The Observational Approach to Populations in Globular Clusters

By IVAN R. KING

Astronomy Dept., University of California, Berkeley, CA 92720-3411, USA

This introductory chapter discusses the observations on which our understanding of globular clusters lies. Successive sections deal with photometry, chemical abundances, the details of color–magnitude diagrams, the distance scale, luminosity and mass functions, and the lower end of the main sequence. An appendix treats the dynamical role of binaries in globular clusters.

Astronomy aims at an understanding of the facts and phenomena that we see, and the processes by which they came about—and in the best of possible cases, the recognition of why they had to be this way and could not have been otherwise. The first stage in this endeavor is to see what is there, and, to the extent that we can, how it became that way.

Globular clusters can in many ways be considered the crossroads of astronomy. They have played a central role in the unfolding of our astronomical understanding, to which they bring two singular advantages: first, each cluster (with a possible rare exception) is a single and specific stellar population, stars born at the same time, in the same place, out of the same material, and differing only in the rate at which each star has evolved. Such a group is much easier to study than the hodgepodge that makes up the field stars of the Milky Way. Second, globular clusters are made up of nearly the oldest—perhaps the very oldest—stars of the Universe, and as such they give us an unparalleled opportunity to probe the depths of time that are the remotest to reach.

This chapter introduces the facts about globular clusters—occasionally relying also on well-established pieces of understanding to help organize the facts and hold them together, but concentrating on the facts. These facts are presented in the context of the observations, and indeed the specific processes by which the observations have been gathered, and continue, in this vitally living science, to be gathered. It is in that spirit that this chapter is presented.

1. Photometry in Globular Clusters

1.1. *The evolution of techniques*

It has been known for a very long time that globular clusters (GCs) have color–magnitude diagrams (CMDs) that differ in appearance from the traditional diagram that was first drawn by Hertzsprung and by Russell. (It is of course this difference that was central to Baade's [1944] original breakthrough in defining stellar populations.) Even back in the 1930's and 40's there were a few good CMDs of GCs (e.g., Greenstein 1939, Hachenberg 1939), but the creation of CMDs did not become a real industry until the 1950's; it may be said that its start was centered on Sandage's thesis work on M3 (Sandage 1954), which was among the first to show the main-sequence turnoff clearly. This was of course a cornerstone of the physical understanding of stellar evolution that was awakening at that time, and that led so quickly to the level of the Vatican conference of 1957 on stellar populations (O'Connell 1958).

What had held the field back, and what was still slowing it down, was the cumbersome and inaccurate method that was needed in order to derive good stellar magnitudes on

photographic plates. The basic problem is that the response of the photographic plate is non-linear and difficult to calibrate, whereas magnitudes require such a calibration. It had helped a great deal when photoelectric photometers were introduced. Their response is linear, and could be used to set up sequences of stars of known magnitude. For decades, however, photoelectric devices had been almost the exclusive province of those few who could build them and keep them working—pre-eminently Stebbins and Whitford (whose work was too voluminous for specific references here). Then, after the end of the Second World War, photomultiplier tubes became available, and it was then possible for the ordinary astronomer to carry out good photographic photometry—but even so, by a technique that to us seems almost incredibly laborious. One would first use the photoelectric photometer, with its linear response, to set up a sequence of standard stars of known relative magnitude. Then one would put each photographic plate, with its notoriously non-linear response, into some sort of device where the stars could be pointed at one by one, to measure the sizes of their images. Relative magnitudes were established by using the photoelectric standards to calibrate the curve of star-image size against magnitude. Finally, on some perfectly clear ("photometric") night, some of the standards would be compared with one of the few internationally recognized magnitude sequences in the sky, so that a zero point could be put on the photographically measured magnitudes. No wonder progress was slow.

What makes photoelectric photometry itself slow and laborious is that it observes only one star at a time. What was needed was an imaging photometric device that had a linear response. A number of such were developed and tried, but the modern era of photometry in globular clusters can be said to have begun with the wide availability of charge-coupled devices (CCDs), after about 1980.

The CCD has three great advantages over the photographic plate. First, obviously, is its linearity, which allows the immediate determination of relative magnitudes of all the stars in the frame, with only a simple zero-point comparison needed in order to make the job complete. Second, the CCD has a much higher quantum efficiency; it registers photons 10 to 100 times as well as a photographic plate. The third advantage is much less obvious; it arises from the fact that the photographic image of a star consists of a very limited number of grains of metallic silver, in the small area covered by the seeing disk of the star. The number of grains of silver limits the S/N ratio of a faint star on top of the sky background; by contrast, a CCD can accumulate a large number of photons, even in this limited area, and thus detect fainter stars. Their greater sensitivity indeed makes CCDs *faster*, but it is this escape from discrete grains that allows them to go fainter.

One other advantage of electronic registration of images is obvious but is worthy of note nevertheless. There is only one original photograph, but a CCD image can easily be copied by co-investigators without loss of fidelity.

In quantitative terms, the quantum efficiency (i.e., the fraction of incident photons that achieve a detectable effect) of the best photographic plates is about 1%, whereas a good CCD can reach a peak QE of 80% at its most sensitive wavelengths.

Only in one respect does the photographic plate retain an advantage: at a well-corrected telescope focus, it can cover a much larger area. This hardly matters for faint stars, where even a small CCD can gather thousands of stars in a single field of a globular cluster; but the smallness of CCD fields has definitely held back one area of globular-cluster research: the study of the stars in the upper part of the CMD. These stars are few in number, and therefore the study of adequate numbers of them requires a wide area—sometimes the whole cluster, which usually covers a considerable fraction of a square degree. While the present-day explorers of the lower main sequence have been glorying in fainter and

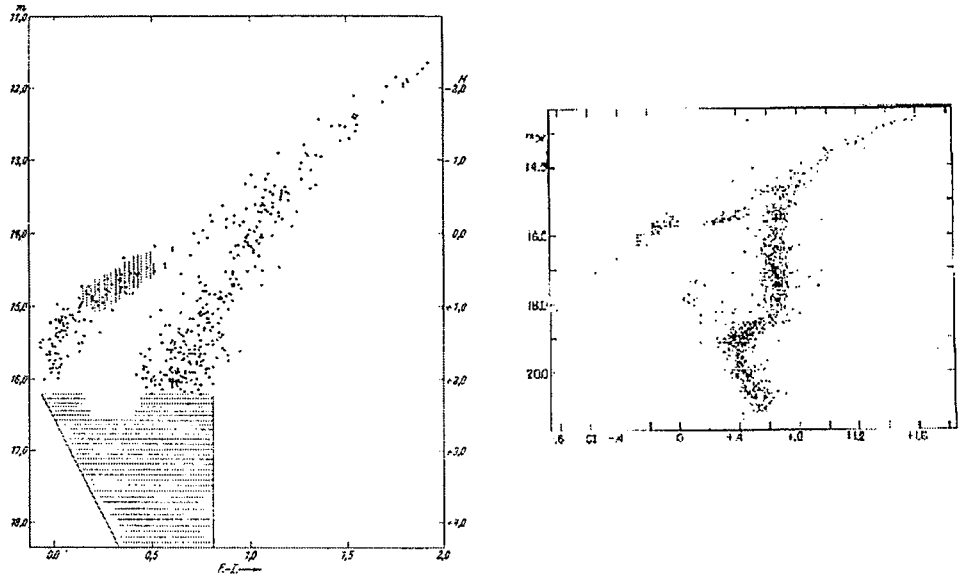

FIGURE 1. CMD of M92 by Hachenberg (1939, left) and of M3 by Sandage (1954, right), showing the progress after photoelectric calibration became relatively easy. The second diagram reaches 5 magnitudes fainter.

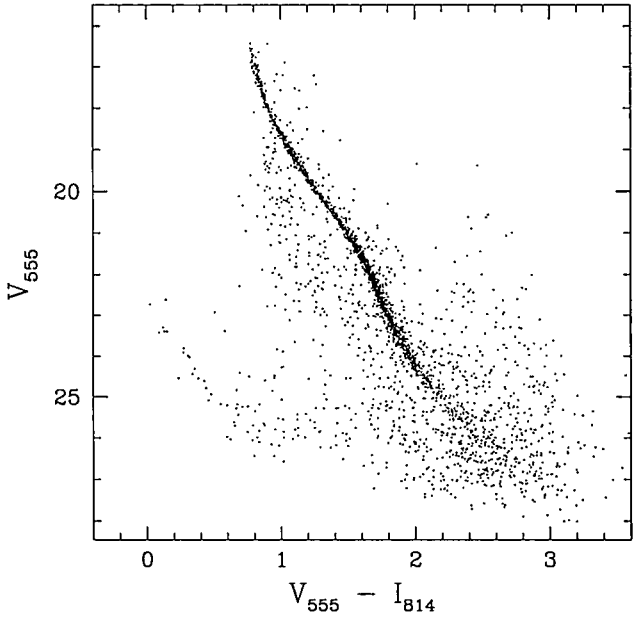

FIGURE 2. A recent CMD made from HST CCD imaging with the WFPC2 camera. The brightest data shown are at the main-sequence turnoff (since stars brighter than that were saturated even on the short exposures). The contrast with the previous figure shows how much difference CCDs and HST have made: another 6 magnitudes of faintness, and higher accuracy. (From Cool, Piotto, & King 1996.)

fainter magnitude limits, many theoreticians have been languishing for data at CCD accuracy for the evolved branches of the CMD. Only recently have a few observers begun to mount CCDs on smaller telescopes (for the larger field that they have) and take the multiple short exposures that are needed to measure the brighter stars over large regions of a cluster.

In specific numbers, the photographic foci of the Kitt Peak and Cerro Tololo 4-meter telescopes take photographs with a usable diameter of 20 cm. The earlier CCDs of the 1980's were postage-stamp size, registering fields a few hundred picture elements (pixels) across, covering only a few minutes of arc. CCDs are now larger, 2048 × 2048 formats being common now (giving fields of 10 or 15 arcmin), with even larger chips, and sometimes arrays of them, coming into use in a few places.

1.2. *The use of CCDs*

I will turn now to a brief description of CCDs themselves, from the point of view of the user, and will move finally to the particular problems of CCD photometry of globular clusters.

A CCD consists of an ultra-thin wafer of silicon, one side covered with a grid of electrodes that define its individual pixels. The thinness itself is a factor, in several ways. First, in most CCDs the light must pass through the silicon, which is not highly transparent at the shorter wavelengths. Hence CCDs achieve their highest quantum efficiency at red and near-infrared wavelengths, while they are less sensitive, or in many cases hopelessly insensitive, in the ultraviolet, or in many cases even in the blue. The second problem of thinned chips is uniformity of thickness, which is most difficult to achieve for the chips that are thinnest. The pixels in regions of the chip that are thinner have less silicon and therefore less sensitivity. Moreover, the transparency problem makes this effect wavelength-dependent. Some thin chips also fail to be completely flat, making the quality of star images vary over the field.

On the whole, however, non-uniformity tends to be less serious than lack of sensitivity at short wavelengths. Since non-uniformity is a problem for other reasons also, it is normal for reduction systems for CCD observations routinely to include procedures for correction of non-uniformity ("flatfielding" the image), whereas limited sensitivity at short wavelengths can be an immutable drawback.

Each pixel of a CCD has an electrode, and it is the arrangement of the electrodes that fixes the pixels, whose boundaries are rather well defined. Except for the "bleeding" effects referred to below, very few of the electrons released by the incident photons are counted into a pixel other than the one into whose discrete boundaries the photon fell.

The CCD accumulates charge in each of its pixels during an exposure that typically lasts 1 to 15 minutes; then it is read out. Readout consists of shifting the batches of accumulated charge down each column from row to row, and reading each successive row (simultaneously in each column) as it reaches the bottom. The readout process is imperfect in several ways, however. First, a pixel that has an outstandingly large number of counts (from a bright star) will affect neighboring rows in its column, "bleeding" into them. Second, a CCD can suffer from inadequate "charge transfer efficiency," in which the readout process loses electrons during the shifting. Finally, the emerging electrons are not counted individually, but go into an analog-to-digital converter, which introduces noise into the counts. This "readout noise" is quoted as an r-m-s value, so that a readout noise level of 5 means, for example, that on the average each pixel has 25 noise electrons added to the output of the readout. (Readout noise is almost always more serious than thermal dark current.) Given this noise level, it would be wasteful of data space to record the exact number of electrons; instead each number of electrons that comes from

the A/D converter is divided by a reduction factor called the DN (digital number). In the above case the readout might be set to record an output number that is only 1/10 of the actual number of electrons ("DN = 10").

The use of a DN has the effect of increasing the dynamical range of a CCD. Although the "full-well" capacity of most CCDs is greater than 100,000, the data systems that record their outputs are usually limited to 15 or 16 bits (32,767 or 65,535 counts). Even so, the centers of bright stars saturate, and observers generally include short exposures on which they can be successfully measured.

Stellar magnitudes can be measured on CCD images (i.e., in the digital array that represents the image) by aperture photometry, a procedure in which one simply adds up the light in a small number of pixels around the center of each star image; but more accurate results are usually achieved by use of some procedure that involves fitting to each star image a correctly chosen multiple of the point spread function (PSF), which should be determined from a number of well-exposed stars in the same image. Thus the result of PSF-fitting, for each star, is a position and a brightness factor that fit the star.

Many ready-made routines are available for PSF-fitting. The most commonly used is DAOPHOT (Stetson 1987), which is available as one of the the many sets of tasks in the massive IRAF package that has been developed by the National Optical Astronomy Observatories, with some important additions from the Space Telescope Science Institute. A number of other packages are available; besides, many users have written their own software for stellar photometry, often to cope with particular situations that they have encountered.

1.3. *The role of HST*

Rather than digress on photometric methods, however, I prefer here to concentrate on the problems that are peculiar to globular-cluster photometry, and in particular to the results that I will be emphasizing later in this chapter, that is, photometry with the Wide Field/Planetary Camera 2 (WFPC2) of the Hubble Space Telescope (HST).

In globular-cluster work, HST has gone several magnitudes fainter than previous ground-based observations, and this has made a big difference to the scientific results. (In principle, the largest ground-based telescopes, Keck and VLT, are capable of going almost equally faint; but for various reasons they have not yet done so.)

HST gets its excellent faintness limit, relative to ground-based telescopes of equal or even greater aperture, mainly from its high resolving power. The faintness limit at which stars can be detected and measured is set not purely by the number of photons collected from the star, but rather by the ratio of the signal from the star to the noise of the sky background in the area that the star image covers. Although HST does in fact see a somewhat darker sky than the ground-based observer does (*much* darker in the infrared, in fact), the more important factor is that the HST star image is so small, compared with a ground-based seeing disk. Thus a star in an HST image competes with a very much smaller area of sky, and we are able to observe fainter stars. The difference between the limiting magnitudes of HST CMDs and those observed from the ground is, in fact, about 4 magnitudes.

There are some important practical problems in WFPC2 photometry of globular clusters, some of them peculiar to HST and others connected more with the nature of the clusters themselves.

It should be noted first that WFPC2 has four 800×800-pixel CCDs, which record adjacent pieces of sky that are re-imaged onto them by transfer optics, from the focal plane of HST. The re-imaging for one chip, PC1, is such that its pixel size is 0.045 arcsec, while the other three, WF2, WF3, and WF4, have a scale of 0.100 arcsec/pixel. The

image that HST forms of a star is comparable in size to a PC1 pixel. Since the Nyquist criterion states that we need at least two samples per resolution element, even PC1 is undersampled, and the WFs are grossly undersampled.

In spite of the undersampling, a PSF can be determined reasonably well, by combining measurements of many stars that are situated in different positions with respect to pixel boundaries. But applying this PSF to a single star image is a different matter; here the undersampling really hurts. The best cure for this is dithering—displacements of otherwise identical images by well-chosen fractions of a pixel; the average of photometry from a set of well-dithered images is more accurate than photometry from images that are identically placed.

It is beneficial for dithering to include, at the same time, displacements by a small integral number of pixels in addition to the fractional displacements. The reason for this is that the flat-field corrections for the chips of WFPC2 are known only on a rather smooth scale; individual pixels have sensitivities that may differ by a few per cent (equivalent to a few hundredths of a magnitude). Dithering by a few pixels randomizes these errors and gives a much better average.

In fact, the most accurate globular-cluster photometry that has been published (Cool, Piotto, & King 1996, Rubenstein & Bailyn 1997) relied on well-dithered images. (It is interesting to note that a globular-cluster CMD offers an automatic index of its accuracy: the color width of the main sequence. Intrinsic widths of GC main sequences are too small to be measured; hence the observed width reflects the photometric accuracy. Main sequences have been produced with color widths of only 0.02–0.03 mag.; cf. Fig. 2.)

1.4. *Crowding problems*

If there is one characteristic that most distinguishes globular-cluster photometry from photometry of any other objects, it is crowding, which hinders photometry in several ways. The most obvious effect of crowding is the partial overlapping of star images. There are two basically different ways of dealing with the interaction of neighboring images with each other. DAOPHOT (like some of its competitors) works simultaneously on a group of star images each of which partially overlaps at least one other member of the group. It makes a fit of the PSF to each star in turn, while holding each of the others fixed at a previously determined level; the process is iterated till it converges. A quite different approach is neighbor subtraction (Yanny et al. 1994). In this method, the magnitudes of all the stars are first measured by some method (e.g., DAOPHOT). Then they are all subtracted from the image, by subtracting at the location of the star a replica of the PSF scaled to fit the magnitude of that star. The final measurements are made by adding each star back into the image, temporarily—an operation which has the effect of restoring all of the pixels of that star to their original values—and measuring it alone, in surroundings from which all its neighbors have been removed.

Another effect of crowding is more subtle; it comes not from the visible neighbors but from the invisible ones. Part of the background that surrounds a star comes not from the sky or from scattered light of bright stars, but rather from the combined light of the stars that are below the threshold of detection. Since these are subject to random fluctuations of number, the background around a star is noisier than it would otherwise be. This interference makes the magnitudes of stars in crowded regions, even when a star appears well isolated, less accurate than they would be in regions that are less crowded.

Scattered light of bright stars was mentioned above. In a crowded cluster field there is so much scattered light from the halos of the images of the numerous bright stars that its sum is equivalent to an increase in the sky background, making it difficult or impossible to detect and measure the faint stars. At the dense center of a rich globular

cluster, even if the faint stars are resolvable the limiting magnitude may be considerably less faint—sometimes by magnitudes—than the level that can be reached farther out in the cluster.

One other method of coping with crowding is worthy of mention (Anderson 1997). If one's objective is to achieve photometry with the highest level of accuracy (for judging the color width of the main sequence, for example), one can use some objective criterion, such as the distance to the nearest neighbor, to select a subset of stars that are, *a priori*, likely to be measured with higher accuracy than the others.

1.5. *The choice of color bands*

Finally, a word about color systems. For the photography-based study of globular clusters, the color bands that were convenient and natural were the blue B and yellow V for which photographic materials had good sensitivity; and it became standard to study globular clusters in B and V. As already noted, most CCDs have inferior sensitivity at B; quite to the contrary, they have excellent sensitivity in the near infrared, with the result that modern CCD studies are tending more and more to be made in the V and I bands. These questions will be taken up in more detail in Section 3, which goes into the details of the photometric results.

2. Chemical Abundances in Globular Clusters

One of the prominent differences between globular clusters is in the chemical mix of the material of which their stars are made. Early in the history of stellar populations it had become evident that the spectra of the kindred stars of the halo population (as we now call it—in those days they were the "high-velocity stars", because the halo stars that are passing through our neighborhood have velocities that differ so much from our near-solar reference standard)—had noticeable differences from those of stars like the Sun (Lindblad 1922, Roman 1950); the differences were traced to a deficiency in the abundances of the elements heavier than helium (Chamberlain & Aller 1951). Because of the faintness of even the brightest globular-cluster giants, it was not until much later, with much greater effort, that two globular-cluster giants were observed and found indeed to be deficient in heavy elements (Helfer, Wallerstein, & Greenstein 1959).

The elements beyond helium are generally referred to as the "metals," even though the most abundant of them include C, N, O, Ne, Si, etc. Since among the more abundant heavy elements iron is the most prominent in stellar spectra, the practice has generally been to cite the "metallicity,"

$$[\mathrm{Fe/H}] = \log_{10} \frac{\mathrm{Fe/H}}{(\mathrm{Fe/H})_\odot},$$

with the implicit assumption that all the other heavy elements go in lock step with iron. It has now become evident that this assumption was unjustified, but [Fe/H] has been used as a metallicity index for many decades, and the differing behavior of other elements is still insufficiently explored. I will therefore use the index [Fe/H] throughout most of this section, shifting only at the end to the modern picture of greater complexity.

Baade's original population distinction had merely contrasted the color–magnitude diagrams, but the spectroscopic studies just mentioned showed that the globular clusters were metal-poor relative to the Sun. The weakening of metal lines was a strong indication; but, interestingly, the most graphic demonstration of low metallicity came from theory. In their epoch-making first calculation of an evolutionary track, Hoyle & Schwarzschild (1955) found that with solar chemical abundances they produced the red-giant branch of

the open cluster M67, while they could match the RGB of the globular cluster M3 only by lowering the abundance of heavy elements by two orders of magnitude.

Differences in metallicity among globular clusters soon became apparent. In Morgan's study of the spectra that Mayall had taken to measure radial velocities of globular clusters (Morgan 1956, 1959), the Great Classifier found that he could distinguish eight different degrees of metal-line strength. (In those papers the distinction between disk and halo globulars was also first made.) Morgan's subjective classes were soon replaced, however, by the quantitative [Fe/H].

Metallicities in globular clusters have been estimated by every method that astronomers have been able to conceive of: line strengths in integrated spectra, various types of color indices or spectral indices, or spectra of individual stars. Only the last of these can provide a direct measure of chemical abundances; the other methods each require calibration onto this true abundance scale.

Even individual spectra must be used with care, because of the complexity of interpreting line strengths in a stellar atmosphere. Low-dispersion spectra are dangerous to use, because their information is contained in the strong spectral lines. These, by the very nature of their strong absorption coefficient, are formed in the upper layers of the star's atmosphere, whereas information about the temperature of the atmosphere, which governs excitation and ionization, comes from the deeper photospheric layers, around unit optical depth. This means that the temperature of the layer in which the strong lines are formed is known only by extrapolation of the photospheric temperature via a model atmosphere, any errors in which will be reflected in abundance errors. The interpretation of strong lines also requires use of a curve of growth, whose characteristics also depend on the model atmosphere; furthermore, it is in the nature of the curve of growth that strong lines are less sensitive to abundance than are weak lines.

The most reliable abundance measurements—perhaps the only reliable ones—come from measurements of weak lines, because, having an absorption coefficient that differs only a little from that of the neighboring continuum, they are formed in the atmospheric layer whose temperature we know the best. The price of this advantage is the cost of obtaining high-dispersion spectra—even higher dispersion than one might think, because of the need to detect even weaker lines, neglect of which might cause the continuum level to be set incorrectly. Getting high-dispersion spectra of stars as faint as those in a globular cluster costs valuable time on large telescopes, but the price is worth it.

Abundance measurements are discussed in more detail by Gratton in Chapter 4.

Values of [Fe/H] in globular clusters range from close to zero down to about -2.3. The best abundance determinations have accuracies of about ± 0.15, or in common parlance, 0.15 dex. Some of the photometric abundance indices are capable of comparable accuracy, but have the danger of systematic errors due to inadequate calibration of the index against high-dispersion spectroscopic measures.

Although globular clusters differ from one another, nearly every cluster contains stars of a uniform chemical abundance, as one would expect from a single event of star formation. An outstanding exception, however, is Omega Centauri, where spectroscopy of red giants and subgiants has revealed an abundance spread of nearly half a dex. The stars of the main sequence of the cluster are too faint for easy spectroscopic study, but it has been recognized for some time that the main sequence has a color width indicative of a range of metallicities, and a recent Berkeley thesis (Anderson 1997) has found an actual split of the middle main sequence into two parallel sequences.

Smaller abundance differences have also been observed in M22 (Anthony-Twarog, Twarog, & Craig 1995, which also gives references to earlier studies).

In recent years it has become evident, however, that it is an oversimplification to

describe the chemical composition of a globular cluster, or of any other stellar population, by a single parameter such as [Fe/H]. The complexity first showed up as an overabundance of oxygen relative to iron (stated as [O/Fe]) in some clusters. A similar behavior had been seen in field stars, and there turned out to be a systematic behavior there, in that [O/Fe] was most enhanced in the stars of lowest [Fe/H] (by about 0.5 dex) but gradually decreased with increasing [Fe/H] until it leveled off at zero (i.e., a solar ratio of oxygen to iron) for values of [Fe/H] greater than −1.0. A recent preprint (Chiappini et al. 1998) gives abundant references to this problem, as does the general review of abundances by Wheeler, Sneden, & Truran (1989).

This behavior of the abundances received a ready explanation, which notes that shortly after the birth of a population the only supernovae that appear are of Type II, because their progenitors are upper-main-sequence stars that have rather short lifetimes. These stars produce a lot of α-process elements, and relatively little iron. (Oxygen is an α-process element, and further study has shown that the heavier α-process elements are also enhanced along with oxygen.) It is only after a billion years or so that supernovae of Type I begin to appear, and then the iron abundance goes up. Thus the initial build-up of metallicity emphasizes the α-process elements, while iron also increases slowly; then the larger production of iron takes over.

Unfortunately recent discoveries in clusters near the Galactic bulge have shattered the uniformity of this picture (Barbuy et al. 1999). NGC 6553, with [Fe/H] = −0.5, has a striking level of α-enhancement, whereas 47 Tuc, with a very similar metallicity but in a different part of the Galaxy, does not show this peculiarity at all. (The same also seems to be true of NGC 6528, another bulge cluster [Barbuy, personal communication].) As Bernard Pagel was so kind as to point out to me, in a discussion shortly after the end of the Winter School, what this means is that there was a difference, for the material that was to go into the stars of each of these two clusters, in the length of time that elapsed in building up a level of [Fe/H] that is similar in both clusters. In the material that was to make up NGC 6553 (and presumably the other bulge clusters), element enrichment by ejection from Supernovae II was so rapid that [Fe/H] was already up to −0.4 in less than 1 Gyr, i.e., before any Supernovae of Type Ia could form and decrease the early [α/Fe] level of 0.4–0.5. For 47 Tuc, by contrast—and apparently for other clusters outside the bulge—enrichment to the same level of [Fe/H] took a much larger number of years, so that there the Sne Ia were able to play their role of bringing up the level of iron relative to the α elements.

This difference in the rate of enrichment of pre-cluster material leads to two different possible scenarios. If all enrichment began at the same time, then the bulge clusters are older than those outside the bulge. On the other hand, if we wish to view all globular clusters as having similar ages, we would then have to postulate that element enrichment began earlier outside the bulge than in it—but with a time lapse in the bulge that contrived to be just long enough for cluster formation to begin simultaneously everywhere.

In any case, it is an observational fact that the chemical-abundance history has been different in different locations in the Galaxy, and the specter thus threatens us that α-enhancement may be a quantity that has to be measured in each individual cluster, quite separately from [Fe/H].

Finally, I leave it to others to speculate on the significance of the fact that CMD properties associated with this sort of abundance difference have also been observed in the stars of the M31 halo (see Sect. 3.3).

3. The Morphology of Color–Magnitude Diagrams of Globulars

The vehicle for describing the populations of globular clusters has always been the color–magnitude diagram (CMD).

Note first that the CMD, a plot of the absolute magnitudes of the stars against their colors, is the observational expression of the HR diagram of the theoreticians, whose coordinates are instead $\log L$ and $\log T_e$, where L is the energy output of the star added up over all wavelengths and T_e is its effective temperature, defined by the relation

$$L = 4\pi R^2 \cdot \sigma T_e^4,$$

where σ is the Stefan–Boltzmann constant. The relation between the observational and the theoretical coordinates of a given star is by no means obvious; the determination of such relationships has consumed a large amount of discussion and effort on the part of both observers and theoreticians. Its details will not be discussed here; we merely note that it is an important problem in the comparison of observations with theory, for which the theoreticians need to produce not just an L and a T_e but a complete spectral energy distribution for each star.

In the present section the coordinates of the CMD will be taken to be M_V and $(B-V)_0$, where B and V are two of the bands of the Johnson system, defined in principle by Johnson & Morgan (1953) and operationally by the standard stars of Landolt (1992). The zero subscript on $B-V$ is intended to remind us that all colors must be corrected for interstellar reddening, just as the definition of absolute magnitude implies that it corresponds to a true distance, corrected for interstellar absorption.

Again, it is well to mention in passing that the conversion between observational and theoretical coordinates depends on the details of the bandpass, so that the conversions will not be the same for, e.g., the standard Johnson V and the V defined by the filter and detector in use at a particular observatory at a particular time. It is a common practice of observers to transform their observed magnitudes to a standard system, as best they can; but in fact this can be a damaging step. In comparisons with theory it is essential instead to integrate the theoretical results over the specific passband in which the observations were made.

Note also that observations tend now to be made increasingly in quite different bandpasses, such as the V and I that are increasingly favored by observers who use CCDs, either from the ground or with HST, or the J and K that are appearing in an increasing number of papers by observers in the near infrared. The present section, which discusses CMDs from the point of view of morphology alone, will express everything in the more traditional B and V, with the understanding that the conversion from other bands can be made in principle, although (as just emphasized) detailed comparisons between observation and theory should always be made with the bandpass actually used in the observations.

Figure 3 is a schematic plot of a globular-cluster color–magnitude diagram, showing the sequences (and parts of sequences) that are generally recognized. Other features are often described and referred to, but the sequences labeled in Fig. 3 are the regions of the CMD that will be discussed in this section. The figure is intended only to be schematic; it can indeed be no better than that, because the positions of the sequences differ from cluster to cluster, and some branches are totally or partially lacking in some clusters.

3.1. *The main sequence(MS)*

The main sequence extends from the turnoff at its top (which will be discussed in Subsection 3.1.1) to the limit of hydrogen burning (discussed in Section 6). Aside from the locations of these two limits, clusters differ in the colors of their MSs at a given ab-