BROADBAND SEARCH TECHNIQUES FOR PERIODIC SOURCES OF GRAVITATIONAL RADIATION

Jeffrey Livas
M. I. T. Physics Department
Building 20F-001
Cambridge, MA 02139
USA

ABSTRACT. Full exploitation of the potential of a laser interferometric gravitational wave detector to detect sources emitting periodic gravitational waves will require the development of specialized search techniques and place unprecedented demands upon computer memory and speed. Traditional approaches to detecting small amplitude sinusoidal signals of unknown period in the presence of broadband noise must be modified to accomodate the unusual nature of the received signal from a gravitational wave source. This paper outlines two methods that have been developed and presents the results of an application of these methods to data taken with the M.I.T. 1.5 meter prototype. Based on these results, a combined hardware/software solution to the data analysis problem is needed before full scale gravitational wave astronomy with periodic sources becomes practical.

1. INTRODUCTION

1.1 The Ideal Search

The rapid developement of gravitational wave antennas has raised the possibility of a new branch of astronomy that could provide information about strongly interacting astrophysical systems that cannot be obtained any other way. One type of detector, the laser interferometric antenna, has two special properties: it has a broadband spectral response and a non-directional spatial response. Full use of the antenna to detect sinusoidal signals is thus a two dimensional problem because both the frequency of the signal and the direction from which it comes are unknown. The ideal search technique would be a method that could recover both the frequency and the location of the source.

The search is complicated by the fact that the relative motion between the source and the detector modulates the signal in a manner that depends explicitly on the location of the source, even for a source emitting monochromatic radiation. Traditional signal processing methods for extracting a sinusoid buried in broadband noise are not applicable without modification. In addition, the sheer volume of data generated by a wide bandwidth detector presents special problems for collection, storage, and analysis.

After a brief discussion of typical sources of periodic gravitational radiation and the model of the source adopted for this discussion, two different techniques for conducting a broadband search will be presented. The first is a method for performing an ideal two dimensional search, although with a finite limit on the sensitivity. The second method has no such limit on its sensitivity, but reduces the sky coverage to a small area around a single direction.

1.2 Possible Astrophysical Sources

Source of gravitational waves may divided into three basic categories: periodic, impulsive, and stochastic. Several recent reviews describe each of these categories in detail.[1,2] This discussion is concerned only with periodic sources, which may be divided into two (overlapping) classes: binary star systems and rotating neutron stars.

1.2.1 *Binaries.* A binary star system is perhaps the simplest of all periodic sources because it has an intrinsically time-varying quadrupole moment that is easy to calculate. Indeed, the famous binary pulsar PSR 1913+16 is the only physical system discovered so far in which gravitational radiation provides a conclusive explanation of the observations.[3] Furthermore, it has been estimated that more than half of the stars in the galaxy have at least one stellar companion. However, the shortest known binary period is ~ 685 seconds, or $\nu = 1.5 \times 10^{-3}$ Hz, which is much lower can be achieved with ground based gravitational wave antennas. The best present prototypes have useful sensitivities down to $\sim$ 100 Hz, and even the most optimistic predictions for future suspension designs estimate that ground motion will begin to dominate the spectrum below 10 Hz. Thus ordinary binaries are not an important source for ground based antennas.

However, a binary system may be a catalyst for an important high frequency periodic source. Wagoner[4] has proposed a model of a rotating neutron star in a close binary system in which accretion drives the rotation of the neutron star in a steady state in which the viscous damping timescale matches the timescale for the growth of gravitational wave instabilities. He estimates that monochromatic emission of gravitational radiation could occur with 200 Hz $\geq \nu \geq$ 800 Hz.

1.2.2 *Rotating Neutron Stars.* The best known example of this type of source is the pulsar. A pulsar is thought to be a rotating neutron star whose symmetry axis does not coincide with the axis of rotation. Such a system has a time changing quadrupole moment and could radiate gravitational waves. The fastest known pulsar spins at 642 Hz (1.5 msec), which is well within the bandwidth of an earth based detector.

There are approximately 450 known pulsars in the galaxy, and the total population has been estimated[5] as $70(\pm 17) \times 10^3$. The number of fast (msec) pulsars in the galaxy has been estimated[6] as 30. Thus, if pulsars produce gravitational radiation, there should be several sources.

1.2.3 *Source Model.* For the purposes of this analysis, the standard source assumed to be a simple narrowband sinusoidal emitter with a $\delta(f)$ frequency spectrum. The detection strategy is simply to integrate for as long as possible increase the signal to noise ratio (SNR). To see that a delta function source is a reasonable model, assume that the search bandwidth is 10 kHz and the integration time one year. Then if the delta function is to be a good approximation, the source must have a $Q \geq f_{max}/\Delta f_{res} \approx 3 \times 10^{11}$, where Δf_{res} is the frequency resolution bandwidth, $\sim 1/T_{int}$. Table 1 shows estimated Q's for a few periodic sources.
Taken at face value, the delta function source approximation is quite reasonable.

Table 1

Typical Q's of some periodic sources.

Process	Estimated Q
Binary systems	$2\pi\left(\frac{1}{8\,\mathrm{hr}}\right)(8\times10^{16}\,\mathrm{sec}) = 2\times10^{13}$
Pulsars (in general)	$2\pi\left(\frac{1}{.6\,\mathrm{sec}}\right)(10^{7}\,\mathrm{yr}) = 3\times10^{15}$
Pulsars (msec)	$2\pi(642\,\mathrm{Hz})(3\times10^{8}\,\mathrm{yr}) = 4\times10^{19}$
Wagoner's model	$2\pi(500\,\mathrm{Hz})(10^{7}\,\mathrm{yr}) = 1\times10^{19}$

2. ANALYSIS TECHNIQUES

Two periodic search techniques will be discussed, both based on the same strategy: estimate the power spectrum of the data and look for large peaks.

2.1 The Naive Approach

2.1.1 *Periodogram Estimate of the Power Spectrum.* Consider a one dimensional time series of K points $y_i(t) = y(t_i)$. The naive way to find a sinusoidal signal of unknown frequency in noise would be to construct a model of the time series as a sum of sinusoids:

$$y_{model}(t) = \sum_{m=0}^{M-1} A_m e^{-i2\pi f_m t}. \tag{1}$$

where the A_m are the unknown amplitudes. The model can then be fit to the data by forming the χ^2:

$$\chi^2 = \frac{1}{K}\sum_{k=0}^{K-1}\left(y(t_k) - y_{model}(t_k)\right)^2 \tag{2}$$

and then minimizing with respect to the A_m:

$$\frac{\partial\chi^2}{\partial A_m} = 0,$$

which leads to

$$A_n = \sum_{k=0}^{K-1} y(t_k)e^{-i2\pi f_n t_k}. \tag{3}$$

This is the expression for the Discrete Fourier Transform (DFT), which can be efficiently computed with the Fast Fourier Transform (FFT) algorithm if the data has been sampled at equally spaced intervals in time.

The periodogram is just the square of the DFT

$$P_N(m\Delta f_0) = \left|\frac{1}{N}\sum_{n=0}^{N-1} y(n\Delta t_0)e^{-i2\pi nm/N}\right|^2 \tag{4}$$

and is an estimate of the power spectrum. Here Δt_0 is the sampling period and $\Delta f_0 = (N\Delta t_0)^{-1}$ is the frequency resolution of the transform. For convenience the square root of the power spectrum $S_N(m) = (P_N(m))^{1/2}$ is usually used because it represents displacement and hence strain directly. In the discussion which follows, the terms periodogram and power spectrum are used interchangeably. When it is necessary to make the distinction, the square root of the power spectrum, $S_N(m)$, will usually be used.

The periodogram is useful because the signal to noise ratio (SNR) in power of a sinusoidal signal which appears in a single frequency bin of the transfrom increases directly with the length of the transform, or by the square root of the length in amplitude.

2.1.2 *Statistical Peak Detection.* The purpose of statistical peak detection is to provide a quantitative criterion for distinguishing possible signal peaks from chance fluctuations due to noise. Each frequency resolution bin m of the periodogram contains an estimate $s = S_N(m)$ of the the square root of the amount of power at that frequency in the data. The estimate fluctuates because of the presence of noise, and the value of the estimate may be regarded as a random variable. The idea behind statistical peak detection is that the probability distribution of $S_N(m)$ when a signal is present in the data is different from the probability distribution with just noise. Figure 1 shows an example of the differential probability distribution of $S_N(m)$ with and without a signal.

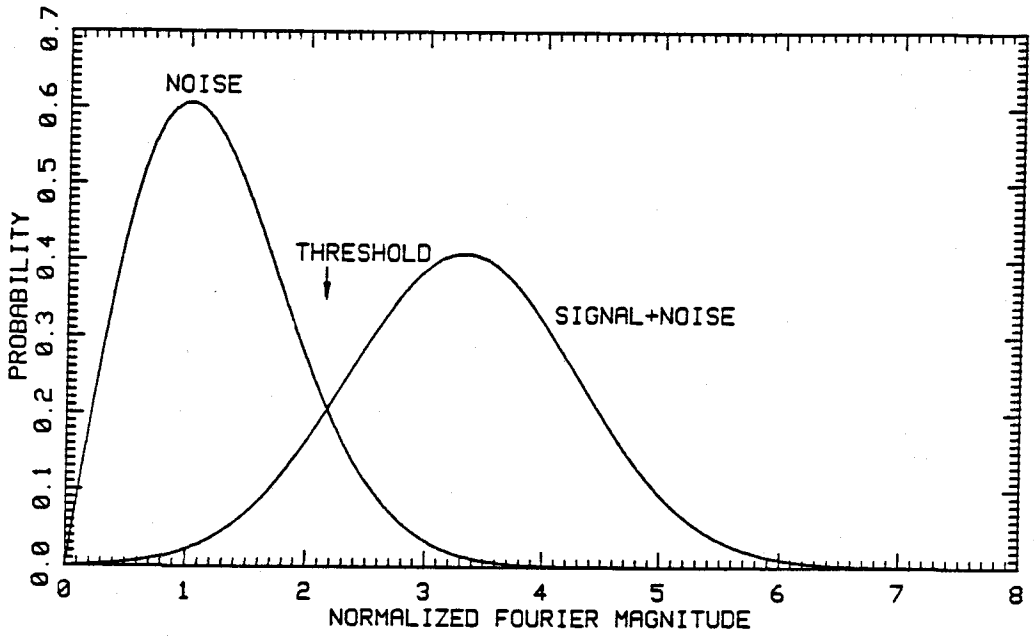

Figure 1

Illustration of the concept behind statistical peak detection. The probability distribution for the amount of power in a single bin of a periodogram changes when a signal is present. The threshold is chosen to minimize the false alarm probability.

With just noise, the $S_N(m)$ follow a χ^2 distribution with two degrees of freedom (a Rayleigh distribution). The presence of a signal means that there is a non

zero expectation value for one of the DFT coefficients, so the most probable value of that coefficient (i.e. the peak of the differential probability distribution) increases. The difference between the distribution with pure noise and the distribution with a signal plus noise can be conveniently summarized with a single threshold. Any frequency bin containing more power than the threshold is judged more likely to belong to the signal plus noise distribution than the noise distribution and is flagged as a possible signal. The probability of a false alarm, that is, the probability that a chance fluctuation of the noise exceeds the threshold, is just given by the area under the noise probability curve from the threshold to infinity:

$$p_{fa} = \int_{x_{th}}^{\infty} p_n(x')dx'. \tag{5}$$

Similarly, the probability of missing a signal, which is the probability that $S_N(m)$ is less than the threshold, is given by

$$p_{ms} = \int_{0}^{x_{th}} p_{s+n}(x')dx'. \tag{6}$$

The probability of detecting a signal is just

$$p_{det} = \int_{x_{th}}^{\infty} p_{s+n}(x')dx'. \tag{7}$$

In principle one would choose the threshold to maximize the detection probability while simultaneously minimizing the false alarm probability. In practice, the threshold is chosen to achieve an acceptable p_{fa}.

2.2 Problems with the Naive Approach

There are three fundamental problems with the naive approach. The first is that real data is not continuous and sampled at evenly spaced intervals in time as required by the FFT algorithm. The second is that the expected signals at the antenna, even for a purely sinusoidal source, are not pure sinusoids. The third problem is that the probability density function used for statistical peak detection depends on the variance σ^2 of the data, and σ in general is an unkown function of frequency.

2.2.1 *Real Data.* Real data is not a continuous stream of samples taken at evenly spaced intervals in time. The limitations of computer memory and hardware, as well as problems with the apparatus, combine to eventually cause gaps. The seriousness of a gap depends on its length. A data collection system which is triggered from an external reference clock will keep the phase of the samples across gaps, so that the data set will in general consist of blocks of data taken at equally spaced intervals in time separated by gaps of varying widths. A very small gap could be replaced by the average value of the data without much problem. A gap that is a large fraction of the length of the data stream must be dealt with in some other way.

In the MIT prototype data collection system, there are two guaranteed types of gaps. The first type of gap is the result of having only one 1600 BPI magnetic tape drive connected to the system. After 15 minutes, a tape is full and data collection

must be stopped while the tape is rewound and a new one is loaded. The gap is approximately 5 minutes. There is nothing fundamental about this type of gap. The 15 minute continuous data limit could be extended by purchasing a higher density drive, and two drives could collect data continuously.

The other type of gap is the result of ground motion induced by traffic on the street outside the laboratory. In this case the gap is approximately 16 hours, as the best time to collect data is between 22:00 and 06:00. A more remote location for the antenna and a better suspension system would increase the length of the available collection time. However, it is difficult to imagine collecting a month of continuous data without gaps of some sort.

The effect of the gaps is to limit the integration time and hence the sensitivity of the search by reducing the size of the FFT that can be performed. There is no possibility of calculating the DFT directly without using the FFT to circumvent the equally spaced samples in time requirement because the computation time required for the DFT scales as N^2. If the computation were performed on a machine capable of 10^8 operations per second, a $N = 2^{24}$ point transform (15 minutes of data on the MIT prototype) would require 2.8×10^6 seconds, or 32 *days* of cpu time. The FFT algorithm would require 4 seconds under the same conditions.

A more serious problem is the loss of information caused by the gaps. One way to think of the gaps conceptually is to consider a second time series $w(t)$ that is unity when data exists and zero in the gaps. The observed data set is then the product of a hypothetical continuous data set and $w(t)$. By the convolution theorem, a product in time space is a convolution in frequency space. The Fourier transform of the data is thus the convolution of the transform of the continuous data set and the trnsform$W(f)$ of the $w(t)$ function:

$$X_{obs}(f) = X_{cont}(f) * W(f). \tag{7}$$

$W(f)$ is called the window function. The effect of the window function is to smear out the information contained in the hypothetical continuous spectrum. For a general sequence of gaps, the smearing can be quite severe, and because the window function has zeros, deconvolution is not possible. Thus even if it were possible to perform the DFT directly, the resolution of the resulting spectrum would be degraded.

2.2.2 *Expected Signals.* The signal expected at the detector from a stationary, purely sinusoidal gravitational wave source is not sinusoidal. It is modulated by the relative motion of the source and the antenna in two distinct ways: FM and AM.

FM. The received signal is frequency modulated by the ordinary relativistic Doppler shift. The frequency observed by the antenna can be written explicitly as

$$f_{obs} = f_{emit} \cdot \gamma_0 \cdot (1 - \frac{\vec{v}}{c} \cdot \hat{r}), \tag{8}$$

where γ_0 is the usual relativistic parameter $(1 - (v/c)^2)^{-1}$, $\hat{r}$ is a unit vector from the antenna to the source, and $\vec{v}$ is the relative motion of the source and antenna. The two major contributions to the relative motion are the daily rotation of the earth, which causes a frequency modulation of $\Delta f/f \approx 10^{-6}$, and the orbital motion around the sun, for which $\Delta f/f \approx 10^{-4}$. In principle there is also a contribution from any motion of the source, but in the simplified model adopted for this discussion, source motion is neglected.

The effect of FM on the spectrum of the received signal is shown schematically in Figure 2. The emitted signal has a delta function spectrum centered at some frequency f_0. The FM spreads the power in that delta function out over a bandwidth of order $2(\Delta v/c)f_0$.

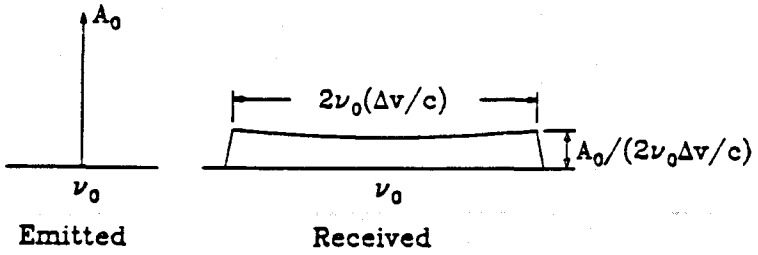

Figure 2
Schematic representation of the effect of FM on the spectrum of a sinusoidal signal.

The smearing of the signal by FM is significant only when the bandwidth over which it is spread becomes larger than the frequency resolution bandwidth of the transform. A rough idea of just when this occurs can be calculated using the following formula:

$$\underbrace{\underbrace{2f_{max}\frac{v_{rot}}{c}\frac{t_{int}}{1\,day}}_{\frac{\Delta v_{rot}}{c}} + \underbrace{2f_{max}\frac{v_{orb}}{c}\frac{t_{int}}{1\,yr}}_{\frac{\Delta v_{orb}}{c}}}_{\text{FM Bandwidth}} > \underbrace{\frac{1}{t_{int}}}_{\Delta f_{res}} \tag{9}$$

Solving for t_{int} and putting in numbers,

$$t_{int} \geq 30\,\text{min} \tag{10}$$

The Doppler effect is not important for data records shorter than 30 minutes. However, data collected at different times will contain the signal at different center frequencies.

AM. The amplitude modulation of the received signal is also caused by the relative motion of the antenna and the source. Although the spatial sensitivity of the antenna is non-directional, it is not isotropic. The amplitude modulation is simply the result of the different sensitivity lobes of the antenna sweeping across the source position as the earth moves.

The effect of the AM on the spectrum of the received signal is illustrated by Figure 3. The emitted signal, with a delta function frequency spectrum centered at frequency f_0, develops sidebands spaced at 1/24 hours (=12 μHz). The AM begins to exceed the resolution bandwidth of the transform for integration times longer than 6 hours.

2.2.3 *Unknown* $\sigma(f)$. The real and imaginary amplitudes of a single DFT follow an independent Gaussian distribution for each frequency. The distribution for

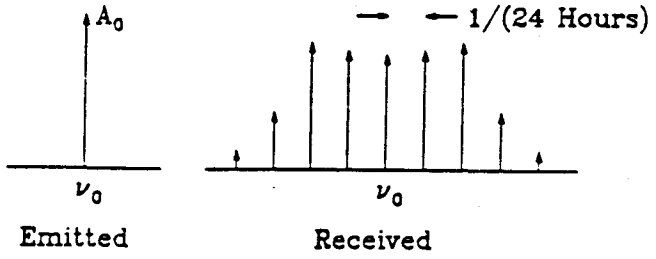

Figure 3
Schematic representation of the effect of AM on the spectrum of a sinusoidal signal.

the amplitude in a bin of the square root of the power spectrum, $s = S_N(m) = \sqrt{(real)^2 + (imag)^2}$, is given by

$$p_{noise}(s) = \left(\frac{s}{\sigma^2}\right) e^{-\frac{s^2}{2\sigma^2}}. \tag{11}$$

This is a χ^2 distribution with two degrees of freedom (Rayleigh distribution) and it depends on one only parameter, σ. To use this distribution function, σ must be known, and in general σ is a function of frequency because the spectrum is not white. The statistics of an individual frequency bin still follows a Rayleigh distribution, but the variance changes from point to point. The statistical peak detection method cannot be used to separate signals from noise unless σ can be estimated reliably.

2.3 Solutions

A solution to each of the three fundamental problems will be outlined in general in the following sections, then in discussed in more detail in Section 3 in the context of an analysis of data taken in June 1985 with the MIT prototype gravitational wave antenna.

2.3.1 *Gaps.* The obvious solution to the problem of gaps in the data stream that are small compared to the length of the data stream was mentioned above – simply fill in the gap with the average value of the data.

The obvious solution to the problem of large gaps in the data is to analyze the continuous pieces separately, then combine the pieces with an r.m.s. (root mean square) average. This is easy to do, but has the disadvantage that the power SNR increases only with the square root of the integration time and not directly with the time as could be achieved with continuous data. Adding more data still improves the SNR, but not as quickly as if the data were continuous.

The gap problem for the June 1985 data taken with the MIT prototype was much more severe than it should be for a real antenna. It is worthwhile exploring methods to deal with the gaps, however, because any real data set will always have gaps of some size if it is large enough. The gap problem may become moot for another reason. The computational problems associated with large data sets may impose a separate constraint on the length of the data set that may be easily ob-

tainable with a better site and a more sophisticated data collection system than the one used by the MIT prototype.

2.3.2 *Expected Signals.* Two different solutions to the problem of a direction dependent received signal have been developed. One solution approaches the ideal of a broadband full sky search, but with a limited sensitivity. The other solution can achieve a higher sensitivity, but only in one direction at a time. Details of the implementation of these ideas will be discussed in Section 3 of this article.

Full Sky Search. A full sky search can be achieved by making use of the FM and AM modulation as signatures to distinguish signals from local disturbances. As discussed in the Section 2.2.2, the FM modulation becomes important only for integration times longer than 30 minutes. Thus, if the data is analyzed in continuous pieces shorter than 30 minutes in length, a real signal will still appear as a delta function in frequency. However, the signal will appear at different frequencies in different pieces, and this difference in center frequency can be used to localize the source to a region of the sky. For a data set consisting of pieces spanning a week baseline, the dominant Doppler shift is that due to the orbital motion of the earth. Figure 4 shows the expected Doppler variation for a source located at the center of the galaxy in early June 1985. The vertical scale is marked off in units of the frequency resolution of a 15 minute FFT, or $\approx$ 1.5 mHz. Since the daily Doppler shift is small, the Doppler shift for the week may be characterized by a single slope, in units of Hz/Hz/sec, for each direction. Furthermore, the slope has a maximum physical value that corresponds to $\pm v_{orb}/c$. Figure 5 is a map of the sky showing areas of constant Doppler slope as a grey scale. A known value of the slope identifies a locus of points on the sky as possible source locations.

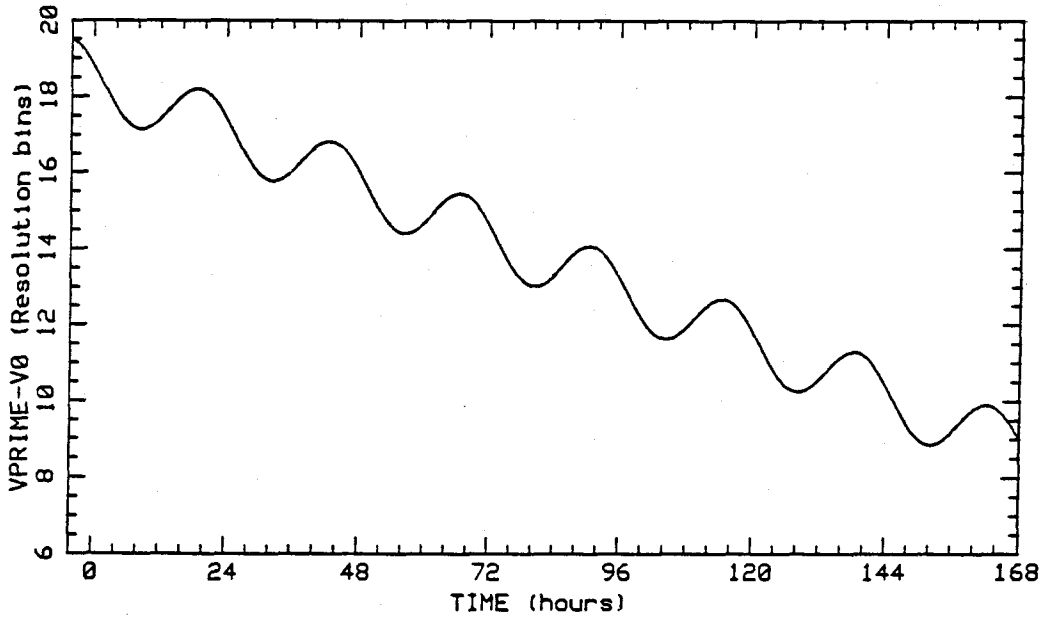

Figure 4
Doppler variation of the received frequency for a sinusoidal source at the center of the galaxy. T=0 for June 4, 1985.

The AM modulation can be used to further localize a source. As with the FM,

the modulation per continuous piece is small, but the signal strength difference between pieces can be significant. The AM can be used to predict the relative signal strengths in each piece of continuous data as a function of the position of the source on the sky. A candidate signal that has been detected in more than one piece must have a physically reasonable Doppler slope and must also match the AM response pattern before it can be considered a possible astrophysical source.

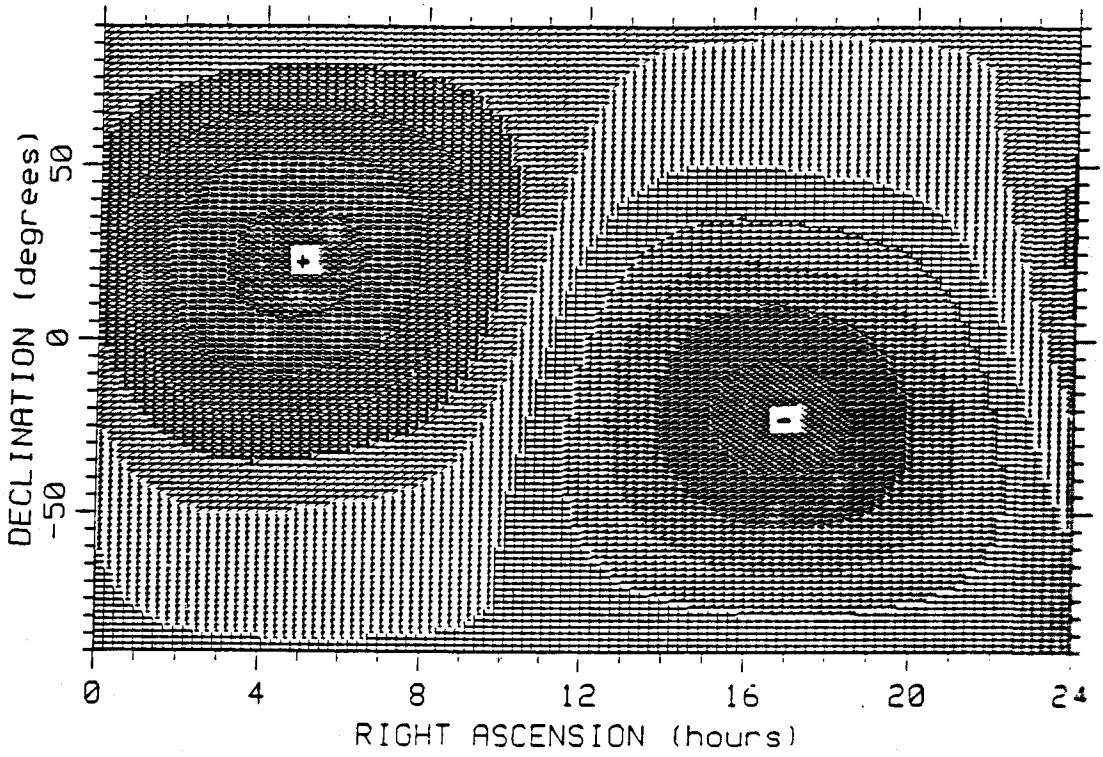

Figure 5
Doppler slope as a function of position of source on the sky, expressed as a grey scale from $+v_{orb}/c$ to $-v_{orb}/c$.

The drawback with this procedure is that the sensitivity is limited by the Doppler modulation to what can be achieved with an integration time of 30 minutes or less. Since the signal appears at different frequencies in different piece of data, averaging does not help.

Single Direction Search. Better sensitivity can be achieved by demodulating the FM, so that a candidate signal would always appear at the same frequency in different pieces. The separate pieces could then be r.m.s. averaged together to improve the SNR. The length of each piece could be much longer than 30 minutes in principle.

The idea is simple. The general form of the signal is:

$$s(t) = A \cdot \cos \varphi(t), \tag{12}$$

where φ is the phase. From the expression for the Doppler shift, the observed frequency is

$$f_{obs} = f_{emit} \cdot \gamma_0 \cdot (1 - \frac{\vec{v}}{c} \cdot \hat{r}). \tag{13}$$

The phase of the observed signal is then

$$\varphi = 2\pi \int_{t_0}^{t} f(t')dt' \tag{14a}$$

$$= 2\pi f \gamma_0 \left(t - t_0 - \frac{\hat{r}}{c} \cdot \int_{t_0}^{t} \vec{v}(t')dt' \right) \tag{14b}$$

$$= 2\pi f \gamma_0 t' \tag{14c}$$

where

$$t' = t - t_0 - \frac{\hat{r}}{c} \cdot \int_{t_0}^{t} \vec{v}(t')dt' \tag{15}$$

is called the rescaled time and depends explicitly on the position of the source. Equation (14c) is the relationship between frequency and phase if the frequency is constant in time. One rescaled time function demodulates all frequencies simultaneously, but for only one direction.

In practice, what is needed for the FFT computation is a time series that is sampled at evenly spaced intervals in rescaled time. The recorded data is a time series taken at evenly spaced intervals in laboratory time. If the recorded data is treated as though the samples were collected at unequal intervals in rescaled time and then resampled to equal intervals by interpolation, the desired result is accomplished. Figure 6 depicts the process schematically. The solid curve is the original data which was recorded at equally spaced intervals in laboratory time. It is plotted with each data point mapped into its equivalent rescaled time value, so the points are not evenly spaced. The dotted curve represents the new time series that is generated by interpolating between the unevenly spaced samples in rescaled time of the solid curve to get an evenly spaced set of samples.

Interpolation is not a clean procedure. The rescaled time and laboratory time can slip by several sampling periods over the course of 30 minutes. Since the error in interpolation increases as the point to be estimated falls further from the known point, the interpolation error puts a low frequency modulation in amplitude and phase on the original time series. The effects of the modulation may be reduced by using a more complicated interpolation algorithm or by oversampling the original time series. Neither is desirable. Oversampling increases the data acquisition and storage problems, and a complicated interpolation algorithm is computationally very slow.

2.3.3 *Unknown* $\sigma(f)$. One solution to the problem of the unknown variance as a function of frequency is to determine the variance from a local average of the data. The power spectrum is divided into pieces referred to as the averaging bandwidth. The variance of the spectrum is assumed to be constant over that bandwidth, and an estimate of σ is derived from the mean value of the spectrum in that bandwidth. Explicitly,

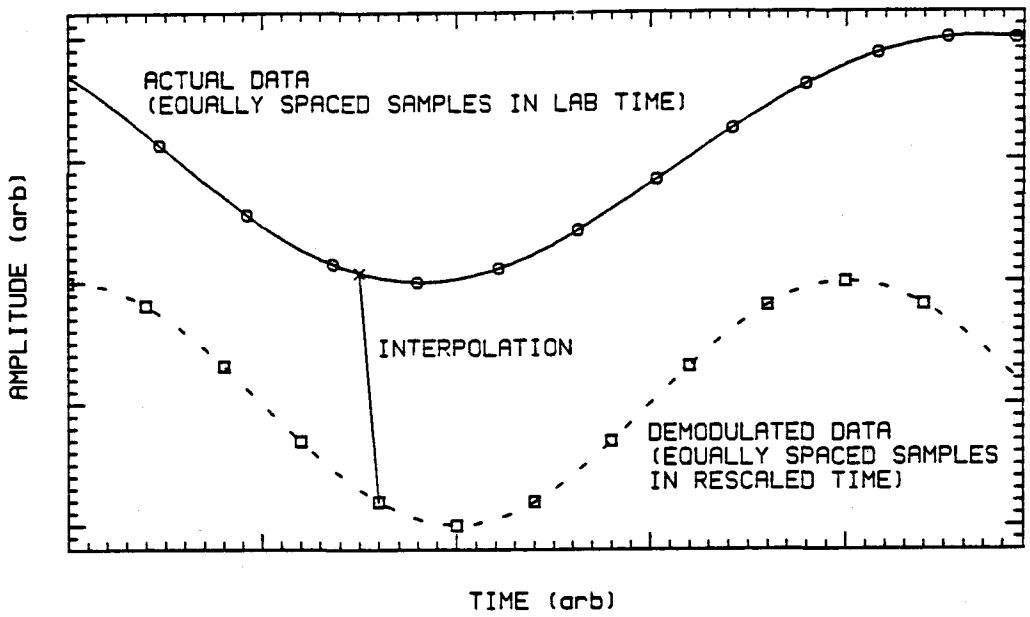

Figure 6

Schematic representation of the resampling procedure to produce a new time series sampled at equally spaced intervals in rescaled time.

$$< s^2 > = \int_0^\infty s^2 p_{noise}(s) ds$$
$$= 2\sigma^2 \qquad (16)$$

where $s = S_N(m)$ is the amplitude of a bin in the square root of the periodogram and p_{noise} was taken from equation (11). The estimated value of σ can be used to normalize each bin in that averaging bandwidth. The normalized values of the whole spectrum may then be compared on an equal basis.

2.3.4 *Summary.* To quickly summarize the techniques discussed so far, the desired signal is assumed to be a high-Q, essentially delta function in frequency. The strategy is then to integrate for as long as possible to increase the SNR. The fact that the received signals are modulated even if emitted by monochromatic sources is handled two different ways. A full sky search can be conducted by analyzing the data in pieces and using the Doppler FM and the AM between pieces as a signature for which to search. An enhanced sensitivity search can be conducted by performing a demodulation for one specific direction at a time. The demodulation collapses the received signal back into a delta function in frequency which can then be averaged to improve the SNR. The statistical variance is estimated as a function of frequency by using local averages over a specfied bandwidth. These techniques will now be discussed in the context of a specific example – the analysis of data taken in June 1985 with the 1.5m MIT prototype antenna.

3. EXAMPLE: ANALYSIS OF JUNE 1985 DATA

3.1 Data Set

The data analysis scheme for periodic sources has been outlined conceptually. An application of these ideas to a real data set illustrates not only how well the ideas work in practice, but also points out some of the practical problems which will have a profound impact on the way data analysis for a large antenna should be conducted.

The MIT 1.5m prototype antenna and data collection system have been described in detail elsewhere.[7] The raw data consists of the demodulated output of the interference fringe lock servo, digitized at either a 20 kHz rate, or a 6 kHz rate. The digitization is triggered from a rubidium standard that is sychronized to WWV. As mentioned above, limitations in the recording hardware restrict the length of a continuous data stream to approximately 15 minutes ($N = 2^{24} \approx 16 \times 10^6$ points).

The data collected in June 1985 consists of approximately 50 magnetic tapes collected over 6 days, representing roughly 10% coverage of the total possible time. One half of the tapes were sampled at 20 kHz, the other half at 6 kHz. Eight of the 25 tapes sampled at 20 kHz were selected for the final analysis.

3.2 Full Sky Search

The first step in the analysis is to choose a value of the averaging bandwidth. For the statistical peak detection scheme to be valid, the statistics of the power in each resolution bin must agree with the theoretical Rayleigh distribution The choice of an averaging bandwidth was made by constructing an experimental probability distribution for several different values of the averaging bandwidth and choosing the value which gave the best least squares fit to the theoretical distribution.

Figure 7 shows the fit with a good choice of the averaging bandwidth. The experimental distribution is computed by normalizing the value of the each frequency bin in the square root of the periodogram to the local σ. The normalized points are then binned in a histogram which counts the number of points with a given normalized amplitude versus that normalized amplitude. The solid curve is the theoretical distribution for pure noise, $p(\xi) = \xi e^{-\xi^2}$, integrated over each bin in the histogram. A χ^2 fit of the experimental distribution to the theoretical distribution was performed, and the value of the averaging bandwidth used in the final analysis was that which gave the lowest value of the χ^2 for the fit. For this data set, that value was 0.35 Hz.

The arrow is set at a threshold of 4.97σ, which represents a false alarm probability of $p_{fa} = 5 \times 10^{-6}$. A frequency bin containing enough power to place it to the right of the threshold is a candidate signal. For this particular distribution, with $\approx 2 \times 10^6$ total resolution bins, a spectrum containing pure noise would contain on average 11 peaks above the threshold. In this case, the threshold was chosen conservatively so as not to miss any potential signals, and it was also placed at the point where the experimental distribution exhibits a non-Rayleigh tail.

Once the value of the averaging bandwidth has been chosen, the first step in the analysis is straightforward: the spectrum of each piece of data is binned separately, and candidate signals are identified with the threshold criterion. A list

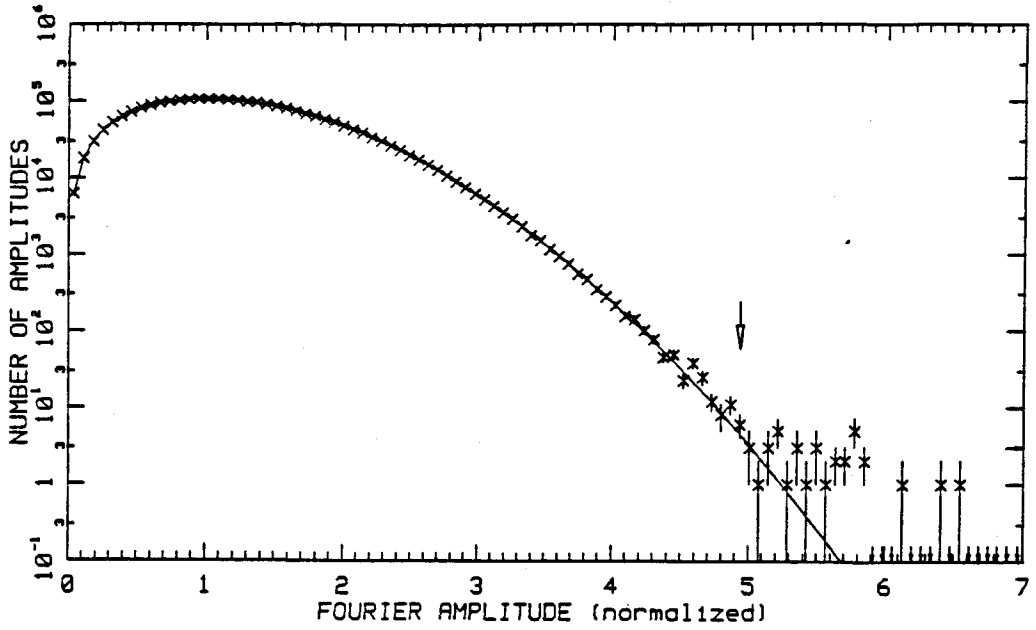

Figure 7

Experimental fit to the theoretical probability distribution for pure noise (solid curve) for a particular value of the averaging bandwidth. Normalized to $\approx 2 \times 10^6$ *total points. The arrow indicates a* 4.97σ *threshold which gives a false alarm probability* $p_{fa} = 5 \times 10^{-6}$.

of the possible signals in each piece is compiled, keeping track of the frequency of the signal, its magnitude in absolute units, and its value relative to the local σ are recorded, along with an estimate of the number of adjacent frequency bins that also qualify as possible signals (a measure of the width of the peak).

The complete spectrum of each piece extended from 0-10 kHz, but only the region from 2-5 kHz was searched. The low frequency cutoff was chosen at a point where the noise spectrum of the antenna begins to increase rapidly with decreasing frequency, and hence the sensitivity to gravitational radiation becomes small. The high frequency cutoff was chosen to avoid problems with the interpolation method used in the single direction search.

These lists of possible peaks are then examined for a series of peaks whose center frequency changes over the course of the week. The method for finding such a series is called a Doppler sieve. Figure 8 describes the operation of the sieve schematically. One piece of data is selected as the reference. If a peak in the reference list is actually an astrophysical signal, then it should appear in the lists of the other pieces, but with its center frequency shifted by a known amount. The sieve is applied by dividing the allowable range of Doppler slopes into discrete values. Each peak in the reference list is used to estimate the center frequency for each of those slopes that that peak would have if it were observed at the time the other data was collected. To get the center frequency of a peak in the j^{th} tape:

$$f_{center}(j) = f_{center}(ref) \cdot m_D \cdot (t_j - t_{ref}) \qquad (16),$$

where m_D is the Doppler slope in units of Hz/Hz/sec. The actual list of candidate signals from the j_{th} tape is examined for anything within a given capture width Δf_{cap} to the projected center frequency. If at least one other tape has a peak which

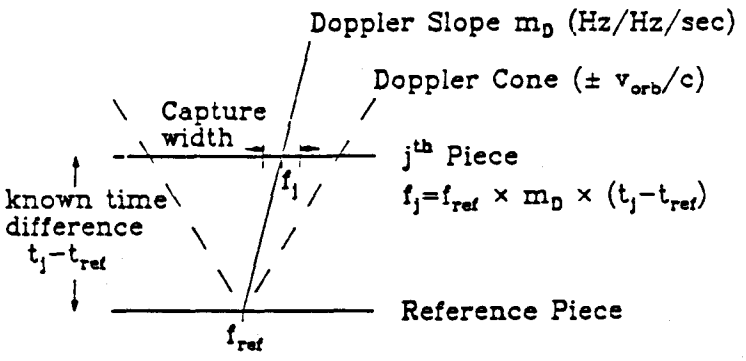

Figure 8
Schematic operation of the Doppler sieve.

satisfies this condition, the reference peak is recorded. Note that the condition for acceptance of a peak is the weakest possible. A single pair of peaks that falls along a physically plausible Doppler slope is enough to qualify. In general, there could be more than two peaks, but the received signals are also AM modulated. A weak signal is therefore most likely to be seen only in the two most favorable observation windows.

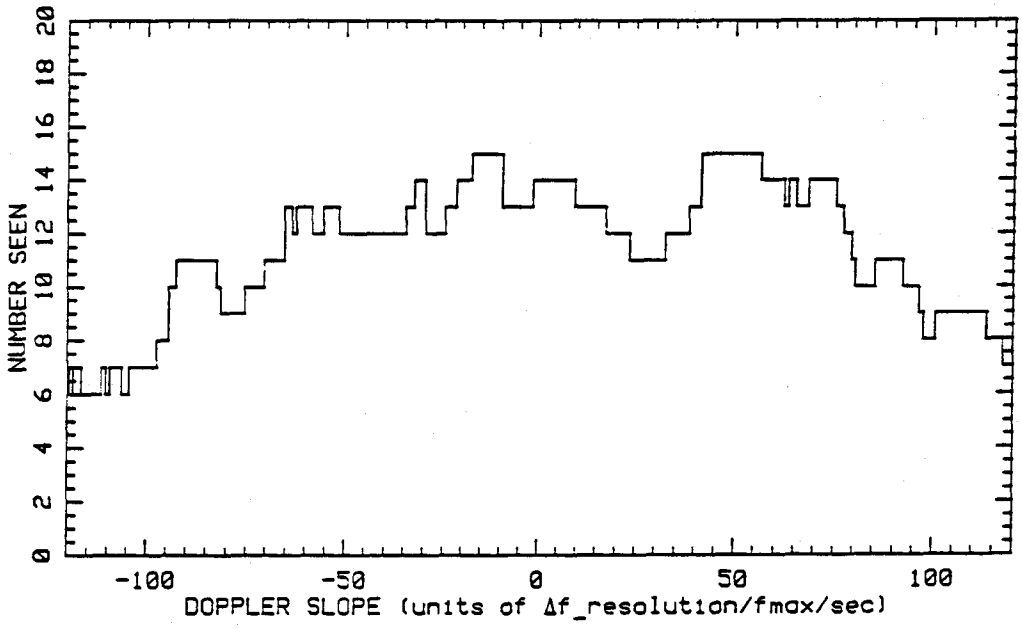

Figure 9
Typical histogram of the number of peaks in a reference list which match at least one other peak in another list versus the Doppler slope for which there is a match. Only slopes between ±50 units are physical.

For each reference piece, the peaks that fall through the Doppler sieve are binned in a histogram of the number of peaks with a given Doppler slope versus

the slope. Figure 9 shows a typical histogram, with the horizontal axis in units of the quantized Doppler slope. A larger range of slopes than is physically possible is plotted to estimate the "background". The physical region in Figure 9 is between $\pm$ 50 units.

Peaks caused by local disturbances will not exhibit a Doppler shift over time and will thus appear in the histogram with a slope near zero. For this reason, any candidate peak appearing in coincidence with a peak in another list connected by a Doppler slope of $m_D = 0$ in quantized units, was eliminated from consideration. Figure 10 shows the histogram of Figure 9, but with the zero-slope bin deleted. A real source near the zenith of the detector would also exhibit no doppler shift, so the exclusion of the zero-slope bin could in principle be eliminating some real sources. If the sources are uniformly distributed on the sky, only 8% of the source would be eliminated. This is a small price to pay for the elimination of a large number of spurious coincidences. The change in the vertical scale indicates that some of the peaks eliminated appeared in the histogram with more than one Doppler slope.

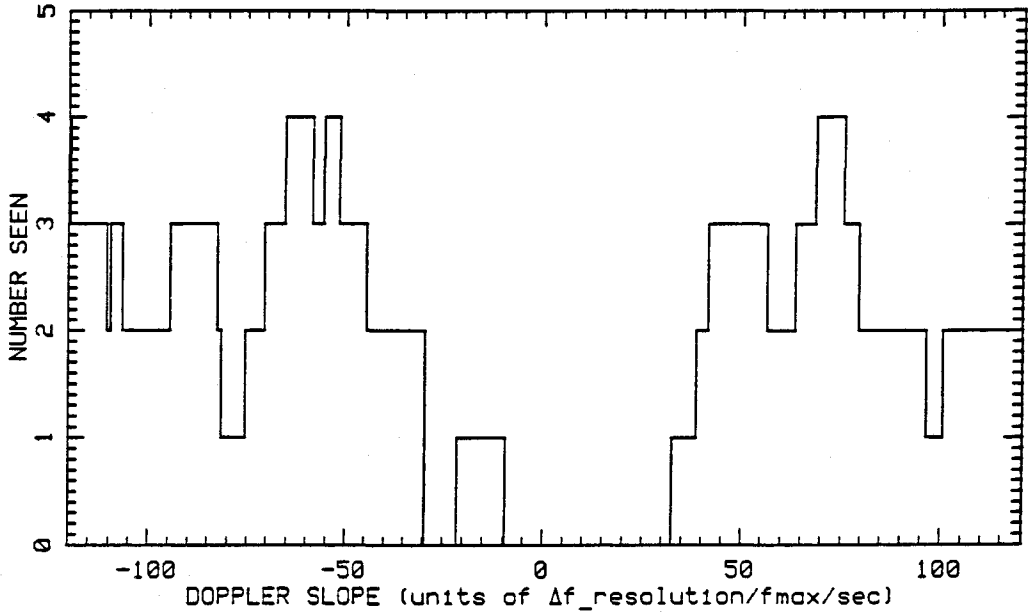

Figure 10

Histogram of the Figure 9 with peaks in the zero-slope bin deleted. Only slopes between $\pm$50 units are physical.

The peaks that manage to fall through the Doppler sieve are then associated with a region of the sky using Figure 5. The sky is divided into discrete bins, and for each bin in the region picked out by the Doppler slope, the expected pattern of signal strengths is determined. Figure 11 is a map of the sky showing, for each bin, which pair of the particular data tapes used in the analysis should have the strongest signals based on the AM response to a source in that portion of the sky. Figure 11 is for a plus polarized wave. A similar map can be made for the cross polarization. The actual strengths can be compared against the predicted pattern.

Figure 12 summarizes the full sky search procedure. A total of three peaks out of $\approx 2 \times 10^6$ frequency bins actually passed both the Doppler sieve test and the AM signal strength pattern. None of the peaks is a good candidate signal for a number

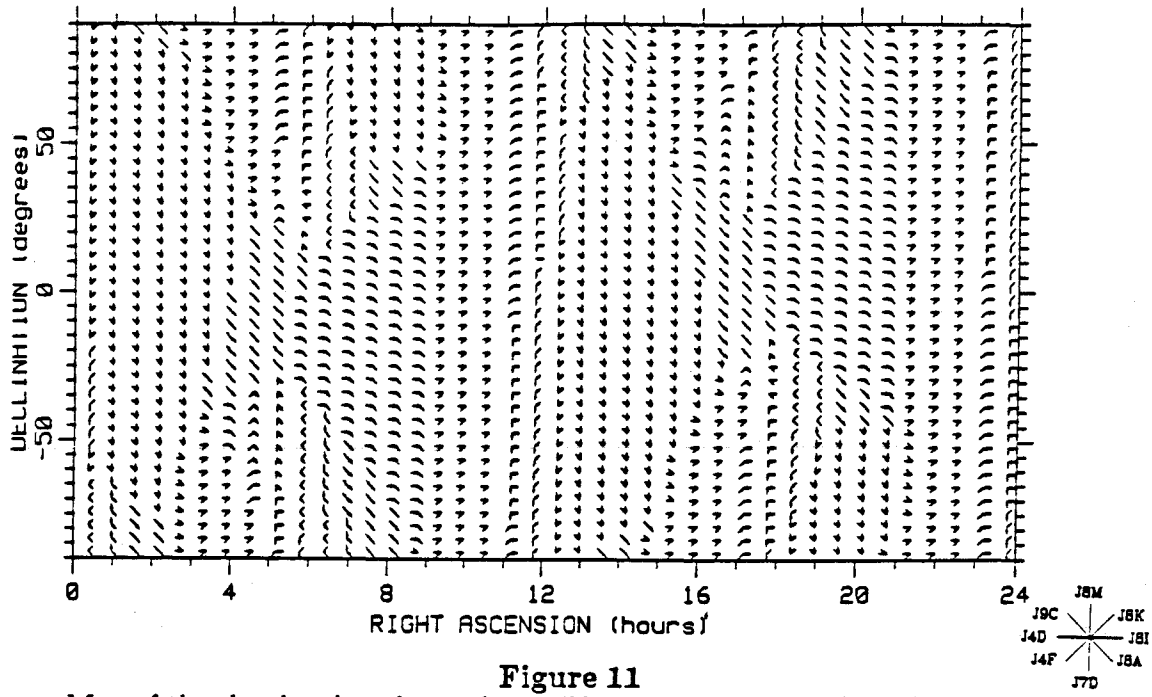

Figure 11

Map of the sky showing, for each possible source location, the pair of tapes in which the signal should be most strongly received. See key at right.

of reasons, but the primary reason is that all occured in a region of the spectrum in which there was appreciable structure in the spectrum on the same scale size as the averaging bandwidth. The local estimate of σ is therefore not a good estimate, and the statistical peak detection method fails because the experimental distribution is not a good approximation to the theoretical Rayeigh distribution.

There are several strategies that could be adopted to circumvent the failure of the statistical peak detection method. One could use a more sophisticated fit to each averaging bandwidth: instead of assuming a flat spectrum, one could allow a slope plus a constant. Another alternative would be to allow the size of the averaging bandwidth, which is chosen by a global optimization of the statistics for the entire spectrum, to vary over the spectrum and select it by some local optimization procedure. Individual peaks could also be modeled and subtracted from the spectrum. Of course, the best solution is to eliminate the noise sources in the antenna to produce a flat spectrum with which to start.

3.3 Single Direction Search

The single direction search procedure is similar to the full sky search in the initial stages of analysis. The particular direction to be searched with this data set was chosen to be the center of the galaxy. Each piece of data was demodulated according to the algorithm described in section 2.3.2, using a piecewise continuous spline interpolation procedure. An average spectrum was computed by r.m.s. averaging each averaging bandwidth, weighted by the σ for tha bandwidth, with the corresponding bandwidths in other pieces of data. The net improvement in SNR, if all of the local σ were identical, should be $\sqrt{8}$. A reference r.m.s. spectrum was also computed by following the same procedure without the Doppler demodulation.

The best value of the averaging bandwidth is chosen as before with a global

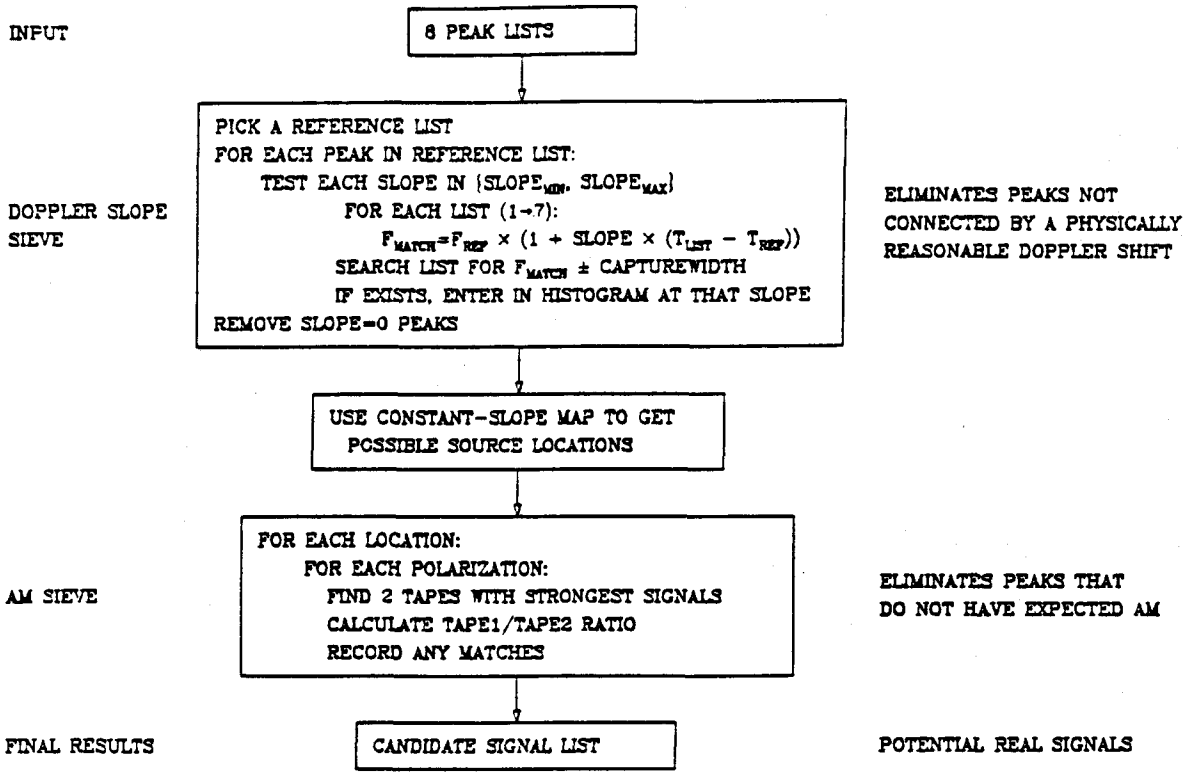

Figure 12
Summary of the full sky search procedure.

minimization of the fit of the experimental differential probability distribution to the theoretical distribution. The Doppler demodulation procedure did not change the value of the best fit averaging bandwidth. Both average spectra are then be subjected to the statistical peak detection analysis and a list of candidate peaks in each spectrum is generated. The expected probability distribution is slightly different because of the r.m.s. averaging. The expected distribution for the amplitude in a single bin of the square root of the power spectrum with only noise present in the data is a χ^2 distribution with $2q$ degrees of freedom, where q is the number of pieces combined in the average.

Real astrophysical signals should have been enhanced by the demodulation procedure, and thus any new peaks in the demodulated spectrum are potentially signals. However, local noise driven resonances which appear at the same frequency in each piece will get smeared by the demodulation procedure. To separate truly new peaks in the demodulated spectrum from local peaks that have been smeared, the two lists of peaks were compared with a procedure very similar to an inverse of the Doppler sieve. First, any peaks that appeared at exactly the same frequencies in both lists were eliminated on the grounds that such signals were not physical. Next, any peaks in the doppler corrected spectrum that had a corresponding peak in the uncorrected spectrum offset by the frequency shift expected for one of the component piece of data was eliminated. The peaks that were left over were truly

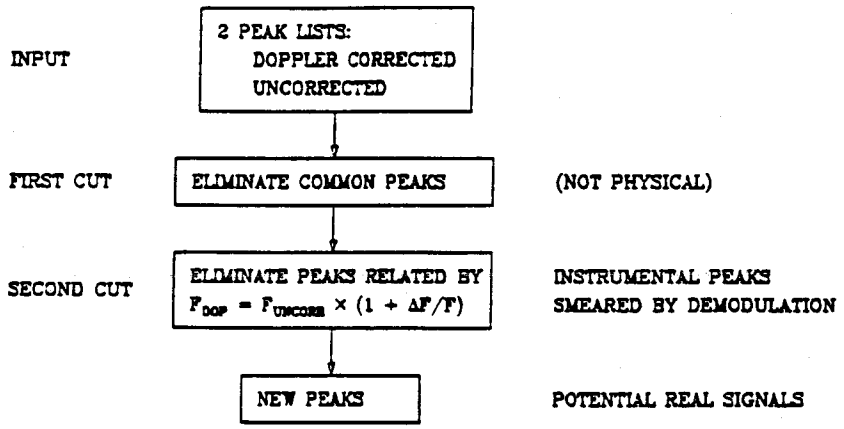

Figure 13
Summary of the single direction search procedure.

new. Figure 13 summarizes the procedure.

Only 22 peaks met the criteria for "new" peaks. For the size data set considered, 2×10^6 frequency bins, a pure noise spectrum should have generated 17^{+9}_{-5} peaks above threshold. The error bars reflect a 1% uncertainty in the value of the threshold. The results are thus consistent with a pure noise spectrum. A closer examination of the 22 peaks would require more data.

3.2.4 Results

The final results of the two searches for periodic signals can be stated rather simply. The full sky search did not see any periodic signals which produced a *measured* strain in the antenna above a level

$$h_{rms} = 4.6(\pm 0.9) \times 10^{-17} \left(\frac{2000}{f}\right)^2 \tag{17}$$

in a frequency band from 2-5 kHz. The indicated error is one standard deviation. The conversion of this measured strain into a source strength depends on the polarization and the direction of the source. For a source at the center of the galaxy, the limits translate into

$$h_{+,rms} = 2.1(\pm .6) \times 10^{-16} \left(\frac{2000}{f}\right)^2 \tag{18a}$$

$$h_{\times,rms} = 7.4(\pm 4.) \times 10^{-17} \left(\frac{2000}{f}\right)^2 \qquad (18b)$$

The corresponding limits for the single direction search are

$$h_{+,rms} = 1.0(\pm 0.1) \times 10^{-16} \left(\frac{2000}{f}\right)^2 \qquad (19a)$$

$$h_{\times,rms} = 5.2(\pm .8) \times 10^{-17} \left(\frac{2000}{f}\right)^2 \qquad (19b)$$

The sensitivity gain for the single direction search is slightly less than the $\sqrt{8} = 2.8$ expected with eight RMS averages. Some of the discrepancy can be attributed to the unequal weighting of the individual spectra, and some is due to a choice of direction for the search that is perhaps not the best for exhibiting the best SNR increase for the method. With more data, the increase in SNR could be much larger.

3.3 Practical Implications for Future Research.

The prospects for data analysis techniques to discover periodic sources of gravitational waves may seem unduly encouraging. Two different methods have been presented to perform broadband searches. However, there are several severe practical difficulties with these methods which will make any attempt to extend these methods difficult. Both are rooted in the computational difficulties associated with large data sets.

It might be expected that single direction broadband search technique could simply be repeated for many different directions. However, the demodulation procedure correctly demodulates signals from a finite patch of sky around the desired direction. An estimate of the patch size can be generated by treating the received frequency in equation (13) as a function of the position angles α, δ, where α is the right ascension and δ is the declination. Then

$$\Delta f = \left(\frac{\partial f}{\partial \alpha}\right) \Delta\alpha + \left(\frac{\partial f}{\partial \delta}\right) \Delta\delta. \qquad (20)$$

The change in either angle must be small enough to keep $\Delta f \leq \Delta f_{res}$, where Δf_{res} is the frequency resolution bandwidth of the transform. Equation (20) can be inverted to and simplified to obtain an expression for either angle of the form

$$\Delta(\alpha, \delta) = \xi(\alpha, \delta) \left(\frac{1}{f_{max}}\right) \left(\frac{1}{T}\right)$$

where $f_max = 1/2\Delta t_{sample}$ is the Nyquist frequency for a given sampling rate Δt_{sample}, and T is the length of the transform, and ξ is a function of α, δ. For a source at the Galactic Center, a 20 kHz sampling frequency, and a 15 minute data record, the patch size on the sky is

$$\Delta\alpha \approx 9^\circ \left(\frac{f_{max}}{10\,\mathrm{kHz}}\right)^{-1} \left(\frac{T}{900\,\mathrm{sec}}\right)^{-1} \tag{21a}$$

$$\Delta\delta \approx 15^\circ \left(\frac{f_{max}}{10\,\mathrm{kHz}}\right)^{-1} \left(\frac{T}{900\,\mathrm{sec}}\right)^{-1} \tag{21b}$$

If this patch size were constant over the whole sky (which it is not), this implies $\sim$ 500 patches must be searched. This does not seem like an excessive number, except that it scales as T^2. The computation time for the analysis that was performed on the MIT data is summarized in Table 2. Unpacking refers to an operation that converted the data (stored in condensed form) into Cray-2 real numbers. The volume of data was such that it could not be easily stored in unpacked form on the available disk space. Net computational speed is approximately 7.5 Mflops for the FFT.

Table 2

Computation times for the periodic source search. Computations were performed on a Cray-2.

Operation	CPU time (secs)	% of Total	CPU time (secs)	% of Total
Unpacking	225.4	46.5	225.4	32.2
Doppler Demod.	—	—	216.0	30.8
FFT	71.3	14.7	71.3	10.2
Apodization	8.8	1.8	8.8	1.3
Peak Search	179.8	37.0	179.8	25.5
Total	485.3	100.0	701.3	100.0

Based on these execution times, a complete sky search with extended sensitivity would require $500 \times 700 = 3.5 \times 10^5$ seconds, or $\sim$ 97 hours of CPU time. This is a large amount of computation time, and any attempt to extend the analysis techniques much beyond what has already been done will require even more time. Of course, there may be physical reasons for selecting a direction in which to search, and then the number of bins required to cover the sky may not be important.

There is another computational bottleneck that is not apparent from the table. The maximum size FFT that can be computed is determined by the size of the core memory of the machine. A Cray 2 has the largest core memory of any machine currently in existence, $2^{28} \approx 268 \times 10^6$ words. Based on the measured execution time of a 2^{24} point FFT, a 2^{28} point FFT should require 1331 seconds, or 22.2 minutes of CPU time. The estimated total CPU time for a doppler demodulation analysis of this size is 3.2 hours, with 123,000 bins required to cover the sky. Any attempt to perform larger FFT's changes the problem dramatically. The computation must be done in pieces, and it becomes immediately limited by the I/O transfer rate to the disk.

The point of this discussion is that a complete solution to the data analysis problem for periodic sources is not in hand. A different approach will be needed to

achieve integration times of 10^6 seconds, and faster general purpose computers are not the answer. A real solution will require a combined hardware/software solution and perhaps different analysis algorithms.

ACKNOWLEDMENTS

This work was supported in part by NSF grant 85-0486-PHY, which was augmented by a separate grant of computer time at the University of Minnesota Supercomputer Center. The author would also like to thank the people of the MIT Gravity Wave group, in particular Rai Weiss, for support and encouragement.

REFERENCES

1) Thorne, Kip S., in *300 Years of Gravitation*, eds. Hawking and Israel, Cambridge, (1987).
2) Schutz, B. F., this volume.
3) Taylor, J. and Weisberg, J. *Ap. J.*, **253**, 908 (1982).
4) Wagoner, R. *Ap. J.*, **278**, 345, (1984).
5) Lyne, A. et al. *M.N.R.A.S.*, **213**, 613, (1985).
6) Henrichs, H. and van den Heuval, E. *Nature*, **303**, 213, (1983).
7) Livas, et. al. in *Proceedings of the Fourth Marcel Grossman Meeting*, ed. R. Ruffini, 591, (1986).