Cell Calcium (1990) 11, 93–109
© Longman Group UK Ltd 1990

0143-4160/90/0011-0093/$10.00

Practical design criteria for a dynamic ratio imaging system

R.Y. TSIEN and A.T. HAROOTUNIAN

Department of Pharmacology, School of Medicine, University of California at San Diego, La Jolla, California, USA

Abstract — Problems encountered and solutions devised during the construction of a productive Ca^{2+}-imaging system are discussed. Many of these relate to the rapid and interactive nature of experiments on cytosolic Ca^{2+}. The emphasis on accurate photometric quantitation of dynamically changing images contrasts with that of most image-processing software packages, which concentrate either on leisurely massage of static images or on descriptions of lateral motions of objects without concern for their brightnesses. Particularly important goals for Ca^{2+}-imaging include real-time ratioing, psychophysically effective display formats, ease of experiment annotation, mass storage of image sequences, automated extraction of time courses and population statistics, aids to presenting images for seminars and publications, and program modifiability.

The recent surge of interest in fluorescence ratio imaging, especially of cytosolic free Ca^{2+} concentrations, has encouraged a large number of individual laboratories and companies to begin development of different instrumentation/computer systems to perform such measurements. Unfortunately, the construction of a good system requires close familiarity not only with image sensors, processors, and computer software, but also with the requirements of actual biological experiments, data analysis, seminar presentation, and publication. Combining the engineering and biological expertise is a challenge, particularly because both the technology and the biological needs are continually evolving. This paper gives our view of the major performance requirements that a ratio imaging system should fulfill, together with some of the solutions and compromises that were adopted in building a system and painfully refining it over the last three years while using it on a wide variety of tissues [1–18]. It is hoped that this discussion will not only help other biological users to define their own needs and select an appropriate system, but also help system builders to distinguish pressing needs from frills. In this way the wasteful duplication of multiple groups independently struggling to create separate incompatible systems might be slightly reduced.

By now the basics of fluorescence ratio imaging are widely known: a fluorescent indicator is loaded into living cells and viewed through a microscope with a sensitive video camera. The indicator dye has been designed to shift its wavelength distribution of fluorescence excitation or emission upon binding the species of interest [19]. Images are obtained at two appropriate wavelengths; by ratioing the intensities at corresponding picture elements, a map can be constructed showing the local analyte concentrations throughout the field of view. Since the process is nondestructive, the imaging can be repeated at

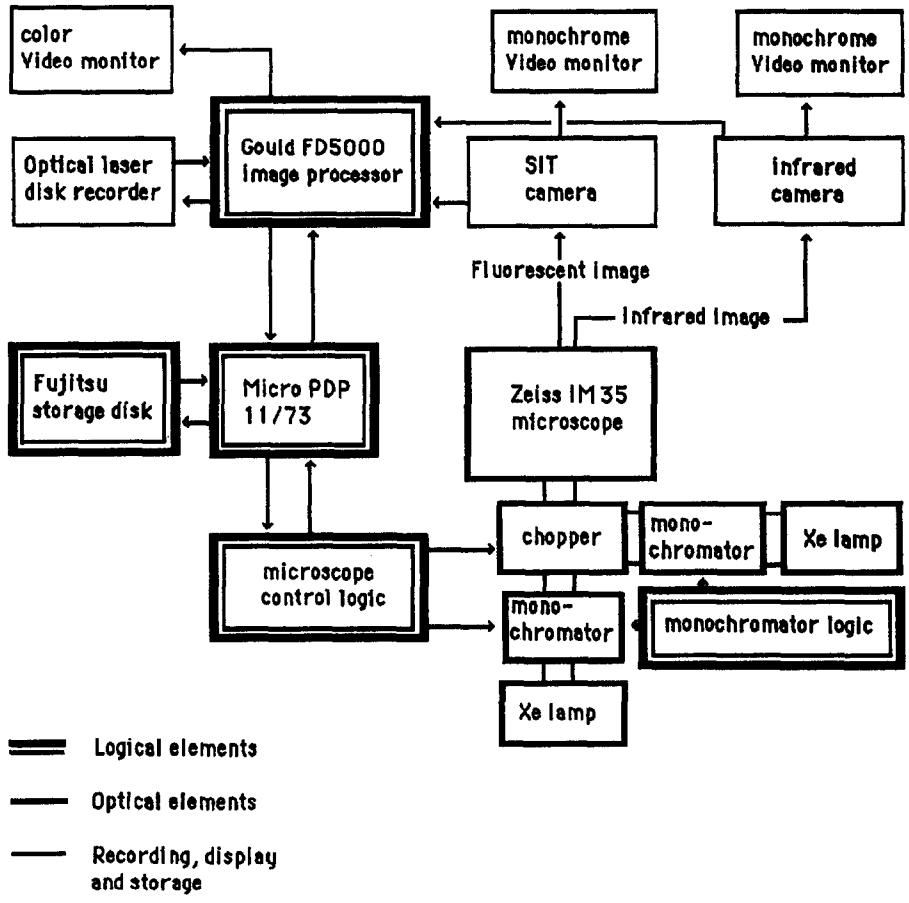

Fig. 1 Block diagram of the hardware in our current image-processing system. The Micro PDP 11/73 is the host computer; the chopper, monochromators, xenon lamps, and control logic in the lower right corner are from Spex Industries, Edison NJ, USA, modified as described in [22]

frequent intervals to trace out the time course of the cellular responses. The virtues of observing at more than one wavelength band and ratioing the intensities were first demonstrated millions of years ago in the evolution of color vision, a ratiometric system of remarkable accuracy [20]. In fluorescence measurements, ratioing compensates for variations in the thickness of the specimen, local concentration of indicator, and the overall absolute level of illumination and sensitivity of the detector [1, 21, 22]. Criteria for choosing appropriate indicators and whether to use excitation or emission ratioing have been reviewed elsewhere [15, 23, 24]. For imaging measurements of $[Ca^{2+}]_i$, Fura–2 with excitation ratioing is by far the most popular choice at present.

Hardware choices

A block diagram of the image acquisition and processing system is shown in Figure 1.

1. Microscope

The first choice is between an upright and an inverted microscope. We chose the inverted configuration because it greatly facilitates viewing living cells in saline using high numerical-aperture oil- or glycerin-immersion objectives. High numerical aperture ($\geq$ 1.25) is extremely valuable because the light-collecting efficiency varies approximately with the square of the numerical aperture. An inverted microscope copes easily with

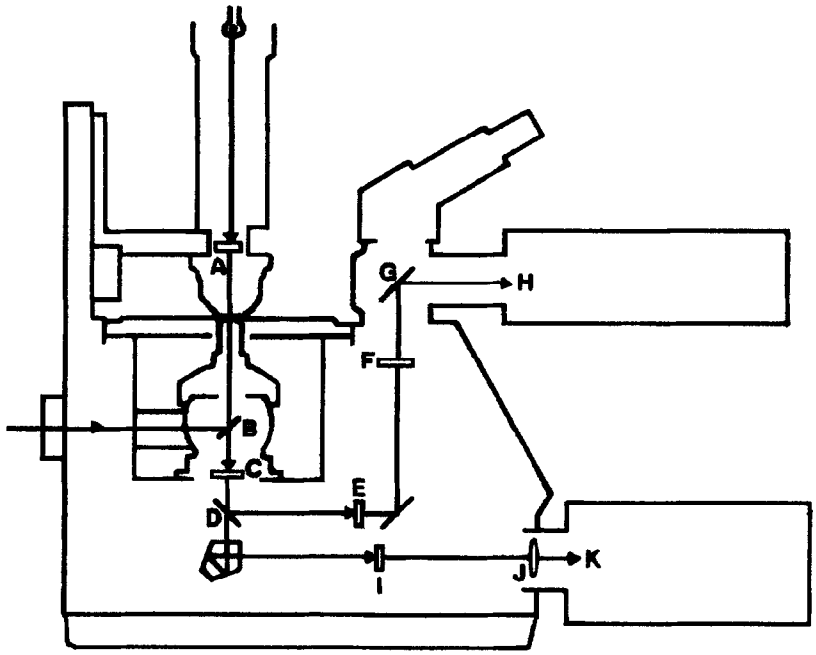

Fig. 2 Useful modifications to a Zeiss IM-35 inverted microscope. Only those optical components that were added or that are necessary to understand the deflection of the beams are shown. **A**, bandpass filter to pass 730–770 nm infrared from the tungsten lamp above; 750WB40 from Omega Optical (Brattleboro, VT, USA). **B**, dichroic mirror to reflect below 405 nm and transmit above 405 nm; DC405LP from Omega Optical. **C**, long-pass barrier filter to block wavelengths below 420 nm or 455 nm; GG420 or GG455. **D**, dichroic mirror to transmit below 700 nm and reflect above 700 nm (Omega Optical DR700SP), mounted on slider to alternate with a completely reflecting mirror of matching dimensions. **E**, optional polarizer for use with Nomarski differential interference contrast. **F**, optional phase ring for phase contrast observation. This ring was homemade using photoresist to capture the image of the phase annulus in the condensor, followed by vacuum deposition of MgF_2 and chromium. **G**, photochanger mirror to direct light either into the eyepieces or to camera H; a trinocular head could also have been used. **H**, infrared-sensitive Newvicon camera (Cohu Inc., San Diego, CA, USA). **I**, filter to contribute additional blocking of infrared wavelengths; BG39 (Omega Optical). **J**, convex lens to decrease the size of the projected image by focusing it less far outside the microscope. **K**, green-sensitive silicon-intensified-target (SIT) camera; model 66, Dage-MTI, Michigan City, IN, USA

the short working distance and non-aqueous immersion fluid demanded by such objectives, and allows the top of the chamber to be left open for easy solution changing or insertion of micropipets. An upright microscope demands either a sealed chamber with cells stuck to the roof, or water-immersion or dry objectives of much inferior light-collecting efficiency. We like the Zeiss (Thornwood, NY, USA or Oberkochen, FRG) research inverted microscopes because: (a) The stage is very rigid and vibration-free. (b) Even while the objective focuses up and down, both the dichroic mirror and the axis of the input excitation beam remain fixed in height above the table; this eases interfacing choppers, filter wheels, photolysis beams, lasers, etc with the excitation beam. (c) The transmitted-light condenser above the stage is particularly easy to pivot fully out of the way to permit changing or manipulating the specimen, yet returns easily to its prefocused height. (d) An infrared-reflecting dichroic mirror and a phase ring or polarizer can be placed in the optical path to permit differential interference or phase contrast observation at infrared wavelengths at the same time as measurement of visible fluorescence. Figure 2 shows the components added to our microscope. Because the fluorescence does not pass through either the polarizer (E) or phase ring (F), the latter do not attenuate the fluorescence as they would if placed at their conventional positions near or inside the objective. Somewhat analogous modifications have recently been described by Foskett [25].

Because infrared is so innocuous, we find this sort of continual transmitted light monitoring extremely useful to check the morphology of cells, the placement of micropipets, and the accuracy of focusing, especially while the fluorescence excitation beam has been blocked to minimize bleaching or photodynamic damage. Unfortunately, at present the modifications for simultaneous infrared are not commercially available, so custom component construction is necessary.

To maintain the cells at other than room temperature, we find that the most effective solution is to thermostat not just the holder for the specimen chamber but also the immersion objective, which is a large thermally conductive mass in intimate contact with the cells. We thermostat the objective by circulating water through a spiral of copper tubing wrapped snugly around the objective. Commercial stage thermostats or air-curtain incubators seem rather ineffective with immersion objectives, and large boxes to heat much or all of the microscope seem too cumbersome to permit frequent manipulation of the specimen in near-total darkness.

One significant problem that we have encountered with our inverted microscope is corrosion of the substage parts and mechanisms due to occasional leaks of saline or culture medium. These seem to be regrettably inevitable with inexperienced users, open chambers, and micromanipulations of cells in the dark. Unfortunately, the manufacturers mostly ignore the problem, assuming that inverted microscopes are used only with cells in sealed leakproof tissue-culture vessels.

2. *Objective*

The current opinion in our laboratory is that the objectives with the best combination of high ultraviolet transmission down to 340 nm, low autofluorescence, high light-gathering efficiency, and reasonable price are the UV-CF series from Nikon (Garden City, NY, USA), especially the 40x, 1.3 numerical aperture lens. In our system, this objective gives the video camera a field of view about 150 μm on a side, or about 0.3 μm per pixel, comparable to the limit set by diffraction. A higher-power objective such as 100x would simply give empty magnification with no better spatial resolution or light-gathering power but with several-fold fewer photons per pixel and a reduced field-of-view, limiting the number of cells that can be simultaneously monitored. Expensive flat-field apochromats are not appropriate, since they have too many internal elements for good UV performance, and fluorescence from an indicator properly loaded in the cytoplasm is inherently diffuse and fuzzy.

3. *Illumination source*

Xenon arc lamps are generally superior to mercury for excitation ratioing because xenon gives a relatively uniform spectral output, without the prominent emission lines of mercury. If one of the desired excitation wavelengths coincides with a mercury line, than it tends to be much more intense than the other wavelength to be ratioed, so that attenuation may be needed to avoid exceeding the dynamic range of the detector. If neither desired excitation wavelength matches a mercury line, then the excess radiation at those lines places greater demands on the wavelength selectivity of the monochromators and interference filters. Though mercury lamps put out more total UV energy, most of that is at 365 nm, which is close to where Fura–2 is least sensitive to Ca^{2+} [19]. A 75 or 150 watt xenon lamp aimed through reasonably efficient optics puts out more than enough energy for Fura–2 measurements. One application where the strong 365 nm emission of mercury would be highly valuable is in pulsed photochemistry [15–17, 26], but paradoxically only xenon is available in flashlamps.

The two main systems for selecting wavelengths, interference filters vs. grating monochromators, each have their pros and cons [22]. Interference filters are typically assembled into a wheel or multiposition changer, which is much cheaper and more compact than monochromators, which are typically alternated by a rotating chopper in which mirror vanes alternate with blank transmitting sectors. Filters give illumination which tends to be spatially more uniform than that from monochromator slits; moreover, any residual spatial variation tends to be matched between the two

wavelengths, whereas it is easy for two separate monochromators to become misaligned so that each preferentially illuminates a different part of the field of view. Filters in a wheel can readily give three or more sequential wavelength bands, which would be useful for nearly-simultaneous monitoring of more than one ion, whereas chopper systems usually only alternate between two beams. On the other hand, monochromators are freely and independently adjustable in wavelength and bandpass. If one or other wavelength needs to be attenuated, as is often needed to minimize bleaching or photodynamic damage or to minimize the dynamic range seen by the camera, narrowing the monochromator bandpass is a nice way to do the attenuation, since it is smoothly variable and improves rejection of stray light at other wavelengths, whereas insertion of neutral density filters in series with individual interference filters is clumsier and gives no side benefit of stray light rejection. Monochromators permit spectral scanning through the microscope, which is sometimes valuable to investigate whether the dye is behaving properly; they also permit closer comparison with cuvet measurements. A chopper-based system also gives more flexibility in interfacing other light sources such as photolysis flashlamps [16, 17]; however, it should also use the minimum number of separate mirror vanes (preferably just one) to minimize variations in reflectivity from one vane to another, which translate to artifactual variations in observed ratio.

Hybrid wavelength-selecting systems are also possible, e.g. chopping between two beams that each contain an interference filter, or wobbling the grating of a single monochromator. The very fastest and fanciest way of changing wavelengths is with an acousto-optical tunable filter (AOTF) [27], though the best acousto-optic materials work only in the visible, not the UV. It remains to be seen whether the microsecond response time of the AOTF outweighs its expense, poor commercial availability, and other optical limitations.

The excitation must be shuttered by the operating program so that illumination is delivered only when actually needed for a measurement or upon specific override from the experimenter. All dyes, including Fura–2 [28], bleach at some rate, and tissue exposure to potentially damaging short-wavelength irradiation must be minimized.

4. Low-light-level imaging detector

Currently the fundamental choice here [29] is between low-light television cameras and non-intensified, cooled charge-coupled devices (CCDs), which have readout rates considerably slower than standard video. The cooled CCDs [30] are expensive, obtainable mainly from one vendor at present (Photometrics Inc., Tucson, AZ, USA), and poorly compatible with the RS-170 video input standard of most image processors, but they have the highest quantum efficiency and photometric accuracy and least geometric distortion of any available camera. For applications in which relatively static images are to be acquired then subjected to intensive computational refinement, they are clearly unsurpassed [31]. For ratio imaging of rapidly-responding cells, the incompatibility with standard video is a serious problem. If one cannot use mass-produced image processors, one either has to spend a lot more money or cut down the number of pixels in the image (i.e. reduce the spatial resolution) or accept painfully slow production of ratio images. Technology is progressing in this area, and perhaps cooled CCDs will get faster and cheaper. Meanwhile, most experimenters have chosen the cheaper alternative of intensified television cameras, typically silicon intensified target cameras (SITs), further-intensified SITs (ISITs), or microchannel plate image intensifiers coupled by lenses or fiber optics to conventional TV cameras or non-cooled CCDs [32]. This is a fairly competitive and shifting area of commercial instrumentation. Here are some guidelines.

Many people assume that the lower the rated light level necessary to produce an image, the better the camera. Unfortunately, rated sensitivities are difficult to compare between manufacturers, partly because criteria for a 'usable' image differ, partly because most ratings are performed with broadband incandescent illumination of 2856°K color temperature, which contains predominantly red and infrared. A camera with extended red sensitivity, as found in microchannel plate intensifiers with S-20 photocathodes, rates much better with such illumination than a camera which has equal

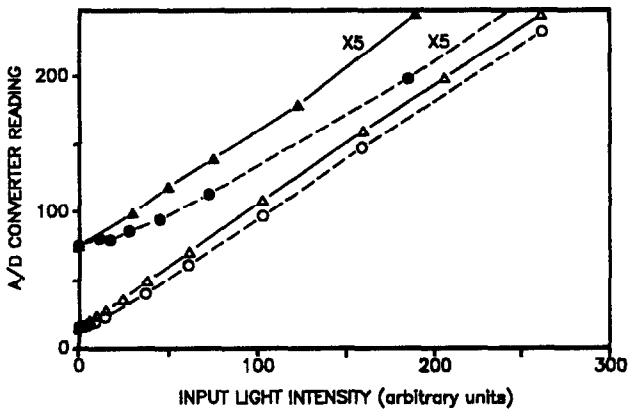

Fig. 3 Response linearity of SIT camera (model 66, Dage-MTI) before and after modification. Circles connected by dashes are from the camera as originally received and with gamma set to 1.0. Triangles connected by solid lines are after modification to remove the effects of a slightly nonlinear black clipping circuit at the output. Open symbols are the original data, closed symbols show the abscissae and ordinates multiplied by 5 in order to expand the plot at the lowest light levels. The abscissae were read from a photomultiplier monitoring the input light, the ordinates from temporally- and spatially-averaged readings of the camera output by the image processor. Control experiments with artificial inputs to the analog-to-digital (A/D) converter indicated that it was linear enough not be the cause of the initial nonlinearity. The modifications consisted of connection of a 2.7 kΩ resistor from the base of Q21 to the emitter of an added PNP transistor (2N4121 or nearly any equivalent that can withstand 15 V from emitter to collector). This transistor acted as an emitter follower, whose base was connected to the junction of CR14 and R53, and whose collector was connected to the -12V supply (bus D). Also the positive end of R83 was disconnected from the +12V supply (bus A) and transferred to the junction of VR6 and R88

sensitivity in the green (where Fura–2 emits) but less red sensitivity. SITs and ISITs are examples of cameras whose peak spectral sensitivity peaks near the Fura–2 emission but have poor red response. An even more important yet less obvious consideration is that the rated ability to produce some picture (albeit fairly noisy) at very low numbers of photons/s is less important than the ability to produce a fairly quiet picture with the fewest number of integrated photons. We can choose to deliver those photons either at higher intensities for a shorter time or at low intensities for a long time. Our impression from side-by-side comparisons of a SIT and an ISIT is that though the ISIT can produce a picture at much lower light levels, the output is so noisy as to require very lengthy signal averaging. Moreover the noise occupies a significant portion of the allowable dynamic range of the video signal, so that to avoid clipping the noise peaks, the signal dynamic range must be considerably restricted. If one accepts that light levels have to be high enough to give a reasonably quiet picture within say 1 s of digital accumulation of frames, in our experience the SIT does as well or better than the ISIT, perhaps because the extra amplification stage in the latter introduces noise even when the gain of each individual stage has been reduced to avoid overall saturation. Where the ISIT's greater sensitivity in photons/s or footcandles really matters is when a dim light source cannot be pulsed, i.e. when observing chemiluminescence or outdoor scenes at night. In fluorescence the option of increasing the illumination while shortening the duration enables cheaper, more rugged cameras and quicker measurements with the same total irradiation of the specimen.

We currently use a Dage-MTI (Michigan City, IN, USA) model 66 SIT camera with a Heimann (rather than RCA) tube selected for high photocathode sensitivity rather than a low number of blemishes. The overall sensitivity is such that with grating monochromators at reasonable bandpasses of about 5 nm, the autofluorescence of cells and the glass optics are usually visible at maximum camera gain. Since for meaningful results the dye signal has to exceed autofluorescence by a considerable margin anyway, the sensitivity would appear to be adequate. The photometric response linearity of the camera was tested at length using illumination from the tungsten illuminator, fed by a well-regulated constant current supply. The light passed through various neutral density filters onto a 0.25 mm diameter pinhole acting as a specimen. This was viewed through a 16x objective followed by a 500–530 nm interference filter so that the spectrum would match the peak emission of Fura–2. The resulting image of a green glowing circle in a black surround simulated a large cell but was much more stable, and was extensively averaged while being simultaneously monitored with a beamsplitter and photomultiplier, which was assumed to have a

perfectly linear response. Even though the camera was set to a gamma of 1.0 to give a nominally linear response, it initially gave a response in which the smallest signal levels were compressed nonlinearly (Fig. 3). Other makes of SIT camera such as Hamamatsu (Hamamatsu City, Japan) have also shown this nonlinearity. In the Dage, we found that we could get a much more linear response (Fig. 3) by adding a resistor and a transistor to the output black-level clamping circuit. It is well worth checking the camera linearity, since the accuracy of ratioing depends strongly on the assumption that twice as much input light will give twice the incremental video signal above the voltage level for the dark background. One sign of nonlinearity is that thick portions of cells (where signal levels are high) tend to give systematically different ratios (and therefore apparent $[Ca^{2+}]_i$ levels) than thin portions where both excitation wavelengths give weaker signals. The resulting images are physiologically implausible, especially in resting cells, but ironically they are sometimes featured in advertisements because the multicolored bulls-eye patterns are visually exciting to the naive.

The main camera types except CCDs have significant 'lag' or afterimages. The time courses of the responses to both light onset and offset can be nonexponential and require hundreds of milliseconds to settle to sufficient accuracy. Therefore, an image newly presented to the camera should not be digitized until after an appropriate period of waiting for the quantitative intensity to stabilize.

5. Infrared TV camera for transmitted light viewing

This camera is far less critical than the low-light fluorescence-monitoring camera. We use an inexpensive Cohu (San Diego, CA, USA) camera with a Newvicon tube, whose response peaks at 750 nm and which is considerably more sensitive than most ordinary CCDs or vidicons. 750 nm happens to be a reasonably optimum wavelength, since shorter wavelengths are less easy to separate from fluorescence emission, but longer wavelengths are not correctly filtered by the sheet polarizers commonly used for differential interference contrast. Reasonably good sensitivity is desirable, again because it reduces the necessary infrared intensity and potential crosstalk with the fluorescence signal. Response linearity is irrelevant, though ability to subtract a background light level and boost the contrast is highly valuable. The image processor should be programmed to permit easy insertion of images from this camera into the stream of recorded fluorescence images; for some experiments in which the cell morphology is changing rapidly, it is helpful if the system can be instructed to pair every fluorescence ratio with a matching transmitted light view automatically.

6. Image processor

Here too there is a bewildering variety of commercial systems from which to choose. We would offer the following two guidelines: (a) as in most computer purchases, it is usually a mistake for the end user to chase after the latest and most glamorous new hardware, especially if it requires writing new software. It is far more cost-effective to buy models that are a few years old but for which software libraries or, better yet, entire application programs are available and debugged. (b) Assuming one has chosen a system that accepts standard video input at 25 (European) or 30 (North America and Japan) frames per second, the image processor should not only be able to digitize and add and subtract frames at that rate but also contain multiple look-up tables that can apply a variety of arbitrary transformations at full video rate. The reason is that it is enormously valuable for the experimenter to be able to see the ratio images in almost real time, or within a second or so of completing the acquisition of the excitation wavelength pair. Such rapid display enables the experimenter to judge the progress of the experiment, decide when to perform the next manipulation on the cells, or recognize when the cells are dead and it is time to clean up. With a video-rate 16-bit-wide pipeline processor and sufficient lookup tables, the necessary background and shading corrections and ratioing can be done in a few frame times or a fraction of a second (see section 7 under *Software requirements* below). Since this is approximately the time required anyway to change wavelengths and outwait camera lag, it causes no additional time penalty. By contrast, many seconds are typically required if the

host computer has to do the arithmetic, or if the video processor's lookup tables are too small and their contents have to be repetitively exchanged. If ratioing is slow, one has the unpleasant choice of either potentially missing important cellular events, or simply recording the input images for separate off-line processing, which ruins the interactive nature of the experiment and is wasteful of videotape or videodisk space. We have found no need for hardware convolution or histogram generators or vector processors, etc; these accessories are mainly of use to people who want to analyze the forms and shapes of images rather than ratio their intensities.

We use an FD5000 image processor from Gould (Fremont, CA, USA). Though this unit has some intrinsic design idiocies and is now fairly obsolete, it also has some unusual features that are useful for this application, such as a built-in 'intensity reference scale', which we use to generate the color bar that calibrates pseudocolor in terms of ratio, $[Ca^{2+}]_i$, or pH_i; 12-bit pseudocolor, or the ability to display 4096 simultaneous colors from a palette of 16 million possibilities; and three separate large lookup tables (12 bits input, 16 bits output) to transform both the input operands and the output of the arithmetic logic unit. Nearly all sectors of these lookup tables are used heavily.

7. Host computer and operating system

Obviously the main criterion here is compatibility with the chosen image processor and its library of support software. The primary task of the host computer is not to do a lot of computation but to serve as real-time interface between the experimenter, the image processor, and auxiliary equipment. Because several processes often need to be overlapped, e.g. changing excitation wavelengths, driving the image processor to do a ratio, commanding a laser disk recorder, and accepting input from the operator, it is highly convenient to use a time-sharing or multitasking operating system. However, this needs to be a 'real-time' operating system as well, since some tasks must be given enough priority to execute or respond within externally fixed time limits, others must wait for the former types to finish, whereas yet others should be done only when the system has nothing more pressing to do. An advantage of a real-time multitasking system is that other users can be checking files, graphing data or editing programs concurrently with the live imaging experiment. Popular operating systems such as DOS or Unix or the Macintosh are not designed for real-time multitasking, since the mass computer market does not have this specialized requirement. Of course, even without a favorable operating system, a skilled programmer can write a single huge program to juggle the multiple demands on the system, but that is essentially re-inventing multitasking. We use a Digital Equipment Corp. (Maynard, MA, USA) PDP-11/73 minicomputer running the RSX-11M operating system, both of which are old products but well-suited to coordinating multiple tasks with hierarchical priorities. Six tasks run concurrently to do an on-line experiment: the master routine that runs the acquisition and ratioing, a task for on-line data analysis, and tasks that interface with the chopper or filter wheel, the magnetic disk, the optical disk recorder, and user input via keyboard and footpedals. These total about 3500 lines of FORTRAN. For experiments that combined imaging with electrophysiology, yet another task was required to do electrical stimulation and recording.

8. Mass storage for video images

Bitter experience shows that the demand for storage capacity is almost insatiable. A full image of 512 pixels horizontal x 484 pixels vertical x 8 bits/pixel occupies nearly a quarter million bytes. One may choose just to store the final ratio image or the two raw pictures at different wavelengths; the latter choice obviously doubles the transfer time and storage space, but gives the user greater flexibility to reprocess the data, especially to change the range of ratios that are coded in the output picture. Depending on the experiment, several hundred to several thousand such ratio images or wavelength pairs are easily generated in a single successful run, and a good day will produce several such runs. If all users are exhorted to be frugal and to store only the pictures that show a visible change from the previous one, the number of pictures can be reduced

a few fold, but the resulting time courses of $[Ca^{2+}]_i$ responses show sparse and unevenly-spaced data points and occasional large gaps.

We started with a 30 megabyte hard disk and soon had to buy a much larger one of 270 megabyte formatted capacity. This was adequate for a few runs per day, but all the data still had to be transferred to magnetic tape once or twice a day, a time-consuming process that was too finicky to leave totally unattended or run concurrently with experiments. Also, to review or analyze any experiment more than a day old, a chunk of space on the disk had to be cleared and the data transferred back from tape. The tapes themselves were expensive (about $30 for a cartridge holding 220 images) and imperfectly reliable. In principle, tapes are erasable for re-use, but we find this is rarely practical, since most tapes have at least a few pictures that one wants to keep. Reviewing an entire tape to pick out and transfer those that deserve preservation is sufficiently laborious that one tends to use new tapes rather than reuse old ones.

Eventually we have settled on two complementary solutions to the storage problem. One is based on the recognition that many experiments do not need to preserve full spatial information but merely require enumerating the ratio or $[Ca^{2+}]_i$ in each of several pre-designatable cells or subregions in the field of view as a function of time. Once the cells have stabilized on the microscope stage, the user draws circles or polygons on the screen to outline the regions of interest. From then on, most of the images are not stored as such but rather only the ratios suitably averaged over all the pixels in each region of interest, plus any other statistical information desired. This listing occupies only tens to hundreds of bytes per time point rather than hundreds of thousands. This approach becomes impractical if the cells move much during the run, since it takes a good deal of time to realign the circles or polygons. For experiments where the zones of interest are shifting or unknown in advance or detailed spatial gradients are to be examined, it is still necessary to record complete images, which we now do on a TQ-2028F monochrome high-resolution optical disk recorder from Panasonic (Secaucus, NJ, USA). Each disk holds 16,000 black-and-white images; data is transferred at video rate in analog form. For a ratio image, we store the pair of constituent images at the two wavelengths, since this gives more reprocessing flexibility than just storing the pseudocolor ratio image, whose digital coding in our system would be incompatible with the TQ-2028F anyway. 8,000 ratio pairs directly stored and randomly accessible on an optical disk costing slightly over $100 in quantity is considerably superior in speed, convenience, and cost to most digital magnetic tape systems, especially since it takes a long time to access data at the end of a tape. The inability to erase the optical disks once written is not a problem in this application, as explained above. However, the optical disk video recorder does have a number of limitations that the operating program must overcome. (a) Though the linearity of the recorder's response is excellent over most of the dynamic range, the video signal played back in general undergoes a slight change in gain and offset voltage in comparison to the original input. The offset change is handled by storing the dark background on the laser disk recorder, so that subtraction of the replayed background from the desired fluorescence image cancels out the recorder offset along with the camera black level voltage or setup level. The gain change of the recorder is canceled by the ratioing between the two wavelengths. The need for such cancellation is another reason why we do not store already-ratioed images on the optical disk recorder. Fortunately, the noise level is adequately low if the playback is averaged for 4 frames or more. (b) The TQ-2028F has no mechanism. for storing digital label information along with each image. Also, it has no inherent file system, only absolute track numbers 00001 to 16000. Users, who are working in a darkened room and have delicate manipulations to make on their cells, cannot be expected to keep accurate notes of where the images are being stored at considerable speed. Therefore the operating program must construct image directories on magnetic disk which document the circumstances leading up to each image and show its track number(s) on the optical disk, so that retrieval should merely require specifying the nth image from a named run. An optical disk system that stored digital data in files rather than analog video images

in tracks might overcome the need for directories on magnetic disk, but we were unable to find a cost-effective system that could transfer data fast enough (0.25 – 0.5 megabyte in 0.25 – 0.5 s) directly from the image processor without tying up the host computer. However, this may well change.

Software requirements

1. Mode of ratio representation

One of the most fundamental decisions is how ratio images are to be displayed. Plate 1 contrasts several common display formats; the specimen happens to be REF-52 fibroblasts stimulated by depolarization together with the hormone vasopressin. This condition induces $[Ca^{2+}]_i$ oscillations [16], which are asynchronous in the individual cells, so that any instant some are high in $[Ca^{2+}]_i$ and some are low. Panels 1A and 1B show the images obtained with 350 nm and 385 nm excitation respectively; note that from these images it is not difficult to see the boundaries between adjacent cells. (340 nm would actually be slightly better than 350 nm for Fura–2, but our microscope and objective originally had poor transmission at 340 nm, so we standardized on 350 nm at an early stage.) The most common ways of displaying ratios are as increasing height on a projected three-dimensional graph (panel 1C), or as monochrome brightnesses (panels 2A–C), or as pseudocolor hues (panels 3A–C and 4A–B). The relationship between ratios and hues can either be according to some completely arbitrary palette or in spectral order. We chose spectral order because ratio and $[Ca^{2+}]_i$ are inherently continuous measures, and a spectrum is a logical way to arrange hues in a continuous and intuitive manner. There seems no reason to use pseudocolor scales that insert arbitrary discontinuities into the ratio scale. We chose blue to be represent low $[Ca^{2+}]_i$ because of the cultural bias that blue and green are restful whereas red and magenta are 'hot' and activated, which would appropriately match high $[Ca^{2+}]_i$. High and low pH were assigned blue and red respectively because nearly all wide-range pH test papers work that way. Given that ratios are coded by spectral hues, those hues can either be displayed always at full or zero intensity (panels 3A–C) or with smoothly graded intensities (panels 4A–B).

An important related question is how to cope with regions (such as between cells) where the numerator and denominator of the intensity ratio are both nearly zero, so that the ratio is meaningless and noisy. The most common solution is to ask the user to set an intensity threshold, so that pixels whose numerator or denominator or mean intensity fails to reach the threshold are displayed as some sort of background, for example zero height in panel 1C, or a color such as black (panels 2A–C and 3A–C) or blue. Such thresholding is quite essential, as shown by the indecipherability of panels 2A and 3A in which no thresholding was applied. Unfortunately, it is often difficult to set a single threshold that cleanly separates background garbage from desired cells. In panels 2B and 3B, a low threshold was applied, which only slightly helps to distinguish neighboring cells from another. In panels 2C and 3C, the threshold was higher, which enables the central cells to be separated but nearly wipes out dimmer but still interesting cells. Threshold setting is a standard requirement of most ratio imaging systems but is uncongenial to us, since it forces the user to spend time making an arbitrary distinction between pixels that are deemed fully significant from neighboring ones that are treated as completely worthless. If the threshold is set too low or too high, important features may become unrecognizably imbedded in noise or be completely suppressed respectively. Instead, we prefer to retain the mean intensity (the average of the two wavelengths) and use it to smoothly modulate the brightness of the hue representing the ratio. The higher the mean intensity, the more significant the data are likely to be, and the more brightly the color is displayed. Retention of intensity information enables one to intuitively judge where cells are thick or thin, where dye is accumulated or excluded, when dye may be leaking out or bleaching. Moreover the human visual system recognizes the shapes and borders of objects primarily from intensity rather than hue information; psychophysics shows that displays of constant luminance in which only hue varies are visually unsettling and hard to decipher quickly [33]. Originally [1–13, 15] we used five out of the eight bits constituting each pixel to code for

hue, which therefore had 32 possibilities. The remaining three bits coded for 8 intensity levels. Panel 4A illustrates this scheme, though only about 3 intensity levels were actually attained in this example. Even if one discounts the desirability of including intensity information, there is little point in reserving all 8 bits for ratios or hues, since the eye cannot discriminate 256 distinct hues and the noisiness of the ratios usually exceeds the quantization increment even with only 32 steps. More recently [13, 16–18] we have switched to 12-bit pseudocolor (panel 4B), in which hue and intensity are each coded by 6 bits or 64 possibilities. Both are now coded finely enough for the quantization to be essentially unnoticeable. Therefore at the edges of cells the pixels fade smoothly into blackness, giving a more realistic appearance than the modes of representation with all-or-nothing (1-bit) intensity coding (panels 3A–C). Indeed, only panel 4B is comparable to the pre-ratioing raw images (panels 1A, B) in the ease with which cell boundaries can be recognized.

The opinion is sometimes expressed that only monochrome representation is really justifiable, whereas pseudocolors are just frills for people with too much equipment money. Although we acknowledge that pseudocolors are expensive to reproduce on paper and are unfair to the color-blind, we are unable to conceive of a better way to display mean intensity and ratio simultaneously for a quarter million pixels. They have the advantage of being relatively easy to discriminate and remember in absolute categories such as blue, green, yellow, red etc; by comparison, it is more difficult to estimate the $[Ca^{2+}]_i$ of any given cell when one is limited to monochrome brightness comparisons as in panels 2A–C. Moreover it is hard to judge whether a cell is getting brighter or dimmer unless the successive images are displayed rapidly without any intervening breaks. Three-dimensional graphs such as panel 1C are very time-consuming to compute and sacrifice lateral resolution. It is visually difficult to compare heights across the width and breadth of the image, and mountains at the front of the picture conceal valleys in the rear. If one must have a hardcopy that can be photocopied, the digital coding in panel 4C might deserve consideration, since it has the same lateral resolution as in panel 1C but is much faster to compute and makes it easy to read off the quantitative ratio or $[Ca^{2+}]_i$ of any zone.

2. Mode of entering commands and changing parameters

Some people like mice, others prefer typing English-like commands, others like dedicated function keys, others like typing short codewords. Though a system with mice and multiple menus overlaying each other may feel user-friendly to the inexperienced, one should remember that real experiments are done in a dark room usually under considerable time pressure once the first stimulus has been applied to the cells. The other parts of the apparatus and the experiment are sufficiently difficult to make it unlikely that casual users could just drop in and get meaningful results, so that considerable training and practice will be necessary in any case. Therefore we chose to make the user interface fast and efficient for experienced experimenters rather than verbose and multilayered for naive newcomers. The approximately thirty parameters that control image acquisition and display are all displayed on screen at once, and the user can change any one of them by a single alphanumeric character followed by the new value. Such alteration can be specified at any time, not just when the program is between operations; this is important particularly when free-running repetitive acquisition of ratio images has been launched. Processing commands are specified with mnemonic letters such as 'b' for taking a pair of background pictures, 'c' for taking color ratio pictures continuously, 'w' for writing the image onto the disk, etc. This works all right for the simplest and most common functions, but more arbitrary letters had to be used for the later-developed operations, but they still seem a little easier to remember than commands such as Alt-F4 or Ctrl-F8.

3. Annotation

It must be easy for all relevant manipulations of the cells to be documented and time-stamped in the computer record. This is needed because it is generally too dark to write in one's experimental

notebook, and one wants both the manipulations and the resulting images to be time-stamped by the same clock (obviously the computer's rather than one's watch) so that accurate latencies can be deduced. If it is not easy to add or revise running comments, people in the excitement of an experiment will neglect to record their actions precisely, and the whole run may be wasted for subsequent analysis. We enter time stamps in two ways: if two people are available, one person gets ready to make the manipulation while another types its description into the comments field. At the moment of execution, the first person says 'Now!' and the second types a dollar sign, which the program immediately replaces by the system clock time at that moment. But often one is doing the experiment alone, particularly at odd hours, and needs both hands to make the solution change, re-focus the microscope, push the stimulus button, etc. Then during a quiet preparatory moment, one types in all the anticipated manipulations, each prefixed by a '{' symbol and containing an imbedded dollar sign. These latent comments are held in reserve and are entered into the record only when one steps on a footpedal at the actual moment of execution. Each pedal press actuates the oldest remaining latent comment and converts its dollar sign into the relevant clock time. This feature is also quite useful when several manipulations must be made in quick succession without enough time for typing in between. As a further aid to lone experimenters, a separate footpedal is available to trigger image acquisition and/or recording onto disk. (Two footpedals are as much as most people can cope with.)

4. Background and shading corrections

Even when there are no cells or dye on the microscope stage, there will be some voltage output from the camera, consisting of autofluorescence from the optics plus tube dark current plus a black level setup voltage inserted by the camera electronics under the control of the user. This output level may vary with position within the image and with excitation wavelength (due to the autofluorescence component); it should be adjusted (using the 'black level' control) to give a small positive reading in the analog-to-digital converter over the entire image even including noise excursions. (If negative peaks of noise fall below the zero level of the A/D converter, they are clipped off, causing the average to become biased in a nonlinear manner.) The background is then stored as a pair of complete images, one at each wavelength, to be subtracted from any subsequent fluorescence image. Ideally the background would be recorded with the chamber filled with its normal solution and with the sensitivity of the camera matching the settings to be used with the dye-loaded cells. If the necessary sensitivity is not known in advance, it is helpful to be able to store (on magnetic and/or optical disk) several backgrounds covering a range of camera gain settings, so that once the cells are added and the necessary sensitivity is determined, one can retroactively choose which of the several backgrounds is most appropriate. The working pair of background images is then held in two banks of the image processor memory for instant use.

One other correction image needs to be acquired, an image showing the relative intensities of the two excitation wavelengths over the entire field of view. This map, which we call the 'shading' image, is obtained by placing a droplet of dye sandwiched between two cover slips onto the specimen stage. This dye sample is assumed to be spatially uniform, should have a reproducible excitation ratio that is comparable to the ratios that one expects to observe in the cells, and should be concentrated enough to give a high signal/noise ratio at both wavelengths. For $[Ca^{2+}]_i$ measurements we use several mM Fura–2 free penta-anion in 100 mM EDTA. Fura–2 at high Ca^{2+} would not be a suitable standard at our customary wavelengths of 350 and 385 nm because its 385 nm signal is so weak relative to 350 nm that it is difficult to fit both within the dynamic range of the camera and maintain good signal-to-noise. Also, the excitation ratio at high Ca^{2+} is extremely sensitive to slight degradation of the Fura–2 upon aging. Fura–2 at zero Ca^{2+} is much more reproducible and suitable as a standard, except that its ratio tends to drift upwards, probably due to leaching of Ca^{2+} from the glass cover slips. For this reason a huge concentration of EDTA (which is less pH-sensitive than EGTA) is used, but even so the 'shading' image should be recorded from a

freshly-made sandwich of dye. If monochromators are being used, their slits are adjusted so that the 350 nm image is about half as bright as the 385 nm; then at typical cytosolic $[Ca^{2+}]_i$ levels, the two images are roughly comparable in brightness (as in Plate 1, panels 1A and 1B), which minimizes the dynamic range with which the camera must cope. The 'shading' image corrects the subsequent ratio images for the relative intensity of the two excitation beams, which can vary over the field of view. If this correction is absent or incorrectly estimated, the usual result is that the excitation ratio and $[Ca^{2+}]_i$ appear to differ systematically from one side of the image to the other; if the stage is rotated or translated, cells or subregions seem to change $[Ca^{2+}]_i$ as they are moved from one part of the field to another.

5. Black level setup voltage and gain setting

It is essential that the camera output fall within the dynamic range of the analog-to-digital (A/D) converter of the image processor at all times, both for acquiring background and shading corrections and for the actual images from cells. To aid the user in adjusting the black level setup voltage to fit the input range of the A/D converter, it is helpful to alternate wavelengths repetitively while the image processor temporarily feeds the live input directly into the color output via lookup tables that accentuate the extremes of the A/D converter range. For example, we set the lookup tables so that an A/D output of 0 gets displayed as a bright blue; A/D counts of 1–4 are converted to yellow-orange. Most of the scale, levels 5 – 239, are displayed as undistorted gray-scale intensities, whereas high readings of 240–254 are shown as lime-green, and the saturating level of 255 becomes bright magenta. The color discontinuities were desired here to give sharp warnings. Thus the black level setup voltage is adjusted so that the image is entirely very dark gray with a few flecks of yellow-orange but no blue; the camera gains for the shading sample and for the actual cells are adjusted until the display shows at most a few zones of lime green but no magenta. This type of set-up display enables the entire image to be checked at a glance, and eliminates the need for a separate oscilloscope or video analyzer, whose calibration with respect to the A/D converter would have to be determined separately.

6. Averaging of input frames

We prefer straight unbiased averaging of 2^N frames at each wavelength because that seems to make better statistical use of the noisy signal than schemes such as exponentially weighted rolling averages. For N = 1 to 8, or 2 – 256 frames averaged, straight summation of 8-bit input needs up to 16 bits of video-rate accumulator and image memory. Because of the finite time required to change wavelengths and overcome camera lag, all the frames at one wavelength are accumulated in one contiguous time segment, followed by a wait period, then all the frames at the other wavelength. The temporal separation of the two wavelength opens the possibility of misregistration artifacts if a cell moves significantly between the two exposures. For example, if the shorter wavelength precedes the longer wavelength of Fura–2 excitation, the leading edge of a moving object will artifactually appear to have a low $[Ca^{2+}]_i$ and the trailing edge to have a high $[Ca^{2+}]_i$. The hallmark of this motion artifact is that its sign reverses if the temporal order of wavelengths is inverted. Such inversion should be easy for the user to specify, and may optionally be made automatic between every acquisition cycle. While it is important to help the user recognize the motion artifact, it is not clear that it is worth trying to calculate corrections for the effect, since that would require assumptions about the smoothness of the motion.

7. Ratioing algorithm

In our system, the sum of 2^N frames at each wavelength is right-shifted by N binary positions to normalize it back to 8 bits. If storage on the optical disk recorder is desired, the pair of averaged monochrome images is transmitted at this point. No extra D/A converter is needed, because the recorder shares one of the color outputs (red, green, or blue) with the color monitor. Momentarily the normal color image is replaced by a pair of black-and-white images, but the interruption is so brief as to be acceptable, and even gives reassurance that the

recording is in progress. Then the backgrounds (likewise stored as 8 bits) at each wavelength are subtracted from the corresponding fluorescence images. Negative numbers are occasionally generated in this subtraction, since due to noise some pixels may fall below the background intensity. Negative results are characterized by all ones in the high 8 bits of the 16-bit difference and are filtered out (i.e., converted to true zero) by a lookup table just after the adder/subtractor. The two 8-bit background-corrected images are then simultaneously passed through separate complicated lookup tables. The shorter-wavelength image gets transformed so that its most significant 5 bits are transferred to bits 10–14 of the output, whereas its suitably-scaled logarithm fills bits 0–8 of the output. A table of logarithms cannot extend down to zero input; our logarithm table is truncated at an input of 4, so that for $0 \leq x \leq 4$, the output is zero, whereas for $4 \leq x \leq 255$, the output is $512 \cdot (\log_2 x - 2)/6$, which fills the range of 0 up to 511 in output bits 0–8. The transformation of the longer-wavelength image is identical except that the logarithm table is inverted about 512, so that for $4 \leq y \leq 255$, bits 0–8 become $512 - [512 \cdot (\log_2 y - 2)/6]$, or a range of 512 down to 1. The error in handling inputs of 0 to 3 is acceptable because if both wavelengths have such small intensity values, the intensity coding will show the ratio with zero luminance (i.e. black) anyway; if one wavelength is bright enough to raise the mean intensity enough to make the pixel visible, then the ratio would be extremely low or high and would be recognized to be imprecise anyway. The two lookup table outputs are then added; the result in bits 10 – 15 is the 6-bit mean intensity, range 0 to 62, while bits 0–9 contain $512 + 512 \cdot \log_2(x/y)/6$; in this 10-bit logarithm, ratios (x/y) of 1/64 to 64 are mapped into the range 0 to 1024. Each integer increment in the output represents an increase in ratio by a factor of $2^{6/512} = 1.0082$, so ratios are calculated to an accuracy of better than one percent. These transformations to obtain the mean intensity and log ratio may sound complicated but are actually executed in one frame time (33 ms); the lookup tables are calculated by the host processor and loaded just once when the program is first initialized. The 10-bit log ratio then receives its shading correction by subtracting the 10-bit log ratio previously obtained from the uniform dye sandwich. The ease of this ratioing of ratios illustrates the advantage of using logarithms. Finally, a user-chosen subdomain of the log ratios (for example, ratios between 1 and 12 times that of the Fura–2/EDTA shading sample) is linearly mapped into the 6 bits of output hue. Ratios below or above the specified limits are forced to hue 0 or 63 respectively. The final composite pseudocolor consists of 6 bits of mean intensity and 6 bits of hue representing log ratio. This picture appears on the color monitor about 0.3 – 0.4 s after completion of acquisition at the second wavelength, and remains there during acquisition of the subsequent ratio pair. If 16 video frames are averaged at each wavelength, the system can generate a ratio image about every 2 s. If laser disk recording is specified, that adds about 1 s per ratio pair, mainly due to the slowness of the RS-232 handshaking, but it can continue until the optical disk is full (8000 ratio pairs). We prefer each gradation in hue to represent an equal increment of log ratio rather than linear ratio, since the log scale expands the lower ratios representing $[Ca^{2+}]_i$ values near resting levels, where the dye is operating near its dissociation constant and gives relatively precise readings, whereas a linear ratio scale devotes too much of its span to high $[Ca^{2+}]_i$ values, which cannot be precise anyway because the dye is nearly Ca^{2+}-saturated at those ratios. Of course, another reason for preferring log ratios is to avoid having to compute the antilogarithm. We also do not bother trying to apply the shading correction to the mean intensities, since the latter are intended more to give a qualitative idea of where the data from the camera is more or less reliable, than as a rigorous quantitative description of dye localization.

8. Data analysis

The most rudimentary quantitative analysis consists of cursors that can be moved around the image and which read out the local ratios and $[Ca^{2+}]_i$ or pH_i values. However, these values are often fairly noisy with large variations between neighboring pixels, so that the demand soon arises [5–14, 16–18] for averaging the ratios or ion concentrations over regions of interest, either during the actual

experiment (as discussed above under *Mass storage*) or from a past sequence of images. The programmer finds rectangular regions much the easiest to cope with, but cells practically never lie in such rectangles nicely aligned with the horizontal and vertical axes. Most disaggregated cells are nearly spherical, so we find that circular regions are the most useful. Occasionally irregular polygons are needed. Once the zones have been set up, the system should permit automatic monitoring of the spatially-averaged behavior in each region as a function of time, i.e. over many successive images. In a few experiments [7] with large populations of small cells such as lymphocytes, we have used as many as a hundred individual zones, but five to twenty would be a more typical number. Each zone needs to be labeled and its position and dimensions recorded, so that if later analysis reveals an interesting response pattern in a particular zone, one can reconstruct which cell was responsible. Unless all the zone boundaries are set very precisely, some form of intensity-based masking is needed to reject contributions from the noisy ratios from the dark regions between cells. Our present choice, which is inherently tolerant of a noisy surround, is to weight the log ratio at each pixel by its mean intensity in forming the spatial average log ratio, which is recorded in tabular form in a disk file. The job of computing the antilogarithms, converting to $[Ca^{2+}]_i$ or pH_i, and plotting as time courses, scattergrams, or frequency-of-occurrence histograms is left to an offline analysis program. This deferral keeps the on-line part of the analysis relatively simple and fast, and permits much greater flexibility in the ratio-to-concentration calibration, including separate retrospective calibrations for each cell if desired. One may question whether pixel averaging should be done on the log ratio rather than the actual ratio or the ion concentration. Providing that the fluctuations within each zone are not too great, it probably matters little at what level the averaging is done. But if the fluctuations are large, then any form of spatial averaging must introduce an element of arbitrariness. For example, most people might prefer to average $[Ca^{2+}]_i$ in nanomolar or micromolar and pH_i in pH units, but why not average $[Ca^{2+}]_i$ in logarithmic units, i.e. pCa, or $[H^+]$ in nanomolar?

9. Public presentation of results

We find the most cost-effective way of capturing still images for presentation at seminars or in publications is simply to put a 35 mm camera on a tripod and photograph the color monitor. An inexpensive telephoto lens has slight pincushion distortion, which counteracts the barrel distortion due to the curvature of the monitor screen. This costs only a few hundred dollars in equipment, compared to thousands for dedicated color hard copy devices. Direct slide film gives much better color contrast and fidelity than print film, partly because the machines that make prints from negatives get confused by artificial pseudocolors. For motion picture sequences, we have preferred a 16 mm animation-type cine camera rather than videotape or videodisk, since color video recording requires an RGB-to-NTSC or -PAL converter. Also, video presentations are particularly prone to embarrassing technical failures in front of a fidgeting audience, are hard to show in a large auditorium unless multiple monitors are hooked together, and suffer from transatlantic incompatibility of video formats. The main problem with a cine camera is the need to engineer the computer interface to bring sequential images onto the monitor and shoot them one by one, a process too tedious to be done by hand for thousands of frames.

10. Creation of montages

Still images often need to be juxtaposed for detailed comparison, for example to contrast appearances before, during, and after stimulation, or agonist vs. control responses, or mutant vs. wild-type cells, etc. When presenting such data in slides, a montage would often be preferable to separate isolated shots, since the latter would require flipping back and forth or relying on the audience's visual memory. Likewise in preparing figures for publication, a pre-formed montage would reduce the printer's labor and perhaps page charges, and ensure that the component images stay together and are printed under identical conditions. We initially tried to assemble such montages by cutting and pasting color prints, but this process was slow and

laborious, and the re-photography generally gave considerable further degradation in color fidelity. A better solution was to write a program to assemble and edit the montage in the image processor. Such a program should permit one to shrink or expand portions of input color or monochrome images, juxtapose them on an output screen, add annotation, and store the results for subsequent modification or replay.

11. Flexibility and modifiability

It should be obvious from the above discussion that systems need to be modifiable to meet individual needs, which often become clear only after considerable trial-and-error. The core sections of our system have been through over 200 rounds of revision. Space does not permit discussion of more specialized modifications, such as interfacing with photochemical activation experiments [16, 17, 26] or electrophysiology [12, 13], or use with non-ratioing Ca^{2+} indicators such as fluo-3 [16, 17], or Na^{+} measurements [11, 14], and yet more will surely continue to be wanted. Systems in which the source code is available and documented therefore are much better for the user than turn-key, black-box systems, which commercial developers prefer to generate but which constrain the user to those functions that the developer has foreseen. Much of the motivation for writing this article has been our frustration at repeatedly being asked to recommend a commercial system for ratio imaging when all the ones that we have seen lack practical features that we have found to be essential. Perhaps with guidance from other experienced users, fully satisfactory systems will evolve, so that we biologists can go back to doing biology.

Acknowledgements

R.Y.T. would like to acknowledge that his first exposures to practical image processing and to modern menu-driven real-time interfaces occurred on a brief visit in 1984 to the laboratories of F.S. Fay (University of Massachusetts) and T.M. Chused (NIH), invaluable experiences that significantly influenced the design of our system. This work was supported by NIH (GM31004, EY04372, and T32 CA 09941) and the Searle Scholars Program (83-K-111).

References

1. Tsien RY. Poenie M. (1986) Fluorescence ratio imaging: a new window into intracellular ionic signaling. Trends Biochem. Sci., 11, 450-455.
2. Poenie M. Alderton J. Steinhardt R. Tsien RY. (1986) Calcium rises abruptly and briefly throughout the cell at the onset of anaphase. Science, 233, 886-889.
3. Hallam TJ. Poenie M. Tsien RY. (1986) Homogeneity of ADP- and thrombin-stimulated rises in $[Ca^{2+}]_i$ in fura-2 loaded human platelet populations revealed by fluorescence ratio image processing. J. Physiol., 377, 123P.
4. Paradiso AM. Tsien RY. Machen TE. (1987) Digital image processing of intracellular pH in gastric oxyntic and chief cells. Nature, 325, 447-450.
5. Schmitt-Verhulst AM. Guizemanes A. Boyer C. et al. (1987) Pleiotropic loss of activation pathways in a T-cell receptor alpha-chain deletion variant of a cytolytic T-cell clone. Nature, 325, 628-631.
6. Poenie M. Tsien RY. Schmitt-Verhulst AM. (1987) Sequential activation and lethal hit measured by $[Ca^{2+}]_i$ in individual cytolytic T cells and targets. EMBO J., 6, 2223-2232.
7. Wilson HA. Greenblatt D. Poenie N. Finkelman FD. Tsien RY. (1987) Crosslinkage of B lymphocyte surface immunoglobulin by anti-Ig or antigen induces prolonged oscillation of intracellular ionized calcium. J. Exp. Med., 166, 601-606.
8. Chen WS. Lazar CS. Poenie M. Tsien RY. Gill GN. Rosenfeld MG. (1987) Requirement for intrinsic protein tyrosine kinase in the immediate and late actions of the EGF receptor. Nature, 328, 820-823.
9. Havran WL. Poenie M. Kimura J. Tsien RY. Weiss A. Allison JP. (1987) Expression and function of the CD3-antigen receptor on murine $CD4^{+}8^{+}$ thymocytes. Nature, 330, 170-173.
10. Ambler SK. Poenie M. Tsien RY. Taylor P. (1988) Agonist-stimulated oscillations and cycling of intracellular free calcium in individual cultured muscle cells. J. Biol. Chem., 263, 1952-1959.
11. Harootunian AT. Eckert B. Minta A. Tsien RY. (1988) Ratio imaging using newly developed fluorescent sodium indicators in rat embryo fibroblasts. FASEB J., 2, A728.
12. Lipscombe D. Madison DV. Poenie M. Reuter H. Tsien RY. Tsien RW. (1988) Spatial distribution of calcium channels and cytosolic calcium transients in growth cones and cell bodies of sympathetic neurons. Proc. Natl. Acad. Sci. USA, 85, 2398-2402.
13. Lipscombe D. Madison DV. Poenie M. Reuter H. Tsien RW. Tsien RY. (1988) Imaging of cytosolic Ca^{2+} transients arising from Ca^{2+} stores and Ca^{2+} channels in sympathetic neurons. Neuron, 1, 355-365.
14. Negulescu PA. Harootunian AT. Minta A. Tsien RY. Machen TE. (1988) Intracellular sodium regulation in rabbit gastric glands determined using a fluorescent sodium indicator. J. Gen. Physiol., 92, 26a.
15. Tsien RY. (1988) Fluorescence measurement and photochemical manipulation of cytosolic free calcium. Trends Neurosci., 11, 419-424.
16. Harootunian AT. Kao JPY. Tsien RY. (1988)

Agonist-induced calcium oscillations in depolarized fibroblasts and their manipulation by photoreleased Ins(1,4,5)P_3, Ca^{2+}, and Ca^{2+} buffer. Cold Spring Harbor Symp. Quant. Biol., 53, In press.
17. Kao JPY. Harootunian AT. Tsien RY. (1989) Photochemically generated cytosolic calcium pulses and their detection by fluo-3. J. Biol. Chem., Submitted.
18. Kawanishi T. Blank L. Harootunian AT. Smith MT. Tsien RY. (1989) Calcium oscillations induced by hormonal stimulation of individual fura-2 loaded hepatocytes. J. Biol. Chem., Submitted.
19. Grynkiewicz G. Poenie M. Tsien RY. (1985) A new generation of Ca^{2+} indicators with greatly improved fluorescence properties. J. Biol. Chem., 260, 3440-3450.
20. Cornsweet TN. (1970) Visual perception, Chapter X especially. New York, Academic Press, pp. 243-249.
21. Tanasugarn L. McNeil P. Reynolds GT. Taylor DL. (1984) Microspectrofluorometry by digital image processing: measurement of cytoplasmic pH. J. Cell Biol., 98, 717-724.
22. Tsien RY. Rink TJ. Poenie M. (1985) Measurement of free Ca^{2+} in individual small cells using fluorescence microscopy with dual excitation wavelengths. Cell Calcium, 6, 145-157.
23. Tsien RY. (1989) Fluorescent probes of cell signaling. Annu. Rev. Neurobiol., 12, 227-253.
24. Tsien RY. (1989) Fluorescent indicators of ion concentrations. Methods Cell Biol., 30, 127-156.
25. Foskett JK. (1988) Simultaneous Nomarski and fluorescence imaging during video microscopy of cells. Am. J. Physiol., 255, C566-C571.
26. Gurney AM. Lester HA. (1987) Light-flash physiology with synthetic photosensitive compounds. Physiol. Rev., 67, 583-617.
27. Kurtz I. Dwelle R. Katzka P. (1987) Rapid scanning fluorescence spectroscopy using an acousto-optic tunable filter. Rev. Sci. Instrum., 58, 1996-2003.
28. Becker PL. Fay FS. (1987) Photobleaching of fura-2 and its effect on determination of calcium concentrations. Am. J. Physiol., 253, C613-C618.
29. Bright GR. Taylor DL. (1986) Imaging at low light level in fluorescence microscopy. In: Taylor DL. Waggoner AS. Murphy RF. Lanni F. Birge RR eds. Applications of Fluorescence in the Biomedical Sciences. New York, Alan R. Liss, pp. 257-288.
30. Connor JA. (1986) Digital imaging of free calcium changes and of spatial gradients in growing processes in single, mammalian central nervous system cells. Proc. Natl. Acad. Sci. USA, 83, 6179-6183.
31. Hiraoka Y. Sedat JW. Agard DA. (1987) The use of a charge-coupled device for quantitative optical microscopy of biological structures. Science, 238, 36-41.
32. Spring KR. Smith PD. (1988) Illumination and detection systems for quantitative fluorescence microscopy. J. Microsc., 147, 265-278.
33. Livingstone MS. (1988) Art, illusion, and the visual system. Sci. Am., 258, 78-85.

Please send reprint requests to Dr R.Y. Tsien, Department of Physiology-Anatomy, University of California, Berkeley, CA 94720, USA