LBL-32634
UC-414

Lawrence Berkeley Laboratory
UNIVERSITY OF CALIFORNIA

Physics Division

Measuring q_o Using Supernovae at $z \approx 0.3$

H.J.M. Newberg
(Ph.D. Thesis)

July 1992

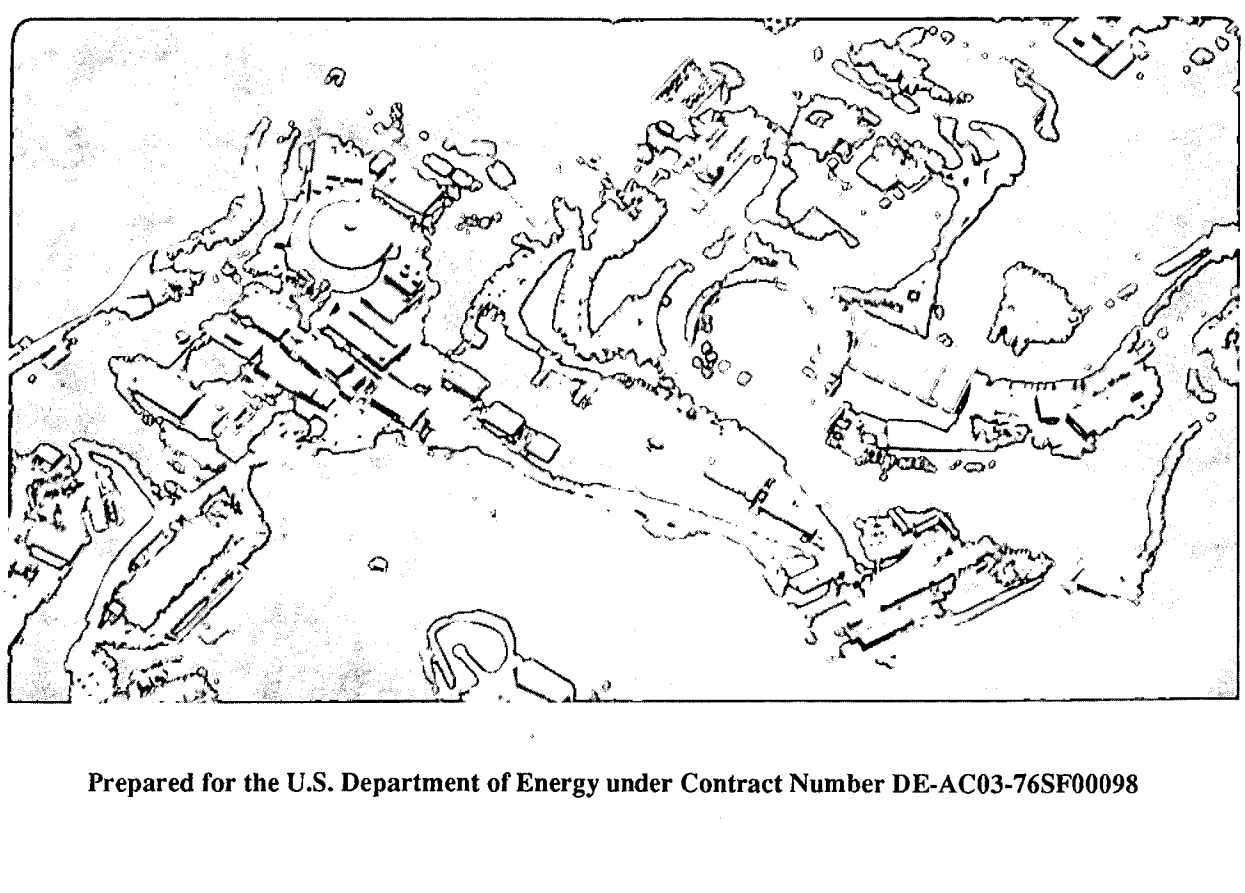

Prepared for the U.S. Department of Energy under Contract Number DE-AC03-76SF00098

DISCLAIMER

LBL-32634
UC-414

Measuring q_o Using Supernovae at $z \approx 0.3$

Heidi Jo Marvin Newberg
Ph.D. Thesis

Department of Physics
University of California

and

Physics Division
Lawrence Berkeley Laboratory
University of California
Berkeley, CA 94720

July 1992

This work was supported by the Center for Particle Astrophysics under National Science Foundation Contract No. AST8809616, and by the University of California Space Sciences Laboratory, and by the Director, Office of Energy Research, Office of High Energy and Nuclear Physics, Division of High Energy Physics, of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

Measuring q_o Using Supernovae at $z \approx 0.3$

by

Heidi Jo Marvin Newberg

Measuring q_0 Using Supernovae at $z \approx 0.3$

by

Heidi Jo Marvin Newberg

Abstract

The measurement of q_0 is extremely important for understanding the quantity of matter in our universe. The measurement of q_0 using supernovae of type Ia as standard candles is appealing because it requires less modeling than other methods using galaxies. The challenge with using supernovae to measure q_0 is in finding enough of them.

In order to find supernovae, we have constructed a very popular f/1 camera for the 3.9m Anglo-Australian Telescope. The camera uses reducing optics that put a 17' x 17' field on a 1024 x 1024 pixel Thomson CCD. Using this system, we image to 23rd magnitude in five minutes. We have developed a software package that uses image subtraction to find supernovae that are approximately magnitude 22.4 or brighter in these images. One field can be processed every 6.6 minutes on a relatively unloaded VAX 6000-6510.

We estimate that this system should find one supernova in every 105-139 images (about two nights of observation on the AAT). Throughout the two years of operation, we observed the equivalent of about four nights with seeing better than two arc seconds. Although we found many candidates, we were unable to confirm any supernovae.

The next generation of this search is currently using the 2.5m Isaac Newton Telescope in the Canary Islands. We have solved many problems encountered with the AAT search by targeting distant clusters of galaxies, by operating at a site that allows

higher resolution imaging, and by scheduling follow-up observations. This system, although still in its infancy, has yielded one event that could be a high redshift supernova.

We estimate that it will require 36 type Ia supernovae, discovered at or before maximum, to find q_0 to within 0.2, and 144 such supernovae to measure q_0 within 0.1. Clearly, the program will have to be expanded to find this quantity of supernovae. In addition, more information about nearby supernovae Ia are necessary to properly model the effects of dust in the host galaxies, and to accurately calculate the redshifted light curves of type Ia supernovae.

I dedicate this thesis to Dr. Charles Bean and Dr. Robert Resnick. Dr. Bean was my undergraduate research advisor at Rensselaer Polytechnic Institute. Under his direction, I wrote the first scientific paper of my career. Dr. Resnick was my undergraduate academic advisor. At Rensselaer, we were fortunate to have full professors teach our recitation sections. As my discussion leader for introductory physics, Dr. Resnick encouraged me to become a physicist. As my advisor, he buoyed my confidence, and helped me to find an undergraduate research project.

Both of these people instilled in me the belief that life will always improve, no matter how bleak the immediate outlook. Their confidence in my ability to succeed and their encouragement to persevere through temporary setbacks have been important to my success in graduate school. I am pleased to take this opportunity to thank them.

Table of Contents

List of Figures

List of Tables

Acknowledgments

This search for high redshift supernovae is the result of the combined efforts of many talented and dedicated people. I would like to thank all of the people who contributed to the effort, and all of those who helped me to become a better scientist in the process.

My thesis advisor, Richard A. Muller, provided an atmosphere of scientific excitement that will be difficult to find anywhere else. Although he taught me many things that are probably more useful in my career, I will always be grateful to him for teaching me how to whistle. Carl Pennypacker has the audacity to believe we can measure the deceleration parameter, and the guts to organize the collaboration. Saul Perlmutter taught me everything I know about image processing and showed me the power of the computer as a labor-saving device. He was also one of the originators of the high-redshift supernova project, and was a source of many good ideas. Warrick Couch and Brian Boyle logged many hours at the controls of the Anglo-Australian telescope, and helped shape the project in its early stages. Gerson Goldhaber insisted we always concentrate on the results.

I would especially like to thank Russell Cannon and the staff of the Anglo-Australian Observatory for their support of our project. Members of their team built the CCD electronics, improved on our filter-shutter design, helped us build our software into their computer system, and otherwise made our collaboration as pleasant and productive as possible.

We were fortunate to have many undergraduates work with us over the course of the project. Special thanks go to Ivan Small, who wrote several of our analysis routines;

Torrey Lyons, who implemented our display routine and wrote the code that reads Figaro files into VISTA; and Jigna Desai and Diane Lowry, who worked whatever hours were necessary to get the data processed in time for follow-up observations.

My thesis committee members, Richard Muller, Alex Filippenko and Buford Price, have made valuable contributions to the content and style of the manuscript.

I would like to thank my parents, Michael Marvin and Nancy Lawson, for supporting me in the esoteric pursuit of a Ph.D. in astrophysics. Last, but certainly not least, I thank my husband, Lee Newberg, who has endured the odd hours of an astronomer, many stressful weeks of job-hunting and thesis-writing, and agreed to proofread this entire manuscript.

This work was supported by the Center for Particle Astrophysics under NSF contract AST8809616, and by UCB Space Sciences Laboratory under contract to the Director, Office of Energy Research, Office of High Energy and Nuclear Physics, Division of High Energy Physics, of the US Department of Energy (DE-AC03-76SF00098).

1. Introduction

1.1. The deceleration parameter, q_0

Since cosmological theories must be supported with a limited amount of experimental data, the strategy thus far has been to make them as simple as possible and to make modifications as contradictions with observations deem them necessary. Although this notion is patently false on scales at least up to the size of clusters of galaxies, it is believed that the universe is homogeneous and isotropic on large scales. This belief is supported by the smooth nature of the cosmic microwave background radiation. Prior to 1929, it was assumed that the universe was also static. However, Hubble (1929) discovered that the galaxies are moving away from each other.

If we wish to preserve the assumption of homogeneity and isotropy in the universe, yet also require that the distance between objects increase, then we must believe that the universe itself is expanding. The framework of general relativity allows us to describe this expansion in terms of a cosmological scale factor, R(t) (see Weinberg 1972, or section 5.1 of this paper for more background). Using this scale factor, the Hubble law can be written as

$$v = H_o d,$$
$$\text{with } H_o \equiv \frac{\dot{R}(t_o)}{R(t_o)}, \tag{1.1}$$

where v is a galaxy's velocity, d is its distance, H_0 is the Hubble constant at the present time, and t_0 is the present age of the universe. The Hubble parameter is written so as to render the scale of R inconsequential.

One of the great mysteries of our universe is the cause of the initial expansion. However, once begun, we expect the expansion to slow due to the gravitational attraction

of all matter (and energy). If the deceleration is large enough that it will eventually cause the cosmological scale factor to decrease, then the volume of the universe is finite. For this case, the universe is closed. If the universe expands indefinitely, then it is an infinite (open) universe. The deceleration parameter, q, is defined as a dimensionless measure of time rate of change of the Hubble parameter:

$$q(t) \equiv -\frac{\ddot{R}(t)R(t)}{\dot{R}(t)^2}. \tag{1.2}$$

Since cosmological models are underconstrained, we will assume the simplest form of Einstein's equations (with zero cosmological constant, Λ_0). Using these, one finds that the deceleration parameter is determined entirely by the mass density of the universe. If we knew the current deceleration parameter, q_0, then we would know the current mass density. If we knew q_0, we could predict whether the universe will eventually begin to contract, or whether it will expand indefinitely. It is this parameter that we wish to measure.

1.2. Other means of measuring q_0

Many attempts have been made to measure the deceleration parameter, q_0. One might assume that the simplest course of action would be to measure the mass density, and infer the deceleration parameter from the expression

$$q_0 = \frac{\Omega_0}{2}, \tag{1.3}$$

which holds in a standard Friedmann universe with $\Lambda_0 = 0$. Here Ω_0 is the ratio of the current mass density to the critical mass density at which the universe would be barely open:

$$\Omega_o = \frac{\rho_o}{\rho_c},$$

$$\rho_c \equiv \frac{3H_o^2}{8\pi G}. \tag{1.4}$$

However, it is difficult to measure the mass density. Mass measurements must generally be made by observations of the interactions of luminous matter. For example, the mass of a galaxy is obtained from the distribution of velocities of its component stars, and the mass of a cluster is inferred from the velocities of its component galaxies. We then determine the mass/luminosity ratio for these systems, and calculate the mass density by integrating the luminosity of the objects in a specified region of space and multiplying by the mass/luminosity ratio. However, this calculation is insensitive to a uniform distribution of non-luminous matter. Calculating the mass density from the mass/luminosity ratio assumes that light is a good indicator of the presence of mass. It has been clearly demonstrated from galaxy rotation curves, for example, that light is at best a weak tracer of matter (Faber and Gallagher 1979, Rubin and Graham 1987).

Even if it were possible to measure accurately the local mass density by kinematic means, we would still have to negotiate the inhomogeneity of space. We have so far found correlations and large-scale clumping of matter on all scales measured. Even the galaxy distribution from the IRAS survey, which covered 87.6% of the sky to 6000 km/sec, is dominated by the Hydra-Centaurus-Pavo-Indus and Perseus-Pisces superclusters (Strauss et al. 1992). Finding a representative sample of the local universe is difficult, indeed. Until we can evaluate the mass of a section of the universe in which individual galaxy clusters are insignificant, these kinematic methods are more useful for constraining the local evolution of our universe than measuring its global properties.

There are three methods for measuring the deceleration parameter directly (Sandage 1975). In each of the three methods, a measurement of a distant object (or objects) is compared with the expected result from a linear Hubble law; the deceleration parameter is calculated from the deviation. The first method is to compare the apparent magnitude of an object of known luminosity with the magnitude we would expect from an object at a

distance determined from the Hubble law. The second method is to compare the angular size of an object of known length with the size expected from that object at a distance determined from its redshift. Or third, our estimation of volume from redshift and solid angle can be compared with an independent measurement of volume. We shall discuss only one representative measurement from each technique.

The magnitude-redshift test

The use of galaxies and their redshifts is the oldest and most tried method of measuring q_0. In general, one assumes that the brightest galaxies in clusters are all approximately the same absolute luminosity. From their apparent magnitudes, one can measure the luminosity distance and compare this with the redshift distance to measure q_0. The success of this method hinges on a small spread in the luminosities of bright cluster galaxies and the lack of a bias in choosing galaxies at high redshift. Unfortunately, it is difficult to evaluate the systematic errors in this method since it is difficult to separate the effects of galaxy evolution, cluster evolution, and biases at high redshift, from the expected results from different values of q_0.

When choosing the sample of galaxies to include, it is important to reduce the systematic bias as much as possible. One example is the Malmquist bias (Rowan-Robinson 1984), which arises from a spread in luminosities of the brightest cluster galaxies. In a magnitude-limited survey, we are likely to choose clusters with intrinsically brighter galaxies at higher redshift. In addition, the magnitudes must be *K*-corrected since the light we are receiving in our filter came from systematically shorter wavelengths in the rest frames of more distant galaxies. The Malmquist bias can be minimized by only including redshifts up to the point at which the dimmest galaxies in the distribution would

be visible. The *K*-corrections must be calculated from galaxy models, or by observations in different wavelength bands.

The greatest difficulty with the galaxy magnitude-redshift test is accounting for galaxy evolution. Figure 1.1 shows the data from several attempts to measure q_0. Although the statistical spread in galaxy magnitudes is small enough to differentiate a closed universe from an empty one, the systematic differences between galaxy evolution models make the measurement entirely uncertain. For example, if we use model 1 to describe the galaxy evolution, the graph shows that Ω_0 is less than zero. If we use model 2, Ω_0 is greater than one.

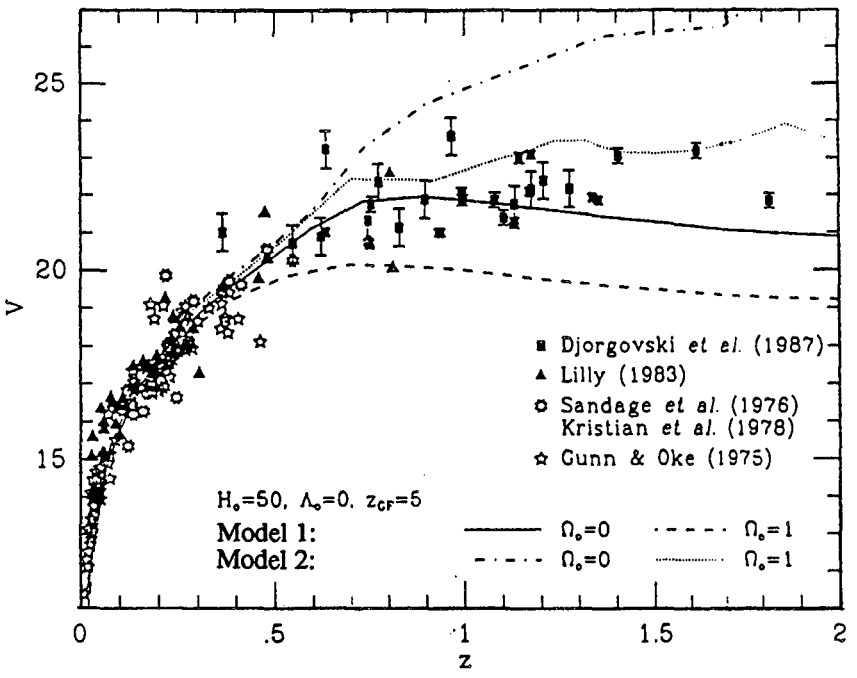

Fig. 1.1. Measuring q_0 using brightest cluster galaxies. This is a visual region Hubble diagram for bright radio galaxies and brightest cluster members (from Spinrad and Djorgovski, 1987). The data include cluster giant ellipticals at the lower redshift ($z<0.76$), and 3CR radio galaxies at the higher redshifts. Superimposed are the simple cosmological models described by $q_0 = \Omega_0 = 0$ and $q_0 = 1/2$, combined with two galaxy evolution models by Bruzual. The individual galaxy models can be found in the paper by Spinrad and Djorgovski. Notice that the individual galaxy models are more important than the density parameter in achieving a good fit to the observations.

Additional concerns arise from the uncertainty in evolution of the clusters themselves. It is possible that the more distant clusters formed early due to large initial density contrasts and thus may have brighter, more massive central galaxies than the clusters at lower redshift. Or, galaxy mergers may make the closer, older galaxies brighter. In any case, galaxy mergers in cluster cores will severely alter the evolution of the brightest cluster galaxies.

The angular-size-redshift relation

Hickson (1977) used the harmonic mean separation of cluster galaxies as a standard ruler to measure q_0. Photographic plates containing ninety-five clusters in the redshift range $0.02 < z < 0.46$ were collected. He then calculated the harmonic mean separation between the 40 brightest galaxies within 3 Mpc of the brightest central cluster galaxy using the formula

$$\lambda = \left[\frac{2}{N\,(N-1)} \sum_{i>j}^{N} \frac{1}{s_{ij}} \right]^{-1}, \tag{1.5}$$

where s_{ij} is the projected separation between the i^{th} and the j^{th} galaxy in Mpc, and $N = 40$. There are $N\,(N-1)\,/\,2$ separations between N galaxies. The harmonic mean separation, λ, was used as a standard ruler from which the curvature of the universe could be obtained. He found $q_0 = -0.9$ before taking into account evolutionary effects, and $q_0 = -0.8$ after taking into account known evolutionary effects, with a formal standard deviation of 0.2. A negative q_0 implies that the expansion of the universe is accelerating, which would necessitate the addition of a repulsive term (Λ_0) in the Einstein field equations. The reliability of this measurement rests on (1) a small dispersion in the harmonic mean galaxy separation in clusters, and (2) the ability to accurately identify clusters, cluster members, and the position of the cluster center.

From numerical simulations of clusters, Hickson found that the dispersion in the harmonic mean galaxy separation, λ, was approximately 20%. However, the more distant clusters could have different intrinsic values of λ due to evolution or selection effects that are hard to quantify. On the average, more distant clusters should appear to be in earlier evolutionary stages. This means the clusters have had less time to collapse into virialized systems, the galaxies may have different types and intrinsic magnitudes, and there may be different galaxy collision rates. In addition, he could be biased towards choosing only very large clusters at high redshift since those are the only ones that would be found at fainter magnitudes in more crowded fields, or because the larger primordial density perturbations would have been the first to form galaxy clusters. Since the observations are generally obtained in one filter, as we look to higher redshifts the galaxies are also sampled at increasingly shorter wavelengths in their rest frames.

The clusters themselves can be difficult to identify in a two-dimensional survey. In sky surveys, chance superpositions of galaxies that are really at quite different redshifts can mimic an actual overdensity, or cluster (Politzer and Preskill 1986). Even if all of the cluster identifications are correct, we still must correct for the expected number of field galaxies that have contaminated our measurement of λ, which Hickson did by subtracting the number of field galaxies in nearby circles of similar radius. He calculated that this operation only increased the uncertainty in measuring λ by a factor of 1.5.

The most serious systematic error in using harmonic mean separations results from the aperture effect. For the purpose of calculating λ, galaxies are assumed to be cluster members if they fall within 3 Mpc of the cluster center. In order to calculate the radius of the aperture on a survey plate that corresponds to a 3 Mpc radius in the rest frame, one must choose a value of q_0. Unfortunately, this introduces a very strong feedback effect, as

explained by Bruzual and Spinrad (1978). As one can see from Fig. 1.2, the value of q_o used to choose the aperture size strongly affects the measured q_o. This is because the harmonic mean distance between bright galaxies increases significantly if even one galaxy is added to the outside perimeter of the cluster. This test serves primarily to measure the aperture enclosing the cluster rather than an intrinsic property of the cluster itself.

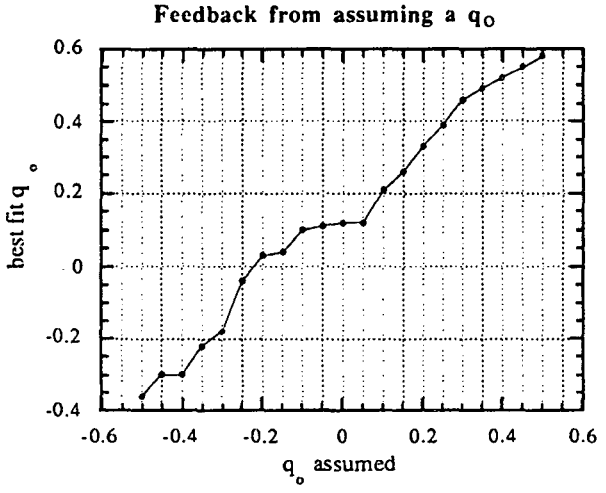

Fig. 1.2. Measuring q_o using clusters as standard rulers. The data for this graph were taken from Bruzual and Spinrad (1978). They demonstrate the feedback problem caused by using q_o to determine the cluster radius, and then using the galaxy separations to in turn calculate q_o. There is no value for which q_o is self-consistent.

The number counts test

A more recent technique for measuring the deceleration parameter uses galaxy counts to estimate the volume of space in different solid angles on the sky. One can calculate a measure of the volume from solid angle and redshift distance, assuming a linear Hubble law. Then the expected number of galaxies is computed from the density of nearby galaxies and compared with the actual number measured. This method was slow to come into use because it requires redshifts for a large number of galaxies at distances greater than

$z = 0.5$. Loh and Spillar (1986) developed a method to photometrically determine redshifts, and thus were able to measure 1000 field galaxies with a median redshift of 0.5. They determined that $0.2 < q_0 < 0.8$ with 95% confidence (Fig. 1.3).

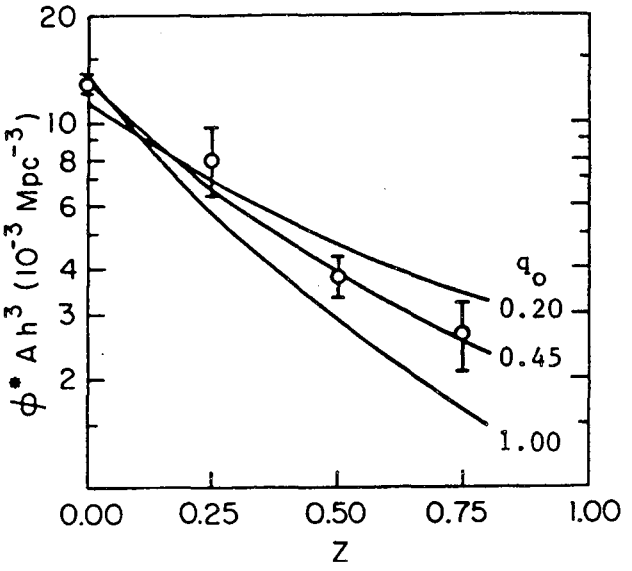

Fig. 1.3. Measuring q_o from galaxy counts. This measurement of the deceleration parameter by Loh and Spillar (1986) used the volume element as determined by galaxy counts. They determined that q_o was between 0.2 and 0.8 with 95% confidence.

In order to estimate the number of galaxies in a volume of given redshift range and solid angle, one must assume that the luminosity distribution of galaxies is known. The galaxies above the limiting magnitude of the image are counted and the number of fainter galaxies is calculated from the galaxy distribution (usually, the Schechter (1976) function). The assumptions here are that galaxies are not created, destroyed, or combined and that they evolve in luminosity such that the fractional change in luminosity with redshift ($d \ln L/dz$) is independent of luminosity at every redshift. The validity of these assumptions has been challenged by Bahcall and Tremaine (1988), who claim that these

systematic effects cause an uncertainty of $\Delta\Omega_0 = \pm 1$ in Loh and Spillar's value of the density parameter.

In addition to these difficulties with using galaxy counts, there are other concerns more directly related to the method of Loh and Spillar. Photometric determinations of redshift may be inaccurate, especially at large redshift where the galaxies could show substantial evolution. The number counts could be inaccurate due to star-galaxy misclassification. Lastly, the photometry of crowded fields can be uncertain when objects overlap.

The key feature of all three of these methods is the existence of an object of known properties that is visible at redshifts exceeding $z \approx 0.2$. The only objects we have so far identified at large redshifts are galaxies, quasars, and supernovae. Galaxies have been used most often in q_0 measurements because, of the three, they are the most readily observable. Unfortunately, galaxy evolution and selection effects have hindered every attempt to quantify their properties at high redshift; a large amount of very uncertain modeling has been employed to wrestle with the problem. Once they can be found in sufficient numbers, supernovae should prove to be more reliable cosmological indicators.

1.3. Supernovae

Observationally, a supernova appears as a very bright new light source, sometimes brighter than the galaxy in which it is found. The spectra of supernovae show characteristic P-Cygni line profiles indicative of high-velocity expansion. Supernovae are divided into categories based on their spectral characteristics. Type II supernovae are defined by the presence of hydrogen emission in their spectra. This class is further subdivided according to the shapes of the supernova light curves. Type I supernovae lack

the hydrogen feature. These supernovae are divided spectroscopically into the subclasses: Ia (silicon absorption in the early spectrum), Ib (no silicon absorption), and possibly Ic (like a Ib with little or no helium). For more information on supernova classification, see Filippenko (1991), Harkness and Wheeler (1990), and Woosley and Weaver (1986) and references therein.

It was originally assumed that supernovae with hydrogen in their spectra (type II's) were associated with the core collapse of very young, massive stars, and that stars which exploded as type I supernovae were the result of explosions of highly evolved, less massive stars which had lost most of their hydrogen envelope. However, as theorists simulated supernova explosions and as more supernovae were observed, it became clear that the classes were not as distinct as they originally appeared. For example, one supernova (1987K) was observed by Filippenko (1988) to make a transition from a type II to a type Ic. It is possible that the type II, Ib, and Ic supernovae form a continuum of hydrogen and helium line strengths, rather than three distinct classes. This could result from a continuum of progenitor masses, or from mass loss in binary systems.

During their evolution, massive stars fuse increasingly heavier elements. This fusion heats the core and creates pressure which keeps the star from gravitationally collapsing. Once the core is converted to iron, however, no more energy can be extracted from fusion, so the core ceases to produce energy. From then on, it is the Fermi pressure of the electrons that keeps the core from collapsing. If the core is more massive than 1.4 solar masses, not even the Fermi pressure can hold it up. It becomes energetically favorable for the electrons and protons in the core to combine to form neutrons, which are now free to fall towards the center of the star. The core is compressed until it reaches nuclear density. What happens next is somewhat uncertain. It is believed that after striking each other, the neutrons recoil to form a shock wave which travels toward the surface. In

order to keep the shock from stalling before it reaches the surface, many theories include scattering of a small percentage of the core collapse neutrinos by the neutrons just behind the shock front. Except at early times, the decay of radioactive elements produced in the explosion (especially ^{56}Ni and ^{56}Co) powers the optical display of these type II supernovae. It has been suggested that type Ib/Ic supernovae are also a result of the core collapse of massive stars that have burned or ejected most of their hydrogen and helium layers. It is unlikely that we will find many core collapse supernovae in a high redshift search, since they are typically one to three magnitudes fainter than supernovae of type Ia (Miller and Branch 1990).

Supernovae of type Ia form a remarkably homogeneous class of extremely luminous objects. Since these supernovae have been observed in elliptical galaxies, they are believed to arise from older, population II stars. Among supernova theorists, the most popular scenario for supernova Ia explosions is the triggered detonation or deflagration of a carbon-oxygen white dwarf (Sutherland and Wheeler 1984, Woosley and Weaver 1986, Arnett, Branch, and Wheeler 1985, Woosley 1989, and Nomoto et al. 1991). In this scenario, matter (possibly from a binary companion) is deposited slowly on the surface of the white dwarf until its mass nears the Chandrasekhar limit. At this point nuclear reactions are ignited in the center of the star. A nuclear burning front moves out from the core, transforming a large portion of the mass to ^{56}Ni. It is the radioactive decay of nickel that is believed to power the light curve of supernovae of type Ia. Models based on this scenario explain the shape of the light curve, the absence of hydrogen, the uniform peak magnitude, and the fact that type Ia supernovae are seen in elliptical galaxies as well as spirals.

If this model is correct, the peak magnitude of supernovae Ia will not vary with redshift. This is because stars at high redshift will presumably be different only in their lower initial heavy element abundances (a heavy element mass fraction of less than one

percent rather than the solar heavy element mass fraction of 2 percent). Since the supernova processes a large portion of its matter into heavy elements, it is unlikely that a small discrepancy in the initial abundances will make a large difference to the explosion energy or the light curve.

Although the supernova rate is not very well known for galaxies significantly more distant than the Coma cluster, we know that they exist due to a discovery by Nørgaard-Nielsen et al. (1989). If type Ia supernovae are, in addition, of uniform light curve and peak magnitude, then they are easily the best standard candles for measuring cosmological distance scales. Although some modeling will have to be applied to correct for systematic effects of extinction and the tendency to discover intrinsically brighter objects with higher efficiency, this will be simple compared to modeling the evolution of entire galaxies or clusters.

1.4. SNe Ia as standard candles

The measurement of q_0 using supernovae of type Ia hinges on the assumption that these supernovae have a consistent absolute magnitude at peak luminosity. Although there is still some contention over the dispersion of peak magnitudes and whether individual SNe Ia do not fit the model (Filippenko et al.1992a, Phillips et al.1992, and Filippenko et al. 1992b), it is clear that the overwhelming majority of SNe Ia are strikingly similar to each other (Minkowski 1964; Kowal 1968; Arnett, Branch, and Wheeler 1985; Cadonau, Sandage, and Tammann 1985; Elias et al. 1985; Leibundgut and Tammann 1990; and Miller and Branch 1990).

The similarity of the light curves is demonstrated by the composite blue light curve of Cadonau, Sandage, and Tammann (1985). They superimposed individual type I light

curves (the difference between Ia and Ib,c was not established until the mid to late 1980's) which had been freely translated so that their peaks coincided in time and magnitude at peak. This shifting is justified by the uncertainty of explosion time and luminosity distance for most supernovae. They concluded that the majority of the discrepancy between different light curves was due to differences in photographic measurements and that the scatter in the blue curves (only $\sigma \approx 0.18$ magnitudes near maximum and $\sigma \approx 0.29$ magnitudes fifty days later) could be explained by observational error in every case known at that time. As was pointed out in their paper, a small error in the subtraction of the background light of the parent galaxy could cause a change in the shape of the light curve. This could account for the increasing dispersion of the light curves with time.

Since we know that some scatter in the peak luminosity will be introduced from gas and dust in the supernova's parent galaxy, perhaps the best wavelength band in which to study the nature of supernova explosions is the infrared. The infrared photometry of Elias et al. (1985) tends to support the idea that supernova Ia light curves are intrinsically the same. From J, H, and K photometry of six Ia supernovae, it was determined that the light curve scatter was at most 0.2 magnitudes. When the Virgo infall velocity was adjusted so that supernovae in the south Galactic hemisphere fell into line with those in the north, the scatter was reduced to 0.1 magnitudes, similar to the earlier result of Elias et al. (1981) for two supernovae found in the same galaxy. This supports the idea that type Ia supernovae are intrinsically the same to better than 10%. The differences between light curves would then be ascribed to gas and dust, or to photometric inaccuracies.

The similarity of the type I light curve shapes is encouraging. However, the critical measurements for our research are the optical absolute magnitude at maximum, and the dispersion of type Ia supernovae. We wish to know the intrinsic dispersion of the supernovae as well as the dispersion we can expect from obscuration by the host galaxy.

The dispersions due to distance scale and photometric error for the nearby supernovae is only important in that it affects the measurement of the average absolute magnitude at maximum. There have been several modern studies of the blue magnitudes of supernovae Ia, and we will discuss each one in turn. However, we will translate all absolute magnitudes to $H_0 = 100$ km/sec/Mpc for comparison purposes.

For twelve well-studied supernovae found in elliptical galaxies, Cadonau, Sandage, and Tammann (1985) found $\mathbf{M_B^{max} = -18.2 \pm 0.4\ mag}$. Unfortunately, all of the supernovae were discovered before 1972 and therefore had no high accuracy photometric measurements. For most, the blue maximum was calculated from the photographic light curve. They find a net dispersion of $\sigma_B \approx 0.45$ magnitudes , but point out that this could all be due to extrapolation of light curves, uncertainty of distance determinations, and internal absorption. (Note that the dispersion, $\sigma_B \approx 0.45$, is not the same thing as the error in determining the average, $\pm$ 0.4 mag.)

Leibundgut and Tammann (1990) determined the blue maximum from six supernovae in the Virgo cluster to be $\mathbf{M_B^{max} = -18.45 \pm 0.12\ mag}$. The scatter of the blue maxima was found to be $\sigma_B \approx 0.18$ mag, which could be purely due to observational errors.

Miller and Branch (1990) used about 40 type I SNe (since many of the galaxies are spirals, the sample could be contaminated with a few SNe Ib) from the updated Asiago Supernova Catalogue (Barbon et al. 1989) to find $\mathbf{M_B^{max} = -18.36 \pm 0.03\ mag}$. For this estimate, the individual peak magnitudes were corrected for extinction by dust in the parent galaxy. The magnitudes of the dimmer supernovae in spiral galaxies were corrected by subtracting 0.8sec(i), where i is the galaxy inclination angle. Although this may seem like an ad hoc way to reduce the dispersion in the measurement, it is clear from their paper

that there is an inclination effect for dust obscuration and that the majority of the supernovae cluster around a peak maximum similar to that of the ellipticals. So, this may be a reasonable procedure. The dispersion in the blue maximum after this corrections is $\sigma_B = 0.27$ mag; again, this could easily be entirely due to observational effects and not due to any intrinsic dispersion of peak magnitudes of type Ia supernovae.

In no case has it been statistically determined that there is *any* intrinsic variation in the peak luminosity of type Ia supernovae (although there is as least one pathological case, see Filippenko et al. 1992a). On the other hand, it has not been shown that they are standard candles to better than 30%, either. For example, it may not be so surprising that these measurements of peak luminosity are similar, since they are largely calibrated from the same set of supernovae in elliptical galaxies. In addition, it has been shown conclusively that extinction in the host galaxy could easily be two magnitudes for a highly inclined spiral galaxy. When Miller and Branch (1990) calculated the peak magnitude before correcting for internal extinction, it was $M_B^{max} = -18.06 \pm 0.07$ mag with $\sigma_B = 0.70$ mag. The dust systematically decreases the maximum and increases the dispersion, creating difficulties for an accurate measurement of q_o.

At high redshift, it will be even more difficult to correct for dispersion due to extinction in the host galaxy. However, it will be much easier to calculate the distance scale; since the peculiar velocity will be small compared to the Hubble velocity at $z = 0.3$, the luminosity distance can be accurately calculated from the redshift. This eliminates one of the largest sources of error in measuring supernova peak magnitudes.

1.5. Overview of experimental technique

It is our goal to use the standard candle properties of type Ia supernovae to measure the deceleration parameter, q_o. Compared to galaxies, supernovae require relatively little modeling to determine their properties at large redshift. The difficulty with using supernovae as distance indicators lies in their transient nature. A type Ia supernova at a redshift of $z \approx 0.3$ will only be visible above magnitude 23 for approximately a month. In order to fit the light curve of the supernova to a standard template, accurate photometry measurements must be completed near maximum light. Due to these constraints, the supernovae must be found rising in brightness, and the discovery data must be reduced immediately so that follow-up observations can be made.

In order to increase our chances of finding distant supernovae, we have built a wide-field f/1 camera for the 3.9m Anglo-Australian Telescope. This camera consists of an hyperboloidal secondary mirror, correcting optics, and a 1024x1024 pixel Thomson CCD with a scale of one arc second per pixel. With this device, we are able to image a 17' x 17' section of the sky to 23rd magnitude in approximately five minutes of exposure.

It is expected that this system could discover approximately one supernova every clear night, if the seeing is two arc seconds or better. If we assume that the nearby supernova rate of one type Ia supernova per 500 galaxy-years (see section 4.1) is applicable to these distant galaxies, that there are approximately 60 useful galaxies (normalized to $10^{10}\ L_{Bsun}$) in each image, and that the surveillance time (the period of time during which, if a supernova exploded, it would be found in the current image) is approximately three weeks, then we would expect to find a supernova in every 140 images. This is comparable to the number of images we would expect to acquire in two and a half

good nights of observation. Of course, only about half of the supernovae will be on the rise.

After a night of observations on the AAT, the images are sent on the next plane from Sydney, Australia, to San Francisco, California, so they can be loaded into the computers at Lawrence Berkeley Laboratory for analysis. Due to travel time, connections, and customs, this trip takes on the average 48 hours from the time the last image has been taken. Once the images are loaded into our computers, it takes two days to search the images for candidates and another day to study the final candidates and make finding charts for those objects that are potentially supernovae. It typically takes five days from the time of initial observation until we are ready for follow-up observations. This delay could easily be reduced to three days or less if we had adequate computer resources at the telescope.

The analysis consists of subtracting, pixel by pixel, a previous reference image of the same field from the newly acquired images. Using convolution, this subtraction is then searched for residuals that would result from a supernova explosion. To eliminate false alarms due to asteroids and cosmic rays, the new image consists of two exposures taken 15 minutes apart. Asteroids can be recognized since they move several arc seconds in fifteen minutes. We can also eliminate cosmic rays, since it is unlikely that they will strike the same pixel in both new images.

We have been granted approximately eight nights a year by the time assignment committee of the Anglo-Australian Telescope. Of these, about half of the nights were clear. In addition, a couple of nights had very poor seeing and several nights were used to collect reference images. Thus, in the two and a half years this project operated at the AAT, there were only an equivalent of four nights during which we should have expected to find

supernovae. Unfortunately, we were unable to prove our ability to find supernovae with this system.

This pilot project failed because we were unable to determine whether or not the variable objects we found were supernovae. There were two reasons for this. First, we were unable to distinguish faint stellar objects from galaxies at a redshift of 0.3 due to the two arc second seeing at the AAT. This made it difficult to distinguish supernovae from the background of variable stars, QSO's, and BL Lacertae objects. Second, we found it difficult to schedule follow-up observations for these faint objects due to weather constraints, moon constraints, and constraints of telescope time assignment. When we were able to obtain follow-up observations (spectra or high resolution images) of our candidates, we found that they were all background sources.

The next generation instrument will avoid the difficulties encountered with this one. The search has been moved to the Isaac Newton Telescope in the Canary Islands. This site is much more conducive to a distant supernova search due to better weather and the availability of large blocks of telescope time. The one arc second seeing makes star-galaxy separation possible, so it will be easier to distinguish the majority of the background variable objects. The large blocks of time and good weather make it possible to schedule follow-up observations with greater certainty. In addition, we have targeted high redshift galaxy clusters, which should increase the number of useful galaxies per image by at least a factor of two.

The remainder of this paper covers the experimental design, analysis techniques, and results of our supernova search on the Anglo-Australian Telescope.

2. Hardware

2.1. Optical Design

The optics for our f/1 camera were designed by Harvey Richardson (Richardson and Morbey 1988) to put as much sky area as possible onto a 1024x1024 pixel Thomson CCD chip. The camera is mounted in the prime focus top end of the 3.9m Anglo-Australian Telescope. (Fig. 2.1). The secondary mirror is a concave hyperboloid which transforms the f/3.25 beam from the primary mirror into an f/1 beam, thus increasing the sky coverage of our CCD chip by a factor of 10.6 over prime focus. The optics could not be made any faster, nor the field any larger, due to the physical constraints of placing correcting optics and a filter/shutter apparatus between the secondary mirror and the focal plane. The equivalent focal length of an f/1 beam from a 391 cm mirror is 391 cm. Therefore, the image scale (scale = 1/f) is 8.8 arc minutes per centimeter. Since the physical dimensions of our CCD are 1.9 cm x 1.9 cm, this gives us a theoretical field size of 16.8' by 16.8' with a pixel scale of 0.98"/pixel.

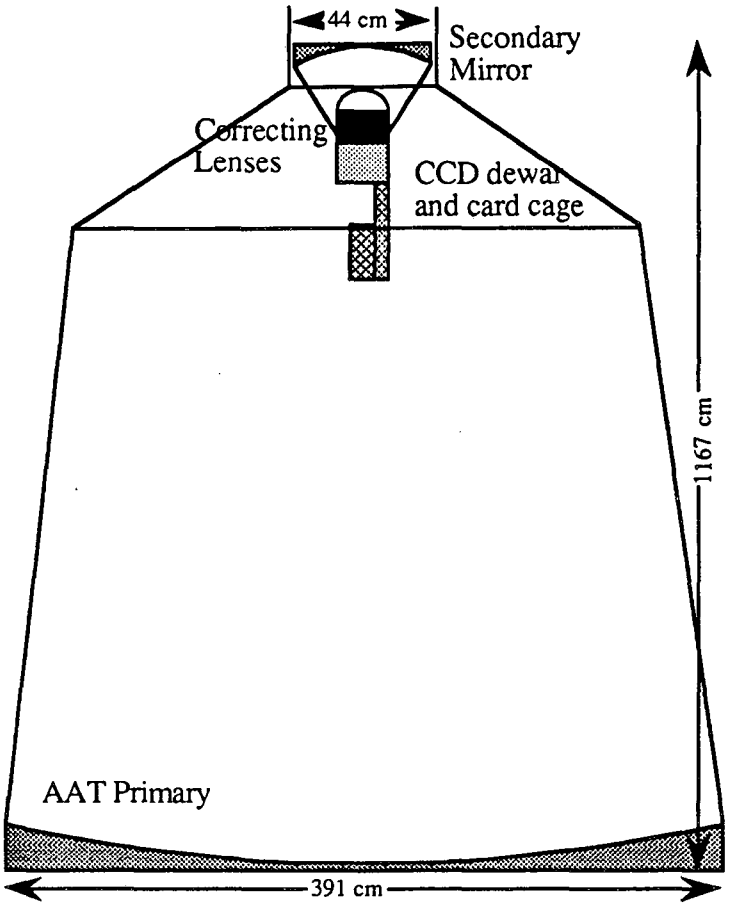

Fig. 2.1. The AAT f/1 system. The secondary mirror, correcting lenses, and CCD constitute the f/1 camera, which sits at the prime focus of the 3.9 meter AAT. This diagram is not to scale.

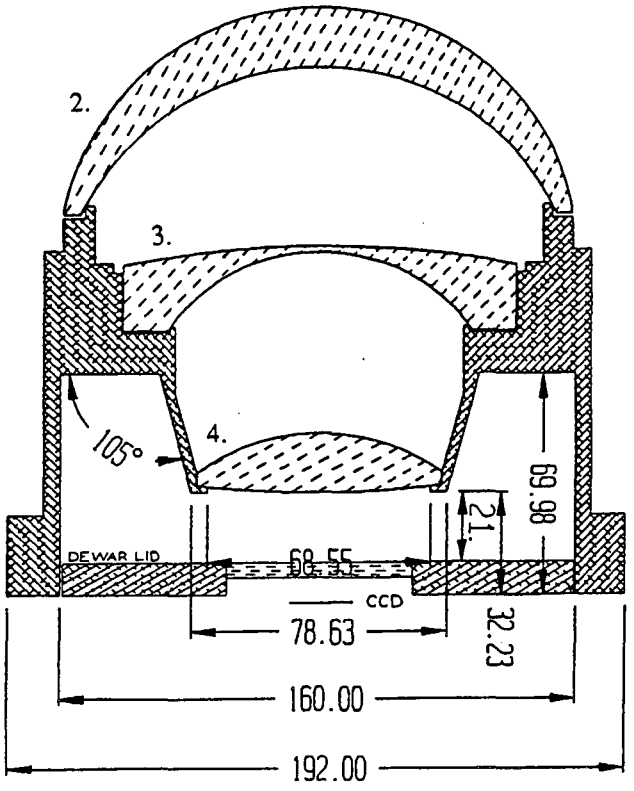

Fig. 2.2. The f/1 focal corrector. This scale drawing (in mm) of the focal corrector shows the three correcting lenses, the dewar window, and the position of the CCD chip itself. The filter-shutter apparatus was designed to fit in the small space (21. mm) between the last lens element and the CCD dewar lid.

1. Concave hyperboloidal mirror:
   Radius of curvature = 977.3
   Conic constant = - square of eccentricity = -2.67
   Clear diameter = 388.

2. Positive lens
   Radius of curvature = 80.72 (convex)
   Thickness = 19.3
   Radius of curvature = 82.09 (concave)
   Clear diameter = 155.

3. Negative lens
   Radius of curvature = 303.3 (convex)
   Thickness = 2.0
   Radius of curvature = 56.87 (concave)
   Clear diameter = 116.

4. Positive lens
   Radius of curvature = 62.79 (convex)
   Thickness = 18.6
   Radius of curvature = -314.3 (convex)
   Clear diameter = 72.

Lens material: BK7 glass
Tolerances: 0.5 wave on mirror, 1 wave on lenses (wavelength = 6328 Angstroms)

Fig. 2.3. Optical parameters for elements of the f/1 focal reducer (in millimeters)

The construction of the optical system (Fig. 2.3) was done by Applied Physics Specialties of Toronto, directed by Martin High. The construction of the supporting structure and correcting lenses was accomplished without a hitch. Because of the low cost and excellent thermal properties of a solid aluminum alloy mirror, this material was chosen over the more traditional glass construction. However, the secondary mirror turned out to be the most troublesome element of the optical system.

The mirror blank is mounted for cutting by fastening three bolts imbedded in the back side through corresponding holes in a rotating vertical plate. A diamond tipped knife is positioned at the center of the rotating mirror, then slowly moved towards the edge, cutting the mirror to the desired shape in the process. Surprisingly, the antiquated computer that calculates the position of the blade does not calculate the positions to as high an accuracy as the blade can be controlled. This results in very tiny annular ridges.

The first time the mirror was cut, the overall figure was not within tolerances and the optics could not pass the Hartmann test at the telescope. The f/1 camera alone would have caused aberations of 1.5 arc seconds in the final image. The second time, the figure was purportedly excellent. However, some cleaning fluid that had inadvertently been left on the diamond turning machine etched a large annulus of the mirror during transport, so the mirror had to be returned for a third cut. The third time, we received an acceptable mirror with a measured reflectivity of 81%. This final mirror was not coated with aluminum or any protective layers. The surface was left to form its own oxide layer, which is relatively transparent in the visible.

a)

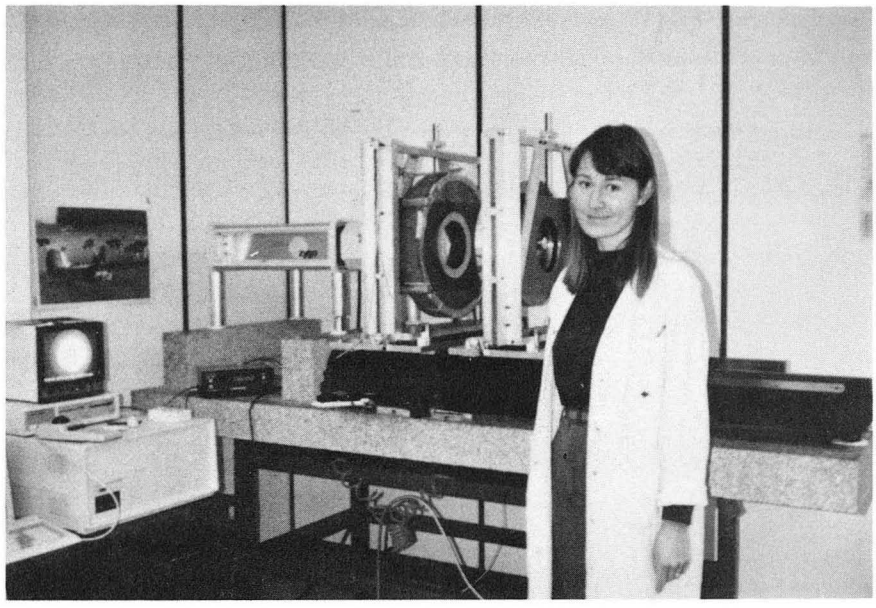

b)

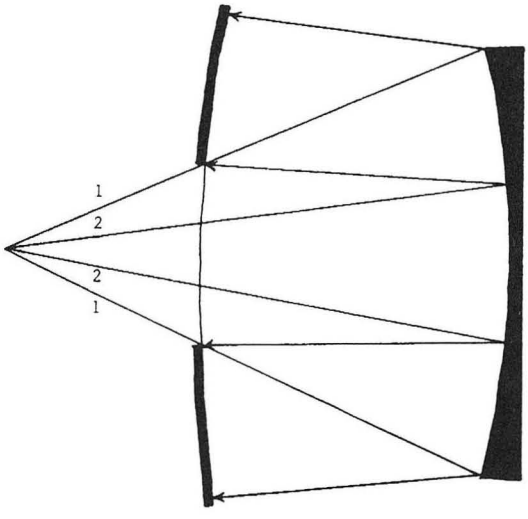

c)

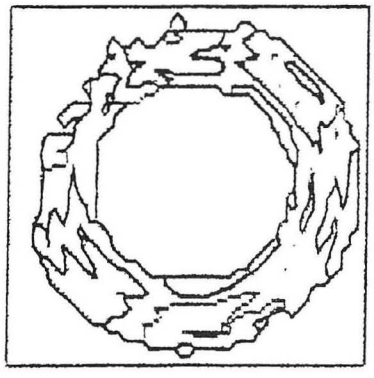

d)

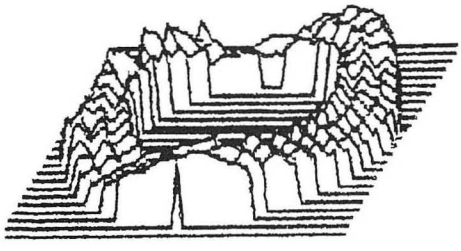

Fig. 2.4. Hindle test. a) The successful interferometer test of the hyperboloidal mirror. b) A scaled ray trace of the Hindle test. A point source is placed at the real focus of the hyperboloidal mirror. It is imaged to a virtual focus behind the hyperboloid. A spherical test mirror is positioned so that its center of curvature coincides with the virtual focus of the hyperboloid. Thus, the rays are relected back along the same path, creating interference fringes. The outermost (1) and innermost (2) rays are shown. c) A contour plot of the corrected deviations from a perfect hyperboloid. d) A surface plot of the corrected mirror deviations. The peak to valley deviation is 0.371 microns.

Figure 2.4 shows the results of the Hindle test on the final mirror surface. These results were corrected for defocusing, which means the best fit spherical wavefront deviation was subtracted in addition to a best fit planar wavefront deviation. Then, the calculated deviation from the spherical test mirror was subtracted and the magnitude of the wavelength discrepancies was divided in half to take into account two bounces off the hyperboloidal mirror.

Notice that the mirror surface has three high regions separated by 120° (Fig. 2.4 d). This astigmatism was probably caused by the three bolts by which the mirror is mounted. It is hard to tell how the mirror shape changes when the mirror is hung facing down rather than facing sideways, as it was for the Hindle test. Some of the spots near the edges were not registered because these fringes moved during the test. The interference test shows the mirror to be a maximum of 0.587 wavelengths (0.371 μm) out of figure, peak to valley. The calculated spot size due to this error is 3.164 μm (Fig. 2.5). This is very small compared to the best seeing expected at the AAT, which is 19 μm (1 arc second).

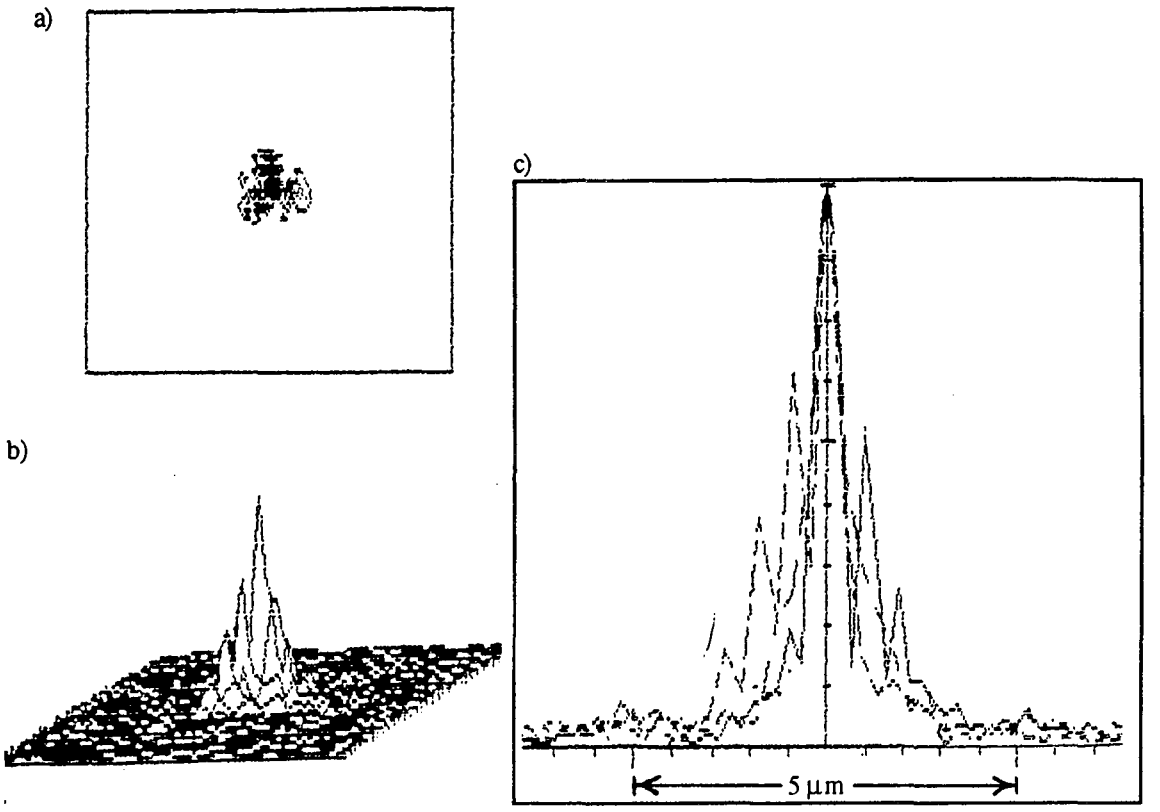

Fig. 2.5. Spot size from Hindle test. This shows the calculated image spot size from the mirror surface found in the Hindle test. a) Calculated spot, showing the three-point structure of the astigmatism. b) Contour of the calculated spot. c) Plot of four slices through the calculated spot. The vertical scale is intensity.

2.2. Filter

There is no part of the optical path in this camera in which the beam is parallel. Therefore, there is no place to insert filters where their surfaces can be perpendicular to every light ray. Any plate of glass inserted into the beam or removed from the beam will cause a distortion at the focal plane. The optical path was designed to include exactly 0.8 cm of glass in the f/1 beam between the last correcting lens and the CCD. Five millimeters are designated for the CCD window and three millimeters are designated for a filter. One of the difficulties of the f/1 camera design was the filter-shutter apparatus, which had to fit in the 2.1 cm gap between the last corrector lens and the window of the CCD camera (Fig.

2.2). There is not enough room within the lens case to move a filter entirely out of the beam in order to slip another one in. A compromise was eventually reached wherein the filters vignette the beam both when they are in the beam and when they are removed, but the vignetting is not severe in either case. The filter-shutter system that was designed by Robert Smits and improved upon by the technical staff of the Anglo-Australian Observatory consists of two levels of moving parts. The lower level contains one filter holder and a round metal plate that is used as a shutter. The upper level contains two additional filter holders. Each of these can be flipped in and out of the beam by energizing one of four solenoids. The filters are temporarily glued into the filter holders before each run, and they cannot be changed except by full disassembly of the instrument.

The best filter to use to look for distant supernovae gives the best signal-to-noise for a SN Ia at a redshift of $z = 0.3$. There is no profit from using a standard filter when looking so far away that many blue photons are redshifted into the visual range. The *K*-corrections (which take into account the light redshifted into or out of a filter) will have to be computed for each supernova at its individual distance, using the wavelength response of our entire system.

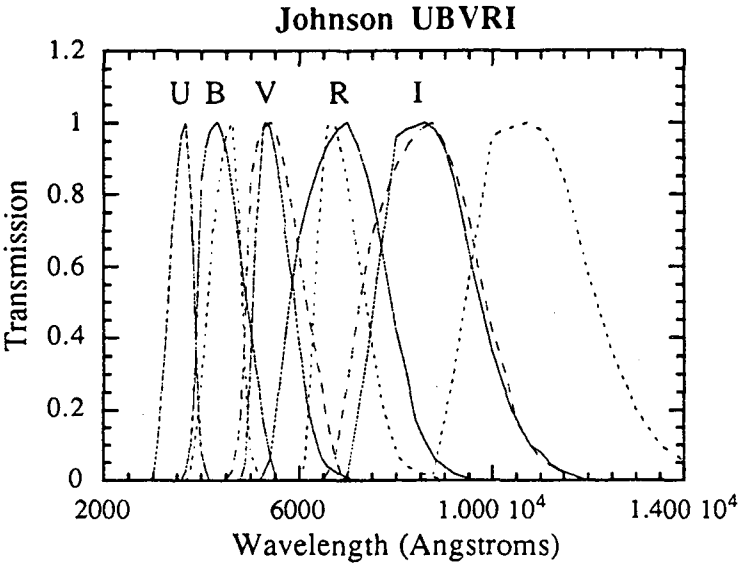

Fig. 2.6. Johnson filters. The solid lines show the transmission curves for the UBVRI Johnson passbands. The dotted lines show the wavelengths to which the photons which would have been detected in each filter are redshifted for z=0.25. The photons in U are redshifted into B, B photons are redshifted into V, etc.

Unfortunately, it is not well known what the colors of a SN Ia at redshift z = 0.3 will be as a function of time. However, due to a fortuitous circumstance, a redshift of z = 0.25 brings one Johnson filter into the next (Fig. 2.6). Using the colors of a SN Ia at maximum from Leibundgut (1991) and the absolute blue magnitude from Leibundgut (1990), we can estimate the apparent magnitudes of a SN Ia at redshift of z = 0.25. Using this information, we can evaluate the relative effectiveness of each of the Johnson filters for finding supernovae at z ≈ 0.25. Although it is not a rigorous calculation of the optimal filter, this exercise illustrates the importance of choosing the right filter.

To calculate the apparent magnitude in the visual, for example, we estimate the flux in the blue at 10 pc to be

$$F_B = 10^{-0.4(M_B - q_B)}\left(\Delta\lambda_{1/2}(B)\right)$$

$$= 4.61 \times 10^{-5} \text{ watts/cm}^2, \qquad (2.1)$$

where $\Delta\lambda_{1/2}(B)$ is the full width at half maximum of the B filter. The normalization, q_B, was calculated from Hellwege (1982), and the absolute blue magnitude, M_B, is taken to be -19.6 ($H_o = 50$ km/sec/Mpc). The flux is multiplied by the portion that will fall within the range of the visual filter (the full width at half maximum of the visual filter divided by the redshifted full width at half maximum of the blue filter). The flux in the visual that reaches the earth from a SN Ia at $z = 0.25$ is then

$$F_V = F_B\left(\frac{\Delta\lambda_{1/2}(V)}{1.25\Delta\lambda_{1/2}(B)}\right)\left(\frac{10\text{pc}}{d_L(z=0.25, H_o = 50\text{km/sec/Mpc})}\right)^2$$

$$= 1.18 \times 10^{-21} \text{ watts/cm}^2. \qquad (2.2)$$

The V magnitude is then obtained from

$$F_V = 10^{-0.4(V - q_V)}\left(\Delta\lambda_{1/2}(V)\right). \qquad (2.3)$$

Using similar calculations to find B and R, we find

$$B = 21.9 \text{ mag},$$

$$V = 21.0 \text{ mag},$$

$$\text{and } R = 20.8 \text{ mag}, \qquad (2.4)$$

for a supernova Ia at a redshift of $z = 0.25$. The visual magnitude is in reasonable agreement with the published results of Leibundgut (1990).

From these magnitudes, we can calculate the number of photoelectrons accumulated in the detector. We correct the magnitudes for reddening in our galaxy using $A_B = 0.20$ mag, $A_V = 0.15$ mag, and $A_R = 0.11$ mag (Hellwege 1982). Then, we calculate the signal from

$$S_B = \varepsilon q A t\left(\Delta\lambda_{1/2}(B)\right)\left(\frac{\lambda_B}{hc}\right)10^{-0.4\,(B + A_B - q_B)}, \qquad (2.5)$$

where ε is the efficiency of the optical system, q is the quantum efficiency of the CCD, A is the area of the telescope primary in cm^2, and t is the length of the exposure. For the three bands,

$$S_B = 0.00211\ \varepsilon qAt \text{ electrons},$$
$$S_V = 0.00283\ \varepsilon qAt \text{ electrons},$$
$$\text{and } S_R = 0.00532\ \varepsilon qAt \text{ electrons}. \quad (2.6)$$

If we assume that the noise is primarily due to the sky background and little is due to the parent galaxy, then

$$N_B = \sqrt{\varepsilon aqAt\left(\Delta\lambda_{1/2}(B)\right)\left(\frac{\lambda_B}{hc}\right)10^{-0.4\,(m_B - q_B)}}, \quad (2.7)$$

where m_B is the magnitude per square arc second of the sky and a is the area of the photometry aperture in square arc seconds. For the following calculation, we use a circular aperture of radius 2.5 arc seconds. At the AAT, the dark sky brightness (at new moon) is approximately B = 22.5 mag, V = 21.5 mag, and R = 20.5 mag (the Johnson R magnitude was estimated from the Kron-Cousins R magnitude, and comparison with Turnrose 1974). Plugging into equation 2.7, we find

$$N_B^2 = 0.0290\ \varepsilon qAt \text{ electrons}^2,$$
$$N_V^2 = 0.0402\ \varepsilon qAt \text{ electrons}^2,$$
$$\text{and } N_R^2 = 0.152\ \varepsilon qAt \text{ electrons}^2. \quad (2.8)$$

The constant εAt is the same for all three filters. The quantum efficiency, q, of our detector is about 0.40 in V and R, and 1/3 as large in B. So, the signal-to-noise ratio for the filters and all possible combinations is

$$\left(\frac{S}{N}\right)_B = 0.0046\ \sqrt{\varepsilon At},$$
$$\left(\frac{S}{N}\right)_V = 0.0089\ \sqrt{\varepsilon At},$$

$$\left(\frac{S}{N}\right)_R = 0.0086\ \sqrt{\varepsilon A t},$$
$$\left(\frac{S}{N}\right)_{B+V} = 0.010\ \sqrt{\varepsilon A t},$$
$$\left(\frac{S}{N}\right)_{V+R} = 0.012\ \sqrt{\varepsilon A t},$$
$$\text{and}\quad \left(\frac{S}{N}\right)_{B+V+R} = 0.012\ \sqrt{\varepsilon A t}. \tag{2.9}$$

One can see from these results that the best wavelength to search for supernovae is in the neighborhood of the V (or R) filter, and that wider filters give higher signal-to-noise ratios.

It is interesting to compare these new-moon results with signal-to-noise ratio estimates for nights with a quarter moon. For this case, we use estimated sky magnitudes at the AAT of B = 21.3 mag, V = 20.8 mag, and R = 20.2 mag. The signal-to-noise ratio in the various wavelength ranges is

$$\left(\frac{S}{N}\right)_B = 0.0026\ \sqrt{\varepsilon A t},$$
$$\left(\frac{S}{N}\right)_V = 0.0065\ \sqrt{\varepsilon A t},$$
$$\left(\frac{S}{N}\right)_R = 0.0075\ \sqrt{\varepsilon A t},$$
$$\left(\frac{S}{N}\right)_{B+V} = 0.0061\ \sqrt{\varepsilon A t},$$
$$\left(\frac{S}{N}\right)_{V+R} = 0.0098\ \sqrt{\varepsilon A t},$$
$$\text{and}\quad \left(\frac{S}{N}\right)_{B+V+R} = 0.010\ \sqrt{\varepsilon A t}. \tag{2.10}$$

Since the moon scatters more light into the B region of the spectrum, the optimal wavelength for finding supernovae has lengthened slightly. However, the broadest band-pass filter is still preferred. The combination of the V and R passbands is consistently at least 30% more efficient than either one separately. Since the signal-to-noise ratio is

proportional to the square root of the exposure time, one would have to expose at least 70% longer in either V or R alone.

At the time this supernova search was begun, it was recognized that the best wavelength range in which to search for supernovae at high redshift was somewhere between V and R. Accordingly, a special filter was created that passed V and the Kron-Cousins R, which is about half as wide as the Johnson R. In retrospect, we probably should have used most, if not all, of the wavelength sensitivity of our CCD detector.

2.3. CCD camera

Our photon detector is a 1024 x 1024 pixel liquid nitrogen cooled Thomson CCD (charged coupled device). It is sensitive to wavelengths in the range 450-900 nanometers with a maximum quantum efficiency of about 40% (Fig. 2.7). The electronics for the chip were expertly built and tested by John Barton of the Anglo-Australian Observatory. The CCD has a gain of five electrons/ADU (Analog to Digital Unit), ten electrons of readout noise, and can be read out in approximately 30 seconds. The hold time of the dewar is about six hours, so it generally has to be filled once in the middle of the night.

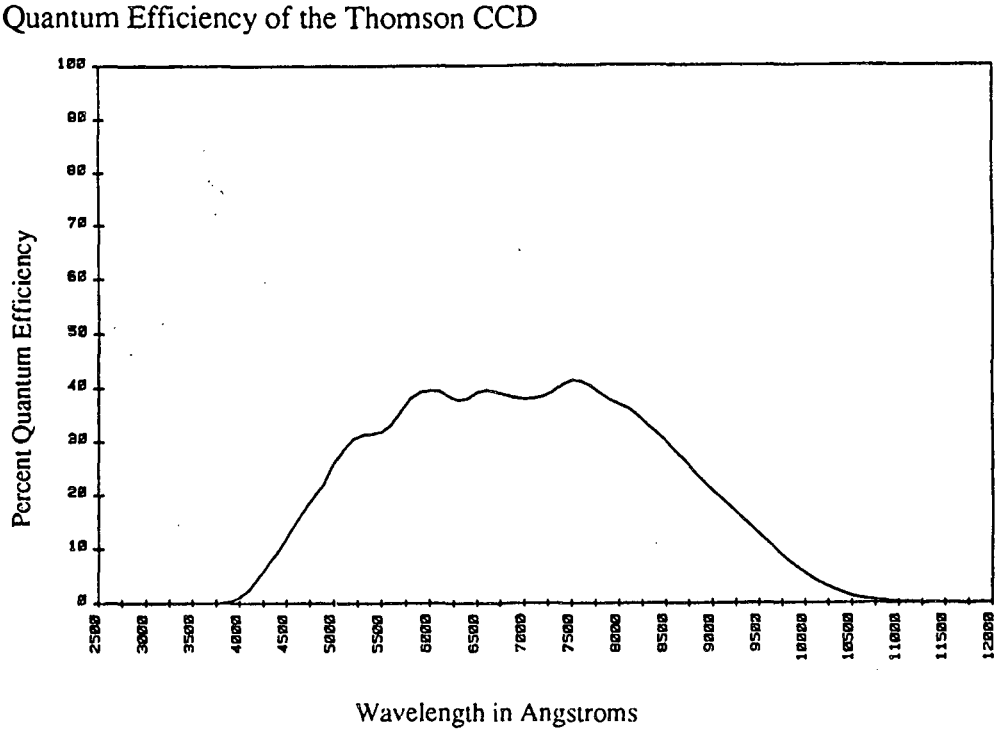

Fig 2.7. Quantum efficiency of the Thomson CCD. The quantum efficiency is the number of photoelectrons produced per incident photon. The bare CCD has a quantum efficiency of greater than 20% in the range 4750-9250 Angstroms, with a maximum quantum efficiency of 40%.

This chip has functioned very well. The bias frames are remarkably flat (Fig. 2.8). The only features are noise, a small periodic ripple (due to 50 Hz standard AC), and cosmic rays. In a study of ten bias frames, we found one cosmic ray of 11,000 electrons and nine cosmic rays in the range 1000-2500 electrons each. The peak to valley of the ripple is 25 electrons. The ripple doesn't stay in the same place from frame to frame, so it cannot be removed by bias frame subtraction. Historically we have had some trouble with the bias level changing in the middle of the night. Since there is no advantage to subtracting the actual bias frame, we subtract from each pixel the average value of the overscan region of the CCD.

Fig. 2.9 shows a CCD image with the bias level subtracted. The image shows the 20-25% vignetting of the corners due to the optics and the filter/shutter apparatus. It also

shows some CCD defects – notably the palm-tree shaped feature in the lower right corner. These defects are not severe; the palm-tree feature has a deficit of only 2%. These features disappear upon flat-fielding (Fig. 2.11).

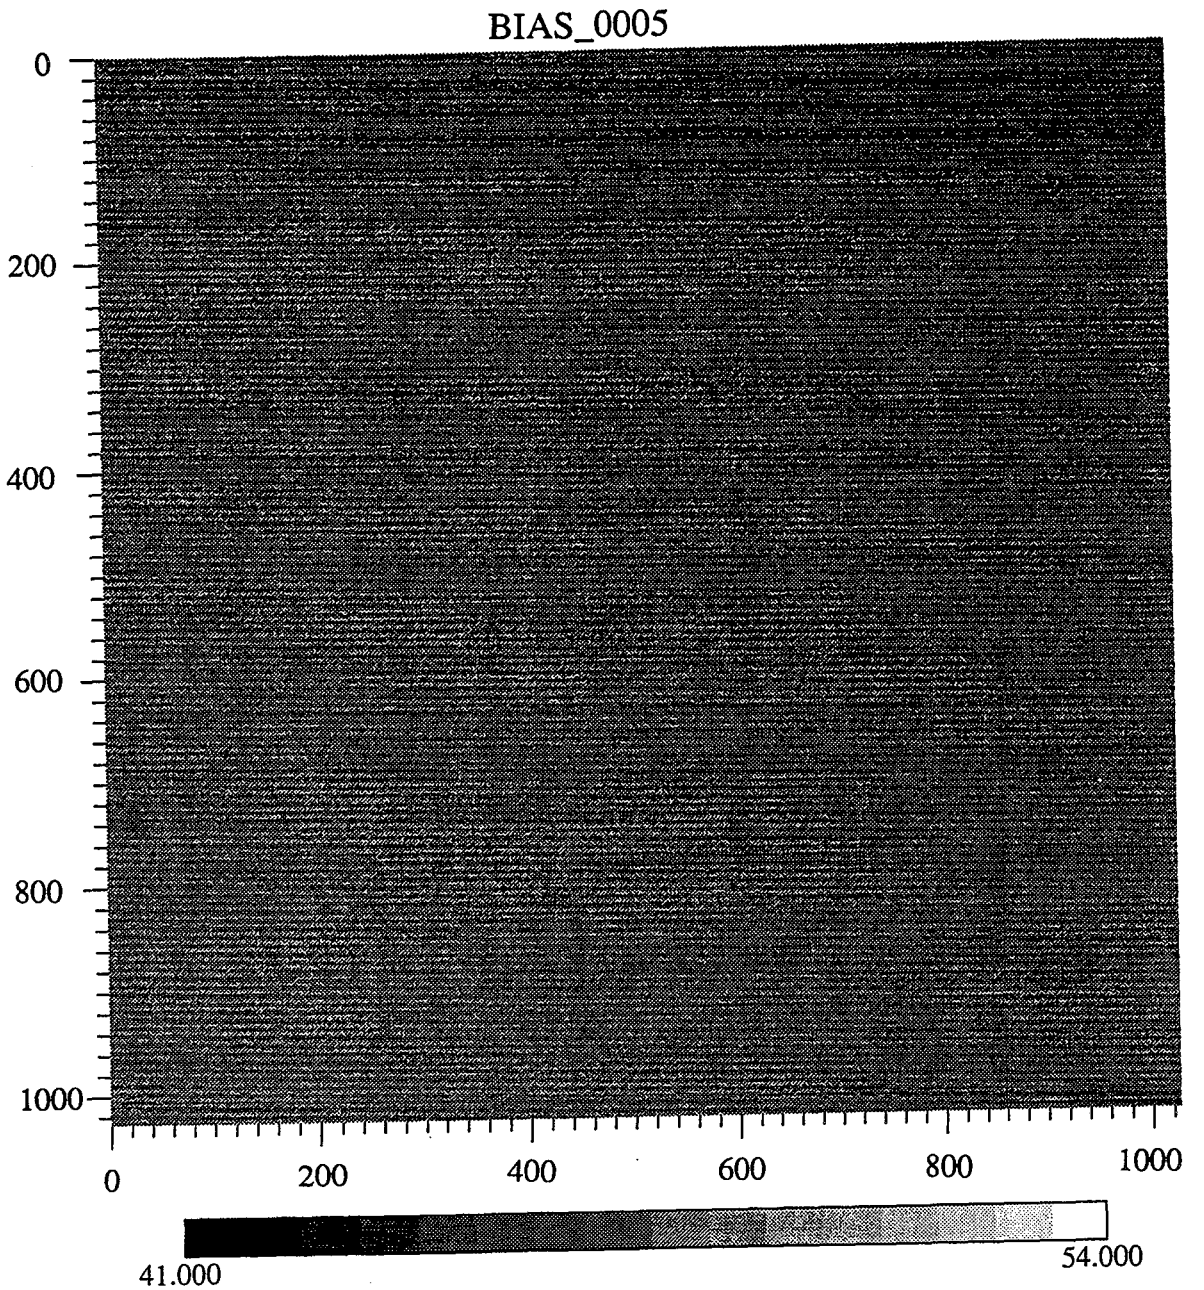

Fig. 2.8. Bias Frame. This bias frame (a 0–second image taken with the shutter closed) was taken on Dec 1, 1989. It shows the 50 Hz ripple (Australia has a 50 Hz standard AC, rather than 60 Hz as we have in the US), and one cosmic ray in the lower right quadrant. The grey scale is given in units of ADU.

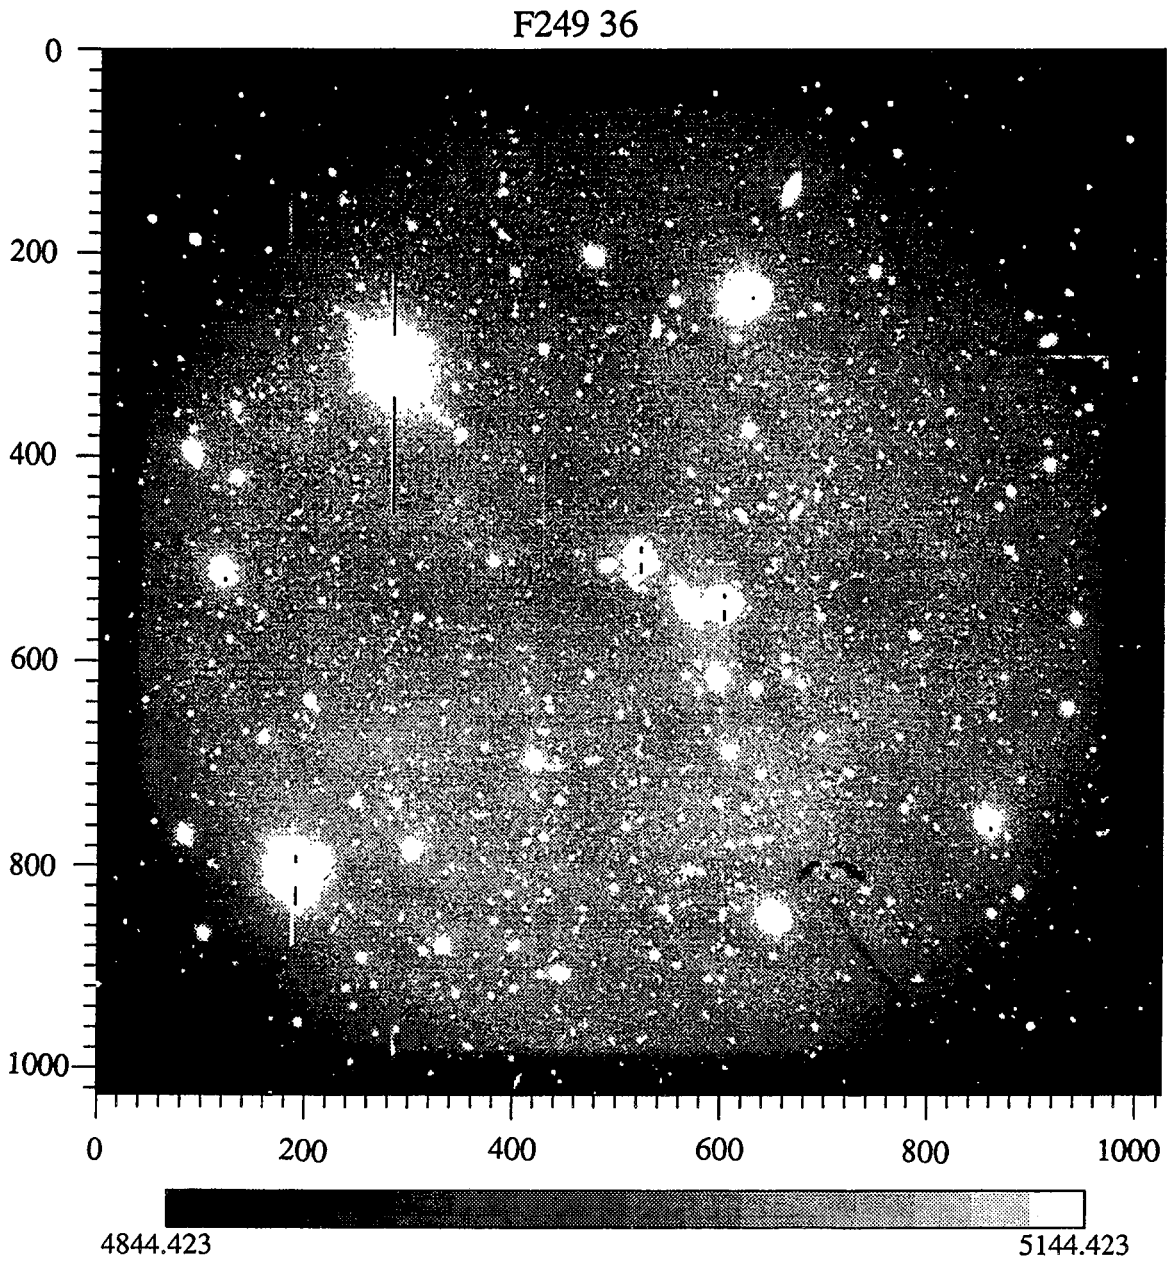

Fig. 2.9. Raw image. This is a raw five minute CCD image from November 30, 1989. There are a few CCD defects such as the "palm tree" in the lower right quadrant and a few darker pixels. Some of the vignetting in the corners is due to the optics, but most is due to the filter/shutter apparatus. Since November, the filter/shutter apparatus has been slightly redesigned so that the vignetting in the corners is 20% rather than 25% as it is here.

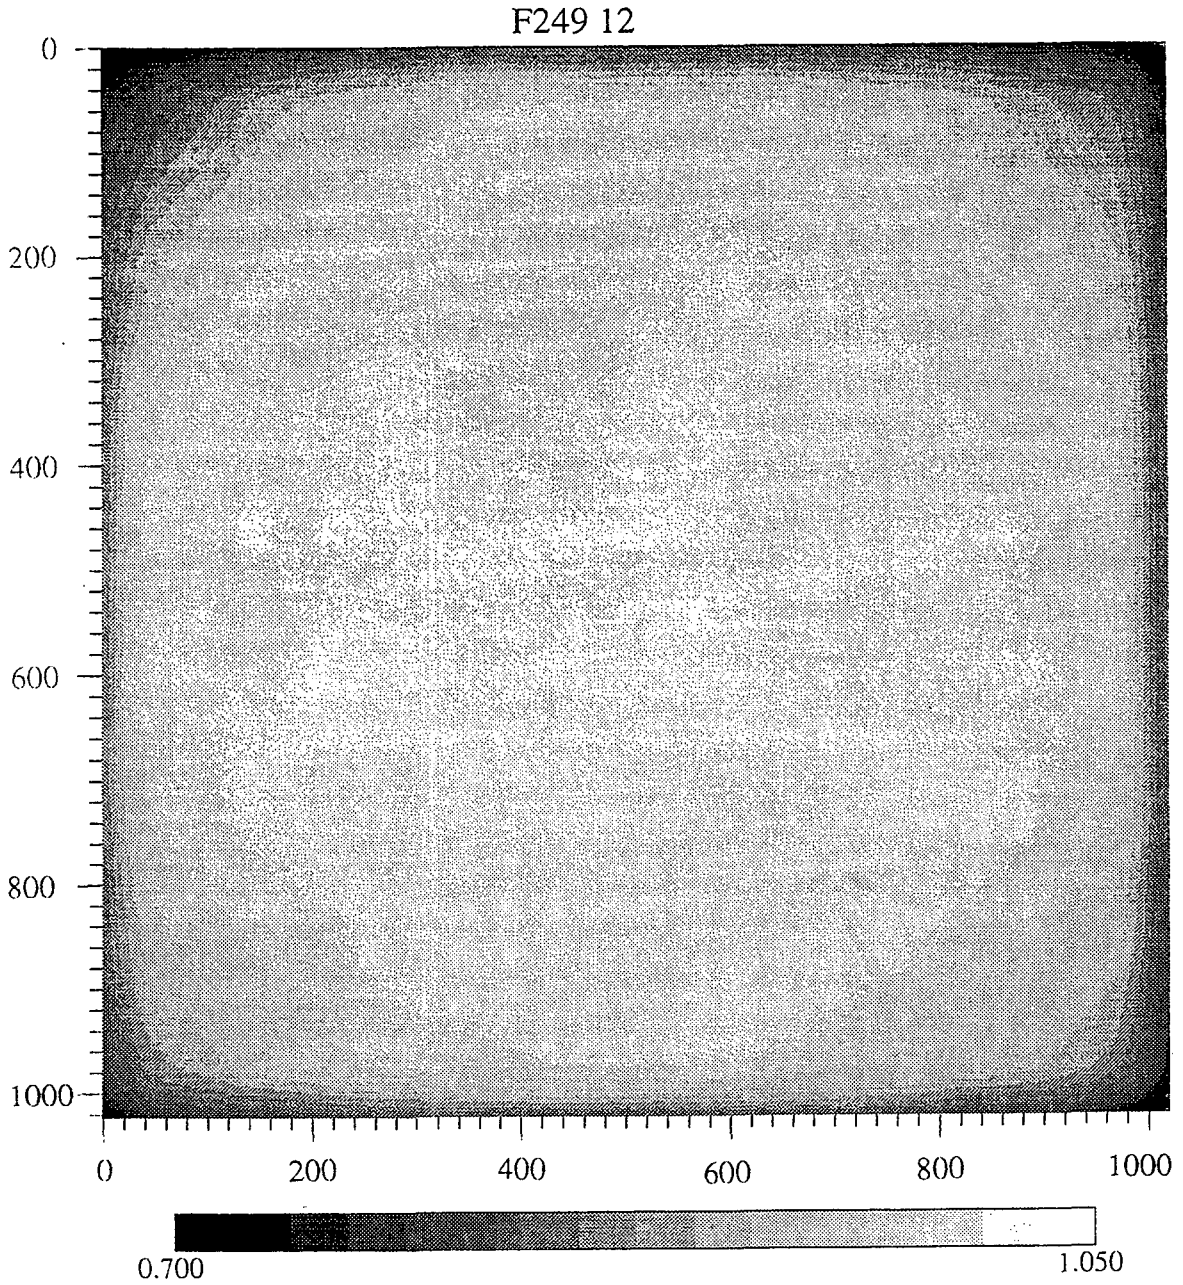

Fig. 2.10. Flatfield. This flat field was generated from the clipped mean of 15 five minute images taken November 30, 1989. All of the effects of the CCD and optics are included in this image, with none of the stars and galaxies.

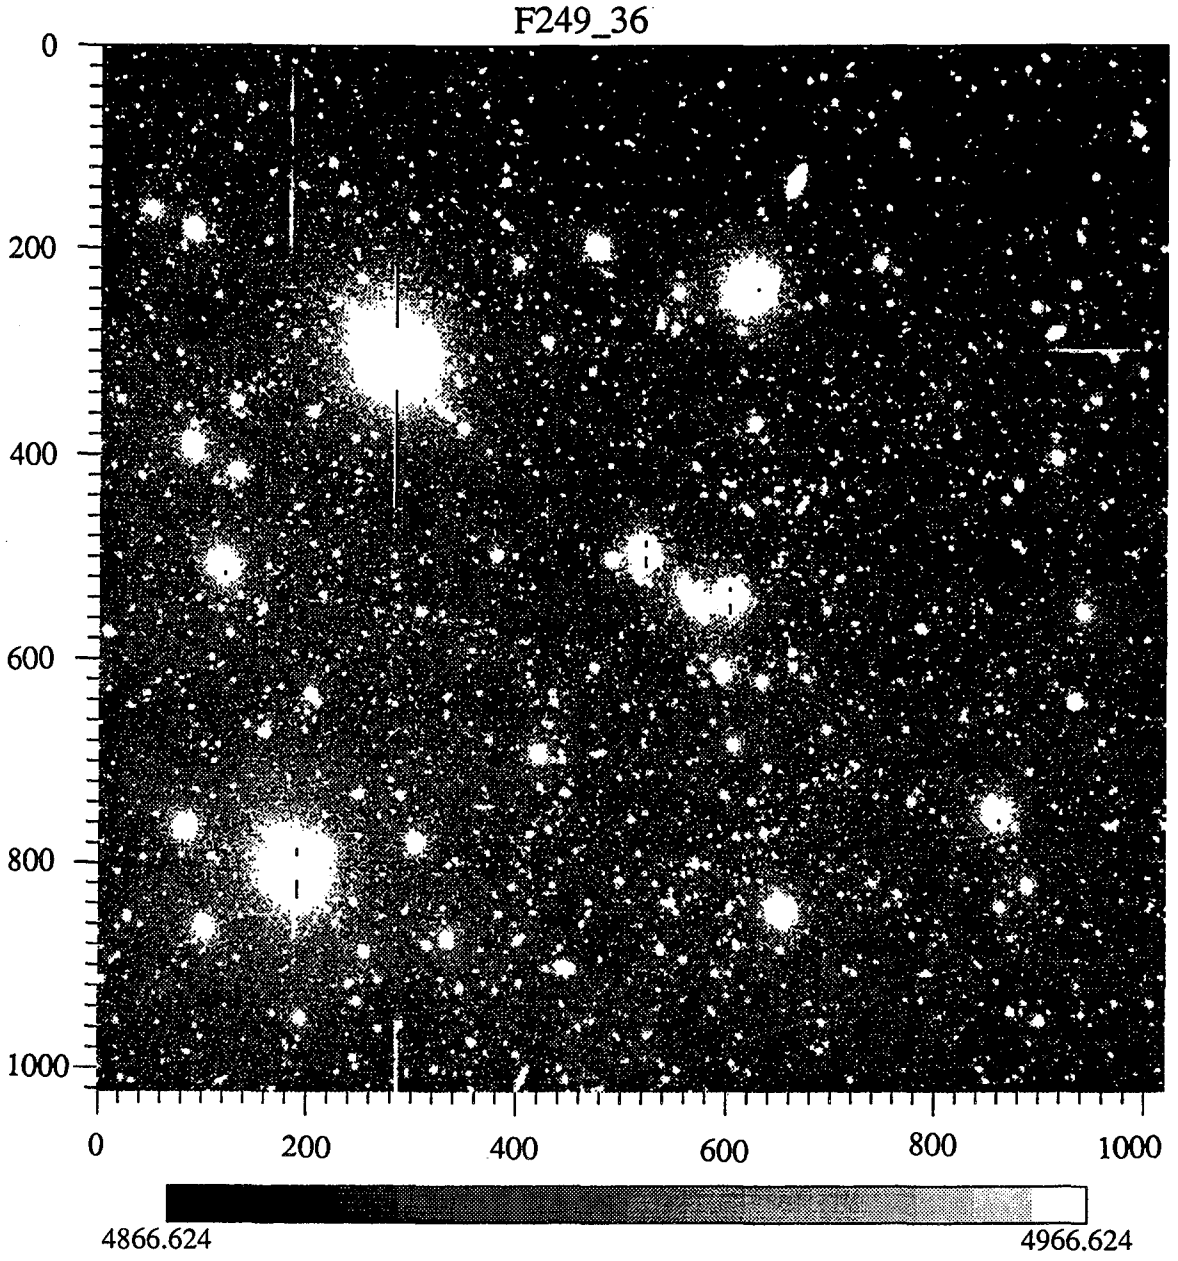

Fig. 2.11. Flatfielded image. This is the flatfielded image of Fig. 2.9.

We have experimented with different ways to make flatfield frames, and settled on a variant of an algorithm suggested by Brian Boyle. For each night, we add together the two 2.5 minute images for each of fifteen different fields. Next, we subtract the bias level (computed from the overscan region) from these images and normalize them so that the central 350 x 350 pixel area of each has the same sky value (where the sky value is determined from the peak of the histogram of pixel values). Then, we calculate the clipped mean value (as defined below) for each pixel. The subroutine that calculates the clipped mean flatfield frame must be supplied with the one sigma pixel noise (*sig*), the number of sigma at which the mean should be clipped (*factor*), and the number of iterations it should use to calculate the median (*iter*). For each pixel of the CCD, the subroutine calculates the median value from fifteen fields. Then it eliminates those pixel values that are farther than *factor*sig* above the median or 2**factor*sig* below it. (The asymmetric interval was chosen because experience showed it gave the best flatfield.) This process of calculating the median and eliminating significantly deviant values is repeated *iter* times. Finally, those values that remain are averaged to find the flatfield value for that pixel. For diagnostic purposes, the average number of pixels used to calculate the mean is printed to the screen.

This method of computing the flatfield has the advantage that it uses images with the actual color of the night sky, rather than the color of dusk or reflected light from the telescope dome. In addition, the CCD is calibrated at the intensity level most useful for our images. There are two pitfalls of this method that must be avoided. First, the effects of the galaxies and stars must be minimized by choosing the appropriate *factor* and *iter*, and by using enough images. For our images, we *iter*ate twice with a *factor* of 1.5. This method cannot be used with very crowded images in which it is likely that any given pixel will be associated with a star or galaxy. If the first median is more than 2**factor*sig* above the optimal value for that pixel, then it is likely that this algorithm will not converge to the

optimal value. It is not hard to show that the approximate number of pixels for which this will happen is given by

$$N \sum_{k=0}^{(n-1)/2} \frac{n!\, p^k (1-p)^{n-k}}{(n-k)!\, k!}, \tag{2.9}$$

where N is the number of pixels in the CCD, n is the number of fields, and p is the probability that any given pixel is below 2**factor*sig* above the local sky value. If n is even, the upper limit of the sum should be rounded up (down) if the median is chosen as the higher (lower) of the two middle values. Since the stars and galaxies increase the values of clumps of pixels, the errant pixels may be highly correlated and may even have somewhat stellar profiles.

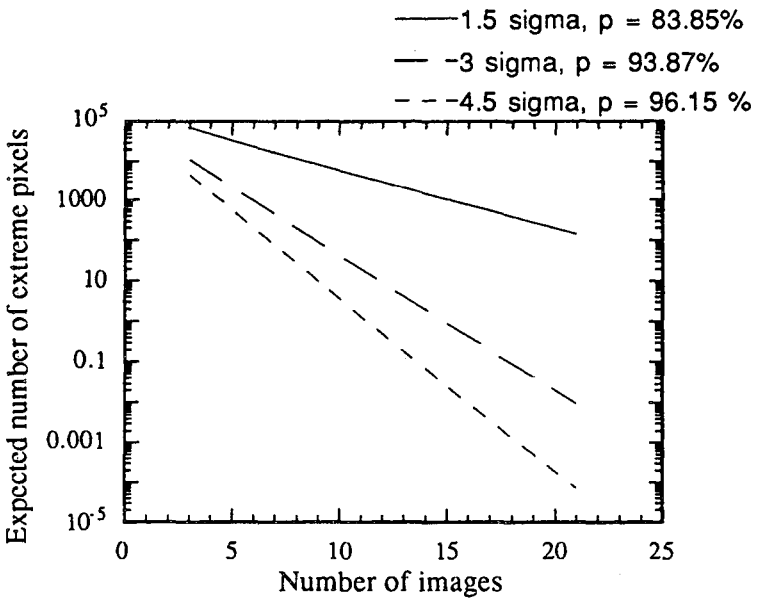

Fig. 2.12. Expected number of flatfielding errors. This graph shows the number of pixels for which the first calculation of the median value will be at least 1.5, 3, and 4.5 sigma high. If the median is 1.5 sigma high, the algorithm will usually recover, but statistical accuracy decreases. If the median is 3 sigma high, the resulting value in the flatfield will probably be high, but may not be noticed in the noise. If the first median is 4.5 sigma above the optimal value for that pixel, it will probably remain high and result in observable flatfielding errors. Since p is close to 1 for our data, the distribution is almost Poisson, which gives us the linear relationships in this plot.

In our sample image, 96% of the pixels are below 4.5 σ above background, 94% are below 3 σ above background, and 84% are below 1.5 σ above background. These values are typical for our fields. Fig. 2.12 shows the consequences of these figures on the first computed median of our clipped mean process. For our 15 images, we expect none of the first medians to be over 4.5 σ high. Only one pixel is expected to have a first median that is 3 σ high, but after the second median, this value will probably be only ~1.5 σ high. 1064 (~0.1%) pixels will have a first median that is 1.5 σ high. The second loop will bring most of these to within 0.5 σ of the optimal pixel value. Using more images to compute the flatfield would make it more robust, but the cost in computing time becomes prohibitive.

The second pitfall is that enough images must be used that the flatfield frame does not significantly affect the noise. We assume that the photons from the night sky fall in a Poisson distribution, and that the CCD will produce a photoelectron with a probability given by its quantum efficiency each time it is struck by a photon. Then, the one sigma pixel noise in a raw image is

$$\sigma = \sqrt{\sigma_{ro}^2 + \frac{sky}{gain}}, \tag{2.10}$$

where σ_{RO} is the readout noise in ADU, sky is the number of ADU in that pixel, and the gain is the number of electrons per ADU. If we create the flatfield frame by averaging N images and then dividing by the background ADU in the center of the CCD, then the noise in the center of the quantum efficiency (flatfield) frame is

$$\sigma_{QE} = \frac{\sigma}{sky_{QE}\sqrt{N}} = \frac{\sqrt{\left(\sigma_{ro}^2 + \frac{sky_{QE}}{gain}\right)/N}}{sky_{QE}}. \tag{2.11}$$

We introduce no noise by subtracting a constant rather than a dark frame. So, the total noise in a flatfielded image is

$$\sigma_{TOT} = \sqrt{\sigma_{ro}^2 + \frac{sky}{gain} + \sigma_{QE}^2 sky^2}\,. \qquad (2.12)$$

For our case, the readout noise is small compared to the Poisson sky noise and sky ≅ sky_{QE}, so

$$\sigma_{QE} = \frac{\sigma}{sky\sqrt{N}}, \text{ and} \qquad (2.13)$$

$$\sigma_{TOT} = \sigma\sqrt{1 + \frac{1}{N}}\,. \qquad (2.14)$$

For example, if we wanted the flatfield frame to add less than 5% extra noise, we would need $N > 9$. This means that 9 pixels must be averaged, not including those clipped out in the clipped mean process. For our example flatfield, 11.7 ± 1.6 pixels were averaged; our estimate of the increase in the noise during flatfielding is 4%.

In general our images flatfield beautifully (Fig. 2.11). The only residual local defects are caused by saturation of the CCD at about 40,000 ADU. When a pixel reaches saturation, it overflows along the column. These overflowed columns sometimes have an extremely low or even negative adjacent column. This is caused by overshoot in the CCD electronics due to the huge potential drop between the saturated pixel and the normal sky background pixel read out immediately afterwards.

2.4. Performance

Figure 2.17 shows the components of the f/1 camera being installed for the first time on the 3.9m Anglo-Australian Telescope. Using this system, we have imaged 12.5 square degrees on the sky with an average of 2 repeat images per field. Some of the fields have no repeat images, and some have up to 9 epochs.

Since our search does not use a standard filter, it takes some effort to calibrate the magnitudes of our objects. We have calibrated our system against the UBVRI Kron-Cousins system standard stars of Graham (1982). The response of our filter plus CCD camera is very close to the sum of the V plus R response of that system (Fig. 2.13). Therefore,

$$r_{V+R} = r_V + r_R, \tag{2.15}$$

where r_{V+R} is the ADU/sec registered with our detector, and r_V and r_R are the count rates for comparable detectors in visual and red passbands, respectively. We wish to use 23rd magnitude as the reference point of our scale, so

$$m_{V+R} = 23 - 2.5\log_{10}\left(\frac{r_{V+R}}{{}^{23}r_{V+R}}\right), \tag{2.16}$$

$$\text{and } {}^{23}r_{V+R} = {}^{23}r_V + {}^{23}r_R, \tag{2.17}$$

where m_{V+R} is the magnitude in the V+R band, and ${}^{23}r_V$, ${}^{23}r_R$, and ${}^{23}r_{V+R}$ are the count rates of a colorless 23rd magnitude object in the V, R, and V+R bands, respectively. The quantities ${}^{23}r_V$ and ${}^{23}r_R$ can be calculated by fitting

$$r_{V+R} = {}^{23}r_V 10^{\frac{(23-V)}{2.5}} + {}^{23}r_R 10^{\frac{(23-R)}{2.5}} \tag{2.18}$$

for stars of varying magnitudes and colors.

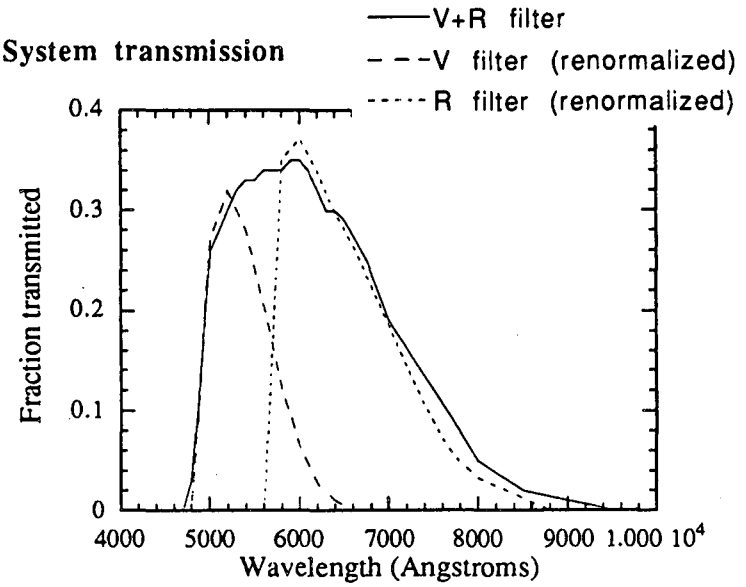

Fig. 2.13. System transmission. The response of our V+R filter is approximately the sum of the V and R passbands of Graham (1982). For comparison purposes, the V and R passbands are shown as 50% and 60% of Graham's transmission, respectively.

In order to find the values ${}^{23}r_V$ and ${}^{23}r_R$, we used a series of nine standard stars imaged on Nov. 30, 1989, one night after a new moon. We measured r_{V+R} from the ADU/sec deposited in an aperture of radius 3 pixels centered on the star. For this night, the full width at half maximum was about 2.5 pixels. This rate was then corrected for atmospheric extinction corresponding to 84% transmission from one atmosphere. From the measured ADU/sec, we calculated two calibration quantities, ${}^{23}y_V$ and ${}^{23}y_R$, from

$$V = 23 - 2.5\log_{10}\left(\frac{r_{V+R}}{{}^{23}y_V}\right)$$
$$\text{and} \quad R = 23 - 2.5\log_{10}\left(\frac{r_{V+R}}{{}^{23}y_R}\right). \tag{2.19}$$

Plugging each of these in turn into equation 2.18, we find

$${}^{23}y_V = {}^{23}r_V + {}^{23}r_R 10^{\frac{(V-R)}{2.5}},$$
$$\text{and} \quad {}^{23}y_R = {}^{23}r_R + {}^{23}r_V 10^{-\frac{(V-R)}{2.5}}. \tag{2.20}$$

The best fit values are

$$^{23}r_V = 1.45 \text{ ADU/sec},$$

$$\text{and } ^{23}r_R = 1.33 \text{ ADU/sec.} \qquad (2.21)$$

Thus, the reference rate for the V+R filter is $^{23}r_{V+R} = 2.78$ ADU/sec. In five minutes, a colorless 23rd magnitude object at zenith will accumulate 834 ADU in an aperture of radius 3 pixels. The magnitudes of our sources are determined from their photometry measurements on November 30, 1989.

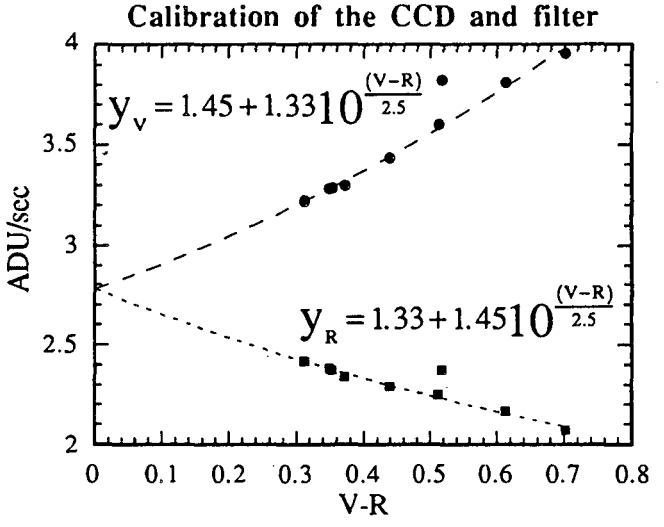

Fig. 2.14. Calibration of the CCD and filter. This graph shows the number of ADU/sec measured by our system for an R=23 object (upper curve) and a V=23 object (lower curve) as a function of its color. The rates have been corrected to sec(z)=1. The y-axis intercept gives us the reference point for calibration of the V+R filter. The one outlying point was not used in the calibration.

In five minutes, our example image accumulated 4090 ADU of sky background. From equation 2.10, the pixel noise is 28.6 ADU. For the dimmest objects, the statistical uncertainty in magnitude is dominated by the sky background. For a photometry aperture of radius 3, the one sigma noise in the center of the CCD is 152 ADU. At the very edges,

the one sigma uncertainty in photometry measurements is approximately 190 ADU. We set the limit for finding objects at 760 ADU, which is 5 sigma for the central region and 4 sigma for the edges. In addition, we masked 23% of the pixels containing the edges of the image, all objects over about 20th magnitude, and reflection streaks. The photometry was executed by first subtracting off a bilinear fit to the background counts (see chapter 3), then finding all of the pixels over 54 (this is half the height of a Gaussian of width 2.5 and integrated volume 760) that were not masked. If the integrated counts in an aperture of radius 3 around the closest centroid to these high pixels was greater than or equal to 760, then that position was stored as the position of a source. Figure 2.15 shows the magnitude distribution of sources in our images.

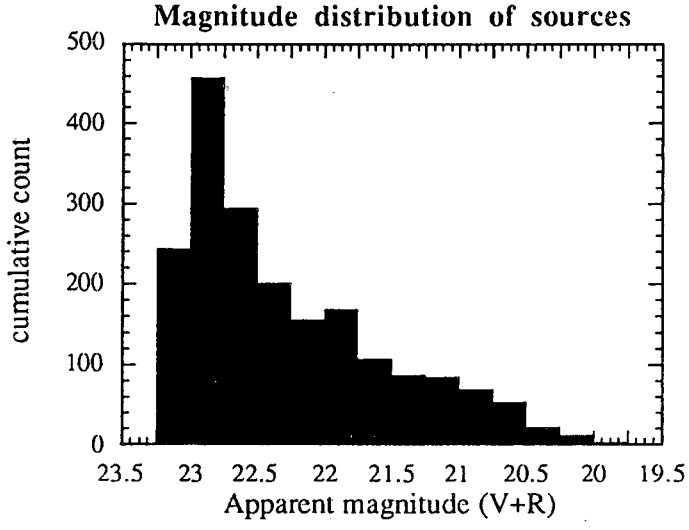

Fig. 2.15. Magnitude distribution of sources. This plot shows the magnitude distribution of 1990 objects found in one f/1 field.

More enlightening than the magnitude distributions is the comparison of the magnitudes of these objects with an image taken one month later (Fig. 2.16). The full width at half maximum was about 30% larger on Dec. 28, 1989, so the aperture was

increased to 3.9 pixels. The one sigma error in the magnitudes of the December sources should be 179 ADU. Notice that there is a group of objects that shows a significant increase, but no corresponding set of objects that decrease. Inspection of a large number of these shows that the majority are not really sources, but rather small noise peaks near the edge of a galaxy. The noise provides the peak, and the integrated counts of the galaxy fringes pushes the magnitude up over 23.1. Since the seeing is worse on December 28, 1989, more light is thrown into the aperture, and the brightness appears to increase. A few of those that increase are close doubles, so the worse seeing throws more of the light from the near object into the photometry aperture. Figure 2.16 shows that if we set the detection threshold at 23rd magnitude, we would find over 100 variable objects in a single image. This is an unacceptable false alarm rate. A more reasonable detection threshold would be 22nd magnitude.

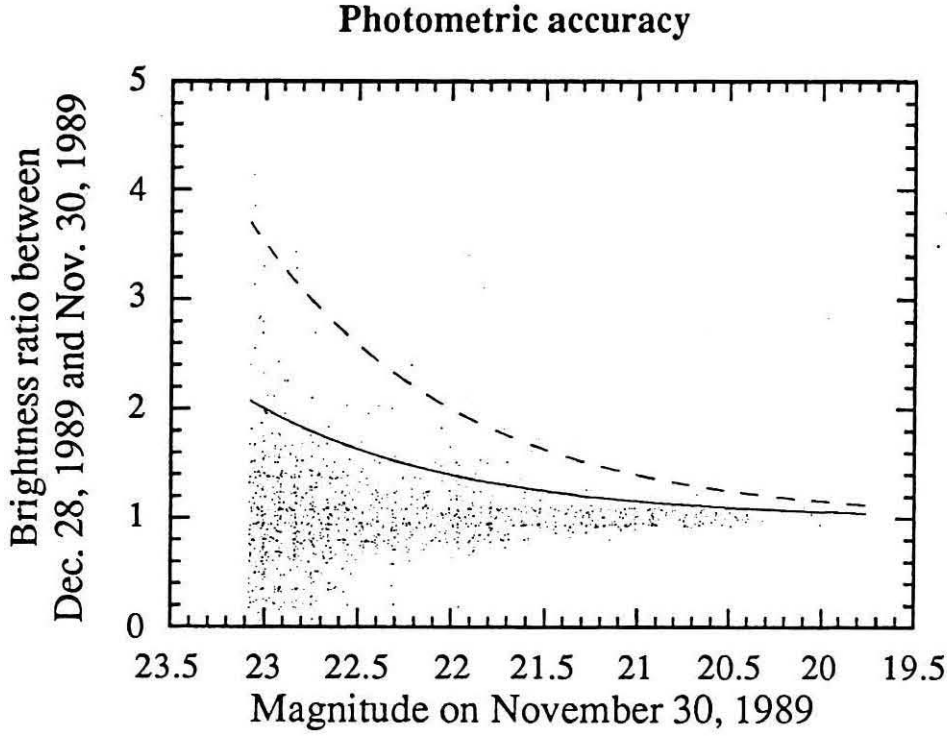

Fig. 2.16. Photometric accuracy. After normalizing the brightnesses on Dec. 28, 1989 and Nov. 30,1989, the brightness ratios of the individual objects is plotted versus the magnitude as determined on Nov. 30,1989. The solid line shows the expected brightness ratio for an increase in brightness equal to 23rd magnitude in (V+R). The dotted line shows the same for a 22nd magnitude increase in brightness. Many of the stragglers in this distribution are not statistical deviations, but result from identifiable errors in object identification.

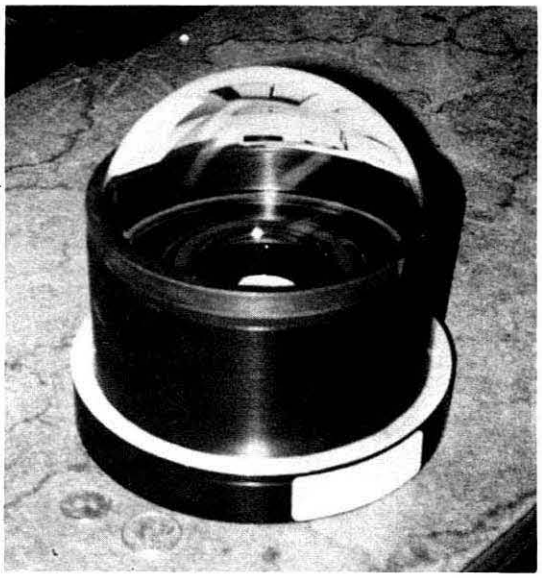

a)

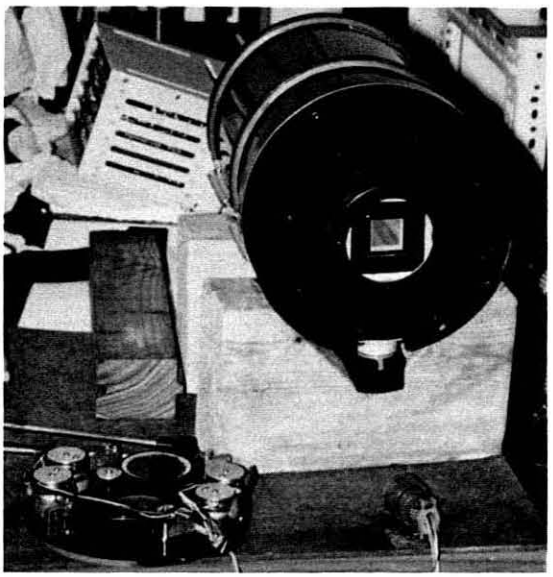

b)

Fig. 2.17. The hardware at first light. a) The correcting lenses are shown with a quarter and an Australian twenty cent piece for scale. b) This shows the CCD card cage (background), the CCD dewar and CCD (right), and the filter/shutter apparatus (lower left). The central grey area on the front of the CCD dewar is the 1.9cm x 1.9cm Thomson CCD. The filter/shutter apparatus screws directly to the top of the CCD dewar. c) For the first time, the f/1 camera is lowered into the prime focus top end of the Anglo-Australian Telescope.

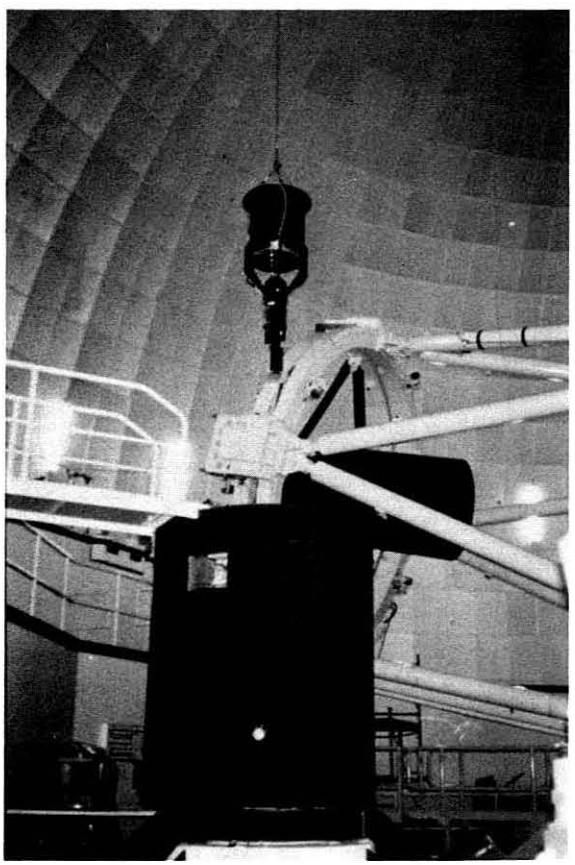

c)

3. Algorithms for finding supernovae

3.1. Overview

The success of a distant supernova search hinges on high quality telescope time, coverage of a large number of galaxies, and fast, reliable software. In the supernova business, we do not have the luxury of analyzing the data at our convenience. In twenty-four days, a z = 0.3 supernova drops in brightness from B = 21.5 mag to B = 23 mag. At maximum, it is difficult to measure the spectrum. At 23rd magnitude, it is nearly impossible. In addition to follow-up spectroscopy for supernova identification, we need to make immediate photometry measurements in order to find the time of maximum and the peak magnitude, which will be used to calculate the deceleration parameter. In order to make measurements near maximum light, the discovery images must be taken before new moon, and the candidates must be found and observed before the approaching full moon makes accurate spectral identification and photometry difficult. Realistically, there must be sufficient telescope time to assure favorable conditions for supernova discovery *and* follow-up within this two week period. The software must reliably identify the candidates within a few days at the most. Preferably, the software would produce the candidates from one night's observations (half a gigabyte of data) in 24 hours.

3.2. Photometry

The first analysis program we were able to produce that met the criteria for speed used list comparison of aperture photometry. Before a night of observing, we created the photometry lists for the reference fields, and transferred them to the computers at the site of the Anglo-Australian Telescope. Half way through the night, enough images were taken that the observer in Australia was able to create a flatfield frame. A scientist in Berkeley,

California, got up early in the morning (which corresponds to the middle of the night in Australia) and began flatfielding the images by remote login to the VAXen in Coonabarabran, Australia. This would be completed by the middle of the day Berkeley time, corresponding to early morning in Coonabarabran. If all went well, the morning backup tapes given to the observers included the raw data in Figaro format (1 gigabyte) and flatfielded images in VISTA compressed format (0.5 gigabytes). These tapes were shipped to Berkeley as quickly as possible, so that if in the next two days we found supernova candidates, we had data to look at for confirmation. (The computers at the AAT are only loosely networked, so it is impractical to send large images over the data links.) While the tapes were being shipped, we were still logged in remotely running the photometry list comparisons and shipping back small files containing information on possible candidates in each image. By the time the data arrived, we had a list of our best candidates to check in the images. Finding charts were generated using the Space Telescope Guide Star Catalog for any candidates that were deemed true variable objects.

Figure 2.16 shows the photometric dispersion for a simple list comparison. The automated program was much more sophisticated, but gave somewhat similar dispersions. The number of candidates was reduced, since 19% of the pixels were rejected around the edges of the images, and objects dimmer than magnitude 22.5 were ignored. Larger areas around bright stars were eliminated as well as any object that was considered "too close" (~10 arc seconds) to another object to get reliable photometry under differing seeing conditions. Supernova candidates were identified as those objects that increased by 20% or more and increased by the equivalent of $m_{V+R} = 22.8$ mag or brighter. In addition, the candidate was eliminated if it differed by more than 10% between the new images or was in the tails of a bright object or an obvious image defect.

Aperture photometry is not the optimal way to find supernovae in distant galaxies. If we do aperture photometry on the distant galaxies, we must include the edges of resolved galaxies and the supernovae which could be anywhere from 0 to 2 arc seconds from the center of the galaxy. This requirement means we must integrate over a large number of pixels regardless of seeing or optimal aperture size from signal-to-noise ratio calculations. In addition, supernovae cannot be found on galaxies below 23rd magnitude even if they are close enough for a supernova Ia to be visible; objects which have neighbors within ten pixels (about 10% of the objects) have to be rejected; and there is no visual feedback from the candidate lists about systematic problems with the images. Some systematics can be found easily by looking for reflections, satellite streaks, or other artifacts at candidate positions. Others are more insidious, like a seeing gradient from one side of the image to the other (Fig. 3.1). It is true that most of these difficulties could be addressed by using a more complicated and thorough photometry algorithm; but, these modifications would slow down the process, and make it more similar to a subtraction. The subtraction algorithm has the advantage of providing immediate visual diagnostic information.

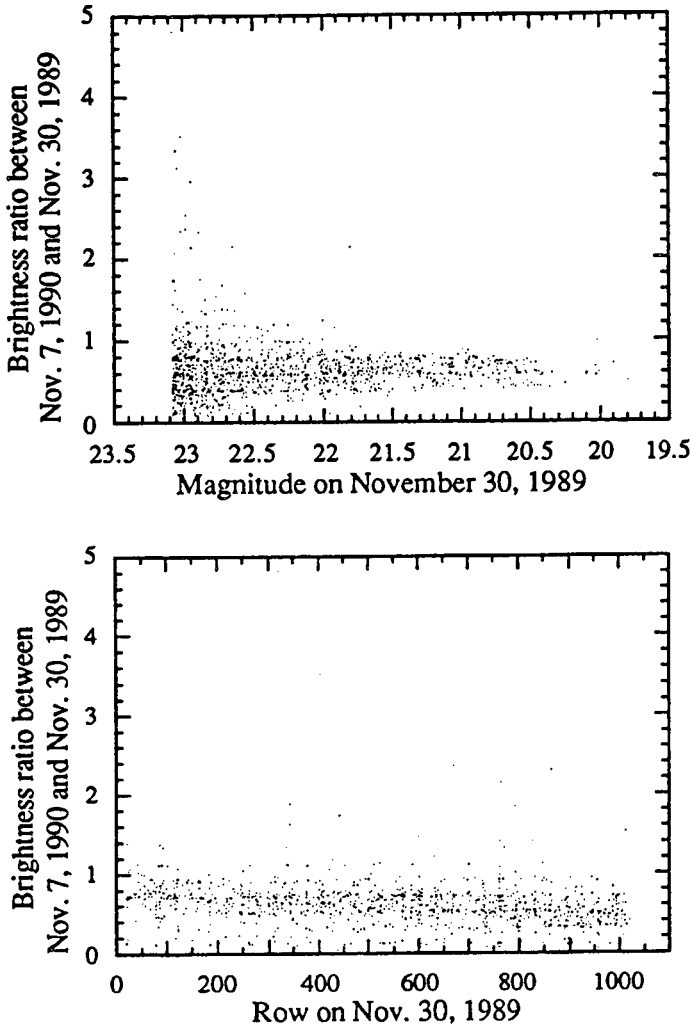

Fig. 3.1. Systematics of photometry. These data show a systematic deficiency of aperture photometry. It is known from subtraction that the stellar widths and brightness ratios can vary across the images. The variation of brightness ratio with row of the CCD for a one year baseline on our example image is shown on the lower graph. Notice that the brightness ratio decreases as row increases. The upper graph shows an unacceptable broadening of our photometry statistics. Notice that the dispersion remains broad even for bright objects, indicating a systematic error in photometry. A similar effect occurs for a comparison of Dec. 28, 1989, and Nov. 30, 1989 (Fig. 2.16), if the aperture radius is decreased from 3" to 2.5". This effect is explained by a small seeing gradient across the image. When a relatively large aperture radius is used, almost all of the light from each object is included in the aperture. If the radius is decreased, then areas of the image with comparatively larger stellar widths in the new image will produce lower brightness ratios. In this way, a seeing gradient produces an apparent gradient in brightness ratio.

3.3. Image Subtraction

For our image analysis, we use the VISTA image processing package, developed at Lick Observatory. This package consists of subroutines which can be invoked from the command line, or by executing procedure files which contain a list of commands. Our

version of the software has been heavily extended over the course of many years of use in the Berkeley Automated Supernova Search (Perlmutter et al. 1988). Extensions include compressing images, manipulating string variables, reading Figaro format images, displaying in a windows system, running external executable programs, and many, many others. In this section, we present the essential features of the VISTA procedure script used in image subtraction for the distant supernova search. In subsequent sections we describe in detail a few of the VISTA extensions that were developed for this distant search and are integral to its operation.

The image subtraction software is written in a menu-driven, modular form that can be run interactively, or as a batch job. The meat of the image analysis is contained in six procedures that (1) read the images into memory, (2) align the images, (3) create a subtraction, (4) identify those areas of the subtraction in which to search for supernovae, (5) find candidates, and (6) eliminate objects that are unlikely to be supernovae.

Read the images into memory

There are two ways to specify the images to be subtracted. Usually, one simply generates a file containing the names of the images to be searched, and the corresponding references. The name of this file is specified when starting up the subtraction procedure. Alternatively, one could specify the run numbers with which to start and end. The procedure will read in the first image, then read in subsequent runs until it finds a second image with the same object name in the header. To facilitate identification of the reference image, the reference images must be pre-processed and named after the object name when using this option. This latter method of specifying the images is useful when processing images as they are being generated at the telescope.

Once the images have been specified, this procedure reads into memory two (2.5 minute exposure) flatfielded images of the same field and a corresponding (5 minute exposure) flatfielded reference image. We have found it useful to remove large scale gradients from each of the images so that the sky background level is uniform. Some of the large scale gradients come from the wings of bright stars, which can cause problems when the reference image is from a night of different seeing conditions. Others are just artifacts of the moon, reflections, or other sky conditions. The subroutine that flattens the images is called *bilinsurf*. As its name implies, it finds the local sky background in a grid of points across the image and fits a bilinear surface. This surface is subtracted from the image pixel by pixel. Fig. 3.2 shows the result of flattening the image in Fig. 2.11.

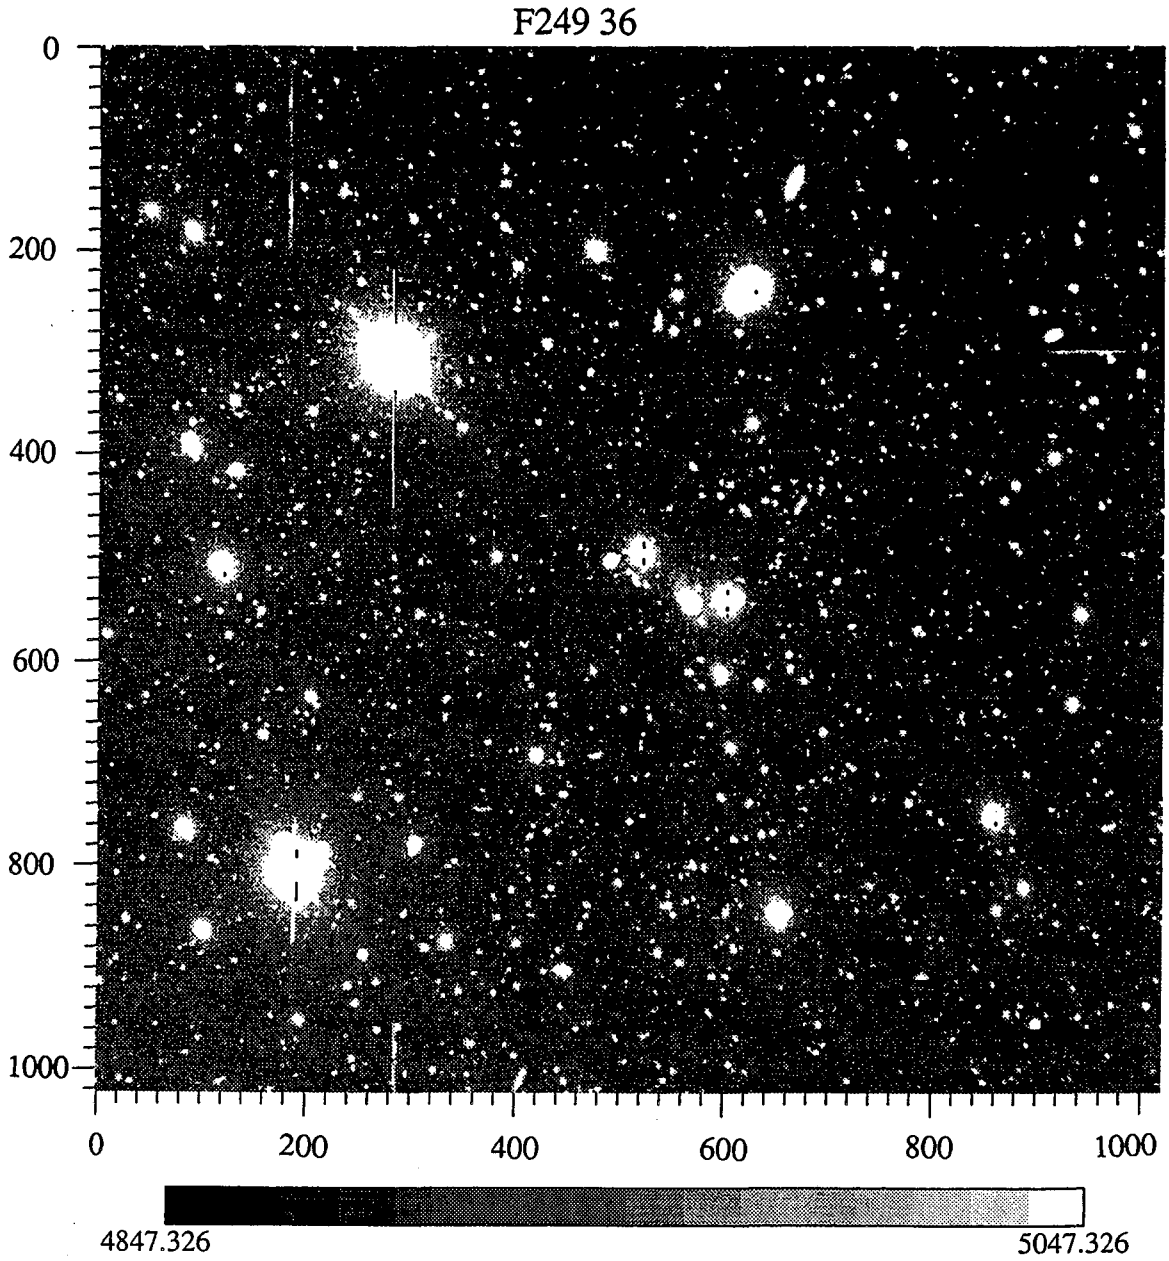

Fig.3.2. Flattened image. In addition to being flatfielded (Fig. 2.11), this image was flattened using the bilinsurf algorithm. In these reproductions of the images, the difference between flattened and unflattened images is subtle. But, notice in Fig. 2.11 that the background in regions adjacent to large saturated stars (like the one in the lower left hand corner, for example) is brighter than the background in regions with fewer bright objects (like the right side of the image). This effect is somewhat removed in Fig. 3.2. The flattening removes large scale gradients in the background that can be caused by tails of real objects in the image or by light leaks.

Align the images

Three images are involved in this process: the two new images and the reference image. First, the two new images are aligned. To accomplish this, we find the positions and magnitudes of bright stars in a region 90% as wide and 90% as tall as the image. In this case, bright is defined as those stars with peak pixels below saturation and above 75 sigma of the pixel noise (calculated from [sky/gain+square of readout noise]$^{1/2}$). Due to constraints of the image matching subroutine, the star lists are limited to no more than 200 entries. The lists from both new images are fed into the *deepmatch* subroutine, which calculates the row and column offsets and the brightness ratio between these two images. The two images are shifted equal but opposite amounts towards each other so that they maintain similar stellar profiles. The sum of these new images is stored in a separate memory allocation, so that the original images are preserved for later analysis. In rare cases, the brightness ratio between the two images is significantly different from one. In this case, the images are adjusted so as to be equal. This will make later candidate elimination easier.

Then, a similar process is applied to the alignment of the summed new image and the reference image. Photometry lists of the brightest stars (with peaks between saturation and 50 times the one sigma pixel noise) are obtained for the summed new image and the reference image. This time, Gaussian fits to the peak pixels in each star are also measured. The lists are sent to *deepmatch*, which fits the row and column offsets, rotation, magnification, brightness ratio, and makes a first guess at the seeing in each image. It has been found that this Gaussian fit to the top three pixels in each direction is unreliable, so better numbers are calculated in the next step, using *fitstars*.

Create a subtraction

The key to this image analysis system is creating the cleanest possible subtraction. The image with the better seeing is shifted and rotated in one operation, which minimizes the distortion of the stellar profiles. The subroutine *fitstars* calculates the best fit two dimensional Gaussians for the (possibly shifted) new image and the (possibly shifted) reference. Since the stellar profiles often are different across these wide fields, *fitstars* actually calculates the best fit Gaussian in each of four quadrants of each image. The results are sent to the subroutine *bilinsmooth*, which is responsible for convolving each pixel of the better image with the appropriate 2-D Gaussian. In practice, the field is broken up into about 100 regions which are all convolved with the same Gaussian to save time. Since it may be necessary to convolve one image in one region and the other image in another region, the actual subroutine is fairly complex. *Bilinsmooth* also calculates the brightness ratio in each quadrant in the convolved images and multiplies each pixel in the new image by the interpolated brightness ratio. When the images are as similar as we can make them, the reference is subtracted from the new image pixel by pixel.

Identify those areas of the subtraction in which to search for supernovae

Because we cannot cleanly subtract sources brighter than twentieth magnitude (and since these objects are much more likely to be stars than galaxies), pixels within a four pixel radius of the centers of these objects are masked out. It is necessary to mask larger regions for brighter stars, since they can cause larger subtraction artifacts. A radius of 30 is used to mask out the brightest saturated objects. If the subtraction routine is being run interactively, the user has the option of masking the pixels on a line. This is often used in cases where a satellite or reflection streak appears in the new image. When ignored, these streaks can swamp the analysis procedure with candidates.

Find candidates

Using our subroutine *locatestars*, we scan the remaining image for high pixels that could potentially be the signal of a supernova. For each of these pixels, we convolve with a Gaussian that is slightly wider than the stellar profile for that image. If the result of the convolution is larger than our threshold for finding supernovae, then that position is marked as a candidate.

This candidate finding routine must work for a variety of input images exposed on nights of widely different conditions. If the threshold for finding supernovae is set too low, many spurious candidates will make it difficult to find a real supernova before it fades. If the threshold is too high, we may not find a supernova that is present in the data. Therefore, the threshold is calculated separately for each image based on the pixel noise in each of the original images and the seeing (see section 3.8).

Eliminate objects that are unlikely to be supernovae

Although we try to be as careful as possible in identifying candidates, there are often cosmic rays, asteroids, and missubtractions among the list of possible supernovae. To eliminate these, we examine the objects in the candidate positions in each of the new and reference images. If the object moves between the two new images, it is assumed to have been an asteroid. If there is more than a three sigma deviation between the brightnesses in the two new images, or the parent in at least one of the new images is less than 30% of the supernova detection threshold, then the candidate is eliminated as a possible cosmic ray. If the increase in brightness is less than 10%, it is likely that the candidate is a missubtraction. (We are unable to reliably identify authentic variable objects at the ten percent level or

below.) After these impostors have been eliminated, the candidate list and the annotated subtraction are printed out for human inspection (Fig. 3.3).

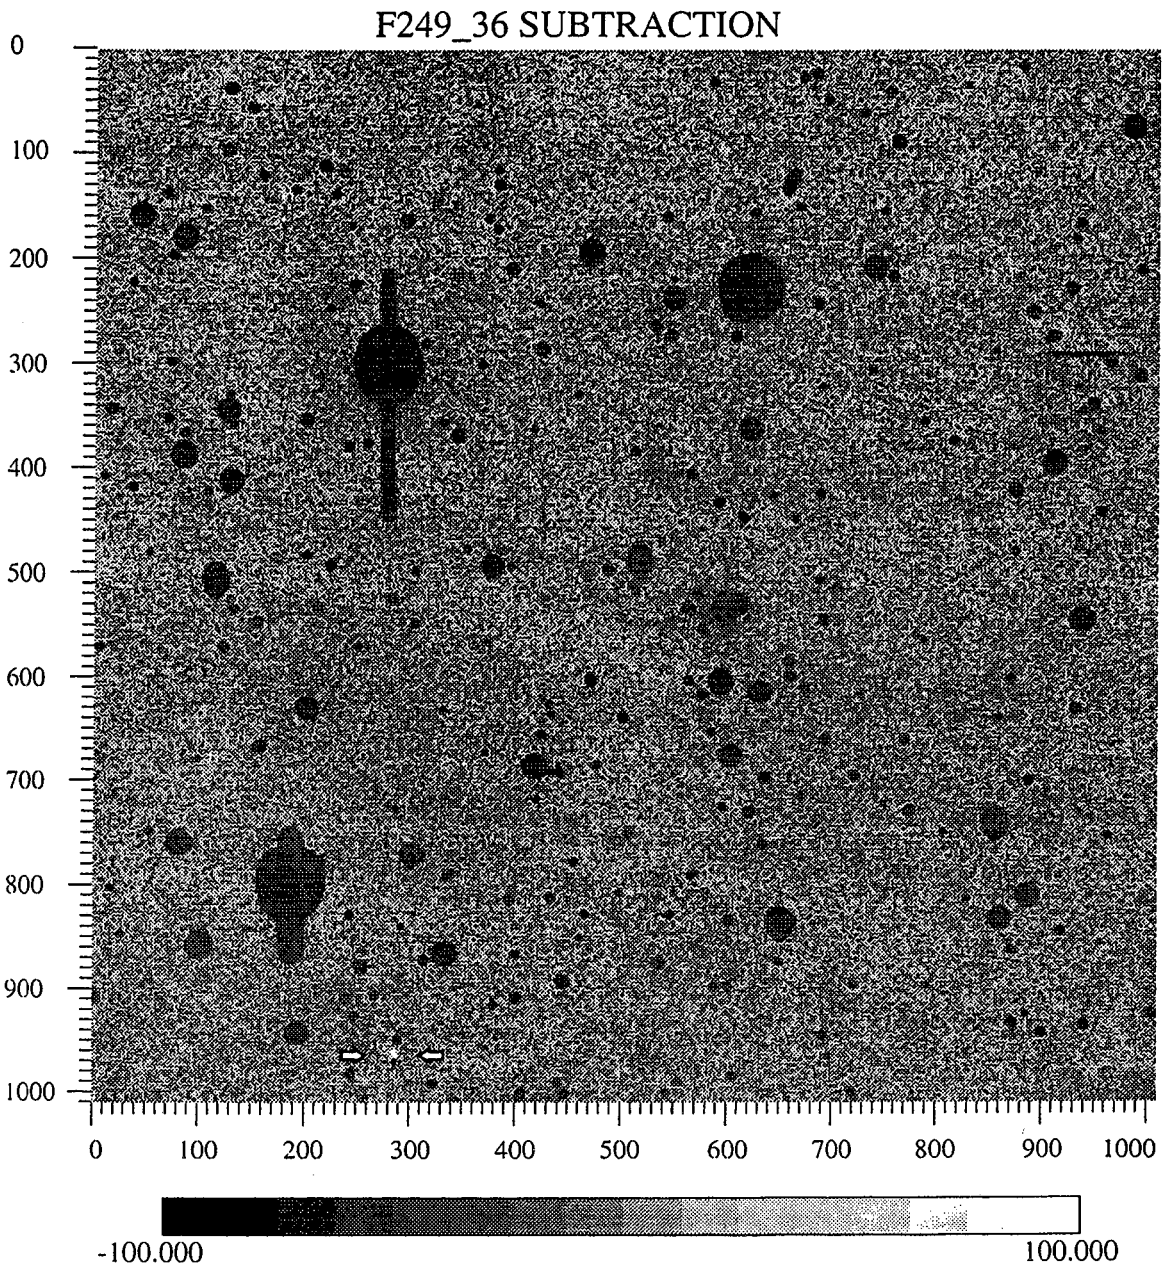

Fig. 3.3. Subtraction. This shows the subtraction of a Nov. 30, 1989 image from the Dec. 28, 1989 image in Fig. 3.2. The dark patches are regions around bright objects that have been automatically masked by the program. The one candidate found in this image, indicated by arrows, is caused by the imperfect subtraction of a reflection streak.

3.4. The *bilinsurf* algorithm

When analyzing an image using photometry, the background is often calculated by finding the average or median value of the sky in an annulus around each object. For crowded fields, it is much faster and potentially more accurate to calculate the background once over the whole image, and subtract it off so that photometry becomes simple integration over an aperture. Accurate background subtraction is also necessary for good subtractions. Local sky background differences with position in the sky, atmospheric effects, and non-Gaussian tails of bright sources can cause subtraction artifacts. Flatfielding should eliminate pixel to pixel fluctuations due to varying quantum efficiency, but it cannot resolve the difficulty of large scale background gradients.

Our flattening routine calculates the sky background in an array of boxes of size *dvsn* which are evenly spaced with a separation of *spacing*. The sky background in each box is assigned to the pixel at its center, producing a grid of background values across the image. For each pixel in the image, the sky background is calculated from bilinear interpolation using the background values of the four closest grid points. The only difficulty in this is handling the edge conditions. To help calculate the background values at the edges, additional background values are calculated in boxes of size *dvsn* located *dvsn*/2 from the edges, and aligned with each row and column. The background is also calculated as close to the corners as possible. The final array of values used for bilinear interpolation is the array of grid points plus the values at the edges (and corners) of the image which are calculated from extrapolation of the edge points and their nearest neighbors on the grid. To achieve a good background subtraction, *dvsn* must be large compared to the scale of the objects we wish to measure, and *spacing* must be small compared to the scale of background variation. Usually, *dvsn* and *spacing* are chosen to be equal. In general if the background variation is small enough that this method works,

sampling with *scale* smaller than *dvsn* will not significantly improve the background fit, and will only take more time. It might sometimes be advantageous to make *spacing* larger than *dvsn* if the background varies very slowly and increased speed is essential. As a word of warning, the overscan rows and/or columns should be removed from the image before calling *bilinsurf*, as these pixels will cause severe edge effects.

Assuming we have a fairly well behaved background to fit, the most important step in this algorithm is the determination of the sky background at each point. After several attempts, it was decided that the peak of a histogram of pixels in the box was the fastest and most accurate measurement of sky background. The trick is to determine what *binwidth* to use for the histogram. If the binwidth is too small, the histogram will be very noisy and it will be difficult to fit the peak. If the *binwidth* is too large, there will be less accuracy in the position of the peak. Many of our images have been compressed for easier disk storage. In this process, we take all of the pixel values in the image and store them as two byte integers. Therefore, if the range of pixel values is larger than 2^{16}=65536, then some of the pixel values are changed to nearby values so that there are only 65536 different values to store. This digitization of pixel values is an insidious problem for detection of a peak, since small changes in the position or width of a bin can cause large variations in the relative number of pixels in adjacent bins.

To combat the problems with digitization, the peak-finding algorithm we use calculates the peak for several positions of the bin, and takes the median peak value. Suppose the *binwidth* was chosen to be ten and the *step* was chosen to be one. A histogram would be created with a bin width equal to one. Since this is probably a noisy histogram as we have discussed, a new histogram is created in the region of the peak by adding together ten adjacent values of the previous histogram. For example, the number of pixels with values 3000, 3001, 3002, ...3009 might be summed, along with 3010-3019,

3020-3029, etc. The peak of this new histogram is found by fitting a parabola to the highest point and the two points on either side. Since the peak could be significantly different for slightly different values of the starting point, we generate nine more histograms. The first one will have bins 3001-3010, 3011-3020, etc.; the second one will have bins 3002-3011, 3003-3012, etc.; and the ninth one will have bins 3009-3018, 3019-3028, etc. The median value of these ten histograms is a fairly robust determination of the background value. In general the value of *step* should be one, although if there are many pixels in the box and there is little digitization, *step* could be increased to save a small amount of time.

3.5. The *deepmatch* algorithm

The *deepmatch* subroutine calculates the mapping between two photometry lists containing the row, column, brightness, and estimated widths of sources from two images. This simplistic routine uses only the row and column positions to calculate the row, column, and angular offsets between the images. As an option, magnification can also be fit. In addition to the offsets, the routine calculates the median widths of the stars and the brightness ratio between the images.

The first attempt to match the star lists starts as an exhaustive search. We attempt to match the first star on list A to the first star on list B, then the first star on list A to the second star on list B, etc. The number of comparisons is reduced by demanding that the matched stars will be within 100 pixels of each other in each direction. For each acceptable pair of stars, the row and column offsets are computed, and the rest of the stars on each list are checked to see if any are within three pixels of matching with this offset. The minimum number of matched stars accepted (*minmatch*) is 25% of the number of stars on the shorter list, but has a lower limit of two and an upper limit of seven. These matching criteria are

only guaranteed to be effective for row and column offsets less than 100 pixels, and small differences in rotation and magnification. The algorithm begins to falter if

$$|\theta| > \frac{3}{N} \quad \text{or} \quad |m-1| > \frac{3}{N}, \tag{3.1}$$

where N is the width of the image in pixels, θ is the rotation offset in radians, and m is the magnification between the two images.

The best values of the row offset, column offset, rotation offset, and magnification are found from the downhill simplex method of Nelder and Mead, as described in Press et al. (1986), pp. 289-293. The starting point for this method is chosen as zero rotation, magnification of one, and row and column offsets given by

$$\text{row offset} = \frac{\sum_{i=1}^{n}\left(A_{row}(i) - B_{row}(i)\right)}{n},$$

$$\text{column offset} = \frac{\sum_{i=1}^{n}\left(A_{col}(i) - B_{col}(i)\right)}{n}. \tag{3.2}$$

The other four vertices of the initial simplex are calculated from

$$\overline{P}_i = \overline{P}_o + \left(\overline{\lambda} \cdot \hat{e}_i\right)\hat{e}_i, \tag{3.3}$$

where $\overline{P}_o$ is the starting point and $\overline{\lambda}$ is given by

$$\overline{\lambda} = 0.05\hat{e}_{row} + 0.05\hat{e}_{col} + 0.0007\hat{e}_{\theta} + 0.0007\hat{e}_{m}. \tag{3.4}$$

The characteristic length scales of rotation and magnification are smaller because their error accumulates across the image. From this starting simplex, the *amoeba* subroutine of Press et al. (1986; downhill simplex method) is used to minimize the sum of the squares of the residual differences between the matched stars. The search is terminated when the decrease in functional value in the last step is fractionally smaller than 0.0001.

Using the best fit parameters, the original star lists are again searched for objects that match within three pixels. Then, the squares of the residuals are minimized for the new list of matched stars. This process is repeated until the number of stars matched is

constant. At this point, the list of matched stars is checked one more time, but this time throwing out stars that appear to match, but have residuals that are larger than two sigma, where sigma is calculated from the residuals of the previous match. The *amoeba* subroutine is run one more time to determine the final row, column, rotation offsets, and magnification.

3.6. The *fitstars* algorithm

The *fitstars* subroutine takes a list of matched stars from the *deepmatch* subroutine and calculates the best guess at the full width in each background-subtracted, aligned image. Since we have found that there is often a variation in stellar widths across the image due to instrumentation, weather, or alignment of images, the row and column full widths are calculated separately in four quadrants of the image. These results are passed to *bilinsmooth*, which uses bilinear interpolation of these values to adjust the stellar widths in each image so that they are equal. For our images, there are not enough bright stars to accurately calculate the widths in boxes smaller than a quarter of the image.

First, we determine which of the matched stars are in each quadrant. We require a minimum of five and a maximum of thirty stars in each quadrant to get a good fit in a short period of time. For the current quadrant, we load the values in fifteen pixel by fifteen pixel boxes around each star in each image into an array.

Next, we calculate the width of the stars assuming their profiles are Gaussians. A Gaussian star has a brightness profile given by

$$z = A\, e^{-\left(\frac{x^2}{2\sigma_x^2}+\frac{y^2}{2\sigma_y^2}\right)}. \qquad (3.5)$$

If we took the natural log of every pixel value, then the cross sections along rows and columns would be the parabolas

$$\ln(z) = \ln[f(y)] - \frac{x^2}{2\sigma_x^2} \quad \text{and} \quad \ln(z) = \ln[g(x)] - \frac{y^2}{2\sigma_y^2}. \tag{3.6}$$

The second derivatives of the cross sections are just $-1/\sigma^2$ for each direction. The second derivative should be constant until the noise dominates, or until the stellar profile becomes significantly non-Gaussian, whichever comes first. Where there are no objects, it should fluctuate around zero. The second derivative of the natural log of the pixel values is found for the pixels in the previously mentioned array from

$$\left.\frac{d^2}{dy^2}\left[\ln[z(x,y)]\right]\right|_x = \ln\left[\frac{z(x,y+1)\ z(x,y-1)}{z^2(x,y)}\right],$$

$$\left.\frac{d^2}{dx^2}\left[\ln[z(x,y)]\right]\right|_y = \ln\left[\frac{z(x+1,y)\ z(x-1,y)}{z^2(x,y)}\right]. \tag{3.7}$$

The σ_X^2 for each pixel are found from the negative reciprocals of the second derivatives in the column direction. The values within 1.5 times the full width at half maximum of the centers of the stars are histogrammed, excluding those values that are unreasonably small fractions of a pixel. Since the full width at half maximum is not known a priori, the histograms are made once using the full width from *deepmatch* and then again using the resulting full width. The peak of the histogram is found in a manner similar to that used in *bilinsurf*. The values are histogrammed ten times, with bin widths of 0.1 and origins different by a step size of 0.01. The median of the peaks of these ten histograms is returned as the column sigma for that quarter image. A similar calculation is done to determine the row sigma.

After the row and column widths have been found for each quadrant of both images, the average row and column widths are passed back as VISTA variables. For each quadrant, the full widths of the Gaussian with which the second image must be convolved to produce the widths in the first image are loaded into a common block so they can be

easily passed to *bilinsmooth.* If the first image has a narrower stellar profile, the value passed is the negative of the width of the Gaussian with which the first image must be convolved to produce the widths in the second image.

3.7. The *bilinsmooth* algorithm

The *bilinsmooth* algorithm does a two dimensional convolution as well as calculating and correcting for the ratio between the images. The program requires two aligned images (a reference image and a new image) and a list of matched stars from the *deepmatch* routine. After running this subroutine, the two images will be ready to subtract (or add) pixel by pixel.

Ideally, we would like to convolve each pixel in the better image with a Gaussian of exactly the right row and column widths. These might be calculated from a bilinear interpolation (extrapolation) of the widths found in *fitstars* for the centers of the four quadrants. The image to be convolved should be chosen on a pixel by pixel basis, since one image might have narrower stellar profiles in some areas and broader stellar profiles in others. However, it turns out to be time-consuming to find the correct width and re-calculate the convolving function on a pixel by pixel basis. Therefore, the image is divided up into a grid of 100 boxes of size nrow/10 by ncol/10. The same convolution is performed for every pixel in a given box.

This algorithm would be unsatisfactory if there were large differences in stellar widths across the image. The method of convolution depends upon there being little difference in the convolving Gaussian from one box to the next. In addition, even our calculation of the convolving function is adversely affected by a large gradient. To determine the full width in the image, we measured stars that were distributed *randomly* in

each quadrant. Then, the most commonly measured full width was assigned to the *center* of the quadrant. If there is much of a difference from one side of the quadrant to the other, the results will depend heavily on the distribution of the bright stars. So, we must be able to assume that the randomly distributed stars are characteristic of the center of the quadrant, that the characteristic length over which widths change is large enough to warrant a bilinear fit to four points, and that the change in width is negligibly small between points separated by one tenth the width of the image.

To implement this algorithm, we extrapolate the widths of the convolving functions from the centers of the quadrants to the corners of the image. Next, we convolve the reference image so as to increase the stellar widths in the row direction. The ten pixelized convolving functions to be used on the first column are calculated. For each pixel in the first column, the convolution in the row direction is performed, skipping pixels that will be convolved on the other image instead. This is repeated for the second column, etc., until we are one tenth of the way across the image. Then, the ten new convolving functions are calculated and we march across columns until we have covered two tenths of the image. The process is repeated until the entire reference image has been convolved in row. Similarly, the reference image is then convolved in column. Finally, all of the steps detailed in this paragraph are repeated to convolve the new image.

At this point, the two images should have the same stellar widths at every pixel, so the brightness of each of the matched stars can be easily calculated using a fixed radius. Using these brightnesses, a brightness ratio is calculated for each quadrant. For every pixel in the new image, the brightness ratio is calculated from a bilinear fit to the four quadrant centers, and the pixel is corrected to match the reference image. This allows us to measure the magnitudes of the candidates from the calibration of the reference images.

3.8. The *locatestars* algorithm

Once the optimal subtraction has been obtained, it is searched for objects that are potentially supernovae. Obviously, one can do better than searching the subtraction for the highest pixel values, since there is additional information in the values of the adjacent pixels. One can also do better than aperture photometry, since the signal-to-noise ratio is higher in the center of an object, and the aperture gives every pixel equal weight. We would like to find a weighted sum of the pixel values that gives us the best chance of finding supernovae.

We chose to convolve with a Gaussian because it is easy to calculate, and with our undersampled pixel scale it is difficult to find a substantially better filter. If the signal is assumed to be a Gaussian with amplitude A and width determined from σ, and the filter is a Gaussian of amplitude K and width determined by Σ, then the total signal in a convolution centered on the peak of the signal is theoretically

$$S = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\left(A\ e^{-(x^2+y^2)/2\sigma^2}\right)\left(K\ e^{-(x^2+y^2)/2\Sigma^2}\right)dxdy$$

$$= \frac{AK2\pi\sigma^2\Sigma^2}{(\sigma^2+\Sigma^2)}. \qquad (3.8)$$

The noise in the measurement is given by

$$N = sK\sqrt{\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\left(e^{-(x^2+y^2)/2\Sigma^2}\right)^2 dxdy}$$

$$= sK\Sigma\sqrt{\pi}, \qquad (3.9)$$

where s is the sigma of the noise in one pixel. The signal-to-noise ratio is a maximum for $\Sigma = \sigma$.

The *locatestars* routine is designed to perform convolution with a Gaussian of width $\Sigma = k\sigma$. The amplitude K is chosen so that the result of the convolution approximates the result from aperture photometry. Putting all this together, the convolution function is

$$\Phi(x,y) = \left(\frac{1}{k^2} + 1\right) e^{-\left(\frac{x^2}{2k^2\sigma_x^2} + \frac{y^2}{2k^2\sigma_y^2}\right)}. \qquad (3.10)$$

It is moderately important for this convolution to be applied at the right position; otherwise, the signal falls off to a fraction

$$e^{-\frac{1}{2}\left(\left(\frac{\Delta x}{\sigma_x}\right)^2 + \left(\frac{\Delta y}{\sigma_y}\right)^2\right)}.$$

For stellar widths in the range of our data, the filter should be applied within 0.2 pixels of the center of the signal.

In order to reduce the running time of this algorithm, we convolve only at positions where there is a peak pixel over a specified level. In order to find a more accurate center of the signal, a centroid of the surrounding pixels is found and the convolution is performed there. In this way, we reduce the number of convolutions from $(1024*5)^2 = 26$ million to less than ten thousand.

In addition to several parameters specific to VISTA, the *locatestars* subroutine accepts the following input options:

(1) The range of pixel values to be searched must be specified. This range is usually chosen so that the minimum is several sigma above the sky background and the maximum is just below the saturation level of the CCD. This way, the algorithm finds the significant peaks most efficiently.

(2) Pixels within the specified range are searched for local peaks. The position of the peak in fractional pixels is found by iteratively finding the centroid of the pixels

within *centrad* of the calculated center, until the position converges. Small values of *centrad* keep the centroid from wandering from object to object in a crowded field, but larger values give a more accurate center for well isolated bright objects. The subroutine also gives the option of rejecting pixels within specified distances from a previously marked object or from a previously masked pixel.

(3) The background level is usually set separately for each object based on the values in a specified annulus centered on the object. This background can be determined by either the average or the median pixel value within the annulus (in both cases, outlying pixel values are discarded). Alternatively, the sky background can be set to zero. This is the optimal choice if the background has been satisfactorily removed (for example by the *bilinsurf* subroutine).

(4) *Locatestars* allows the options of aperture photometry (convolution with a cylindrical function of specified radius) or weighted photometry using convolution with a 2-D Gaussian function. In either case, the subroutine must be supplied with the minimum number of integrated ADU counts that must be found for the peak to be added to the photometry list. Because optics often cause vignetting in a roughly circularly symmetric manner, the subroutine may be given a different threshold exterior to a specified distance from the center of the image.

(5) In the case of Gaussian convolution, *locatestars* must be supplied with the row and column widths of the Gaussian signal for which it is looking. The convolving function will be a Gaussian of width "k" times the width of the expected signal. Additionally, one specifies the radius within which the convolution is calculated. Larger radii give more accurate photometry in sparse fields, but smaller radii are necessary for crowded fields and to reduce the processing time.

In our subtractions, k is set to 1.15. This was found experimentally to find the largest percentage of planted supernovae in our images. For candidate status, the threshold number of ADU counts in the convolution is set at 6.5 times the one sigma noise given by:

$$N = s\left(\frac{1}{k} + 1\right)k\sqrt{\pi\sigma_{row}\sigma_{col}}, \qquad (3.11)$$

which is a variant of equation 3.9. The minimum pixel value is set to 70% of the maximum height of a 2-D Gaussian at the detection limit, where the background value is assumed to be zero. Finally, the radius within which the convolution is calculated is twice the full width of the best fit stellar profiles. Smaller radii reduce the detection efficiency and larger radii slow down the process unnecessarily.

3.9. Software timing

As we have previously stated, it is very important to process the images quickly. The ability of our software system to process data in a timely manner depends heavily on the computer used. The power of VAX CPU's is measured in terms of the VAX 780 computer. On the average, the complete analysis of one field takes 73.7 minutes of VAX 780 equivalent CPU time. At LBL we are fortunate to have access to a VAX 6000-6510 (15 times faster than a VAX 780), which processes a field in 4.85 minutes, and a VAX 6000-610 (40 times faster than a VAX 780), which processes each field in 1.87 minutes of CPU time. Of course, the elapsed time to process each image is highly dependent on CPU availability, speed of disk access, and the size of machine memory space available. VISTA derives much of its speed from being a memory-based rather than disk-based image processing system. However, this means the subtraction process consumes over 25 megabytes of memory space. In general, there is not enough space in physical memory for all of this information, so the process must be paged out to a disk, which can increase the actual time of data processing. On a relatively unloaded VAX 6000-6510 (typically a batch

job run overnight), the average elapsed time to process a single field is 6.6 minutes, with a one sigma deviation of about 0.3 minutes. However, if there is a significant load, the processing may take up to fifteen minutes. Fig. 3.4 shows the elapsed times for various sections of the subtraction process on an unloaded VAX 6000-6510.

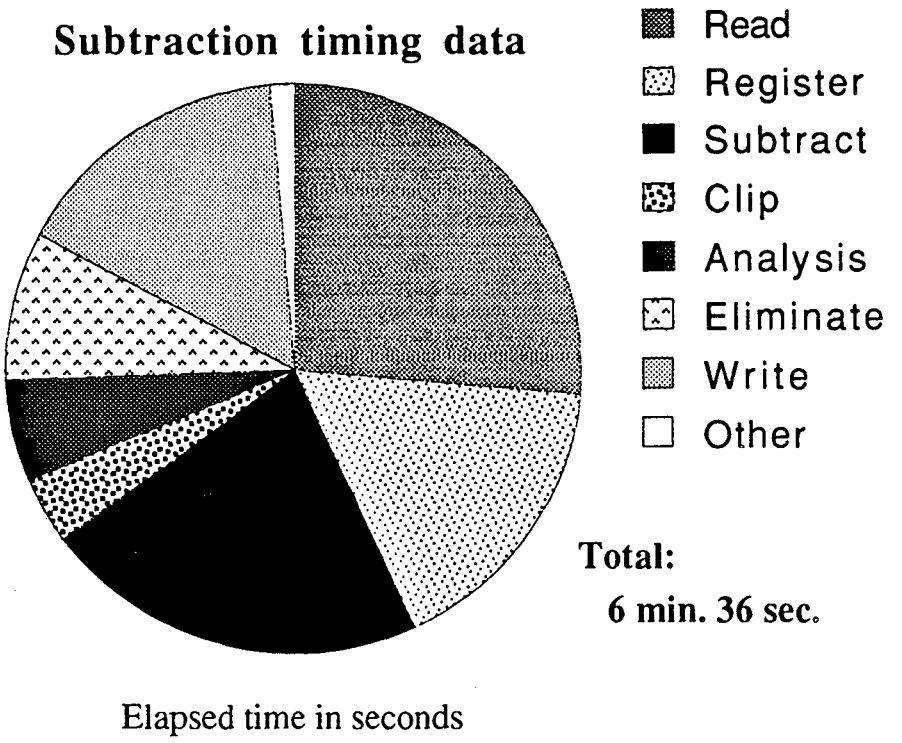

Fig. 3.4. Subtraction timing data. This graph shows the average elapsed times for subtractions on a relatively unloaded VAX 6000-6510 (CPU is 15 times as fast as a VAX 780). The individual sections took the following times: parameters: 1 sec., reading in and flattening images: 105 sec., image registration: 67 sec., creating a subtraction: 86 sec., masking bright stars: 16 sec., subtraction analysis: 24 sec., eliminating spurious candidates: 33 sec., writing the subtraction to disk: 63 sec., and other: 1 sec. The times for image analysis and candidate elimination are highly dependent on the quality of the subtraction. The total elapsed time was six minutes and thirty-six seconds per subtraction, with a corresponding CPU time of 4 minutes and 51 seconds.

The subtraction process has evolved to a state where it is relatively robust, and most images process without error in untended batch jobs. Typically, only about two fields from each run of 50 images will have to be processed interactively, due to the presence of a

bright satellite streak or an unusually bright star, which must be masked out by hand to avoid swamping the process with candidates. If the atmospheric conditions were unusually bad, or if the field center has drifted more than 100 arc seconds from the intended center, or if there are large rotational differences between the new images and the reference, the fields cannot be processed automatically.

3.10. Efficiency for finding supernovae

The challenge in adjusting the subtraction software parameters is to maximize the efficiency for finding supernovae while keeping the number of false alarms to an acceptable level. If the criteria for candidate identification are too loose, we will find events that are caused by noise fluctuations in the CCD, rather than luminosity changes in an astronomical object. With fifty million pixels per night, even if we set the threshold at five times the pixel noise, we would expect to find fifteen purely statistical events per night. Fortunately, we have more information about each object than pure single-pixel statistics, but this example shows a vulnerability of variability studies of a large number of objects. It is important when evaluating an algorithm to include both the number of events which will have to be followed up and the fraction of supernovae we are expected to find.

Subtraction candidates

The candidates from a single night of observations include variable objects and statistical fluctuations, in unknown proportion. The number of candidates produced in each run is highly dependent on the atmospheric conditions and the amount of time that has elapsed between the new images and the reference images. To demonstrate the variation, we will compare the results of subtractions of 25 fields taken on Dec. 28, 1989, and the

results from the same fields taken on Nov. 13, 1990. We will use the reference images taken on November 30, 1989 for both sets of subtractions.

For the Dec. 28, 1989, 24 of the 25 images ran smoothly in batch, and one had to be processed interactively to remove a bright streak produced by a satellite. Twelve of the 25 images produced no candidates, four produced only one candidate, three produced two candidates, and six produced more than two candidates. Of the 74 total candidates, 69 of them were eliminated by glancing at the printed subtractions; twenty-five were on easily identified reflection streaks, ten were clustered around a single bright star that was improperly masked, nineteen surrounded dimmer masked objects, and fifteen were due to improper quantum correction of three pixels. We noticed that these three pixels were poorly corrected because they caused candidates in approximately five images each. In future searching, the new images will be taken with a small offset between them, so that the software can identify these pixel defects from their motion with respect to the stars.

Subtraction of the same fields taken on November 13, 1990, produced 235 candidates with none of the images requiring processing by hand. Using similar criteria for visual elimination of candidates, the total number of candidates was reduced to a pool of 78. To a researcher intent on finding a distant supernova within a day or two at the most, this number of candidates represents a significant obstacle. Most of these candidates result from improper extrapolation of the ratio and stellar widths in *bilinsmooth.* This was discovered by inspection of the subtraction artifacts, but can be easily demonstrated by comparing the number of candidates in the center of the image with the number in the corners. To this end, we divide each image into two regions. One region is described by the interior of a square formed by connecting the midpoints of the image edges. The other region is everything else (in essence, the image corners). The corners produced 87% of the candidates, while only covering 50% of the image. This leaves ten candidates in the central

50% of twenty-five images, as compared with five candidates in the entirety of the same images taken on Dec. 28, 1989. We attribute the factor of four in the rate of finding candidates to an increase in object variability on time scales of a year as opposed to time scales of one month. Future improvements to the subtraction process will concentrate on accurate estimation of the stellar widths of the images and the ratio between the images at *all* locations.

Detection of simulated supernovae

In order to test the efficiency of the subtraction routine, we have placed many fake supernovae in images, then run the subtraction program to determine what fraction of these simulated supernovae we detect. The efficiency for finding supernovae depends on the apparent magnitude of the supernova, the apparent magnitude of the host galaxy, the position on the CCD, the quality of the images, and the quality of the subtraction. An extra complication to calculating the efficiency is the variable threshold used to find candidates.

The tests were performed by adding a fake supernova, a replica of one of the stars in the image, to each of about one hundred objects in each image. The steps in this process were as follows.

(1) The list of host objects was generated by choosing objects in the image which had integrated magnitudes in a chosen range. A rough star-galaxy separation was attempted by insisting that the peak pixel must be less than 50% of the total light from the object. If there were many more than 100 objects in the list, the list was shortened by deleting objects more or less evenly across the image. Using this method, we do not measure the efficiency as a function of position, but rather measure the average efficiency across the whole image.

(2) A bright, unsaturated star was chosen as the model for the fake supernova. The pixel values in a small box containing the star were copied to another location, background subtracted, and scaled down to the desired magnitude of the fake supernova. By using this technique, we assure that the fake supernova has the correct stellar profile and contributes negligibly to the noise. (Note: Although real supernovae would contribute noise, the amount is small compared to the noise contributed by the sky background and the host galaxy. A 22nd magnitude supernova would contribute 20 counts of noise. Typically, we integrate over an effective area of ten pixels, which would contribute at least 89 counts of background noise. When added in quadrature, the total noise is 91 counts. At most, the increase in noise due to a supernova is 2%, so the test is not a significant overestimate of our ability to find supernovae. If we did not scale a bright star, then the noise in the region of our simulated supernova would be too high by a factor of $\sqrt{2}$.)

(3) The subtraction was started, and was run until the two new images were adjusted with respect to each other.

(4) Our fake supernova image was then superimposed on each of the host objects we chose for this image.

(5) The subtraction algorithm was continued until the final candidates were identified.

(6) Host objects on which we could not have detected a supernova, since it was behind a mask or within 10 pixels of the edge of the image, are eliminated.

(7) The efficiency is the number of fake supernovae found divided by the number of host objects on which supernovae were placed.

To illustrate the relative effects of host galaxy magnitude, supernova magnitude, and the supernova's distance from the core of the host galaxy, we show the results of efficiency tests (Fig. 3.5) from our example image (Fig. 2.9, 2.11, and 3.2). The efficiency for finding supernovae is plotted versus supernova apparent magnitude for supernovae placed: (1) in blank areas of the image, (2) directly on top of 21.9 – 22.2 magnitude host galaxies, (3) directly on top of 20.4 – 20.9 magnitude host galaxies, and (4) one arc second south and one arc second east of the nucleus of 20.4 – 20.9 magnitude host galaxies. As expected, it is slightly more difficult to find supernovae on brighter galaxies, and slightly easier if the supernova is displaced from the core. The dominant effect, however, is the apparent magnitude of the supernova itself.

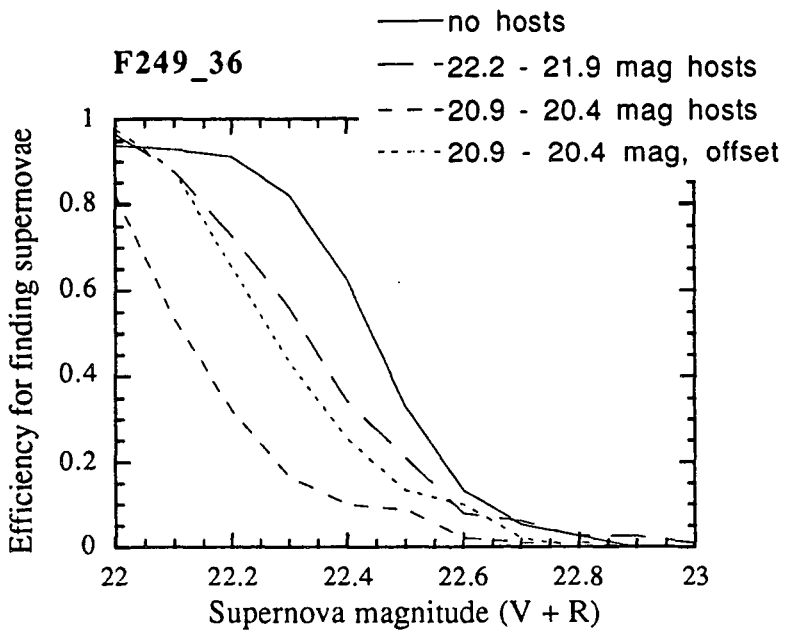

Fig. 3.5. Efficiency vs. supernova magnitude. The efficiency for finding supernovae was determined at intervals of 0.1 magnitude from 22nd to 23rd magnitude for supernovae placed: (1) in blank areas on the image, (2) directly on top of 21.9 – 22.2 magnitude host galaxies, (3) directly on top of 20.4 – 20.9 magnitude host galaxies, and (4) one arc second south and one arc second east of the nucleus of 20.4 – 20.9 magnitude host galaxies. Each determination of the efficiency is based on the results for approximately 100 simulated supernovae.

These efficiency curves can be understood in terms of a simple model. Even though the fake supernovae that are being introduced into a single image are identical, they are being placed on top of noise. Therefore, the number of counts measured for the supernovae in the subtraction should have a Gaussian distribution centered on the number of counts of the fake supernovae, C_{SN}. If the variable threshold in counts is C_T, and the width of the Gaussian distribution is σ, then the expected efficiency of the subtraction is

$$\text{efficiency} = \frac{\int_{C_T}^{\infty} e^{-(C-C_{SN})^2/2\sigma^2} dC}{\int_{-\infty}^{\infty} e^{-(C-C_{SN})^2/2\sigma^2} dC}$$

$$= \frac{1}{\sqrt{2\pi}} \int_{\left[\frac{C_T - C_{SN}}{\sigma}\right]}^{\infty} e^{-z^2/2} dz. \qquad (3.12)$$

For supernovae at exactly the threshold magnitude, the efficiency for finding supernovae should be 50%. As we stated in section 3.8, we set our threshold at the 6.5 σ level. Therefore, equation 3.12 can be rewritten as

$$\text{efficiency} = \frac{1}{\sqrt{2\pi}} \int_{\left[\frac{C_T - C_{SN}}{C_T/6.5}\right]}^{\infty} e^{-z^2/2} dz. \qquad (3.13)$$

Figure 3.6 shows that the single parameter C_T is about 1400 counts (m = 22.5 mag) for this particular image, and that this model works very well for simulated supernovae in blank regions of the image. It would be convenient if the threshold C_T were identical to the variable threshold used by the subroutine *locatestars* in the image analysis section of the subtraction code. Unfortunately, *locatestars* photometry depends heavily on an accurate measurement of the stellar widths of the images on which it operates (see equation 3.8). Since we have already discovered that the calculation of these widths is the weakest point in this algorithm, it should be no surprise that the actual threshold, 1400 counts, is different from the requested threshold, 1050 counts.

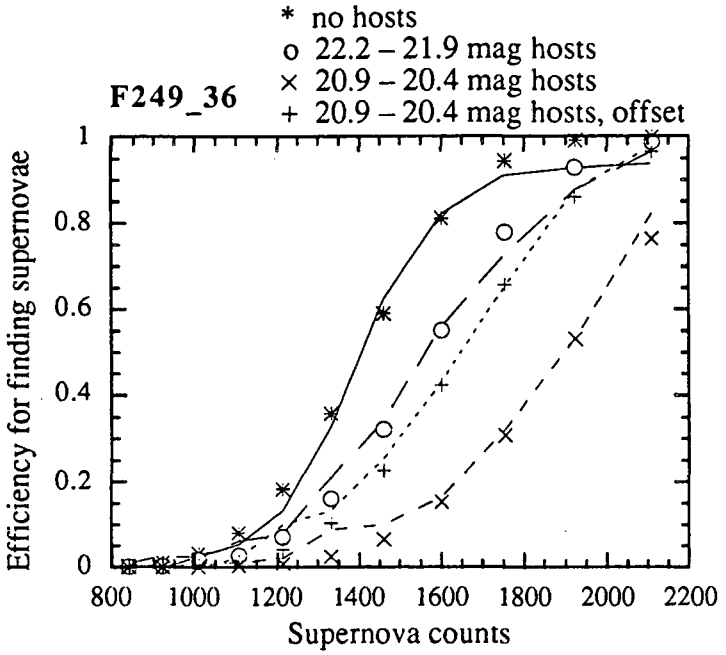

Fig. 3.6. Model fits to the supernova efficiency. The four curves in this diagram correspond to those plotted in Fig. 3.5, except they are plotted vs. counts rather than vs. apparent magnitude. The asterisks, circles, x's and +'s are the fits to the efficiency curves using equation 3.13. The parameter C_T was fit separately for each efficiency curve.

The efficiency for finding supernovae on galaxies is easily understood in terms of the same model. The only difference is that the *locatestars* algorithm is less sensitive to supernovae placed on top of objects. This arises because the galaxies are not perfectly subtracted from each other, leaving a bumpy surface on which the supernova must be detected. The result is a spreading of the stellar profile which reduces the number of counts detected by *locatestars*, and consequently raises C_T. Of course, brighter objects produce bumpier subtractions and higher values of C_T. Objects placed away from the core of the host galaxy will suffer less distortion from the galaxy subtraction.

This method of inserting fake supernovae can be used to calculate efficiency curves for any image we wish to choose. Unfortunately, it is extremely time-consuming and wasteful of computer resources to compute the efficiency for all images. We have found

that a reasonable estimate for C_T for supernova placed on 22nd magnitude galaxies (which provide the majority of our useful luminosity for finding distant supernovae) is obtained by multiplying the *locatestars* threshold by 1.3, which corresponds to a 30% underestimate of the stellar widths in the image. For a good night of observations like Dec. 28, 1989, the average *locatestars* threshold was 1170 counts, but the distribution had a FWHM of about 450 counts. For purposes of estimating the expected supernova rate (section 4.3), we will use an average efficiency curve corresponding to $C_T = 1520$ counts ($m_{V+R} = 22.4$).

4. Expected supernova rate

4.1. The supernova rate

Table 4.1 shows the various determinations of the rate of supernovae of type Ia (also see van den Bergh and Tammann 1991). Several early landmark papers on the supernova rate have been omitted, since they were published before the distinction between supernovae of type I and II had been made. From left to right, the columns give (1) the reference, (2) supernova Ia rate calculated from all galaxies in the survey, in h^2 SNe per $10^{10}L_{Bsun}$ per 100 years, (3) the number of supernovae used to calculate the rate in column 2, (4) the supernova rate in late spirals (typically Sbc-Sd), (5) the number of supernovae used to calculate column 4, (6) the supernova rate in ellipticals and S0 galaxies, (7) the number of supernovae used to calculate column 6, (8) whether type Ia and Ib supernovae were differentiated in calculating the rate, (9) whether or not the supernova rate appeared to depend on galaxy inclination angle, and (10) the limiting magnitude of the survey. Because the methods were so different among different experimenters, an attempt was made to quote only numbers calculated by the authors themselves, and to give information in the table to help the reader evaluate the results. In many cases, though, it was necessary to calculate the table entry from data given in the text; these entries are identified with an asterisk. The uncertainty in supernova rate measurements was not included, since they were usually calculated purely from Poisson statistics, if at all. The systematic error in the measurements is likely to be much more important, though difficult to quantify.

Table 4.1. Supernova Ia rate from various sources

Source	All galaxies		Late spirals		Ellipticals		Diff.	Incl.	m_{lim}
	rate	SN	rate	SN	rate	SN	Ia/Ib?	Eff.?	
Local neighborhood (v<7000 km/sec):									
Barbon (1978)	0.26*	4			0.32*	4	no	no	19.5
Mnatsakanyan et al. (1979)	1.16*	2					no	no	14
Tammann (1982)	1.60*	75*	3.08	52*	0.68*	19	no	yes	?
Tsvetkov (1983)	0.25*	6	0.47	3	0.09	1	no	yes	16
Cappellaro &Turatto (1988)	1.08	26	2.16*	15*	0.68*	8	no	yes	16.5
Evans et al. (1989)	0.28	8.5	0.2	2	0.3	2	yes	no	15
Muller et al. (1992)	0.21	3	0.37	2			yes	no	16.5
High redshift (0.2<z<0.5)									
Couch et al. (1989)	0.15*	1					yes	no	24

The supernova rate is surprisingly difficult to calculate, as is apparent by the large number of correction factors and assumptions used in the various papers. In all of the papers cited here, the supernova rate per galaxy was measured in reference to a "standard" galaxy, where standard is defined as 10^{10} solar blue luminosities. The rates of Tammann (1982) and Cappellaro and Turatto (1988) were adjusted to agree with $H_0 = 100$ km/sec/Mpc as was used in all the other papers. Authors differed on such details as the absolute luminosity of the sun, the wavelength ranges used to calculate absolute magnitude, and whether the galaxy magnitudes should be corrected for internal extinction. In most cases, the surveyed set of galaxies was known, so there was some hope of estimating the

integrated luminosity of each galaxy individually. Mnatsakanyan, Oskanyan, and Lovash (1979), however, merely scanned plates covering 1830 sq. degrees of the sky, so they were forced to estimate the number and magnitudes of the galaxies surveyed.

Two sources of major differences in rate – the differentiation of Ia's from Ib's or Ic's, and the controversial inclination angle corrections – are specified in Table 4.1. Since the difference between supernovae of type Ia and Ib (and the more recent Ic designation) was not recognized until the mid 1980's, earlier supernova searches quoted one rate for all of type I's. It is difficult to determine how much of an effect the Ib's had on the supernova rates, since the Ib's are substantially dimmer. Since Evans et al. (1989), Muller et al. (1992), and Branch (1986) all targeted galaxies at about the same distance, primarily galaxies in the Virgo cluster, we expect that the search with dimmest limiting magnitude would find the highest ratio of type Ib,c's to type Ia's. Evans, van den Bergh, and Mc Clure (1989), with a limiting magnitude of ~15, found that about 71% of their type I supernovae were Ia's (50% in late spirals). By contrast, Muller et al. (1992) found that with a limiting magnitude of ~16.5, 38% of the type I supernovae were of type Ia (29% in late spirals). (Note, however, that Muller et al. could preferentially find Ib,c's due to the redder bandpass of their system.) From published supernova spectra taken in the years 1885-1984, Branch (1986) finds 78% of the I's are Ia's (73% in spirals); spectra were obtained mainly for the brightest supernovae. Of course, all of the supernovae found in elliptical galaxies are of type Ia.

The inclusion of type Ib's in the supernova rates of the late 1970's and early 1980's partially explains the relatively high rates found during this period. However, it is likely that a second source of the inflated rates comes from the correction due to inclination of spiral galaxies. For these photographic searches, it was fashionable to correct the supernova rates of inclined spirals by factors of 2 to 6, since the measured rates appeared to

be artificially low in the highly inclined galaxies. (Tsvetkov calculated the absorption expected in each galaxy individually, rather than multiplying the resulting rate by an ad hoc factor.) The more recent supernova searches of Evans et al. (detector is a human eye) and Muller et al. (CCD detector) do not see this inclination effect, nor do they see the inflated rates. Some of the correction could be justified due to difficulties in photographically detecting supernovae on the higher surface brightness inclined spirals. However, it seems likely that the photographic searches overcorrected for the inclination effect.

It is interesting to compare the galaxies targeted by the various searches. Most of the supernova searches targeted relatively nearby galaxies in standard catalogs, with a large majority coming from the Virgo cluster. The exceptions to this rule are (1) Barbon (1978) targeted only the Coma cluster of galaxies, (2) Mnatsakanyan et al. (1979) targeted large sections of the sky, and (3) Couch et al. (1989) targeted only rich clusters in the range $0.2 < z < 0.5$. Given that the data in the table exhibit quite a wide range of supernova rates and a variety of biases, it is difficult to see any clear difference between the rates in clusters as opposed to field galaxies, or the rate in ellipticals as opposed to spirals. Although it may seem that the supernova rate in distant clusters calculated from Couch et al. is slightly low, it is based on only one supernova discovery, and is thus highly uncertain. In addition, Couch et al. found another supernova which was assumed to be a type II, but they could not rule out the possibility that it, too, was a type Ia.

4.2. Projected discovery rate for SNe Ia

The rate at which we expect to find supernovae at the AAT is dependent on the supernova rate per luminosity, the integrated luminosity of the galaxies in the field, the surveillance time for each galaxy, the efficiency for finding supernovae, and the number of fields observed. Unfortunately, many of these factors are interdependent or hard to

quantify. For example, the efficiency for finding a given supernova depends on its apparent magnitude and on the luminosity of the parent galaxy. The efficiency is not a constant over the nominal surveillance time, since the supernova increases and decreases in brightness. The surveillance time can nominally be calculated from the expected light curve of a redshifted supernova, but these curves are not easy to generate, since detailed knowledge of the Ia spectra as a function of time are required to calculate the *K*-correction.

In order to make the calculation easier, we have chosen to adopt the following simplifications:

(1) We assume that the efficiency function is a step function at the average threshold of $m_{V+R} = 22.4$ mag. In other words, we would find one hundred percent of the supernovae brighter than magnitude 22.4 mag, and none of the dimmer supernovae. Since we generally gain more galaxy years from a dimmer cutoff than we lose if the cutoff is brighter by a similar amount, this simplification artificially decreases our expected supernova rate.

(2) The subtraction algorithm cannot find supernovae that are 10% or less of the luminosity of the host galaxy. Since the majority of the galaxies which are intrinsically bright enough for this to be a problem are excluded because their apparent magnitude is brighter than our limit of $m_{V+R} = 20.4$, we will assume that the difference in surveillance time due to this effect is negligible. This simplification artificially increases our expected supernova rate.

(3) To calculate the surveillance times, we use the composite supernova Ia blue light curve shape from Cadonau, Sandage, and Tammann (1985). The peak absolute magnitude (including about 0.2 magnitudes of extinction in the supernova's host

galaxy) is taken to be -18.1 mag for H_0 = 100 km/sec/Mpc (Miller and Branch 1990). The expected light curve is calculated for a given redshift by shifting the peak magnitude by the appropriate distance modulus, and time dilating the surveillance time by a factor of (1+z). We do not calculate the *K*-correction, which takes into account the shifting of spectral lines into and out of our filter. The number of days the supernova will remain brighter than magnitude 22.4 is tabulated in Table 4.2.

(4) The supernova rate per one hundred years per blue luminosity of $10^{10}L_{sun}$ is conservatively estimated as 0.20.

Using these simplifications, we need only the luminosity density as a function of redshift to calculate the expected supernova rate.

Luminosity density as a function of redshift

From Fig. 2.15, we see that there are as many as 2000 objects in each image to a limiting magnitude of 23. The majority of these objects are galaxies. However, most of the galaxies are very luminous, distant galaxies on which we have no chance of finding supernovae. In order to estimate the supernova rate, we estimate the total galactic luminosity as a function of redshift using the Schechter luminosity function (Schechter 1976). The number of galaxies per unit volume in the luminosity interval from L to L+dL is given by

$$\varphi(L)dL = \varphi^* (L / L^*)^{\alpha} \exp(-L / L^*) d(L / L^*). \qquad (4.1)$$

From Kirshner et al. (1983), the constants are given by

$$\alpha = -1.25,$$

$$\varphi^* = 1.5 \times 10^{-3}\ \mathrm{Mpc}^{-3},$$

$$\text{and } L^* = 1.31 \times 10^{11}\, h\, L_{sun} \qquad (4.2)$$

for field galaxies. These constants are very similar to the constants previously found by Schechter for galaxies in clusters.

We are interested in finding the total luminosity per AAT image as a function of redshift. We have in addition the constraint that the apparent magnitude of the galaxy must be fainter than magnitude 20.4, and, if it is to be seen on the image, it must be brighter than 23rd magnitude. The total luminosity enclosed is given by

$$\text{total luminosity between } z_1 \text{ and } z_2 = \int_{z_1}^{z_2} A(z) \int_{L_1(z)}^{L_2(z)} L\varphi(L)dL\ dz, \qquad (4.3)$$

where A(z)dz is the differential volume element and L_1 and L_2 are the luminosities for which the apparent magnitudes will be 23.0 and 20.4, respectively. From the standard, matter-dominated formulation of cosmology, these functions are given by (Zombeck 1990):

$$A(z) = \frac{c^3(\Delta\Omega)}{H_o^{\,3}(1+z)^3} \frac{\left\{q_o z + (q_o - 1)\left[(1+2q_o z)^{1/2} - 1\right]\right\}^2}{q_o^{\,4}(1+2q_o z)^{1/2}}, \qquad (4.4)$$

$$L_i(z) = L_{sun}\, 10^{(M_{sun} + \mu(z) - m_i)/2.5},\ i = 1,2, \qquad (4.5)$$

$$\text{and } \mu(z) = 5\log_{10}(d_L(z)/10pc). \qquad (4.6)$$

Here, $\Delta\Omega$ is the solid angle subtended by one image, $\mu(z)$ is the distance modulus, and d_L is the luminosity distance as given by equation 5.29. The results of the numerical integration of this function for $q_0 = 0.5$ and $q_0 = 0.1$ are given in Table 4.2 and Table 4.3, respectively.

The expected supernova rate

From here, calculation of the expected supernova rate is easy. The number of normalized galaxies (defined as $10^{10}L_{sun}$) is multiplied by the surveillance time in years for

each redshift interval to produce the number of galaxy-years per image in that interval. Then, the total number of galaxy-years is calculated by summing over all relevant redshift intervals. Tables 4.2 and 4.3 show the results for a limiting magnitude of 22.4 (the average limit for our current search). We acquire between 4.00 and 5.29 galaxy-years for each image observed, depending on the value of q_o. Since we mask about 6% of the image and ignore about 4% of the image around the edges, the expected supernovae rate for $q_o = 0.5$ is

$$\frac{1 \text{ SN}}{500 \text{ h}^{-2}\text{gal.–yrs.}} \frac{(0.90)\ 4.00 \text{ h}^{-2}\text{gal.–yrs.}}{1 \text{ image}} = \frac{1 \text{ SN}}{139 \text{ images}}, \tag{4.7}$$

independent of h. If qo is 0.1, it would only require 105 images. If, on the average night we can gather about 55 images, this results in about one supernova every two to two and a half nights. It is estimated that 20% of these supernovae will occur in galaxies dimmer than 23rd magnitude, so they will occur on what will appear to be a blank spot in the image.

The number of galaxy years per image can be increased by lengthening the exposure time and thus decreasing the threshold for finding supernovae. For example, to reach a limiting magnitude of 23.0, we would have to increase the exposure time by a factor of three, which corresponds to about a factor of 2.5 in real time, since there is overhead in pointing and readout for each exposure. We can see from Tables 4.2 and 4.3 that this coincidentally increases the integrated galaxy years per image by a factor of 2.5, so we would break even on the number of supernovae discovered per night. Since it is currently impossible to spectroscopically verify supernovae more distant than $z \approx 0.4$ (because they are too dim), we have opted for the shorter exposure time, which favors supernovae at a redshift closer to 0.3.

One of the shortcomings of this search is the emphasis on field galaxies rather than clusters. At the time that the reference pictures were taken, it was assumed that there would be an insignificant number of supernovae in cluster galaxies as compared with supernovae

in field galaxies, partly because it was thought the integrated luminosity in field galaxies would be much larger, and partly because it was thought that the supernova rate in cluster galaxies (which are preferentially elliptical) was significantly lower. Neither of these assumptions has been borne out by further research.

In fact, Table 4.2 shows the desirability of including clusters in the fields that are searched. If we include a modest rich cluster at $z = 0.3$ containing $5 \times 10^{11}\ h^{-2}$ solar luminosities (50 normalized galaxies), then we would practically double our expected rate for finding supernovae. In addition, if the supernova were discovered in the cluster, we would have a better means of measuring its redshift. If one assumes a space density of rich clusters of $6 \times 10^{-6}\ h^3\ Mpc^{-3}$ (Bahcall 1988), then one finds there are 5,000 rich clusters in the redshift range $0.24 < z < 0.34$ above a Galactic latitude of 60°. If we chose fields randomly, we expect only one in 66 would contain a rich cluster in this range purely by chance. It is clear that it is possible, and in our best interest, to target high redshift clusters.

Table 4.2 - Galaxy years per image, H_0 = 100 km/sec/Mpc, q_0 = 0.5

z	normalized galaxies (visible)	normalized galaxies (invisible)	# of days visible, m_{lim}=22.4	gal. years on visible galaxies, m_{lim}=22.4	total gal. years, m_{lim}=22.4	# of days visible, m_{lim}=23.0	total gal. years, m_{lim}=23.0
0.18–0.20	1.67	0.36	38	0.17	0.21	49	0.27
0.20–0.22	2.21	0.49	35	0.21	0.26	45	0.33
0.22–0.24	2.83	0.65	32	0.25	0.31	43	0.41
0.24–0.26	3.53	0.83	31	0.30	0.37	40	0.48
0.26–0.28	4.30	1.04	29	0.34	0.42	37	0.54
0.28–0.30	5.12	1.27	28	0.39	0.49	36	0.63
0.30–0.32	5.99	1.54	26	0.43	0.54	34	0.70
0.32–0.34	6.90	1.82	23	0.44	0.55	33	0.79
0.34–0.36	7.83	2.15	19	0.41	0.52	31	0.85
0.36–0.38	8.76	2.50	11	0.26	0.34	29	0.89
0.38–0.40	9.70	2.88	0	0.00	0.00	29	1.00
0.40–0.42	10.62	3.29	0	0.00	0.00	24	0.92
0.42–0.44	11.51	3.73	0	0.00	0.00	21	0.88
0.44–0.46	12.37	4.19	0	0.00	0.00	15	0.68
0.46–0.48	13.18	4.68	0	0.00	0.00	8.8	0.43
0.48–0.50	13.95	5.20	0	0.00	0.00	8.9	0.47
Total	120.47	36.62		3.20	4.00		10.26

Table 4.2 - Galaxy years per image, H_0 = 100 km/sec/Mpc, q_0 = 0.1

z	normalized galaxies (visible)	normalized galaxies (invisible)	# of days visible, m_{lim}=22.4	gal. years on visible galaxies, m_{lim}=22.4	total gal. years, m_{lim}=22.4	# of days visible, m_{lim}=23.0	total gal. years, m_{lim}=23.0
0.18–0.20	2.02	0.44	38	0.21	0.26	49	0.33
0.20–0.22	2.73	0.61	35	0.26	0.32	45	0.41
0.22–0.24	3.56	0.82	32	0.31	0.38	43	0.52
0.24–0.26	4.51	1.07	31	0.38	0.47	40	0.61
0.26–0.28	5.57	1.37	29	0.44	0.55	37	0.70
0.28–0.30	6.73	1.72	28	0.52	0.65	36	0.83
0.30–0.32	7.98	2.11	26	0.57	0.72	34	0.94
0.32–0.34	9.30	2.56	23	0.59	0.75	33	1.07
0.34–0.36	10.67	3.07	19	0.56	0.72	31	1.17
0.36–0.38	12.07	3.64	11	0.36	0.47	29	1.25
0.38–0.40	13.48	4.27	0	0.00	0.00	29	1.41
0.40–0.42	14.89	4.96	0	0.00	0.00	24	1.31
0.42–0.44	16.26	5.71	0	0.00	0.00	21	1.26
0.44–0.46	17.59	6.53	0	0.00	0.00	15	0.99
0.46–0.48	18.85	7.42	0	0.00	0.00	8.8	0.63
0.48–0.50	20.03	8.37	0	0.00	0.00	8.9	0.69
Total	166.24	54.67		4.20	5.29		14.12

5. Calculating q_0 from supernovae

5.1. Finding q_0 in a Friedmann universe

It is not the intent of this section to teach general relativity, but merely to point out the major signposts and to illuminate the assumptions in the derivation of equation 5.30, which describes the measurement of q_0 using standard candles. I follow the notation and derivation of Steven Weinberg in *Gravitation and Cosmology* (1972).

We begin by adopting general relativity and the Cosmological Principle (the hypothesis that the universe is spatially homogeneous and isotropic). We will assume that the Cosmological Principle holds in general, even though we know it breaks down on scales less than or equal to the size of galaxy clusters. One can then derive that the universe must have the Robertson–Walker metric (c = 1):

$$(d\tau)^2 = (dt)^2 - R^2(t)\left\{\frac{(dr)^2}{1-kr^2} + r^2(d\theta)^2 + r^2\sin^2\theta(d\phi)^2\right\}. \qquad (5.1)$$

The critical unknown quantities in this metric are R(t), the cosmological scale factor, and k, a constant which takes on the values -1, 0 , or +1, depending upon whether the universe is open, critical, or closed, respectively. An important property of this metric is that an object at rest with respect to r, θ, and ϕ, will remain at rest at that coordinate position. This results from the absence of mixed temporal and spatial terms in metric tensor (i.e., the coefficients of the (dr dt), (dθ dt), and (dϕ dt) terms are zero), in addition to the coefficient of the $(dt)^2$ term being independent of spatial coordinates. One can define cosmological standard time at each coordinate position by placing synchronized clocks at rest with respect to this coordinate system (dr = dθ = dϕ = 0). The clocks will read the proper

time for a freely falling body at rest in our comoving coordinate system. Notice, however, that the proper distance between two coordinate positions is proportional to R(t).

Since this simple universe expands and contracts under its own self gravity, the evolution of the cosmological scale factor R(t) can be found from the gravitational field equations of general relativity:

$$R_{\mu\nu} = -8\pi G\, S_{\mu\nu},$$
$$\text{where } S_{\mu\nu} = T_{\mu\nu} - \frac{1}{2} g_{\mu\nu} T^{\lambda}{}_{\lambda}. \tag{5.2}$$

This is the simplest form of Einstein's equations (without the extensions of a cosmological constant or the Brans-Dicke theory, for example); in gross terms, it states that the "curvature" of space, $R_{\mu\nu}$, is a result of the energy and momentum sources, $S_{\mu\nu}$, which is in turn derived from the energy-momentum stress tensor, $T_{\mu\nu}$. Applying these equations to the Robertson -Walker metric, we find the simplified equations for evolution:

$$3\ddot{R} = -4\pi G(\rho+3p)R \tag{5.3}$$

$$\text{and } R\ddot{R} + 2\dot{R}^2 + 2k = 4\pi G(\rho-p)R^2. \tag{5.4}$$

These can be written as the first order differential equations

$$\dot{R}^2 + k = \frac{8\pi G}{3}\rho R^2 \tag{5.5}$$

$$\text{and } \dot{\rho}R^3 = \frac{d}{dt}\left\{R^3(\rho+p)\right\}. \tag{5.6}$$

If the universe is dominated by nonrelativistic matter with negligible pressure, then equation 5.6 gives us

$$\rho \propto R^{-3} \quad \text{for} \quad p << \rho. \tag{5.7}$$

Otherwise, for highly relativistic particles,

$$\rho \propto R^{-4} \quad \text{for} \quad p = \frac{\rho}{3}. \tag{5.8}$$

Equations 5.1, 5.5, 5.6, 5.7, and 5.8 serve to define the standard model, the Friedmann universe.

From this standard model, we would like to derive a relationship between the apparent magnitude of our standard candles, the redshift, and q_o. We know from the definition of absolute magnitude and luminosity distance that

$$m = M + 5\log_{10}\left(\frac{d_L}{10\ pc}\right). \tag{5.14}$$

So, we will have our desired relationship if we can find $d_L(q_o,z)$.

Consider a point source at the coordinate center of the universe (r = 0) at time t. Now imagine that there is a detector at coordinate position r_o, time t_o,and with area A_o that detects light from the point source. If the point source radiates a monochromatic pulse of light equally in all directions for a duration Δt, then the fraction of photons that reach the detector is

$$\frac{A_o}{\text{The area of a surface at a distance } R(t_o)r_o}. \tag{5.15}$$

Using the line elements in the Robertson-Walker metric, we can integrate over the surface at $R(t_o)r_o$ to find that

$$\text{fraction of photons} = \alpha = \frac{A_o}{4\pi R^2(t_o)r_o^2}. \tag{5.16}$$

The luminosity of the light seen at $r = r_o$ is the photon energy times the number of photons per second per detector area. We know that due to time dilation, $\Delta t_o = [R(t_o)/R(t)]\Delta t$. In addition, the energy of each photon is redshifted, so the total energy reaching our detector is $E_o = \alpha[R(t)/R(t_o)]E$, where E is the total energy emitted by the source over a time period Δt. Thus, the luminosity is

$$L_o = \frac{E_o}{\Delta t_o\ A_o} = \frac{A_o}{4\pi R^2(t_o)r_o^2}\frac{R^2(t)}{R^2(t_o)}\frac{1}{A_o}\frac{E}{\Delta t} = \frac{R^2(t)}{4\pi R^4(t_o)r_o^2}\frac{E}{\Delta t}. \tag{5.17}$$

The luminosity distance d_L is defined such that the Newtonian inverse square law

$$L_o = \frac{1}{4\pi d_L^{\,2}} \frac{E}{\Delta t} \tag{5.18}$$

works. Thus, we find that the luminosity distance is given by

$$d_L = \frac{R^2(t_o) r_o}{R(t)} = R_o r_o (1+z). \tag{5.19}$$

At this point, we have a choice of expanding d_L in an approximate power series or solving for d_L exactly using a matter-dominated or radiation-dominated model universe. The luminosity distance in powers of z is

$$d_L = H_0^{-1} [z + 0.5 (q_0 - 1) z^2 + ...]. \tag{5.20}$$

This should hold within 10% for redshifts less than 0.5. However, since we need higher accuracy, we will use the "exact" solution for $R_o r_o$ as a function of H_o,q_o and z for the special case of a matter-dominated universe.

Our solutions to the Einstein equations (5.5, 5.6) give us

$$\rho_o = \frac{3}{8\pi G}\left(\frac{k}{R_o^{\,2}} + H_o^{\,2}\right) \tag{5.21}$$

and

$$p_o = -\frac{1}{8\pi G}\left[\frac{k}{R_o^{\,2}} + H_o^{\,2}(1 - 2q_o)\right] \tag{5.22}$$

for the current epoch. For the most accepted assumption that we live in a matter-dominated universe (so that $p_o << \rho_o$), we find that

$$k = [(2q_o - 1)] R_o^{\,2} H_o^{\,2}, \tag{5.23}$$

where

$$q_o = \frac{1}{2}\frac{\rho_o}{\rho_c} = \frac{\Omega_o}{2}. \tag{5.24}$$

Consider again the light ray traveling from r = 0 to us. The Robertson-Walker metric for this case is

$$0 = (dt)^2 - \frac{R^2(t)(dr)^2}{1 - kr^2}. \tag{5.25}$$

Rearranging this,

$$\int_0^{r_o} \frac{dr}{\sqrt{1-kr^2}} = \int_{t_z}^{t_o} \frac{dt}{R(t)} = \int_{R_z}^{R_o} \frac{dR}{R\dot{R}}. \tag{5.26}$$

With the help of equation 5.5,

$$\int_0^{r_o} \frac{dr}{\sqrt{1-kr^2}} = -\frac{1}{H_o R_o} \int_{(1+z)^{-1}}^{1} \frac{1}{x}\left(1 - 2q_o + \frac{2q_o}{x}\right)dx. \tag{5.27}$$

For any of the three values of k, this integrates to

$$r_o = \frac{zq_o + (q_o - 1)\left(-1 + \sqrt{2q_o z + 1}\right)}{H_o R_o q_o^2 (1+z)}. \tag{5.28}$$

Plugging this into equation 5.19 gives us the final result for the luminosity distance in a matter-dominated Friedmann universe using the simplest form of Einstein's equations:

$$d_L = \frac{1}{H_o q_o^2}\left[zq_o + (q_o - 1)\left(-1 + \sqrt{2q_o z + 1}\right)\right]. \tag{5.29}$$

Including factors of c, we can measure q_0 in the standard matter-dominated Friedmann universe from the implicit relation

$$m = M + 5\log\left\{\frac{c}{H_o q_o^2}\left[zq_o + (q_o - 1)\left(-1 + \sqrt{2q_o z + 1}\right)\right]\right\} + 25, \tag{5.30}$$

where the luminosity distance d_L is measured in Mpc, H_0 is measured in km/sec/Mpc, and c is measured in km/sec. The Hubble constant, H_0, should be the same as that used to calculate the intrinsic absolute magnitude of the supernova. In this equation, H_0 is a bit of a red herring, since its contribution to computing M cancels its effect in the equation. The relationship between m and z for supernovae of type Ia is plotted in Fig. 5.1 for several values of q_0.

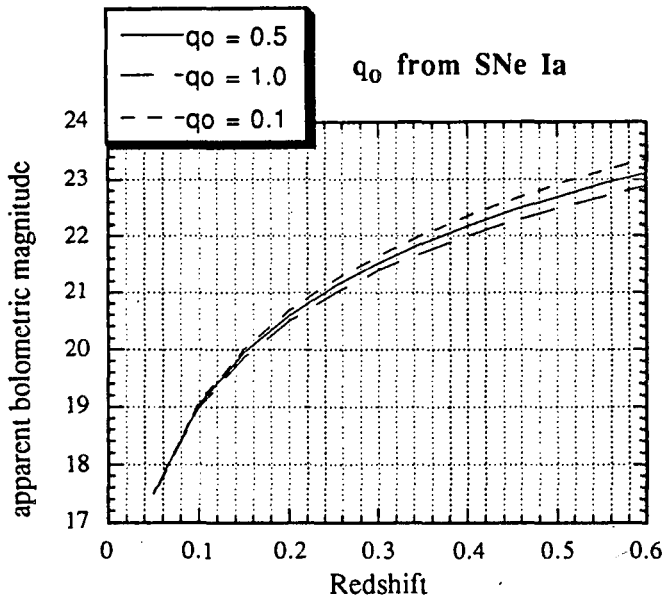

Fig. 5.1. Measuring q_o from SNe Ia. We show the relationship between apparent magnitude and redshift for several values of q_o. The bolometric magnitude for a supernova at maximum is taken to be M=19.9 mag for Ho=50 km/sec/Mpc, as given by Leibundgut and Tammann (1990). For photometry in filters, the K-correction must be included as well, so the actual shape of these curves would be slightly different.

Since the luminosity distance from which equation 5.30 is derived uses the *total* energy generated at the source, the magnitudes are generally specified as *bolometric*. However, the luminosity distance derivation holds as long as we agree to measure the same set of photons that were measured in the rest frame of the source. This requires that the wavelengths measured on earth be redshifted from the wavelengths used at the source to calculate the absolute magnitude. For example, if the distant supernova is measured in a wavelength range $5000 < \lambda < 7000$ Angstroms, the apparent magnitude must be compared with the absolute magnitude in the wavelength range $[5000/(1+z)] < \lambda < [7000/(1+z)]$. There is an additional correction due to the zero point of the magnitude scale being based on an A0V star. Equation 5.30 works for any wavelength range $\lambda_1 < \lambda < \lambda_2$, provided the absolute magnitude, M, is replaced by

$$M\left(\frac{\lambda_1}{(1+z)} < \lambda < \frac{\lambda_2}{(1+z)}\right) + 2.5\log_{10}\left(\frac{\text{abs. lum. of an A0V star, } \lambda_1 < \lambda < \lambda_2}{\text{abs. lum. of an A0V star, } \frac{\lambda_1}{(1+z)} < \lambda < \frac{\lambda_2}{(1+z)}}\right) \quad (5.31)$$

5.2. Error in determination of q_0

From equation 5.30, the only quantities that are necessary to measure q_0 are

(1) the redshift of the supernova

(2) the (corrected) absolute magnitude of supernovae of type Ia at maximum, and

(3) the measured apparent magnitude of the supernova at maximum.

I will discuss the error in the measurement of q_0 contributed from each of these.

Error in redshift measurement (z)

We estimate the error in redshift to be that introduced by the peculiar velocity of the host galaxy, or one percent, whichever is greater. Since peculiar velocities in galaxy clusters can be on the order of ~1000 km/sec and the redshift velocity for z = 0.3 is ~90,000 km/sec, the accuracy of z will be about 1%. To estimate the effect of this error on q_0, we must find $\partial q_0(z,m,M)/\partial z$ by differentiating equation 5.30 with respect to z. This partial derivative is

$$\left.\frac{\partial q_o}{\partial z}\right|_{m,M} = \frac{q_o^2\left\{\sqrt{2q_o+1}-(1-q_o)\right\}}{\left(q_o^2 z-3q_o z+q_o-2\right)+\left\{2-q_o(1-z)\right\}\sqrt{2q_o+1}}. \qquad (5.32)$$

For $q_0 = 0.5$, the value of $\partial q_0/\partial z$ is 54.9, 25.4, and 14.9, for z equal to 0.2, 0.3, and 0.4, respectively. For smaller values of q_0, the derivative is slightly smaller. Therefore, an upper limit on the error in q_0 due to error in cosmological redshift is

$$\Delta q_o = \frac{\partial q_o}{\partial z}\Delta z \approx \begin{cases} 0.18, & z = 0.2 \\ 0.08, & z = 0.3 \\ 0.06, & z = 0.4 \end{cases}. \qquad (5.33)$$

As we shall see, this error is relatively small compared to the uncertainties introduced by the other two sources.

Error in the absolute magnitude of supernovae type Ia at maximum

The measurement of q_0 using supernovae of type Ia is based on the assumption that these supernovae are standard candles. Regardless of whether supernovae Ia are uniform events, the magnitudes observed on earth vary due to absorption by gas and dust in the host galaxy. Since it is very difficult to measure this extinction, it will be necessary to model the effects of dust on the distribution of supernova magnitudes. First, we assume that galaxies at a redshift of 0.3 are approximately the same as those with recessional velocities less than 2000 km/sec (although, see Butcher and Oemler 1984). Then, we can use the distribution of magnitudes from these nearby supernovae to calculate the distribution expected for various values of redshift and detection threshold. The high redshift distribution will differ from that of nearby supernovae, since the magnitude cutoff of our search biases us towards finding less extinguished or intrinsically brighter supernovae.

From studies of nearby supernovae, it is known that the peak magnitude of SNe Ia is stable within 30%, probably within 20%, and possibly within 10% (see section 2.2). However, due to uncertainty in extinction due to dust, it is unlikely that we would be able to correct the magnitude to better than 30%, or 0.3 magnitudes, on an individual supernova. For high redshift supernovae, the expected signal will be calculated by multiplying the system response of our detector by a redshifted standard type Ia spectrum at maximum (see section 5.1). It is reasonable to expect that the consistency of peak magnitude will be about the same in these *K*-corrected wavelength bands as it is in the blue.

In order to estimate the error this introduces to our measurement of q_0, we calculate the first partial derivative:

$$\left.\frac{\partial M}{\partial q_o}\right|_{z,m} = -\frac{5}{\ln 10}\frac{-\frac{2}{q_o}\left(q_o z+(q_o-1)\left(-1+\sqrt{1+2q_o z}\right)\right)+z-1+\frac{z(q_o-1)}{\sqrt{1+2q_o z}}+\sqrt{1+2q_o z}}{q_o z+(q_o-1)\left(-1+\sqrt{1+2q_o z}\right)}. \quad (5.34)$$

Richard Muller made the amazing discovery that for small values of z, this daunting expression reduces to

$$\left.\frac{\partial M}{\partial q_o}\right|_{z,m} \cong z. \quad (5.35)$$

Therefore, the error introduced into the measurement of q_0 is

$$\Delta q_o = \left.\frac{\partial q_o}{\partial M}\right|_{z,m} \Delta M \cong \frac{\Delta M}{z} \cong 1. \quad (5.36)$$

Clearly, we're going to have to beat down the errors by observing many supernovae.

Error in the measured peak magnitudes of supernovae at $z \approx 0.3$

In order to measure the peak magnitude of a supernova, one must identify the time of maximum in addition to measuring the apparent magnitude at various points along the light curve. If the type Ia light curve were completely standard, one could measure the peak apparent magnitude by identifying one of the two major features of the light curve: the peak, or the knee which occurs 30 days later. The knee is the more difficult of the two, since it occurs when the supernova is three magnitudes dimmer than maximum.

Since the expected visual magnitude of a high redshift supernova is 21 or dimmer, we will make the assumption that the sky noise is much more important than any additional Poisson noise due to the supernova or galaxy. Then

$$\frac{S}{N} = L_{SN}\sqrt{\frac{\varepsilon q A t}{\pi\sigma^2 L_{sky}}}, \quad (5.37)$$

where

A ≡ collecting area of telescope,

$\varepsilon \equiv$ efficiency of optics & filter,

$q \equiv$ quantum efficiency of CCD,

$t \equiv$ time of observation,

$\sigma \equiv$ FWHM of a stellar profile in arcsec,

$L_{SN} \equiv$ photons/sec/collecting area from the supernova, and

$L_{sky} \equiv$ background photons/sec/arcsec^2/collecting area.

The knee occurs when the supernova is 1/16 the brightness of maximum light. From equation (5.37), we see that it would require 256 times as long to get the same signal-to-noise ratio, and it would realistically be much worse due to confusion of the signal with the host galaxy. Therefore, it is crucial to find the supernovae at or before maximum light.

The uncertainty in magnitude due to Poisson noise fluctuations is

$$\delta m \sim \sqrt{2}\left(\frac{2.5}{\ln 10}\right)\frac{N}{S}, \tag{5.38}$$

where the $\sqrt{2}$ has been added because we must compare two images to find the increase in light due to the supernova. For the AAT, with our CCD and a V+R filter, we reach $\Delta m<0.1$ mag in 10 minutes. This is similar to the accuracy we should expect from a one hour image with a 1m class telescope located at a 1" seeing site, if we use a filter passband that is about half as wide. Including uncertainty in the time of maximum, it is reasonable to assume we can find the maximum to within $\Delta m<0.2$ mag for a supernova discovered at or before maximum, if we can get a few well calibrated photometric points.

From equations 5.30, and 5.35, we see that

$$\left.\frac{\partial q_o}{\partial m}\right|_{z,M} = -\left.\frac{\partial q_o}{\partial M}\right|_{z,m} \cong -\frac{1}{z}. \tag{5.39}$$

Therefore, the error introduced by measuring the apparent magnitude is approximately

$$\Delta q_o = -\left.\frac{\partial q_o}{\partial m}\right|_{z,M} \Delta m \cong \frac{\Delta m}{z} \cong 0.67. \tag{5.40}$$

The three sources of error are independent, so we can add their effects in quadrature. The total uncertainty in q_o for one supernova will be

$$\Delta q_o = \sqrt{\left(\Delta q_o^{\,2}(z) + \Delta q_o^{\,2}(M) + \Delta q_o^{\,2}(m)\right)} \cong 1.2. \qquad (5.41)$$

The uncertainty decreases as $N^{-1/2}$, where N is the number of supernovae discovered at or before maximum. Given the assumptions made in this paper, I estimate that we will need 36 supernovae to measure q_o within $\Delta q_o = 0.2$, and 144 supernovae to measure q_o within $\Delta q_o = 0.1$. One can can theoretically do better by finding supernovae at higher redshift, but follow-up observations will be much tougher.

There are several important results that must be found from nearby supernovae in order to make this determination possible. One important result we are missing is a characteristic time series of spectra for supernovae Ia. These will be necessary to calculate the *K*-corrected light curves for supernovae at arbitrary redshift. In addition to this, we need a reliable distribution of absolute magnitudes for supernovae Ia that are *uncorrected* for extinction in the host galaxy, but *corrected* in the sense of equation 5.31. This distribution and the limiting magnitude of the supernova search determine the unbiased absolute magnitude expected as a function of redshift. Clearly, the measurement of q_o from supernovae could benefit from a better understanding of nearby supernovae Ia near maximum light.

6. Conclusion

The search for distant supernovae at the Anglo-Australian Observatory has produced a very popular f/1 camera for the 4m Anglo-Australian Telescope, a very sophisticated software system designed to find variable objects by image subtraction, and a set of fields that will be very useful for studying quasar variability in the future. Many lessons were learned in this attempt to find distant supernovae, including the desirability of observing clusters, the need for high resolution images, and the importance of scheduled dark time follow-up observations within a week after supernova searching. It is important to search clusters to increase the probability of finding supernovae. It is important to have images with seeing of about 1" so that galaxies can be distinguished from stars and QSO's, and to increase the chance that a supernova will be resolved from the galaxy core. It is important to have scheduled follow-up at a large telescope, so that we can spectroscopically verify the discovery of a supernova. We have found that without these three elements, our best efforts produce only lists of possibly interesting objects, and no substantiated supernovae.

The second generation of this search, although still in its infancy, has already demonstrated the advantages of solving the first two of these difficulties. Fig. 6.1 shows a possible supernova discovered in March, 1992, on the 2.5 meter Isaac Newton Telescope in the Canary Islands. We can be certain that we are looking at a galaxy, and it is probable that the increase in brightness that we observed is offset from the galaxy core. In addition, since we were able to get a longer block of time and analyze these images within days of the initial observations, we were able to get several images of the galaxy. Due to bad weather and equipment failures, it was not possible for us to get follow-up spectroscopy for this object, even though we had arranged for override time.

Although it may take several attempts to find an efficient system for finding distant supernovae, the value of a measurement of q_0 from supernovae makes it worth pursuing. In preparation for the time when distant supernovae can be found in sufficient numbers, there are at least two results that are needed from studies of closer supernovae Ia: a complete time series of typical spectra of supernovae of type Ia, and a distribution of supernova Ia peak absolute magnitudes, calculated *without* corrections for extinction in the host galaxy. The spectral series will make it possible to make *K*-corrections for supernovae at any distance. The distribution of peak magnitudes will make it possible to model the bias produced by making magnitude-limited observations of objects at many distances. In general, a better understanding of type Ia supernovae will produce a more accurate determination of q_0.

Even those who doubt that the best way to measure the cosmological deceleration parameter is to use type Ia supernovae cannot deny the value of this independent determination of q_0. It is our hope that the advances in image processing and supernova search strategy that we have made with our pilot search will eventually lead to a system capable of finding the supernovae necessary to make a cosmologically interesting measurement of q_0.

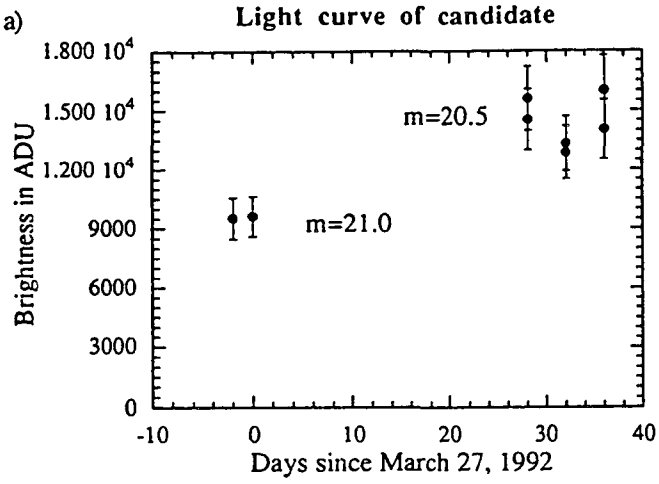

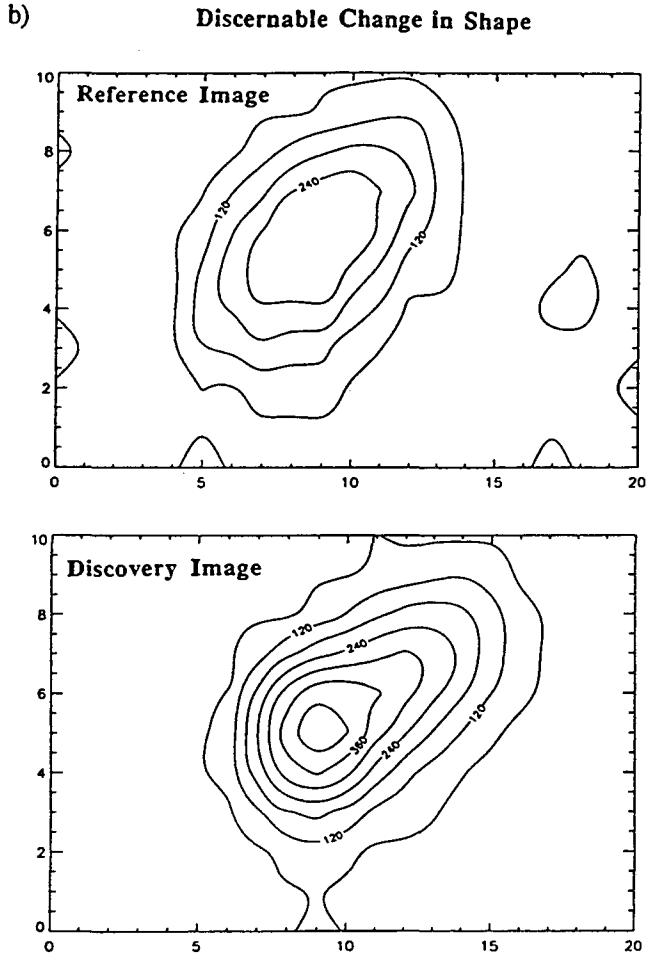

Fig. 6.1. Possible supernova. There was one candidate from the March, 1992 data run on the INT that shows all of the indications of a supernova. The host galaxy is in close association with a galaxy cluster at z=0.28. We show a) a preliminary isophotal R light curve of the galaxy, and b) contour plots of the galaxy in the reference and discovery images. The light curve is consistent with a type Ia supernova explosion in a z=0.28 galaxy, but it does not allow a positive identification. We are currently collecting images of this galaxy to see if it returns to its original luminosity. The contour plots suggest that the increase in brightness is displaced from the center of the galaxy.

7. References

Arnett, W. D., Branch, D., and Wheeler, J. C. (1985). Hubble's constant and exploding carbon-oxygen white dwarf models for Type I supernovae. *Nature*, **314**, 337-338.

Bahcall, N. A. (1988). Large-scale structure in the universe indicated by galaxy clusters. *Ann. Rev. Astron. Astrophys.*, **26**, 631–686.

Bahcall, S. R., and Tremaine, S. (1988). Evolutionary corrections to the redshift-volume measurement of the density parameter. *Ap. J. Lett.*, **326**, L1-L4.

Barbon, R. (1978). Supernovae in the Coma cluster of galaxies. *A. J.*, 83(1), 13-19.

Barbon, R., Benetti, S., Cappellaro, E., Rosino, L., and Turatto, M. (1990). Type Ia supernova 1989B in NGC 3627. *Astron. Astrophys.*, **237**, 79-90.

Barbon, R., Cappellaro, E., Turatto, M. (1989). The Asiago supernova catalog. *Astron. Astrophys. Suppl. Ser.*, **81**, 421-443.

Branch, D. (1986). On the relative frequencies of the kinds of type I supernovae. *Ap. J. Lett.*, **300**, L51-L54.

Bruzual, G., and Spinrad, H. (1978). The characteristic size of clusters of galaxies: a metric rod used for a determination of q_o. *Ap. J.*, **220**, 1-7.

Butcher, H., Oemler, A. (1984). The evolution of galaxies in clusters V - A study of populations since $z \sim 0.5$. *Ap. J.*, **285**, 426-438.

Cadonau, R., Sandage, A., and Tammann, G. A. (1985). Type I supernovae as standard candles. In *Supernovae as Distance Indicators*, ed. N. Bartel, pp. 151-165. Berlin, Heidelberg, New York: Springer–Verlag.

Cappellaro, E., and Turatto, M. (1988). A new determination of the frequency of supernovae. *Aston. Astrophys.*, **190**, 10-16.

Couch, W. J., Hansen, L., Jørgensen, H. E., Nørgaard-Nielsen, H. U., Ellis, R. S., and Aragon, A. (1989). The Danish-Anglo-Australian Distant Supernova Search. In

Particle astrophysics: Forefront experimental issues, ed. E. B. Norman, pp. 192-195. Singapore: World Scientific.

Elias, J. H., Frogel, J. A., Hackwell, J. A., and Persson, S. E. (1981). Infrared light curves of type I supernovae. *Ap. J. Lett.*, **251**, L13-L16.

Elias, J. H., Matthews, K., Neugebauer, G., and Persson, S. E. (1985). Type I supernovae in the infrared and their use as distance indicators. *Ap. J.*, **296**, 379-389.

Evans, R., van den Bergh, S., and Mc Clure, R. D. (1989). Revised supernova rates in Shapley-Ames galaxies. *Ap. J.*, **345**, 752-758.

Faber, S. M., Gallagher, J. S. (1979). Masses and mass-to-light ratios of galaxies. *Ann. Rev. Astron. Astrophys.*, **17**, 135-187.

Filippenko, A. V. (1988). Supernova 1987K: type II in youth, type Ib in old age. *A. J.*, **96**, 1941-1948.

Filippenko, A. V. (1991). The optical diversity of supernovae. In *Supernova 1987A and other supernovae*, ed. I. J. Danziger and K. Kjär, pp. 343-362. Garching: ESO.

Filippenko, A. V., Richmond, M. W., Branch, D., Gaskell, C. M., Herbst, W., Ford, C. H., Treffers, R. R., Matheson, T., Ho, L. C., Dey, A., Sargent, L. W., Small, T. A., van Breugel, J. M. (1992a). The subluminous, spectroscopically peculiar type Ia supernova 1991bg in the elliptical galaxy NGC 4374. *A. J.*, in press.

Filippenko, A. V., Richmond, M. W., Matheson, T., Shields, J. C., Burbidge, E. M., Cohen, R. D., Dickinson, M., Malkan, M. A., Nelson, B., Pietz, J., Schlegel, D., Schmeer, P., Spinrad, H., Steidel, C. C., Tran, H. D., Wren, W. (1992b). The peculiar type Ia SN 1991T: Detonation of a white dwarf? *Ap. J.*, **384**, L15-L18.

Filippenko, A. V. (1988). Supernova 1987K: type II in youth, type Ib in old age. *A. J.*, **96**, 1941-1948.

Graham, J. A. (1982). UBVRI Standard stars in the E-regions. *Publ. A. S. P.*, **94**, 244-265.

Harkness, R. P., and Wheeler, J. C. (1990). Classification of supernovae. In *Supernovae*, ed. A. G. Petschek, pp. 1-29. New York: Springer-Verlag.

Hellwege, K.-H., ed. (1982). Landolt-Börnstein, *Numerical Data and Functional Relationships in Science and Technology*, volume VI/2b, p. 73. Berlin, Heidelberg, New York: Springer-Verlag.

Hickson, P. (1977). The angular-size-redshift relation. II. A test for the deceleration parameter. *Ap. J.*, **217**, 964-975.

Hubble, E.(1929). A relation between distance and radial velocity among extra-galactic nebulae. *Proc. Nat. Acad. Sci.*, **15**, 168.

Kirshner, R. P., Oemler, A., Schechter, P. L., and Shectman, S. A. (1983). A deep survey of galaxies. *A. J.*, **88**, 1285–1300.

Kowal, C. T. (1968). Absolute magnitudes of supernovae. *A. J.*, **73**, 1021-1024.

Leibundgut, B. (1990). Supernova Studies II - The effect of redshift on supernovae Ia light curves. *Astron. Astrophys.*, **229**, 1-6.

Leibundgut, B. (1991). Supernovae Ia as standard candles. In *Supernovae: The tenth Santa Cruz summer workshop in astronomy and astrophysics*, ed. S. E. Woosley, pp. 751-759. New York: Springer-Verlag.

Leibundgut, B., and Tammann, G. A., (1990). Supernova Studies III - The calibration of the absolute magnitude of supernovae of type Ia. *Astron. Astrophys.*, **230**, 81-86.

Loh, E. D., and Spillar, E. J. (1986). A measurement of the mass density of the universe. *Ap. J. Lett.*, **307**, L1-L4.

Miller, D. L., and Branch, D. (1990). Supernova absolute magnitude distributions. *A. J.*, **100**, 530-539.

Minkowski, R. (1964). Supernovae and Supernova Remnants. *Ann. Rev. Astron. Astrophys.*, **2**, 247-266.

Mnatsakanyan, R. G., Oskanyan, A. V., and Lovash, M. (1979). Translation: On the frequency of supernovae. *Astrofizika*, **15**, 413-422.

Muller, R. A., Newberg, H. J. M., Pennypacker, C. R., Perlmutter, S., Sasseen, T. P., and Smith, C. K. (1992). High rate for type Ic supernovae. *Ap. J. Lett.*, **384**, L9-L13.

Nomoto, K., Yamaoka, H., Shigeyama, T., Kumagai, S., and Tsujimoto, T. (1991). In *Lecture series given at the Les Houches Summer School on SUPERNOVAE*, ed. J. Andouze, S. Bludman, R. Mochkovitch, and J. Zinn-Justin, to be published by Elsevier Science Publishers B. V.

Nørgaard-Nielsen, H. U., Hansen, L., Jørgensen, H. E., Salamanca, A. A., Ellis, R. S., and Couch, W. J. (1989). The discovery of a type Ia supernova at a redshift of 0.31. *Nature*, **339**, 523-525.

Perlmutter, S., Crawford, F. S., Muller, R. A., Pennypacker, C. R., Sasseen, T. P., Smith, C. K., Treffers, R., and Williams, R. (1988).The status of Berkeley's real-time supernova search. In *Instrumentation for Ground-Based Optical Astronomy*, ed. L. B. Robinson (New York: Springer), 674–680.

Phillips, M. M., Wells, L. A., Suntzeff, N. B., Hamuy, M., Leibundgut, B., Kirshner, R. P., Foltz, C. B. (1992). SN 1991T: Further evidence of the heterogeneous nature of type Ia supernovae. *A. J.*, **103(5)**, 1632-1637.

Politzer, H. D., and Preskill, J. P. (1986). Poisson clusters and Poisson voids. *Phys. Rev. Lett.*, **56**, 99-102.

Press, W. H., Flannery, B. P., Teukolsky, S. A., and Vetterling, W. T. (1986). *Numerical Recipes*. Cambridge: Cambridge University Press.

Richardson, H., and Morbey, C. L. (1988). Optical Design of an f/1 focal reducer located before the prime focus of large telescopes. In *ESO Conference on very large telescopes and their instrumentation*, ed. Ulrich, M. -H., pp. 1149–1156.

Rowan-Robinson, M. (1984). *The cosmological distance ladder*. New York: W. H. Freeman and Company. p. 340.

Rubin, V. C., and Graham, J. A. (1987). Evidence for a massive component in late-type spirals. *Ap. J. Lett.*, **316**, L67-L69.

Sandage, A. (1975). The redshift. In *Stars and stellar systems, Galaxies and the universe vol. 9*, ed. A. Sandage, M. Sandage, J. Kristian, pp. 761-785. Chicago: University of Chicago Press.

Schechter, P. (1976). An analytic expression for the luminosity function for galaxies. *Ap. J.*, **203**, 297–306.

Spinrad, H., and Djorgovski, S. (1987). The status of the Hubble diagram in 1986. In *Observational Cosmology* (IAU Symposium 124), ed. A. Hewitt, G. Burbidge, and Li-Zhi Fang, pp. 129-141. Dordrecht: Reidel.

Strauss, M. A., Davis, M., Yahil, A., and Huchra, J. P. (1992). A redshift survey of IRAS galaxies IV - The galaxy distribution and the inferred density field. *Ap. J.*, **385**, 421-444.

Sutherland, P. G., and Wheeler, J. C. (1984). Models for type I supernovae: partially incinerated white dwarfs. *Ap. J.*, **280**, 282-297.

Tammann, G. A. (1982). Supernova statistics and related problems. In *Supernovae: A Survey of Current Research*, ed. Rees, M. J., and Stoneham, R. J., pp. 371-403. Dordrecht, Boston: Reidel.

Tsvetkov, D. Y. (1983). Supernova frequency estimates. *Sov. Astron.*, **27(1)**, 22-24.

Turnrose, B. E. (1974). Absolute spectral energy distribution of the night sky at Palomar and Mount Wilson observatories. *Publ. A. S. P.*, **86**, 545-551.

van den Bergh, S., and Tammann, G. A. (1991). Galactic and extragalactic supernova rates. *Ann. Rev. Astron. Astrophys.*, **29**, 363-407.

Weinberg, S. (1972). *Gravitation and Cosmology*. New York: John Wiley & Sons.

Woosley, S. E. (1990). Type I supernovae: carbon deflagration and detonation. In *Supernovae*, ed. A. G. Petschek, pp. 182–212. Berlin: Springer Verlag.

Woosley, S. E., and Weaver, T. A. (1986). The physics of supernova explosions. *Ann. Rev. Astron. Astrophys.*, **24**, 205-253.

Zombeck, M. V. (1990). *Handbook of space astronomy and astrophysics*. Cambridge: Cambridge University Press.

LAWRENCE BERKELEY LABORATORY
UNIVERSITY OF CALIFORNIA
TECHNICAL INFORMATION DEPARTMENT
BERKELEY, CALIFORNIA 94720