

## Chapter 5

# Discovery of 3 New Binary Millisecond Pulsars in the Globular Cluster M62

### Abstract

We report the discovery of three new binary millisecond pulsars in the globular cluster M62 = NGC 6266. These pulsars are the first new objects discovered with the 100 m Green Bank Telescope. Two of the three pulsars were found using a new search method sensitive to binaries whose orbital periods are of the same order as the observation time. With six pulsars, M62 is now the third ranking globular cluster in terms of total known pulsar population. Timing of these new pulsars should provide important information about the host cluster, including useful constraints on its central density and initial mass function.

(Chandler, A. M., Anderson, S. B., Backer, D. B., Jacoby, B. A.,  
Kulkarni, S. R., and Prince, T. A., to be submitted to  
the *Astrophysical Journal*)

## 5.1 Introduction

The discovery of the first millisecond pulsar (MSP) (Backer et al., 1982) quickly led to the suggestion (Alpar et al., 1982) that these objects are formed in low-mass X-ray binaries (LMXBs) when a neutron star gains angular momentum in the process of accreting material from a binary companion. Since the specific incidence of LMXBs in globular clusters is significantly higher than in the Galactic disk, clusters were recognized early on as promising targets for MSP searches. The first cluster MSP discovery came several years later in M28 (Lyne et al., 1987). Over the last 15 years, the known population has grown to at least 67 pulsars in 21 clusters (Kulkarni & Anderson, 1996; Camilo et al., 2000; Possenti et al., 2001; Lyne et al., 2000; Ransom, 2001; S. Ransom, private communication; this paper; Appendix A), now comprising roughly half of the known MSPs.

It is more than just their sheer numbers that make cluster pulsars interesting. Exchange interactions in the dense cluster cores can lead to the formation of exotic pulsar systems such as triples (Backer et al., 1993; Freire et al., 2001), extremely tight binaries (e.g., Camilo et al., 2000; Ransom et al., 2001; D’Amico et al., 2001), double neutron star systems (Prince et al., 1991), and perhaps even neutron star-black hole binaries (not yet discovered). When multiple pulsars are found in a single cluster, they can help elucidate the internal dynamics of the cluster as a whole. Precision timing of cluster pulsars can reveal the effects of acceleration in the gravitational potential of the cluster, and these measurements can be used to place useful limits on the cluster’s central density and mass-to-light ratio. Additionally, pulsar positions can be used to make statistical neutron star mass measurements, and constrain the initial mass function of the cluster stars.

Detecting these pulsars is not easy. Pulsars in globular clusters are typically several times more distant than field pulsars, and their flux densities are therefore roughly an order of magnitude lower. To recoup this lost signal strength, integration times must increase by two orders of magnitude, while maintaining a fast sample rate. Large fast Fourier transforms (FFTs) are therefore required (perhaps as large as  $10^9$

points). To correct for a (usually) unknown degree of interstellar dispersion, and to maintain sensitivity to Doppler-shifted binary pulsars, many hundreds of these large FFTs must generally be calculated in a full search. The potential payoffs usually justify such a computational effort, and these searches have become more tractable in recent years.

All MSP searches benefit from advances in computing power, search algorithms, and data recording speeds. But improvements in telescopes and observing hardware can also drive new discoveries. Installation of the new low noise, wide-band multibeam receivers and high-resolution filterbanks at the Parkes Observatory in Australia has led to the discovery of at least 21 cluster pulsars in the last few years (Camilo et al., 2000; Possenti et al., 2001). Similarly, recent Arecibo upgrades have greatly improved pulsar search sensitivities, particularly at high frequencies. At least 3 new pulsars have been detected in L-band globular cluster searches from Arecibo (S. Ransom, private communication). The recently completed Robert C. Byrd Green Bank Telescope (GBT), however, has the potential to bring about a much more significant increase in the known population of globular cluster pulsars. With a 100-meter reflecting surface, the GBT is the largest fully-steerable telescope in the world, able to see declinations from  $+90^\circ$  down to  $\sim -45^\circ$ . This covers 85% of the celestial sphere, including enough of the southern hemisphere to cover the Galactic center region, encompassing more than 80% of the cataloged globular clusters in the Galaxy (Harris, 1996). Our group recently installed the Berkeley-Caltech Pulsar Machine (BCPM) at the GBT for community use. This flexible digital filterbank is the GBT's first pulsar back end, and has allowed several groups to begin searching for new cluster pulsars under the GBT's Early Science program.

In this paper we report the discovery of 3 new binary MSPs in the globular cluster NGC 6266 = M62. These are the first new objects discovered using the GBT. A preliminary announcement of these results was made by Jacoby et al. (2002). Two of these new pulsars were discovered using a novel technique sensitive to pulsars whose orbital periods are approximately equal to the observation length, traditionally a poorly sampled regime of orbital periods. With the three previously known pulsars

in M62 (D'Amico et al., 2001; Possenti et al., 2001), there are now six known MSPs in this cluster, making it a prime candidate for dynamical studies. M62 is now the third ranking globular cluster in terms of the number of known pulsars; only 47 Tuc and M15 have more, with 20 and 8 respectively (Camilo et al., 2000; Anderson, 1993).

In section 5.2 below, we describe our search observations and analysis methods. We also describe our new binary pulsar search technique in great detail. In section 5.3, we present the results of our search, including the parameters of the newly discovered pulsars in M62. We discuss the importance of the new discoveries and future prospects for M62 and globular cluster searches in general.

## 5.2 Observations and Analysis

### 5.2.1 Data Collection

Installation of the Berkeley-Caltech Pulsar Machine (BCPM) at Green Bank was completed in early August, 2001. The BCPM is a copy of the 96-channel, 170-MHz bandwidth, analog/digital filter bank that was designed and built by a team from Berkeley, NRL, and Caltech with previous installations at Nançay, Effelsberg and Arecibo. A partial technical description is contained in Backer et al. (1997).

The BCPM accepts dual polarization IFs in the 300-500 MHz range. These are divided into 6 analog channels using Mixer/Filter/LO units with independent gain and leveling control. The signals are then sampled and divided digitally 16 ways, for a total of 96 channels in each polarization. The maximum bandwidth per channel is 1.8 MHz, though smaller bandwidths are selectable. The dual polarization data are then brought together, decimated down to the desired sample rate, and (optionally) summed. The mean is then removed and the remainder is quantized into 4 bits. The 4-bit data are formatted and passed to the host workstation (a Sun Ultra-1) over an EDT interface. Data are then passed to disk, and from disk to tape. Monitoring, control, and data acquisition in the host workstation use a version of the Penn State Pulsar Machine (Cadwell, 1997) software, which has been integrated into the GBT

observing system.

Our original search data were recorded on August 16, 2001, with confirmation data taken December 4-9, 2001. For our search and follow up observations, we used the BCPM fed by the L-band receiver, with the central frequency set to 1400 MHz. To limit the effects of dispersion, we used 0.7 MHz channels, for a total bandwidth of 67.2 MHz. We selected a sample rate of  $100 \mu\text{sec}$ , with the two polarizations summed in hardware. Thus, our 4-hour search observation included just over  $2^{27}$  time samples. Much of our analysis was actually carried out with the data resampled at  $200 \mu\text{sec}$  resolution, since the rather large dispersion measure of the M62 pulsars ( $\text{DM} = 114.4 \text{ cm}^{-3} \text{ pc}$ ) causes a single-channel dispersive delay of  $240 \mu\text{sec}$ . The half power beam radius of the telescope (at 1400 MHz) is approximately  $4.3'$  (HWHM), which is nearly 24 times the cluster's core radius, and 3.5 times the half-mass radius. This beam should easily cover all dynamically relaxed pulsar systems in the cluster, as well as some that have been ejected from the core, but remain bound to the cluster (like, for example, PSR B2127+11C, located approximately 13.5 core radii from the center of M15; Anderson, 1993).

### 5.2.2 Search for Isolated Pulsars

We began with a standard search for isolated pulsars (Burns & Clark, 1969; Lyne & Graham-Smith, 1998). Of course, MSPs are almost certainly formed in binary systems, but solitary MSPs are not uncommon. Their existence can be explained by the obliteration of eclipsing companions, tidal disruption of the companion, or ionization of the binary system by near collisions (particularly for cluster pulsars) (van den Heuvel & van Paradijs, 1988; Biggs et al., 1994).

The first step in the data analysis was to shift and add the filterbank channels to create a single dedispersed time series. Since there were already known pulsars in M62 (D'Amico et al., 2001), we were able to use the cluster's previously determined dispersion measure (the column density of free electrons along the line of sight to the cluster)  $\text{DM} = 114.4 \text{ cm}^{-3} \text{ pc}$ . This saved us considerable effort, compared to a truly

blind search involving many trial DMs.

We calculated the FFT of this time series to generate an estimate of the normalized power spectrum. Frequency bins containing statistically significant power were saved as possible pulsar candidates. To take advantage of the short duty cycles typical of radio pulses, incoherent power sums of up to 16 harmonics were also calculated and searched for significant candidates.

The top candidates were then analyzed by human. Visualization techniques included time-domain and frequency-domain analyses. For each candidate, we folded a single dedispersed pulse profile, as well as a two-dimensional folded filterbank profile (to verify that the candidate was broad-band and exhibited the expected amount of dispersion). From the dedispersed time series, we also calculated a two-dimensional time-resolved pulse profile (D’Amico et al., 2001) and a time-resolved power spectrum, in which the data are divided into a number of shorter segments, and each is folded or Fourier transformed individually. These time-resolved techniques are useful for spotting binary motion, though this simple FFT search is expected to be sensitive only to rather long-period binaries. This is discussed in much greater detail in the next section. We note that one of our newly discovered pulsars was found using this straightforward FFT method, despite the fact that it is in a  $\sim 1$  day binary.

Figure 5.1 shows the sensitivity of this search to isolated pulsars. Following standard conventions (e.g., Dewey et al., 1985; Bhattacharya, 1998; sec. 4.2.2), the sensitivity is approximated as

$$S_{\min} = \alpha\beta \frac{T_{\text{sys}}}{G(N_p B \tau_{\text{int}})^{1/2}} \left( \frac{w_{\text{eff}}}{P - w_{\text{eff}}} \right)^{1/2}. \quad (5.1)$$

$S_{\min}$  is the minimum detectable flux density. The factors on the right hand side of Equation 5.1 are characteristic of a particular observing system and search method. The parameter  $\alpha$  is the threshold signal-to-noise ratio (8.0 was used in Fig. 5.1);  $\beta$  is a factor which accounts for various losses in the observing system ( $\beta \sim 1.5$ );  $T_{\text{sys}}$  is the system temperature, the sum of the receiver temperature and the sky temperature (36 K used in Fig. 5.1);  $G$  is the gain of the telescope (1.85 K Jy $^{-1}$ );  $N_p$  is the number

of polarizations summed (2);  $B$  is the bandwidth of the observation (67.2 MHz);  $\tau_{\text{int}}$  is the integration time (13421.8 s);  $P$  is the rotation period of the pulsar; and finally  $w_{\text{eff}}$  is the effective pulse width.

The effective width accounts for the broadening of the intrinsic pulse width  $w_0$  due to the finite sampling time ( $\tau_{\text{samp}} \approx 200 \mu\text{sec}$ ), and the effects of interstellar dispersion and scattering. The dispersion contribution is due to the dispersive delay within a single filterbank channel

$$\tau_{\text{DM}} = 8.3 \times 10^6 b_{\text{MHz}} \nu_{\text{MHz}}^{-3} \text{DM} = 0.24 \text{ ms}, \quad (5.2)$$

where  $b_{\text{MHz}}$  is the channel bandwidth,  $\nu_{\text{MHz}}$  is the central observing frequency, both expressed in MHz, and DM is the dispersion measure in  $\text{cm}^{-3} \text{pc}$ . The dispersive broadening dominates the effect of scattering

$$\tau_{\text{scat}} \approx \left( \frac{1000}{\nu_{\text{MHz}}} \right)^{4.4} 10^{-7.231+0.9255 \log_{10} \text{DM}+0.814613(\log_{10} \text{DM})^2} = 0.003 \text{ ms} \quad (5.3)$$

(Cordes et al., 1991). The effective width is approximated by summing these contributions in quadrature

$$w_{\text{eff}}^2 = w_0^2 + \tau_{\text{samp}}^2 + \tau_{\text{DM}}^2 + \tau_{\text{scat}}^2. \quad (5.4)$$

For Figure 5.1, we assumed an intrinsic pulse width equal to 10% of the period.

The dashed line in Figure 5.1 shows the approximate sensitivity of the previous Parkes search of M62 (D’Amico et al., 2001). Due primarily to the much larger reflecting area of the GBT, our search sensitivity is significantly better ( $\sim 50\%$  lower flux limit at all periods). We note, however, that our sensitivity is severely limited by the number of frequency channels available in the BCPM. Once the GBT “spigot” comes on-line, an additional factor of 2 improvement in sensitivity should be realized.

### 5.2.3 Binary Pulsar Search

More than half of the known pulsars in globular clusters are in binary (or triple) systems. When a pulsar moves in an orbit with one or more companions, the observed,




Figure 5.1: M62 search sensitivity for isolated pulsars. Solid line shows the approximate sensitivity of the current search; dashed line shows that of the previous search from Parkes. The periods and approximate flux densities of the six (binary) pulsars in M62 are represented by the letters A through F on the plot.

Doppler-shifted pulse frequency will change with time. In a simple FFT search, the acceleration of the pulsar causes its power to spread out over a range of spectral bins, making its signature difficult (if not impossible) to detect. A number of techniques have been developed to mitigate this effect, with various advantages, disadvantages, and realms of applicability. These binary search techniques increase the computational complexity of a search, but improvements in computer power have made their use much more practicable. The ratio of binary to solitary detection rates in recent years is nearly 3 to 1, suggesting that selection effects have artificially deflated the cluster pulsar binary fraction, which may be much greater than 50%.

Clearly, sensitivity to binary pulsars in a globular cluster search is vital. In this section, we begin by briefly discussing the existing binary search methods that we utilized. These include “acceleration” techniques which have proven their mettle in previous surveys, as well as the sideband or “cepstrum” method (i.e., spectrum of spectrum method) which shows great promise, but has not as yet yielded new pulsars. We did not discover any new pulsars using these methods in the current search. Next, we present a new, powerful binary search technique that we have developed to complement the other approaches. We describe how the new method works and show that in its intended regime of applicability, it compares quite favorably to the other binary search methods, both in terms of sensitivity and computational effort. Perhaps the most important proof of its efficacy is the fact that two of our newly discovered M62 pulsars were found using this new method.

To determine the most appropriate binary search technique for a given system, we must compare the orbital period,  $P_b$ , to the observation time,  $T_{\text{obs}}$ . The following discussion is broken down into the three regimes  $P_b \gg T_{\text{obs}}$ ,  $P_b \ll T_{\text{obs}}$ , and  $P_b \sim T_{\text{obs}}$ .

$$P_b \gg T_{\text{obs}}$$

If a pulsar spinning with rest frequency  $f_0$  moves in a binary system, the observed, Doppler-shifted frequency is given by

$$f'(t) = f_0[1 + v(t)/c], \quad (5.5)$$

where  $c$  is the speed of light and  $v$  is the radial velocity of the pulsar (the velocity projected along the line of sight). When an observation covers only a very small part of an orbit, the time dependence of the radial velocity of the pulsar can be well approximated by a constant acceleration  $\Delta v(t) \approx at$ . This is equivalent to keeping only the lowest order derivative in a Taylor expansion of the frequency evolution, i.e., tantamount to assuming a constant frequency derivative  $\dot{f}$ .

In a blind search, one corrects the time series for an assumed acceleration before calculating the FFT, harmonic folds, etc., repeating the entire process for each acceleration trial. A full search typically involves  $\sim 10^2 - 10^3$  iterations. Acceleration searches of this sort have proven to be quite successful (e.g., Anderson, 1993; Camilo et al., 2000). The acceleration correction can also be carried out in the frequency domain, using essentially a matched filtering approach (Middleditch et al., 1993; Ransom, 2001). This technique is computationally similar to the time domain method, usually involving roughly the same number of operations to carry out. But it may be more efficient than the time-domain method (for a reasonable degree of signal spreading) when the full FFT is distributed over multiple processors, since only one large FFT is required.

Acceleration searches are most effective for pulsars in very wide orbits. For such long-period systems (depending to a lesser extent on orbital phase), the acceleration method is capable of gathering nearly all of the original signal power back into a single bin of the power spectrum (subject to the sinc function response of the discrete Fourier transform). As the orbital period approaches the observation time, the minimum detectable pulsar flux increases, since only a small fraction of the observation is well fit by the assumed linear frequency dependence. We will return to this issue in section 5.2.4 below.

$$P_b \ll T_{\text{obs}}$$

When an observation contains many orbits, the phase modulation of the pulsar signal results in a family of sidebands around the rest frequency of the pulsar in the Fourier spectrum (Ransom, 2001). The number of sidebands depends on the orbital semima-

major axis and the spin frequency,  $N_{\text{sb}} \approx 4\pi x f_0$  ( $x = a \sin i/c$  is the light travel time of the projected semimajor axis). The known binary pulsars with  $P_b < 10$  d have spin periods in the range [1.607, 1004] ms and orbital semimajor axes in the range [0.010, 20.0] ls. For these systems, the number of sidebands  $N_{\text{sb}}$  covers the range  $\sim [4, 13000]$ . For all but the brightest pulsars, the spectral power distributed over these sidebands becomes buried in noise.

Since these side bands are evenly spaced at the orbital frequency, a convenient method of incoherently summing the power back up is to step through the power spectrum, calculating small FFTs. In general the number of sidebands and their location in the spectrum are unknown, so overlapping FFTs of several sizes must be tried. The family of sidebands acts as a pulse train which, when transformed back to the time domain, can produce detectable power at the orbital period (and at “harmonic” multiples of the period). Once a detection has been made, the orbital parameters can be estimated from the original complex Fourier spectrum.

This is the best method currently available for detecting ultra-short-period binary pulsars. While this method has successfully redetected previously known pulsars, its application has not yet resulted in any *new* detections, but it has only been in use for a short while. The shortest known radio pulsar orbital period currently stands at 96 minutes (Camilo et al., 2000), close to the limits of what can be detected using traditional acceleration searches. Two of the three currently known accreting MSPs are in  $\sim 40$  minute orbits (Markwardt et al., 2002; Galloway et al., 2002), and neutron star binaries with orbital periods as short as 11 minutes have been observed (Stella et al., 1987). The modulation sideband technique finally allows pulsation searches with meaningful sensitivity to such systems.

The sideband search method is most effective when an observation encompasses more than a few orbits. Strictly speaking, the method can work as long as at least one orbit is covered by the data, but the sensitivity degrades rapidly as  $T_{\text{obs}}$  approaches  $P_b$  (when  $T_{\text{obs}} = P_b$ , the spacing of the sidebands is equal to the Fourier step size of the spectrum).

$$P_b \sim T_{\text{obs}}$$

Acceleration and sideband searches cover complementary regimes of  $P_b/T_{\text{obs}}$ , but they leave a sizeable gap around  $P_b/T_{\text{obs}} \sim 1$ . Previous searches have therefore had very little sensitivity to pulsars in this region. Traditionally, the best way to handle such systems would be to divide an observation into shorter segments and perform acceleration searches on each of them separately. This effectively eliminates the usual  $S_{\text{min}} \sim T_{\text{obs}}^{-1/2}$  advantage of accumulating longer data sets. Alternatively, if the observation time can be increased, then the sideband search method can be applied. However, increasing the observation time significantly (so that  $T_{\text{obs}} \gtrsim 2P_b$ ) is not always possible. At Arecibo, for instance, transit times are limited to  $\lesssim 3$  hr, while at other sites, sources are usually up for only 8 – 12 hr. Of course, if a source is circumpolar or if the observatory is in space or (one day) on the moon, then longer observations may be a viable solution. In the next section, we present a third possibility. We describe a new pulsar search technique that we have developed to fill the  $P_b/T_{\text{obs}}$  gap left by the other two binary search methods.

### 5.2.4 The Dynamic Power Spectrum Method

Our new search method began as a candidate visualization technique, already alluded to in section 5.2.2 above. It involves the use of a dynamic power spectrum (DPS). The general idea is to divide an observation into a number of shorter segments, and calculate a power spectrum for each (optionally, harmonic folded). The power spectra are then stacked on top of one another in a two-dimensional array. In this spin frequency vs. time plane, a bright pulsar is easily visible as a contiguous curve of excess power tracing out  $f(t)$ . Our search method relies on the frequency-local nature of these pulsar signatures. Even if none of the segments contain globally significant power, we can detect pulsars by finding a pattern of *locally* significant powers, i.e., barely excess powers in several different time segments within a small range of spin frequencies. In principle, this method is similar to the segmented acceleration search mentioned in the last paragraph, but with a formalized way of lowering the individual

detection thresholds using a built-in coincidence requirement.

The concept of a dynamic power spectrum is not new — applications of DPS techniques have ranged from speech and musical analysis to the study of burst oscillations in LMXBs (Bracewell, 2000; van Straaten et al., 2001). Rudimentary DPS methods have been used in pulsar searches (in M15 by members of our group and in Terzan 5 by Lyne et al., 2000) without success. These earlier applications were very limited in scope, and involved analyzing spectra by eye. In this section, we present a comprehensive method to automate the search process and optimize searches for a wide range of pulsar parameters. The DPS method described here requires far less computation than acceleration searches, and nicely fills in the  $P_b/T_{\text{obs}}$  gap.

We begin by considering the ideal number of segments  $\mathcal{S}$  into which an observation should be divided. This number should be large enough to keep the signal power from spreading out over too many spectral bins, but should also be as small as possible to avoid excessive reduction in the single-segment signal-to-noise ratio ( $\text{SNR} \sim \mathcal{S}^{-1/2}$ ). The independent Fourier step size of the DPS will be  $F = \mathcal{S}/T_{\text{obs}}$ . We will require that the fundamental frequency drift of a pulsar, over the course of a single segment, be less than  $F$ .

For a pulsar in a circular orbit, the projected radial velocity is given by

$$v(t) = -v_0 \sin i \cos \left( \frac{2\pi}{P_b} t + \phi_0 \right), \quad (5.6)$$

where  $v_0$  is the orbital velocity of the pulsar,  $i$  is the orbital inclination (the angle between the plane of the orbit and the plane of the sky), and  $\phi_0$  is the orbital phase at  $t = 0$ . Here we have defined  $\phi$  to be zero at the descending node, where the pulsar is maximally moving away from the earth (i.e., where the pulse frequency is observed to be at its minimum). We can substitute this expression into Equation 5.5 and take a derivative with respect to time to obtain

$$\dot{f}(t) = \frac{2\pi f_0 \beta}{P_b} \sin \left( \frac{2\pi}{P_b} t + \phi_0 \right), \quad (5.7)$$

where we have used  $\beta = v_0 \sin i/c$ . In the worst-case scenario, we have  $|\dot{f}_{\max}| = 2\pi f_0 \beta / P_b$ . We require that the frequency range visited by this pulsar over the course of a single segment be less than the Fourier step size, i.e.,  $\Delta f_{\max} = |\dot{f}_{\max}|(T_{\text{obs}}/\mathcal{S}) < \mathcal{S}/T_{\text{obs}}$ . Solving for the minimum  $\mathcal{S}$ , we obtain

$$\mathcal{S} = \left( \frac{2\pi f_0 \beta T_{\text{obs}}^2}{P_b} \right)^{1/2}. \quad (5.8)$$

For typical pulsar and search parameters, the ideal number of segments ranges from one to several hundred. For certain systems, if  $T_{\text{obs}}/P_b \gtrsim 5$ , the number of segments may become unwieldy, but in such cases, the sideband search method is more appropriate.

In a blind search, the pulsar parameters are not known *a priori*, so several values of  $\mathcal{S}$  must be attempted, with each resulting DPS searched independently. A fairly exhaustive pulsar search might employ 6 orders,  $\mathcal{S} = 4 - 128$  (in powers of two). For non-overlapping segments,  $\mathcal{S}$  is restricted to powers of two so that inefficient non-power-of-two FFTs need not be used. Note that if the segment spectra are harmonically folded (generally a good idea), then the frequency drift of the higher harmonics will be proportionately larger than the drift of the fundamental. In such cases optimal detection may occur with a larger number of segments than Equation 5.8 would indicate (although SNR losses associated with using more segments may very well negate any benefit from increasing  $\mathcal{S}$ ).

Calculating a DPS involves roughly the same number of operations ( $N_{\text{op}}$ ) as calculating a single full-length FFT. In a full search we calculate a DPS for  $N_s$  different values of  $\mathcal{S}$ , for a total of roughly

$$N_{\text{op}} \sim \sum_{\mathcal{S}=\mathcal{S}_1}^{\mathcal{S}_{N_s}} \left[ 5N \log_2 \frac{N}{\mathcal{S}} + N_{\text{search}} \right] \approx 10N_s N \log_2 N, \quad (5.9)$$

where  $N$  is the total number of samples in the full time series and  $N_{\text{search}}$  is the number of operations required to detect a pulsar signal, once the DPS has been calculated. As we shall see below, coincidentally  $N_{\text{search}} \sim N \log_2 N$ . With  $N_s = 6$ , a full DPS

search requires about a factor of 10 – 100 fewer operations than a full acceleration search. We also note that the DPS segments can almost always be processed with in-core, single-processor FFTs.

We now consider the problem of actually detecting the presence of a weak pulsar, having calculated a DPS for a given observation. In many ways, the optimal way to detect the  $f(t)$  pattern of a pulsar is by eye. Since this would be rather tedious for a full DPS search, we have developed a hierarchical search algorithm which we describe below. The method we describe here works reasonably well in terms of false positive rate, false negative rate, and computational complexity, though other algorithms are certainly possible.

We begin with a first pass over the entire DPS, selecting the bins containing the top powers. For this first cut, the threshold is set fairly low, so that  $\sim 100$  “candidates” survive to the next stage of the analysis. We do not expect any of these first cut candidates to be statistically significant, when considering all of the DPS bins as independent trials (of which there might typically be  $\sim 10^8$ ). Essentially all of these hits are expected to be due to noise. With  $N/2$  points in the DPS, this first cut requires  $\sim N$  operations.

Given a first cut candidate in frequency bin  $b$  and segment  $s$ , we set about to look for locally significant powers in the other segments, in the vicinity of  $b$ . If a pulsar has a maximum line of sight acceleration magnitude  $a_{\max} = c|\dot{f}_{\max}|/f_0$ , then the largest possible change in spin frequency from one segment to the next is

$$\Delta f \leq |\dot{f}_{\max}| \Delta t = f_0 \frac{a_{\max}}{c} \frac{T_{\text{obs}}}{\mathcal{S}}. \quad (5.10)$$

Therefore, in segment  $s + 1$  or  $s - 1$ , we should look within a range  $b \pm \Delta b$ , where

$$\Delta b = b \frac{a_{\max}}{c} \frac{T_{\text{obs}}}{\mathcal{S}}. \quad (5.11)$$

In segments  $s+2$  and  $s-2$ , we search  $b \pm 2\Delta b$ , etc. We thus search within an hourglass figure drawn through the first cut candidate point  $(b, s)$ , whose opening angle is

determined by  $a_{\max}$ . A candidate is assigned a statistical score (i.e., significance) based on the powers found and the number of bins searched. As we move up or down through the segments, if we encounter a (locally) significant power, we redraw the hourglass through the current point and continue outward until we reach the first or last segment. This “bootstrapping” technique effectively rewards candidates whose powers follow a contiguous path, by searching fewer points.

To maintain the best sensitivity to a variety of signals, we perform the hourglass search for a number of different  $a_{\max}$  values. The largest  $a_{\max}$  used should cover any signals detected using smaller  $a_{\max}$  trials, but the significance of a detection will be best when the hourglass searched is just large enough to encompass the pulsar’s path through the  $f - t$  plane. A typical search might involve  $a_{\max} = 5, 10, 25, 50, 75, 100, \& 200 \text{ m s}^{-2}$ , and higher values are easily accommodated. With  $\sim 100$  first cut candidates, each  $a_{\max}$  search requires  $\sim 100 \Delta b \mathcal{S}^2$  operations. For typical search parameters, the total number of operations required for all 7  $a_{\max}$  values listed above happens to be of the same order as  $N \log_2 N$ .

As with any pulsar search, the last stage of the process is human inspection of the top candidates.

### DPS Examples

Figure 5.2 shows a section of an example DPS. Here we see a 3 hr observation of a simulated 3.0 ms pulsar in a 2.8 hr orbit about a  $0.1M_{\odot}$  companion (viewed at  $i = 60^\circ$ ). The sinusoidal signal had a single-pulse signal to noise ratio of 0.01. This signal is easily detected, despite the fact that none of the powers are globally significant. The largest power in the feature is consistent with the expected tail of the noise distribution, given the total number of bins in the DPS (in this case,  $N = 2^{24}$ ).

According to Equation 5.8, the ideal number of segments for this pulsar is  $\mathcal{S} = 48.25$ . In this case,  $\mathcal{S} = 32$  yielded a more significant detection than  $\mathcal{S} = 64$ . For this pulsar,  $a_{\max}$  is  $17.95 \text{ m s}^{-2}$ . The dashed lines in Figure 5.2 indicate the initial search hourglass corresponding to  $a_{\max} = 25 \text{ m s}^{-2}$ , centered on the bin containing the largest power.




Figure 5.2: Example dynamic power spectrum. Section of the DPS of a simulated 3.0 ms pulsar in a 2.8 hr circular orbit with a  $0.1M_{\odot}$  companion is shown. The observation time was 3.0 hr. Dashed lines show the  $a_{\max} = 25 \text{ m s}^{-2}$  local search hourglass centered on the first cut hit in segment 13.

To this point our discussion has focused on circular orbits, but the method works for elliptical systems as well. Figure 5.3 shows a non-sinusoidal DPS detection of PSR B2127+11C, the relativistic binary in the globular cluster M15. Four harmonics have been summed in each segment’s spectrum (so neighboring bins are not independent), and because the pulsar is so bright, the DPS is shown in a log scale. This pulsar is readily detected with small  $a_{\max}$  trials, but if it were significantly less luminous, it would be best detected with a rather large  $a_{\max}$  ( $a_{\max} = 200 \text{ m s}^{-2}$  is overlaid on the figure). An acceleration search would have serious trouble detecting such a signal.

### Comparison with Other Binary Search Methods

We now consider a quantitative comparison between the DPS method, the acceleration method, and the sideband search method. In Figure 5.4, we plot the results of simulations showing the minimum detectable instantaneous (i.e., single-pulse) SNR as a function of observation time. All of these simulations were calculated with a 2.0 ms pulsar in a 1.0 hr circular orbit with a mass function  $f = 2.89 \times 10^{-4}$ , equivalent to a companion mass of  $0.1 M_{\odot}$  viewed at an inclination of  $i = 60^{\circ}$ . The points plotted are averaged over the starting orbital phase of the observation. This plot is similar to Figure 4.6 of Ransom (2001), comparing the sensitivities of only the acceleration and sideband search methods.

Also included for reference in the plot is a standard “8-sigma” theoretical sensitivity curve (dotted line). This shows the single-pulse, single-harmonic SNR at which a stationary (or uniformly moving) pulsar would be detected with a single-trial significance equivalent to an  $8\sigma$  Gaussian event. At the risk of being pedantic, we now explicitly derive the equation for this curve to aid the reader in understanding Figure 5.4. In a power spectrum normalized to unity, the probability that the power in a particular bin will exceed some threshold  $P$  is equal to  $\exp(-P)$ . We equate this to the probability of a Gaussian random variable exceeding  $8\sigma$ , and solve for  $P$ :

$$P = -\ln \left[ \int_8^{\infty} \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx \right] = 35.01. \quad (5.12)$$




Figure 5.3: Elliptical DPS example. PSR B2127+11C in M15 is seen passing through periastron. The data are from an 8000 s Arecibo observation from January 1999.




Figure 5.4: Comparison of binary search sensitivities. The curves show the minimum detectable instantaneous signal-to-noise ratio as a function of observation time, for a simulated 2.0 ms pulsar in a 1.0 hr circular orbit with a  $0.1M_{\odot}$  companion. The solid curve represents the DPS sensitivity, the dashed curve is for acceleration searches, and the dot-dashed curve is for the sideband search method. Plotted is the SNR of the fundamental (i.e., lowest harmonic) of our simulated Gaussian waveform (10% FWHM pulse). The simulated acceleration and DPS searches used 4-harmonic sums (see text). For reference, the dotted curve shows the theoretical 8-sigma single-harmonic coherent detection limit.

This is equivalent to an overall (i.e., integrated) SNR of  $\sqrt{35.01} = 5.917$ . For a given  $T_{\text{obs}}$ , this overall SNR is due to a sum of  $(T_{\text{obs}}/P_{\text{spin}})$  individual pulses, so the  $8\sigma$  single-pulse SNR is

$$\text{SNR}_{8\sigma} = 5.917 \left( \frac{P_{\text{spin}}}{T_{\text{obs}}} \right)^{1/2}. \quad (5.13)$$

Thus, longer observations are sensitive to weaker signals, with a coherent detection limit proportional to  $T_{\text{obs}}^{-1/2}$ .

Our simulations did not actually use a simple 8-sigma threshold. Rather, we plot the single-pulse fundamental SNR that is 95% likely to result in a detection with better than 1% significance (by significance, we mean the probability that a candidate was produced by random noise, accounting for all the statistical trials involved). Practically speaking, however, our threshold is within 15% of the 8-sigma curve over the entire range of  $T_{\text{obs}}$  shown in Figure 5.4.

Another important difference between our simulations and those of Ransom (2001) is that we used a more realistic radio pulsar waveform — Gaussian with a 10% full width at half maximum (FWHM), rather than sinusoidal. This benefits the acceleration searches and the DPS searches since these methods can easily take advantage of incoherent harmonic summing. This may be possible with the sideband search method, but would be extremely cumbersome, involving incoherently summing secondary spectra, ideally of different lengths. For a fair comparison between the three methods (i.e., since the sideband method used only a single harmonic), we plot the SNR of only the fundamental harmonic of our simulated pulsar signal. Because of harmonic summing, at the short- $T_{\text{obs}}$  end of the figure the acceleration and DPS sensitivities are below the theoretical *single-harmonic* threshold.

We have checked that our acceleration search simulation results are consistent with analytical predictions (Johnston & Kulkarni, 1991; Evans et al., in preparation). As has been pointed out previously (Johnston & Kulkarni, 1991; Ransom, 2001; Jouteux et al., 2002), acceleration searches are optimal when  $T_{\text{obs}} \approx P_b/7$ , a fact verified by our simulations. From Figure 5.4, we see that once  $T_{\text{obs}} \gtrsim P_b/6$ , the DPS method is more sensitive than the acceleration method (even with its optimal observation

length).

We also note that the acceleration method is most effective when  $\dot{f}$  is most nearly constant. This fact leads to two difficulties. First, many pulsars in tight binaries are eclipsed by their companions for part of their orbit. These eclipses occur at precisely the orbital phase where the acceleration method is most effective, greatly reducing the probability of detection (when the pulsar passes in front of its companion, it can still be detected). The DPS method on the other hand is most likely to hit at orbital phases where the pulse frequency is most nearly constant (i.e.,  $90^\circ$  away from eclipse). This is evident in Figures 5.2 and 5.3. Second,  $\dot{f}$  is most nearly constant when  $f$  is at its extrema. This means that in acceleration searches, pulsars are most likely to be detected when they are experiencing their maximum acceleration. Because acceleration searches are so computationally expensive, they generally cover only a very limited range of accelerations;  $|a_{\max}| = 30 \text{ m s}^{-2}$  is typical. When a pulsar's maximum acceleration exceeds the search cutoff, it is less likely to be detected. The simulated pulsar of Figure 5.4 has a maximum acceleration of  $70.8 \text{ m s}^{-2}$ , so our simulated acceleration search used  $|a_{\max}| = 75 \text{ m s}^{-2}$ . And indeed, the most significant detection of the pulsar almost always occurs near  $|a_{\text{trial}}| \approx 70 \text{ m s}^{-2}$ . The sensitivity of a standard acceleration search would be considerably worse. Finally, we note that large  $a_{\max}$  values can be searched with the DPS method at a far lower cost than moderate  $a_{\max}$  acceleration searches.

It is clear from Figure 5.4 that if long, continuous observations are possible, then the sideband search method has the best sensitivity. But to beat the DPS method, sideband searches require  $T_{\text{obs}} \gtrsim 5P_b$ . Thus, for globular cluster searches, the acceleration method is best suited to detecting binary systems with  $P_b \sim$  days, the DPS method should be the method of choice for pulsars with  $P_b \sim$  hours, while the sideband method is ideal for anything shorter. The DPS method is very similar in computational complexity to the sideband search method (although no large, parallel or out-of-core FFT is required), and is orders of magnitude more efficient than the acceleration method. With the addition of the DPS method, pulsar searches can now cover the complete range of  $T_{\text{obs}}/P_b$  with respectable sensitivity.

Table 5.1: Pulsars in M62<sup>a</sup>

Name	Period (ms)	$P_b$ (days)	$a \sin i$ (ls)	$M_C^{\min}$ ( $M_\odot$ )
J1701-3006A	5.241	3.80	3.48	0.19
J1701-3006B	3.593	0.14	0.25	0.12
J1701-3006C	3.806	0.21	0.19	0.07
J1701-3006D	3.418	1.12	0.98	0.12
J1701-3006E	3.234	0.16	0.07	0.03
J1701-3006F	2.295	0.20	0.06	0.02

<sup>a</sup>Statistics for pulsars A, B, and C are from Possenti et al. (2001)

## 5.3 Results and Discussion

### 5.3.1 Search Results

Our M62 search resulted in the discovery of three new binary millisecond pulsars, PSR J1701-3006D, E, and F. Their initial spin and orbital parameters are listed in Table 5.1, along with the parameters of the three previously known pulsars in M62 (D’Amico et al., 2001; Possenti et al., 2001).

Pulsar D was discovered in a straightforward FFT search, i.e., without using any special binary search techniques. This was aided by the fact that the pulsar is relatively bright ( $S_{1400} \sim 0.2$  mJy) and passes through its descending node during the August 2001 observation. Pulsar A, with its fairly long orbital period, was also detectable without any binary correction. For our acceleration searches, we divided the data into eight 1677.7 s segments. Each segment was searched separately with trial values up to  $|a_{\max}| = 60 \text{ ms}^{-2}$ . The acceleration search did not uncover any new pulsars, but pulsars A, B, C, and D were re-detected. The sideband search also revealed no new pulsars.

Pulsar E was discovered in the August data using the DPS method. It was not detected in the acceleration search of the August data, but was seen in an acceleration search with the follow-up December data. Pulsar F was discovered with the DPS method in the December data. It is visible in a DPS of all five of the December observations, but is not detected by any other method, nor is it visible at all in the

August data. Pulsars E and F therefore establish the validity of the DPS method as a pulsar search tool. (In fact, the DPS method also detected the other 4 pulsars in M62.) In Figures 5.5, 5.6, and 5.7, we show the dedispersed pulse profile and DPS for the discovery observation of each of the three new pulsars. There is strong evidence that pulsar E is eclipsed by its companion’s wind for  $\sim 15\%$  of its orbit (at 1400 MHz).

We have begun regular observations of M62 at Green Bank, with the goal of determining the timing solutions for the six pulsars. After a year of such observations, we will be able to precisely determine the rest spin period  $P$ , period derivative  $\dot{P}$ , sky position, and the five keplerian orbital parameters (orbital period, projected semimajor axis, eccentricity, longitude of periastron, and epoch of ascending node passage) for each pulsar.

### 5.3.2 Cluster Dynamics

Until recently, 47 Tuc and M15 were the only two globular clusters known to contain more than 2 pulsars. With the recent flurry of new discoveries, there are now 4 clusters containing 5 or more known pulsars. These are listed in Table 5.2, along with some of their salient parameters. Some pulsar population statistics for these clusters are listed in Table 5.3. When several pulsars are known within a single cluster, their timing characteristics can be used to determine several important properties of the cluster itself. After the pioneering work of Phinney (1992) and Anderson (1993) with M15, similar dynamical studies have been applied to 47 Tuc (Freire et al., 2001) and NGC 6752 (D’Amico et al., 2002). We expect that M62 will provide another laboratory in which these ideas can be applied.

Cluster pulsar  $\dot{P}$  measurements are contaminated by acceleration in the gravitational potential of the cluster. The effect of this motion on  $\dot{P}$  can be of the same order as, or larger than the intrinsic magnetic spindown of the pulsar. This may make it difficult to determine key derived parameters of the pulsar, such as characteristic age and magnetic field, but it allows us to study the dynamics of the cluster as a whole.




Figure 5.5: Pulse profile and DPS for PSR J1701-3006D.




Figure 5.6: Pulse profile and DPS for PSR J1701-3006E.




Figure 5.7: Pulse profile and DPS for PSR J1701-3006F.

Table 5.2: Globular clusters with the largest known pulsar populations.  $D_{\odot}$  is the cluster’s approximate distance from the sun.  $r_c$  is the core radius; a ‘c’ indicates a core collapsed cluster. Note that M62’s classification as core collapsed is somewhat uncertain (Djorgovski, 1993; Djorgovski & King, 1986), and some have argued that the core of 47 Tuc may have been collapsed in the past (de Marchi et al., 1996).  $M_V$  is the integrated absolute visual magnitude.  $\rho_0$  is the central luminosity density, and  $\rho_c$  is the central mass density derived from pulsar dynamical studies (Harris, 1996; Camilo et al., 2000; Freire et al., 2001; Anderson, 1993; D’Amico et al., 2002).

Cluster	Total Pulsars	$D_{\odot}$ (kpc)	$r_c$ (")	$M_V$ (mag)	$\log_{10} \rho_0$ ( $L_{\odot,V} \text{ pc}^{-3}$ )	$\rho_c$ ( $M_{\odot} \text{ pc}^{-3}$ )
47 Tuc	20	5.0	23	-9.42	4.77	$\gtrsim 4.0 \times 10^5$
M15	8	10.3	2.2c	-9.17	5.38	$\gtrsim 2.7 \times 10^6$
M62	6	6.9	11c	-9.19	5.14	
NGC 6752	5	4.0	10c	-7.73	4.91	$\gtrsim 7.1 \times 10^5$

Table 5.3: Pulsar population statistics of the top globular clusters.

Cluster	Isolated Pulsars	Binary, $P_b < 0.3$ d	Binary, $P_b > 0.3$ d	$P_{\min}$ (ms)	$P_{\text{median}}$ (ms)	$P_{\max}$ (ms)
47 Tuc	7	7	6	2.101	3.590	7.589
M15	7	0	1	4.027	18.636	110.665
M62	0	4	2	2.295	3.506	5.241
NGC 6752	4	0	1	3.266	5.277	9.035

If a negative  $\dot{P}$  is observed for a given pulsar, we can conclude that it lies in the more distant half of the cluster and is accelerating towards the earth. Assuming that the acceleration dominates the magnetic braking contribution to  $\dot{P}$ , we can use the measured acceleration and position to place a lower limit on the cluster's central density and mass-to-light ratio (Phinney, 1992). Using this derived mass-to-light ratio, the observed surface brightness profile, and the pulsar positions, we can determine the maximum line of sight acceleration for the remaining pulsars. This allows us to constrain their intrinsic spindown rates and quantities derived therefrom. We note that 13 of the 25 pulsars previously studied in 47 Tuc, M15, and NGC 6752 exhibit negative period derivatives, and we therefore reasonably expect approximately half of the pulsars in M62 to do the same.

In a dynamically relaxed cluster, the number density of objects of a given mass should follow a Boltzmann distribution, with heavier objects more tightly distributed about the core (Spitzer, 1987). Comparison of the radial distribution of pulsars and the observed distribution of turnoff main sequence stars in M62 should allow a statistical determination of the pulsar masses. Inspection of these distributions should also allow the determination of the dominant mass species in the central region of the cluster, and help constrain the cluster's initial mass function. This will tell us how many neutron stars were formed in M62 and might provide clues as to their retention rate and recycling efficiency. Of course, any pulsars whose characteristic ages do not exceed the core's two body gravitational relaxation time ( $\sim 4 \times 10^7$  yr; Harris, 1996) or pulsars that have obviously been ejected from the core must be excluded from this analysis.

Dynamical analyses such as these have been applied to the 7 core pulsars in M15 (Phinney, 1992; Anderson, 1993), the 15 pulsars in 47 Tuc with coherent timing solutions (Freire et al., 2001), and the 3 core pulsars in NGC 6752 (D'Amico et al., 2002) with great success. Analysis of the pulsar positions and accelerations in these clusters has yielded results that are not attainable by other means, and we can expect the application of these techniques to M62 to be just as fruitful.

### 5.3.3 Cluster Pulsar Demographics

Once there is a large enough sample of populous clusters, we may be able to form clear links between cluster properties and pulsar demographics. With the current census, however, few definitive conclusions can be drawn.

Selection effects have undoubtedly skewed the list of clusters appearing in Table 5.2. Some deserving clusters are certainly absent because their correct dispersion measure has not been found and/or binary selection effects have prevented the full revelation of their pulsar populations. With the advent of new search techniques such as the sideband and DPS methods, binary selection effects should become less of a factor. One might also argue that 47 Tuc appears in first place because of a reverse selection effect — since it is visible at Parkes for long periods when practically nothing else is, it has been observed far more than any other cluster.

The only clear similarity between the clusters listed in Table 5.2 is that they all appear to be fairly dense. Indeed, they are all within the top 80<sup>th</sup> percentile among the Galactic globular population. However, if we were to include the fifth ranking cluster (the very low-density M13, with 4 MSPs; Anderson, 1993, S. Ransom, private communication), this conclusion would fail to hold. In fact, the recycling rate is probably not a very sensitive function of central density (Phinney, 1996).

A glance at Table 5.3 shows that the pulsar populations in these top four clusters are distinctly different. There are, however, some similarities. Both 47 Tuc and M62 contain roughly equal numbers of so-called “normal” and “short-period” binaries. The “normal” binaries are characterized by orbital periods  $P_b \sim 0.4 - 4$  d and companion masses  $M_c \sim 0.2M_\odot$ , and are very similar to a number of systems found in the Galactic disk (see, e.g., Phinney & Kulkarni, 1994). The “short-period” binaries (a.k.a. “eclipsing” binaries, although not all of them exhibit eclipses) have  $P_b \sim 1.5 - 5.5$  hr and  $M_c \sim 0.03M_\odot$  (Camilo et al., 2000). M62 A and D clearly fall into the former category, while M62 C, E, and F are members of the latter group. M62 B, 47 Tuc V, and 47 Tuc W (along with Ter 5 A; Lyne et al., 1990) have very short orbital periods, but somewhat more massive companions ( $\sim 0.1M_\odot$ ), and may be

indicative of an emerging class of their own.

These pulsar systems are probably formed when an NS exchanges into a primordial main sequence (MS) binary, supplanting the lighter star. The longer-period systems are presumably the result of stable mass transfer, once the  $\sim 1M_{\odot}$  companion evolves off the MS and overflows its Roche lobe. Accretion from the giant's expanded envelope spins up the NS (Joss & Rappaport, 1983; Paczynski, 1983; Webbink et al., 1983). The short-period systems cannot be produced by direct exchange of an NS into an extremely tight primordial binary, because the liberated binding energy would most likely eject the new NS-MS pair from the weak gravitational field of the cluster. Rather, the short-period systems may have resulted from the exchange of an NS into a binary system with a more massive ( $\sim 1-3M_{\odot}$ ) MS companion (Rasio et al., 2000). In this case, when the secondary enters the giant stage, a common envelope phase can occur, making the orbit more compact. Further orbital decay due to gravitational radiation may be followed by another episode of mass transfer once the white dwarf companion again overflows its Roche lobe, resulting in the very low mass, short-period systems observed. Clearly the initial exchanges that produce such systems must occur within the first  $\sim 10^9$  yr of the cluster's life, when massive MS stars still exist. This fact indicates that 47 Tuc and M62 may have had similar early dynamical histories.

But 47 Tuc also contains a roughly equal number of isolated pulsars, whereas M62 has none, despite the fact that there are no systematic selection effects against detecting solitary pulsars. In fact, M62 is currently the only cluster with more than two pulsars, not containing an isolated pulsar. Based on the 47 Tuc population, we might expect  $\sim 3$  isolated pulsars in M62. Detecting zero would be only a  $\lesssim 2\sigma$  chance occurrence. In other words, we cannot confidently reject the hypothesis that the two populations are drawn from the same underlying distribution.

If the disparity is real, however, it is not easily explained. The large emerging population of short-period binaries (Deich et al., 1993; Camilo et al., 2000; Ransom et al., 2001; D'Amico et al., 2001) argues against companion ablation on short timescales, possibly indicating that these companions are no longer degenerate (van den Heuvel & van Paradijs, 1988) and may never completely evaporate (Robinson et al., 1995).

Companion destruction scenarios might therefore require systems even more compact than those currently observed, but it is unclear why these may have formed in 47 Tuc and not in M62. Complete tidal disruption of the mass donor is also an unlikely explanation for the isolated pulsars in 47 Tuc. In this case, the “companion” is disrupted during tidal capture and forms an accretion disk around the NS (Krolik, 1984). Significant spin-up is not expected from such systems, however (Verbunt et al., 1987), and this process is unlikely to form true MSPs. It appears that the most likely formation channel for the isolated pulsars in 47 Tuc involves collisions resulting from three-body encounters (Verbunt et al., 1987; Sigurdsson & Phinney, 1995). The lack of isolated pulsars in M62 might indicate that these catastrophic binary-single star encounters are less common (despite M62’s larger density), or preferentially eject the NSs from the cluster for some unknown reason.

Similarities are also apparent between M15 and NGC 6752. These post-core-collapse (PCC) clusters each contain a single binary pulsar system, both of which seem to have been ejected from the central regions of their respective clusters. This is indicative of a high rate of encounters in these clusters’ cores. This fact is also supported by the observation that the remaining pulsars in these clusters are all isolated, indicating that they have each likely experienced at least one exchange or collision in the past.

The differences between these two clusters are striking, however. The derived central mass-to-light ratio of NGC 6752 is 3 – 4 times larger than that of M15, suggesting a much flatter IMF, and a much larger concentration of compact objects in the core of NGC 6752 than in M15. PSR B2127+11C in M15 is a highly eccentric double NS binary, likely the result of an interaction in which the pulsar exchanged into a binary containing its current companion (Prince et al., 1991). The recoil associated with this exchange ejected the new system to the cluster outskirts. Despite the fact that PSR J1911-5958A in NGC 6752 is a more normal MSP binary ( $P = 3.3$  ms, eccentricity  $e < 10^{-5}$ , companion mass  $\sim 0.2M_{\odot}$ ), its dynamical history may have been quite unusual. Colpi et al. (2002) have suggested that a scattering event involving a massive black hole binary at the cluster’s center ejected the system to its

current location, over 3 half mass radii from the cluster center. In any case, these two binary systems are markedly different.

The isolated pulsar populations in M15 and NGC 6752 are also very different. NGC 6752 contains only bona fide MSPs, while the distribution of periods in M15 is much broader. The long-period pulsars may be the result of destructive tidal capture (Tavani, 1992), but the shorter-period pulsars are most likely the result of spin up in a binary system, followed by ionization or a destructive stellar collision.

Clearly, a larger sample of highly populated clusters will be required before the study of cluster pulsar demographics can move beyond speculation. There are hints of similarities and differences among the current top clusters, but their explanation is still elusive. For example, it is unclear why M62 should be so similar to 47 Tuc and so different from the other PCC clusters in Table 5.2.

This vexing situation may get worse before it gets better, but one certainty is that the list of highly populated clusters will indeed keep growing, perhaps by leaps and bounds. The recent resurgence of cluster pulsar discoveries will undoubtedly continue, aided by new and improved observing hardware and search techniques.

We are grateful to the NRAO-Green Bank staff for all of their support. We also thank Scott Ransom and Ingrid Stairs for freely providing information on unpublished search results. The National Radio Astronomy Observatory is a facility of the National Science Foundation, operated under cooperative agreement by Associated Universities, Inc. Access to the Hewlett-Packard V2500 computer, located at the California Institute of Technology, was provided by the Center for Advanced Computing Research. This work was supported by NSF grants 0005-1-000024 and 00040-1-000262.