they were acquired. The histogram in Figure 1-13 was generated months after the data were taken. Years ago, even after people had gotten used to having computers attached to their flow cytometers, they used to make a big fuss about acquiring data in list mode. There may have been some flimsy excuse for that attitude before mass storage media such as recordable CDs became available; today, there is simply no reason not to collect data from every run in list mode. Period. All currently available instruments have the necessary software for list mode data storage, and can write files compliant with one or another revision of the **Flow Cytometry Standard (FCS)** established by the Data File Standards Committee of the **International Society for Analytical Cytology (ISAC)**, an organization to which most serious flow cytometer users either belong or should. The standard makes it possible for analysis software from both cytometer manufacturers and third parties to read data from any conforming instrument.

As to the actual data in Figure 1-13, we notice that the histogram, like the DNA histograms in Figures 1-8B and 1-9A, is **multimodal**, meaning not that it has multiple identical maxima, but that it contains multiple peaks. Only one of these, that to the far left, would even be suspected of being a needle rather than a haystack. From the labels in Figure 1-13, it can be surmised that there is good reason to suspect that the peaks at increasingly higher values of 90° scatter represent lymphocytes, monocytes, and granulocytes; we can even go back to page 7 and look at Figure 1-2 to convince ourselves that this is the case. However, just as we can't readily separate the G_0/G_1 cells from the S cells, or the S cells from the G_2/M cells, by looking at a DNA histogram alone,

we can't readily separate the lymphocytes from the monocytes and the monocytes from the granulocytes by looking only at the histogram of 90° scatter.

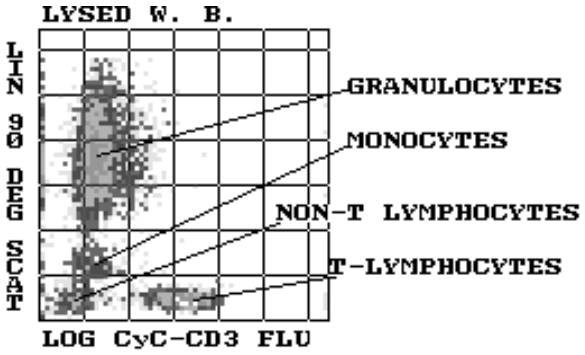

Figure 1-14. Bivariate distribution of anti-CD3 antibody fluorescence intensity vs. 90° (side) scatter for the same leukocyte population shown in Figure 1-13.

The picture gets a lot clearer when we look at the **bivariate distribution**, or **two-parameter histogram**, shown in Figure 1-14. The raw data in this distribution came from the same **list mode** file used to compute the histogram of 90° scatter shown in Figure 1-13; meaning that, thanks to the ready availability of computers and data storage media, we are able to look at the same cells from many different points of view. Figure 1-14 shows clearly identifiable **clusters** of cells; it provides a much clearer separation of lymphocytes, monocytes, and granulocytes than one could obtain using 90° scatter alone, and it also clearly separates the lymphocytes into those that bind the anti-CD3 antibody, i.e., the T cells, and those that do not, most accurately identified as "non-T" lymphocytes.

While a similar separation of cell clusters would be discernible on a dot plot, the bivariate distribution provides a more detailed picture of the relationship between two measured parameters, because the distribution provides an indication of the number of cells and/or the fraction of the cell population sharing the data values corresponding to each point in the two-dimensional measurement space, whereas the dot plot only indicates that one or more cells share the data values corresponding to a point in that space.

A bivariate distribution is computed by setting aside n^2 storage locations, where n is the number of bits of resolution desired for the data. Obviously, n cannot be greater than the number of bits of resolution available from the ADC; in practice, a lower value is typically used, for two reasons. First, the memory requirements are substantial. If each parameter has values ranging from 0 to 1,023, it is necessary to use 1,048,576 storage locations for a single distribution; this requires 2 megabytes if each location uses two bytes, or 16 bits, which would allow up to 65,535 cells or events to be tallied in any given location. If each location uses four bytes, or 32 bits, 4 megabytes of storage are required, but the maximum number of cells that can be tallied in a location is increased to over 4 billion.

While the issue of memory requirements for distribution storage would seem moot at a time when a computer can be equipped with a gigabyte of RAM for a couple of hundred dollars, a second consideration remains. When a two-parameter histogram is computed at high resolution, it is usually necessary to include a very large number of events in order to have more than a few events in each storage location; computing at a lower resolution may actually make it easier to appreciate the structure of the data from smaller cell samples.

For a relatively long time, it was common to compute two-parameter histograms with a resolution of 64 × 64; these require 4,096 storage locations per histogram, which was a manageable amount of memory even in the early days of personal computers. Now, resolutions of 128 × 128 (16,384 storage locations) and 256 × 256 (65,536 storage locations) are widely available. The distribution displayed in Figure 1-13 has 64 × 64 resolution; values on a 1,024-channel scale, such as would be produced by a 10-bit ADC, would be divided by 16 to produce the appropriate value on a 64-channel scale, while the 8-bit (256 channels) values yielded by the lower-resolution converters found in older instruments would be divided by 4.

The data presentation format used in the display of Figure 1-14 is that of a **gray scale density plot**; the different shades of gray in which different points are displayed denote different numbers of cells sharing the corresponding data values. There is an alternative display format for density plots in which different frequencies of occurrence are represented by different colors instead of different shades of gray; this type of plot is described as a **chromatic** or **color** density plot. One can think of the gray levels or different colors in density plots as analogous to the scales that indicate different altitudes on topographical maps. Unfortunately, although the altitude scale is displayed on almost every published topographical map, the analogous scale of cell numbers or frequencies rarely finds its way into print alongside cytometric density plots.

Since computers now used for flow cytometric data analysis have color displays and color printers, chromatic plots are more common than gray scale plots. However, although color pictures are eye-catching and useful for presentations and posters, they can cost you money when included in publications. Those of us who run on lower budgets can almost always use a well-chosen gray scale for published displays without losing information; those lucky enough to not have to think about the cost of color plates might want to choose color scales that will not become uninformative when viewed by readers with defects in color vision.

Figure 1-15 (next page) displays the two-parameter histogram data of Figure 1-14 in an **isometric plot**, or **three-dimensional projection**, also commonly called a **peak-and-valley plot**; Figure 1-16 (next page) shows the same histogram as a **contour plot**.

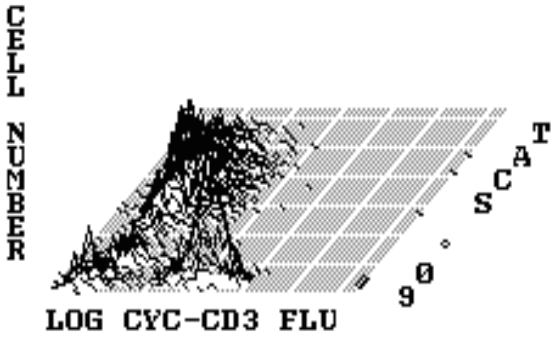

Figure 1-15. The two-parameter histogram of Figure 1-14 shown as an isometric or "peak-and-valley" plot.

In a peak-and-valley plot, a simulated "surface" is created; the apparent "height," or z-value, corresponding to any pair of x- and y-coordinates is made proportional to the frequency of occurrence of the corresponding paired data values in the sample. In a contour plot, a direct indication of frequency of occurrence is not given for each point in the x-y plane. Instead, a series of **contour lines**, or **isopleths**, are drawn, each of which connects points for which data values occur with equal frequency. A contour plot, like a density plot, resembles a topographic map; a peak-and valley plot is more like a relief map. The fact is, though, that there is really no more information in one type of bivariate histogram display than in another. Take this as a mantra.

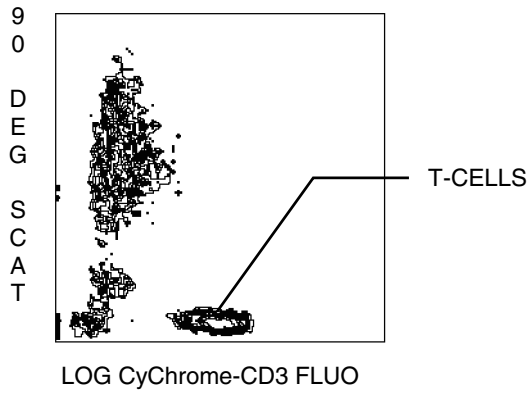

Figure 1-16. The two-parameter histogram of Figures 1-14 and 1-15 displayed as a contour plot.

Peak-and-valley plots seem to have largely fallen out of favor; none too soon, say I. One of their major disadvantages lies in the *a priori* unpredictability of where peaks and valleys will turn up. Big peaks in the simulated "front" block the view of smaller peaks in the simulated "back," unless you use the 3-D graphics capabilities of your computer display to tilt and rotate the display. I suspect it's the dedicated computer gamers who are saving peak-and-valley plots from extinction.

Contour plots require more computation than either chromatic or peak-and-valley plots. Although contour plots may appear to have higher resolution, this appearance is deceiving, resulting as it does from the necessity to **smooth** the data, i.e., average over neighboring points, in order to get the plot to look respectable. Contour plots also do not normally show single occurrences, although these can be superimposed as dots on a contour plot; some such adaptation is essential when dealing with rare events. I think people tend to use contour plots in publications because they look more detailed than dot plots and often reproduce better than gray scale density plots.

I have always favored density plots, using chromatic plots for primary computer output and presentations, and, as a rule, gray scale plots for publication. I routinely use a **binary logarithmic intensity scale**, with one color or gray level indicating single occurrences, the next 2-3 occurrences, and subsequent colors indicating 4-7, 8-15, 16-31, 32-63, 63-127, and more than 127 occurrences. This makes it very easy to spot cells that occur with frequencies of less than one in 10,000.

Although commercially available flow cytometers are now equipped to display sixteen or more parameters (which would typically include light scattering at two angles and fluorescence in twelve spectral regions, with the balance possibly made up of different characteristics of the same pulses, such as width or height and area, and/or of ratios of the heights or areas of two signals from the same cell), almost all analysis is done using two-dimensional histograms or dot plots of two parameters in various combinations.

Three-Dimensional Displays: Can We Look at Clouds from Both Sides? No.

Humans aren't very good at visualizing spaces of more than three dimensions, but you'd certainly expect that, with everybody doing five- and six-parameter measurements in flow cytometry, three-dimensional displays would be commonplace. Some software packages produce a "three-dimensional dot plot," which I have called a **cloud plot** (see Figure 5-11, p. 241). Cloud plots have the same disadvantage as peak-and-valley plots; when one cloud gets in front of another, you have to recompute and change the viewing angle to see what's where. A few people have gone so far as to generate stereo pair images to improve the three-dimensional quality of the displays; they may be the same folks who have kept peak-and valley plots alive.

Three-parameter histograms are problematic for several reasons. First, even a 64 × 64 × 64 3-parameter histogram requires 262,144 storage locations, although there are some tricks that can reduce the storage requirements. Once you do, though, there's still a problem with how to display the data. Isometric plots would require four dimensions, which is out, and contour and chromatic plots demand x-ray vision on the part of the observer. As a result, what people have generally done when they need to represent something analogous to a 3-parameter histogram is show 2-parameter

histograms in "slices," with the resolution along the "sliced" axis often lower than that of the 2-parameter histograms, so that there might be four to sixteen 64 × 64 histograms rather than sixty-four. "Slicing" a histogram, if you stop to think about it, is exactly equivalent to defining a series of rectangular gating regions along the z-axis. And, speaking of gating regions, the slicing technique just mentioned is about the only practical way of setting gates in a three-dimensional space.

There's a lot of information to deal with when you're just looking at two-parameter displays; three-parameter displays could quickly get you to the point of information overload. If we are dealing with n parameters, the number of possible two-parameter displays, counting those showing the same two parameters with the x- and y-axis switched, and omitting those in which the same parameters are on both axes, is $n \times (n\text{-}1)$, while the number of possible three-parameter displays, counting those showing the same three parameters with axes switched and omitting those with the same parameter on two or three axes, is $n \times (n\text{-}1) \times (n\text{-}2)$. For the five-parameter data we have been looking at, we have 20 possible two-parameter displays and 60 possible three-parameter displays; for 16-parameter data, we would have 240 possible two-parameter displays, which is frightening enough, and a mind-boggling 3,360 three-parameter displays. It could take months to run an analysis on a single tube if we had to look at all of them. So, as usual, it is best to get our heads out of the clouds.

Identifying Cells in Heterogeneous Populations: Lift Up Your Heads, Oh Ye Gates!

Most of the interesting applications of flow cytometry involve identifying cells in heterogeneous populations; what varies from case to case is the basis of the heterogeneity. We have already noted several varieties of heterogeneity in our brief examination of DNA content analysis. Cells in a presumably pure, clonally derived, unsynchronized culture will contain different amounts of DNA because they are in different stages of the cell cycle. A DNA aneuploid tumor contains stromal and tumor cells with different G_0/G_1 DNA contents, and both stromal and tumor cells may be in different cell cycle phases. Sperm differ in DNA content depending on which sex chromosome is present. Heterogeneous populations of microorganisms such as are encountered in seawater contain many different genera and species, each with its characteristic genome size. In all of the above cases, it is possible to identify cell subpopulations based on differences in DNA content.

In the widely studied heterogeneous cell populations that comprise blood, the majority of cells are neither DNA aneuploid nor progressing through the cell cycle. Thus, when the **problem** is the identification of different cell types in blood, DNA content is generally not a **parameter** of choice. Figure 1-17, on the next page, illustrates the use of several better suited parameters and of multiple gating methods in one of the most common clinically relevant applications of flow cytometry, the identification of T lymphocytes bearing CD4 and CD8 antigens in human peripheral blood.

If we simply stained cells with a combination of differently colored acidic and basic dyes, as Paul Ehrlich, who developed the basic technique, did in the late 1800's, we would be able to use transmitted light microscopy (with relatively strongly absorbing dyes at high concentrations) or fluorescence microscopy (with fluorescent dyes, probably at lower concentrations) to do a classical differential white blood cell count. The presence or absence of cytoplasmic granules would let us distinguish the granulocytes from the mononuclear cells (monocytes and lymphocytes). The relative amount of staining of those granules by the acidic and the basic dye would allow us to identify eosinophilic (acidophilic to Ehrlich), basophilic, and neutrophilic granulocytes. The size of the cells, amount of cytoplasm, and nuclear shape would allow us to distinguish most of the monocytes from most of the lymphocytes. But that's about as far as we would get. A typical peripheral blood lymphocyte is a small, round cell with a relatively thin rim of cytoplasm surrounding a compact, round nucleus. The nucleus, like the nuclei of all cells, stains predominantly with the basic dye (one of the methylene azure dyes in a typical Giemsa or Wright's stain), which is attracted to the acidic phosphate groups of the nuclear DNA. The basic cytoplasmic proteins attract some of the acidic dye (eosin in the mixtures commonly used for staining blood), but RNA in the cytoplasm also attracts the basic dye. And the staining pattern of most peripheral blood lymphocytes is pretty much the same, whether they are B lymphocytes or T lymphocytes, and, if T lymphocytes, whether they bear the CD4 or the CD8 antigen (although both antigens are present on developing T lymphocytes in the thymus, almost all of the T lymphocytes present in the peripheral blood have lost one or the other).

The optical flow cytometers used for differential white cell counting in hematology laboratories, which typically measure forward and side scatter, can distinguish lymphocytes from monocytes and granulocytes using these measurements alone, but cannot thereby distinguish different types of lymphocytes. However we have already seen from Figures 1-13 through 1-16 that the combination of side scatter measurements and measurements of fluorescence of cell-bound antibodies allows us to distinguish T lymphocytes from other lymphocytes. It should therefore come as no surprise that the **probes**, or reagents, that allow us to define lymphocyte subpopulations, and most other subclasses of cells in the blood, bone marrow, and organs of the immune system, are **antibodies**, and that we detect antibodies bound to cells by the **fluorescence** of labels attached, usually covalently, to the antibody molecules. Flow cytometry greatly facilitated the development of monoclonal antibody reagents, and flow cytometry has since been indispensable for defining the specificities of these reagents and, thereby, allowing their routine use for cell classification in clinical and research laboratories.

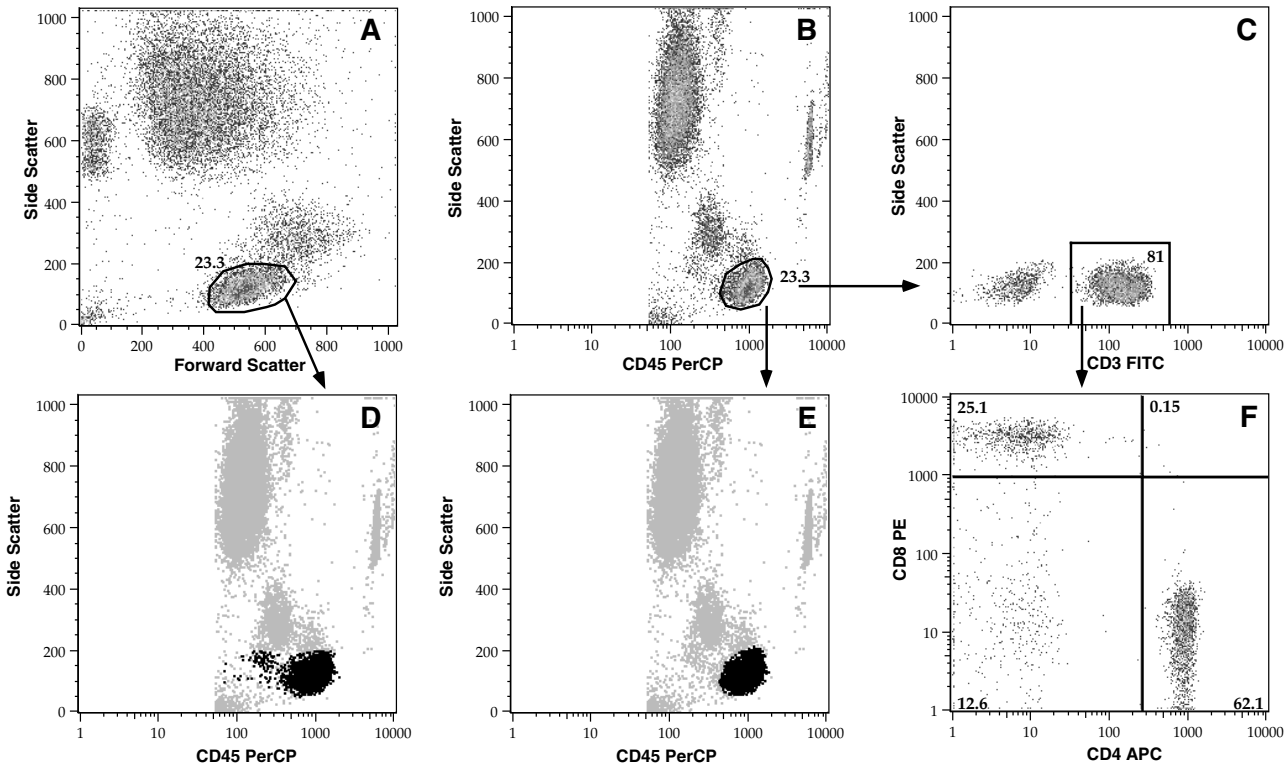

Figure 1-17. Identification of human peripheral blood T lymphocytes bearing CD4 and CD8 antigens. Data provided by Frank Mandy; analysis and displays done by Jennifer Wilshire using FlowJo software (Tree Star, Inc.).

The displays in Figure 1-17 show an older and a newer gating method for defining a lymphocyte population. The gates are drawn with the aid of a mouse or other pointing device. Flow cytometric software typically provides for several types of **bitmap gating**, which allows the user to define more or less arbitrarily shaped gating regions on a dot plot or two-parameter histogram. Almost all programs allow the user to draw **polygons** to define the boundaries of gating regions; most also allow definition of regions bounded by **rectangles**, **ellipses**, or **free-form curves**. While most cell sorters make use of no more than four gating regions at any given time, data analysis software typically provides for a larger number, to facilitate deriving counts of a reasonable number of cell subpopulations in heterogeneous samples such as are obtained from blood.

Cluster Headaches

The objective of gating is the isolation, in the measurement space, of a **cluster** of cells. The term **cluster** is used in flow cytometry (and in multivariate data analysis in general) to denote any relatively discernible, reasonably contiguous region of points in a bivariate display; that may sound imprecise, but there isn't any more precise definition. You're supposed to know a cluster when you see one.

In panel A of Figure 1-17, a polygonal gate is drawn around a cluster of cells with intermediate values of forward scatter and low values of side scatter; it was established by sorting experiments in the 1970's that most of the cells in such a cluster were lymphocytes[157], and lymphocyte gating was incorporated into analysis of lymphocyte subsets at a fairly early stage in the game[175-6]. However, there was some concern that cells other than lymphocytes might be found in the gate. If one were interested only in T lymphocytes, it would be possible, as Mandy et al demonstrated in the early 1990's[1027], to define a well isolated cluster of these cells on a display of anti-CD3 antibody fluorescence vs. side scatter (look back at Figures 1-13 to 1-16). This did not satisfy the HIV immunologists, who wanted to know not only the absolute number of CD4-bearing T cells per unit volume of blood, but also what percentage of total lymphocytes the CD4-bearing T cells represented. The current practice for defining a lymphocyte cluster uses a two-dimensional display of anti-CD45 antibody fluorescence vs. side scatter, as shown in panel B of Figure 1-17, taking advantage of the fact that lymphocytes have more CD45 antigen accessible on their surfaces than do monocytes and granulocytes[1251].

Painting and White- (or Gray-) Washing Gates

The gates in panels A and B of Figure 1-17 have been painstakingly drawn so that each includes 23.3 percent of the total number of events (where events include cells, doublets, debris, and the counting beads added to the sample). We have decided to accept the CD45/side scatter gate in panel B as the "true" lymphocyte gate; the question then comes up as to whether the forward scatter/side scatter gate in Panel A contains cells that would not fit into this gate.

In order to answer this question, we need a way of finding the cells in the forward scatter/side scatter gate on a display of CD45 vs. side scatter. Most modern data analysis programs incorporate the means to do this; the user can associate a different color with each gate set, thus allowing cells falling within that gate to be distinguished on plots of parameters other than those used to set the gates. Becton-Dickinson's "Paint-A-Gate" program was one of the first to provide this facility.

When you are working on a low budget, and restricted to monochrome displays, you can always emulate Whistler and use shades of gray instead of colors, as has been done in panels D and E of Figure 1-17. In this instance, the cells from the gates in panels A and B have, respectively, been shown in black in panels D and E; all of the other cells appear in light gray. The panel D and E displays also use a convenient feature found in the FlowJo program; the dots can be, and here are, made larger. This can be useful when one tries to show very small subpopulations in dot plots, and, indeed, we see that there are a few cells from the forward scatter/side scatter gate of panel A that show up in panel D outside the "true" lymphocyte gate as defined using CD45/side in panel B. Of course, the cells from the gate in panel B remain in the same positions in panel E; since we started out assuming that the gate in panel B was the true gate, you can't really call that a whitewash.

Moving right along, in this case to panel C, we will look only at the cells from the lymphocyte gate defined in panel B, on a plot of anti-CD3 antibody fluorescence vs. side scatter. We can now draw a rectangular gate around those that bear the CD3 antigen; these are the T cells.

The Quad Rant: Are You Positive? Negative!

In panel F, the T cells defined by the gate in panel C are shown on a plot of anti-CD4 antibody fluorescence vs. anti-CD8 antibody fluorescence. This plot is broken into **quadrants**, i.e., four rectangular gating regions that intersect at a single central point. The percentages of events that fall in each of the quadrants are indicated. There are clear clusters of events with high levels of CD8 and low levels of CD4 and of events with high levels of CD4 and low levels of CD8, a small but respectable number of events with low levels of both, and a few events with high levels of both. At first, it seems as if all's right with the world. But maybe there's a problem with our world view.

Dividing measurement spaces into quadrants is, in part, a throwback to the old days of flow cytometry without computers, when gates were implemented using hardware, and it was much easier to make them rectangular than it was to make them any other shape. Quadrants work best when the data fall neatly into rectangular regions, and when cells either have a lot of a particular antigen or other marker, making them **positive**, or very little or none, making them **negative**. The CD4-CD8 distribution of peripheral blood lymphocytes is about as good an example of this situation as can be found, but, even here, we see that, while the events divide clearly into positives and negatives on the CD4 axis, there are some events with intermediate levels of CD8.

If we were looking at cells from the thymus, quadrants wouldn't work well at all, because there are a lot of immature T cells in the thymus that have both CD4 and CD8, some of which are acquiring the antigens and some of which are losing them, and where one draws the quadrant boundaries is pretty much arbitrary. But problems with immunofluorescence data go beyond that, and beyond quadrants.

Deals with the Devil: Logarithmic Amplifiers and Fluorescence Compensation

The need and desire to measure immunofluorescence have motivated much of the development of modern flow cytometry. However, two problems associated with immunofluorescence measurement, and the less than satisfactory techniques applied to their solution, have been frustrating to beginners and experts alike.

The first problem is that of making and representing the results of measurements encompassing a large **dynamic range**. The first flow cytometers used to make immunofluorescence measurements weren't very sensitive. The green fluorescent dye **fluorescein** was used to label antibodies, and fluorescence was measured through color glass long pass filters, which, in addition to fluorescein fluorescence, let through **cellular autofluorescence**, probably due primarily to intracellular **flavins**. The filters themselves also emitted some fluorescence when struck by stray laser light. This made it impossible to detect fewer than several thousand antibody-bound fluorescein molecules on an unstained cell. However, since the maximum number of antibody-bound fluorescein molecules present on a cell might be a million or more, it was desirable, even before sensitivity was increased, to have some useful way of expressing results that varied over the three decade range between 1,000 and 1,000,000.

One obvious technique was to report and display results on a **logarithmic scale**. You can see examples of this in Figure 1-17, if you look at the numeric values and the positions of the tick marks on the axes of panels B, C, D, E, and F. Although the linear scales shown for forward and side scatter measurements (as in panel A) are accurate, the logarithmic scales may only be approximate. When analog data are digitized to relatively high resolution (16 to 20 bits), it is possible to convert signals accurately from a linear to a four decade (range 1 to 10,000) logarithmic scale and back using a digital computer; some modern cytometers employ this technique. However, in the 1970's and 1980's, the high-resolution ADCs needed to implement this procedure simply weren't available. The stopgap solution, which is still in use by some manufacturers, was to employ **logarithmic amplifiers**, commonly if not affectionately known as **log amps**.

A log amp is an analog electronic circuit that, in principle, puts out a voltage or current proportional to the voltage or current at its input. So far, so good. The bad news is that the proportionality constant may vary with time, temperature, input voltage, and, I suspect, the experimenter's astro-

logical sign. A log amp isn't a log table, or even approximately like one. The worse news is that nobody much cared how bad log amps were until the late 1980's, when people got interested in trying to make quantitative measurements of immunofluorescence and got really screwed up trying to convert from logarithmic to linear scales and back. We can expect the trend toward digital processing will continue, allowing logarithmic amplifiers to be replaced or, alternatively, monitored and calibrated; either approach should result in increased accuracy of representation of measurements on logarithmic scales.

A different set of complications was introduced by the development of antibody labels that enabled flow cytometers with a single illuminating beam (488 nm) to be used to make simultaneous measurements of immunofluorescence from several cell-bound antibodies. The first of these labels was the yellow fluorescent **phycoerythrin (PE)**, a protein found in the photosynthetic apparatus of algae. By attaching dyes to this molecule, making what are called **tandem conjugates**, it is possible to obtain fluorescence emission at longer wavelengths. When the rhodamine dye **Texas red** is attached to phycoerythrin, the resulting conjugate emits in the orange spectral region (620 nm); phycoerythrin with the **indodicarbocyanine** dye **Cy5** attached emits in the red (670 nm). Tandem conjugates of phycoerythrin with the cyanine dyes **Cy5.5** and **Cy7** emit even farther in the red or near infrared, at 700 and 770 nm. Some flow cytometers now in commercial production can be used to make simultaneous measurements of cells labeled with fluorescein, phycoerythrin, and all of the tandem conjugates just mentioned; most allow fluorescence in at least three spectral regions to be measured. The raw measurements, however, will not leave us in a state of conjugate bliss; we still have to contend with the problem of **compensation** for **fluorescence emission spectral overlap** between the labels, which only gets worse as the number of labels excited at a single wavelength increases. Figures 1-18 and 1-19 (pages 37 and 38) should provide some understanding of the problem and its solution.

Most fluorescent materials emit over a fairly broad range of wavelengths. When we describe fluorescein as green fluorescent, what we really mean is that if you look at it under a fluorescence microscope, the fluorescence looks green, and that if you measure the spectrum in a spectrofluorometer, the emission maximum is in the green spectral region. When we try to measure fluorescein fluorescence in a flow cytometer, we typically use a detector fitted with a green filter that passes wavelengths between 515 and 545 nm. However, as can be seen in Figure 1-18, the emission spectrum of fluorescein doesn't start abruptly at 515 nm and stop abruptly at 545 nm; it extends out well beyond 600 nm, although the fluorescence at the longer wavelengths is considerably less intense. There's quite a bit of emission from fluorescein in the 560-590 nm spectral region that we call yellow, and in which the emission maximum of phycoerythrin lies. That means that if we were to stain cells or other particles with fluorescein and nothing else, and measure them in a flow cytometer with both green and yellow detectors, we'd pick up a strong signal in the green detector, and also detect some signal in the yellow detector.

The phycoerythrin emission spectrum also extends well beyond the 560-590 nm yellow wavelength range we use for measurements of phycoerythrin fluorescence. There is some emission below 560 nm, in the 515-545 nm green region, and considerably more above 580 nm. If we put cells stained with phycoerythrin and nothing else into the cytometer, we'd get the strongest signals from the yellow detector, and some signals from the green, orange, and possibly the red detectors as well. The same argument holds for the orange and red fluorescent tandem conjugates; each of these will definitely be detectable in the detector intended to measure the other, and signals from the orange conjugate will show up at the yellow detector as well, and possibly also in the green one.

If we put a cell sample stained with antibodies labeled with fluorescein, phycoerythrin, and the orange and red PE-Texas red and PE-Cy5 conjugates into the machine, the signal we get from the green detector is going to be comprised mostly of fluorescein fluorescence, with a smaller contribution from phycoerythrin fluorescence, possibly a wee bit from the orange conjugate, and some from cellular autofluorescence. The signal from the yellow detector will represent mostly phycoerythrin fluorescence, with substantial contributions from fluorescein and the orange conjugate, possibly some from the red conjugate, and some from autofluorescence. And so on for the signals from the orange (600-620 nm) and red (660-680 nm) detectors. Now, how much is "some," "a substantial contribution," or "a wee bit"?

That will depend on the gain settings used for the various detectors. Once these are set, it is fairly simple to quantify the degree of spectral overlap. For example, suppose we measure cells or beads stained only with fluorescein, and they produce signals with a mean intensity (peak height or area) of 5 V from the green detector and signals with a mean intensity of 1 V from the orange detector. If we were then to measure cells stained with fluorescein and phycoerythrin, and we wanted to remove the fluorescein contribution from the orange detector signal, we could subtract 1/5 of the green signal intensity. If a doubly stained cell yielded a signal of 1 V from the green detector, we'd subtract 0.2 V from the orange signal, no matter what the value of the orange signal was; if the cell yielded a 4 V signal from the green detector, we'd subtract 0.8 V from the orange signal, and so on. Well, actually, we'd also have to do the reciprocal calculations to figure out how much of the orange signal to subtract from the green signal to remove the contribution due to phycoerythrin. In principle, though, we could figure out the whole business using high school algebra, by solving **simultaneous linear equations**. Linear equations… aye, there's the rub.

All of the operations involved in fluorescence compensation must be performed on **linear** signals. You have to make the measurements on a linear scale to determine the fractions of fluorescence signal at each detector due to each fluores-

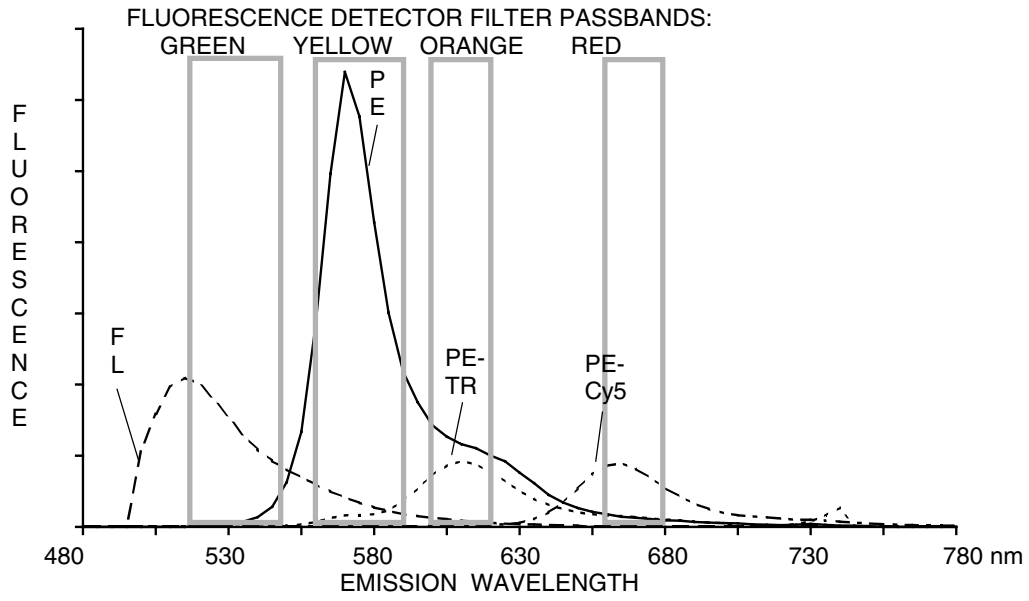

Figure 1-18. Why fluorescence compensation is necessary. The emission spectra shown are for equal concentrations (mg antibody/mL) of mouse anti-human IgG directly conjugated with fluorescein (FL), phycoerythrin (PE), and tandem conjugates of phycoerythrin with Texas red (PE-TR) and Cy5 (PE-Cy5), with excitation at 490 nm. Boxes demarcate the passbands of the green (530 nm), yellow (575 nm), orange (610 nm), and red (670 nm) filters used on the fluorescence detectors. Spectra are corrected for PMT responsivity differences at different wavelengths.

cent label, and the subtractions needed to make the necessary corrections also have to be done in the linear domain. But, as you remember, we usually tend to feed signals from immunofluorescence measurements through logarithmic amplifiers. How, then, do we introduce the fluorescence overlap compensation?

What happens in most older flow cytometers is that yet another analog circuit is built in between the preamplifier outputs and the log amp inputs. The circuit is something like an audio mixer, except that it subtracts signals instead of adding them; the operator adjusts one knob to determine the amount of green signal to subtract from yellow, another to determine the amount of yellow signal to subtract from green, etc. For two colors, this isn't all that hard to do. For three colors, you need six knobs, although you can get away with four if you ignore the green-orange and orange-green interactions. For four colors, you should have twelve knobs, though you might get away with eight. Each knob, of course, is attached to a **potentiometer**, or **variable resistor**, which, as was noted in the discussion of window comparators on p. 28, is basically a volume control. That starts to add up to a lot of electronic circuitry. Things may look neater if you let a computer control the compensation using **digital-to-analog (D-A) converters**, but you still end up with a lot of electronics at the input of your log amps.

Now, the whole reason we bother using log amps is to get a large dynamic range for our measurements. If we want a four decade dynamic range, with the top of the highest decade at 10 V, we end up with the top of the next highest at 1 V, the top of the next highest at 100 mV, and the lowest decade encompassing signals between 1 and 10 mV. If you want to process signals between 1 and 10 mV, you have to keep the noise level below 1 mV. I've measured noise in a number of older flow cytometers from a number of manufacturers, and I haven't run across one with noise below 1 mV at the preamplifier outputs. The more electronic components you stick in the circuit, the more opportunities there are to increase the noise level, and my considered opinion is that it is unlikely that a system that implements four-color compensation in electronics will be able to maintain the low noise level needed to insure a true four decade dynamic range.

Quite aside from all that, though, most people can't solve simultaneous linear equations in their heads, and those few who can probably can't manage to solve equations and twiddle knobs on compensation circuitry at the same time. You have a reasonable chance of getting two-color compensation close to right by eye; three-color compensation gets a little tougher, and you're kidding yourself if you think you can do four-color compensation correctly without solving equations. As far as I know, the manufacturers have capitulated completely on the subject of compensation for more than four colors; the knobs are gone.

There was really no choice. If you keep all the electronic measurements linear, using an A-D converter with 16 or more bits' precision, you can dispense with 1) all of the knobs and their associated electronics, 2) all of the log amps, and 3) the semiempirical process of knob twiddling for fluo-

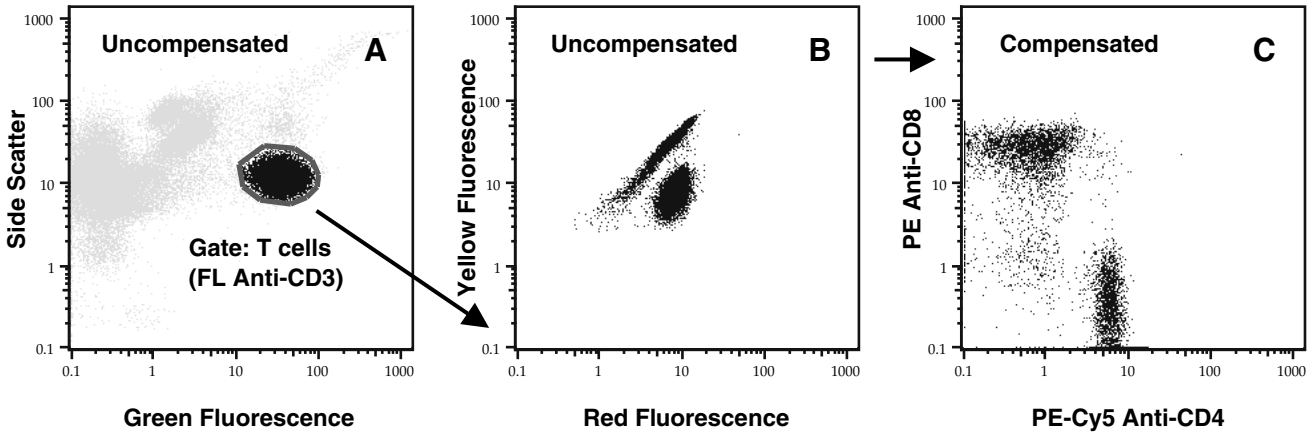

Figure 1-19. How compensation gets data to fit into quadrants.

rescence compensation. The simultaneous linear equations can be solved using digital computation, which can also do highly accurate conversions to a logarithmic scale. Once you have access to high-resolution digitized data, the logarithmic scale is only really needed for display, anyway; any statistical calculations that need to be done can be done on the linear data.

A significant advantage of high-resolution digitization of data is that you can go back to data that were not properly compensated when they were collected, transform them from a log scale to a linear one, if necessary, and redo the compensation. There are flow cytometry software packages that will let you play this game with log scale data that were digitized using 8- or 10-bit ADCs, but you end up with plots that have "holes" in the clusters due to the substantially larger, unavoidable digitization errors associated with lower resolution converters. The plots can be, and usually are, made more lovely to look at by **dithering**, adding random numbers to the data values. This technically degrades the quality of the data, but not by that much. I used to disapprove of it; I am now willing to accept it as yet another of the many deals with the devil that have to be made at the current state of the art. Within a few more years, almost all of the instruments in use will have higher resolution data analysis, and the plots of flow cytometric data will look pretty without benefit of dithering and diddling.

Evils of Axes: Truth in Labeling Cells and Plots

Mislabeling of axes, usually unintentional (I hope), is seen all too often in plots of flow cytometric data. Beginners and old-timers do it, and the mislabeling gets by journal reviewers, editors, and proofreaders. With the aid of Figure 1-19, which illustrates the effects of compensation, we can consider why mislabeling may occur and how to avoid it.

The data in Figure 1-19 were taken from a sample of whole blood stained with fluorescein anti-CD3 antibody, phycoerythrin anti-CD8 antibody, and phycoerythrin-Cy5 anti-CD4 antibody. Erythrocytes in the sample were lysed, and the sample was fixed with a low concentration of formaldehyde, before analysis. Panel A shows a dot plot of green fluorescence vs. side scatter, with both parameters displayed on a 4-decade logarithmic scale. A polygonal gate is drawn around a cluster I claim are T cells; the cells (events, if we want to be more precise) in this gate are plotted in black, while the remainder of the population is plotted in light gray.

If you look at Figure 1-14 (p. 31), you will notice that it is also a plot, in this case, a two-dimensional histogram, with anti-CD3 fluorescence on the x-axis and side scatter on the y-axis. In Figure 1-14, the y-axis is explicitly labeled as linear, and the x-axis as log, since one cannot tell whether the scale is log or linear simply by looking at the superimposed grid. The logarithmic scales on the axes of the panels in Figure 1-19 provide us with tick marks that would tell us that the scale was logarithmic even without the associated numbers, which are simply arbitrary indicators of intensity.

However, the x-axis of Figure 1-14 is labeled as "Log CyC-CD3 Flu," which means that this axis represents the intensity of fluorescence, on a logarithmic scale, of an anti-CD3 antibody, labeled in this case with PE-Cy5, with CyC being an abbreviation for one of the trademarked versions of this tandem conjugate label. The x-axis in panel A of Figure 1-19 is labeled "Green Fluorescence." What's the difference?

The difference is that fluorescence compensation has been applied to the data in Figure 1-14, but not to the data in panel A (or panel B) of Figure 1-19. So what is displayed on the x-axis in panel A is really green fluorescence, most of which is from the fluorescein label on the anti-CD3 antibody, but some of which is from the PE anti-CD8 and PE-Cy5 anti-CD4 antibodies. And some is probably from cellular autofluorescence, but we'll neglect that for the time being. We can get away with drawing a T cell gate using the uncompensated data because the fluorescein fluorescence pretty much dominates the uncompensated signal.

The situation is quite different when we look at panel B of Figure 1-19. The cells in this dot plot are only those with side scatter and fluorescence values falling within the T-cell gate shown in panel A. The axes of panel B are labeled as

showing red and yellow fluorescence, both on logarithmic scales, and I labeled them that way because the data are not compensated. There are two major clusters of cells/events visible in panel B, but points in each display significant intensities of both red and yellow fluorescence.

Panel C of Figure 1-19 shows the same cells, i.e., those in the original T cell gate, after compensation has been applied. What compensation has done is solve three linear equations in three unknowns; this gives us the fluorescence intensities of the fluorescein anti-CD3, PE anti-CD8, and PE-Cy5 anti-CD4 antibodies, which can now be plotted as such, allowing the x- and y-axes of panel C to be labeled "PE-Cy5 anti-CD4" and "PE anti-CD8." The major clusters of cells, representing CD-4 bearing T lymphocytes (often described as $CD3^+CD4^+$ cells, where the superscript "+" denotes positive) and CD-8 bearing T lymphocytes ($CD3^+CD8^+$ cells), are clearly visible, and could be fit nicely into quadrants.

Now, it would probably be perfectly legitimate to label the x-axis of panel C as "PE-Cy5 CD4," or even just "CD4," and the y-axis as "PE CD8," or just "CD8." However, if you want to be picky, what you are looking at is antibody bound to the cells. There's little doubt that almost all of the anti-CD4 antibody bound to T cells is bound to CD4 antigen on the cell surfaces, or that almost all of the anti-CD8 antibody bound to T cells is bound to cell surface CD8 antigen. On the other hand, both Figure 1-14 and panel A of Figure 1-19 show apparent binding of anti-CD3 antibody to monocytes and granulocytes; this is almost certainly **nonspecific binding**, which can occur via a variety of different mechanisms, and if we haven't got "truth in labeling" for the cells, we won't have it for the axes.

Some labels for axes should get the axe right away. The first candidates on my hit list are "FL1," "FL2," "FL3," etc., which usually mean green (515-545 nm), yellow (564-606 nm), and red (635 to about 720 nm, by my guess, limited by the characteristics of the 650 nm long pass filter at the short end and by the fading response of the detector at the long end) fluorescence. These were the fluorescence measurement ranges in the Becton-Dickinson FACScan, the first really popular benchtop 3-color fluorescence flow cytometer. The fluorescence filters in this instrument could not be changed, so at least FL1, FL2, and FL3 always meant the same thing – to FACScan users. However, in the B-D FACSCalibur, which has replaced the FACScan, while FL1 and FL2 still represent the same wavelength ranges, FL3 is different for 3- and 4-color instrument setups (650 long pass for 3-color; 670 long pass for 4-color). I think it's perfectly appropriate to use, for example, "green fluorescence," "515-545 nm fluorescence" (indicating the approximate range), or "530 nm fluorescence" (indicating the center wavelength), or even "Green (530-545 nm) fluorescence)," but let's lose FL1, FL2, FL3, etc. If you're using a long pass filter, then say, for example, ">650 nm fluorescence." If the data come from a flow cytometer with multiple excitation beams, then you might want a label like "UV-excited blue fluorescence," or "355→450 nm fluorescence." And also remember that the fluorescence color designation or bandwidth range is only really appropriate if you're displaying or talking about uncompensated data; the whole point of compensation is to get you a new set of variables that represent the amounts of probes or labels in or on the cells, rather than the measurement ranges in the cytometer.

I've already been through the labels once, in the discussion in the previous column about whether to use the antigen name or the antibody name as an axis label. However, I will return to this area to skewer the next victim on my hit list, which is "FITC." Almost everybody uses this; I have done so myself, but I have seen the error of my ways. "FITC" is a perfectly valid abbreviation for fluorescein isothiocyanate, which is the most popular reactive fluorescein derivative used to attach a fluorescein label to antibodies and other probe molecules. Once the FITC reacts with the antibody, it isn't FITC anymore, and one typically dialyzes the fluorescent antibody conjugate, or runs it over a column, in order to remove free fluorescein (the FITC is pretty much all hydrolyzed by the time you finish, anyway). Oh, yes, FITC can also be applied directly to cells, to stain proteins; once again, what you end up with bound to the proteins is fluorescein, not FITC. It would seem simple enough to use "FL" as an abbreviation for fluorescein, the way we use "PE" for phycoerythrin. I guess the problem here is that nobody wants to describe a fluorescent antibody as, say, "FL anti-CD3," rather than "FITC anti-CD3," because that might get it confused with "FL1," "FL2," "FL3," etc. Well, after I take over the world, we won't have that problem.

Then there are the scatter signals. "Forward Scatter," "Small Angle Scatter," "FALS," and "FSC" are all acceptable as axis labels; however, unless you have calibrated your measurement channel and have derived a cell size measurement from forward scatter, "Cell Size" is really inappropriate. In the same vein, I'd use "Side Scatter," "Large Angle Scatter," "90° scatter," "RALS," or "SSC" without much hesitation, but avoid "Granularity." People knowledgeable about flow will know what you are measuring; if your audience is uninitiated, you should provide a brief explanation.

It's also about time that people stopped referring to data collected with flow cytometers as "FACS data" instead of "flow cytometry data." "FACS" is the abbreviation for "Fluorescence-Activated Cell Sorter (or Fluorescence-Activated Cell sorting"), originally used by Herzenberg et al, and has been a Becton-Dickinson trade name since B-D commercialized their instrument in the 1970's. All FACSes are flow cytometers, but not all flow cytometers are FACSes, and some FACSes, such as FACScans and FACSCounts, aren't even Fluorescence-Activated Cell Sorters.

And, finally, as long as I'm ticked off, I should remind you that the tick marks on the log scale will almost certainly not represent the real scale if the instrument uses log amps without compensating for their deviations from ideal response.

When Bad Flow Happens to Good Journals

Well, you might ask, does it really matter that much whether the axis labels are absolutely correct? Won't the more egregious mistakes be picked up before manuscripts get accepted and published? Unfortunately not; there has been a great deal of weeping and wailing in the cytometry community of late about this issue, because we see a lot of bad cytometry data presentation in a lot of the more prestigious general interest and cell biology journals, and even in some of the tonier titles in hematology and immunology.

To be sure, flow cytometry may not be the only technical area in which there are such problems. A typical paper with ten or more authors might include data from gel electrophoresis, gene array scanning, confocal microscopy, etc., as well as flow data. It will probably have been reviewed by no more than three people, and they can't know all of the methodology in detail. There may be gel curmudgeons and array curmudgeons out there grumbling at least as loudly as the flow curmudgeons and the confocal curmudgeons.

In preparing this edition of *Practical Flow Cytometry*, I asked several people to send me corrected versions of data displays that appeared in papers dealing with significant refinements in technology that were critical to the biological or medical applications discussed. The referees didn't pick up the original mistakes; neither did the authors, who were good sports about responding to my requests.

Most of the time, bad flow data presentations, or even minor errors in interpretation, don't invalidate the principal conclusion(s) of a paper. When they do, the obvious remedy is for the original authors to correct their errors, or for some other people to produce another paper using better technique to reach the right conclusion. But it's much better all around if the mistakes are corrected before the manuscripts get sent in.

Meanwhile, it is incumbent upon us all to maintain a certain level of vigilance, not only when preparing cytometric data for presentation and publication, but when looking at data that others have presented or published. If it's important to you to know the details of an experiment, either because you want to duplicate it and/or adopt the methodology or because its conclusions form part of the foundation for something you want to do, work through the details. These days, it's not that uncommon to find multiparameter flow data in a paper in which little details such as the source of the antibodies used, or even which antibodies had which labels, are omitted from the "Materials and Methods" section.

Now, in an ideal world, in which everything has been done correctly, it shouldn't matter that much; I've already come out in favor of simple axis labels such as "anti-CD4" or "CD4," and, assuming that the reagents and cell preparation, initial measurements, gating, and compensation were not flawed, it shouldn't matter which antibody or label was used in an experiment. But it does. If the details you need aren't in the published paper, contact the author. That's why the e-mail address, and the snail mail address, are there. There is also an increasing likelihood that there will be another option; the journal and/or the authors may maintain a web site from which you can get technical details that were omitted from the published work.

Sorting Sorting Out

Flow sorting extends gated analysis to isolate pure populations of viable cells with more homogeneous characteristics than could be obtained by any other means. If you can get the cells that interest you into a gate in your multiparameter measurement space, you can get them into a test tube, or into the wells of a multiwell plate. Flow sorting is especially useful in circumstances in which further characterization of the selected cells requires short- or long-term maintenance in culture or analytical procedures that cannot be accomplished by flow cytometry.

A flow cytometer is equipped for sorting by the addition of a mechanism for diverting cells from the sample stream and of electronics and/or computer hardware and software that can determine, within a few microseconds after a cell passes by the cytometer's sensors, whether the values of one or more measurement parameters fall within a range or ranges (called a **sort region**, or **sort gate**) preset by the experimenter, and generate a signal that activates the sorting mechanism. The selected cells can then be subjected to further biochemical analysis, observed in short- or long-term culture, or reintroduced into another biological system (as was mentioned on p. 26, a substantial number of animals and more than a few babies have been conceived from sorted sperm).

The range of particles that can be sorted has been extended substantially in recent years; laboratory-built[2325] and commercially available instruments are now in routine use for sorting *C. elegans* nematodes and *Drosophila* embryos, while laboratory-built microfluidic apparatus has been used to sort bacteria[2326] and could, in principle, sort DNA fragments, other macromolecules, or viruses[2327]. Sorting of beads, rather than cells, has also come into use for various applications of combinatorial chemistry; the work of Brenner et al[2328-9] on gene expression analysis presents a good example.

The first generation of practical sorters accomplished cell separation by breaking the sample stream up into **droplets**, applying an electric charge to the droplets containing the selected cells, and passing the stream through an electric field, which would divert the charged droplets into an appropriate collecting vessel. A few older, and some newer, instruments use mechanical actuators to collect cells from a continuous fluid stream; while such mechanical sorters operate at lower rates (hundreds versus thousands of cells/s) than droplet sorters, their closed fluidic systems are better adapted for work with potentially infectious or otherwise hazardous materials that might be dispersed in the aerosols inevitably generated by droplet sorters. Large-particle sorters are typically mechanical, but not all of them have closed fluidic systems.

In general, sorting larger objects limits you to lower sorting speeds. If you're sorting lymphocytes, or something smaller, in a droplet sorter, you can use a 50 µm orifice, and generate droplets at rates of 100,000 droplets/s. If you're sorting pancreatic islets, which may be a few hundred µm in diameter, you'll need a 400 µm orifice, and you probably won't be able to go much above 1 kHz for a droplet generation frequency. If you're sorting *Drosophila* embryos, using a mechanical sorter (they're probably a little too big for a droplet sorter), you can measure your sort rate in dozens per second, rather than thousands.

Since cells arrive at the observation point at random times, at least approximately following Poisson statistics, there is always some probability of coincidences, which, as was noted on pp. 17 and 20-21, can pose some problem in flow cytometric analysis. Coincidences pose a fairly obvious problem in sorting, as well; they can result in your getting cells you don't want in the same droplet/well/tube as cells you do want. If the sorter is operated in the so-called **coincidence abort** mode, in which a wanted cell accompanied by an unwanted cell is not sorted, the **purity** of sorted cells is maintained, but the **yield** is decreased, while if wanted cells coincident with unwanted ones are sorted, yield is maintained at the expense of purity. All other things being equal, working at higher cell analysis rates ultimately ends up increasing the likelihood of coincidences, but there may be times when the best strategy is to sort twice, first for enrichment of a rare subpopulation, and then to increase purity of the cells recovered during the first sort.

In many cases in which flow sorting comes to mind as an obvious way of answering questions about a cell subpopulation, multiparameter analysis may allow the desired information to be obtained expeditiously without physically isolating the cells. Since the 1990's, most flow cytometry is multiparameter flow cytometry, as should be obvious from the content of the past dozen or so pages. Things were different in the bad old days.

In the 1970's, a method that was likely to come to mind for determining the distribution of DNA content in a lymphocyte subpopulation defined by the presence of a particular cell surface antigen involved staining cells with the appropriate fluorescent antibody, and then flow sorting to isolate those cells bearing the surface antigen. The sorted cells would subsequently be stained with a DNA fluorochrome such as propidium iodide; the restained sorted cells could then be run through the flow cytometer once more to determine the DNA content distribution.

This procedure was actually followed when Ellis Reinherz and Stuart Schlossman wanted to know whether there was any difference in DNA synthetic patterns between CD4- (then T4-) and CD8- (then T8-) bearing T cells; cells were stained with fluorescein-labeled monoclonal antibodies, sorted on a Becton-Dickinson FACS fluorescence-activated cell sorter, then sent to my lab, stained with propidium iodide, and analyzed on my recently built flow cytometer, which, at that time, wasn't sensitive enough to measure immunofluorescence. The chart recorder attached to my "Cytomutt" duly produced histograms of DNA content for the CD4-positive and CD8-positive cells and the antigen-negative cells, which had also been sorted.

The technically demanding and tedious exercise just described, which required at least an hour's combined use of the two instruments, did get the desired results. However, it would have been much easier to stain the entire cell population with both the fluorescent antibody and the DNA fluorochrome, making correlated multiparameter measurements of antibody fluorescence and DNA fluorescence in each cell, and using gated analysis to compile the DNA content distributions of antibody-positive and antibody-negative cells, eliminating the sorting. There was even an instrument available to us that could have done the job.

To be fair, most immunologists, faced with the same problem today, would instinctively look toward multiparameter measurement for the solution. When some colleagues and I recently had occasion to revisit the issue of DNA content of peripheral blood CD4-positive and CD8-positive T cells in the context of HIV infection and response to multidrug therapy, it was reasonably simple to deal with cells simultaneously stained for CD3, CD4 or CD8, and DNA (and RNA) content[2330].

However, those of us who have been in the flow cytometry and sorting business for a long time are likely to experience a sense of *déjà vu* when the cell and molecular biologists and geneticists bring in samples to be sorted on the basis of expression of *Aequorea* green fluorescent protein (GFP) or, more likely, one of its variants. When I wrote the previous edition of this book, Martin Chalfie et al[1648] had just demonstrated the use of GFP as a reporter of gene expression; as far as he or I knew, nobody had yet done flow cytometry on cells transfected with GFP. Most cell sorting involved selection of cells bearing one or more surface antigens. Today, people who run sorting facilities tell me that a substantial amount of their time is now spent sorting samples for cells expressing GFP or its relatives. And they also mention that the people who bring in those samples often initially contemplate sorting the cells, staining them again to measure some other parameter, and reanalyzing them.

So, although multiparameter cytometry is now old hat for the immunologists, there are some other folks out there who haven't made it that far along the learning curve. I hope the above cautionary tale, and the lengthy discussion of multiparameter cytometry that has preceded it in this chapter, will help prevent unnecessary sorting. When in doubt, work with your sorter operator and facility manager.

The nuts and bolts details of sorting will be covered at length in Chapter 6; I'll devote the rest of this discussion to what is probably the most important step in designing a sorting experiment: doing the math. A lot of people think they know that state-of-the-art high-speed cell sorters can analyze at least 16 parameters and sort (into four streams) at rates of 100,000 cells/s. However, when I polled a select group of people who actually run state-of-the-art high-speed

sorters in various labs at universities, medical facilities, and biotech and pharmaceutical companies in the Boston area, I found that nobody had done more than 8-parameter analysis, and that, while a few people had run 40,000 cells/s on occasion, 20,000 cells/s was a more typical analysis rate. Ger van den Engh, who has played and continues to play an important role in high-speed sorter development, recommends that experimenters assume analysis rates no higher than 10,000 cells/s when assessing the feasibility of proposed experiments.

Now, a lot of people want to use sorting to isolate cells that make up a very small fraction of the population being analyzed. Gross et al[2331] showed that it was possible to detect and sort cells from a human breast cancer line seeded into peripheral blood mononuclear cells at frequencies ranging from 1 cell in 10^5 to 1 cell in 10^7; they reported 40% yield and 22% purity for the sorts of cells at the lowest frequency. The raw numbers may be more impressive; a sample of 1.2×10^8 cells, which should have contained 12 cancer cells, was analyzed, giving rise to 23 sort decisions, of which 5 yielded cancer cells identifiable as such by microscopy. That sounds encouraging; even at 10,000 cells/s, it would only take about 3 hours to get 5 cells. Or about 6 hours to get 10 cells. And if you wanted to get 1,000 cells, you'd have to sort for about 25 days, 24/7.

You may have noticed that, when you're looking for cells present at low frequencies, while it is advantageous to be able to analyze at high speeds, there isn't much need for a high-speed sorting mechanism. In the above example, the sort frequency was 8/hr. There are a lot of people taking up time on very expensive, multiparameter high-speed sorters doing low frequency sorts based on one- or two-parameter measurements; sooner or later, somebody is going to make money selling simpler instruments for those jobs. Of course, if there is a method of enriching the population for the cells of interest before you start sorting – immunomagnetic separation, for example – you should take advantage of it.

A surprisingly large number of folks seem not to be doing the math before they write and submit grant applications involving sorting, which, for example, propose to isolate 10^6 cells initially present at a frequency of 1 cell/10^7. Even if you had a 100% yield, that would require analysis of 10^{13} cells *in toto*, and, even if you ran the high-speed sorter at 10^5 cells/s, it would take 10^8 seconds, or a little over three years, to do the sort. And, amazing though it seems, some of these cockamamie proposals actually get funded. A grant application is typically reviewed by a few more people than review a manuscript, but, if there are enough other high-tech gimmicks in the application, there may not be a reviewer who knows enough about sorting to ask the right questions. So, do the math. Whether as an applicant or as a reviewer, you could save the taxpayers some money.

Parameters and Probes II: What is Measured and Why

Most flow cytometers used for research, and the majority of such instruments used in clinical immunology applications, measure only three physical parameters, namely, forward (or small angle) and side (or large angle) light scattering and fluorescence, even if they measure 16 colors of fluorescence using excitation from four separate light sources. A few instruments can also measure light loss (extinction), or sense electronic impedance to measure cell volume. The remainder of the discussion of parameters and probes in this chapter will deal only with scatter and fluorescence measurements; Chapter 7 is more ecumenical and more comprehensive.

In the course of introducing cytometry in general and flow cytometry in particular, I have already covered DNA content determination using various fluorescent dyes and the identification of cells in mixed populations using fluorescently labeled antibodies. If you will flip back to Table 1-1 (p. 3), you will see that there are a great many parameters and probes about which I have, thus far, said nothing at all. However, DNA stains, on the one hand, and labeled antibodies, on the other, do represent two fundamentally different types of probes.

Probes versus Labels

The chemical properties of the DNA dyes themselves determine the nature and specificity of their interactions with the target molecule. The nature and specificity of interactions of labeled antibodies with their targets is, ideally, determined solely by the structure of their combining sites; labels are added to facilitate detection and quantification of the amount of bound antibody based on the amount of fluorescence measured from the label. Under various circumstances, the labels themselves may decrease the specificity of antibody binding; this is always at least slightly disadvantageous and may be intolerable. DNA dyes can fairly be classified as probes; molecules such as fluorescein more often serve as labels. But, as usual, there are gray areas.

Fluorescein diacetate (FDA), actually diacetylfluorescein, was discussed on pp. 24-27; this is an example of a **fluorogenic enzyme substrate**. The nonfluorescent, uncharged FDA diester freely crosses intact cell membranes; once inside cells, it is hydrolyzed by nonspecific esterases to produce the fluorescein anion, which is highly fluorescent and which leaves intact cells slowly. Since most cells contain nonspecific esterases, FDA is not terribly useful as an indicator of enzyme activity; other nonfluorescent fluorescein derivatives can be used as probes for the activity of more interesting enzymes, such as beta-galactosidase. Different derivatives of fluorescein and other dyes can be introduced into cells and cleaved by esterases to produce indicators of pH, oxidation-reduction (redox) state, and the concentration of sulfhydryl groups or of ions such as calcium and potassium. So the best I can do to clarify the status of fluorescein is to say that it is a label when it is used covalently bound to a relatively large molecule such as an antibody, oligonucleotide, or protein ligand for a cellular receptor, and a probe when introduced into cells in a slightly chemically modified, low molecular weight form. The detailed discussion of probes in Chapter 7 provides examples of when this distinction breaks down.

We will now embark on a quick tour of selected parameters and probes for their measurement. Details and spectra appear in Chapter 7. It is appropriate to mention that the single most useful reference on fluorescent probes is the *Handbook of Fluorescent Probes and Research Products*[2332], edited by Richard P. Haugland; this is the catalog of Molecular Probes, Inc. (Eugene, OR). The latest printed version is the 9th Edition, which appeared in 2002. A CD-ROM version is available as well, and all the information in the handbook, and more, with updates, can also be found at Molecular Probes' Web site (www.probes.com).

Living and Dyeing: Stains, Vital and Otherwise

Before getting down to specific (and not-so-specific) stains, it's probably a good idea to define some terms relevant to staining cells and what does or does not have to be done to the cells in order to get them to stain. A dye or other chemical that can cross the intact cytoplasmic membranes of cells is said to be **membrane-permeant**, or, more simply, **permeant**; a chemical that is excluded by intact cytoplasmic membranes is described as **membrane-impermeant**, or just **impermeant**. Because permeant dyes stain living cells, they (the dyes) are also described as **vital dyes**, or **vital stains**. You will occasionally find an opposite, incorrect definition of a vital stain as a stain that does not stain living cells; don't believe it. This seems to be one of the few urban legends of cytometry.

There are numerous transport proteins that concentrate certain chemicals in, or extrude other chemicals from, cells. Many commonly used dyes, including Hoechst 33342, serve as substrates for the glycoprotein pump associated with multidrug resistance in tumor cells, and may not readily stain cells in which this pump is active; the general lesson is that the action of transporters may make it appear that a permeant compound that is efficiently extruded is impermeant. Microorganisms may have a broader range of transporters than do mammalian cells, making it risky to assume that they will handle dyes in the same way.

Staining cells with impermeant dyes requires that the membrane be **permeabilized**. This can be accomplished in the context of **fixation** of the cells. "Fixation" originally described a process that made tissue tough enough to section for microscopy and prevented it from being autolyzed by internal hydrolytic enzymes and/or chewed up by contaminating microorganisms. Most fixatives act either by denaturing proteins (e.g., ethanol and methanol) or by cross-linking them (e.g., formaldehyde and glutaraldehyde); since this is likely to change the structure of cell-associated antigens, it is common practice to stain with fluorescent antibodies before fixing cells. In general, the fixation procedures used for flow cytometry are relatively mild; one principal objective is to kill HIV and other viruses that may be present in specimens, and another is to allow samples to be kept for several days before being analyzed. In recent years, the real pathologists have been using microwave radiation as a fixative or adjunct; I have not run across reports of its use for flow cytometry.

Permeabilization without fixation can be accomplished using agents such as the nonionic detergents Triton X-100 and Nonidet P-40; permeabilizing agents may also be added to a mixture of one or more fixatives to make cytoplasmic membranes permeable to fluorescent antibodies while retaining cellular constituents, allowing staining of intracellular antigens. Several proprietary mixtures, some of which include red cell lysing agents, are available from manufacturers and distributors of antibodies.

Most sorting is done with the intention of retrieving living cells, so fixation is not an option. However, there are procedures, such as lysolecithin treatment and electroporation, which can transiently permeabilize living cells, allowing otherwise impermeant reagents to enter while preserving viability of at least some of the cells in a sample. In this context, it is important to remember that a permeant "vital" stain may eventually damage or kill cells. It is always advisable to establish that measurement conditions do not themselves perturb what one is attempting to measure.

Nucleic Acid (DNA and RNA) Stains

Although a large number of fluorescent dyes can be used to stain DNA and/or RNA, relatively few of them are specific for DNA, and most of these are sensitive to base composition (A-T/G-C ratio). **DAPI** (4', 6-diamidino-2-phenylindole), **Hoechst 33258**, and **Hoechst 33342** increase fluorescence approximately 100 times when bound to A-T triplets in DNA. All these dyes are excited by UV light (325-395 nm), and emit in the blue spectral region with maxima between 450 and 500 nm.

Chromomycin A_3 and **mithramycin** exhibit increased fluorescence on binding to G-C pairs in DNA; they are excited by violet or blue-violet light (400-460 nm) and emit in the green between 525 and 550 nm. The combination of Hoechst 33258 and chromomycin A_3 has been used with dual excitation-beam flow cytometers to discriminate the majority of human chromosomes based on differences in DNA base composition, and to demonstrate differences in base composition among bacterial species. **7-aminoactinomycin D** (**7-AAD**[735]) also enhances fluorescence (maximum around 670 nm) on binding to G-C pairs in DNA; although it is best excited by green light (500-580 nm), it can be excited at 488 nm.

Dyes such as **ethidium bromide** (**EB**) and **propidium iodide** (**PI**), both excitable over a range from 325 to 568 nm and emitting near 610 nm, increase fluorescence on binding to double-stranded nucleic acid, whether DNA or RNA, and the latter property is shared by a large number of asymmetric cyanine nucleic acid stains (e.g., the **TO-PRO-** and **TOTO-** series (impermeant), **SYTO-**series (permeant), **Pico Green**, etc.) introduced by Molecular Probes. These dyes can be used to stain total nucleic acid in cells; specific staining of DNA requires RNAse treatment. Many of the cyanine nucleic acid dyes increase fluorescence several thousandfold; they have been used for detection of DNA fragments[1144,2327,2333-4] and viruses[2335-7].

Until recently, Hoechst 33342 was the only dye that could be used reliably to determine DNA content in living cells. However, in 1999 and 2000, Smith et al[2338-9] reported that **DRAQ5**, an anthraquinone dye with an excitation maximum around 650 nm and an emission maximum near 700 nm when bound to DNA, could also provide a reasonably good DNA content histogram. DRAQ5 can also be excited at 488 nm, albeit somewhat inefficiently.

DRAQ5 does not increase fluorescence significantly on binding to DNA; it stains nuclei because it is present in higher concentrations in association with nuclear DNA than elsewhere in the cell, and the quality of staining is thus relatively more dependent on relative concentrations of dye and cells than is the case for most other DNA dyes. **Acridine orange** (**AO**), like DRAQ5, does not increase fluorescence on binding to either DNA or RNA, but stains by virtue of its concentration on the macromolecules.

Darzynkiewicz et al showed, beginning in the mid-1970's, that, after cell membrane permeabilization and acid treatment, AO could be used for stoichiometric staining of DNA and RNA in cells[262-3,525,1348-9]. On excitation with blue light (488 nm is eminently suitable), the DNA-bound monomer fluoresces green (about 520 nm); the RNA-bound dye forms red (>650 nm) fluorescent aggregates. The combination of DNA and RNA staining allows the cell cycle to be subdivided into stages that are not distinguishable on the basis of DNA content alone, permitting discrimination between G_0 and G_1 cells.

Relatively specific staining of double-stranded (predominantly ribosomal) RNA in cells can be achieved using a combination of **pyronin Y** (excitable at 488 nm with emission in the yellow around 575 nm), which stains RNA, with one of the Hoechst dyes, which binds to DNA and prevent DNA staining by pyronin Y. In a dual excitation-beam instrument (UV and 488 nm), DNA and RNA content in living cells can be estimated simultaneously from pyronin Y and Hoechst 33342 dye fluorescence, providing information that is substantially equivalent to what could be obtained using AO (Fig. 1-2, p. 27) without requiring that the cells be sacrificed[113]. Cells stained with this dye combination have been sorted with retention of viability[2340-2].

Toba et al[2343-5] found that DNA and RNA could be measured in permeabilized cells using the combination of 7-AAD and pyronin Y in a system with a single 488 nm excitation beam; Schmid et al modified the staining conditions and reported improved precision and reproducibility[2346].

Fluorescence and Fluorescent Labels

Because the fluorescent label on a probe is usually not intended to interact directly with the structure to which the probe binds, labels are developed and/or synthesized predominantly for their desirable spectral characteristics.

In order for an atom or molecule – or part of a molecule; the all-inclusive term would be **fluorophore** – to emit fluorescence, it must first absorb light at a wavelength shorter than or equal to the wavelength of the emitted light, raising an electron to an excited state. Absorption requires only about a femtosecond. In order to have a high likelihood of fluorescing, a material must have a high likelihood of absorbing the excitation light; the likelihood that a molecule will absorb is quantified as the **absorption cross-section** or the **molar extinction coefficient**.

Fluorescence results from the loss of at least some of the absorbed energy by light emission. The period between absorption and emission is known as the **fluorescence lifetime**; for organic compounds, this is typically a few nanoseconds. Some of the absorbed energy is almost always lost nonradiatively, i.e., unaccompanied by emission, by transitions from higher to lower vibrational energy levels of the electronic excited state. The fluorescence emission will then be less than the energy absorbed; in other words, emission will occur at a wavelength longer than the excitation wavelength. The difference between the absorption and emission maxima is known as the **Stokes shift**, honoring George Stokes, who first described fluorescence in the mid-1800's. Stokes shifts are typically only a few tens of nanometers.

Fluorescence is an intrinsically quantum mechanical process; the absorbed and emitted energy are in the form of photons. The **quantum yield** and **quantum efficiency** of fluorescence are, respectively, the number and percentage of photons emitted per photon absorbed; they typically increase with the cross section and extinction coefficient, but are also dependent on the relative likelihoods of the excited molecule losing energy via fluorescence emission and nonradiative mechanisms. The quantum yields of some dyes used in cytometry are quite high, above 0.5, but it is important to note that quantum yield, particularly for organic fluorophores, is affected by the chemical environment (i.e., the pH, solvent polarity, etc.) in which the molecule finds itself. If an excited molecule that might otherwise fluoresce instead loses energy nonradiatively, for example, by collision with solvent molecules, it is said to be **quenched**; once returned to the electronic ground state, it can be reexcited. However, there is usually a finite probability that light absorption will be followed by a change in molecular structure, making further cycles of fluorescence excitation and emission impossible; this is called **(photo)bleaching**.

In principle, increasing the illumination intensity can increase the intensity of light scattering signals without limit. However, this is not even theoretically possible for fluorescence signals, because, at some level of illumination, all the available molecules will be in excited states, leaving no more to be excited if illumination intensity is further increased. This condition of **photon saturation** is often reached in cytometers which use laser powers of 100 mW or more; bleaching, which may also make the dependence of emission intensity on excitation intensity less than linear, is noticeable at power levels of tens of milliwatts. Saturation and bleaching are discussed at length by van den Engh and Farmer[1130].

When an excited fluorophore is in close proximity (typically no more than a few nanometers) to another fluorophore, nonradiative energy transfer (**fluorescence resonance**

energy transfer, or **FRET**) from the excited (donor) molecule to the nearby acceptor molecule may occur, followed by fluorescence emission from the acceptor in its emission region. The probability of energy transfer increases with the degree of overlap between the absorption spectrum of the second fluorophore and the emission spectrum of the first. I have said "fluorophore" rather than "molecule" here because energy transfer can occur between different structures within the same molecule. An accessible review of FRET is provided by Szöllösi et al[2347].

In the intact photosynthetic apparatus of algae and cyanobacteria, absorbed blue-green and green light is utilized for photosynthesis by a series of intra- and intermolecular energy transfers via **phycobiliproteins** to **chlorophyll**, without subsequent emission. In 1982, Oi, Glazer, and Stryer[114] reported that extracted algal phycobiliproteins could be used as highly efficient fluorescent labels with large Stokes' shifts. As you might have noticed from the extensive previous discussion, it has become common practice to attempt to improve on nature by conjugating dyes to phycobiliproteins to add an additional phase of energy transfer and further shift the emission spectrum of the **tandem conjugates**. The first such tandem conjugate, described by Glazer and Stryer in 1983[306], was made by linking **phycoerythrin** (**PE**) to **allophycocyanin** (**APC**), a phycobiliprotein which absorbs relatively efficiently, although not maximally, at phycoerythrin's yellow (575 nm) emission wavelength and which emits maximally in the red at 660 nm.

Until both flow cytometers and monoclonal antibodies became widely available in the early 1980's, the most widely used fluorescent label was **fluorescein**, usually conjugated to proteins as the **isothiocyanate** (**FITC**); second labels were only infrequently needed. Fluorescein is nearly optimally excited at 488 nm, and emits in the green near 525 nm. While rhodamine dyes had been used for two-color immunofluorescence analysis by microscopy, they were not suitable for 488 nm excitation. A small number of studies were done with yellow-excited dyes, which needed a second excitation beam, making flow cytometers substantially more expensive. Phycoerythrin (PE), which emits in the yellow near 575 nm, is maximally excited by green light but absorbs reasonably well at 488 nm. Its extinction coefficient is high enough to make the fluorescence signal from PE-labeled antibody substantially higher than that from an equivalent amount of fluorescein-labeled antibody (Fig. 1-18, p. 37).

We have already encountered tandem conjugates of PE suitable for 488 nm excitation (PE-Texas red, emitting near 610 nm; PE-Cy5, near 670 nm; PE-Cy5.5, near 700 nm; PE-Cy-7, near 770 nm). Allophycocyanin absorbs maximally in the red near 650 nm, and is well excited by red diode (635-640 nm) and He-Ne (633 nm) lasers. Tandem conjugates of APC with Cy5.5 and Cy7 emit in the far red and near infrared, as do the PE conjugates with the same dyes.

A principal disadvantage of phycobiliproteins as fluorescent labels is their large size; with a molecular weight near 240,000, PE binding increases the molecular weight of an immunoglobulin G antibody by about 150 percent. This may not be an issue when labeled antibodies or lectins are used to stain cell surface structures, but becomes one when it is necessary to use labeled reagents to demonstrate intracellular constituents. A number of lower molecular weight labels have been developed for this purpose. The symmetric cyanines[1361-4] include Cy5, Cy5.5, and Cy7, and their shorter wavelength absorbing cousins, e.g., Cy3, which can be excited at 488 nm and emits in the same region as PE; we have already run across them as acceptors in tandem conjugates. Molecular Probes has recently developed the Alexa series of dyes[2348] (also see the Molecular Probes handbook/Web site[2322]); different members of this series are excitable at wavelengths ranging from the UV to the near infrared. Alexa dyes, used alone or as acceptors in tandem conjugates, are reported to have better fluorescence yields and photostability (resistance to bleaching) than more commonly used labels with similar spectral characteristics, and seem to be coming into wider use. Low, rather than high, molecular weight labels are almost always used on oligonucleotide probes, which allow demonstration and quantification of specific nucleic acid sequences in cells or on beads or solid substrates (e.g., in gene arrays).

As was mentioned previously, it is the probe, not the label, that confers specificity; dyes must be derivatized into forms that contain a functional group, such as an isothiocyanate or sulfonyl chloride, that will allow the **reactive dye** to bind covalently to the probe. FITC, applied to cells, will stain accessible proteins. Staining of intact cells will be limited to the cell surface; in fixed or permeabilized cells, both surface and intracellular proteins will be stained.

Binary Fishin': Tracking Dyes Through Generations

Otherwise nonspecific, but persistent fluorescent staining of cellular proteins or lipids has recently been put to good use in studying cell proliferation. Since cellular proteins and lipids are apportioned more or less equally to each daughter cell during cell division, analysis of the fluorescence of cells after staining with a so-called **tracking dye** should allow determination of how many cycles of division have occurred since its ancestor was stained. The dye first widely used for such studies was **PKH26**[1551-5], a yellow fluorescent cyanine dye with long alkyl side chains that incorporates itself tightly enough into lipid bilayers that it is not readily lost from cells. It was called a tracking dye because it could also be used to follow cells that had been removed from animals, labeled, and reinjected. Estimation of the numbers of cells in various daughter generations after PKH26 labeling requires application of a mathematical model[1555].

An alternative to PKH26, **carboxyfluorescein diacetate succinimidyl ester** (**CFSE**)[2349], is a nonfluorescent fluorescein ester that enters cells and is hydrolyzed to a reactive dye by nonspecific esters; the end result is that fluorescein molecules are bound covalently to intracellular protein. Distributions of CFSE fluorescence in proliferating populations

usually show peaks indicating the positions of cells in different daughter generations; these can be analyzed with mathematical models, but it is also possible to combine sorting with CFSE labeling to isolate cells from different generations[2350], which cannot be done reliably when PKH26 is used as a tracking dye.

Membrane Perturbation: A Matter of Life and Death?

The integrity of the cytoplasmic membrane is essential to cell function. Although at least some cells can survive transient small breaches of the membrane, longer-term and/or larger defects may deprive the cell of materials it would normally accumulate, and may also expose it to toxins it would normally exclude. Thus, we tend to think that cells with a demonstrable loss of membrane integrity are dead.

Trypan blue has been the preferred probe for a **dye exclusion test for "viability,"** i.e., retention of membrane integrity, performed by visual inspection of cells under the microscope; the Bio/Physics Systems Cytograf, made in the early 1970's, measured extinction and scattering using a red He-Ne laser source, and could detect trypan blue uptake by cells. These days, people who want to do dye exclusion testing by flow cytometry typically use impermeant nucleic acid dyes such as **propidium iodide** or **7-aminoactinomycin D**, both excitable at 488 nm, and, emitting, respectively, at about 620 and about 670 nm, or the red-excited dye **TO-PRO-3**, emitting at about 670 nm. Cells that take up the dye and become fluorescent are considered to be nonviable.

Fluorescein is anionic, and, therefore, relatively impermeant; when produced intracellularly by hydrolysis of fluorescein diacetate (FDA), it leaves cells slowly, giving us a **dye retention test for "viability."** Cells with intact membranes accumulate and retain fluorescein after exposure to FDA and become (green) fluorescent; cells with membrane damage do not retain fluorescein and do not fluoresce. The fluorescein derivative **calcein**, produced in cells by esterase action after exposure to the **acetoxymethyl ester**, **calcein-AM**, is retained much more effectively than fluorescein and is now preferred for dye retention tests.

The problem with dye exclusion and retention tests is that, while the methodology works well for cells that are killed by freezing or heat or by interaction with cytotoxic T or NK cells, all of which inflict early and usually lethal damage on the cytoplasmic membrane, cells that are killed by other means, e.g., those rendered reproductively nonviable by such agents as ionizing radiation, may retain membrane integrity for days after exposure. Uptake of impermeant dyes is therefore a better indicator of nonviability than retention is of viability, but there are situations in which impermeant dyes can end up in viable cells[2351].

One can, of course, combine dyes, for example, propidium iodide and calcein-AM, which will result in cells with intact membranes exhibiting green cytoplasmic fluorescence while cells with damaged membranes show red nuclear fluorescence, but this does not solve the basic problem. And, in part thanks to cytometry, we can now distinguish one kind of death (**necrosis**) from another (**apoptosis**), making the issue of viability assays even more contentious. Darzynkiewicz et al have discussed the cytometry of cell necrobiology in detail[2352]. Disturbances in membrane organization in apoptosis, resulting in the exposure of phosphatidylserine, are usually looked for using fluorescently labeled **annexin V**[2353-4].

When viability is not an issue, measurements of fluorescence of cells over time after exposure to fluorescent dyes, drugs, or labeled drug analogs can be useful in detecting the presence of various transport proteins. Uptake or efflux kinetics in themselves can only suggest a mechanism; when the transporter or pump being investigated has been well characterized, establishing that known substrates and inhibitors affect fluorescence kinetics as predicted is critical for confirmation of the initial hypothesis.

Cytoplasmic/Mitochondrial Membrane Potential

Electrical potential differences are present across the cytoplasmic membranes of most living prokaryotic and eukaryotic cells, and also between the cytosol and the interior of organelles such as chloroplasts and mitochondria. Membrane potential ($\Delta\Psi$) is generated and maintained by transmembrane concentration gradients of ions such as sodium, potassium, chloride, and hydrogen.

Changes in cytoplasmic $\Delta\Psi$ play a role in transmembrane signaling in the course of surface receptor-mediated processes related to the development, function, and pathology of many cell types. Cytoplasmic $\Delta\Psi$ is reduced to zero when the membrane is ruptured by chemical or physical agents; mitochondrial $\Delta\Psi$ is reduced when energy metabolism is disrupted, notably in apoptosis. In bacteria, $\Delta\Psi$ reflects both the state of energy metabolism and the physical integrity of the cytoplasmic membrane.

Flow cytometry can be used to estimate membrane potential in eukaryotic cells, mitochondria *in situ*, isolated mitochondria, and bacteria[424,2355]. Older methods, using lipophilic cationic dyes such as the symmetric cyanines **dihexyloxacarbocyanine** [$DiOC_6(3)$] and **hexamethylindodicarbocyanine** [$DiIC_1(5)$] or **rhodamine 123**, or lipophilic anionic dyes such as **bis (1,3-dibutyl-barbituric acid) trimethine oxonol** [$DiBAC_4(3)$] (which is often, incorrectly, referred to as **bis-oxonol**), can detect relatively large changes in $\Delta\Psi$, and identify heterogeneity of response in subpopulations comprising substantial fractions of a cell population. All of the dyes just mentioned can be excited at 488 nm and emit green fluorescence, with the exception of $DiIC_1(5)$, which is red-excited and emits near 670 nm. Newer techniques that use energy transfer and/or ratios of fluorescence emission at different wavelengths allow precise measurement of $\Delta\Psi$ to within 10 mV or less[2356-7].

Since, in most eukaryotic cells, $\Delta\Psi$ across mitochondrial membranes is larger than $\Delta\Psi$ across cytoplasmic membranes, exposure of cells to lipophilic cationic dyes results in higher concentrations of dye in the cells than in the suspending

medium, and higher concentrations in mitochondria than in the cytosol. If cells are washed after being loaded with dye, staining of the cytosol may be minimized while mitochondrial staining persists. This is the basis for the use of $DiOC_6(3)$, $DiIC_1(5)$, rhodamine 123, and other cationic dyes to estimate mitochondrial $\Delta\Psi$; the procedure has become commonplace for studies of apoptosis, in which early increases in mitochondrial membrane permeability result in loss of $\Delta\Psi$. **JC-1**, a cyanine, exhibits green fluorescence in monomeric form and red fluorescence when aggregated at higher concentrations[1681-2], and has become popular for work on mitochondria in apoptosis.

Among other factors, action of efflux pumps, changes in membrane structure, and changes in protein or lipid concentration in the medium in which cells are suspended can produce changes in cellular fluorescence which may be interpreted erroneously as changes in $\Delta\Psi$. For example, it was observed in the 1980's that hematopoietic stem cells were not stained by rhodamine 123, and some people concluded that this reflected low mitochondrial $\Delta\Psi$; it was later found that the dye was being actively extruded by a glycoprotein pump. Getting good results from cytometric techniques for estimation and measurement of $\Delta\Psi$ demands careful control of cell and reagent concentrations and incubation times and selection of appropriate controls.

Indicators of Cytoplasmic [Ca^{++}]: Advantages of Ratiometric Measurements

The importance of calcium fluxes in cell signaling was appreciated when flow cytometry was in a relatively early stage of development, but it was not until some years later that suitable probes became available[2358]. The first probes exhibited differences in the intensity of fluorescence in the presence of low and high intracellular [Ca^{++}], but did not change either their fluorescence excitation or emission spectral characteristics to a significant degree. Since the distribution of fluorescence intensity from cells loaded with the probes was typically quite broad (a problem also associated with membrane potential probes), it was possible to appreciate large changes in cytoplasmic [Ca^{++}] affecting all or most of the cells in a population, which would shift the entire distribution substantially, but not to detect even a large change in cytoplasmic [Ca^{++}] involving only a small subpopulation of cells. This came as a disappointment to immunobiologists who hoped to use flow cytometry to detect calcium responses associated with activation of lymphocytes by specific antigens.

Roger Tsien and his colleagues, who had developed some of the earlier calcium probes, came to the rescue in 1985 with **Indo-1**[858]. This, like other probes, is a selective calcium chelator, but does not significantly perturb cellular calcium metabolism. Its fluorescence is excited by UV light; wavelengths between 325 and about 365 nm, which pretty well covers the range of UV sources available for flow cytometry, are suitable. Indo-1's attraction, however, is due primarily to the fact that there are substantial differences in emission spectra between the free dye, which shows maximum emission at about 480 nm, and the calcium chelate, which emits maximally at about 405 nm. The **ratio** of emission intensities at 405 and 480 nm in cells loaded with Indo-1 [it is introduced as an acetoxymethyl (AM) ester] can, therefore, provide an indication of cytoplasmic [Ca^{++}]. The **ratiometric** measurement cancels out many extraneous factors, most notably including the effect of cell-to-cell variations in dye content, which plague older techniques for calcium measurement and for measurement of $\Delta\Psi$. Effects of uneven illumination and of light source noise also are eliminated by virtue of their equal influences on the numerator and denominator of the ratio. This advantage, it should be noted, is common to other ratiometric measurements (e.g., of $\Delta\Psi$ and of intracellular pH) in which both parameters used in the ratio are measured at the same time in the same beam.

If aliquots of loaded cells are placed in solutions with various known Ca^{++} concentrations and treated with a **calcium ionophore** such as **A23187** or **ionomycin**, it is possible to calibrate the fluorescence ratio measurement to yield accurate molar values of cytoplasmic [Ca^{++}]. Indo-1 is widely used, at least by people with UV excitation sources in their flow cytometers[862, 1714-8].

Since there are probably more than 10,000 fluorescence flow cytometers out there that don't have UV sources, that's small comfort. Luckily, there are alternatives. In 1989, Tsien and his collaborators described a series of fluorescein- and rhodamine-based calcium indicators suitable for use with 488 nm excitation[1719]. The most widely used of these is **Fluo-3**, which has the spectral characteristics of fluorescein, but which is almost nonfluorescent unless bound to calcium. Unlike Indo-1, Fluo-3 does not exhibit a spectral shift with changes in calcium concentration. A Fluo-3 fluorescence distribution is a haystack; if you're stimulating a cell population, the haystack moves to the right when the cytoplasmic [Ca^{++}] goes up and back to the left when it goes back down. However, there is another dye, **Fura red**, also suitable for 488 nm excitation, which exhibits high fluorescence when free in solution (or cytosol) and low fluorescence when bound to calcium; a Fura red haystack moves in the opposite direction from a Fluo-3 haystack with changes in cytoplasmic [Ca^{++}]. More to the point, the ratio of fluo-3 to Fura red fluorescence provides a precise, calibratable indicator of cytoplasmic [Ca^{++}] that can be used in the majority of fluorescence flow cytometers[2358]. Both Fluo-3 and Fura red, like Indo-1, are loaded into cells as AM esters.

Finding Antigen-Specific Cells Using Tetramers

While ratiometric probes did improve the precision of intracellular calcium measurements, they did not get them quite to the point of being able to detect specific responses of very small numbers of lymphocytes to antigens. As it turned out, a more direct approach was to succeed. In 1996, Altman et al[2359] described identification of antigen-specific cytotoxic ($CD3^+CD8^+$) T cells using a fluorescently labeled complex containing four each of 1) a class I major histo-

compatibility complex (MHC) α chain, 2) β_2-microglobulin, and 3) an antigenic peptide. Since that gets to be a lot to write or say, the probes are now universally described as **tetramers**.

Antigen presentation to T cells requires binding of antigenic peptides associated with HLA proteins (class I proteins for cytotoxic [$CD3^+CD8^+$] T cells, class II proteins for helper [$CD3^+CD4^+$] T cells) on the antigen presenting cell to the T cell receptor; attempts to bind a labeled monomeric complex (1 each) of α chain, β_2-microglobulin, and peptide to cytotoxic T cells were unsuccessful because the binding affinity of the monomers was too low. Tetramers did the trick, and have come into wide use since they were originally described[2360-2]. We now have not only **class I tetramers**, reactive with cytotoxic T cells, but also **class II tetramers**, which contain MHC class II proteins, and react in an antigen-specific fashion with helper T cells[2363-4]. They're not available at the corner store quite yet, but rumor has it that it was tetramers and their possibilities that made one of the major instrument companies decide to stay in the fluorescence flow cytometry business.

Hip, Hip Arrays: Multiplexing on Slides and in Bead Suspensions

If you have been keeping up with biology at all over the past few years, it's unlikely that you have not run across gene array technology[2365-7], which allows the expression of hundreds or thousands of genes to be studied by, for example, hybridizing different colors of labeled cDNA derived from the same cells grown under different circumstances to a slide on which the requisite genetic sequences have been synthesized or deposited in small spots. The slides are then scanned, allowing differences in expression to be detected by color differences resulting from the presence of different amounts of the cDNAs on each spot. The array concept has taken off; we have gene arrays, protein arrays, cell arrays, and even tissue microarrays, which allow high-throughput molecular profiling of tumors[2368].

Multiplex analysis allows flow cytometry to accomplish some of the same tasks for which gene arrays are now used. It occurred to various people in the mid-1980's[1820-34] that various types of ligand binding assays could be done in a flow cytometer by using fluorescence measurements to quantify binding to appropriately coated beads. By using a different size and/or color bead for each of a number of assays, it would be possible to perform all of them at once on a single sample in a single tube[2369].

The latest incarnation of multiplex analysis uses a small, dedicated flow cytometer capable of identifying as many as 100 different colors of beads, and has been applied successfully to both protein[2370-1] and nucleic acid[2372-4] analysis. In a study monitoring multiple pathogenesis-related genes simultaneously in chemical-treated and control *Arabidopsis* samples, Yang et al[2375] reported that a multiplexed flow cytometric assay they developed yielded results comparable to those obtained from a slide-based gene array.

GFP and Its Relatives: Mild-Mannered Reporters

The 1994 report by Chalfie et al[1648] on the use of *Aequorea* green fluorescent protein (GFP) as a reporter of gene expression quickly spawned a growth industry. GFP mutants are now available with cyan, green, and yellow fluorescence and with excitation characteristics far better suited to flow cytometry (and imaging, confocal microscopy, etc.) than the wild type protein. Moreover, GFP variants have been engineered to behave as sensors of such functional parameters as intracellular (or intracompartmenal) pH, [Ca^{++}], etc., and, using energy transfer between molecules with different spectra, for quantitative measurements of protein-protein interactions[2376-8]. The mild-mannered reporter has shed glasses and business suit and emerged from the phone booth as Supermolecule. I have already mentioned that sorting for fluorescent protein expression now seems to account for a significant amount of flow facilities' time; this trend can be expected to continue.

Beyond Positive and Negative: Putting the -Metry in Cytometry

If you spend most of your flow cytometer time doing immunofluorescence analysis, you can pick up some bad habits. Given an instrument that often costs upwards of a hundred thousand dollars, is full of fancy electronics, has its own computer attached to it, and can probably detect a few hundred molecules of fluorescent dye in or on a cell, it does seem that we underutilize its capacities when we report the results of highly sensitive and precise fluorescence measurements as "positive" and "negative."

To be sure, sometimes "positive" and "negative" are good enough to get the job done. In the previous examples of counting various types of T lymphocytes in human peripheral blood, we defined the subpopulation of T cells by their scattering characteristics and by the presence of the CD3, CD4, or CD8, cell surface antigens, and, in general, the cells we're looking at either have a substantial amount of the antigen or have little or none. When we look at our "CD3-positive" cells, they either do or do not have substantial amounts of the CD4 and CD8 antigens. We don't need to be experienced in flow cytometry to know "positive" and "negative" when we see them in these contexts, and, using these concepts, we can obtain a satisfactory answer to the question, "What are the relative proportions of ($CD3^+CD8^+$) and ($CD3^+CD4^+$) T cells in this blood sample?"

However, if the question we are asking is, "What proportion of ($CD3^+CD4^+$) T cells are activated?," we may need to extend our conceptual framework somewhat, both in terms of biology and in terms of cytometry. "What is an activated lymphocyte?," "What is a cancer cell?," and "What is a dead cell?" are major quasitheological questions guaranteed to provoke debate among analytical cytologists for a while to come. But let's suppose we have decided to define activation in terms of expression of the CD25 antigen, which is the cell surface receptor for the cytokine interleukin-2. Well, then,

we can just gate the T cells, further gate the CD4-positive cells, and then count the CD25-positive and negative cells, right? Unfortunately not. The number of CD25 molecules on an inducer T cell seems to range from hundreds or less to many thousands; the problem in defining "positive" and "negative" is that there is no clear breakpoint.

Well, then, perhaps we could say that a cell with more than 5,000, or 10,000, or some other seemingly arbitrary number of molecules of CD25 on its surface is activated. That might work, provided we had a way of determining the number of molecules from the immunofluorescence measurement. As it turns out, this can be done, but it isn't always as easy as it looks.

The hematology counters are ahead of the fluorescence flow cytometers in this department. They all report **red cell indices**, including erythrocytes' **mean corpuscular volume (MCV)** in femtoliters and **mean corpuscular hemoglobin (MCH)** in picograms. Every instrument in every lab everywhere uses the same units. Way back in 1977, I suggested that we should have "**white cell indices**," which didn't go over resoundingly well in the Dark Ages of polyclonal antisera. The proposition has been better received of late, for several reasons. The need is more apparent, our apparatus and reagents are better, and there are people interested in developing and testing standardized materials that will make it possible for everyday users of flow cytometry to do quantitative immunofluorescence measurements. Figure 1-20 illustrates one technique, which uses beads with known numbers of antibody binding sites as standards.

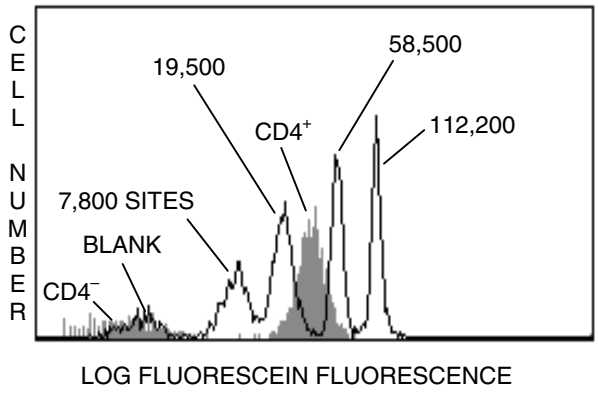

Figure 1-20. Fluorescence intensities of CD4-positive and negative cells (plotted as bars) compared with intensities of beads bearing known numbers of antibody binding sites, stained with the same fluorescein-anti-CD4 antibody as was used to stain the cells.

What Figure 1-20 shows is that most of the CD4-positive cells in the same lysed whole blood sample as is depicted in Figures 1-13 through 1-16 exhibit fluorescein fluorescence intensities consistent with there being somewhere between 19,500 and 58,500 antibodies bound to the cell surface. Is that a good number? Actually, it's probably a little low; people who've done the experiments carefully seem to come up with an average of about 50,000 molecules of CD4 per CD4-positive cell. I may have come up with the lower number because there wasn't enough antibody added to the blood sample to bind to all of the available CD4 molecules; I didn't **titrate** the antibody, i.e., determine whether adding more antibody would have increased the cells' fluorescence intensities. So, as I said, it isn't always as easy as it looks.

However, there has been a great deal of work done on improving quantitative fluorescence measurement since the last edition of this book was written; for now, it's probably enough to mention that an entire issue of the journal *Cytometry* was devoted to the topic in October, 1998[2379].

1.5 WHAT'S IN THE BOX: FLOW CYTOMETER ANATOMY, PHYSIOLOGY, AND PATHOLOGY

It may have occurred to you that I have spent a great deal of time dealing with history, data analysis, parameters, and probes without getting into the details of how a flow cytometer works. That fits in with my idea that what we should be concerned with, first and foremost, is what information we want to get out of the cells and what we have to do to the cells to get it. It is now fairly clear that, although we can derive some information about cell size and morphology from light scattering signals, getting the details about biochemistry and physiology will require treating the cells with one or more fluorescent probes. We are now ready to consider more of the details of how the fluorescence of those probes is measured.

Light Sources for Microscopy and Flow Cytometry

There are substantial differences in time scale between flow cytometry and microscopy. A human observer at a microscope moves different cells into and out of the field of view at a rate that is, under any circumstances, much slower than the rate at which cells are transported through the observation region (or, if you prefer, past the "interrogation point," which always seems to me to describe a "?") of a flow cytometer. The response time of the human observer is pretty long, i.e., hundredths of seconds, or tens of thousands of microseconds. That's why movies and television work; changing the picture a few dozen times a second produces the illusion of continuous motion. In flow cytometry, a cell passing through the apparatus is typically illuminated for somewhere between one and ten microseconds. This disparity in observation times means, among other things, that flow cytometers need more intense light sources than are commonly used in microscopes.

Both the **sensitivity** (i.e., how much light can be detected) and **precision** (i.e., how reproducibly this can be done) of light measurements are functions of the **amount of light**, i.e., the **number of photons**, reaching the detector. The human eye is an extremely sensitive photodetector; when properly dark-adapted, a person with good eyesight may well perceive single photons emitted from weakly fluo-

rescent or luminescent objects. The quantum nature of light obviously does not allow for any improvement upon this level of sensitivity in the electro-optical photodetectors used in flow cytometers.

Therefore, to make a flow cytometer comparable in sensitivity to a human observer, we would expect to have to get approximately the same amount of light from the observation region of the flow cytometer in a few microseconds as is collected by the observer at the microscope in a few milliseconds. Since the amount of light collected is, in general, directly dependent on the intensity of illumination, a cytometer needs a light source approximately a thousand times as bright as would be needed in the microscope.

The term **brightness**, when used in a technical sense, denotes the amount of light emitted from or through a unit surface area or solid angle, rather than the total amount of light emitted from a source. By this criterion, the 800 μW **laser** in a supermarket bar code scanner is brighter than the sun, and practically any laser can potentially be used as a light source for flow cytometry. The requisite brightness is also found in some kinds of **arc lamps** (high-pressure mercury and xenon lamps, sometimes specified as "short arc" lamps).

The majority of fluorescence flow cytometers now in use are benchtop models with a single blue-green (488 nm) illuminating beam, derived from an air-cooled argon ion laser. If a benchtop apparatus has a second illuminating beam, it is usually red (nominally 635 nm), coming from a diode laser. Larger instruments, such as high-speed sorters, use water-cooled argon and krypton ion lasers, which can be tuned to produce emission at a variety of UV (350-364 nm) and visible wavelengths; some systems obtain UV emission at 325 nm from an air-cooled helium-cadmium laser. Typical laser powers range from 10 to 25 mW in benchtop cytometers and up to hundreds of milliwatts in larger systems.

Instrument Configurations: The Orthogonal Geometry

Flow cytometers using arc lamp sources have been and still may be built around upright or inverted microscopes, simply by placing the **flow cell** or **flow chamber** in which cells are observed where the slide would normally go. Most modern fluorescence flow cytometers, however, use laser sources, and employ a different optical geometry, which is shown schematically in the intimidating but informative Figure 1-21 (the uncaptioned color version of the figure on the back cover may be helpful). The cytometer shown in the figure is designed to measure light scattering at small and large angles and fluorescence in four spectral regions.

The figure is a top view. If you look carefully along the left side, about halfway up from the bottom, you'll see the cell, which is, or at least should be, the *raison d'être* for the instrument and for our mutual efforts. The **core** or **sample stream** of cells would pass through the system in a direction perpendicular to the plane of the drawing, and the axes of the sample stream, the focused laser beam used for illumination, and the lens used to collect orthogonal scatter signals are all at right angles to one another, which is why the cytometer is described as having an **orthogonal** geometry. For the time being, we won't go into the details of how the cell gets into the center of the rectangular quartz cuvette in which the measurements are made.

Laser Beam Geometry and Illumination Optics

The beam coming out of the laser is radially symmetric, but the intensity varies with distance from the axis of the beam. If you plotted intensity versus distance from the axis, you'd come up with the familiar bell-shaped **Gaussian** or **normal** distribution.

It helps our cause to illuminate the cell and as little of the region surrounding it as possible. Most cells that are subjected to flow cytometry are less than 20 μm in diameter, so it would be advantageous to focus the illuminating beam to a spot not much bigger than this. This could be done using a single convex spherical lens. However, problems arise due to the Gaussian intensity profile of the laser beam and to the vagaries of fluid flow.

In order to measure scatter and fluorescence signals from cells with a precision of a few percent, it is necessary that illumination be uniform within that same few percent over the entire width of the sample or **core** stream. As long as the sample is flowing, we know that cells will get through the plane, defined by the intersection of the axes of the illuminating beam and the collection lens, in which the observation point lies. However, while, under ideal conditions, we'd like to have the cells strung out along the axis of flow like beads on a string, in practice, there's apt to be some variation in lateral position of cells in the core stream. If the beam is focused to a very small spot, the variation in intensity of illumination reaching cells at different positions will be too high to permit precise measurements.

Calculations show that if the diameter of the focused beam is about 100 μm, there will be only about 2% variation in intensity over the width of a 20 μm sample stream. There are, however, good reasons not to use a 100 μm round spot. If cells travel through the apparatus at velocities in the range of 2-5 m/sec, it will take 20-50 μs for a cell to traverse a 100 μm beam. During this time, most of the beam will be illuminating things other than the cell, and any scatter and fluorescence signals from these things will increase background noise levels.

Since variations in intensity over the Gaussian profile of the laser beam along the axis of fluid flow aren't a problem, because each cell goes through the whole beam, it makes sense to use a relatively small focal spot dimension in the direction of the axis of flow. A spot size of 20 μm allows cells to traverse the beam in 4-10 μs, increasing illumination of the cells during their **dwell time** in the beam and decreasing background as well. If the spot is made smaller than a cell diameter, say 5 μm, cells of different sizes spend different lengths of time in the beam – everybody isn't famous for the same number of microseconds – and **pulse width** can be used to measure cell size.

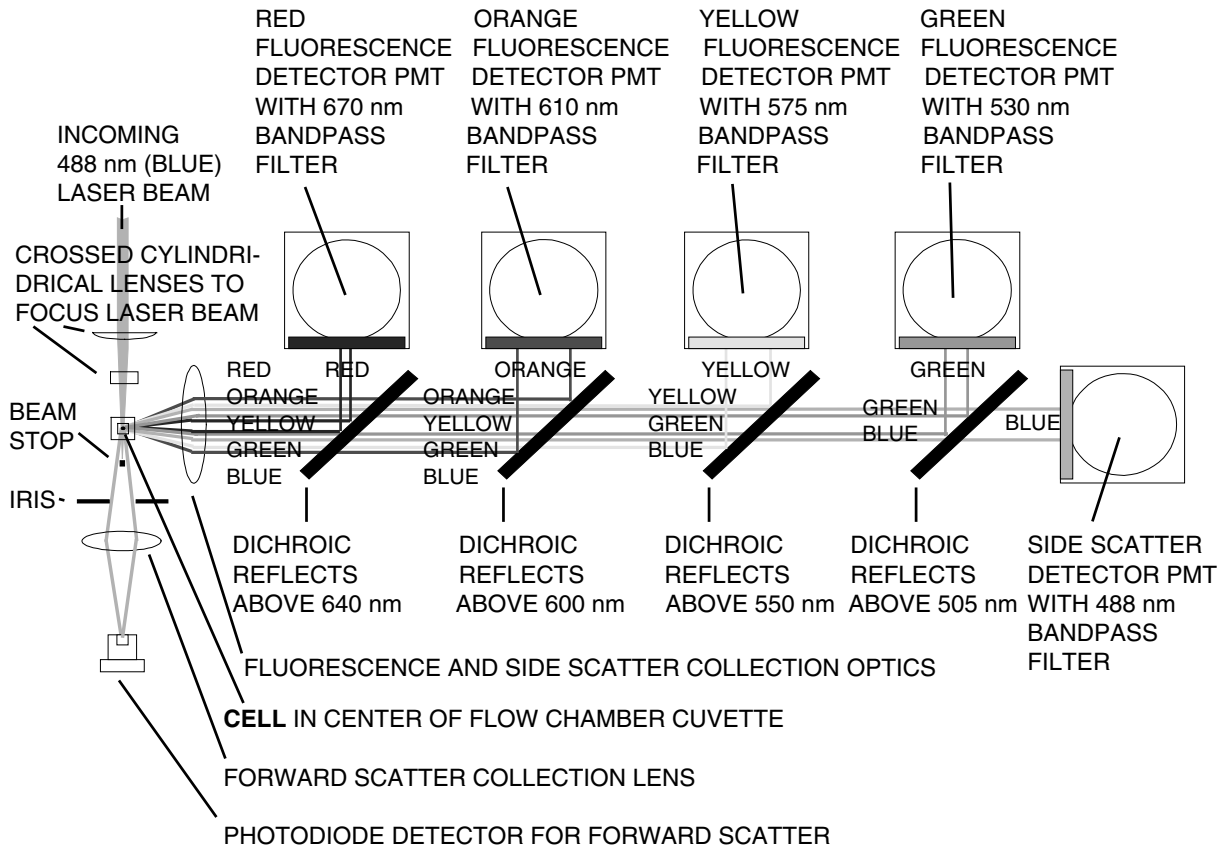

Figure 1-21. Schematic of the optical system of a fluorescence flow cytometer.

Using a really small spot, say 2 µm, you can extract a substantial amount of information about cell shape and structure by digitizing the signal at very high rates. Until recently, the processing electronics required for this technique, which is called **slit-scanning** flow cytometry, were too complex and expensive to be widely used, but the hardware and software are now more accessible should a compelling application come along. Current conventional instruments settle for **elliptical focal spots** 5-20 µm high and about 100 µm wide; these are obtained using **crossed cylindrical lenses** of different focal lengths, each of which focuses the beam in only one dimension. The crossed cylindrical lenses are shown at the left of Figure 1-21, above the cuvette. The lens closest to the cuvette is placed one focal length away from the sample stream, and focuses the beam in the dimension perpendicular to the plane, which is why you can't see the lens's curvature. The other lens, in this diagram, is placed so that its focal point is at the **beam stop**, which is a component of the forward scatter collection optics.

Flow Chamber and Forward Scatter Collection Optics

Earlier instruments examined cells in cylindrical quartz capillaries, or in a cylindrical **stream in air** following passage of fluid through a round orifice; the observation point in most cell sorters is still in a stream in air. However, in the benchtop instruments that are most widely used, observation is done in **flat-sided quartz cuvettes** with a square or rectangular cross section. The internal dimensions of the cuvettes are typically 100-200 by 200-400 µm; they are essentially small spectrophotometer cells and are, not surprisingly, produced for the flow cytometer manufacturers by the same companies that make spectrophotometer cells for other purposes. Cylindrical capillaries or streams in air themselves act like cylindrical lenses, and refract substantial amounts of light from the illuminating beam, which greatly increases the background noise level in scatter measurements and may also interfere with fluorescence measurements. Flat-sided cuvettes scatter relatively little of the incident light, minimizing such interferences.

The **beam stop** in the cytometer shown here is a vertical bar; we're looking at its cross section in the top view. What a beam stop needs to do is block the illuminating beam, once the beam has traversed the cuvette, so that as little of the beam as possible will reach the forward scatter detector and interfere with the measurement of light scattered by the cell at small angles to the beam. In an instrument in which observation is done in a round capillary or in a stream in air, the beam stop has to be horizontal, to block light refracted by the capillary or stream; the forward scattered light that is

detected is light scattered "up and down," i.e., out of the plane of Figure 1-21. The laws of physics that govern focused laser beams end up dictating that we can collect light scattered at smaller angles using a flat-sided cuvette and a vertical beam stop than we can using a capillary or round stream and a horizontal beam stop.

The actual range of angles over which small-angle or forward scatter signals are collected varies considerably from instrument to instrument. The lower end of the range is set by the placement and dimensions of the beam stop; in many flow cytometers, the upper end of the range is adjustable by manipulating an **iris diaphragm**, shown below the beam stop at the left of the figure. The light that gets around the beam stop and through the diaphragm is converged by the **forward scatter collection lens**, which, in the apparatus shown in the figure, is bringing the light to a focus at the **forward scatter detector**.

The detector illustrated here is a **photodiode**, a silicon solid-state device that takes photons in and puts electrons out, usually at the rate of about 5 electrons out for every 10 photons in, giving it a quantum efficiency of 50 percent. The actual sensing area of the detector is in the neighborhood of 1 mm^2. When you make the same kind of silicon chip with a larger surface area, you can get some fairly serious electric currents out of the resulting **solar cell**. The photodiodes used as forward scatter detectors in most flow cytometers typically have output currents of a few microamperes, not because they're smaller than solar cells, but because there aren't enough photons, even in the relatively strong forward scatter signal, to produce higher currents. When you're trying to measure forward scatter signals from relatively small particles, e.g., bacteria, a photodiode may not be up to the job, and it may be better to use a more sensitive detector, such as a **photomultiplier tube (PMT)**. These are used for side scatter and fluorescence detection, but are larger, more complicated, and - probably most important from the commercial point of view - more expensive than photodiodes. In an ideal world, the flow cytometer manufacturers would offer a high-sensitivity PMT forward scatter detector option on all models; turn on the news if you still think ours is an ideal world.

Fluorescence and Side Scatter Optics

The really hairy part of Figure 1-21, and of the average flow cytometer, is the part that deals with the collection of fluorescence and side scatter signals and the diversion of light in different spectral regions to the appropriate photomultiplier tube detectors. The first task is to collect the light. I have shown a single, simple **collection lens** for fluorescence and side scatter, but the optics actually used are somewhat more complicated.

As was noted in Figure 1-1, light is scattered, and fluorescence emitted, in all directions, i.e., over a solid angle corresponding to the entire surface of a sphere. In principle, we'd like the lens to collect light over as large a solid angle as possible, so we can collect as much of the fluorescence as possible. One way to do this is to use a high-N.A. microscope lens to collect the light; this is done in many instruments, some of which even use a functional equivalent of oil immersion to get the highest possible N.A. Another is to place the collection lens at its focal distance from the sample stream. Various experimenters have used parabolic or ellipsoidal reflectors and high-N.A. fiber optics for light collection in attempts to increase the total amount of light collected.

As has already been suggested in the discussion of forward scatter detectors, ideal solutions are hard to come by. Every decision made in the design of a flow cytometer involves tradeoffs. In the case of light collection optics, the problem we run into is usually that, as we collect more light, we have less control over where we collect it from. What we really need to do is collect as much light from the immediate region of the cell, and as little from elsewhere, as possible, because any light we collect from elsewhere will only contribute to the background or noise. Thus, the all-important **signal-to-noise ratio** will decrease, even though the signal itself increases. Flow cytometer designs using ellipsoidal or parabolic reflectors or fiber optics for light collection have, so far, run into this problem.

The simple collection lens shown in the figure is illustrated as producing a **collimated beam** of light, i.e., one in which rays entering the lens at all angles come out parallel, with a so-called "focus at infinity." In most real flow cytometers, the light collected from the collection lens is either not collimated or is converged by a second lens, and then passes through a small aperture, or **field stop** (see p. 9), that lets most of the light collected from the region near the cell through and blocks most of the light collected from elsewhere. Some instruments incorporate an additional lens behind the field stop to recollimate the collected light, because there is some advantage in presenting a collimated beam to the **dichroics** and **optical filters** used direct light collected at different wavelengths to different detectors.

Optical Filters for Spectral Separation

The lens that collects the fluorescence emitted from, and the light scattered at large angles by, cells transmits light encompassing a range of wavelengths. Most of the light is scattered laser light, at 488 nm; much of the rest should be fluorescence from the cells, which will of necessity be at wavelengths above 488 nm. The choice of wavelength regions for fluorescence measurements is based on the fluorescence emission spectral characteristics of the available fluorescent probes or labels that can be excited at 488 nm.

The apparatus illustrated in Figure 1-21 is designed to detect fluorescence in four spectral regions, which we call **green (515-545 nm)**, **yellow (560-590 nm)**, **orange (600-620 nm)**, and **red (660-680 nm)**. It also detects scattered light at the excitation wavelength, 488 nm. Each of the detectors is a photomultiplier tube, and all of the detectors are fitted with **bandpass optical filters** that transmit light in the appropriate wavelength ranges.

There are basically two kinds of optical filters that can be used for wavelength selection; they are **color glass**, or **absorptive**, filters and **dielectric**, or **interference**, filters. Color glass filters are made of glass or plastic impregnated with **dyes** that absorb light in the unwanted wavelength regions and transmit most of the light in the desired regions. Dielectric filters are made by depositing thin layers of dielectric materials on a glass or quartz substrate; within some wavelength range, which is determined by the thickness of these layers, there will be **destructive interference**, resulting in light of these wavelengths being reflected from, rather than transmitted through, the filter.

Filters can be made with several kinds of transmission characteristics. There are **edge** filters, which may be either **long pass** or **short pass** types; long pass filters block shorter and transmit longer wavelengths and short pass filters block longer and transmit shorter wavelengths. Long pass and short pass filters are usually specified by the wavelength at which their transmission is either 50% of the incident light or 50% of their maximum transmission. There are **bandpass** filters, which block wavelengths above and below the desired region of transmission; they are specified by the wavelength of maximum transmission and by the **bandwidth**, which defines the range of transmission, usually expressed as the range between the points below and above the peak at which transmission is 50% of maximum. There are also **notch** filters, which are designed to exclude a narrow range of wavelengths.

Absorptive filters can be very effective at getting rid of light outside their desired **passbands**, i.e., those regions in which they transmit light (many transmit less than 0.01% outside the passband), and can also be made to have good (>90%) light transmission in the passband. However, the dyes incorporated into the filter to absorb the unwanted light may fluoresce; this phenomenon can (and did, in the earlier fluorescence flow cytometers) interfere with the detection of weak fluorescence signals from cells. As a result, most modern instruments now use interference filters, which reflect rather than absorb unwanted light.

Real interference filters used as long pass or bandpass filters frequently incorporate an absorptive layer behind the dielectric layers to get rid of the last little bit of unwanted light, because it's difficult to get rid of more than 99% of it by interference and reflection alone. Fluorescence in these filters is not a big problem because the interference layers get rid of most of the light that might excite fluorescence before it hits the absorptive layer – provided, that is, that you mount the filter **shiny side out**, that is, with the interference layers facing where the light's coming from and the colored absorptive side facing where it's going.

Dichroics, also called **dichroic mirrors** or **dichroic beamsplitters**, are interference filters, usually without an added absorptive layer. They can be made with either **long reflect** (i.e., short pass) or **short reflect** (i.e., long pass) characteristics, and both kinds are used in flow cytometers. As is the case with other types of interference filters, it's easier to make a filter that reflects 97% of unwanted light than it is to make one that transmits 90% of wanted light. When flow cytometers measured fluorescence in only two spectral regions, they only needed one dichroic (maybe two, if you count one to reflect blue (488 nm) light to the orthogonal scatter detector and keep it away from the fluorescence collection optics). When you start measuring fluorescence in three or four regions, it becomes advisable to do careful calculations to make sure you don't lose a lot of the light you want in the dichroics. The Devil, as we all know, is in the details, and more deals with the Devil are made in the details of dichroics and filters than in most other areas of flow cytometer design.

The layout shown in Figure 1-21 assumes that the strongest signal, or the one with the most light we can waste, is the blue orthogonal scatter signal, and that the green, yellow, orange, and red fluorescence signals are progressively weaker. Even if all of the dichroics transmit 90% of the incident blue light, only 65% of the light coming through the collection lens will reach the filter in front of the orthogonal scatter detector PMT. About 70% of the green fluorescence will make it to the filter in front of the green detector PMT, while 77% of the yellow, 86% of the orange, and 96% of the red fluorescence will get to the filters in front of the detectors for those spectral regions. We therefore lose the least light from the weakest signal.

There are other ways to improve light transmission; one is to ditch the in-line arrangement of PMTs shown in the figure, instead first splitting the red/orange and the blue/green/yellow regions, so that the green fluorescence signal passes through two dichroics and the others through only a single dichroic. Another, which I routinely use in the "Cytomutt" flow cytometers I build, is to place a second fluorescence collection lens at 180° from the first one, so that each lens collects light for at most three detectors.

The spacing between the dielectric layers of interference filters and dichroics determines the wavelengths at which interference will occur, and, therefore, the wavelengths that will be transmitted or reflected by these components. The distance between the layers changes with the angle at which light hits the filter (remember trigonometry?), and, as a result, the passband of the filter changes with the angle of incidence of the light. In theory, light should be collimated before it gets to the dichroics and filters; this is generally not done because the light coming from the collection lens is contained within a fairly small solid angle. Problems with dichroics and filters are more likely to result from using the wrong filters or from mounting filters incorrectly. Dielectric filters also degrade over time, as moisture gets in between the dielectric layers, but, when this occurs, the filters tend to look ugly enough so that you'd think about ordering new ones.

I hope, by now, to have conveyed the impression that dichroics and filters are among the most critical parts of a flow cytometer; not surprisingly, the right – or wrong – selection of dichroics and filters can also make a big difference when

you're doing fluorescence microscopy, by eye or with image analyzers, etc. A few hundred dollars spent on good filters may dissuade you from smashing tens or hundreds of thousands of dollars worth of instrument to smithereens out of frustration.

Multistation Flow Cytometers

Before going on to a discussion of detectors and electronics, I will point out that, whereas most flow cytometers have a single excitation beam, and you can have any color you want as long as it's 488 nm, there are systems available that offer a wider choice of excitation wavelengths. Some of these can use two or more illumination beams, separated by a small distance in space. A good way to conceptualize such a **multistation flow cytometer** might be to imagine two or more copies of Figure 1-21 stacked one on top of another. Because the beams in a multistation instrument are separated by a short distance, it takes a short time for cells to travel from one beam to another, and the signals are therefore separated in time. Since the velocity of cells through the system is approximately constant, the time interval between signals from different beams is also approximately constant.

In flow cytometers that form an image of the sample stream, as most now do, it is customary to form separate images of the intersections of two or more beams with the sample stream, and divert light from each observation point to the appropriate detectors. In instruments in which no image is formed, and in which light from multiple observation points reaches all the detectors, a **time-gated amplifier** is used. This allows signals from the detectors that measure events at the downstream observation point to reach the signal processing electronics only at a set time interval after signals are detected at the upstream observation point.

Multistation instruments have also been built that incorporate electronic volume sensors as well as laser or arc lamp illumination; cell sorters are also multistation instruments, as are cell "zappers" or **photodamage cell sorters**. These use a high energy pulsed laser beam downstream from the measurement beam and switch the beam on to destroy cells with selected characteristics.

Flow cytometers with multiple illumination beams are used primarily for multiparameter measurements involving probes that cannot be excited at the same wavelength. For example, sorting human chromosomes stained with combinations of dyes that preferentially stain A-T and G-C rich regions of DNA requires separated ultraviolet (325-363 nm) and blue-violet (436-457 nm) illuminating beams. Other applications use ultraviolet and 488 nm beams and 488 and red (633 or 635 nm) beams; as many as five beams have been used in a single apparatus. The current trend is toward multiple illumination beams, even in benchtop instruments.

Photomultipliers and Detector Electronics

A **photomultiplier tube (PMT)**, like a photodiode, takes in photons and puts out electrons. However, whereas a plain photodiode never does much better than 7 electrons out for every 10 photons in, a PMT may get as many as a few hundred thousand electrons out for each photon that reaches its **photocathode**. PMTs, like cathode ray television tubes and the tubes favored by audiophiles and rock musicians who can't see the trees for DeForest, are among the last survivors of the vacuum tube era. They incorporate a photocathode, which is placed behind a glass or quartz window so light can reach it, a series of intermediate electrodes, or **dynodes**, and another electrode called the **anode**. A voltage is applied to each electrode; the photocathode is at the lowest voltage, with each dynode at a successively more positive voltage and the anode at the most positive voltage of them all – which is usually ground, because the photocathode is generally a few hundred to a couple of thousand volts negative.

Photons hitting the photocathode result in **photoelectrons** being emitted from the photocathode, and accelerated toward the first dynode by the electric field resulting from the difference in electric potential (voltage) between these electrodes. The electrons acquire energy during this trip, so, when they whack into the dynode, they dislodge more electrons from it, which are accelerated toward the next dynode, and so on. The bigger the difference in potential, i.e., applied voltage, between stages, the more energy is imparted to the electrons at each stage, and the more electrons are released from the receiving electrode. This gives the PMT a mechanism for **current gain** that is relatively noise-free. The PMTs used in most flow cytometers have current gains as high as 10^6. However, the **quantum efficiency** of PMT photocathodes is typically lower than that of photodiodes, with peak values of 25% (i.e. 25 electrons out for 100 photons in) in the blue spectral region, and, usually, much lower values in the red. Detector quantum efficiency is important because the sensitivity and precision with which fluorescence (or any other optical signal) can be measured ultimately depend on the number of electrons emitted from the detector photocathode.

Why is it that at detectors, we measure success one electron at a time? Because detection is subject to the same Poisson statistics we ran into on p. 19. When you count (or detect) n of anything, including photoelectrons, there is an associated standard deviation of $n^{1/2}$. When you detect 10,000 photoelectrons, the standard deviation is $10{,}000^{1/2}$, or 100, and the coefficient of variation (CV) is 100 × (100/10,000), or 1%. When you detect 10 photoelectrons, the standard deviation is $10^{1/2}$, or about 3.16, and the CV is 100 × (3.16/10), or 31.6%. I am talking about photoelectrons, rather than photons, here, because, while the detector, whether diode or PMT, "sees" photons, if you will, all the electronics lets us "see" is electrons.

If we had reliable low-noise amplifiers with gains of several million, we'd always be better off with the 50-70 electrons we could get out of the photodiode for every 100 photons hitting it than we would with the 8-25 electrons emitted from the PMT cathode under the same conditions; all the gain in the PMT doesn't get around the imprecision

introduced by the lower number of electrons it starts with and, in fact, there is also a statistical aspect to the PMT's gain mechanism.

Unfortunately, the high-gain, low noise amplifiers we'd need to use photodiodes as sensitive fluorescence detectors don't exist. There are, however, solid-state devices called **avalanche photodiodes (APDs)**, which combine high quantum efficiency with a mechanism that can produce gains as high as a few thousand when a voltage is applied across the diode. While APDs are now used for both scatter and fluorescence detection in some commercial flow cytometers, they do not match the sensitivity of PMTs.

The photodetectors we have been talking about are sources of **electric current**. A **preamplifier**, which is the first stage in the **analog signal processing electronics**, converts the current output from its associated detector to a voltage. The preamplifier also accomplishes the important task of **DC baseline restoration**.

An ideal flow cytometer is something like an ideal dark field microscope; when there's no cell in the observation region, the detector shouldn't be collecting any light at all. In practice, there's always some small amount of light coming in. In the case of the scatter detectors, most of this light is stray scattered light from the illuminating beam; in the case of the fluorescence detectors, the light background may come from fluorescence excited in various optical elements such as the flow chamber, lenses, and filters, from fluorescence due to the presence of fluorescent materials in the medium in which cells are flowing, and from **Raman scattering**, which produces light at frequencies corresponding to the difference between the illumination frequency and the frequencies at which absorption changes molecular vibrational states. In flow cytometry, the major interference due to Raman scattering results from scattering by water; when 488 nm illumination is used, this scattering occurs at about 590 nm, and may interfere with detection of signals from probes labeled with phycoerythrin, which fluoresces near this wavelength.

The net result of the presence of all of the abovementioned stray light sources is that there are some photons reaching the detectors in a flow cytometer even when there isn't a cell at the observation point, producing some current at the detector outputs. There may also be some contribution from the so-called **dark current** of the detector, which results from the occasional electron breaking loose from the cathode due to thermal agitation. There are some situations in which performance of photodetectors is improved by refrigerating them to reduce dark current; flow cytometry in the contexts we're discussing isn't one of them. Even with the detectors in liquid nitrogen, we'd have to deal with the background light, which will contribute a signal with an average value above zero to whatever signal we collect from the cells.

The background signal can be considered as the sum of a constant **direct current (DC)** component and a variable **alternating current (AC)** component, representing the fluctuations due to photon statistics and to other sources of variation in the amount of stray light reaching the detector. One important source of such variation may be **light source noise**, i.e., fluctuations in the light output of the laser or lamp used for illumination; in some circumstances, particularly scatter measurements of small particles, source noise can be the major factor limiting sensitivity.

What we'd like to measure when a cell does pass by the observation station is the amount of light coming from the cell, not this amount plus the background light. We can do this, to a first approximation, by incorporating an electronic circuit that monitors the output of the detector and uses negative feedback to subtract the slowly varying component of the output from the input, thereby eliminating most of the DC background signal, and restoring the **baseline** value of the preamplifier output to ground.

In practice, baseline restorers will keep their voltage outputs within a few millivolts of ground when no cells are coming by. When a cell does arrive, it will scatter and probably emit small amounts of light, which will be collected and routed to the various detectors, producing transient increases, or **pulses**, in their output currents, which will result in voltage **pulses** at the preamplifier outputs. At this point, as was noted on p. 17, all of the information we wanted to get from the cell resides in the heights, areas, widths, and shapes of those pulses; we will ultimately convert these to digital values, in which form they can be dealt with by the computers that are almost universally used for data analysis in flow cytometry. However, before we get into the details of how pulse information is processed, we ought to consider the only element of Figure 1-21 that has been neglected to this point, namely, the cell flowing through the apparatus, and how it gets there.

Putting the Flow in Flow Cytometry

Figure 1-21 describes the cell as being in the center of the cuvette, and I have already talked about a **core** or sample stream of cells that is about 20 μm wide, while mentioning that the internal dimensions of the cuvette are on the order of 200 by 200 μm. The space between the core and the inner walls of the cuvette is occupied by another stream of flowing fluid, called the **sheath**. How the core and sheath get where they are can be appreciated from a look at Figure 1-22.

Fluid mechanics tells us that, if one smoothly flowing stream of fluid (i.e., the core stream) is injected into the center of another smoothly flowing stream of fluid (i.e., the sheath stream), the two streams will maintain their relative positions and not mix much, a condition called **laminar flow**. There are generally differences in fluid flow velocity from the inside to the outside of the combined stream, but the transitions are even. If the velocities of the two streams are initially the same, and the cross-sectional area of the vessel in which they are flowing is reduced, the cross-sectional areas of both streams will, obviously, be reduced, but they will maintain the same ratio of cross-sectional areas they had

at the injection point. If the sheath stream is flowing faster than the core stream at the injection point, the sheath stream will impinge on the core stream, reducing its cross-sectional area. In the flow chamber of a flow cytometer, both mechanisms of constricting the diameter of the core stream may be operative.

The core stream, which contains the cell sample, is injected into the flowing water or saline sheath stream at the top of a conical tapered region that, in the flow chamber shown in the figure, is ground into the cuvette. The areas of both streams are reduced as they flow through the tapered region and enter the flat-sided region in which cells are observed. Core and sheath streams may be driven either by gas pressure (air or nitrogen), by vacuum, or by pumps; most instruments use air pressure. **Constant volume pumps**, e.g., **syringe pumps**, which, if properly designed, deliver a known volume of sample per unit time with minimum pulsation, provide finer control over the sample flow rate. Since knowing the sample flow rate makes it easy to derive counts of cells per unit volume, flow cytometric hematology analyzers incorporate constant volume pumps; why fluorescence flow cytometers, in some cases made by the same manufacturers, do not remains something of a mystery.

The overall velocity of flow through the chamber is generally determined by the pressure or pump setting used to drive the sheath. If the sheath flow rate is increased with no change in the core flow rate, the core diameter becomes smaller and the cells move faster; if the sheath flow rate is decreased under the same circumstances, the core diameter becomes larger and the cells move more slowly. In some circumstances, it is desirable to adjust sheath flow rates; if cells move more slowly, they spend more time in the illuminating beam, receive proportionally more illumination, and they therefore scatter and emit proportionally more light. If the amount of light being collected from cells is the limiting factor determining sensitivity, slowing the flow rate can improve sensitivity, allowing weaker signals to be measured.

This aside, it is generally preferable to be able to control the core diameter, and therefore the volume of sample and number of cells analyzed per unit time, without changing the velocity at which cells flow through the system. This is done by leaving the sheath flow rate constant and changing the driving pressure or pump speed for the core fluid. More drive for the core results in a larger core diameter; more cells can be analyzed in a given time, but precision is likely to be decreased because the illumination from a Gaussian beam is less uniform over a larger diameter core. Less drive for the core gives a smaller core diameter and a slower analysis rate, but precision is typically higher. When the cytometer is being used to measure DNA content, precision is important; when it is being used for immunofluorescence measurement, precision is usually of much less concern.

The use of **sheath flow** as just described has proven essential in making flow cytometry practical. Without sheath flow, the only way of confining 10 μm cells within a 20 μm diameter stream would be to observe them in a 20 μm diameter capillary or in a stream in air produced by ejecting the cells through a 20 μm diameter orifice. This would very quickly run afoul of Shapiro's First Law (p. 11). As a matter of fact, even with sheath flow, Shapiro's First Law frequently came into play when cell sorters were typically equipped with 50 μm orifices. That orifice size was fine for analyzing and sorting carefully prepared mouse lymphocytes, but people interested in analyzing things like disaggregated solid tumors might encounter mean intervals between clogs of two minutes or so. With the larger cross-sectional areas of the flow chambers now used in most flow cytometers, clogs are not nearly the problem they once were.

Clogs, however, are not the only things that can disturb the laminar flow pattern in the flow chamber. Air bubbles perturb flow, as do objects stuck inside the chamber but not large enough to completely obstruct it. In the first commercial cell sorters, the standard method for getting rid of air bubbles was to remove the chamber from its mount while the apparatus was running, and turn it upside down; the bubble would rise to the top and emerge from the nozzle along with a stream of sheath and sample fluid that would spray all over the lab. This technique became inappropriate with the emergence of AIDS in the 1980's. Now, even droplet sorters incorporate an air outlet (which I have referred to elsewhere as a "burp line") for getting rid of bubbles. In some flow cytometers with closed fluidic systems, the air bubble problem is minimized by having the sample flow in

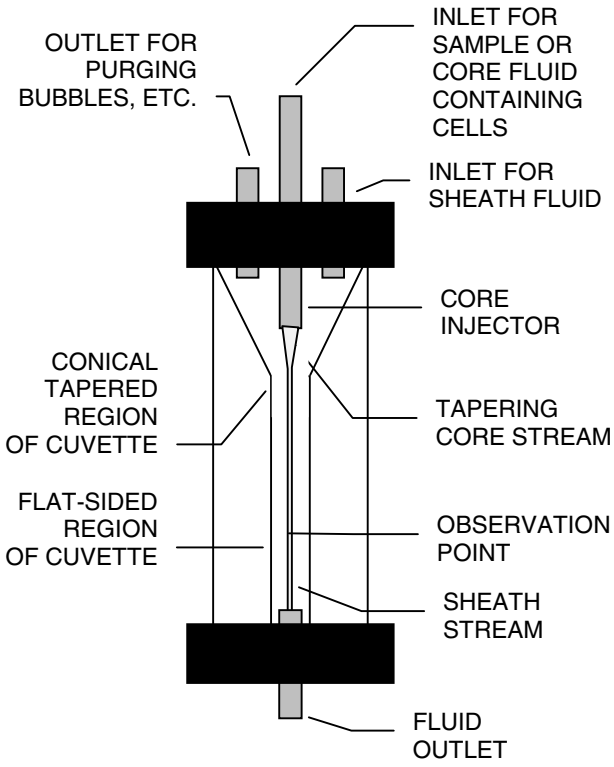

Figure I-22. A typical flow chamber design.

at the bottom and out at the top, essentially turning Figure 1-22 upside down; bubbles are more or less naturally carried out of the flow chamber.

Disturbances in laminar flow, whether due to bubbles or junk, often result in the core stream deviating from its central position in the flow chamber and in differences in velocity between different cells at different points within the core. Turbulent fluid flow is now described mathematically using chaos theory; you can recognize turbulent flow in the flow chamber by the chaos in your data.

For the present, we will assume the flow is laminar, the optics are aligned, and the preamplifiers are putting out pulses with their baselines restored, and consider the next step along the way toward getting results you can put into prestigious journals and/or successful grant applications.

Signal Processing Electronics

We have already mentioned that a cell is going to pass through the focused illuminating beam in a flow cytometer in something under 10 µs, during which time the detectors will produce brief current pulses, which will be converted into voltage pulses by the preamplifiers. Using analog peak detectors, integrators, and/or pulse width measurement circuits, followed by analog-to-digital conversion, or, alternatively, rapid A-D conversion followed by digital pulse processing (p. 21), we will reduce pulse height, area, and width to numbers, at least some of which will, in turn, be proportional to the amounts of material in or on the cell that are scattering or emitting light. But which numbers?

First, let's tackle the case in which the focal spot, in its shorter dimension, along the axis of flow, is larger than the cell, meaning that there is some time during the cell's transit through the beam at which the whole cell is in the beam. Because the beam is Gaussian, the whole cell may not be uniformly illuminated at any given time, but intuition tells us that when the center of the cell goes through the center of the beam, we should be getting the most light to the cell and the most light out of it. The preamplifier output signal, after baseline restoration, is going to be roughly at ground before the cell starts on its way through the beam, and rise as the cell passes through, reaching its **peak value** or **height** when the center of the cell is in the center of the beam, and then decreasing as the cell makes its way out of the beam. Since the whole cell is in the beam when the pulse reaches its peak value, this value should be proportional to the total amount of scattering or fluorescent material in or on the cell.

Things get a little more complicated when the beam is the size of the cell, or smaller. In essence, different pieces of the cell are illuminated at different times as the cell travels through the beam. In order to come up with a value representing the signal for the whole cell, we have to take the **area**, or **integral**, rather than the height of the pulse. There are two ways to do this with analog electronics. One is to change the frequency response characteristics of the preamplifier, slowing it down so that it behaves as an integrator, in the sense that the height of the pulse coming out of the slowed-down preamplifier is proportional to the area or integral of the pulse that would come out of the original fast preamplifier. Putting the slowed pulse into a peak detector then gives us an output proportional to the area or integral we're trying to measure. Alternatively, we can keep the fast preamplifier, and feed its output into an analog **integrator** instead of a peak detector.

If we decide to do digital pulse processing, we have to digitize the pulse trains from the preamplifier outputs rapidly enough so that we have multiple samples or "slices" of each pulse. We can then add the values of a number of slices from the middle of the pulse to get an approximation of the area, or integral; eight slices will do, but sixteen are better. This works pretty well. However, if we're only taking eight or sixteen slices of a pulse, we may not get as accurate a peak value or a pulse width value as we could using analog electronics.

The peak value we get from digital processing is simply the largest of our eight or sixteen slices. These provide us with only a fairly crude connect-the-dots "cartoon" of the pulse, thus, while there is a substantial likelihood that the largest digitized slice is near the peak value, there is a relatively low probability that the digitization will occur exactly when the peak value is reached.

Similarly, if we estimate pulse width from the number of contiguous slices above a set threshold value, we will have a fairly coarse measurement; if the digitization rate gives us at most sixteen slices, our range of pulse widths runs from 1 to 16, with each increment representing at least a 6 percent change over the previous value. If we had fast enough analog-to-digital converters to be able to take a few hundred slices of each pulse, and fast enough DSP chips to process the data, we could get rid of analog peak detectors and pulse width measurement circuits, but we're not there yet. The digital integrals are already good enough to have been incorporated into commercial instruments.

Is It Bigger than a Breadbox?

I have been referring to benchtop flow cytometers and big sorters, but I haven't shown you any pictures. Now's the time to fix that.

Figure 1-23, on the next page, shows the Becton-Dickinson FACScan, the first really successful benchtop flow cytometer, introduced in the mid-1980's. It uses a single 488 nm illuminating beam from an air-cooled argon ion laser, and measures forward and side scatter and fluorescence at 530 and 585 and above 650 nm. The data analysis system is an Apple Macintosh personal computer, shown in front of the operator.

Figure 1-24 (courtesy of Cytomation) shows that company's MoFlo high-speed sorter. The optical components, including two water-cooled ion lasers and a large air-cooled helium-neon laser, are on an optical table in front of the operator. Most of the processing electronics are in the rack to the operator's left; the two monitors to her right display data from an Intel/Microsoft type personal computer.

Figure 1-23. FACScan Analyzer (Becton-Dickinson)

Figure 1-24. MoFlo High-Speed Sorter (Cytomation)

Figure 1-25. Microcyte Cytometer (Optoflow)

Neither the FACScan nor the MoFlo risks being mistaken for a breadbox. However, the Microcyte analyzer shown in Figure 1-25 (photo courtesy of Optoflow AS) comes close. It is a two-parameter instrument with a red diode laser source, and measures medium angle scatter using a photodiode and fluorescence using an avalanche photodiode. As you might guess, it can be run on batteries.

Flow Cytometer Pathology and Diagnostics

As the benchtop flow cytometer starts to look more and more like a "black box" (okay, a "beige box," "gray box," or whatever from some manufacturers), it becomes increasingly important for a user to know how to verify that the instrument is running properly. It is, of course, equally important to know when a big sorter is and is not running properly, but the larger instruments tend to make their operators aware of problems.

An instrument in proper alignment, running particles through an unobstructed flow system at a rate within the manufacturer's specifications, should get nearly identical measurements from nearly identical particles. There are now several companies producing nearly identical particles in the form of plastic microspheres, i.e., beads, impregnated with fluorescent dyes. If everything's right, one ought to be able to make scatter and fluorescence measurements of such particles with high precision, meaning coefficients of variation no higher than a few percent. The only biological objects that are likely to yield CVs in that range are noncycling cells, such as peripheral blood lymphocytes, stained with a fluorescent DNA stain; most people stick with beads.

An instrument in which optical alignment is adjustable by the operator will typically yield the lowest measurement CVs at the point at which signal amplitudes are maximized. However, optical misalignment is not the only potential cause of poor measurement precision. Fluctuations in the power output of the light source will decrease precision, as will the presence of cell aggregates, large pieces of debris, and/or gas bubbles in the flowing stream. These create turbulence, resulting in the measured particles being distributed over an excessively large portion of the stream and/or traveling at different velocities; under these conditions, nearly identical particles will obviously not produce nearly identical signals.

Sensitivity, which, in the context of flow cytometry, basically means the degree to which fluorescence distributions from dimly stained cells (or beads) can be discriminated from distributions from unstained (control) cells or (blank) beads will usually be degraded if precision falls substantially short of the mark. Loss of sensitivity may also be due to degradation and/or incorrect choice or installation of optical filters.

Precision of instruments should always be determined using beads carrying fairly large amounts of dye, to minimize the contribution to variance from photoelectron statistics. Determination of instrument sensitivity virtually demands that at least some of the test objects used produce low-

intensity signals. Beads used for sensitivity testing typically come in sets containing an undyed or blank bead and beads loaded with four or more different levels of fluorescent dye.

Flow cytometer manufacturers and third parties also supply beads that can be used to optimize fluorescence compensation settings, and, as was previously noted in the discussion of quantitative fluorescence measurements, beads that allow the scale of the instrument to be calibrated in terms of numbers of molecules of a particular probe or label.

1.6 ALTERNATIVES TO FLOW CYTOMETRY; CYTOMETER ECOLOGY

In order to use flow cytometry to study characteristics of **intact cells** from solid tissues or tumors, or of cultured cells that grow attached to one another and/or to a solid substrate, various methods are used to prepare **single cell suspensions** from the starting material. Flow cytometry itself can provide a good indication of the efficacy of such preparative procedures. In a similar fashion, the technique can be very useful in **monitoring bulk methods** for purifying cell subpopulations, e.g., **sedimentation** and **centrifugation** techniques and **affinity-based separations**. If large cell yields are more important than high purity, bulk separation with flow cytometric monitoring may be preferable to sorting as a preparative method.

We have learned and can probably continue to learn a great deal by dissociating tissues and even organisms into suspensions of intact cells that can be characterized in flow cytometers, sorted, and subsequently studied in culture. However, the procedures used for cell dissociation, by nature, have to remove most of what holds the cells together. Since such **adhesion molecules** are probably as important as anything else for our understanding of cells' behavior, it is inevitable that there will come a point at which we won't be able to answer critical questions using cells stripped of these essential components.

It will make sense, at that point, to find instrumental alternatives to flow cytometry in a new generation of image analyzers and scanning cytometers, designed with an emphasis on preserving cell viability, which allow us to use the armamentarium of analytical techniques and reagents, in the development of which flow cytometry has played a major role, to study cells in organized groups.

We may also, of course, run up against the limits of flow cytometry simply by developing a desire to measure something repeatedly in one cell over an interval greater than a few hundred microseconds. This can be accomplished by combining static cytometry with **kinetic analysis techniques**, such as **flow injection analysis**, adapted from analytical chemistry.

I am reminded that one of the Mayo brothers said that a good surgeon had to know when to stop cutting and when not to cut; a good analytical cytologist will have to know when to put aside flow cytometry. Not now, though. Keep reading.

When we do consider the alternatives to flow cytometry, and even the availability of different types of flow cytometers, we run into something of an information gap. In the last edition of this book, I described flow cytometry as having been a growth industry since about 1985, based on census data compiled by Kit Snow of Beckman Coulter Corporation and shown in Figure 1-26.

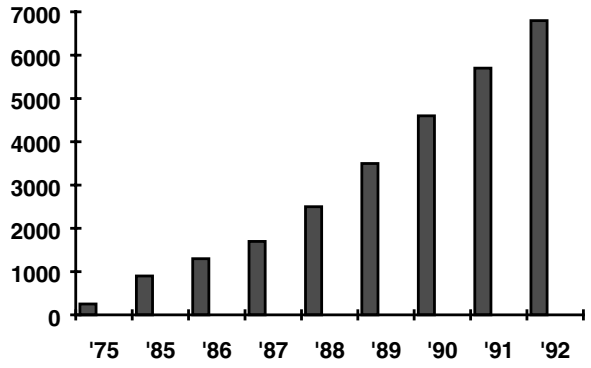

Figure 1-26. Estimated numbers of fluorescence flow cytometers in use worldwide, 1975-1992.

I tried to get updates on these numbers from various manufacturers, and even expanded my search to look for data about scanning laser cytometers, confocal microscopes, etc. Nobody's talking. The best I could do was come up with numbers that nobody would say were way too high or way too low. So here goes.

The great majority of fluorescence flow cytometers now in use are benchtop models similar to the one shown in Figure 1-23; they use low-power, air-cooled argon ion laser light sources operating at a fixed emission wavelength of 488 nm, and measure forward and orthogonal light scattering and fluorescence in three or four (green, yellow and/or orange, and red) spectral regions. Most of these systems have been designed for ease of use, with the needs of the clinical laboratory market foremost in mind. Newer instruments in the same class have added features such as a second (red) laser and closed fluidic sorting systems. The estimate is that there are somewhere between 12,000 and 20,000 such flow cytometers in use worldwide.

There are also probably around 2,000 larger, more elaborate fluorescence flow cytometers, which may use one or more air-cooled or water-cooled laser sources, can be equipped to measure eight or more parameters, and offer droplet sorting capability. These instruments are typically used in research laboratories rather than in clinical settings. Then, there are several hundred commercially produced fluorescence flow cytometers using arc lamp rather than laser sources, at least an equal number of instruments designed for multiplexed assays on beads, and one to two hundred laboratory-built flow cytometers.

The confocal microscopy folks seemed happy with the estimate that there are between 3,000 and 5,000 confocal systems worldwide; only two or three hundred of these are equipped for multiphoton excitation.

The area of relatively low-resolution scanning laser cytometry[2380] has gotten more active in recent years. The CompuCyte Laser Scanning Cytometer (LSC), developed by Lou Kamentsky[2047], is generating an increasing number of interesting publications[2381-2], about which I will say more later. I estimate that there are 100 to 250 LSCs now in circulation, and probably a similar number of volumetric capillary cytometers[1365], built by Biometric Imaging, now part of B-D. And there is at least one promising scanning system that hasn't yet made it into production but is worth watching[2383].

1.7 THE REST OF THE BOOK

In Chapter 2, I will point you toward some sources of information that may be of use to you in learning more of the details of cytometry, flow and otherwise, discussion of which began in this Chapter and will continue in Chapter 4. Chapter 2 will also provide brief descriptions of a bunch of books on cytometry and related topics that have appeared since the last edition of this tome.

I have devoted the intermediate Chapter 3 to the history of flow cytometry, because I think that an appreciation of how things came to be as they are is as important to further progress as is an understanding of the science and technology. Chapters 5 and 6, respectively, provide additional material on data analysis and flow sorting.

Parameters and the probes used for their measurement are discussed in Chapter 7, which also presents some basic applications of flow cytometry and of some alternative methods. Chapter 8 considers flow cytometers, software, and related accessories now available from commercial manufacturers, and criteria that may influence buying decisions.

Chapter 9 briefly discusses the option of building flow cytometers; although the details on the construction, care and feeding of "Cytomutts" featured in the earlier editions have been omitted, some material that may help users understand their apparatus better has been retained.

Current and proposed applications of, and alternatives to, flow cytometry in biomedical research and laboratory medicine are considered in Chapter 10. Chapter 11 lists "Sources of Supply," while Chapter 12 is an Afterword, containing afterthoughts, aftershocks, and late breaking news. That's all I wrote. Well, almost.

Lis(z)t Mode

When cells are in such altered states
You don't know where to set the gates,
It's best to minimize the risk
And store them all on your hard disk.
If there's a clog before you're done,
You'll save some data from a run,
And, thus, you may stay out of jams
You'd get in with live histograms.

List mode, just work in list mode;
When you consider all the options, it's the only thing to do.
This mode, and only this mode,
Lets you make sense of samples that, at first, leave you without a clue.

Once we're in list mode, anyway,
With prices as they are today,
It isn't putting on the Ritz
To digitize to sixteen [or more] bits.
It's clear that, once we've made this change,
We'll have enough dynamic range
To transform data digitally,
So log amps will be history.

List mode, we'll work in list mode,
And go from linear to log and back without the log amps' ills.
Once we've got list mode, our only pissed mode
May be when we try pinning down which agencies will pay the bills.

List mode can help us analyze
How many molecules of dyes
And antibodies will be found
On each cell type to which they're bound.
At long last, different labs can see
Results compared objectively,
Advancing science as a whole
And aiding quality control.

List mode, by using list mode,
We'll all get heightened sensitivities and much reduce the fears
And trepidation of calibration,
Although the folks who make the particles may have us by the spheres.

From East to West, from South to North,
We'll send our data back and forth,
Why, we'll soon have it in our reach
To run our samples from the beach.
But, unless they've been well prepared,
When they are run, we'll run them scared,
List mode or not, there's still no doubt
That garbage in gives garbage out.

List mode, we all need list mode,
Though there are ends for which list mode itself can never be the means.
Even with list mode, there won't exist code
That gets good data from bad samples and/or misaligned machines.

("List Mode" © Howard Shapiro; used by permission. The music is derived from Liszt's Hungarian Rhapsody No. 2.)