PHYSICAL REVIEW D VOLUME 47, NUMBER 8 15 APRIL 1993

ARTICLES

Pulsar search using data compression with the Garching gravitational wave detector

T. M. Niebauer,* A. Rüdiger, R. Schilling, L. Schnupp,† W. Winkler, and K. Danzmann
Max-Planck-Institut für Quantenoptik, D-8046 Garching bei München, Germany
(Received 25 July 1991)

We describe briefly the standard problems associated with the detection of a periodic signal of astrophysical origin. We present a method for data compression when the bandwidth of the expected signal is much smaller than the detector bandwidth. The reduced data set can be corrected for the Doppler shift due to Earth's orbital and rotational motion even when the expected signal is close to the Nyquist frequency of the original data rate. Standard matched filtering techniques and Fourier transform processing can also be implemented on the compressed data stream. We applied this procedure to 100 hours of data taken with the Garching gravitational wave prototype detector. The results of this test reveal that the amplitude statistics of our prototype in the frequency domain follow closely the expected Rayleigh distribution over all 100 h. This analysis puts a constraint of 9×10^{-21} on possible gravitational wave strains produced by periodic sources in SN 1987A.

PACS number(s): 04.80+z, 95.75.Pq, 95.85.Sz, 97.60.Gb

I. INTRODUCTION

In March 1989, we undertook a 100 h coincidence run with the Garching and Glasgow prototype laser interferometric gravity wave detectors in order to verify their reliability over long time periods. In this paper, we present the data analysis methods and results from a pulsar search on this continuous 100 h using data from the Garching prototype only.

The goal of this investigation was to search for a pulsar within a restricted bandwidth of a few Hz near 2 and 4 kHz in the direction of the 1987A supernova. These parameters were chosen to agree with the observations of a submillisecond pulsar in SN 1987A reported by Kristian *et al.* [1].

An efficient scheme was developed using a digital complex heterodyne technique (CHT) for data compression. The original data were reduced by a factor of 1250 which greatly eased the computational burden of a Fourier transform over 100 h of data. The CHT has an added advantage that the Doppler shift correction and matched filtering can subsequently be performed on the reduced data set.

The results of the data analysis revealed that the frequency domain statistics follow the Rayleigh probability distribution that is expected to result from random white noise. The statistics from the entire 100 h of data allowed us to set a detection limit for signals in our restricted bandwidth, and originating from SN 1987A: strains of 9×10^{-21} for either polarization, with 95% confidence against false alarms.

II. THE 100 HOUR DATA RUN

The high sensitivity of the prototype laser-interferometric gravitational wave detectors has now led several groups to propose building full scale astrophysical observatories [2–4]. These plans require that the prototypes evolve from laboratory setups used to test new optical measurement techniques into stable astrophysical instruments. This led to the decision by the Glasgow and Garching gravitational wave groups to run their prototype detectors in coincidence for 100 h between March 2–6, 1989. The purpose of the data run was to demonstrate the ability of the interferometric prototypes to operate in a stable condition over time periods of many hours. Data were collected for later analysis of the interferometer performance, to look for possible signal correlations between the two prototype detectors, and to search the data for possible signals of astrophysical origin.

The data-taking run clearly showed that there are no fundamental obstacles for continuous operation of the planned full scale detectors. The duty cycle of both prototype interferometers was quite high. For example, the Garching prototype maintained a servo lock duty cycle of about 99% and required very little manual intervention. The overall duty cycle of the Garching prototype was, however, reduced to about 90% due mostly to the lack of a dual tape drive in the data acquisition hardware. Similar duty cycles were obtained with the Glasgow prototype [5].

*Present address: AXIS Instruments Company, 6400 Lookout Road, Bldg 105, Boulder, CO 80301.
†Deceased.

III. THE GARCHING DATA ANALYSIS

There were three distinct parts of the Garching data analysis. First a careful study has been made [6] of the time domain amplitude statistics and possible correlations with the various "housekeeping data" which monitored servo conditions and seismic noise. Second, we developed an efficient data compression algorithm ideally adapted to searching a narrow bandwidth for periodic signals. We applied the technique to the 100 h of data to examine both the statistical characteristics of the noise and to set strain limits on gravitational radiation from a restricted class of periodic signals in SN 1987A. Finally, the data from both prototype detectors are currently being examined for possible correlations [7].

The results of the first analysis have yielded a great deal of prototype-specific information, some of which was new. For example, we learned that the system loses lock for short periods (less than 1 min) during automatic argon refills in the laser system. Many harmonics of the 50 Hz line frequency were also observed which indicate that better electrical isolation would be helpful. In addition, we found that the laser cooling system introduced periodic short-lived spikes. Such observations have been one immediate benefit of the 100 h data run.

Shown in Fig. 1 is a sample Fourier noise spectrum of the Garching prototype during the test run. The curve was made by averaging 10 000 contiguous individual power spectra of length 50 ms each near the beginning of the data taking run. The very high spectral contributions below 1 kHz, due largely to seismic noise [8], have been partly removed by an analog high-pass filter (with a corner frequency of about 300 Hz) to avoid amplifier saturation caused by low-frequency drifts. The noise floor above 1 kHz becomes constant (white) and approaches the theoretical shot noise limit for the setup [8, 9]. The high-frequency portion of the spectrum was attenuated by two (four-pole) electronic low-pass filters with a corner frequency of about 4 kHz. These filters provided sufficient antialiasing to ensure a good signal-to-noise ratio (SNR) up to 4 kHz when the interferometer was sampled at 10 kHz.

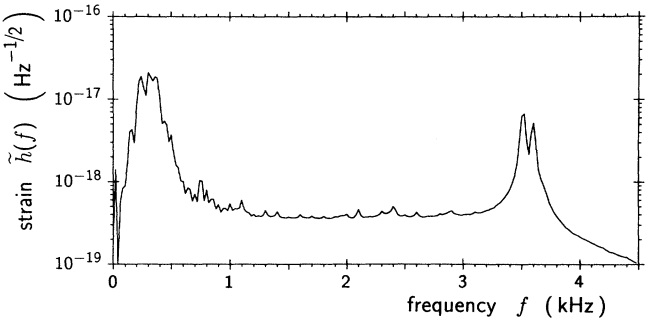

FIG. 1. Measured noise spectrum of the Garching GW interferometer. The spectrum is given as *linear* spectral density $\tilde{h}(f)$ of the apparent strain $h(t)$, in units of $\mathrm{Hz}^{-1/2}$. The noise spectrum is flat (and near the shot noise limit) in the frequency range from 1 to 3 kHz, and again above 3.7 kHz; the apparent drop is due to the 4 kHz antialiasing filter which, however, attenuates signals and noise in identical fashion. The double peak near 3.5 kHz is due to mechanical resonances in the setup.

The noise amplitude varied slowly as a function of time by as much as a factor of 2 or 3 over a span of several hours. The general trend was for the noise to become larger as time progressed, due to a slow misalignment of the beam. Toward the middle of the 100 h run the interferometer was manually readjusted, which helped bring the system back to the original noise level. At the time of the data run, the origin of the slow misalignment was unknown. Recently, however, a component responsible for beam steering was identified as having poor mechanical stability. Changing this component has reduced the beam steering drift quite markedly. Finding this problem was another example of how the prototype benefited from the data taking run. In the planned full scale detectors, however, the problem with beam steering will be eliminated by automatic servo controlled alignment systems.

IV. A PULSAR SEARCH IN SN 1987A

A. General motivation

This paper describes the results of that part of the Garching data analysis designed to search for periodic signals of astrophysical origin. The original motivation for this analysis was provided by the reported optical sighting of a submillisecond pulsar in the remnants of supernova 1987A by Kristian *et al.* [1]. They observed an optical pulsar of frequency 1968.629 Hz with significant power in higher harmonics. They also reported that the pulsar observations were improved by a transformation to barycentric coordinates to remove effects of the Doppler shift. In the same paper, they speculated that obscuring matter could explain the fact that it was not observed 13 days later. Efforts made to duplicate the observation ultimately failed [10], and it is now believed to have been caused by instrument error [11].

Although the reported frequency and its second harmonic (2 and 4 kHz) are well within the measurement bandwidth of the present gravitational-wave (GW) prototypes, any reasonable estimates of the expected strength from such a pulsar are still several orders of magnitude smaller than the sensitivity of these detectors. If, for a rough estimate, one assumes that the rotational energy of the reported pulsar (estimated by Kristian *et al.* [1] as about 6×10^{45} J) would be converted into GW radiation over only five years, the expected strain amplitude at the earth would be as low as 5×10^{-25}. This limit is still about three orders of magnitude too small to be seen using 100 h of data with the Garching prototype (assuming a constant noise background of about $10^{-19}\,\mathrm{Hz}^{-1/2}$ reported by Shoemaker *et al.* [8]).

Despite the poor chances of detecting gravitational radiation from pulsars in SN 1987A, we decided that it would nonetheless be useful to search for possible signals using the 100 h data from the Garching detector to see

what practical limit could be obtained. In addition to any physical significance of such a limit, we could also gain valuable experience and develop analysis techniques which may prove useful when the full scale observatories are operational. The pulsar search also provided a practical framework in which the frequency domain statistics of the noise from the prototype detectors over long observation times could be studied and compared with the expected statistical distribution.

B. Overview and goals

The pulsar search was made in a bandwidth of about 4 Hz around both the fundamental frequency 1968.629 Hz reported [1], and its second harmonic. The relatively large bandwidth was chosen to allow for a possible frequency drift between the time of the optical sighting and the time of our 100 h run. This bandwidth also allows one to make a very simple Doppler shift correction over the entire 100 h as will be explained later. We restricted the pulsar search to the direction of supernova 1987A when applying the Doppler shift corrections.

An important part of this analysis was to incorporate a data compression (decimation) algorithm which reduced the volume of the data by a factor given by the ratio of the original bandwidth to that required for signal detection. In this case, the original data were sampled at 10 kHz corresponding to a Nyquist bandwidth of 5 kHz. The data were compressed by 1250 resulting in a 4 Hz bandwidth for the pulsar search. The fast Fourier transformation (FFT) over the entire 100 h was quite easily managed with normal memory requirements on a Cray-XMP in less than 30 s of CPU time. In fact, the large majority of the total CPU time (about 45 min) was used for the data compression algorithm itself. Therefore some effort was spent trying to optimize the technique and the program code.

It was also important to develop a technique which would allow flexible signal processing on the reduced (compressed) data set. In particular, it was necessary to be able to correct the signal for amplitude and phase modulation caused by Earth's motion relative to the source.

The data compression scheme which we developed (and will describe in the next section) satisfies both goals of being efficient in terms of CPU time and memory and being flexible in terms of subsequent signal processing. The compression requires roughly one complex multiplication and addition per data point. The technique is applied in the time domain and thus does not require large memory stores often needed for convolution techniques in the frequency domain. Furthermore, we will show that signals in the compressed data stream can be simply (and with low computer cost) amplitude or phase modulated. Thus, one can apply a great variety of different analyses to the same set of compressed data. This may be useful, for example, if one wants to look for pulsars over the entire sky where a different Doppler shift correction is required for every source direction.

In order to emphasize the importance of data reduction for broadband laser interferometry, let us consider the 100 h data set from the Garching prototype. The main interferometer was sampled at 10 kHz. The various *housekeeping data* channels raised the overall data rate to 20 kHz. Over 100 h about 10^{10} samples were written onto tape. The Glasgow prototype, for an identical time resolution, had an even higher overall sampling rate of 60 kHz.

Therefore it is important to keep the number of operations made on the full data set small to avoid excessive computational time. Even more problematic is the large amount of memory that is required to perform the efficient FFT routines. This limitation would rule out the possibility of making an FFT over the 100 h without the data reduction mentioned above. It is, of course, possible to perform large FFT's in small pieces [12] but only at the expense of severe input-output limitations. Certainly, large memories exist and faster computers will be developed, but this does not diminish the utility of developing efficient algorithms or data compression techniques for GW detection.

The pulsar analysis presented below can be split into four distinct categories: data compression, Doppler shift corrections, polarization dependent demodulation, and, finally, statistical peak detection. The latter three topics have been discussed by Livas in connection with a data analysis of the MIT prototype [13,14]. However, the data compression and especially our method of Doppler shift correction has not yet been discussed for GW detection.

V. DATA COMPRESSION

A. General discussion

The sampling theorem states that any continuous real function with a bandwidth limit B can be completely specified by discrete values sampled with a frequency $f_S \geq 2B$. Defining the *Nyquist* frequency f_N as half the sampling frequency, the sampling theorem is valid if $f_N \geq B$. One can extend this concept to digital data streams where the original sampling rate is higher than is needed to record some signal of interest. The originally sampled data stream can be digitally filtered to preserve desired signal frequencies and then resampled at a rate compatible with the necessary signal bandwidth. The data compression can be optimal in the limit where the digital filters are square in the frequency domain.

It is natural to combine data compression with any data analysis that has a bandwidth much smaller than that contained in the original data stream. Data reduction is especially important in cases like this pulsar search where the bandwidth ratio is quite large (1250).

There are several techniques which can be used to compress the data as well as several different types and realizations for the required digital filters. In the next subsections, a brief introduction to the concepts and the filter requirements for three different data compression techniques will be given. These will be called single quadrature heterodyning, aliasing, and complex heterodyning techniques. The advantages of the latter will be explained, and we will show that this scheme is efficient and also makes possible a very simple way to perform

the necessary Doppler shift corrections on the reduced data set.

A discussion of the digital filtering requirements, techniques for filter realization, and the computational demands for the different types of digital filters can be found in Appendix A.

B. Compression techniques

1. *Single quadrature heterodyne technique (SQHT)*

One simple approach to data compression is to apply a bandpass filter centered on the desired signal frequency (not necessarily near dc). This can be followed by a multiplication (heterodyne) with a sinusoidal signal, at a frequency f_m equal to the lower edge of the filter. The difference frequencies will be shifted near dc and the sum frequencies can be removed by a low-pass filter before resampling at a lower rate. The first (bandpass) filter is required to keep noise from below f_m to be heterodyned to the same place as the corresponding signal frequencies above the heterodyne. The final low pass is needed to avoid aliasing from the sum frequencies. (Notice that the initial bandpass could be replaced by a high-pass filter with the corner frequency at the lower edge of the signal bandwidth.)

2. *Aliasing technique (AT)*

A better approach might be to apply a bandpass filter centered on the desired signal frequencies (as described above) and simply resample without heterodyning or low-pass filtering [15]. This technique "aliases" the lower edge of the passband to dc if the low-frequency cutoff is chosen to be a multiple of the resampling frequency. The advantage of the AT is that one avoids the heterodyne and the final low-pass filter needed in the SQHT.

3. *Complex heterodyne technique (CHT)*

The most flexible data compression scheme is the CHT. The CHT requires a heterodyne (multiplication) of the original data with both a sine and a cosine function, with frequency f_m chosen to lie at the center of the signal bandwidth. These two heterodyned series are then low-pass filtered (for reasons of antialiasing) before resampling.

We will refer to the real functions resulting from the heterodyning of the original data X_t as the cosine and sine quadratures

$$\begin{aligned} C_t &= X_t \cos \omega_m t\,, \\ S_t &= X_t \sin \omega_m t\,, \end{aligned} \tag{1}$$

where t is an integer index denoting the original sampling times. It is usually more convenient to consider these two real quadratures as a single complex time series, the *analytic signal*, where the cosine and sine quadrature are taken to be the real and imaginary components, respectively, i.e., $A_t = X_t \exp(i\,\omega_m t)$. We will use these same names for the heterodyned functions directly after multiplication with the original data, after low-pass filtering, and also after the resampling. This should cause no confusion since the functions contain the same signal information after each step.

The analytic signal produced by the CHT is simply a frequency shifted version of the original data according to standard Fourier transform theory. Any signal $x_s(t) = x_{s0} \exp(-i\,\omega_s t)$ will be shifted to $x_{s0} \exp[-i(\omega_s - \omega_m)t]$. The *mirror* frequencies equidistant on either side of the heterodyne, which would be mapped to the same place in the SQHT, are instead separated into positive and negative frequencies using the CHT. This can be understood in the frequency domain since the cosine quadrature contains the sum of two terms coming from the mirror frequencies of the original data. The sine quadrature, on the other hand, contains the difference of these two contributions (with the imaginary factor $-i$). The complex sum $C + iS$ (the analytic signal) separates these contributions neatly.

One consequence of the above discussion is that no filter is needed before the complex heterodyning. Thus, the CHT requires only one complex multiplication and a digital low-pass filter before the resampling process. The simpler filtering requirement can be an important advantage over the other two methods we have discussed. The main difference is that for the SQHT and the AT very sharp filters are required because of the high corner frequencies relative to the bandwidth. For example, this pulsar search would require digital filters with cutoff frequencies near 4 kHz which begin to roll off after 0.1 Hz if the SQHT or the AT is used. The CHT, on the other hand, requires a single low-pass filter with a corner frequency of about 2 Hz which rolls off after about 0.1 Hz. We will see later that a high quality low-frequency digital filter can be constructed that uses only about one addition for each original data sample.

At first glance, it might seem that the data compression is worse by a factor of 2 for the CHT because of the need to store both the real and imaginary components (cosine and sine quadratures) for the resampled data stream. This is not the case, however, because the positive and negative frequency components of the complex data stream are unique. As mentioned above, the frequencies above and below the heterodyne are separated into the negative and positive frequencies. This is in contrast to the components of the real data streams produced by the SQHT or AT in which the positive and negative frequencies are redundant. This means that the required resampling rate for the CHT is only $\frac{1}{2}$ that needed for the SQHT or the AT. Thus, any of these techniques can be used to achieve the same data compression factor. In practice, the compression is limited only by the quality of the digital filtering and the required bandwidth.

One of the most useful advantages of the complex heterodyning (CHT) is that the phase of the heterodyne frequency can easily be varied using a *second* complex heterodyne in the time domain by multiplying the complex analytic signal with the phase function $e^{i\phi}$. This phase factor simply adds to the initial heterodyne func-

tion $\exp(i\omega_m t)$. The phase adjustment ϕ can have a time dependence (and may also have a frequency dependence), thus allowing the heterodyne to follow a phase (or a frequency) modulated signal. The most important point, however, is that when the phase modulation is slow compared to the resampling rate, these corrections can be made on the reduced (resampled) data set.

This technique is useful for correcting pulsar signals for the phase modulation caused by the Doppler shift. We will see later that this technique is especially helpful in the case where the pulsar signal is sparsely sampled in the original data stream. It is also convenient that the Doppler correction can be made for sources in many different parts of the sky using the same compressed data set.

C. Data compression on the Garching data

The data compression was made for two different frequencies which we will call the 2 and 4 kHz heterodyne channels. More precisely, the center (or heterodyne) frequencies are given by the fractional values $\frac{2010}{1021}$ kHz (1968.658 Hz) and its second harmonic. These frequencies were chosen to be near those reported by Kristian *et al.* [1], and to be rational so that the functions could be stored in a look-up table. The original data X_t were multiplied by a sine and cosine heterodyne function as shown in Eq. (1) for the 2 and 4 kHz channels.

The last step of the CHT required low-pass filtering of each quadrature for each channel prior to the resampling. This was broken down into three successive stages of filtering and resampling. Since each stage operated on a successively smaller input data stream, we were able to use increasingly sophisticated filters. The result was that we could have a very simple filter for the first stage and still end up with a sharp filter for the last stage with almost no computational increase.

The first filter stage was a simple moving boxcar average of 50 data points. These low-passed data were resampled such that neighboring averages overlapped by 50%. This choice of resampling moved the Nyquist frequency of the resampled data stream outside the main lobe of the $\mathrm{sinc}(x)=\sin(x)/x$ produced by the boxcar low-pass filter (see Appendix C). The first step compressed each quadrature (for each channel) by a factor of 25, from 10 kHz to 400 Hz.

The second stage compressed each quadrature by another factor of 25 producing time series resampled at 16 Hz. The filter chosen for this step was a two-pole Butterworth infinite impulse response (IIR) digital filter with a cutoff frequency f_c of about 2.53 Hz. The measured transfer function agreed very well with the usual Butterworth form

$$H^2(f) = \frac{1}{1+(f/f_c)^{2n}}, \tag{2}$$

where n is the number of poles.

The final stage compressed the data by another factor of 4 so that each quadrature was sampled at 4 Hz. This corresponds to a reduction in sampling rate by a factor of 2500. However, since we have to store two quadratures (sine and cosine) the data reduction factor is only 1250. The antialiasing filter for this final step was an IIR six-pole Butterworth digital low-pass filter with a cutoff frequency of about 1.4 Hz. The cutoff frequencies for both the second and third stages were chosen so that the attenuation was about 20.0 dB at the resampled Nyquist frequency.

The first boxcar filter and resampling stage required only about one complex addition per original data point (for each channel). Furthermore, since this first filter was not recursive it could take full advantage of the Cray vectorization. The following two stages, although more complex, required less computation because they operated on compressed data sets. In fact, the first filter stage dominated the total computational CPU time associated with the entire pulsar search.

As discussed in Appendix A, the IIR filters can be realized with about n multiplications and additions on each input data point when the data reduction factor is greater than the number of filter poles n. However, since an IIR filter is recursive, it cannot be vectorized on the Cray. Therefore for the last two stages, the code was interleaved so that both quadratures of both channels were handled in parallel.

Finally, we note that the IIR low-pass filters can suffer significant roundoff error if the data reduction factor is too high relative to the precision of the arithmetic. This was noticed in single precision (four byte) arithmetic when trying to compress the data by the full factor of 1250 using only one filter and resampling stage on our mainframe computer. The roundoff problem was no longer observable after the process was broken into the three steps described above. Furthermore, since the pulsar search was ultimately run on the Cray-XMP with 64 byte floating point arithmetic, the roundoff error in the digital filters was insignificant.

This three step low-pass filter and resampling procedure had the computational efficiency of a simple moving average but resulted in a filter shape corresponding to a six-pole Butterworth filter. Each compressed quadrature was finally sampled at 4 Hz giving an overall data reduction factor of 1250 compared to the original data stream. The bandwidth of the pulsar search was restricted to about 3.6 Hz centered on each heterodyne frequency in order to avoid looking at frequencies containing significant aliasing. The final bandwidth could have been increased somewhat by using even better low-pass filtering in the last stage. However, trying to gain more bandwidth by using sharper filters gives diminishing returns once the unaliased bandwidth is close to the Nyquist frequency of the compressed data.

VI. DOPPLER SHIFT CORRECTION

Relative motion between the source and the detector gives rise to the well-known Doppler shift. Earth's rotation and the orbital velocity both combine to modulate the frequency (or phase) of possible extraterrestrial signals. In the case of a periodic signal, the modulation splits the signal into many frequency components separated by diurnal and annual periods. This seriously

degrades the SNR in any single frequency bin when the observation time is long enough to resolve these frequencies.

One can remove the Doppler shift caused by Earth's motion by correcting the laboratory measurements by the time it would take a signal to travel from the detector to some rest frame — the barycentric frame. The time delay between the barycentric and lab frames can be calculated simply as the dot product of the Earth-barycenter vector with a unit vector in the direction of the source star. Over 100 h this can be described adequately by a time delay δt given by a sum of two sinusoidal terms:

$$\delta t = a_1 \cos(\Omega_1 t + \phi_1) + a_2 \cos(\Omega_2 t + \phi_2)\,, \tag{3}$$

where Ω_1 and Ω_2 are the orbital and rotational angular frequencies, respectively. The amplitudes and phases depend on the location of source and detector.

The time difference between the Garching detector and the barycenter for a signal originating from the direction of SN 1987A was calculated using a standard routine EARTH [16] which includes perturbations of the barycenter due to the planets. The results of the program were fitted to Eq. (3) giving the parameters: $a_1 = 31.43\,\mathrm{s}$, $a_2 = 5.00\,\mathrm{ms}$, $\phi_1 = 214°$, and $\phi_2 = 313°$. The time was referenced to the starting of our 100 h run at 16:00 local time on March 2, 1989. The residuals after fitting the time delays calculated using EARTH to Eq. (3) were found to be less than $100\,\mu\mathrm{s}$ (the time between 10 kHz samples).

A signal $s(t)$, which is a pure sinusoid with frequency ω_s and amplitude s_0 in the barycentric frame, acquires a phase modulation in the lab frame due to the time correction δt of

$$s(t) = s_0 \sin \omega_s (t + \delta t)\,. \tag{4}$$

Using Eq. (3) for δt, and expanding Eq. (4) we can estimate the observation time one can have before the Doppler shift becomes important. The constant and linear terms in a power expansion of the time correction (as a function of time) can be ignored for this pulsar search. The constant term causes only a phase shift and the linear term introduces only a frequency offset (small compared to the bandwidth of the search). The quadratic term, however, causes the frequency to change with time. The frequency can change by one bin (given by the reciprocal of the observation time) in 11 min for the orbital term and 4 min for the rotational term for a signal frequency of about 5 kHz.

The spread in frequency due to the modulation can also be understood by expanding Eq. (3) in the frequency domain. The result is that the signal frequency ω_s is split into many sidebands with separations given by Ω_k where $k = 1, 2$ for orbital and rotational frequencies, respectively. The amplitude of the nth sideband, with frequency $\omega_s + n\Omega_k$, is given by $J_n(a_k\omega_s)$ and is significant until n becomes greater than the argument of the Bessel function. This quantity $a_k\omega_s$, the so-called modulation index, is quite large for both the orbital and rotational terms for signal frequencies in the kHz range.

A. Interpolation

Compensation for the Doppler shift due to Earth motion is accomplished by using the barycentric time instead of the laboratory time for the Fourier transform variable. This transformation provides the correct demodulation for signals of any frequency. A practical problem arises, however, when trying to perform the discrete Fourier transforms (DFT) on a data stream which is not sampled at equal intervals in the barycentric frame. In theory this can be solved by interpolating the data consisting of equally spaced data in the lab frame to equally spaced samples in the barycentric frame. The sampling theorem states that a function x which is band-limited to B and sampled with a period $T < \frac{2}{B}$ can be interpolated *precisely* using the relation

$$x(t) = \sum_{n=-\infty}^{\infty} x(nT)\,\mathrm{sinc}\left[\pi\left(\frac{t}{T} - n\right)\right]\,, \tag{5}$$

where $\mathrm{sinc}(x) = \sin(x)/x$ and n is an integer. In practice, however, Eq. (5) converges very slowly. Moreover, the interpolation would have to be done on the original data stream which increases the computational burden.

B. Rebinning

A very common procedure for coarse interpolation is to "rebin" the data into the nearest of the equally spaced time bins in the barycentric frame. This method, called "binning," consists of omitting or repeating a data point whenever the lab frame slips relative to the barycentric frame by one sample. The time scale for the slip between the lab and barycentric frames by one 10 kHz sample is about 10^7 samples. This limit comes from the quadratic term of the Doppler shift due to Earth's rotation.

At first glance, it might appear that omitting one data point out of every 10^7 should not seriously affect the data evaluation. In fact, however, binning fails quite badly when the signal is only sparsely sampled in the original data. This is clearly the case for our 4 kHz channel in which the signal would be sampled only 2.5 times during one period. Omitting a data point, therefore, introduces a sudden phase change of 144° causing a severe loss of signal.

In order to appreciate the magnitude of the error which may be introduced by binning, consider the simple case of a sinusoidal signal with a constant Doppler shift. Then the time delay between the lab and reference frame changes with a constant velocity. The effect of binning can be seen by omitting one data point in the middle of a record containing a pure sinusoid. This can be expressed analytically as the sum of two square time windows one of which is time shifted by one sample. The resulting signal is distributed over sidebands with amplitudes A_n given by

$$A_n = \mathrm{sinc}\left(\frac{\pi n}{2}\right) \cos\left[\pi\left(\frac{n}{2} + \frac{f_s}{f_S}\right)\right] \quad \text{for } n = 0, \pm 1, \pm 2, \ldots \tag{6}$$

where f_s is the signal frequency and f_S the sample frequency. Thus, for our case of data sampled at 10 kHz, binning would reduce the amplitude of the largest sideband of a 2 and 4 kHz signal to 80% and 70% of the original value, respectively. The amplitudes drop further as the observation time is increased so that many slips occur between the lab and barycentric frames. Notice that this result is independent of the length of the data record.

We have shown with this simple example that binning can introduce serious errors for signal frequencies which are not small against the sampling frequency. These results are only marginally improved for these frequencies if linear interpolation is used instead of binning. As stated above, it is not practical to consider higher-order interpolations using Eq. (5) since the series converges too slowly. It has been suggested [14, 17] that the data can be oversampled to reduce the binning error. However, this solution is clearly not in line with the our stated purpose of reducing the data rate.

C. Frequency tracking

An alternative to interpolation of the input data is to slowly vary the phase of the heterodyne frequency to follow the Doppler induced phase shift. The correction is made by simply multiplying the analytic signal with $\exp(i\,\omega_s \delta t)$, where ω_s is the anticipated signal frequency near 2 or 4 kHz and δt is the time correction between the lab and barycentric frames. This can also be understood as a slow rotation of the complex plane so that the Doppler shifted signal remains stationary. One important advantage of this procedure over interpolation is that the phase shift can be calculated explicitly so that the Doppler shift correction can be made exact. Perhaps of even greater importance is the fact that the correction can be performed on the reduced data.

The time dependent phase factor is a function of the source and detector locations (because of δt) and also depends linearly on the expected signal frequency ω_s. In our particular pulsar search (SN 1987A) the frequency ω_s was assumed given from astronomical observation. But in a general survey ω_s is not known in advance. It is possible, however, to use the heterodyne frequency ω_m (instead of ω_s) to correct for the Doppler shift of all signals inside a limited bandwidth. This is equivalent to making a second complex heterodyne (in the time domain) on the reduced data, with the function $\exp(i\,\omega_m \delta t)$ where δt is calculated for a given source location in the sky. The phase error associated with this procedure is given by the product of the time correction δt and the difference $\omega_m - \omega_s$ between the heterodyne and signal frequencies. A numerical calculation of the phase error in the Doppler shift correction using a frequency of about 4 kHz and the time function given in Eq. (3) integrated over 100 h constrains the bandwidth to about 5 Hz. This sets the scale for the largest reasonable resampling rate such that the Doppler shift can be made on the reduced set with a constant frequency over the given observation time.

VII. POLARIZED GRAVITATIONAL WAVES

Gravitational radiation emitted from a pulsar spinning with constant frequency has a polarization which is dependent on the angle between the spin axis of the source and the observation point [18]. Ignorance of the spin axis therefore introduces an uncertainty in the polarization state of the gravity wave. In general, the gravity wave is a linear combination of two polarizations denoted as h_+ and $h_\times$. In addition, there is an unknown constant phase difference between the two states. Furthermore, an arbitrary rotation of the observation axis in the direction of polarization will mix the two possible linear polarization states [19].

Interferometric GW detectors have different sensitivities to the two linear polarizations. The sensitivity to a given polarization changes as the plane of the detector rotates around the propagation direction of the gravitational wave. Therefore, the measured signal $V(t)$ of an unknown pulsar can be written as a linear combination of the two polarizations multiplied by the corresponding sensitivity of the detector to a certain polarization,

$$V(t) = A_+ S_+(t)\cos(\omega_s t + \phi_0) + A_\times S_\times(t)\cos\omega_s t\,, \qquad (7)$$

where A_+ and $A_\times$ are the strengths of the orthogonal polarization states, ω_s is the frequency of the gravity wave, and ϕ_0 is the relative phase between the two components. The sensitivities S_+ and $S_\times$ are in general time dependent functions of the relative position between the (fixed) source and the detector attached to a moving earth.

The sensitivity to the two orthogonal linear polarizations can be derived by rotating the reference frame of the interferometer into the plane perpendicular to the propagation direction of the signal. In this frame, the spatial components of the h_+ and $h_\times$ strain matrices [20] are

$$h_+ = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad \text{and} \quad h_\times = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \qquad (8)$$

where the z axis is aligned with the propagation direction.

The results written in terms of the source location (propagation direction) and the detector coordinates are

$$S_+ = 3\sin^2\alpha\,\sin^2\theta\,\cos 2\psi + \sin 2\alpha\,\sin 2\theta\,\cos 2\psi\,\cos\eta - 2\sin 2\alpha\,\sin\theta\,\sin 2\psi\,\sin\eta + (1+\cos^2\alpha)(1+\cos^2\theta)\cos 2\psi\,\cos 2\eta - 2\,(1+\cos^2\alpha)\cos\theta\,\sin 2\psi\,\sin 2\eta\,, \qquad (9)$$

$$S_\times = 4\sin\alpha\,\sin\theta\,\sin 2\psi\,\cos\eta + 2\sin\alpha\,\sin 2\theta\,\cos 2\psi\,\sin\eta + 4\cos\alpha\,\cos\theta\,\sin 2\psi\,\cos 2\eta + 2\cos\alpha\,(1+\cos^2\theta)\cos 2\psi\,\sin 2\eta\,. \qquad (10)$$

Here α and β denote the declination and right ascension of the source, respectively. The detector has a colatitude and east longitude of θ and λ, respectively, and ψ is the angle between one arm of the interferometer and the local northerly direction. The time dependence has been combined in the function $\eta(t) = \lambda - \beta + \Omega_E t$, where Ω_E is the earth rotation and t is the local sidereal time. One can see from Eqs. (9) and (10) that the sensitivity of the detector to either polarization has both diurnal (η) and semidiurnal (2η) terms. In addition there is a time independent sensitivity to the + polarization.

One can show in general that there is a 90° phase shift between the two polarizations for both the diurnal and semidiurnal terms. Thus, the integral over time of the product of the two sensitivities is zero — irrespective of the location of source and detector. This is a manifestation of the orthogonality of the two GW polarization states. Either polarization can therefore be selected by multiplication of the detector output with one of the desired sensitivity functions S_+ or $S_\times$ before signal detection.

In the case of pulsar detection, a signal can be expressed as the sum of two amplitude modulated sinusoidal signals (one for each polarization) as is shown in Eq. (7). The amplitude modulation S_+ and $S_\times$ split each sinusoidal signal into as many as five sidebands due to the dc, diurnal, and semidiurnal terms. This naturally leads to a reduction in SNR if the pulsar search is made for a single frequency. The power in these sidebands can be restored to the carrier frequency by multiplying the detector output by one of the sensitivity functions shown in Eqs. (9) and (10). As stated above, this procedure also selects one of the two polarizations due to the orthogonality. If the signal polarization state is known, this procedure will give optimal signal detection because it simply reduces to matched filtering the signal.

If the polarization of the signal is not known, however, examples can be found where filtering each polarization can either increase or decrease the chance of signal detection when scanning the Fourier transform for a significant peak. Consider, for simplicity, a detector placed on the north pole receiving GW radiation from a source directly overhead. The GW signal given by Eq. (7) reduces to

$$V(t) = A_+ \cos 2\Omega_E t \cos(\omega_s t + \phi_0) + A_\times \sin 2\Omega_E t \cos \omega_s t . \quad (11)$$

If only one polarization is present (i.e., A_+ or $A_\times$ is zero) then the signal has two sidebands at $\omega_s \pm 2\Omega_E$. Here it is helpful to demodulate the signal so that the signal in both sidebands is added and transformed to the original carrier at ω_s. On the other hand, the amplitude modulated signal is already single frequency when both linear polarizations are present with equal amplitudes and a phase difference $\phi_0 = 90°$ (circular polarization). Then it would be best to search for a pulsar without first demodulating for a particular polarization.

Typically the signal is more complicated than the example given above because the amplitude modulation produces more than two sidebands for each polarization. Thus, it is usually better to look for both polarizations. In this pulsar search we will first search the spectrum without any polarization corrections and then include them to isolate the strain amplitudes for both polarizations independently.

VIII. STATISTICAL PEAK DETECTION IN THE FREQUENCY DOMAIN

The last step in the analysis is to search the frequency domain for statistically significant peaks. Naturally, a knowledge of the noise statistics in the frequency domain is needed to specify the probability distribution for the power spectrum. One must find a statistical model which correctly describes the actual measurements. Once this is known the confidence that a peak will be found above a given threshold can be calculated. It is important to consider the statistics for the case where a signal is present as well as that given only by noise.

In the following discussion, we will assume white noise, Gaussian in the time domain, although some of the results can be derived for more general distributions. The main advantage of the Gaussian model is that simple analytic formulas can be derived for the probability distribution of the power spectrum. We will see that this model describes very well the frequency domain statistics for the 100 h Garching data run.

A. Gaussian noise

Let us assume that every sample $V(t)$ of a noise process can be treated as an independent random variable given by a Gaussian probability distribution in the time domain

$$p(V) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(V-V_0)^2/2\sigma^2} , \quad (12)$$

where V_0 and σ^2 are the mean and variance of the distribution, respectively. In the case of no signal, the mean V_0 is zero. The variance σ^2 is constant if the noise is *stationary* over time.

The discrete Fourier transform (DFT) of the noise signal $V(t)$ yields a noise spectrum $\tilde{V}(\omega)$ given by

$$\tilde{V}(\omega) = \sum_t V(t) e^{-i\omega t} . \quad (13)$$

It is well known that arbitrary linear combinations of Gaussian random variables also produce Gaussian random variables [21]. Thus, the real and imaginary parts at every frequency are also given by Gaussian random variables. The spectrum is frequency independent or *white* as long as the samples are uncorrelated in the time domain. It is also easy to show that the real and imaginary parts are statistically independent when the noise is stationary. In this case, the distribution of the frequency domain amplitudes can be found by integrating the joint probability distribution of the components over all phase angles. The result is a white amplitude spectrum with the well-known Rayleigh distribution. The squared amplitude spectrum or *power spectrum* with components $P = \tilde{V}\tilde{V}^*$ has an even simpler probability distribution given by an exponential function

$$p(P) = \frac{1}{P_0} e^{-P/P_0}, \qquad 0 \le P \le \infty, \tag{14}$$

where $P_0 = 2\sigma^2$ is the average noise power. As stated above, Eq. (14) is independent of frequency.

1. Adding power spectra

Groth [22] has pointed out that it is sometimes useful to add the power from different frequency components either in the same spectrum or in different spectra to accumulate the signal power:

> As examples, the sum might consist of (a) one bin from each of n distinct spectra, all corresponding to the frequency of an expected periodic signal; (b) n bins corresponding to n harmonics of an expected periodic signal; or (c) n bins located within the peak of a quasi-periodic signal whose width is greater than or the order of n bins [22].

In this paper, we will be most interested in the first example (a) where several power spectra are added to improve the SNR.

The probability distribution for the sum of n power spectra, each with a constant noise power P_0 and no signal, follows a χ^2 distribution with $2n$ degrees of freedom,

$$p_n(P) = \frac{1}{(n-1)!} \frac{P^{n-1}}{P_0^n} e^{-P/P_0}. \tag{15}$$

This equation is easily proven inductively by convolving Eqs. (14) and (15) to find $p_{n+1}(P)$. The assumed stationarity of the noise guarantees that each frequency component is independent with the same mean value P_0 which can be estimated by averaging the power spectrum over all frequencies.

The mean and variance of the distribution in Eq. (15) are

$$\langle P \rangle = nP_0 \quad \text{and} \quad \langle P^2 \rangle - \langle P \rangle^2 = nP_0^2. \tag{16}$$

Thus, the average noise power grows with the number n of spectra whereas the width of the distribution increases only with $\sqrt{n}$. Since the signal power will also increase linearly with the number of spectra, the overall sensitivity increases as $\sqrt{n}$. In linear measure, the SNR increases as the fourth root of the number of added power spectra. This also means that the sensitivity increases as the fourth root of the total observation time used to accumulate the spectra.

2. False alarm threshold

The probability $\mathcal{P}_0$ of finding a power greater than some threshold P_T at any given frequency is found by integrating the probability distribution $p_n(P)$:

$$\mathcal{P}_0 = \text{Prob}(P > P_T) = \int_{P_T}^{\infty} p_n(P)\, dP. \tag{17}$$

The probability of detecting a false signal originating from the noise in any of N frequency bins can be found using the Poisson distribution with a mean $\mu = N\mathcal{P}_0$. This uncertainty $\mathcal{P}_1$ is sometimes called the *error of the first kind* and is given by

$$\mathcal{P}_1 = \text{Prob}(P > P_T) = \mu e^{-\mu}. \tag{18}$$

One can use this equation to set the threshold for the pulsar search by equating $\mathcal{P}_1$ to the desired probability for "false alarms."

B. Gaussian noise plus signal

The *error of the second kind*, $\mathcal{P}_2$, on the other hand, refers to the probability that a real signal with power $P_s > P_T$ is not detected in the presence of noise. To calculate this error one needs the probability distribution $p_n(P; P_s)$ for the sum of n power spectra consisting of noise and signal P_s.

For a single spectrum, $\mathcal{P}_2$ can be calculated using the Rayleigh distribution displaced from the origin by the signal strength. The result for the sum of n spectra is more difficult to derive and we simply quote the result obtained by Groth [22] where each individual power spectrum has been normalized to unity (i.e., $P_0 = 1$)

$$p_n(P; P_s) = \left(\frac{P}{P_s}\right)^{(n-1)/2} e^{-(P+P_s)} I_{n-1}\left(2\sqrt{PP_s}\right), \tag{19}$$

where I_n is the modified Bessel function of order n.

It is important to realize that P_s in Eq. (19) is the total accumulated signal power. For the special case where the signal power s^2 in each normalized spectrum is identical, the total signal power is given by $P_s = ns^2$. The mean and variance of Eq. (19) then become

$$\langle P \rangle = n(1 + s^2) \quad \text{and} \quad \langle P^2 - \langle P \rangle^2 \rangle = n(1 + 2s^2). \tag{20}$$

The average power in the frequency bin containing the signal increases by ns^2 over the mean power in the bins with no signal, and the distribution spreads as $\sqrt{n(1+2s^2)}$. For small signals, this variance agrees with that found in Eq. (16) for the distribution of frequency components containing no signal.

The error of the second kind, $\mathcal{P}_2$, can be obtained by integrating the probability distribution of signal plus noise from zero to the threshold P_T:

$$\mathcal{P}_2 = \text{Prob}(P < P_T) = \int_0^{P_T} p_n(P; P_s)\, dP. \tag{21}$$

The error is a function of the signal strength and becomes negligible when the signal power is large compared with the average noise power. Shown in Fig. 2 is the probability that a real signal is below a threshold (and therefore missed) as a function of the difference between the signal and threshold *amplitudes*. Two curves are drawn, for thresholds of 1σ and 5σ (where a threshold amplitude $x\sigma$ corresponds to a normalized power $P = x^2/2$). The curves are valid for a single spectrum ($n = 1$), and similar curves could be given for averages of n spectra. In the case $n = 1$, the probability of an error of the second kind is down to 1% for signals about 1.5σ above the threshold.

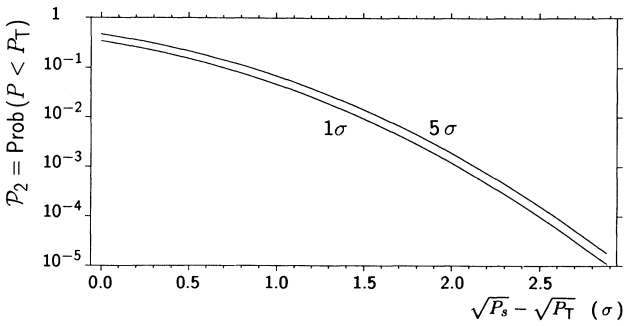

FIG. 2. Probability $\mathcal{P}_2$ of missing a signal due to noise (error of the second kind) for a single spectrum. This error depends strongly on the amount by which the signal amplitude $s = \sqrt{P_s}$ exceeds the threshold $\sqrt{P_T}$, but is almost independent of the value of the threshold itself ($\sqrt{P_T} = 1\sigma$, 5σ).

IX. THE NARROW-BAND PULSAR SEARCH ON THE GARCHING DATA

The CHT data compression was applied to the 100 h of data taken with the Garching interferometer. First, the original 10 kHz sampled data were complex heterodyned (sine and cosine) using two independent frequencies of 1968.658 and 3937.316 Hz. The channels were low-pass filtered as described above and resampled at 4 Hz. Each channel, after data compression, consisted of about 1.5 million complex time samples over the 100 h. A schematic diagram of the data analysis is given in Fig. 3.

In the course of the analysis, several different time scales were used for the integration time of the FFT. We made the evaluation first with many 9 min time slices. The length of the FFT was then increased to 1, 9, 36, and finally 100 h. The shorter integration periods have the advantage that the corrections due to Doppler shift or changing antenna sensitivity are less important. In addition, the problems of inter-tape gaps and changing sensitivity could be avoided using time scales less than or equal to an hour. Shorter FFT's are also better when searching for unknown signals whose strength or frequency is time dependent. Of course for truly periodic signals, the sensitivity is the best when the FFT is made over the entire 100 h.

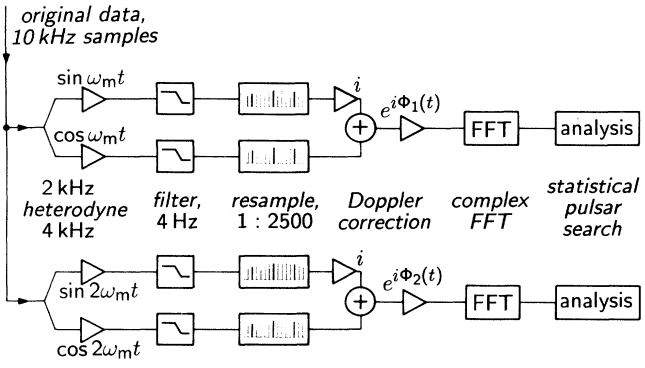

FIG. 3. Block diagram of the data analysis. The original data, sampled at 10 kHz, are processed in two separate channels, for the two heterodyne frequencies near 2 and 4 kHz. The data, resampled at 4 Hz, then undergo a channel dependent Doppler shift phase correction by $\Phi(t)$ before their complex Fourier transforms are analyzed for possible signals. The triangular symbols signify multiplication with the (time dependent) values indicated.

In general, the noise was found to decrease as the square root of the integration time for single FFT's and no unexpected difficulties were encountered with the introduction of gaps, Doppler shift correction, or demodulation for the two orthogonal polarizations. Therefore, we will present only the results of the 9 min and 100 h integration periods.

X. 9 MINUTE INTEGRATION TIMES

The 100 h data stream after data compression consisted of a complex time series sampled at 4 Hz. First we artificially subdivided the data into many time slices made up of 2048 points which corresponds to a real time of 8.53 min. The number of resampled data points was chosen to be an integer power of 2 which facilitated use of the standard complex FFT subroutine CRFFT2 available on the Cray.

This time period was short enough that the effect of Doppler shift and changing sensitivity to linearly polarized gravitational waves could be neglected. The time slices were chosen so that no gaps were included. Furthermore, the noise level of the interferometer remained roughly constant over this time scale.

A. Whitening the spectrum

The 9 min spectra produced from the compressed data clearly show the attenuation from the digital low-pass filters. In Fig. 4(a) we show the average amplitudes of 100

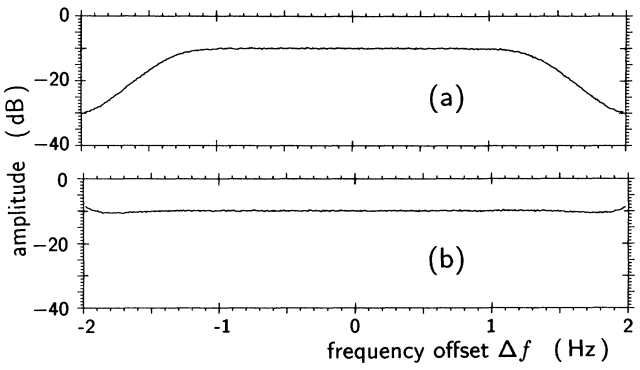

FIG. 4. (a) Amplitude spectrum of the compressed data, in the 2 kHz channel, averaged over 100 spectra. Amplitudes (in arbitrary units) are shown in logarithmic scale (dB), versus frequency offset from heterodyne frequency, Δf, in Hz. (b) Same spectrum, after *whitening*. See text.

power spectra for the 2 kHz channel. The spectrum is centered on the heterodyne frequency and has a bandwidth of 4 Hz. These spectra can be "whitened" by multiplying with the inverse of the calculated filter attenuation of the last two Butterworth filters. (The initial boxcar filter did not contribute significantly to the attenuation in this bandwidth.) The whitened spectrum Fig. 4(b) is frequency independent except in the extreme edges where some extra noise due to aliasing can be seen. The magnitude of this noise excess agrees with the estimated aliasing from the digital filter transfer functions.

As a check on this hypothesis we compared the spectra made with two sets of heterodyne frequencies (offset by about 1 Hz). In the overlap regions near the center of the bands we found a correlation of 1.0 which decreased toward the extreme edges of the bands. The poorer correlation near the edges was due to different noise components which were aliased by the offset heterodyne frequencies. The results again were completely consistent with the calculated aliasing produced by the digital filters. The remarkable agreement between the whitened spectrum and the original data is a clear reflection of the excellent filter response that can be obtained with IIR digital filters.

The statistical evaluation of the resampled power spectrum is simplest if the aliased components near the edges of the bandwidth are not used. As a conservative approach we discarded 5% of the components on each side. The remaining 3.6 Hz centered on the 2 and 4 kHz channels were then searched for possible signals as will be described below.

B. Statistics for individual spectra

A histogram was made for the normalized power spectrum of each independent 9 min time slice. The bins were chosen equally spaced in power from zero to a maximum value P_{M}. The last bin contained all frequencies with power greater than or equal to P_{M}. As discussed above, white noise, Gaussian in the time domain, should lead to a power spectrum with independent components distributed as a negative exponential. Thus, a graph of the natural logarithm of the bin content against the power for each bin should lead to a straight line with a slope of -1. The last bin, however, gives a measure of the cumulative probability to find a power larger than P_{M}. The statistics can be improved by accumulating the power spectra of many individual time slices in one histogram. In this case, the assumption is made that all frequency components in any time slice are statistically independent.

Shown in Fig. 5 are the statistics using 450 independent spectra taken from different 9 min time slices. The bottom and top curves show the statistics for the 2 and 4 kHz channels, respectively. All components were binned without regard to frequency or the time slice from which they were derived. The highest bin was chosen such that (without accumulation) on average one frequency component in any given time slice would fall into the last bin. The calculated distribution is graphed as a solid line, where the last point represents the theoretical cumulative probability for $P \geq P_{\mathrm{M}}$. The measured statistics

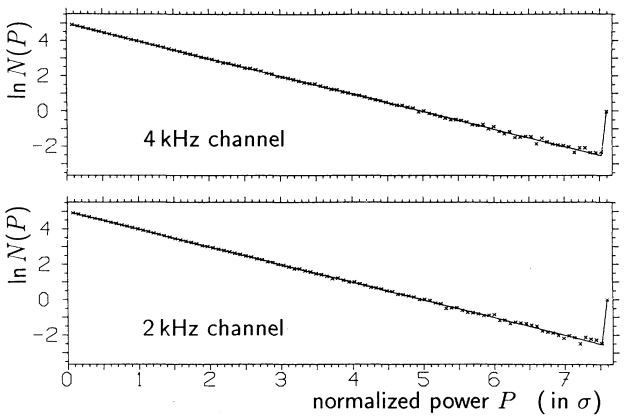

FIG. 5. Power histogram from 450 spectra (of 9 min each): natural logarithm of bin content, $\ln N(P)$, against normalized power P. The bin contents $N(P)$ follow closely the theoretical distribution, as borne out by the straight line with slope of -1 and the accumulated content for $P \geq P_{\mathrm{M}}$ in the top bin.

are observed to coincide closely with the expected probability distribution. Even the number contained in the last bin agrees very well with the assumption of random noise components. Note, that if there were a real signal large enough to fall in the last bin in each spectrum, the content of the last bin would have been significantly higher than statistically allowed.

It is interesting to display as a function of frequency those components of the power spectrum that fall into the highest bin of the histogram. An unusual (nonwhite) distribution of these high amplitude events could indicate the presence of a signal containing several frequency components even if no individual signal component were statistically significant. For example, a pulsar in a binary system (as was indeed suggested by Kristian *et al.* [1]) would have a frequency modulation due to the Doppler shift that could give a cluster of frequency components centered on the pulsar frequency.

Figure 6 graphs the power of all components falling into the last bin against the associated frequency of the component. The plots for both the 2 and 4 kHz channels are displayed. All power spectra were normalized so that the mean value was unity before accumulating the amplitude statistics. The mean of all the power spectra corresponds to an average linear strain amplitude of 2.5×10^{-20} for an observation time of 8.5 min.

There is no apparent clustering near any particular frequency. Furthermore, there is no power greater than 14 times the mean value. In linear amplitude, this translates into a strain limit of about 9.4×10^{-20}. The statistical significance of any single peak with amplitude P can be calculated from the Poisson distribution Eq. (18) for finding one event out of the available components (in this case 204 800) each with an individual probability of e^{-P}. For example, the probability of finding a frequency with power greater than 15 is about $\frac{1}{17}$.

We again caution, however, that to change the sim-

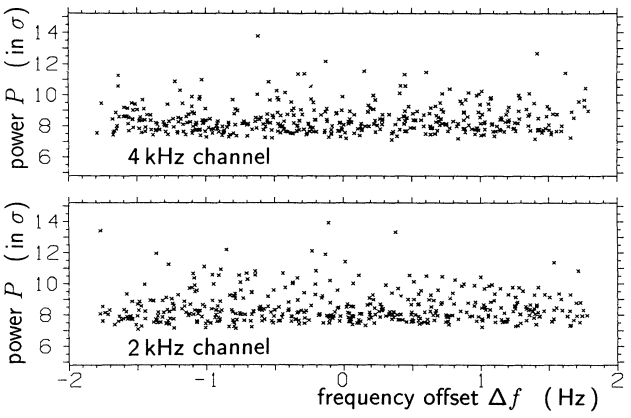

FIG. 6. Scatter plot for events in the top bin ($P \geq 7.5\sigma$): power P versus frequency offset Δf, for the 4 and 2 kHz channels. No clustering at a particular frequency (indication of a signal) is observed.

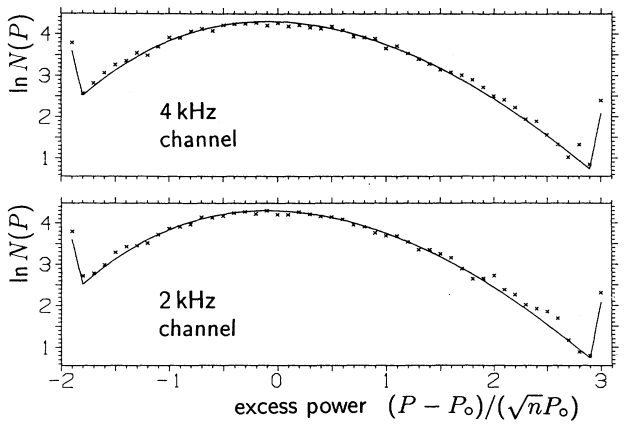

FIG. 7. Power histogram from a total of 450 of the Doppler-corrected 9 min spectra. Natural logarithm of bin content, $\ln N(P)$, against normalized "excess" power, $(P - P_0)/(\sqrt{n}P_0)$.

ple statement that no signals were found above a given threshold into a hard limit on signal strength one must state the associated error of the second kind for any given threshold. This error is about 50% for signals near the levels quoted above and falls off quickly for larger signal strengths.

The above spectra were made without including the Doppler correction, mainly because its effect on 9 min time slices is still tolerable. We note, however, that we also repeated the search after including the Doppler shift for SN 1987A and obtained similar null results.

C. Combining power spectra

The SNR of the above search was limited by the noise in a given frequency bin of a single 9 min time slice. Knowledge about the relationship of the signal in different time slices can be used to increase the detectability of the signal by combining the power spectra in the appropriate way.

For example, if the *frequency relationship* of the signal between various time slices is known, *power* spectra may be added to increase the signal-to-noise ratio. As was discussed in Sec. VIII A 1, the SNR increases, in linear measure, only as the fourth root of the observation time.

On the other hand, if the *phase relationship* of the signal between different time slices is used, the signal *amplitudes* can be accumulated and then the linear SNR will increase with the square root of the observation time. A straightforward way of accomplishing this for periodic signals is to make a single large FFT over the entire 100 h, as will be discussed in Sec. XI.

However, before changing the integration time of the FFT, it is still interesting to compute the statistics resulting from the addition of many 9 min spectra. This leads to increased sensitivity for signals with a bandwidth less than about 2 mHz ($\frac{1}{9\,\mathrm{min}}$). Furthermore, we can test the statistical model which assumes that frequency components from neighboring time slices are uncorrelated. In this case, we would expect that the average of n power spectra should be distributed as a χ^2 with $2n$ degrees of freedom [Eq. (15)]. This model can be checked directly.

In order to combine spectra in a meaningful way, we included the Doppler correction for a source located in SN 1987A. This was accomplished using the second complex heterodyne on the reduced (resampled) data set described above. Shown in Fig. 7 are the statistics obtained by summing 50 independent power spectra. Nine independent sets of 50 summed spectra were binned to increase the number of statistically independent frequencies.

The agreement between the observed statistics and the calculated χ^2 distribution is again quite good. This provides another demonstration that the statistical laws discussed above are appropriate. The limit on strain improved only marginally over that obtained using single 9 min slices. This is due mainly to the observed increase in noisiness as the experiment went along, but also to the slow improvement in sensitivity with only the fourth root of the observation time. We can, however, gain much more in sensitivity by making an FFT over the total 100 h of data.

XI. 100 HOURS INTEGRATION TIME

The procedure for handling the 100 h time series was more complicated than the 9 min analysis for several reasons. The 100 h stream contained gaps and had a noise level which varied over time. Furthermore, one can no longer neglect the Doppler shift correction and the correction due to the changing sensitivity of the antenna to linearly polarized gravitational waves for long integrations. A minor technical change was also made in the FFT algorithm so that the length was not constrained to an integer power of 2, which otherwise would have limited the integration to only 73 h. This was made by

breaking the 100 h into a few contiguous pieces (each with length corresponding to a power of 2) and combining the smaller FFT's. This straightforward procedure has been described in detail by Hocking [12].

A. Gaps and changing noise levels

The 100 h data stream analyzed here contained gaps due to the data acquisition hardware. During these periods no data were taken. The most common gap occurred when the magnetic tape was rewound and changed. This took typically 5 min every hour. There were also two unforseen failures in the data taking hardware. A locked tape hub caused a 98 min interruption and a blown fuse in the computer gave a down time of 48 min. There were no other gaps longer than 10 min. The total data loss from these sources amounted to about 9%.

In addition to times when the data acquisition was halted, there were also periods where the apparatus was either out of lock or had noise levels which were too great to be included in the data analysis. The out-of-lock periods occurred very seldom and could to a large extent be traced to argon refills of the laser. These gaps were clearly marked in the "housekeeping data" and rarely lasted longer than a minute in a 1 h tape.

Other gaps were artificially introduced in the post processing whenever the noise level rose by more than 9 dB above that obtained with good beam alignment. This threshold served to remove short spikes which were not large enough to throw the system out of lock, but were clearly times of instability. There were also larger sections which were removed with this threshold when the system performance was stable but extremely poor due to misalignment of the interferometer. The total amount of data removed by this procedure (16%) together with the gaps introduced by the data taking hardware (9%) totaled about 25 of the 100 h.

Even after the very poor data were removed, the output noise level still varied slowly over time. Since the electronic output is a calibrated measure of the path length difference in the interferometer, the sensitivity for real GW signals is constant even if the noise level changes due to varying fringe contrast. In order to obtain optimum SNR for real GW signals it is optimum to weight the data so that periods of low noise make the highest contribution to the overall SNR. It can be easily shown that for data containing a periodic signal and varying noise, the best SNR is realized if the data are weighted by the inverse square noise before the Fourier transform is applied. This factor is commonly used for weighted least squares [23].

The squared noise level was determined by integrating the sum of the squares of the cosine and sine heterodyne quadratures of the resampled data (4 Hz) over 2048 resampled data points (about 8.5 min). The inverse of this average was used as a weighting factor before the final FFT was applied.

The gaps and the time dependent weighting factor together form a *window* function which has the value of zero during gaps and is equal to the weight otherwise. This window multiplies the signal in the time domain which results in a convolution of the signal and window Fourier transforms.

For sinusoidal signals, the single frequency component is split into sidebands with amplitudes given by the Fourier transform of the window function. As long as the window function is dominated by the dc term, the effect of the window is a simple scaling by the dc amplitude.

The effect of a window function on the noise is more difficult to describe because the resulting noise must be described by a nonstationary process. Assuming that the noise before application of the window is uncorrelated and normalized with a variance of unity, the correlation function of the noise $N(t)$ after multiplication with the window function $W(t)$ can be written:

$$\langle N(t)\,N(t')\,\rangle = W^2(t)\,\delta_{t,t'}\,, \tag{22}$$

where the angle brackets denote an ensemble average. This equation verifies the intuitive notion that the resulting noise is given by the original noise amplitude modulated by $W(t)$.

We have shown [9] that the power spectrum of such a noise process remains white, with an amplitude given by the rms time average of the window function. Thus, even though the windowed noise is not stationary in the time domain, the noise is still "white" in the frequency domain. This means that a window function will not produce features mimicking periodic signals, even if the the window contains periodic gaps.

Amplitude modulated noise, on the other hand, may have unequal quadrature components at subharmonic frequencies ($\frac{\omega}{2}$) found in the Fourier transform of the squared window function [9]. This means that spurious signals can be produced in subsequent matched filtering that is both phase and frequency sensitive. However, pul-

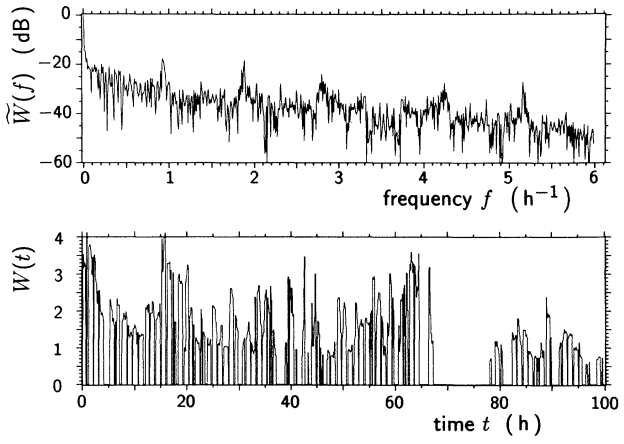

FIG. 8. Time series $W(t)$ (lower plot, linear) and Fourier transform $\widetilde{W}(f)$ (upper plot, logarithmic), of the "window function". The abscissas are linear: time t in hours, frequency f in cycles per hour. This window function $W(t)$ is used for weighting the (resampled) input data stream. $W(t)$ is normalized to have unity time average and a unity dc term $\widetilde{W}(0)$ in the spectrum.

sar searches with unknown phase of the signal are phase insensitive, and in this case one does not have to worry about spurious signals created by the action of a window function on random noise. Furthermore, as long as the Fourier transform of the window function is dominated by the dc component, any spurious features of this nature will be quite small.

Figure 8 shows the time dependence and the FFT of the window function described above, including gaps and the weighting proportional to the inverse squared noise. The bottom curve shows the window in the time domain normalized so that the average over time is unity. This curve shows that the average noise level in the original data was almost twice that obtained during short periods when the interferometer produced the low noise levels. As discussed above, this is believed to have been caused by misalignment in the beam steering. The upper curve shows the FFT of the window. Clearly recognizable are the harmonics produced by the periodic gaps when the tape was changed, on average about once every 64 min. Note also that the component at dc is more than 15 dB larger than that found in any other frequency bin.

B. Corrections due to earth motion

1. *FM demodulation — Doppler shift*

The complex data stream produced by heterodyning, filtering, and resampling was corrected for the Doppler shift for signals originating in the direction of SN 1987A with a frequency equal to the heterodyne frequency. This correction was made on all the data once using the second complex heterodyne with the complex demodulation function $\exp[i\omega_m \delta t(t)]$, the time delay $\delta t(t)$ given by Eq. (3). The constant and first derivative terms were removed since they change only the phase and frequency offset of the observed signal but do not modulate the signal.

2. *Amplitude demodulation — Polarization correction*

A GW signal of arbitrary polarization can be decomposed into a superposition of the two orthogonal polarizations usually denoted h_+ and $h_\times$. As described above, a rotation of the antenna in the plane perpendicular to the propagation direction gives rise to amplitude modulated signals in the interferometer. The combined signal from these two contributions depends critically on the phase between the two polarizations as well as the location of the antenna and source. In some cases (e.g., circularly polarized gravitational waves and the detector situated on the north pole) it is best to look at the mixture. In most cases, however, the best signal detection algorithm is to search independently for the two signals by amplitude demodulation using the sensitivity function of the antenna to the desired polarization.

Thus, we elected to make three different searches. Namely, first we searched for a signal by making only the Doppler shift correction but no demodulation for the linear polarization. This was followed by two independent searches applying the amplitude modulation given by Eqs. (9) and (10). In all three cases, the power spectrum was "whitened" before searching for threshold crossings in the middle 90% of the 4 Hz bandwidth as was described in the 9 min analysis.

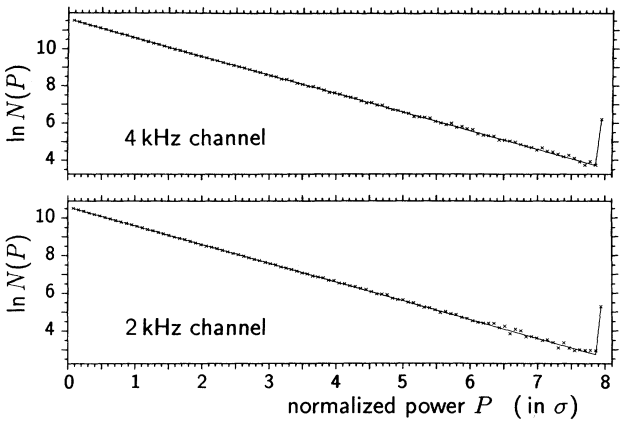

FIG. 9. Power histogram of a 100 h FFT of weighted, Doppler-corrected data. Again, it shows excellent agreement with theory (slope; contents of last, accumulative, bin).

The whitened spectra were featureless and had amplitude statistics which followed very closely those predicted by the Rayleigh distribution. The mean strain amplitudes were 1.26×10^{-21} for the case where first no polarization was extracted, and 2.18×10^{-21} or 2.22×10^{-21} for plus or cross polarization, respectively.

Shown in Fig. 9 are the amplitude statistics for the case where no polarization was extracted. The top and bottom curves correspond to the 4 and 2 kHz channels, respectively. The components falling into the last bin form the basis of the scatter plots shown in Fig. 10. A normalized power level of $P/P_0 = 17$ on the y axis cor-

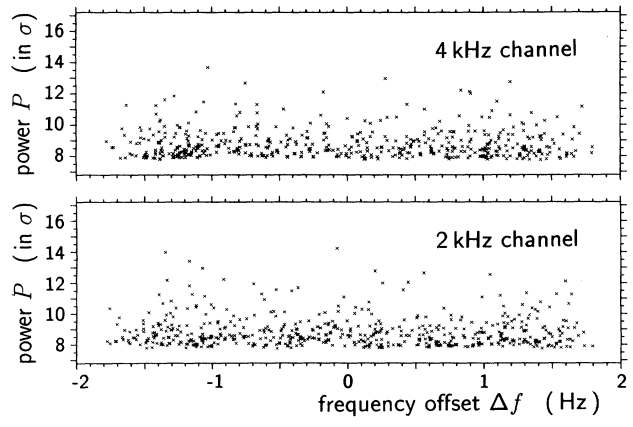

FIG. 10. Scatter plot for events in the top bin ($P \geq 7.9\sigma$) in the 100 h spectrum: power P vs frequency offset Δf, for the 4 and 2 kHz channels. Here also, no clustering at specific frequencies can be observed.

responds to a 95% confidence level that no *false alarms* due to the noise would occur in any one of the 1 441 792 bins. Similar graphs were obtained for the + and × polarization.

No significant signals above the level of $P/P_0 = 17$ were found. This threshold corresponds to limits which are larger by a factor of 4 on the strain amplitude than the mean values quoted above, or 9×10^{-21} for either polarization. In order to have a 99% confidence level for the error of the second kind (that of *losing* events), one must relax this limit by another 1.5σ, to 1.3×10^{-20}.

XII. CONCLUSIONS

We have introduced the complex heterodyne method which shows promise as a computer efficient technique for data reduction in the case where the detector bandwidth is greater than the expected signal bandwidth.

The efficiency of data reduction may facilitate a full sky search over the full bandwidth of the detector for long observation times, which would otherwise be ruled out because of the computational cost and memory limitations. The analysis can be accomplished by subdividing the full bandwidth into many smaller frequency channels which can be processed sequentially or in parallel.

The technique also allows for FM demodulation which accomplishes the Doppler shift correction on the reduced data set. This procedure avoids rather serious binning errors without requiring oversampling of the original data stream.

The data compression was applied to a 100 h data stream taken recently with the Garching prototype. The data volume was reduced by a factor of 1250 in a search for periodic signals within a 4 Hz bandwidth of two different frequencies near 2 and 4 kHz.

The observed amplitude statistics in the frequency domain were well described by the Rayleigh distribution for the case of single power spectra. The statistics also followed the more general χ^2 distribution when several power spectra were added. These results are expected for a noise source that is Gaussian in the time domain.

Finally, the data were searched using different integration times, partly to check that the different phases of the data analysis were operating properly, and partly to look for signals with unknown frequency modulation. The result of all these searches revealed no statistically significant periodic signal. The most stringent strain limit for periodic signals coming from SN 1987A was 9×10^{-21}.

ACKNOWLEDGMENTS

We wish to thank J. Hough, H. Ward, and B. Schutz for discussions which stimulated this pulsar search. We also thank H. Ögelman for his help with the calculation of the barycentric time correction. This work was supported by the Bundesministerium für Forschung und Technologie (BMFT).

APPENDIX A: DIGITAL FILTERING

The quality of the digital filtering generally limits the resampling rate due to aliasing. The process of resampling can be viewed in the frequency domain as a convolution of the original data (after digital filtering) with an infinite sequence of delta functions at integer multiples of the sampling frequency [24]. Thus, any frequencies above the Nyquist frequency for the resampled data stream which are not adequately removed by the digital filter are aliased to low frequencies. This adds additional noise to these corrupted frequencies.

The ideal filter shape is, of course, square in the frequency domain with a cutoff at the Nyquist frequency of the resampled data stream. In practice this cannot be achieved and any deviations lead to aliasing. Naturally, the computational cost of the digital filter necessarily increases as the quality of the filter increases so that one must find a compromise between good filtering and computational effort. We will indicate how such filters can be designed and implemented, in order to get a rough idea of the computational requirements for the three data compression schemes (SQHT, AT, and CHT) discussed in Sec. V B.

There are two broad classes of digital filters called the finite impulse response (FIR) filters and the infinite impulse response (IIR) filters. The transfer function of an FIR filter of length N has a z transform given by [25]

$$H(z) = \sum_{k=0}^{N-1} a_k z^{-k} . \tag{A1}$$

The constants a_k can also be thought of as simple weighting factors for a moving average in the time domain. The simplest example of a low-pass filter is the *boxcar* moving average where all weighting factors are equal to $\frac{1}{N}$.

The more general class is the IIR filter, with a transfer function having both zeroes and poles. Its z transform is given by

$$H(z) = \frac{\sum_{k=0}^{M-1} a_k z^{-k}}{1 + \sum_{k=1}^{N} b_k z^{-k}} , \tag{A2}$$

where $M \leq N$. The IIR filters use both past inputs and outputs to generate new output values. The feedback can result in an unstable filter if the coefficients are not selected properly.

In practice, it is quite easy to design stable IIR filters by using standard analogue filter designs and transforming them into digital filters. This requires a relation between the variable z in the z transform and s in the Laplace transform. A design which produces excellent results uses the *bilinear transform* defined by

$$z = \frac{1+s}{1-s} . \tag{A3}$$

Typically the IIR filters require far fewer terms in the z transform than the FIR filters for a given set of filter specifications. This means that the IIR filters can be realized in the time domain with few operations and very small memory requirements. The IIR filters avoid the need to make a lengthy FFT and inverse FFT for every output data point.

An IIR filter can be efficiently programmed using many

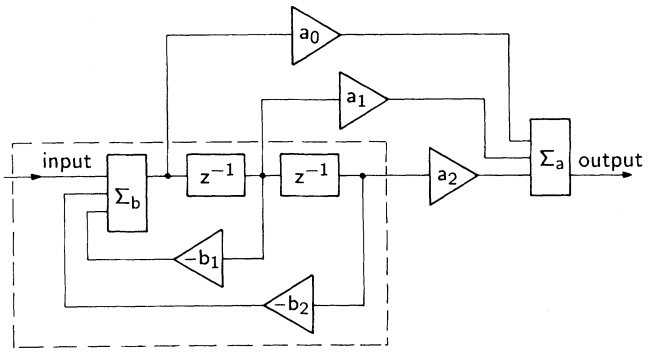

FIG. 11. Schematic diagram of the canonical realization for IIR filters. The boxes denoted by z^{-1} represent a time delay by one sample. The triangular units are multiplications with the coefficients a_k and b_k in Eq. (A2). The dashed box contains the input and update operations that are required for each new data point, whereas the other operations (a_k, Σ_a) are carried out only in those steps where filter output is required.

different schemes. Fig. 11 shows the *canonical realization* for a two-pole IIR filter. In general, this technique uses $n+1$ memory cells and requires $2n-1$ multiplications and additions per input data point, where n is the number of poles in the filter.

For this pulsar search, we used low-pass IIR digital filters designed using computer code given by Cappellini [26]. We rewrote these routines to make better use of the Cray's vectorization. We tested these IIR filters using random data and also using several artificial signals with many frequency components (e.g., a frequency chirped signal and also one made containing a sum of all sinusoidal terms in a Fourier expansion). In every case the measured transfer behaved quite well and followed the classic Butterworth attenuation function, Eq. (2).

In this paper, IIR low-pass filters were used for the purpose of antialiasing before the data were resampled at a lower rate. Thus, the filter output could be calculated at a lower rate than the input data stream. As shown in Fig. 11, the canonical realization uses half of the operations to update the memory and the other half to calculate the output. This means that for input cycles which require no filter output, the number of operations per input data point are reduced by a factor of 2 to n multiplications and additions for an n-pole IIR digital filter. These operations are indicated for a two-pole IIR filter inside the dashed box in Fig. 11.

APPENDIX B: TESTING THE CHT

We tested our procedure for data compression using the CHT and the effectiveness of removing phase modulation with a second complex heterodyne on the reduced data set in many ways. The test described in this appendix used a data set consisting of four artificial periodic signals added to the real interferometer data with a sampling rate of 10 kHz. Three of the "fake signals" were sinusoidal terms with equal amplitudes of 20 dB (in arbitrary units) with the frequencies 1969.0, 3936.0, and 3938.0 Hz. Arbitrary phases were also assigned to each signal. The fourth artificial signal (at the lowest frequency, 1968.0 Hz) was phase modulated with modulation index equal to 3.0 and a modulation frequency of 40 mHz.

The data were then complex heterodyned in two channels at 1968.658 and 3937.316 Hz. These complex data were low-pass filtered and resampled at 4 Hz using the same programs which were used to compress the data for the pulsar search.

The spectra for the 4 and 2 kHz channels are shown in the top two curves in Fig. 12. Each spectrum consists of 2048 compressed complex values corresponding to about five million original data. The total bandwidth is ±2 Hz centered on the corresponding heterodyne frequency. We can see that the four signals have been shifted to the appropriate frequencies by the complex heterodyne and that the amplitudes are also correct. The frequency modulated signal was shifted to −0.658 Hz relative to the 2 kHz heterodyne and occupied about 0.3 Hz bandwidth due to the sidebands caused by the modulation.

A second complex heterodyne was made by multiplying the 2 kHz complex *compressed* data stream with $\exp[i\phi(t-t_d)]$, where $\phi(t)$ matched the input phase modulation applied to the lowest signal frequency. The time t_d was introduced to account for the delay caused by the digital filter. We found that a constant time delay of

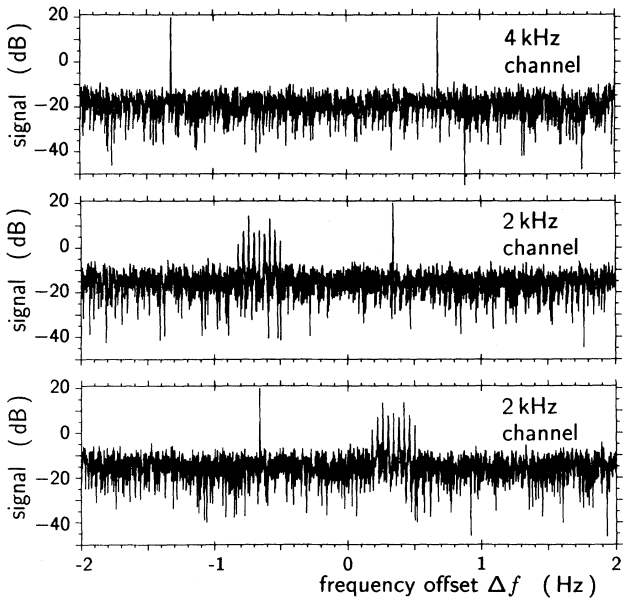

FIG. 12. Artificial signals added to the digitized interferometer data. Top curve: Spectrum of the 4 kHz channel of compressed interferometer data, with two monochromatic signals added. Middle curve: Spectrum of the 2 kHz channel, with one sinusoidal, one modulated signal. Bottom curve: The 2 kHz channel after demodulation ("Doppler correction"): the sidebands of the modulated signal are recombined into one peak; the sinusoidal signal becomes spread out.

about 0.6 s was necessary to give good demodulation of the signal. Note, however, that this delay is negligible compared with the time scales that are important for the Doppler shift corrections made in the pulsar search. More important than the size of the constant time delay of the digital filter, however, was the fact that excellent demodulation was obtained without having to introduce a complicated frequency dependent phase shift.

The "demodulated" spectrum is shown in the bottom curve of Fig. 12. It can be seen that the demodulation was very effective since all the energy in the sidebands has been reconverted into a single frequency component having the same amplitude (20 dB) which it was originally given. Thus, this procedure results in almost no loss in SNR for modulated signals even when they are offset from the heterodyne by a substantial fraction of the compressed bandwidth. As a consequence of the demodulation process the originally unmodulated frequency at 1969 Hz (0.342 Hz above the heterodyne frequency) now appears modulated.

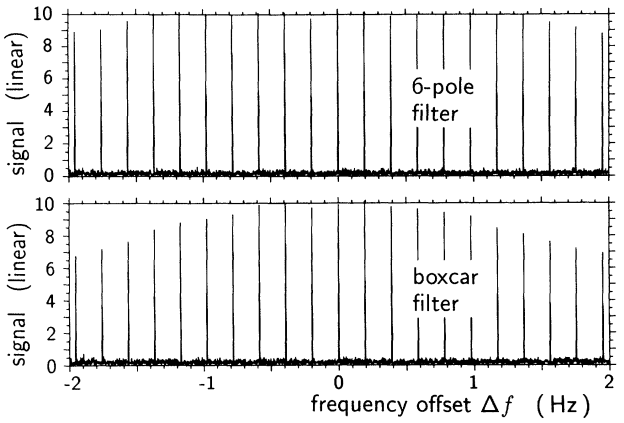

FIG. 13. Loss in SNR as function of frequency offset Δf, compared for two filtering schemes. Top curve: the CHT as described in this pulsar search (six-pole filter), applied to original interferometer data to which signals of equal strength (+10 in arbitrary units) had been added at equidistant frequencies. Both axes in linear scale. Bottom curve: a CHT, using, however, a simple boxcar average filter, applied to the same data, showing a more pronounced loss of signal as one nears the band edges.

APPENDIX C: THE CHT WITH SIMPLER FILTERS

The CHT can be used with a wide variety of filtering and resampling schemes. The choice used for this pulsar search was made rather arbitrarily as a compromise between computing effort and effective antialiasing. One can make more effective use of the bandwidth by making the low-pass antialiasing filter sharper.

It is also interesting, however, to see how it works using simpler filter-and-resample schemes. The most straightforward way to reduce the quantity of data (and also its bandwidth) would be to store fixed length averages without any overlap between the neighboring averages. This is equivalent to a boxcar average with no overlap. In the frequency domain, this looks like a sinc filter with the first zero occurring at the resampling frequency.

These filters work well near dc since all signal components which would be aliased to dc occur at the zeros of the sinc function. As one approaches the Nyquist frequency, however, the SNR is reduced according to the sinc function. This happens because the signal follows the sinc function whereas the noise turns out to remain constant, since aliasing from noise components above the Nyquist frequency compensate the drop due to the sinc function. On the other hand, this filter is quite simple to implement. Furthermore, the resulting background noise is white so that "whitening" by the inverse of the filter function is not necessary.

Shown in Fig. 13 is a comparison of the two techniques. The data were made using real interferometer output for background noise to which artificial sinusoidal signals were added. The top graph shows the response of the three-stage filter-and-resample scheme described for the pulsar search. The noise has been "whitened" by the inverse filter function. The bottom graph shows the outcome using a boxcar low-pass filter and resampling with no overlap. It is clear that this scheme works well close to dc but begins to lose in SNR as one nears the edge of the resampled bandwidth.

[1] J. Kristian, C. R. Pennypacker, J. Middleditch, M. A. Hamuy, J. N. Imamura, W. E. Kunkel, R. Lucino, D. E. Morris, R. A. Muller, S. Perlmutter, S. J. Rawlings, T. P. Sasseen, I. K. Shelton, T. Y. Steiman-Camero, and I. R. Tuohy, Lett. Nature **338**, 234 (1989).

[2] J. Hough, B. J. Meers, G. P. Newton, N. A. Robertson, H. Ward, G. Leuchs, T. M. Niebauer, A. Rüdiger, R. Schilling, L. Schnupp, H. Walther, W. Winkler, B. F. Schutz, J. Ehlers, P. Kafka, G. Schäfer, M. W. Hamilton, I. Schütz, H. Welling, J. R. J. Bennet, I. F. Corbett, B. W. H. Edwards, R. J. S. Greenhalgh, and V. Kose, "Proposal for a joint German–British interferometric gravitational wave detector," Max-Planck-Institut für Quantenoptik, Report No. MPQ 147, 1989.

[3] A. Brillet, A. Giazotto, and M. Jacqemet, "VIRGO, Final Conceptual Design," Proposal to CNRS/France and INFN/Italy, 1992.

[4] R. E. Vogt, R. W. P. Drever, F. J. Raab, K. S. Thorne, and R. Weiss, "Laser Interferometer Gravitational-Wave Observatory (LIGO)," Proposal to NSF, California Institute of Technology, 1989.

[5] H. Ward, "Workshop contribution to GR12," reviewed by J. Hough, in *General Relativity and Gravitation*, Proceedings of the Twelfth International Conference on General Relativity and Gravitation, Boulder, Colorado, 1989, edited by N. Ashby, D. F. Bartlett, and W. Wyss (Cam-

bridge University Press, Cambridge, England, 1990).
[6] R. Schilling (in preparation).
[7] B. Schutz (private communication).
[8] D. Shoemaker, R. Schilling, L. Schnupp, W. Winkler, K. Maischberger, and A. Rüdiger, Phys. Rev. D **38**, 423 (1988).
[9] T. M. Niebauer, R. Schilling, K. Danzmann, A. Rüdiger, and W. Winkler, Phys. Rev. A **43**, 5022 (1991).
[10] H. Ögelman, C. Gouiffes, T. Augusteijn, H. Pedersen, F. Gutierrez, G. Hasinger, J. Melnick, W. Pietsch, and C. Sanyini, Astron. Astrophys. **237**, L9 (1990).
[11] J. Kristian, Lett. Nature **349**, 747 (1991).
[12] W. K. Hocking, Comput. Phys. **3** (1), 59 (1989).
[13] J. Livas, in *Gravitational Wave Data Analysis*, edited by B. F. Schutz (Kluwer, Dordrecht, 1989), pp. 217–238.
[14] J. Livas, in *The Detection of Gravitational Radiation*, edited by D. Blair (Cambridge University Press, Cambridge, England, 1991).
[15] This technique has been proposed by Norman McKenzie.
[16] P. Stumpff, FORTRAN code EARTH, export version 3.0, Max-Planck-Institut für Radioastronomie, 1981.
[17] S. Smith, Phys. Rev. D **36**, 2901 (1987).
[18] P. C. Peters, Phys. Rev. **136**, B1224 (1964).
[19] K. S. Thorne, in *300 Years of Gravitation*, edited by S. W. Hawking and W. Israel (Cambridge University Press, Cambridge, England, 1987).
[20] R. L. Forward, Phys. Rev. D **17**, 379 (1978).
[21] P. F. Panter, *Modulation, Noise, and Spectral Analysis* (McGraw-Hill, New York, 1965).
[22] E. J. Groth, Astrophys. J. Suppl. **29**, 285 (1975).
[23] P. R. Bevington, *Data Reduction and Error Analysis for the Physical Sciences* (McGraw-Hill, New York, 1969).
[24] E. O. Brigham, *The Fast Fourier Transform* (Prentice-Hall, Englewood Cliffs, New Jersey, 1974).
[25] R. A. Roberts and C. T. Mullis, *Digital Signal Processing* (Addison-Wesley, London, 1987).
[26] V. Cappellini, A. G. Constantinides, and P. Emiliani, *Digital Filters and Their Applications* (Academic, London, 1978).