Available online at www.sciencedirect.com
SCIENCE DIRECT®
Clinical Immunology 110 (2004) 277–283
CLINICAL IMMUNOLOGY
www.elsevier.com/locate/yclim

New approaches to fluorescence compensation and visualization of FACS data

James W. Tung,* David R. Parks, Wayne A. Moore, Leonard A. Herzenberg, and Leonore A. Herzenberg

Department of Genetics, Stanford University Medical School, Stanford, CA 94305-5120, USA

Received 7 November 2003; accepted with revision 12 November 2003

Abstract

The Fluorescence Activated Cell Sorter (FACS) is an invaluable tool for clinicians and researchers alike in phenotyping and sorting individual cells. With the advances in FACS methodology, notably intracellular staining for cytokines, transcription factors and phosphoproteins, and with increases in the number of fluorescence detection channels, researchers now have the opportunity to study individual cells in far greater detail than previously possible. In this chapter, we discuss High-Definition (Hi-D) FACS methods that can improve analysis of lymphocyte subsets in mouse and man. We focus on the reasons why fluorescence compensation, which is necessary to correct for spectral overlap between two or more fluorochromes used in the same staining combination, is best done as a computed transformation rather than using the analog circuitry available on many flow cytometers. In addition, we introduce a new data visualization method that scales the axes on histograms and two-dimensional contour (or dot) plots to enable visualization of signals from all cells, including those that have minimal fluorescence values and are not properly represented with traditional logarithmic axes. This "Logicle" visualization method, we show, provides superior representations of compensated data and makes correctly compensated data look correct. Finally, we discuss controls that facilitate recognition of boundaries between positive and negative subsets.

Keywords: FACS; Flow cytometry; High-Definition FACS; Hi-D FACS; Spectral overlap; Fluorescence compensation; Fluorescence-minus-one (FMO); Logicle visualization; Quantile contour plot

Introduction

The Fluorescence Activated Cell Sorter (FACS) is designed and developed to measure the expression of sets of genes on individual cells and sort cell populations of interest based on the expression of these genes. FACS and other flow-cytometry instruments have become invaluable in identifying lymphocyte subsets in human, mouse, and other species. Several advances have contributed to the expanded use of FACS instruments. New FACS methodology has been developed for staining surface markers and intracellular molecules (cytokines, transcription factors, phosphoproteins [1,2]). Multiple fluorochromes with distinct excitation and emission properties have become available, and High-Definition (Hi-D) FACS instruments have been developed to simultaneously measure fluorescence signals from 11 (or more) fluorochromes and to additionally measure light scatter signals that provide and index of cell size and granularity [3,4].

Because multiple fluorochromes are used with Hi-D FACS instruments, it is necessary to correct for spectral overlap among the fluorochromes when interpreting the signals obtained, that is, the FACS data. The need for this correction, commonly called fluorescence compensation, was recognized in the analysis of early single laser two-color FACS experiments in the mid-1970s [5] and is a central step in FACS data analysis. In the FACS, a fluorochrome associated with a cell is excited by a laser focused on the cell stream. The fluorescence emitted by the fluorochrome passes through a bandpass filter (which only allows certain wavelengths to pass through) and is collected by the detector assigned to that fluorochrome. However, since the emission spectrum of each fluorochrome is not necessarily restricted to its assigned channel (bandpass filter plus detector), some of the fluorescence it emits may be collected by channels assigned to other fluorochromes. In multilaser systems, nonzero excitation of a dye by a laser other than the one intended for that dye

* Corresponding author.
E-mail address: Tung@darwin.stanford.edu (J.W. Tung).

1521-6616/$ - see front matter © 2004 Elsevier Inc. All rights reserved.
doi:10.1016/j.clim.2003.11.016

similarly results in signals on detectors assigned to other dyes (''cross-laser'' signals). These ''spectral overlap'' signals, which are proportional to the fluorescence signal the fluorochrome gives in its assigned channel, must be subtracted to accurately evaluate the fluorescence signal of each fluorochrome in its assigned channel. In a properly compensated sample, cells with no dye corresponding to a particular signal channel will appear in that channel distributed symmetrically around their mean autofluorescence value.

Fluorescence compensation, that is, the process of subtracting the spectral overlap signals from the overall fluorescence detected in each channel, can be accomplished either by adjusting the analog electronic compensation circuitry present on most older flow cytometry instruments or by employing a computed transformation on the primary measurement data. Until recently, most laboratories, including our own, used analog electronic compensation during data collection. However, we and several other laboratories have now shifted to routinely collecting uncompensated data when doing FACS analyses and correcting for spectral overlap (compensating the data) before analysis with software that has a compensation utility (e.g., FlowJo, TreeStar, Inc, San Mateo, Calif.). We will discuss this more fully later in this chapter.

Data resulting from computed compensation often includes cell populations that range into very low and negative values in some dye dimensions. Such populations cannot be represented adequately in a logarithmic display, so we developed a new method for defining plot axes that allow such populations to be displayed well without losing the advantages of log display for other populations. The axes constructed with this ''Logicle'' visualization method approximate logarithmic scaling at higher signal levels but include an approximately linear region that extends through zero to negative values and allows visualization of cells whose compensated signals fall near or below zero.

In addition, we discuss methods for determining the fluorescence boundaries separating negative cells from those with small amounts of valid fluorescence signal. This problem, which complicates many analyses, must be solved by adding appropriate controls to identify boundaries for individual subsets of cells in the analyzed sample. Thus, in this chapter, we outline the full set of controls necessary for Hi-D FACS studies and discuss how these controls can be used effectively.

FACS data collection and fluorescence compensation

For many years, FACS data collection was typically based on the use of analog (hardware) ''compensation'' to correct for overlap in the fluorescence emission spectra of fluorochromes excited by the same laser [5]. However, these analog compensation electronics introduce nonlinear systematic errors into the evaluations. Furthermore, they cannot resolve between-laser (excitation) spectral overlaps.

These problems can be avoided by collecting uncompensated data and using computed, rather than analog, compensation to correct for spectral overlap. Computed compensation is available in real time on some newer instruments and is supported for offline analysis by most current cytometry software (with varying levels of user friendliness). However, most instruments still use analog logarithmic amplifiers in their measurement circuitry. While computed compensation can be done with these instruments, the accuracy of the computed compensation is dependent on the accuracy of the log amp calibration, which is difficult to determine. Therefore, the most reliable method is to use high dynamic range linear digital electronics with computed compensation. When digital electronics are not available, the preferred method is to collect uncompensated data and use off-line computed compensation. The logical data visualization tool (see below) provides improved displays of data resulting from computed compensation that lend themselves to correct interpretations.

Compensation computations, the spectral matrix, and dye signal estimates

The process of fluorescence compensation transforms a set of color signal measurements on each cell into a set of dye signal estimates. This can be best represented as multiplication of color signal vectors by a compensation matrix to yield dye signal vectors. The matrix is obtained by analyzing single-stained compensation control samples.

In practice, a population of labeled cells is identified in each single stain sample, and a corresponding negative population is identified. The difference in mean signal of the positive and negative populations is evaluated in each of the fluorescence measurement channels for each single stain. For a particular dye, the differences in all fluorescence channels normalized to the difference on the measurement channel optimized for that dye constitute a row in what may be called the ''spectral matrix''. Taken together, the elements of the spectral matrix express how each detector responds to each dye.

In a four-dye case, the relationship between dyes and signals can be expressed as

$$\begin{pmatrix} D_1 \\ D_2 \\ D_3 \\ D_4 \end{pmatrix} \begin{pmatrix} 1 & M_{12} & M_{13} & M_{14} \\ M_{21} & 1 & M_{23} & M_{24} \\ M_{31} & M_{32} & 1 & M_{34} \\ M_{41} & M_{42} & M_{43} & 1 \end{pmatrix} = \begin{pmatrix} S_1 \\ S_2 \\ S_3 \\ S_4 \end{pmatrix}$$

or $D * M = S$, where D_1 to D_4 represent amounts of the four dyes, M is the spectral matrix, and S_1 to S_4 are the expected signals on the four detectors.

Of course, in fluorescence compensation, we have the observed signals (S) and need to obtain estimates for the amounts of the dyes (D). In this case, we want

$$D = S^{*}M^{-1}$$

where M^{-1} is the inverse of the spectral matrix, which has been called the ''compensation matrix''. Since computers have no problem inverting matrices, we need only concern ourselves with the (noninverted) spectral matrix.

Since the accuracy of each measured signal is limited at least by statistical uncertainty in the number of photons detected, dye signal estimates have corresponding statistical uncertainties that depend not only on the signal from the dye of interest but also on the signals from other dyes with significant spectral overlaps. A cell population that is essentially negative for a particular dye label but which has positive staining with one or more spectrally overlapping dyes will have near zero (compensated) mean for the dye of interest. However, statistical variation in correcting for the other dyes may result in considerable spread in dye estimates for individual cells so that some cells (no more than 50%) may have negative values for the dye estimate. There is, of course, no such thing as negative dye, but the negative estimates serve to balance elevated positive estimates so that the population mean is correct.

When computed compensation is applied to data from instruments with analog logarithmic amplifiers, the accuracy of the results is critically dependent on how faithfully the computed scaling corresponds to the actual log amp behavior. The construction of the spectral matrix from the compensation control samples assures that compensation will be correct for uncompensated signal levels corresponding to those of the compensation control samples, but systematic over- or undercompensation will occur for signals much higher or lower than the control levels to the extent that the computed scaling is inaccurate. Analog compensation avoids this sensitivity by operating on signals before they go to the logarithmic amplifier, but the typical analog compensation systems become nonlinear as full compensation is approached and generally lead to some level of systematic overcompensation.

Compensation controls

For computed compensation, data must be collected from single-stained ''compensation'' samples and processed to derive the appropriate spectral matrix. For some fluorochromes, for example, fluorescein, one single-stained sample is sufficient to enable compensation of all samples stained with that fluorochrome. For other fluorochromes, particularly tandem dyes such as Cy7PE, Cy5PE, Cy7APC, etc., one single-stained sample for each lot of the fluorochrome is needed since the spectra of individual lots may vary (see Fig. 1).

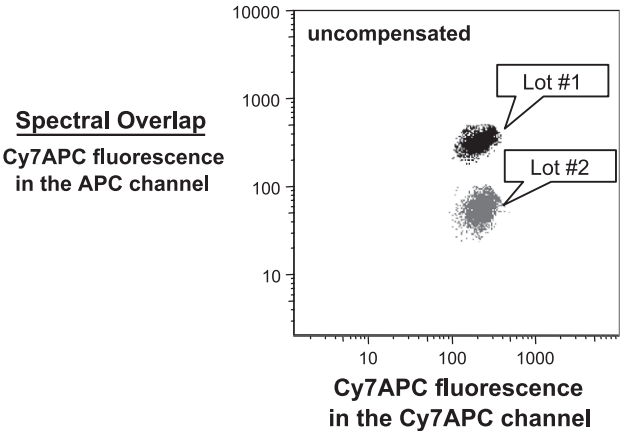

Fig. 1. Different lots of tandem dyes have different fluorescence properties. Differences in the spectral overlap in the APC channel of two different lots of Cy7APC conjugated to anti-IgM are shown in the figure. The differences mean that the compensation settings for the two lots differ and hence that a compensation control must be run for each if the two lots are used in the same experiment.

Examining the uncompensated data in the channel assigned to the dye and the channel detecting spectral overlap for a set of dyes can help to determine whether separate compensation controls are necessary. If the location of the ''diagonal'' obtained in this way for two or more antibodies is the same, a single control may be usable for those antibodies. The real test, however, is to examine the spectral matrix elements obtained with each single stain control. If these are essentially identical for two or more reagents, a single compensation control can be used for all of them. The conservative and usually appropriate approach is to use a compensation control sample for each different reagent used in an experiment.

FACS sorting and fluorescence compensation

Analog hardware compensation must still be used for most multicolor sorting on instruments that do not have real-time computed compensation. The compensation controls used to do the analog compensation should be the same as those used for computed compensation. It is important to collect data for all test and control samples at the same hardware settings. As a general reminder, reanalysis of some of the sorted cells under the same instrument conditions used in the sort is a key control in almost any cell sorting experiment.

Visualization of FACS data: contour plots and dot plots

Although dot plots are still used by many immunologists, they are inherently poor for displaying FACS data. This is because they are highly sensitive to local saturation. If two cells fall at the same location on the grid used for the plot, they will appear as a single dot. In fact, if 100 cells fall at

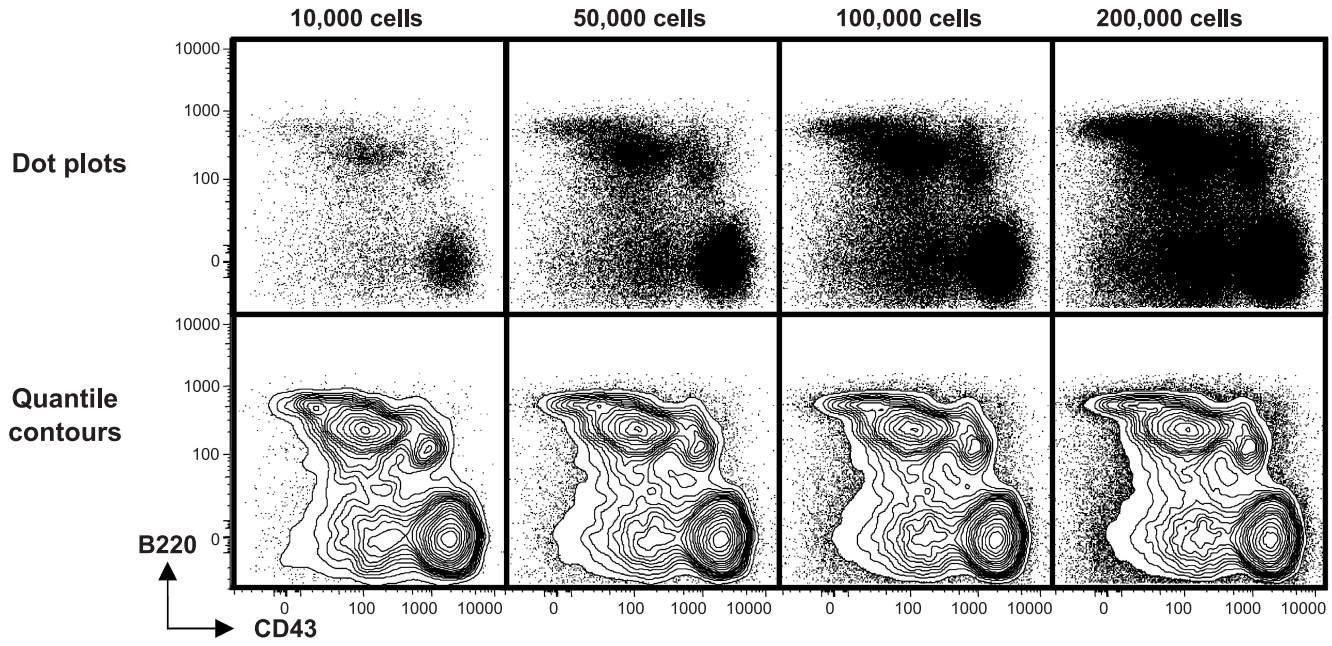

Fig. 2. Impressions of data vary depending on the number of cells displayed in dot plots. Dot plots tend to lose visible structure when cell numbers are too high or too low. Cell subsets that are clearly distinguishable on one dot plot may not be resolved on another that displays higher or lower numbers of cells from the same analysis. In contrast, contour plots (with outliers) maintain the discrimination of subsets regardless of cell numbers.

that location, only one dot will be placed on the plot. Furthermore, when many cells fall close to one another, the dots representing them become contiguous, resulting in a large black area with no local geography to indicate locations of peaks within the area. Thus, as the number of cells analyzed increases, the plots become progressively more saturated and features of the data landscape are obscured (see Fig. 2).

In contrast, the quantile contour plots that are available in most FACS analysis software packages barely change as cell numbers increase over a 20-fold range (see Fig. 2). The contours in these plots are computed so that an equal percentage (usually 5%) of the cells are bounded by each contour. Thus, the contour levels on the plots provide a good indication of the frequency of cells in given region. Contour plots, however, are poor for depicting very low density regions and are uninformative about the location of the last quantile, typically 5% of cells (the outliers). Therefore, the best data representation can usually be obtained with a combination plot in which the standard quantile contour plot is augmented by visualization of the outliers as dots on the plot. For data including narrow features that are not well represented in contour plots, color dot displays in which color is used to indicate the extent of event pileup can be quite effective.

Visualization of FACS data: Logicle axes

The contour and dot plots that are currently used by most laboratories have standard four-decade logarithmic axes that provide a wide dynamic range for display of FACS data. However, the absence of a zero point and negative values on these logarithmic axes introduces major problems, particularly for visualizing cells with little or no associated fluorescence. This interferes with visualizing compensated data, since the subtraction of spectral overlap during compensation is designed to return cells with no associated fluorochrome to background values. Statistical variation in the number of photoelectrons detected typically results in ‘‘negative’’ cell populations with more spread in compensated data values than would be observed for the same set of cells completely unstained. In such circumstances, some cells commonly receive negative data values that are simply part of the overall distribution for the population. If compensation values are appropriately set, compensated data values for a cell population that is negative for a particular dye can be expected to distribute symmetrically around a low value representing the autofluorescence of the cells in that dye dimension. Logarithmic displays, however, cannot accommodate zero or negative values. This situation can be understood as follows: on a logarithmic scale, all values below the lowest decade must either be discarded (not acceptable) or ‘‘piled up’’ at the lowest point on the scale. The pile-up obscures the true center of the compensated distribution. Furthermore, it often breaks the distribution artificially into what appears to be two subsets, one centered on the pile-up (the lowest point on the scale) and the other centered higher than the true center of the compensated population (see FITC-positive cells in Fig. 3, center panel). This data display artifact often results either in misinterpretation of the higher ‘‘population’’ as a weakly positive

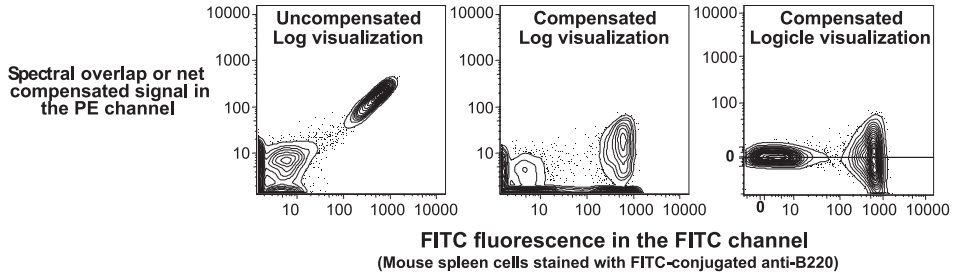

Fig. 3. Logicle visualization facilitates validation of compensation. Mouse spleen cells singly stained with B220-FITC are shown in two-dimensional plots (FITC vs. PE) on all panels. Roughly half of the cells in the spleen are B220-positive (B cells). Left panel: uncompensated data showing the spectral overlap of FITC fluorescence in the PE channel. Middle panel: correctly compensated data displayed with standard logarithmic axes. Note that an indeterminate number of cells in both the B220-positive and B220-negative subsets are "piled up" on the axes. In addition, note that the peak in the display of the FITC-positive cells in the PE channel does not match that of the FITC-negative subset. Right panel: compensated data displayed with Logicle visualization. Note that both FITC-positive and FITC-negative subsets are distributed symmetrically around a near-zero level in the PE channel.

subset or in serious overcompensation of the entire data set due to attempting to force this "population" down to the axis.

The Logicle data display (David Parks, Wayne Moore et al., manuscript in preparation) addresses these problems by enabling visualization of FACS data on mathematically defined axes that are asymptotically linear in the region just above and below zero and asymptotically logarithmic at higher (positive and negative) values. Thus, compensated values that fall either above or below zero can be correctly displayed. Note that Logicle visualization does not change the data. It merely allows lower data values to be properly represented and allows peaks in the region around zero to be in their proper position.

Fig. 3 illustrates how the Logicle display makes it easy to confirm the accuracy of fluorescence compensation. This figure shows data for a cell sample stained only with an FITC reagent. This stain divides the cell sample into two subsets. One subset is not stained by the FITC reagent while the other has a high FITC signal with significant spectral overlap detected on the PE channel (Fig. 3, left panel). In a properly compensated sample involving only PE and FITC staining, the spectral overlap will be subtracted from the fluorescence collected on the PE channel, and the signals for all popula-

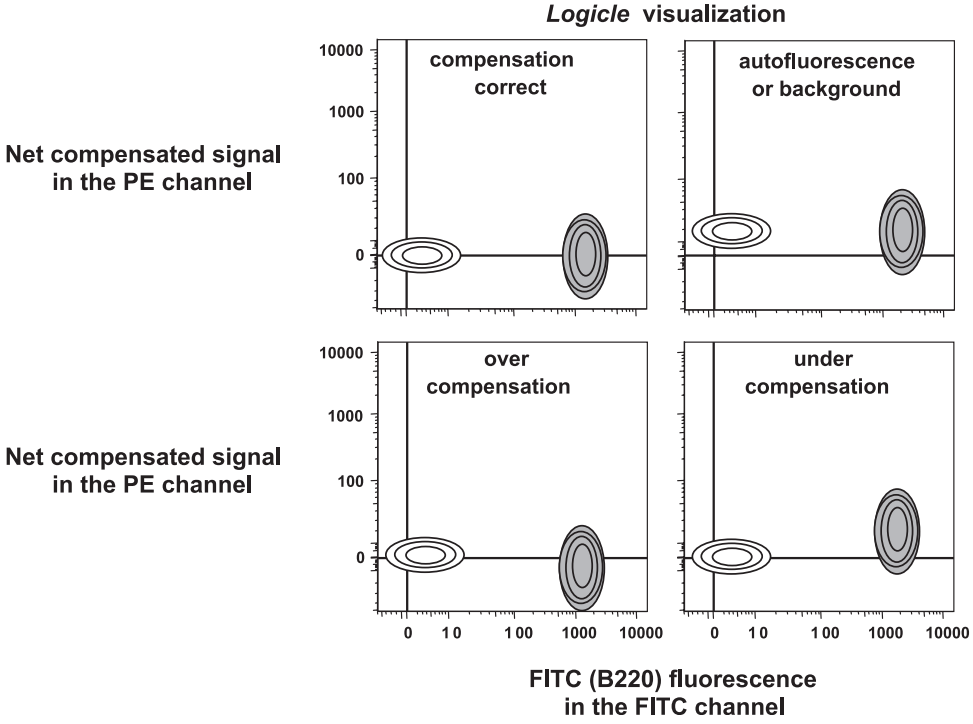

Fig. 4. Diagrammatic representation of compensation status of subsets viewed with Logicle axes essentially linear for small positive or negative data values in the PE dimension. Upper left panel: correctly compensated cells. Upper right panel: correctly compensated cells with high intrinsic autofluorescence or background in the PE channel. Lower left panel: overcompensation of the FITC-positive cells (too much fluorescence subtraction). Lower right panel: undercompensation of the FITC-positive cells (too little fluorescence subtraction).

tions on the PE channel will be distributed symmetrically around the autofluorescence value for the cells in the sample (Fig. 3, right panel). When multiple fluorochromes are involved, the compensation calculations are more complex, but the end result is the same: the spectral overlaps are corrected and the distribution representing cells that do not bind the fluorochrome detected in a given channel winds up in a peak centered on their mean autofluorescence value.

The diagram in Fig. 4 shows the expected Logicle plots for cells that are properly compensated, overcompensated, undercompensated, or autofluorescent. Note that overcompensation drives the peak for the FITC-positive population below the mean autofluorescence in the PE channel while undercompensation fails to bring this population to equivalence with the FITC-negative population. For cells that are equally autofluorescent in the PE channel, both the FITC-positive and the FITC-negative cells will be distributed symmetrically around the mean PE channel autofluorescence value.

Evaluating background fluorescence: fluorescence-minus-one controls

Spectral overlap and the statistics of fluorescence compensation can differentially affect subsets of cells in that different subsets in a multicolor stain may have different levels of autofluorescence and typically show different amounts of variation in compensated background signal in any given detection channel. Therefore, it is often difficult to establish a reliable boundary for distinguishing cells that express low levels of a determinant from cells that do not express detectable levels of the determinant. In essence, this requires determination of a threshold that distinguishes fluorescence due to the presence of a particular fluorochrome from background fluorescence, which in multiparameter and Hi-D FACS studies is a composite of autofluorescence and variation due to statistical uncertainty in correction of spectral overlaps that will in general differ among different subsets of cells. Thus, the necessity to identify the subsets by staining with multiple fluorochrome-coupled reagents introduces major problems in identifying a boundary separating positive from negative cells in individual subsets. In particular, it is important to recognize that distribution of compensated background values for a given subset becomes broader as the extent of spectral overlap that must be corrected by compensation increases. For example, consider a sample that contains two subsets of cells, one that expresses a determinant detected by an FITC-coupled reagent and another that does not. After compensation, both subsets will be distributed symmetrically around a central value close to zero in the PE channel. However, because of the statistics of the compensation subtraction process, the FITC-negative subset will have a substantially smaller spread in the PE channel than the FITC-positive subset (see Fig. 3, right panel).

In multiparameter and Hi-D FACS experiments, the situation is further complicated because the expression of multiple markers by individual subsets necessitates simultaneous compensation for several fluorescence signals. As a practical matter, this means that the best way to determine the boundary between no detectable expression and low expression of a particular determinant is to measure it directly in a sample stained with all of the reagents used

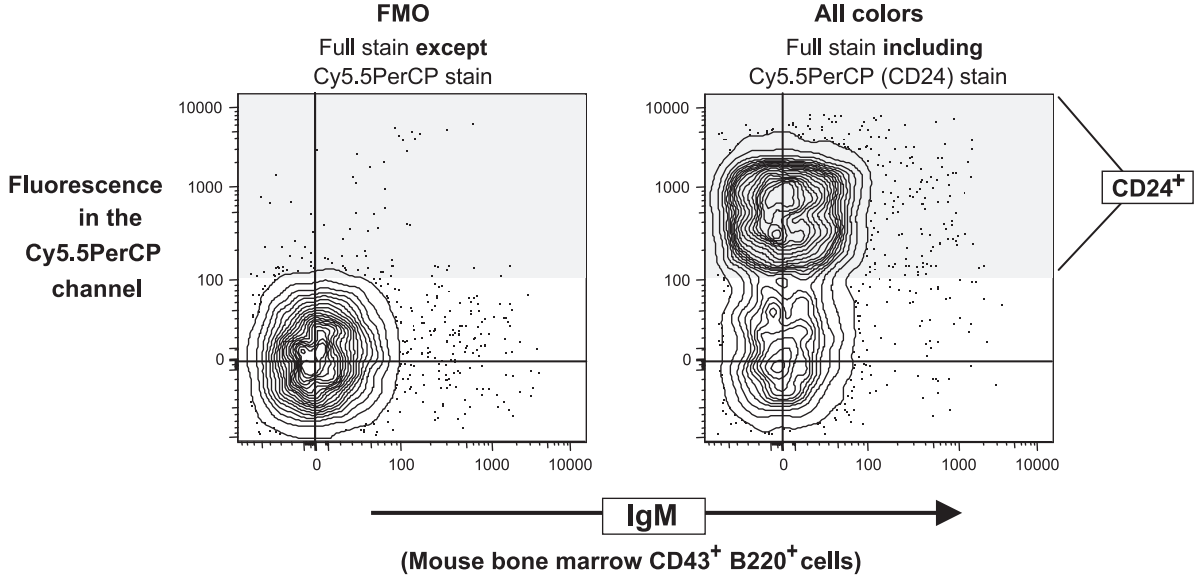

Fig. 5. Fluorescence-minus-one (FMO) staining and Logicle display reveal the boundary between positive and negative cells. Data are shown for live, size-gated B220-positive CD43-positive mouse bone marrow cells in an 11-color analysis. To determine the boundary for CD24-positive cells, the Cy5.5PerCP anti-CD24 stain was omitted from the FMO control and included in the full stain. The upper bound of the FMO fluorescence observed in the Cy5.5PerCP channel identifies the boundary separating the positive from the negative cells.

except for the one that registers in the channel of particular interest [3,6]. This ‘‘fluorescence-minus-one’’ (FMO) [6] control allows the subset to be gated and the background level for the subset to be determined (see Fig. 5). Under the conditions illustrated in Fig. 5, an unstained sample used as a negative control would indicate a positive signal threshold of 30–40 units (data not shown) instead of the proper 100-unit threshold derived from the FMO control.

Conclusion

In this chapter, we have introduced Logicle visualization methods for FACS data. We have demonstrated the use of these methods to provide displays that assist in correct data interpretation and provide clear visual confirmation of correct compensation. Improved displays are also helpful in establishing FMO boundaries that distinguish positive from negative cells. In addition, we have discussed the kinds of controls necessary for correctly setting compensation values and have underscored the advantages of collecting uncompensated FACS data and doing the compensation with software designed for the purpose. Although these considerations are central to doing good FACS work with Hi-D FACS instruments, they are relevant to all FACS studies and, if generally applied, should help improve the overall quality of FACS analysis and sorting.

Acknowledgments

We wish to thank Ometa Herman and Sue Sheppard for their technical assistance. We would like to acknowledge funding in part from the NIH in the form of Grant number EB000231 from the National Institute of Biomedical Imaging and Engineering. We would also like to thank John Mantovani for his administrative expertise.

References

[1] O.D. Perez, S. Kinoshita, Y. Hitoshi, D.G. Payan, T. Kitamura, G.P. Nolan, J.B. Lorens, Activation of the PKB/AKT pathway by ICAM-2, Immunity 16 (2002) 51–65.
[2] O.D. Perez, G.P. Nolan, Simultaneous measurement of multiple active kinase states using polychromatic flow cytometry, Nat. Biotechnol. 20 (2002) 155–162.
[3] N. Baumgarth, M. Roederer, A practical approach to multicolor flow cytometry for immunophenotyping, J. Immunol. Methods 243 (2000) 77–97.
[4] S.C. De Rosa, M. Roederer, Eleven-color flow cytometry. A powerful tool for elucidation of the complex immune system, Clin. Lab. Med. 21 (vii) (2001) 697–712.
[5] M.R. Loken, D.R. Parks, L.A. Herzenberg, Two-color immunofluorescence using a fluorescence-activated cell sorter, J. Histochem. Cytochem. 25 (1977) 899–907.
[6] M. Roederer, Spectral compensation for flow cytometry: visualization artifacts, limitations, and caveats, Cytometry 45 (2001) 194–205.