Proceedings of the NATO Advanced Research Workshop: Gravitational Wave Data Analysis, edited by B.F. Schutz (Reidel, Dordrecht, 1987).

DATA ANALYSIS AS A NOISE DIAGNOSTIC: LOOKING FOR TRANSIENTS IN INTERFEROMETERS

Daniel Dewey
Massachusetts Institute of Technology
Room 3-253
Cambridge, MA 02139
U.S.A.

ABSTRACT. In addition to the ultimate goal of detecting and analyzing source waveforms, data analysis provides a framework for identifying important noise sources. Only by developing and implementing data analysis programs designed to search for expected astrophysical sources will we uncover all relevant instrumental noise sources and have a true picture of an antenna's performance. In particular, the characterization of an interferometric antenna solely by its noise spectrum hides the possible existence of weak and infrequent noise transients. A search for such transients using simple matched filter templates was carried out with the MIT prototype. An analysis of 1.7 hours of data (250 Mbytes) showed an agreement with Gaussian statistics out to an amplitude signal-to-noise ratio of 5.5. Above this signal-to-noise level, transients identified with electrical feedthrough in our data taking system were recorded at a rate of $\simeq$70 per hour, and transients of unidentified origin and various natures occurred at a rate of $\simeq$22 per hour.

1. PHILOSOPHY AND INTRODUCTION

Data analysis strategies should be used at the prototype stage to uncover all relevant noise sources.

Astrophysical theory suggests two broad classes of sources: periodic sources and transient sources. This dichotomy of source types leads naturally to (at least) two noise diagnostics: the (long-time averaged) power spectrum and the pulse-height distribution (PHD) of discrete transient events. Measurement of an interferometer's output spectrum (typically presented as a log-log plot of equivalent strain noise ($h/\sqrt{Hz}$) vs. frequency) has been used almost exclusively by interferometer groups as the specification of instrument performance with little investigation of transient events; Forward's aural data analysis[1] is a

more-appropriate-than-it-first-appears exception. Conversely, bar antenna groups have tended to present bar performance in terms of the distribution of transient events with little direct measurement of the spectral properties of their noise. To some extent these specializations reflect the difference in instrumental bandwidths and philosophies.

This work presents results obtained by applying the matched filter technique to search an interferometer's output for a broad range of transient events. The basic idea of applying matched filters has been described previously[2] and is briefly recounted in Section 2. Section 3 presents specific details of the data collection and data analysis and the results of analyzing 1.7 hours of data are given in Section 4. Finally, Section 5 contains some brief comments on the technique and a discussion of future work.

2. THE MATCHED FILTER AND A SET OF TEMPLATES

2.1. The Matched Filter

The optimum signal-to-noise ratio (SNR) for the detection of a known signal $h(t)$ in the presence of white, Gaussian noise characterized by a spectral density $\tilde{h}$ is a well-known result of detection theory[3,4,5] and is given by:

$$(S/N)_{opt} = \left[\frac{2\int h^2(t)dt}{\tilde{h}^2} \right]^{1/2} \qquad (2.1)$$

How is a detection yielding this optimum (i.e. maximum) SNR achieved? Again the result is well-known: it is obtained through matched filter correlation in which a template of the known signal is correlated with the input waveform to produce an output waveform; this output is essentially the dot product of the template with the input waveform as a function of template position.

Because the input noise to the linear correlation filter is Gaussian (assuming no signal is present) the resulting output time series is again Gaussian. Thus, ideally, a pulse-height distribution (PHD) of the output series will be Gaussian and given by the probability density:

$$p(x) = \frac{1}{\sqrt{2\pi}} e^{-(x/\sigma)^2/2} \frac{dx}{\sigma} , \qquad (2.2)$$

where σ is the rms value of the filter output. Gross deviations from this distribution represent non-Gaussian noise, perhaps due to a non-linear component in the signal chain; other deviations, especially 'tail events', may be

caused by the existence of transient noise sources.

It should be mentioned that these matched filter results could have been equivalently expressed in the frequency domain (i.e. using $h(\omega)$ in place of $h(t)$); however a time-domain presentation seems a more appropriate view for the understanding of discrete, transient events. Additionally, with modern digital signal processing techniques it is often more convenient to work in the time domain.

2.2 Effect of a Mismatch of Template and Signal

An initial problem in implementing the matched filter technique is the fact that the shape of the signal is not known. The effects of template-signal mismatch can be easily evaluated in the following framework. If the signal and template are viewed as made up of some number of half-cycle components N_{hc} (defined by the zero-crossings of the waveform) then a mismatch between signal and template can be viewed in terms of a mismatch of: 1) the shape of each component pulse, 2) the amplitude of each component pulse, 3) the spacing of the component pulses, and 4) the total number of component pulses. In general, if a template $h_1(t)$ is used to detect a signal of shape $h_2(t)$ the reduction in SNR from an ideal matched filter is

$$\frac{(S/N)}{(S/N)_{opt}} = \frac{\int h_1(t)h_2(t+\tau)dt}{[\int h_1^2(t)dt \cdot \int h_2^2(t)dt]^{1/2}}, \qquad (2.3)$$

where τ is chosen to maximize the numerator. Note that because of the symmetry of Eq. (2.3) the roles of template and signal may be interchanged. Through Eq.s (2.1) and (2.3) the effect of the four types of mismatch can be determined.

2.2.1 Shape mismatch. In general a shape mismatch of component pulses will have only a slight effect on the SNR for detection. For example, detecting a triangular waveform using a square wave gives $(S/N)/(S/N)_{opt} \simeq 87\%$; for a three-level waveform (Fig. 1b) and a sine wave, it's 96%.

2.2.2 Amplitude mismatch. It can be seen from Eq. (2.1) that a single component pulse of amplitude h_{cp} makes a contribution to the integrand that is proportional to $h_{cp}^2 \Delta t$, where Δt is the length of the component pulse. Thus, only the largest amplitude component pulses are important.

2.2.3 Number mismatch. A mismatch in the number of component pulses leads to a reduction in SNR given by:

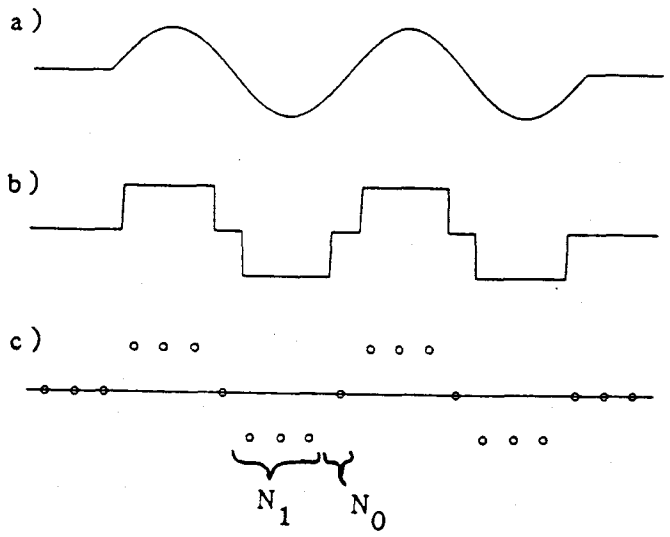

Figure 1. From a canonical burst to a digital template. The four half-cycle burst in a) is replaced by a three level approximation in b) and finally a discrete-time template in c); this template is specified by the three integers N_1=3, N_0=1, and N_{hc}=4, more compactly written as ^{1}314.

$$\frac{(S/N)}{(S/N)_{opt}} = \left[\frac{\min\{N_{hc,1}, N_{hc,2}\}}{\max\{N_{hc,1}, N_{hc,2}\}} \right]^{1/2} . \qquad (2.4)$$

For example, detecting a 5 half-cycle waveform with a 6 half-cycle template results in $(S/N)/(S/N)_{opt} \simeq 91\%$.

2.2.4 Spacing (or frequency) mismatch. The sensitivity of the SNR to a mismatch in the spacing or frequency of component pulses is a function of the number of half-cycles in the template; the bandwidth decreases with increasing N_{hc}. For templates with N_{hc}=2 a frequency range of $\pm 30\%$ is covered with a minimum $(S/N)/(S/N)_{opt}$ of 80%, while the N_{hc}=6 templates cover only a $\pm 12\%$ frequency range when maintaining a comparable sensitivity.

From the mismatch discussions above it becomes clear that to ensure a high SNR for detection the parameters it is most important to match between template and waveform are the frequency and number of half-cycles.

2.3 The Set of Templates

Many of the expected waveforms from gravitational wave burst sources can be well approximated by several (1 to 8) half-cycles of a sine wave at a constant frequency[2]; one such canonical burst, made up of four half-cycles, is shown in Fig. 1a. The frequency scale for burst sources (e.g.

stellar collapse and particle-black hole encounters) is predicted to be in the range of 1 kHz or less to as high as 10 kHz. Here, the goal in choosing a set of templates is to reduce the chance of a "false dismissal" by choosing a small set of templates that covers the space of frequency and number-of-half-cycles without excessive SNR reduction.

In order to avoid (computationally expensive) multiplications in the correlation operation, a three level approximation to the sine wave[7], as shown in Fig. 1b, is used. This introduces little reduction in sensitivity: the SNR for detecting the original burst by the clipped version is 96% of optimum. This three-level waveform is then converted to the digitized form as shown in Fig. 1c where N_1 specifies the number of digital points per half-cycle during which the template takes the value +1 (or -1) with an equivalent definition of N_0; N_0 is chosen to be of order $N_1/3$ to obtain the best SNR. The center frequency of the template is a function of the digital sample rate and is roughly given by:

$$f_{template} \simeq \frac{f_{sample}}{2(N_1+N_0)} = \frac{f_{Nyquist}}{N_1+N_0} \quad . \qquad (2.5)$$

Thus, the range in frequency spanned by a given template or set of templates depends on the digital sample rate. Note that larger values of N_1 and N_0 allow for finer frequency resolution in the location of templates; however, there is a tradeoff as smaller values reduce the number of points per template, reduce the amount of required computation, and make full use of the available digital bandwidth.

The set of digital templates used here are listed in Table 1., categorized by their frequency and number of half-

TABLE 1. Set of templates used in the analysis, organized by frequency and number of half-cycles (based on $f_{Nyquist}$=10 kHz).

Center frequency	Template Code ($N_1N_0N_{hc}$) N_{hc}=1	N_{hc}=2	N_{hc}=4	N_{hc}=6
0.9 kHz		832	834	
1.0 kHz				736
1.3 kHz			624	626
1.4 kHz		522		
1.7 kHz	601		424	426
2.0 kHz				416
2.5 kHz		312	314	316
3.3 kHz	301		214	216
5.0 kHz		202	204	206
≃10 kHz	101	102	104	

cycles. This set was chosen to cover a frequency range of 900 Hz to 5 kHz (when used with data sampled at 20 kHz) and N_{hc}=1-8 with an SNR response better than 80% of maximum for most bursts in these ranges. The increased frequency density of larger N_{hc} filters needed to ensure adequate overlap of response is visible in the table. The N_{hc}=1 templates have response at DC and were included primarily for their computational use as mentioned in Section 3.2.3. Note that template 101 is the identity template in that filtering by it reproduces the original waveform.

3. DATA COLLECTION AND ANALYSIS

3.1. The Data Taking System

The data were taken spaced throughout the night of June 9, 1985 and consist of ten tapes each recording 15 minutes of interferometer output signal at a rate of 20 ksamples/s giving a Nyquist frequency of 10 kHz. The data were taken with a system comprising an A/D with an analogue anti-aliasing filter (four poles with -6 dB at 6.5 kHz), a dual-port memory, and tape drive all under control of an 11/23 computer[8].

One detail that will become important in the results is the operation of the data taking system as it records data to tape. Following standard dual-port operation two buffers are set aside in memory, and the A/D data are written into one buffer while the previously filled buffer is transferred to tape. One complete cycle of the two buffers is here called a 'chunk' of data and consists of 32,608 A/D samples and 160 words of housekeeping information and represents $\simeq$1.6 seconds of data. Thus, for each chunk of data the tape drive performs two data transfers, one at the beginning of the chunk and one at the middle of the chunk.

3.2. Processing of the Data

The data processing steps for the data analysis are briefly outlined below. Typically these steps were performed on one fourth of a tape at a time, producing four PHDs (Section 3.2.4) and four lists of high-SNR events (Section 3.2.5) from each tape.

3.2.1. Check for proper antenna operation. Using either the interferometer output itself (i.e., checking for A/D saturation) or through housekeeping information, times during which the interferometer is unlocked are determined and not analyzed.

3.2.2. Highpass filter the data. Because the interferometer noise is non-white and rises steeply with decreasing frequency below a noise corner frequency of about 1 kHz[8], the data are digitally highpass filtered at 1 kHz using a 61 point FIR filter apodized with a Kaiser window[9]; this is essentially a pre-whitening step.

3.2.3. Template filter the data. The highpass filtered data are then filtered by all 22 templates of Table 1, resulting in 22 output time series. Because the templates are three level the correlation can be performed with only additions and subtractions. Additionally, by saving past values of the templates it is often possible to generate a new template output simply by adding two previously determined values. For example, to generate the i^{th} value of templates 301, 312, 314 and 316 from the new data point h(i) the following four additions are performed:

$$\begin{aligned} 301(i) &= 201(i-1) + h(i) \\ 312(i) &= 301(i) - 301(i-4) \\ 314(i) &= 312(i) + 312(i-8) \\ 316(i) &= 312(i) + 314(i-8) \quad . \end{aligned} \tag{3.1}$$

3.2.4. Form pulse-height distributions. The ouput values of the 22 templates are binned in 22 pulse-height distributions for later examination. These PHDs contain all points output from the template, however once a point is binned its identity is lost.

3.2.5. Save location of high-SNR events. In order to be able to go back and examine the time series of events with high values of SNR, a list for each template is formed of the location and SNR of the 128 events with the highest SNR values detected by that template. This scheme for saving events, thus, does not set a predetermined SNR threshold but rather floats the threshold. If the data were completely Gaussian in character (with no non-statistical tail events), the 128 events in $\simeq$4 minutes of data would set an SNR threshold of approximately 4.2; in practice the threshold varied from 4.5 to 5 becauses of tail events.

3.2.6. Exclude brief high event rate regions. By plotting the location in time of the high-SNR events, striking departures from a uniform distribution of event arrival times may show up. In particular, occasional clumping of events on a time scale of tens of seconds was seen in the data analyzed here and suggests that the clump of events was due to momentary excitation of a component of the antenna, e.g., excitation of the string resonances of the mass suspension wires. These 'regions of excitation' were noted and high-SNR events in them were vetoed from further processing.

3.2.7. Remove multiple detection of high-SNR events. If a single high-SNR burst appears in the data it is likely that it will be detected several times by a single template and, furthermore, be detected by more than one template. For example, a burst of five half-cycles at 1 kHz having an SNR of 10 may be detected by template 736 (see Table 1) with SNR≃9.1 (according to Section 2.2.3) at each of the two times one half-cycle apart when the template 'covers' the burst. When the template-burst overlap is four half-cycles a detection at an SNR of ≃8.2 will be made, etc. Additionally, templates 834, 624, and 626 will also detect the burst at significant SNR levels. Thus these multiple detections of a single burst must be removed from the data to present a PHD in which *tail events correspond to single independent bursts*. This is accomplished by 1) searching all templates for the event with the largest SNR value and saving it, then 2) removing all other events in all other templates which fall within ±30 points (≃3 ms) of this event and 3) continuing this process until the largest SNR values of remaining events are 5 to 6; thus ensuring that only the tail events are subject to this processing.

3.2.8. Combine all saved events into one distribution. Finally, with multiple detections removed, all high-SNR events from all templates can be combined into one PHD which gives information on the agreement of the data with Gaussian statistics. Because the multiple detection removal has essentially classified events according to template, the PHD of events from each template alone can be viewed and together these provide a rough classification of the types of bursts that were detected.

Using non-optimized code the steps above take ≃30 hours of VAX/730 CPU time (typically as four ≃8 hour overnight runs) to process 15 minutes of data, or 6 ms per data point. Much of this time is spent performing the digital highpass filter correlation (essentially 61 real multiply-accumulates per point) and could be avoided through the use of an analogue (or dedicated digital) highpass filter before data recording.

4. RESULTS

The pulse-height distribution (PHD) obtained by applying the data analysis steps outlined above to the 1.7 hours of data is presented in Fig. 2a, plotted as log number-of-events vs. SNR^2 so that the Gaussian distribution of Eq. 2.2 is a straight line. Here many non-Gaussian tail events are visible above SNR≃5.6. Examining the individual PHDs showed that most of these tail events were detected (at the highest

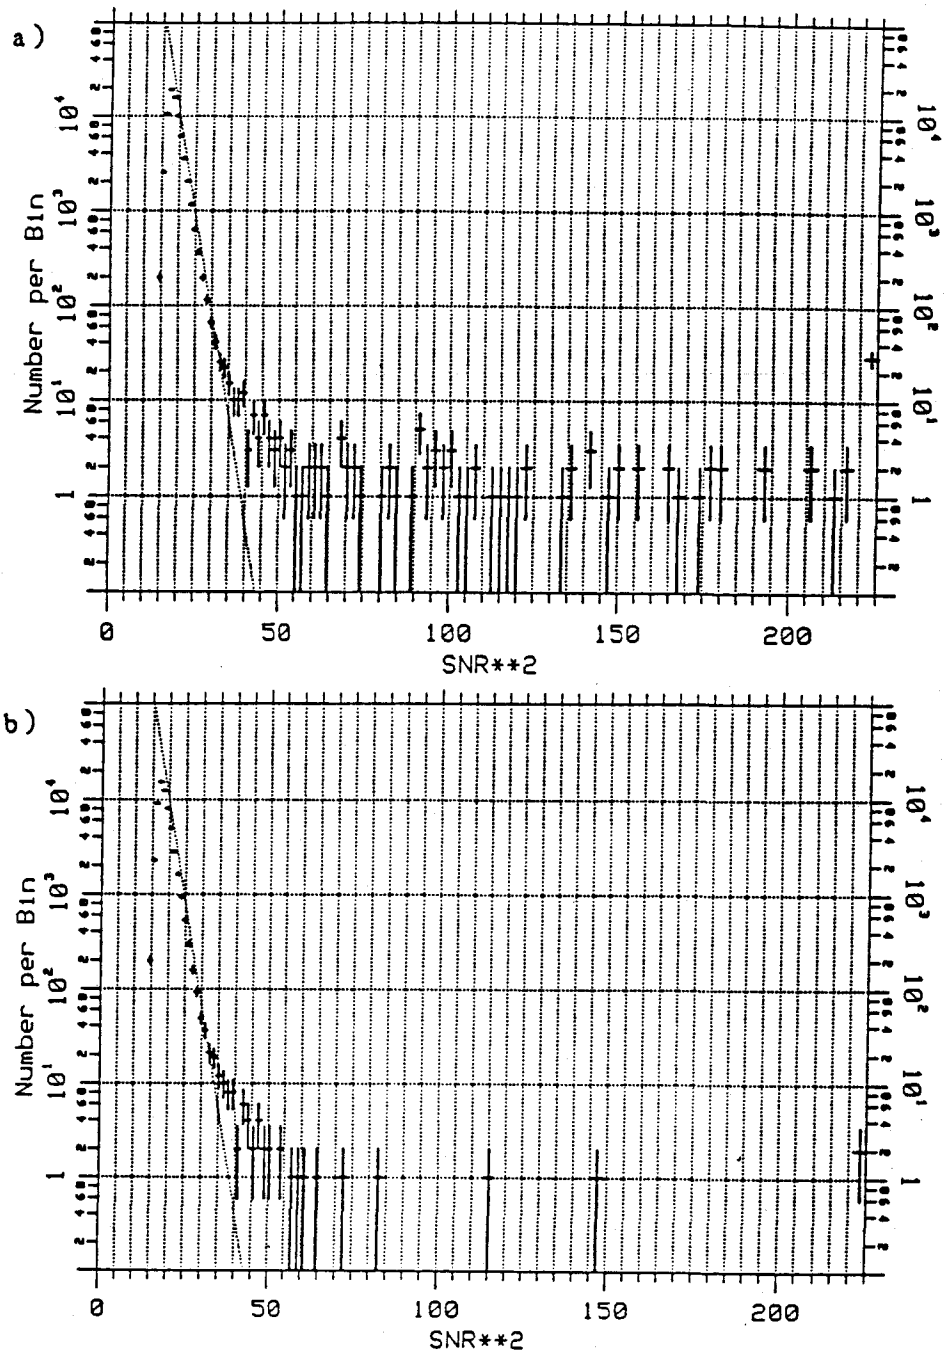

Figure 2. Pulse-height distributions of events. In a) all high-SNR events from all templates are binned. In b) the events from template 102 and the N_{hc}=1 templates (which showed no tail events) have been removed leaving a pulse-height distribution indicative of the instrument operation. In both cases, multiple detection of high-SNR events have been removed (Section 3.2.7).

SNR level) by template 102; if the events detected by template 102 (and those of templates with $N_{hc}=1$) are excluded, the PHD of Fig. 2b results. The template-102 events and the Fig. 2b events are described in the following sections.

4.1. The Template-102 Events

As the difference between Fig. 2a and Fig. 2b shows, the template-102 events dominate the total pulse-height distribution. Looking at the time series of a '102 event, Fig. 3a, reveals that it consists of a single displaced data point. Because of the presence of the analogue anti-aliassing filter this noise source most likely originates at or after the A/D input. A variety of hypotheses regarding the noise source were explored, many of these (single bit errors in the recording of data, A/D conversion errors) made specific predictons about the size of the pulses as recorded on tape (i.e., a power of two in A/D units); this was not supported by the data.

However, when a distribution of the location in time of the template-102 events was made modulo one chunk of data, Fig. 3b, it became clear that the events were most likely associated with feedthrough from the tape drive as it started each write-to-tape operation; these operations occurring twice per chunk of data.

4.2. The Other Events

The remaining non-statistical events, those in Fig. 2b, were examined and many were found to occur at nearly the same time (within $\pm$25 ms) and, like the more extensive clumps discussed in Section 3.2.6, these are viewed as being caused by a single disturbace. There were, however, several events which were isolated in time; most of these were detected at the highest SNR level by the $\simeq$1 kHz 736 template. The largest of these isolated events was detected with an SNR of 23.5 (!) and is shown in Fig. 4a as it appeared to the template filters, i.e., after the highpass filtering. The cause of these events has not been determined, however a gravity-wave origin is unlikely as the amplitude of the tail events detected corresponds to $h_o > 5 \cdot 10^{-14}$ which would require an efficient supernova at a distance closer than that of Proxima Centauri.

To summarize, all non-statistical, non-template-102 events observed in the 1.7 hours of data are categorized in Table 2. An example of a glitch within a region excluded from pulse-height binning is shown in Fig. 4b; this may represent a temporary or near unlocking of the interferometer; however, its origin is also unknown.

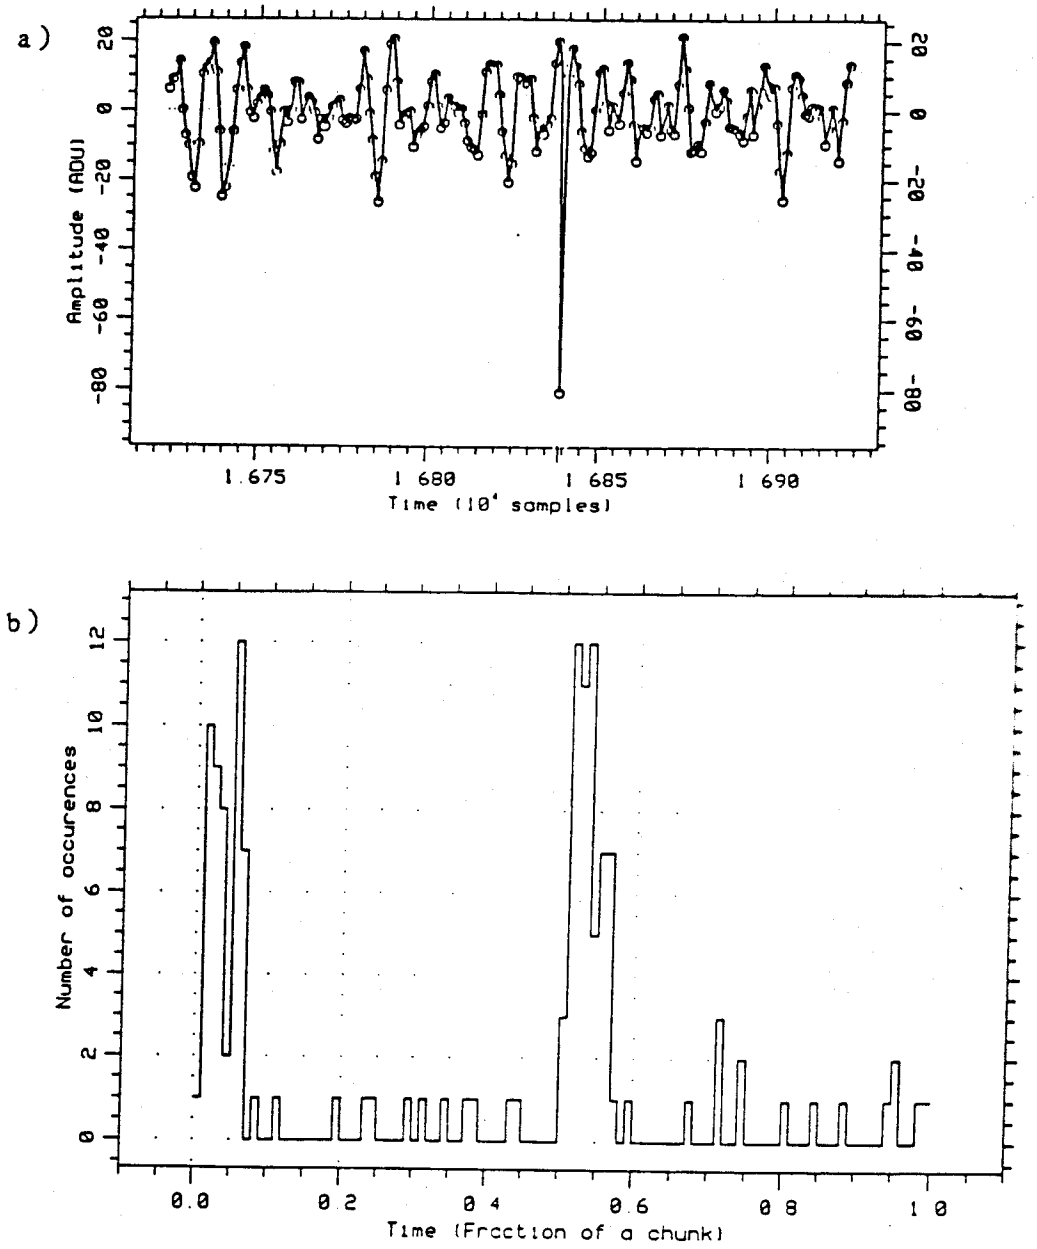

Figure 3. Template-102 events. In a) a typical high-SNR event detected by template 102 is shown. In b) template-102 events have been binned by occurance time within one data taking cycle, a chunk. The two peaks suggest that these events are due to feedthrough originating during write operations to the tape drive which occur twice per chunk.

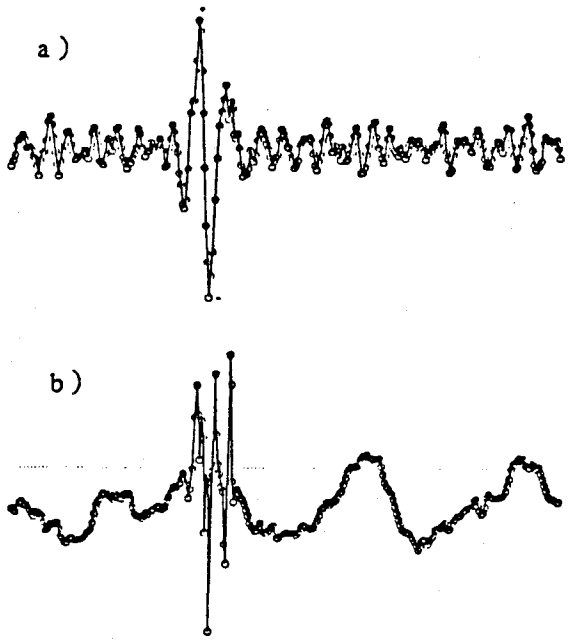

Figure 4. Some transients in the data. Shown in a) is a pulse that was detected by the 736 template at an SNR of ≃23.5; this time series is seen after the digital highpass filter, as it appeared to the template filters. In b) is shown one of several glitches which appeared in the output of the interferometer and were excluded from the pulse-height distribution, in this case shown before digital highpass filtering.

5. CONCLUSIONS AND FUTURE DIRECTIONS

5.1. Conclusions

The results of the application of matched filters to the MIT interferometer are encouraging: the instrument operated at a

TABLE 2. Summary of the Results. All non-statistical occurrences in the 1.7 hours of data are summarized below:

10	Regions excluded from pulse-height binning, each is ≃10-20 seconds of data (Sec. 3.2.6)
12	Isolated events ($h > 5 \cdot 10^{-14}$)
15	Regions (<1 second) with several events
37	Total non-statistical 'events' in 1.7 hours, or ≃ 500 per day

constant rms noise level over the course of one night and the filter output distributions showed agreement with Gaussian statistics out to an amplitude signal-to-noise ratio of 5.5. As listed in Table 2 the actual non-statistical event rate (where a generous interpretation of an event is taken) is of order 500 per day. Finally, the technique succeeded in detecting, isolating, and identifing a transient noise source associated with our data collection hardware, this source would not have been apparent in spectra from the instrument.

5.2. Chirp Sources and Improved Algorithms

The templates described here were chosen to obtain near optimal detection of many of the predicted burst sources of gravitational radiation. Templates for a second class of transient sources, the chirp sources[2], were not explicitly considered because chirps from, say, a system of two 1.4 $M_\odot$ objects, will consist of only a few cycles above the MIT antenna corner frequency of 1 kHz and would therefore be adequately detected by the burst templates.

With interferometric antennas now operating at noise corner frequencies of 500 Hz or less, and given the exciting astrophysical information chirp sources can provide[10], much effort is now directed to detection schemes for these sources. Unlike the burst souces, chirp sources are expected to emit 100 to 1000 or more cycles of detectable waveform; thus, naive matched templates will be very long, requiring large amounts of processing time or custom digital filters. With innovative thinking, for example, the short-time FFT approach[11] or schemes similar to Eq. (3.1), efficient algorithms and hardware will be developed to detect these sources.

5.3. Diagnostic Philosophy

In the process of searching for real signals, data analysis techniques can also detect a variety of noise sources. For example, using a spectrum analyzer we (implicitly) search for continuous sources and identify some noise sources: shot noise, seismic noise, mirror resonances, etc.; in this case the strain noise $h(f)$ is a measure of antenna performance.

With a philosophy similar to that of spectral analysis, the matched filter analysis described here can be viewed as a diagnostic for transient noise sources. In this context very little data-taking hardware is required to acquire transient statistics (for example, neither accurate absolute timing nor continous data-taking are needed). Through this transient analysis one might uncover a variety of potential transient noise sources; some examples are: scattered light noise (which may mimic chirp signals), outgassing bursts in

the vacuum, local electrical disturbances, local mechanical disturbances, laser intermittency, and the up-conversion of low frequency motions to high-frequency transients (already seen and important in bar systems[12]). Because bar detector groups have focussed on searches for transient sources there is a chance to learn from and contribute to their research efforts; for example, plots of (single-antenna) event arrival-time differences[13] would have identified the template-102 events without requiring the hunch that lead to Fig. 3b.

Finally, it is important to emphasize that by searching for both continous and transient sources at the prototype stage will we ensure that all noise sources are understood in LIGO-scale antennas.

ACKNOWLEDGEMENTS

I'd like to thank Bernard Schutz for his organization of this workshop, the members of the Groupe de Recherches sur les Ondes de Gravitation at Orsay for their pre-workshop hospitality, and Elisabeth Busch at MIT for making travel arrangements. This work was supported in part through NSF grant 85-04836-PHY.

REFERENCES

[1] R.L. Forward, Phys. Rev. D 17, 379 (1987).
[2] D. Dewey, in *Proceedings of the Fourth Marcel Grossmann Meeting on General Relativity*, edited by R. Ruffini (Elsevier, Amsterdam, 1986),p. 581.
[3] M. Davis, in this volume.
[4] D.O. North, Proc. of the IEEE 51, 1016 (1963).
[5] C.W. Helstrom, Statistical Theory of Signal Detection (Pergamon Press, Oxford, 1969).
[6] D. Dewey, PhD Thesis, MIT Physics Dept. (1986).
[7] A.R. Whitney, A.E.E. Rogers, H.F. Hinteregger, C.A. Knight, J.I. Levine, S. Lippincott, T.A. Clark, I.I. Shapiro, and D.S. Robertson, Radio Science 11, 421 (1976).
[8] J. Livas, R. Benford, D. Dewey, A. Jeffries, P. Linsay, P. Saulson, D. Shoemaker, and R. Weiss, in *Proceedings of the Fourth Marcel Grossmann Meeting on General Relativity* (Ref. 2),p. 591.
[9] S.M. Bozic, Digital and Kalman Filtering (Edward Arnold Ltd., London, 1979).
[10] B. Schutz, Nature 323, 310 (1986); and in this volume.
[11] S. Smith, in this volume.
[12] P.F. Michelson, J.C. Price, and R.C. Taber, Science 237, 150 (1987) and references therein.
[13] W. Johnson, in this volume.

Proceedings of the NATO Advanced Research Workshop: Gravitational Wave Data Analysis, edited by B.F. Schutz (Reidel, Dordrecht, 1987).

DATA ANALYSIS AS A NOISE DIAGNOSTIC: LOOKING FOR TRANSIENTS IN INTERFEROMETERS

Daniel Dewey
Massachusetts Institute of Technology
Room 3-253
Cambridge, MA 02139
U.S.A.

ABSTRACT. In addition to the ultimate goal of detecting and analyzing source waveforms, data analysis provides a framework for identifying important noise sources. Only by developing and implementing data analysis programs designed to search for expected astrophysical sources will we uncover all relevant instrumental noise sources and have a true picture of an antenna's performance. In particular, the characterization of an interferometric antenna solely by its noise spectrum hides the possible existence of weak and infrequent noise transients. A search for such transients using simple matched filter templates was carried out with the MIT prototype. An analysis of 1.7 hours of data (250 Mbytes) showed an agreement with Gaussian statistics out to an amplitude signal-to-noise ratio of 5.5. Above this signal-to-noise level, transients identified with electrical feedthrough in our data taking system were recorded at a rate of $\simeq$70 per hour, and transients of unidentified origin and various natures occurred at a rate of $\simeq$22 per hour.

1. PHILOSOPHY AND INTRODUCTION

Data analysis strategies should be used at the prototype stage to uncover all relevant noise sources.

Astrophysical theory suggests two broad classes of sources: periodic sources and transient sources. This dichotomy of source types leads naturally to (at least) two noise diagnostics: the (long-time averaged) power spectrum and the pulse-height distribution (PHD) of discrete transient events. Measurement of an interferometer's output spectrum (typically presented as a log-log plot of equivalent strain noise ($h/\sqrt{Hz}$) vs. frequency) has been used almost exclusively by interferometer groups as the specification of instrument performance with little investigation of transient events; Forward's aural data analysis[1] is a

more-appropriate-than-it-first-appears exception. Conversely, bar antenna groups have tended to present bar performance in terms of the distribution of transient events with little direct measurement of the spectral properties of their noise. To some extent these specializations reflect the difference in instrumental bandwidths and philosophies.

This work presents results obtained by applying the matched filter technique to search an interferometer's output for a broad range of transient events. The basic idea of applying matched filters has been described previously[2] and is briefly recounted in Section 2. Section 3 presents specific details of the data collection and data analysis and the results of analyzing 1.7 hours of data are given in Section 4. Finally, Section 5 contains some brief comments on the technique and a discussion of future work.

2. THE MATCHED FILTER AND A SET OF TEMPLATES

2.1. The Matched Filter

The optimum signal-to-noise ratio (SNR) for the detection of a known signal h(t) in the presence of white, Gaussian noise characterized by a spectral density $\tilde{h}$ is a well-known result of detection theory[3,4,5] and is given by:

$$(S/N)_{opt} = \left[\frac{2\int h^2(t)dt}{\tilde{h}^2} \right]^{1/2} \qquad (2.1)$$

How is a detection yielding this optimum (i.e. maximum) SNR achieved? Again the result is well-known: it is obtained through matched filter correlation in which a template of the known signal is correlated with the input waveform to produce an output waveform; this output is essentially the dot product of the template with the input waveform as a function of template position.

Because the input noise to the linear correlation filter is Gaussian (assuming no signal is present) the resulting output time series is again Gaussian. Thus, ideally, a pulse-height distribution (PHD) of the output series will be Gaussian and given by the probability density:

$$p(x) = \frac{1}{\sqrt{2\pi}} e^{-(x/\sigma)^2/2} \frac{dx}{\sigma} \qquad (2.2)$$

where σ is the rms value of the filter output. Gross deviations from this distribution represent non-Gaussian noise, perhaps due to a non-linear component in the signal chain; other deviations, especially 'tail events', may be

caused by the existence of transient noise sources.

It should be mentioned that these matched filter results could have been equivalently expressed in the frequency domain (i.e. using $h(\omega)$ in place of $h(t)$); however a time-domain presentation seems a more appropriate view for the understanding of discrete, transient events. Additionally, with modern digital signal processing techniques it is often more convenient to work in the time domain.

2.2 Effect of a Mismatch of Template and Signal

An initial problem in implementing the matched filter technique is the fact that the shape of the signal is not known. The effects of template-signal mismatch can be easily evaluated in the following framework. If the signal and template are viewed as made up of some number of half-cycle components N_{hc} (defined by the zero-crossings of the waveform) then a mismatch between signal and template can be viewed in terms of a mismatch of: 1) the shape of each component pulse, 2) the amplitude of each component pulse, 3) the spacing of the component pulses, and 4) the total number of component pulses. In general, if a template $h_1(t)$ is used to detect a signal of shape $h_2(t)$ the reduction in SNR from an ideal matched filter is

$$\frac{(S/N)}{(S/N)_{opt}} = \frac{\int h_1(t)h_2(t+\tau)dt}{[\int h_1^2(t)dt \cdot \int h_2^2(t)dt]^{1/2}}, \qquad (2.3)$$

where τ is chosen to maximize the numerator. Note that because of the symmetry of Eq. (2.3) the roles of template and signal may be interchanged. Through Eq.s (2.1) and (2.3) the effect of the four types of mismatch can be determined.

2.2.1 Shape mismatch. In general a shape mismatch of component pulses will have only a slight effect on the SNR for detection. For example, detecting a triangular waveform using a square wave gives $(S/N)/(S/N)_{opt} \simeq 87\%$; for a three-level waveform (Fig. 1b) and a sine wave, it's 96%.

2.2.2 Amplitude mismatch. It can be seen from Eq. (2.1) that a single component pulse of amplitude h_{cp} makes a contribution to the integrand that is proportional to $h_{cp}^2 \Delta t$, where Δt is the length of the component pulse. Thus, only the largest amplitude component pulses are important.

2.2.3 Number mismatch. A mismatch in the number of component pulses leads to a reduction in SNR given by:

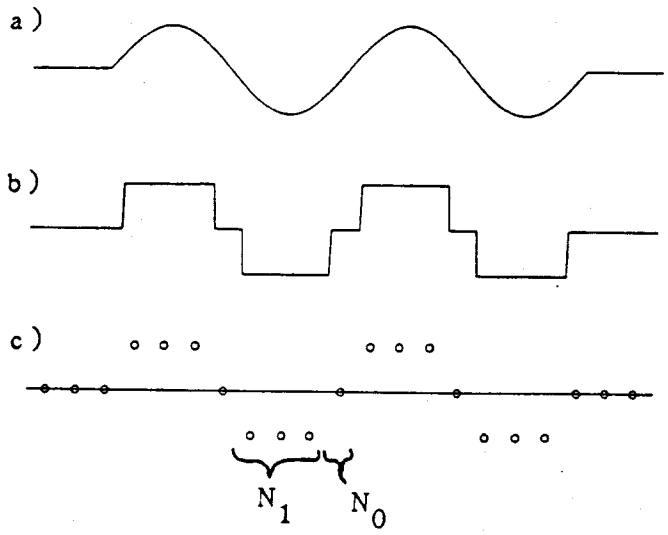

Figure 1. From a canonical burst to a digital template. The four half-cycle burst in a) is replaced by a three level approximation in b) and finally a discrete-time template in c); this template is specified by the three integers N_1=3, N_0=1, and N_{hc}=4, more compactly written as $^1_2$314.

$$\frac{(S/N)}{(S/N)_{opt}} = \left[\frac{\min\{N_{hc,1}, N_{hc,2}\}}{\max\{N_{hc,1}, N_{hc,2}\}} \right]^{1/2} . \qquad (2.4)$$

For example, detecting a 5 half-cycle waveform with a 6 half-cycle template results in $(S/N)/(S/N)_{opt} \simeq 91\%$.

2.2.4 Spacing (or frequency) mismatch. The sensitivity of the SNR to a mismatch in the spacing or frequency of component pulses is a function of the number of half-cycles in the template; the bandwidth decreases with increasing N_{hc}. For templates with N_{hc}=2 a frequency range of ±30% is covered with a minimum $(S/N)/(S/N)_{opt}$ of 80%, while the N_{hc}=6 templates cover only a ±12% frequency range when maintaining a comparable sensitivity.

From the mismatch discussions above it becomes clear that to ensure a high SNR for detection the parameters it is most important to match between template and waveform are the frequency and number of half-cycles.

2.3 The Set of Templates

Many of the expected waveforms from gravitational wave burst sources can be well approximated by several ($\frac{1}{2}$ to 8) half-cycles of a sine wave at a constant frequency[2]; one such canonical burst, made up of four half-cycles, is shown in Fig. 1a. The frequency scale for burst sources (e.g.

stellar collapse and particle-black hole encounters) is predicted to be in the range of 1 kHz or less to as high as 10 kHz. Here, the goal in choosing a set of templates is to reduce the chance of a "false dismissal" by choosing a small set of templates that covers the space of frequency and number-of-half-cycles without excessive SNR reduction.

In order to avoid (computationally expensive) multiplications in the correlation operation, a three level approximation to the sine wave[7], as shown in Fig. 1b, is used. This introduces little reduction in sensitivity: the SNR for detecting the original burst by the clipped version is 96% of optimum. This three-level waveform is then converted to the digitized form as shown in Fig. 1c where N_1 specifies the number of digital points per half-cycle during which the template takes the value +1 (or -1) with an equivalent definition of N_0; N_0 is chosen to be of order $N_1/3$ to obtain the best SNR. The center frequency of the template is a function of the digital sample rate and is roughly given by:

$$f_{template} \simeq \frac{f_{sample}}{2(N_1+N_0)} = \frac{f_{Nyquist}}{N_1+N_0} \quad . \qquad (2.5)$$

Thus, the range in frequency spanned by a given template or set of templates depends on the digital sample rate. Note that larger values of N_1 and N_0 allow for finer frequency resolution in the location of templates; however, there is a tradeoff as smaller values reduce the number of points per template, reduce the amount of required computation, and make full use of the available digital bandwidth.

The set of digital templates used here are listed in Table 1., categorized by their frequency and number of half-

TABLE 1. Set of templates used in the analysis, organized by frequency and number of half-cycles (based on $f_{Nyquist}$=10 kHz).

Center frequency	Template Code ($N_1N_0N_{hc}$)			
	N_{hc}=1	N_{hc}=2	N_{hc}=4	N_{hc}=6
0.9 kHz		832	834	
1.0 kHz				736
1.3 kHz			624	626
1.4 kHz		522		
1.7 kHz	601		424	426
2.0 kHz				416
2.5 kHz		312	314	316
3.3 kHz	301		214	216
5.0 kHz		202	204	206
≃10 kHz	101	102	104	

cycles. This set was chosen to cover a frequency range of 900 Hz to 5 kHz (when used with data sampled at 20 kHz) and N_{hc}=1-8 with an SNR response better than 80% of maximum for most bursts in these ranges. The increased frequency density of larger N_{hc} filters needed to ensure adequate overlap of response is visible in the table. The N_{hc}=1 templates have response at DC and were included primarily for their computational use as mentioned in Section 3.2.3. Note that template 101 is the identity template in that filtering by it reproduces the original waveform.

3. DATA COLLECTION AND ANALYSIS

3.1. The Data Taking System

The data were taken spaced throughout the night of June 9, 1985 and consist of ten tapes each recording 15 minutes of interferometer output signal at a rate of 20 ksamples/s giving a Nyquist frequency of 10 kHz. The data were taken with a system comprising an A/D with an analogue anti-aliasing filter (four poles with -6 dB at 6.5 kHz), a dual-port memory, and tape drive all under control of an 11/23 computer[8].

One detail that will become important in the results is the operation of the data taking system as it records data to tape. Following standard dual-port operation two buffers are set aside in memory, and the A/D data are written into one buffer while the previously filled buffer is transferred to tape. One complete cycle of the two buffers is here called a 'chunk' of data and consists of 32,608 A/D samples and 160 words of housekeeping information and represents $\simeq$1.6 seconds of data. Thus, for each chunk of data the tape drive performs two data transfers, one at the beginning of the chunk and one at the middle of the chunk.

3.2. Processing of the Data

The data processing steps for the data analysis are briefly outlined below. Typically these steps were performed on one fourth of a tape at a time, producing four PHDs (Section 3.2.4) and four lists of high-SNR events (Section 3.2.5) from each tape.

3.2.1. Check for proper antenna operation. Using either the interferometer output itself (i.e., checking for A/D saturation) or through housekeeping information, times during which the interferometer is unlocked are determined and not analyzed.

3.2.2. Highpass filter the data. Because the interferometer noise is non-white and rises steeply with decreasing frequency below a noise corner frequency of about 1 kHz[8], the data are digitally highpass filtered at 1 kHz using a 61 point FIR filter apodized with a Kaiser window[9]; this is essentially a pre-whitening step.

3.2.3. Template filter the data. The highpass filtered data are then filtered by all 22 templates of Table 1, resulting in 22 output time series. Because the templates are three level the correlation can be performed with only additions and subtractions. Additionally, by saving past values of the templates it is often possible to generate a new template output simply by adding two previously determined values. For example, to generate the i^{th} value of templates 301, 312, 314 and 316 from the new data point h(i) the following four additions are performed:

$$\begin{aligned} 301(i) &= 201(i-1) + h(i) \\ 312(i) &= 301(i) - 301(i-4) \\ 314(i) &= 312(i) + 312(i-8) \\ 316(i) &= 312(i) + 314(i-8) \end{aligned} \qquad (3.1)$$

3.2.4. Form pulse-height distributions. The ouput values of the 22 templates are binned in 22 pulse-height distributions for later examination. These PHDs contain all points output from the template, however once a point is binned its identity is lost.

3.2.5. Save location of high-SNR events. In order to be able to go back and examine the time series of events with high values of SNR, a list for each template is formed of the location and SNR of the 128 events with the highest SNR values detected by that template. This scheme for saving events, thus, does not set a predetermined SNR threshold but rather floats the threshold. If the data were completely Gaussian in character (with no non-statistical tail events), the 128 events in $\simeq$4 minutes of data would set an SNR threshold of approximately 4.2; in practice the threshold varied from 4.5 to 5 becauses of tail events.

3.2.6. Exclude brief high event rate regions. By plotting the location in time of the high-SNR events, striking departures from a uniform distribution of event arrival times may show up. In particular, occasional clumping of events on a time scale of tens of seconds was seen in the data analyzed here and suggests that the clump of events was due to momentary excitation of a component of the antenna, e.g., excitation of the string resonances of the mass suspension wires. These 'regions of excitation' were noted and high-SNR events in them were vetoed from further processing.

3.2.7. Remove multiple detection of high-SNR events. If a single high-SNR burst appears in the data it is likely that it will be detected several times by a single template and, furthermore, be detected by more than one template. For example, a burst of five half-cycles at 1 kHz having an SNR of 10 may be detected by template 736 (see Table 1) with SNR≃9.1 (according to Section 2.2.3) at each of the two times one half-cycle apart when the template 'covers' the burst. When the template-burst overlap is four half-cycles a detection at an SNR of ≃8.2 will be made, etc. Additionally, templates 834, 624, and 626 will also detect the burst at significant SNR levels. Thus these multiple detections of a single burst must be removed from the data to present a PHD in which *tail events correspond to single independent bursts*. This is accomplished by 1) searching all templates for the event with the largest SNR value and saving it, then 2) removing all other events in all other templates which fall within ±30 points (≃3 ms) of this event and 3) continuing this process until the largest SNR values of remaining events are 5 to 6; thus ensuring that only the tail events are subject to this processing.

3.2.8. Combine all saved events into one distribution. Finally, with multiple detections removed, all high-SNR events from all templates can be combined into one PHD which gives information on the agreement of the data with Gaussian statistics. Because the multiple detection removal has essentially classified events according to template, the PHD of events from each template alone can be viewed and together these provide a rough classification of the types of bursts that were detected.

Using non-optimized code the steps above take ≃30 hours of VAX/730 CPU time (typically as four ≃8 hour overnight runs) to process 15 minutes of data, or 6 ms per data point. Much of this time is spent performing the digital highpass filter correlation (essentially 61 real multiply-accumulates per point) and could be avoided through the use of an analogue (or dedicated digital) highpass filter before data recording.

4. RESULTS

The pulse-height distribution (PHD) obtained by applying the data analysis steps outlined above to the 1.7 hours of data is presented in Fig. 2a, plotted as log number-of-events vs. SNR^2 so that the Gaussian distribution of Eq. 2.2 is a straight line. Here many non-Gaussian tail events are visible above SNR≃5.6. Examining the individual PHDs showed that most of these tail events were detected (at the highest

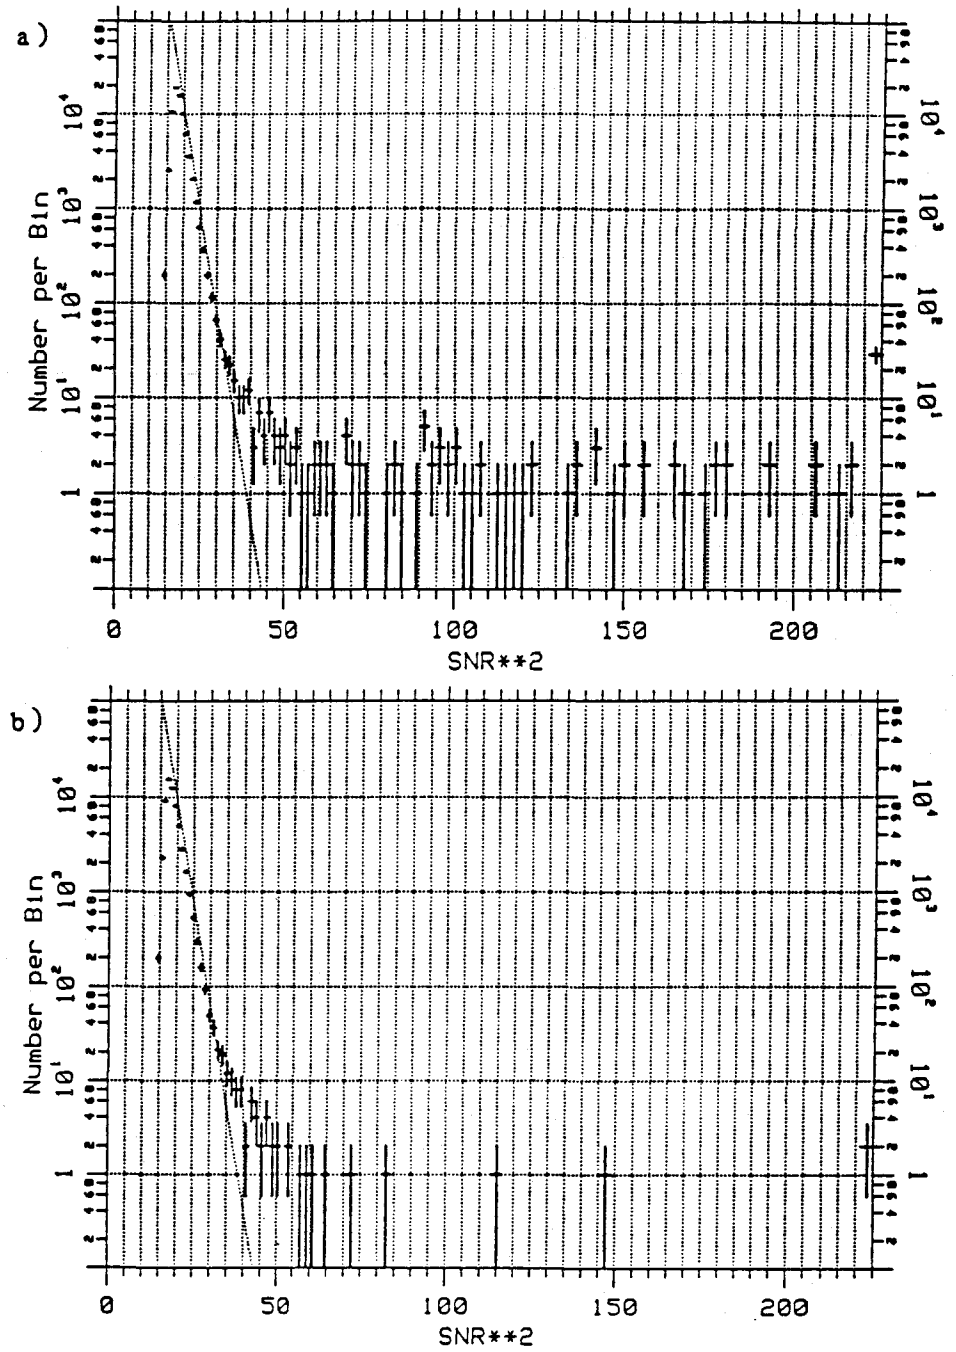

Figure 2. Pulse-height distributions of events. In a) all high-SNR events from all templates are binned. In b) the events from template 102 and the N_{hc}=1 templates (which showed no tail events) have been removed leaving a pulse-height distribution indicative of the instrument operation. In both cases, multiple detection of high-SNR events have been removed (Section 3.2.7).

SNR level) by template 102; if the events detected by template 102 (and those of templates with N_{hc}=1) are excluded, the PHD of Fig. 2b results. The template-102 events and the Fig. 2b events are described in the following sections.

4.1. The Template-102 Events

As the difference between Fig. 2a and Fig. 2b shows, the template-102 events dominate the total pulse-height distribution. Looking at the time series of a '102 event, Fig. 3a, reveals that it consists of a single displaced data point. Because of the presence of the analogue anti-aliassing filter this noise source most likely originates at or after the A/D input. A variety of hypotheses regarding the noise source were explored, many of these (single bit errors in the recording of data, A/D conversion errors) made specific predictons about the size of the pulses as recorded on tape (i.e., a power of two in A/D units); this was not supported by the data.

However, when a distribution of the location in time of the template-102 events was made modulo one chunk of data, Fig. 3b, it became clear that the events were most likely associated with feedthrough from the tape drive as it started each write-to-tape operation; these operations occurring twice per chunk of data.

4.2. The Other Events

The remaining non-statistical events, those in Fig. 2b, were examined and many were found to occur at nearly the same time (within $\pm$25 ms) and, like the more extensive clumps discussed in Section 3.2.6, these are viewed as being caused by a single disturbace. There were, however, several events which were isolated in time; most of these were detected at the highest SNR level by the $\simeq$1 kHz 736 template. The largest of these isolated events was detected with an SNR of 23.5 (!) and is shown in Fig. 4a as it appeared to the template filters, i.e., after the highpass filtering. The cause of these events has not been determined, however a gravity-wave origin is unlikely as the amplitude of the tail events detected corresponds to $h_o > 5 \cdot 10^{-14}$ which would require an efficient supernova at a distance closer than that of Proxima Centauri.

To summarize, all non-statistical, non-template-102 events observed in the 1.7 hours of data are categorized in Table 2. An example of a glitch within a region excluded from pulse-height binning is shown in Fig. 4b; this may represent a temporary or near unlocking of the interferometer; however, its origin is also unknown.

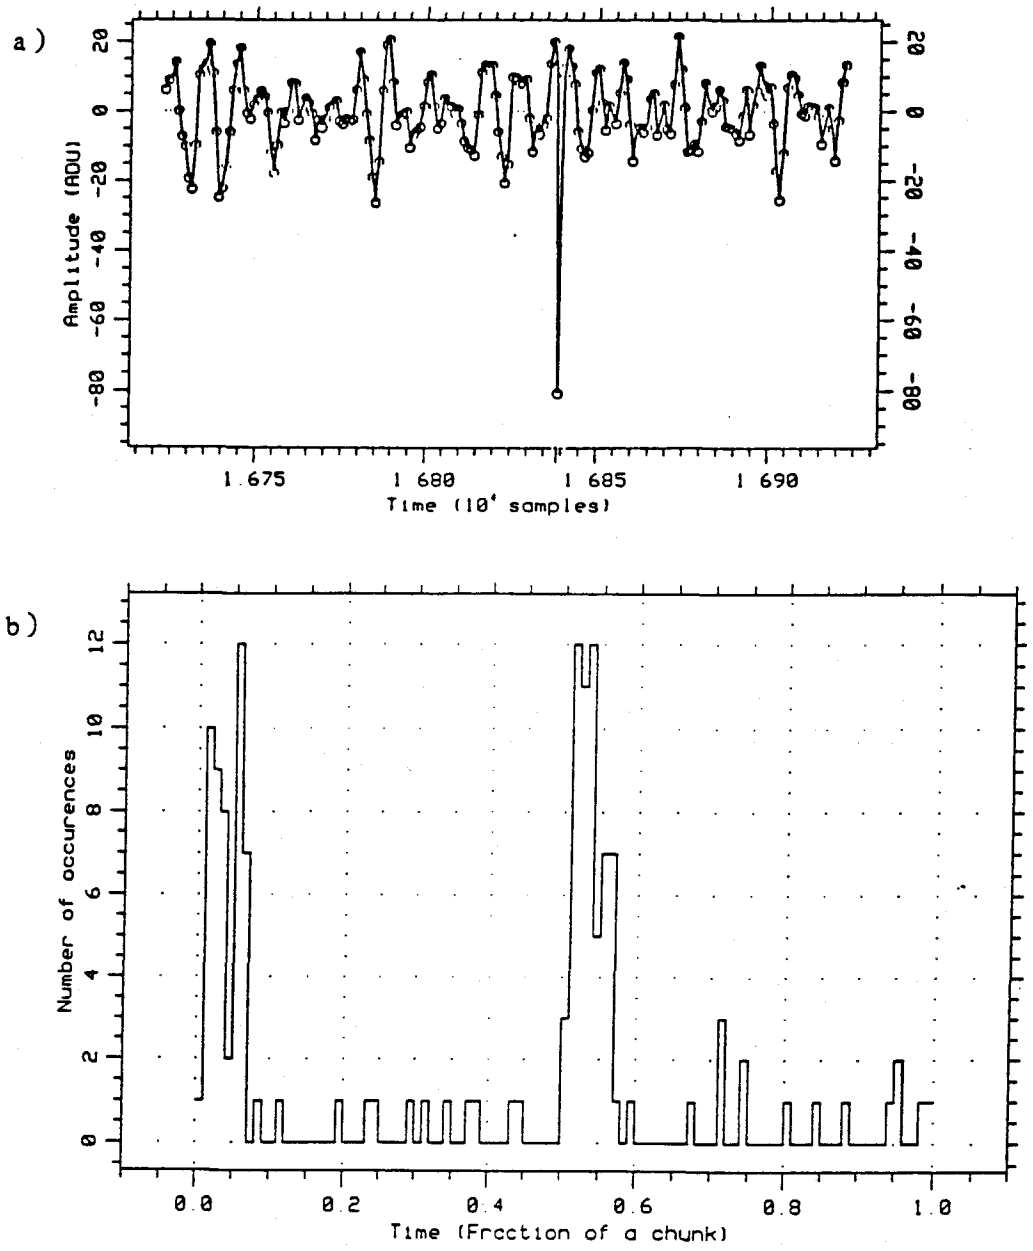

Figure 3. Template-102 events. In a) a typical high-SNR event detected by template 102 is shown. In b) template-102 events have been binned by occurance time within one data taking cycle, a chunk. The two peaks suggest that these events are due to feedthrough originating during write operations to the tape drive which occur twice per chunk.

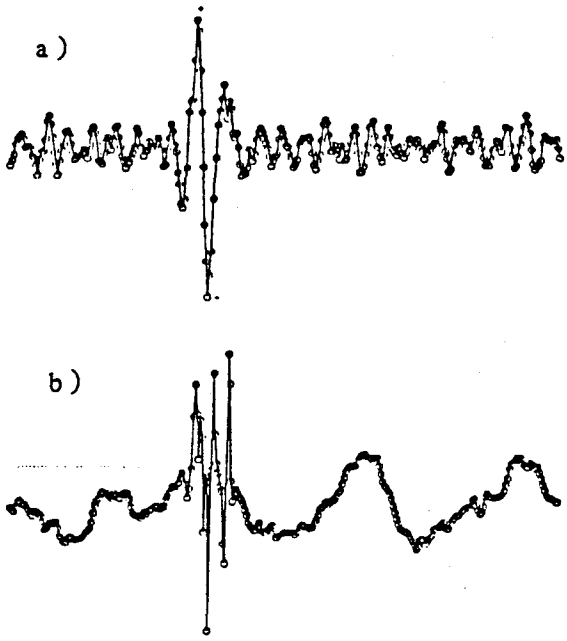

Figure 4. Some transients in the data. Shown in a) is a pulse that was detected by the 736 template at an SNR of ≃23.5; this time series is seen after the digital highpass filter, as it appeared to the template filters. In b) is shown one of several glitches which appeared in the output of the interferometer and were excluded from the pulse-height distribution, in this case shown before digital highpass filtering.

5. CONCLUSIONS AND FUTURE DIRECTIONS

5.1. Conclusions

The results of the application of matched filters to the MIT interferometer are encouraging: the instrument operated at a

TABLE 2. Summary of the Results. All non-statistical occurrences in the 1.7 hours of data are summarized below:

10 Regions excluded from pulse-height binning, each is ≃10-20 seconds of data (Sec. 3.2.6)
12 Isolated events ($h_9 > 5 \cdot 10^{-14}$)
15 Regions (<1 second) with several events

37 Total non-statistical 'events' in 1.7 hours, or ≃ 500 per day

constant rms noise level over the course of one night and the filter output distributions showed agreement with Gaussian statistics out to an amplitude signal-to-noise ratio of 5.5. As listed in Table 2 the actual non-statistical event rate (where a generous interpretation of an event is taken) is of order 500 per day. Finally, the technique succeeded in detecting, isolating, and identifing a transient noise source associated with our data collection hardware, this source would not have been apparent in spectra from the instrument.

5.2. Chirp Sources and Improved Algorithms

The templates described here were chosen to obtain near optimal detection of many of the predicted burst sources of gravitational radiation. Templates for a second class of transient sources, the chirp sources[2], were not explicitly considered because chirps from, say, a system of two 1.4 $M_\odot$ objects, will consist of only a few cycles above the MIT antenna corner frequency of 1 kHz and would therefore be adequately detected by the burst templates.

With interferometric antennas now operating at noise corner frequencies of 500 Hz or less, and given the exciting astrophysical information chirp sources can provide[10], much effort is now directed to detection schemes for these sources. Unlike the burst souces, chirp sources are expected to emit 100 to 1000 or more cycles of detectable waveform; thus, naive matched templates will be very long, requiring large amounts of processing time or custom digital filters. With innovative thinking, for example, the short-time FFT approach[11] or schemes similar to Eq. (3.1), efficient algorithms and hardware will be developed to detect these sources.

5.3. Diagnostic Philosophy

In the process of searching for real signals, data analysis techniques can also detect a variety of noise sources. For example, using a spectrum analyzer we (implicitly) search for continuous sources and identify some noise sources: shot noise, seismic noise, mirror resonances, etc.; in this case the strain noise $h(f)$ is a measure of antenna performance.

With a philosophy similar to that of spectral analysis, the matched filter analysis described here can be viewed as a diagnostic for transient noise sources. In this context very little data-taking hardware is required to acquire transient statistics (for example, neither accurate absolute timing nor continous data-taking are needed). Through this transient analysis one might uncover a variety of potential transient noise sources; some examples are: scattered light noise (which may mimic chirp signals), outgassing bursts in

the vacuum, local electrical disturbances, local mechanical disturbances, laser intermittency, and the up-conversion of low frequency motions to high-frequency transients (already seen and important in bar systems[12]). Because bar detector groups have focussed on searches for transient sources there is a chance to learn from and contribute to their research efforts; for example, plots of (single-antenna) event arrival-time differences[13] would have identified the template-102 events without requiring the hunch that lead to Fig. 3b.

Finally, it is important to emphasize that by searching for both continous and transient sources at the prototype stage will we ensure that all noise sources are understood in LIGO-scale antennas.

ACKNOWLEDGEMENTS

I'd like to thank Bernard Schutz for his organization of this workshop, the members of the Groupe de Recherches sur les Ondes de Gravitation at Orsay for their pre-workshop hospitality, and Elisabeth Busch at MIT for making travel arrangements. This work was supported in part through NSF grant 85-04836-PHY.

REFERENCES

[1] R.L. Forward, Phys. Rev. D 17, 379 (1987).
[2] D. Dewey, in *Proceedings of the Fourth Marcel Grossmann Meeting on General Relativity*, edited by R. Ruffini (Elsevier, Amsterdam, 1986), p. 581.
[3] M. Davis, in this volume.
[4] D.O. North, Proc. of the IEEE 51, 1016 (1963).
[5] C.W. Helstrom, Statistical Theory of Signal Detection (Pergamon Press, Oxford, 1969).
[6] D. Dewey, PhD Thesis, MIT Physics Dept. (1986).
[7] A.R. Whitney, A.E.E. Rogers, H.F. Hinteregger, C.A. Knight, J.I. Levine, S. Lippincott, T.A. Clark, I.I. Shapiro, and D.S. Robertson, Radio Science 11, 421 (1976).
[8] J. Livas, R. Benford, D. Dewey, A. Jeffries, P. Linsay, P. Saulson, D. Shoemaker, and R. Weiss, in *Proceedings of the Fourth Marcel Grossmann Meeting on General Relativity* (Ref. 2), p. 591.
[9] S.M. Bozic, Digital and Kalman Filtering (Edward Arnold Ltd., London, 1979).
[10] B. Schutz, Nature 323, 310 (1986); and in this volume.
[11] S. Smith, in this volume.
[12] P.F. Michelson, J.C. Price, and R.C. Taber, Science 237, 150 (1987) and references therein.
[13] W. Johnson, in this volume.