SHORT TERM DISTORTION IN DYNAMIC NOISE-FILTERS

by

John Nelson Wright

Submitted in Partial Fulfillment

of the Requirements for the

Degree of Bachelor of Science

at the

Massachusetts Institute of Technology

May, 1976

Signature of Author.. Signature redacted

Department of Electrical Engineering, Date

Signature redacted

Certified by...................

Thesis Supervisor

Signature redacted

Accepted by..

Chairman, Departmental Committee on Theses

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

I would like to thank both Dr. Amar Bose and Mr. Joe Veranth for their help and quidance throughout the preparation of this thesis. Thanks is also due to Dr. William Henke for his assistance with the MITSYN computer system. Finally, I would like to thank my wife Linda, without whose support my education would have been infinitely more difficult.

I ABSTRACT

The concept of dynamic noise-filtering has seen a revival recently, and several such filters have been described and marketed. The principle of operation in all such schemes involves modulating the cuttoff frequency of a low pass filter with a program-derived control signal.

The non-linearity of such a system has not, to our knowledge, been treated quantitatively in the literature.

In this thesis, we present:

1) descriptions of dynamic filtering systems that have been proposed
2) the computer implementation of a simple dynamic filter
3) spectral analysis of the filter's output signal when presented with simple test signals

We conclude that any spectral components generated by the filter's action are of small enough magnitude to be below our measurement system's resolution.

II INTRODUCTION

The Problem of Noise in Audio Reproduction

One of the fundamental limitations in recording and reproducing audio material has been, since the birth of the art, noise in the transmission medium. The dynamic range of the human ear can extend to 120 dB or more, and one would be hard pressed to find any system with a comparable range. The current crop of professional tape machines, for comparison, can barely approach a signal to noise ratio of 70 dB.

A relevant question here, of course, is whether such a transmission channel need duplicate the range of the ear, and the answer is, thankfully, no. Even the quietest recording environments have a noise level around 20 dB SPL, and a very quiet listening room has a typical background level of 40 dB SPL. If such a system were to reproduce peaks of 120 dB SPL, comparable to those in many live musical performances, then a signal to noise ratio of 80 dB would be required. The dynamic range of commercially available tapes, discs, and broadcasts rarely exceeds 60 dB, and the prospects of their improving dramatically in the near future are not good.

The time-honored approach to such a problem, as an alternative to fixing the channel, is to "fix" the signal. That is, subject the signal to some sort of processing, at one or both ends of the channel, in some manner that reduces the perceived noise. Historically, this approach led to the pre-emphasis/de-emphasis schemes currently standard in all tapes, discs, and broadcasts. Several more exotic schemes have been proposed and used, and they all involve processing that is basically non-linear.

The Advent of Non-linear Systems for Noise Reduction

There are, of course, many types of noise that a system is subject to; the two most common being random (Gaussian) noise and impulse-like noise, such as record scratches. Although some specialized systems have been designed to deal with the latter, [1],[2] most noise-reduction approaches address themselves primarily to the former, including the approaches this paper is henceforth concerned with.

In the area of non-linear processing there are two distinct approaches, single ended systems and two ended systems; that is, systems designed to operate at only one end of the channel and those designed to operate at both ends.

In the latter category, the Dolby system [3] is the best known and the most important historically. Also in this class are such classical compressor/expanders as the dbx Noise Reduction System and the Burwen Labs Noise Eliminator.

The failing of this approach is that it leaves untouched a huge body of existing recordings, all degraded to some extent by noise. Clearly, a desirable system would be one which could process such recordings to reduce the perceived noise level with little or no effect on the program material. Such a system would necessarily be a single ended one. A linear filter fails because it invariably intrudes on the program material if it is to be effective on the noise. Non-linear approaches have included level-dependant gating and expansion in the low end of the dynamic range, but the most successful approach has been dynamic filtering.

III THE DYNAMIC FILTER AS A NON-LINEAR SYSTEM

A Generalized Description

Conceptually, the operation of a dynamic filter is simple. Its heart is a low pass filter, whose cutoff frequency is positioned in a dynamic manner such that, at any instant, all the major spectral components of the program material are below it and any higher frequency noise components are attenuated. In other words, by looking at the input signal, it determines the optimum position for a low pass noise filter and then assumes that position.

The simplest block diagram looks like:

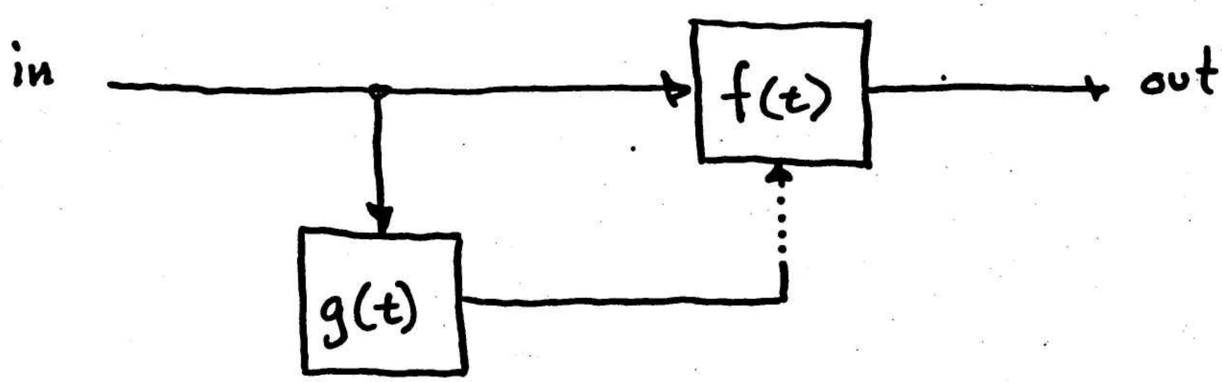

(where f(t) represents the actual filter and g(t) represents the control for the filter).

The system fails to meet linearity criteria; it is neither homogeneous nor additive. That is, if:

$$F[x(t)] = y(t)$$

then, in general:

$$F[A\,x(t)] \neq A\,y(t)$$

(where F represents operation by the filter and A is an arbitrary constant). Also, if:

$$F[x_1(t)] = y_1(t)$$

and:

$$F[x_2(t)] = y_2(t)$$

then, in general:

$$F[x_1(t) + x_2(t)] \neq y_1(t) + y_2(t)$$

The former case can be illustrated by assuming x(t) to be a low amplitude high frequency signal. Below some point, the control section will assume it to be noise, keep the filter closed, and x(t) will be attenuated; but if x(t) is increased to some large amplitude, the filter will assume that it constitutes program material and will open up to allow it through essentially unchanged.

The latter case can be seen by considering a similar situation, with $x_1(t)$ a low amplitude, high frequency signal and $x_2(t)$ a high amplitude, high frequency signal. When the two imput signals are summed, the output component due to $x_1(t)$ is larger than $F[x_1(t)]$ because the magnitude of $x_2(t)$ has caused the filter to open up.

It should be stressed here that the system is not time variant in the formal sense; that is, if:

$$F[x(t)] = y(t)$$

then:

$$F[x(t+\tau)] = y\ (t+\tau)$$

(where τ is an arbitrary time interval). In other words, a specfic input generates the same response regardless of when the input is applied. This is an important point in any detailed analysis of such a filter's input-output characteristics, as this thesis intends to present, and it is a point that is consistantly evaded in the literature. Indeed, virtually all the descriptions of such devices make the implicit assumption that they are (or can be treated as) linear, time varying systems. There is, to our knowledge, no a priori justification for such an assumption.

Common Audible Problems

The reason for this conceptual distortion arises from the fundamental limitations, from a psychoacoustic standpoint, of the whole dynamic filtering concept. The most widely recognized problem (and one shared with virtually any non-linear noise reduction scheme) is that of background noise modulation, usually refered to as "pumping", "breathing", etc.

The problem stems from the fact that, as the filter opens up to track the program spectrum, the level of noise will necessarily increase. If, due to filter mistracking or excessively noisy program material, this increase in noise is not sufficiently masked, then it will be audible as a hiss that varies along with the program material. The audibility of background noise thusly modulated is often higher than the same unmodulated noise, and most critical listeners agree that it is more irritating. With some marginal sources the problem appears not as breathing but as a slight fuzzy sound riding along with high frequency signals that often sounds like modulation distortion, such as would be produced by phono cartridge mistracking.

This problem is illustrated graphically in figure 3.1, using the model and system described in part V; time domain data is on the left and frequency domain data on the right. Input to the filter is white noise. As the filter opens up (reading from bottom to top in the illustration), the spectrum and time domain character of the noise can be seen to modulate.

The designers of these systems, recognizing this problem, have concentrated on producing schemes to minimize it. The resulting focus on the filter's characteristics as they change in time has apparently obfuscated the fact that, from a

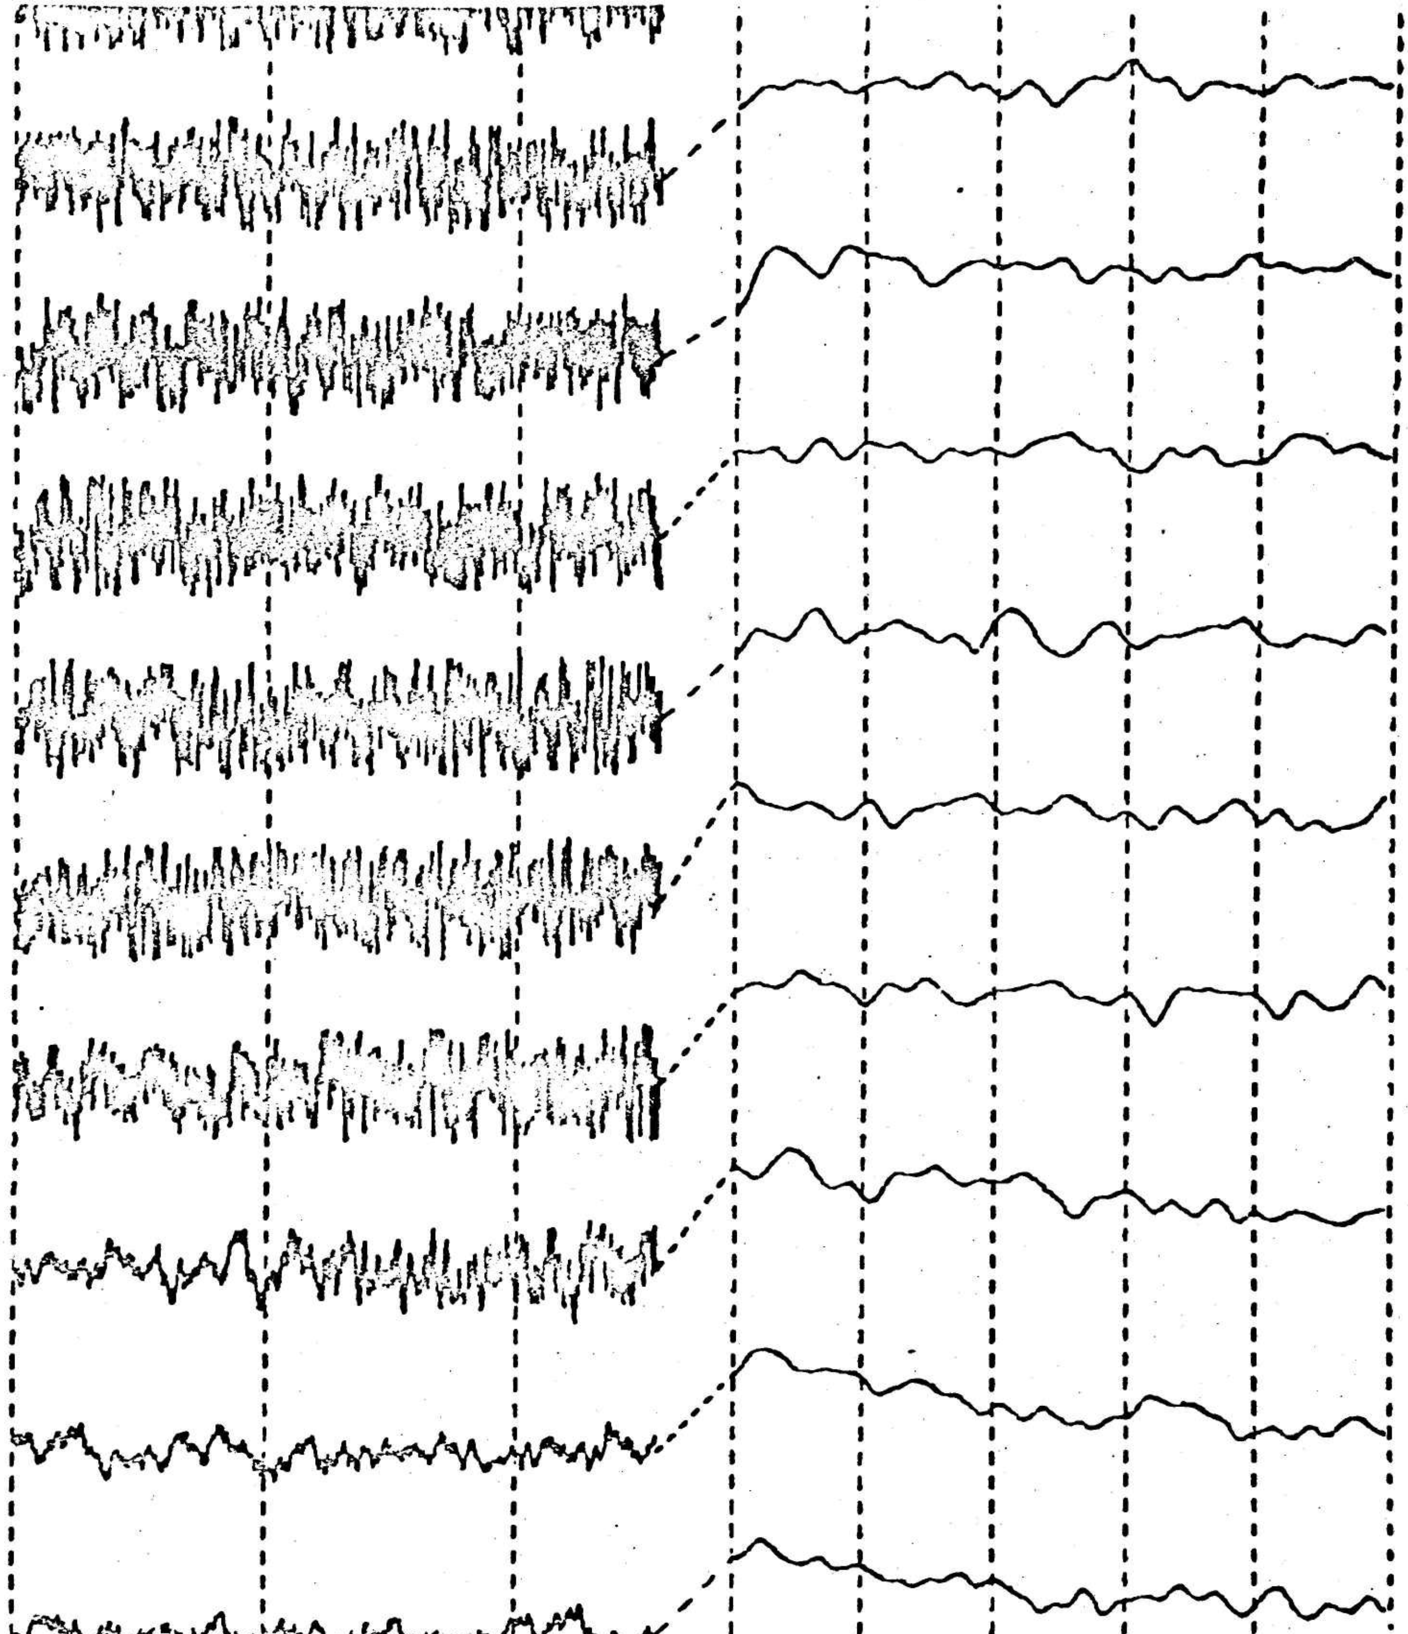

TO= 0.030 W=0.0256 DT= 0.020 CCOF=0.0050 DATE= 0 GDF=1000. 50 300 150

noise modulation by a dynamic filter

Figure 3.1

rigorous standpoint, the characteristics are not time varying but program varying.

This may seem like an academic point, and, indeed, for solving the above problem it provides little insight. However, the recognition of the dynamic filter as a non-linear device raises the possibility of spectral distortion arising in the program material due to its action.

The Possibility of Artifacts Resulting from the Filter's Non-linearity

This can be approached from the standpoint of homomorphic signal theory. If the output of the filter can be described as:

$$y(t) = A(t)B(t)$$

(where A(t) is the modulatory component and B(t) is the vibratory component) then the design effort to date has been focused on defining and optimizing the characteristics of A(t) in order to minimize the problem of noise modulation. It has been implicitly assumed that such modifications will not affect B(t).

It is the purpose of this thesis to explore that assumption. In the succeeding sections we will describe in some detail dynamic filters that have been built and marketed, the computer implementation of a simple dynamic filter for this thesis, and the spectral output of that filter when presented with simple test signals. Our purpose is to search for spectral components that are added to the signal by the filter's action, and thus constitute non-linear distortion.

IV IMPLEMENTATION APPROACHES

H. H. Scott's System

The first dynamic filter was described by H. H. Scott in 1947.[4] It was marketed briefly and has become a classic of sorts. The low pass filter was built around a reactance tube acting as a voltage variable capacitor. The control portion of the scheme involved a bandpass filter centered around 2.5 kHz, followed by a rectifier and a filter. A block diagram is shown in figure 4.1. The rationale was that energy in the octave around 2.5 kHz implies the existance of higher harmonics, and the filter should then open up.

Burwen's System

More recently, Burwen [5],[6] has described and marketed (in professional and consumer versions) a dynamic filter that is the direct descendant of the Scott scheme. While it is considerably more complex in detail, it is conceptually a very similar system. The low pass filter is implemented with operational amplifiers and a two quadrant multiplier in a feedback loop. It is a single pole filter with a 6dB/octave rolloff.

The control section consists of a weighted high pass filter, a full wave rectifier, a peak detector, a non-linear

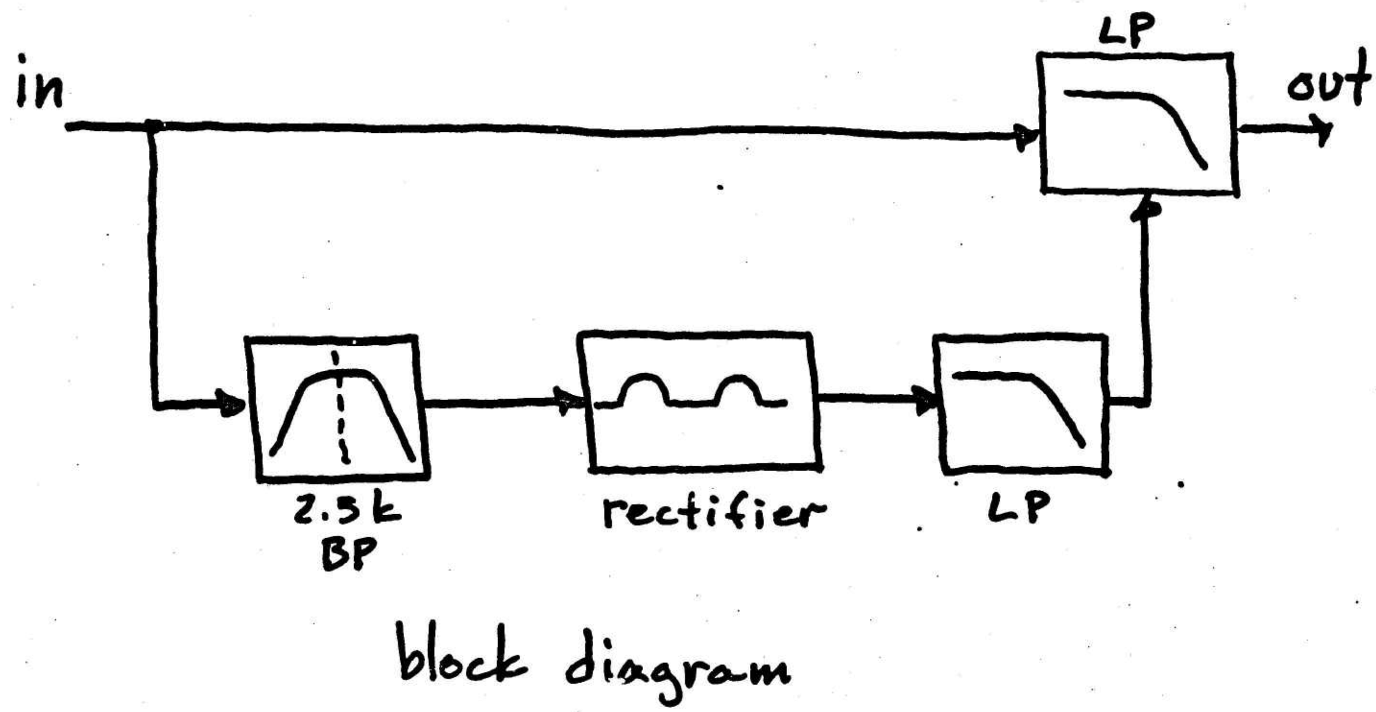
in
LP
out
2.5k
BP
rectifier
LP
block diagram

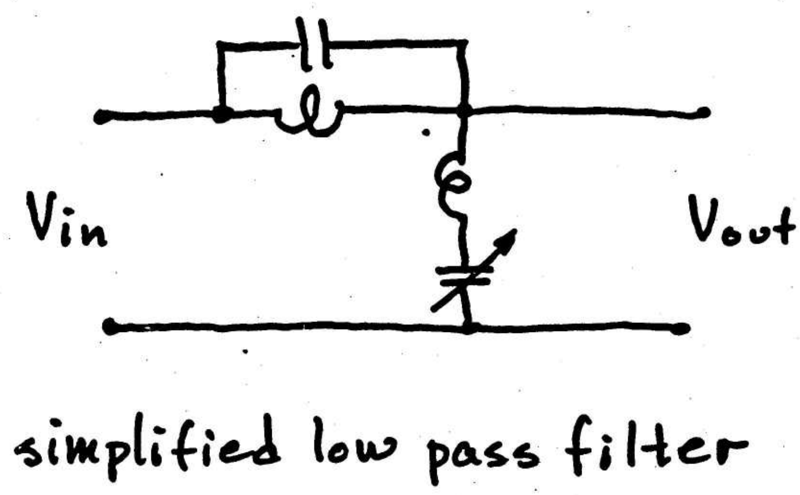
Vin
Vout
simplified low pass filter

Figure 4.1

filter, and a limiter. The non-linear filter works in such a manner that, when presented with a large $\frac{dv}{dt}$, it exhibits a short time constant and fast rise time, and in steady state provides good filtering characteristics to keep the control voltage free of A.C. components. The limiter controls the maximum and minimum frequencies to which the filter will open; in the consumer version these are 500 Hz and 20 kHz. The attack time constant is approximately 0.5 ms, the decay time constant 50 ms. A block diagram is shown in figure 4.2.

The assumption implicit in the operation of the control section is that the energy in the weighted high pass filter corresponds directly to the proper cutoff point of the filter.

Ives' System

A conceptually more involved approach has been proposed and implemented by Fred Ives.[7],[8] The basic idea is to derive the control signal not from the energy in a fixed band, but rather from the characteristics of the output signal itself. The result is a system employing negative feedback to force the filter to track the upper corner of the program spectrum, and Ives has dubbed it dynamic spectral filtering.

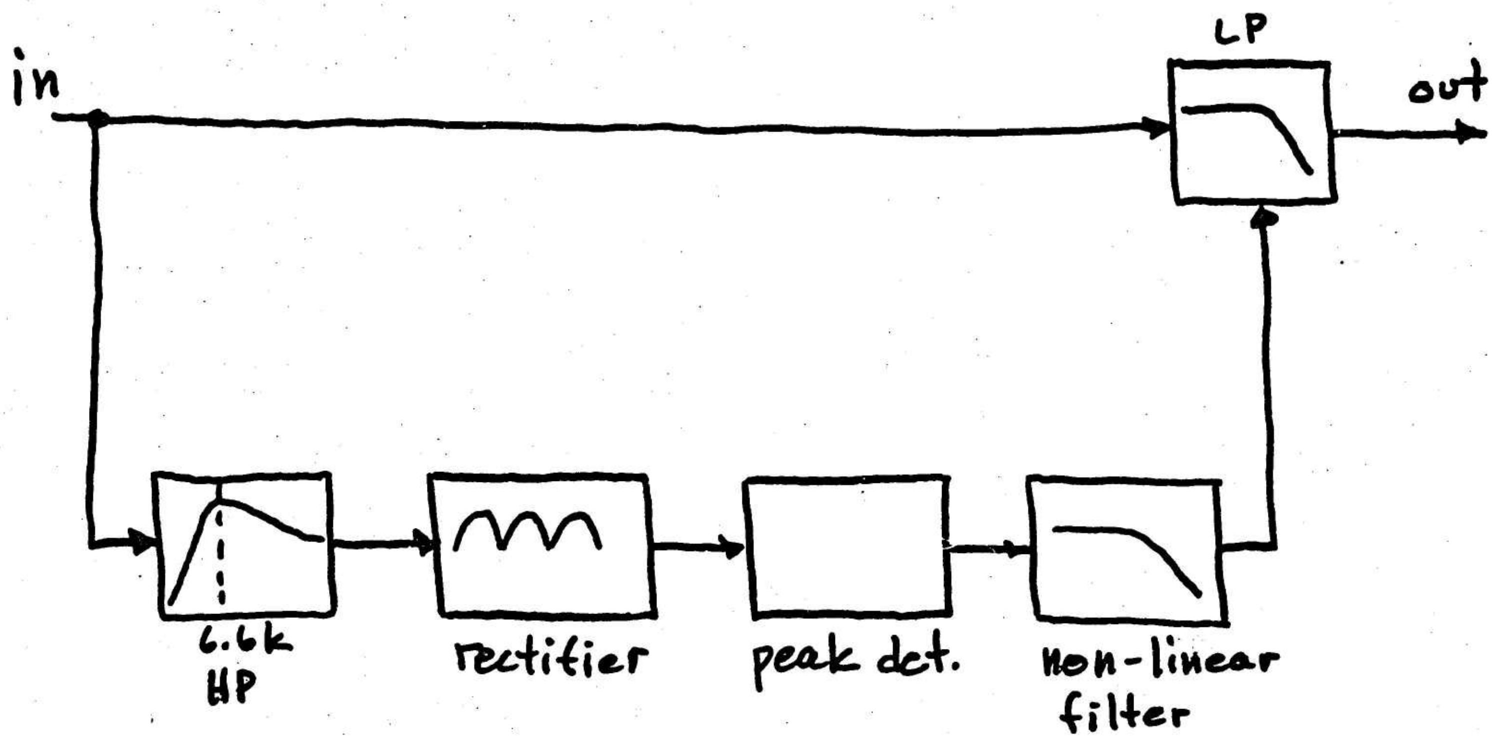

block diagram

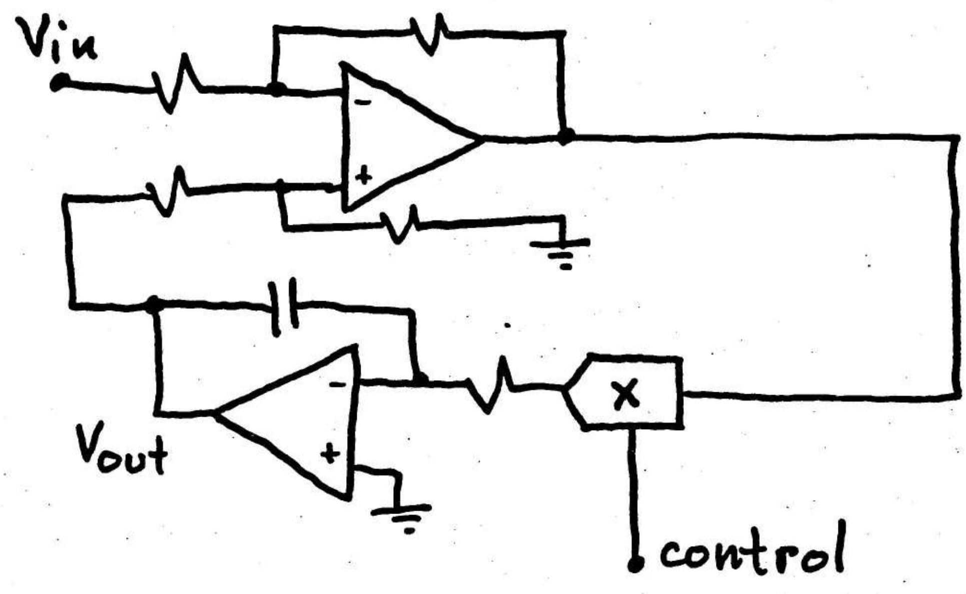

simplified low.pass filter

Figure 4.2

The filter is a state variable configuration, using operational amplifiers and pulse width modulation multipliers in a feedback loop. The filter provides low pass, band pass, and high pass outputs, the latter two being used in the control section.

The control section passes these two outputs through rectification and peak detection then, using a log ratio circuit, takes the ratio of the energy in these two bands. The feedback loop works to keep this ratio constant. A limiting circuit controls the two extremes of the filter's cutoff points. The attack time constant is 10 ms. and the decay time constant is 50 ms. A block diagram is shown in figure 4.3.

The assumption in this scheme is that the shape of the spectrum of the program material remains constant, and that a cutoff point thusly determined is optimum regardless of the level of the program.

The Phase Linear System

The most recent dynamic filter to be described and marketed is the Phase Linear system.[9],[10] The actual filter is fairly complex in design, but from the simplest conceptual viewpoint it remains a program controlled low pass filter.

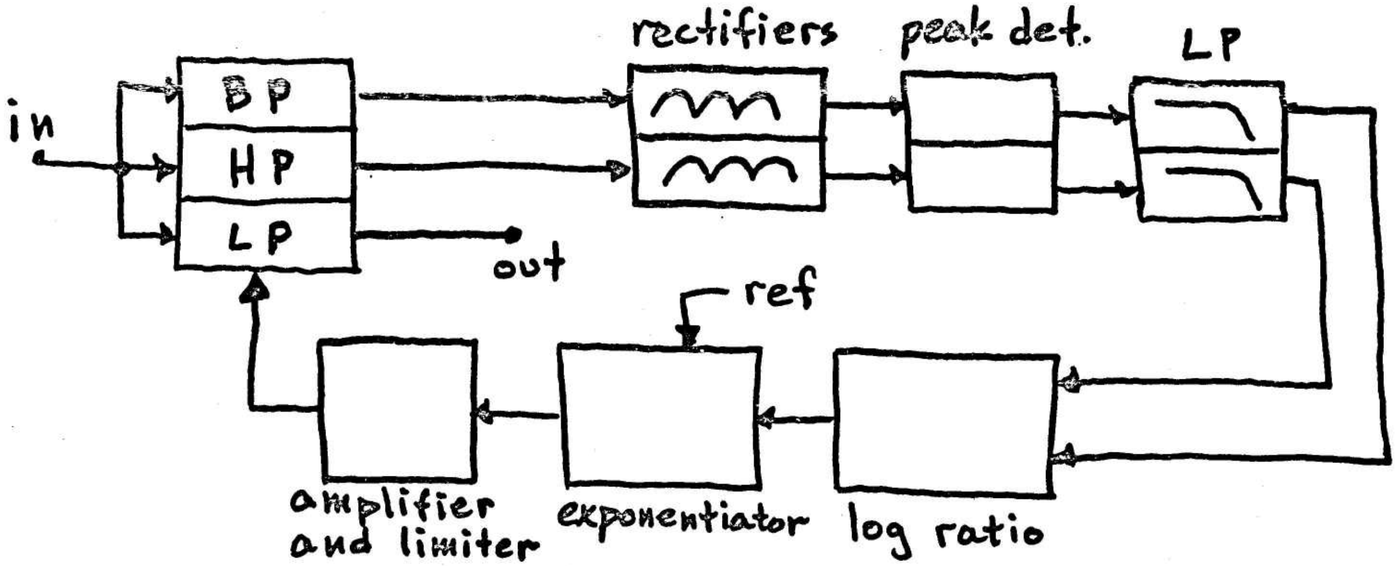

block diagram

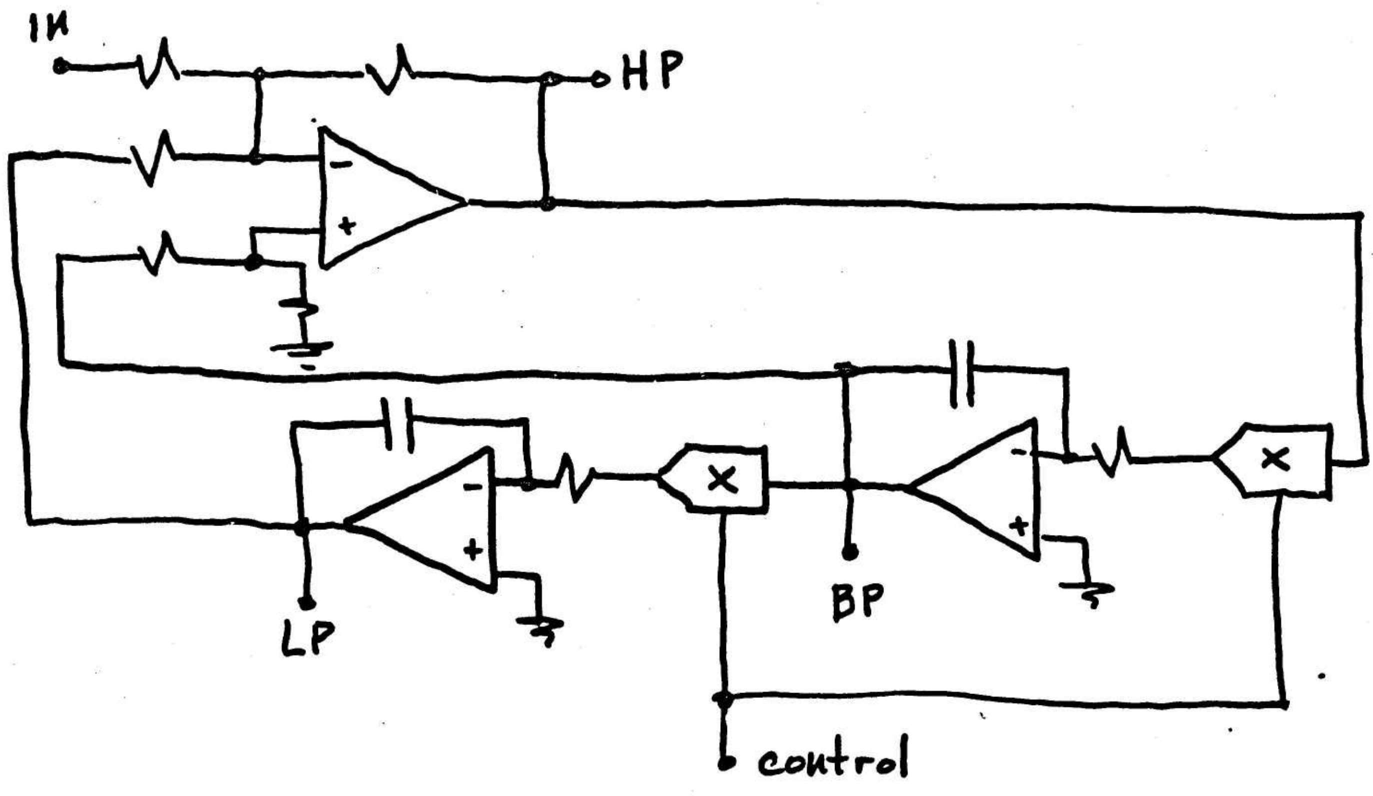

simplified filter

Figure 4.3

Four bandpass filters, centered at 1.3 kHz, 2.5 kHz, 5 kHz, and 10 kHz form the heart of the system. These filters are of the constant voltage type, so that summing their outputs, and adding the signal remaining in the base band below 900 Hz, forms a unity gain system. This condition corresponds to a fully open filter. The shape of the filter is controlled by separately expanding the signal from each of the bandpass filters, so that low signal levels correspond to a relatively large amount of attenuation, and large signals result in unity gain. The control sections for each bandpass output operate separately, with the exception of an interconnect that keeps the gain of any output at least as great as the next higher frequency bandpass output. The resulting system is at any instant a low pass filter with a monotonic rolloff, but the slope of the rolloff as well as the cutoff frequency is determined by the details of the program material.

The details of the control circuitry are quite complex, and beyond the scope of this paper. Briefly, though, each bandpass output is multiplied by a control signal composed of the rectified, filtered, and logarithmically weighted bandpass output. The attack time constant is 3 ms., the decay time constant is 50 ms. A block diagram is shown in figure 4.4.

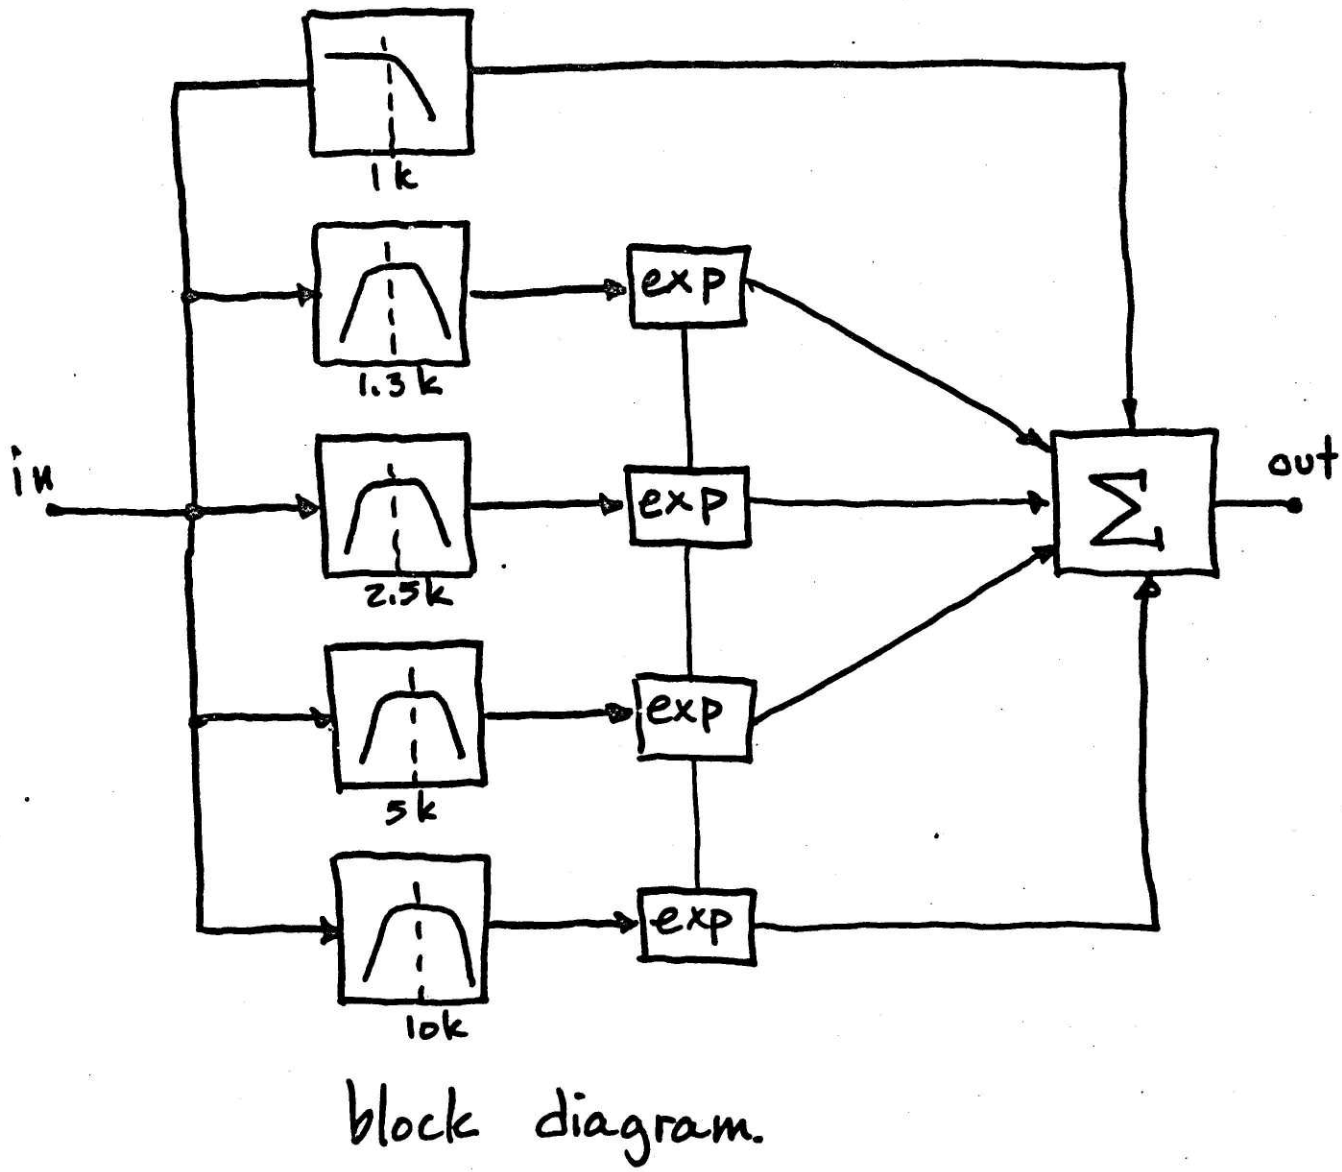

Figure 4.4

The Simple Model Used for the Experiments

It was desired to implement a simple dynamic filter to study the effects of its action on the spectral content of the program. The requirements were judged to be:

1) that it display characteristics similar to all the available systems when presented with simple single frequency test signals

and

2) that it be simple and general enough that the results obtained could be applied to any similar future system.

The simplest variable filter, in network terms, would be an RC section using a variable R or a variable C.

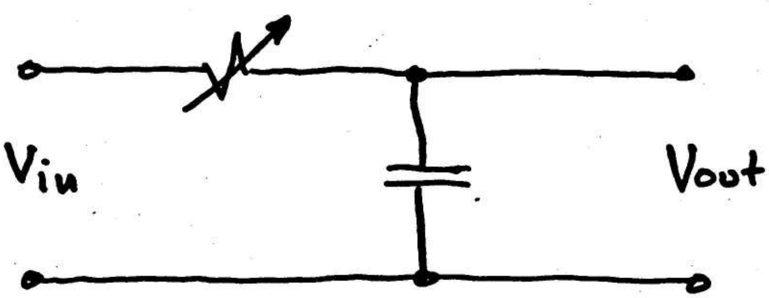

or

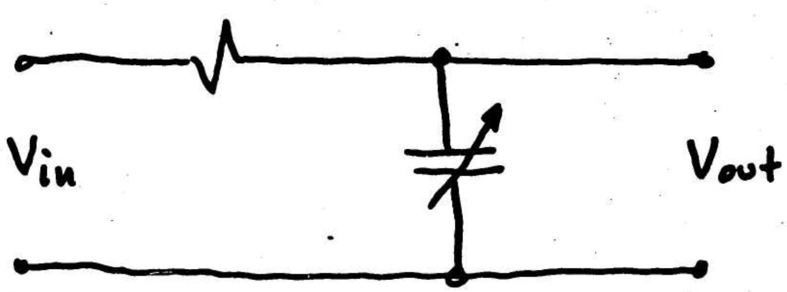

There is, however, an important difference between the two approaches. This can be seen by considering the filter in D.C. steady state, such that $V_{in} = V_{out}$, and having a step input into the control line. In the case of a variable R, V_{out} will remain unchanged. In the case of a variable C, though, V_{out} will undergo a step change followed by a decaying exponential. The initial step is necessitated by conservation of charge on the capacitor, such that:

$$Q = C_1 V_{out_1} = C_2 V_{out_2}$$

A review of all the implementations previously described reveals that only Scott's system employs the variable C approach. He apparently recognized the proplem, because he spends some effort in his design to keep thumps from the action of the control section from entering the program material.

The state variable approach used in the Burwen and Ives system is of the variable R type, as can be seen by considering the same D.C. steady state experiment:

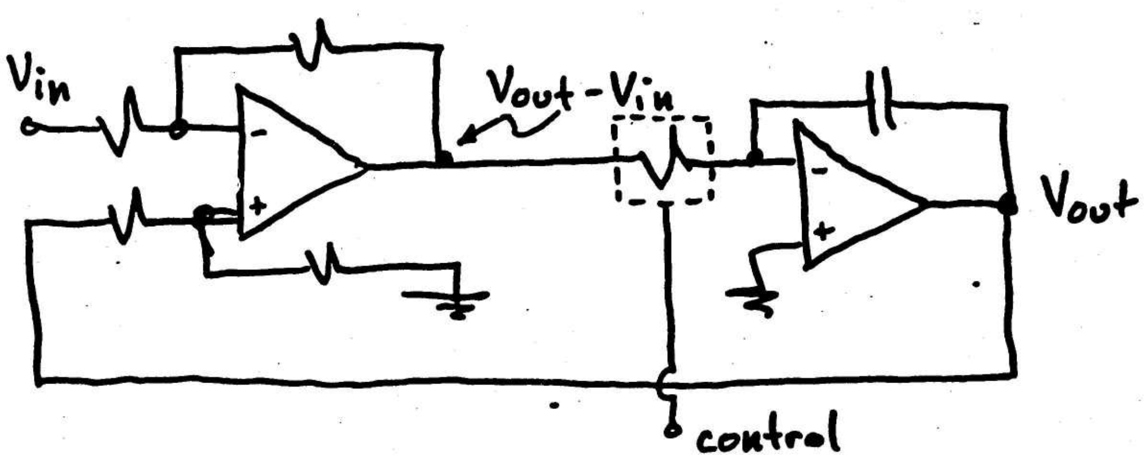

With a D.C. input, the quantity $V_{out} - V_{in}$ will be driven to zero by the feedback loop, and a step on the control line should thus have no effect on the equilibrium of the system.

The Phase Linear scheme is not amenable to the same sort of simple network analysis, since it obtains its low pass characteristics by a different approach. However, considering the same test as before, it can be seen that it is, at least in some sense, equivalent to the variable R approach. The bandpass filters will keep any D.C. components from entering the expanders, so a step on the control line should have no effect on the output.

From this, it was decided that the model for the experiments should be of the variable R type.

The actual computer-implemented model is a single pole filter, whose cutoff point is determined by the energy in the test signal. The control section includes rectification, peak detection, and filtering, followed by a functional block to map the resulting parameter to a suitable filter time constant. A block diagram is shown in figure 4.5. The similarities of the model to the Burwen system in particular are obvious.

The time constants associated with the envelope of the test signal and the filter attack time can be readily varied.

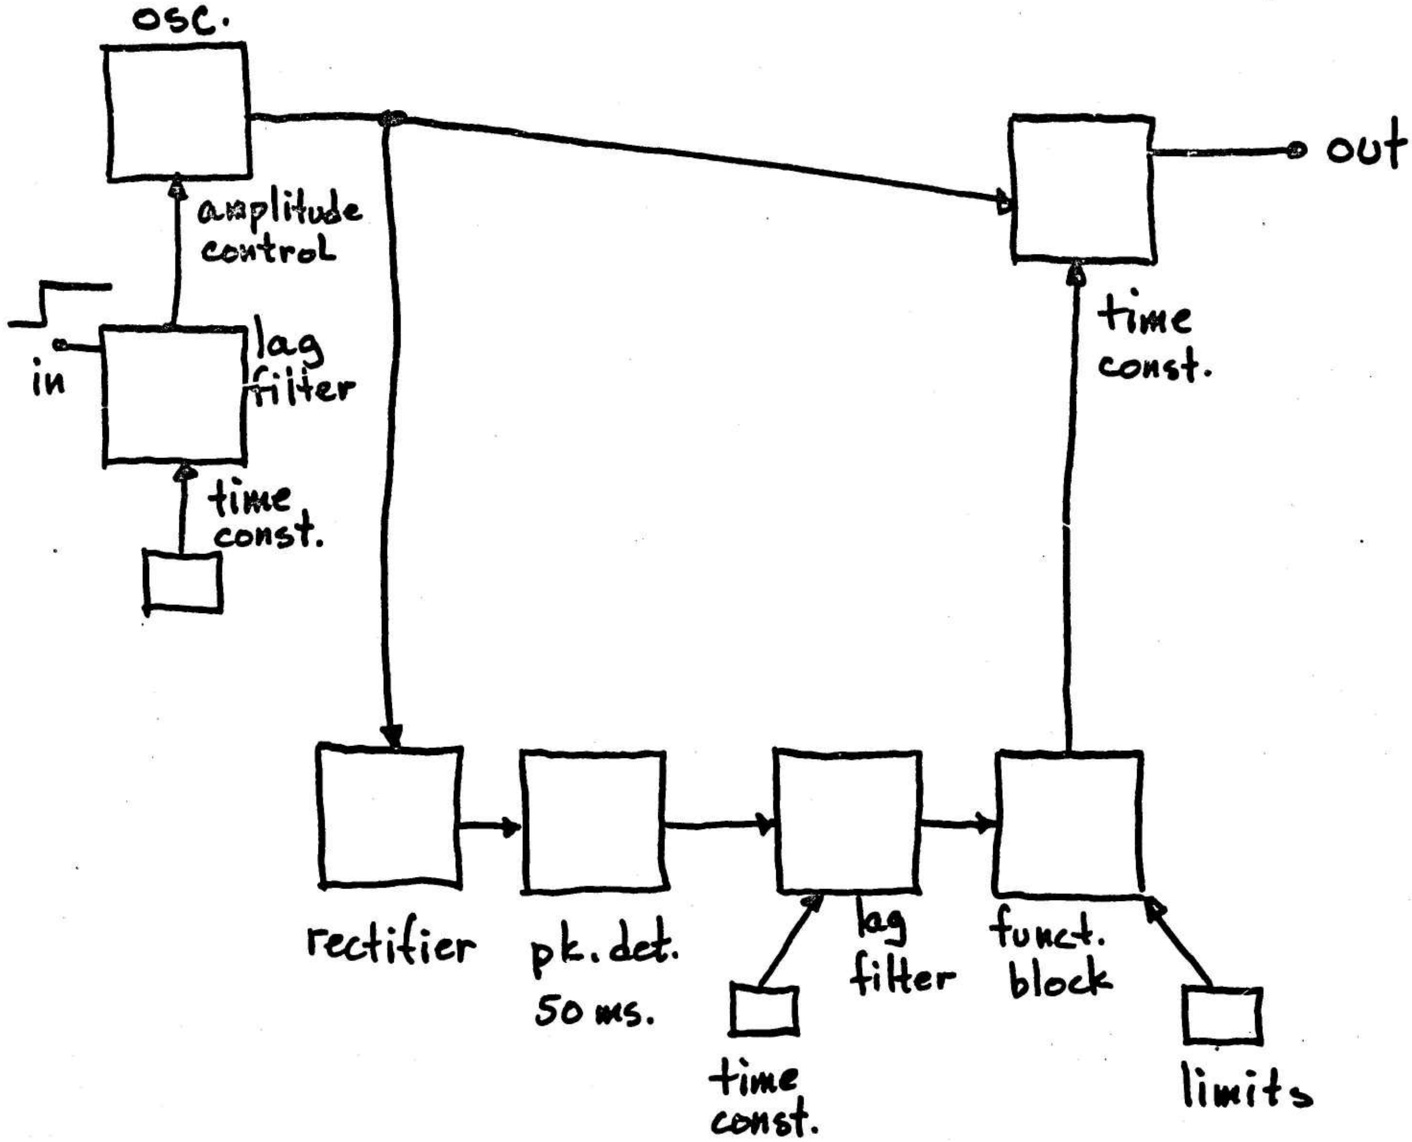

block diagram

Figure 4.5

The decay time constant, though not an important parameter in the tests, was set at 50 ms.

At this point, it is instructional to step through the operation of the model.

The first lag filter (a simple first order low pass, as are all the filters in the model) is presented with a step input. Ths shape of the output is an exponential, with the desired pre-programmed time constant. This controls the amplitude envelope of the sine oscillator. If desired, this can be a step envelope rather than an exponential.

The signal is then routed separately to the filter sections and the control section. In the control section, the signal is full-wave rectified and peak detected. The peak detector has a time constant of 50 ms, which controls the decay characteristics of the control signal. The signal then passes through another lag filter, and, since the peak detector has virtua lly instantaneous rise time, this filter controls the attack characteristics of the control signal. The attack time constant is pre-programmed to the desired value.

At this point, the control signal represents a time weighted measure of the energy in the input signal. The subsequent functional block is specified such that, when this signal is zero, it delivers a time constant to the program lag filter to keep its cutoff point at 500 Hz. When the signal goes to unity, as it will in sinusoidal steady state, it delivers a time constant corresponding to a cutoff of one-half octave above the test frequency. Thus changing the test frequency necessitates a change in one of the functional block parameters, but this inconvenience is offset by the simplicity of the approach. Between its two limits, the functional block provides a linear mapping.

The expected output will thus be a sinusoid, with some phase modulation, and with an envelope that is some function of the original envelope and the control section's attack characteristics. The general form of this output should correspond well to the output obtained from any of the other dynamic filters described; they all control the shape of their response characteristics by rectifying, peak detecting, and filtering the program material to obtain a measure of the energy in some portion of the spectrum. The exact details of any particular scheme may not coincide with the model, but we expect that, for the simple class of test signals with which we are concerned, the gross behavior will be very similar to that of our model and any results obtained could be generalized, within limits, to any dynamic filter.

Expected Results

For even the simplest test signals, such as a sine wave multiplied by a step, any analytic expression for the filter's transfer function is difficult or impossible to obtain due to the non-linearity of the system. However, we can get an estimate of the shape of the output spectrum by assuming an exponential envelope in the time domain, controlled by the filter's attack time, and some phase modulation, causing some spectral smearing around the fundamental.

For example, an input signal of the form:

$$U_{-1}(t)\ \sin(2\pi \times 2500)t$$

should yield an output of the form:

$$U_{-1}(t)\ (1-e^{-t/.01})\ \sin(2\pi \times 2500)t$$

when the attack time constant is 10 ms. The phase modulation caused by sweeping the angle of the program filter from -90° to 0° at 2.5 kHz may be estimated by assuming the change in phase to be constant over 10 ms.:

$$\frac{90°}{10\text{ms.}} = \frac{360°}{40\text{ms.}}$$

corresponding to sidebands in the vicinity of:

$$\frac{1}{40\text{ms.}} = 25 \text{ Hz}$$

from the fundamental.

The spectra of these functions are sketched in figure 4.6.

One point immediately obvious is that the sketch of the output shows substantially reduced energy in the upper frequencies. This is hardly surprising, as it correlates well with an intuitive view of one of the effects of the dynamic filter, which is its dulling of transients.

It should be emphasized, however, that this sketch is only an estimate. While the general shape should be correct, the only way to look in detail for any artifacts of this non-linear operation is to perform a detailed spectral analysis on an actual output signal.

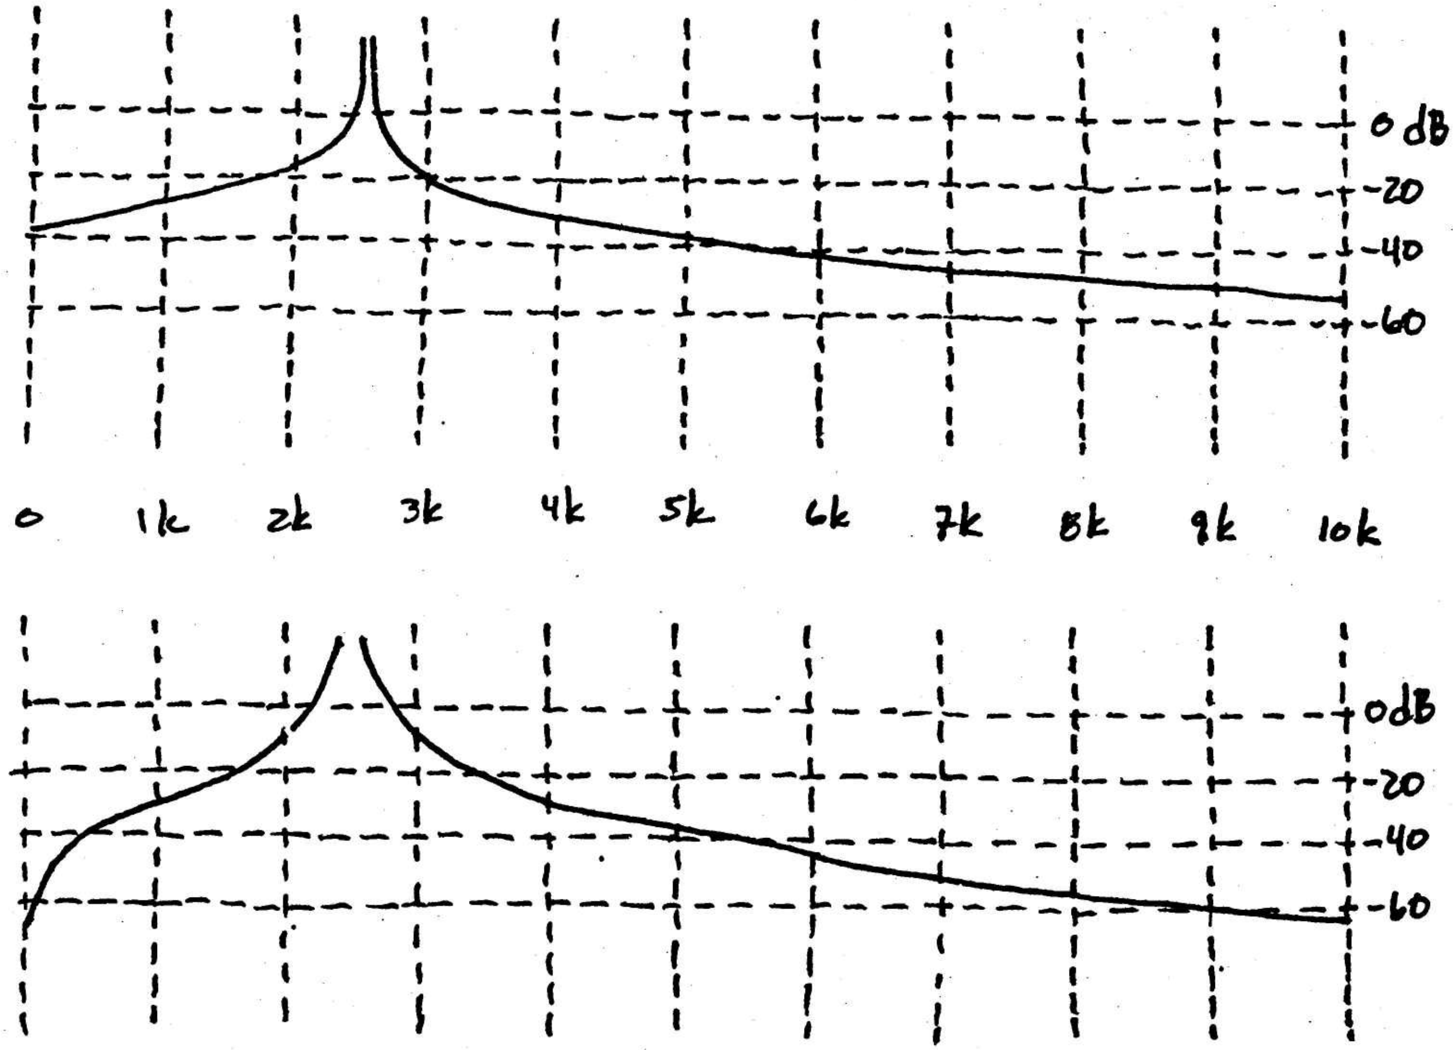

figure 4.6

V THE EXPERIMENT

Computer Implementation of the Model

The model described in the previous section was implemented on a PDP-9 computer at the computation facility of the Speech Communication Group at MIT. The language used was MITSYN, developed by Dr. William Henke of that group. MITSYN is a high level interactive language used for time and frequency signal processing, and implemented with graphic and sonics support. [11]

The model was implemented graphically, in block diagram form, in virtually the same configuration as figure 4.5. For each experimental run, the system would process the configuration and deliver the input and output signals, in digital form, to a signal store buffer. From there, a discrete Fourier transform was performed on the data, and the spectra were displayed and subsequently printed out as hard copy, along with their respective time domain blocks.

Limitations on Spectrum Analysis

The area of digital signal processing is an enormously rich one, though it is relatively young, and one of its areas of concern is the errors caused by examining finite blocks of data. Frourier theory, in both the discrete and continous cases, converges only in the limit, so there are errors necessarily encountered in almost any real world application that involves time domain data of finite length.

The system implementation of the Fourier analyzer provided a maximum time window of 512 samples. It was found that, for most of the data taken, a 20 kHz sample rate was necessary to avoid aliasing. The time windows thus obtained were 25.6 ms duration.

Discrete Fourier analysis can theoretically give exact results for a finite length window if that window contains an integral number of cycles of a periodic signal; in any other case the computed spectrum will be smeared to some extent. This is illustrated in figures 5.1 and 5.2.

Figure 5.1 is a window of 25 cycles of a 500 Hz signal. The sample rate was 5120 Hz and the window was 256 samples. (Actually, due to some asynchronism in the system, the window is slightly less than 25 cycles, but this will not alter the point.) The log magnitude spectrum is similar to the impulse at 500 Hz that we would expect. The abnormality around 2100 Hz is below the system's nominal dynamic range, and can probably be ignored as an artifact of the system.

Figure 5.2 shows the results of presenting a 502 Hz signal for analysis under the same conditions. It can be seen that the calculated spectrum differs markedly from the

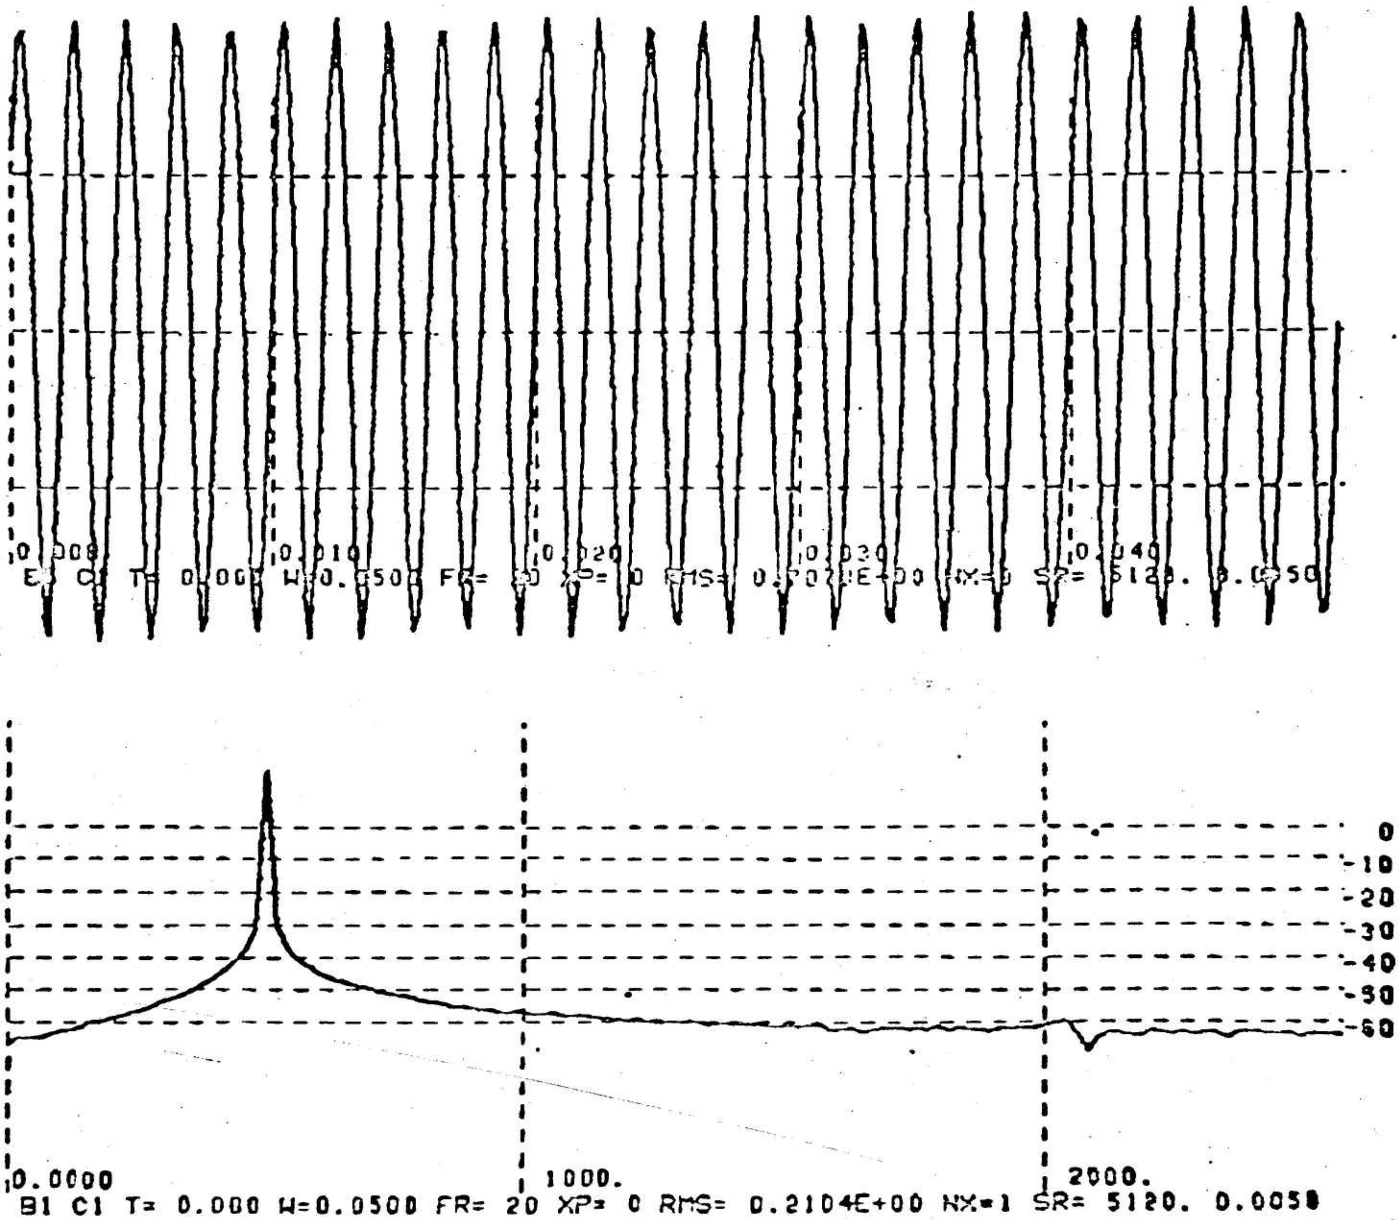

500 Hz test signal

Figure 5.1

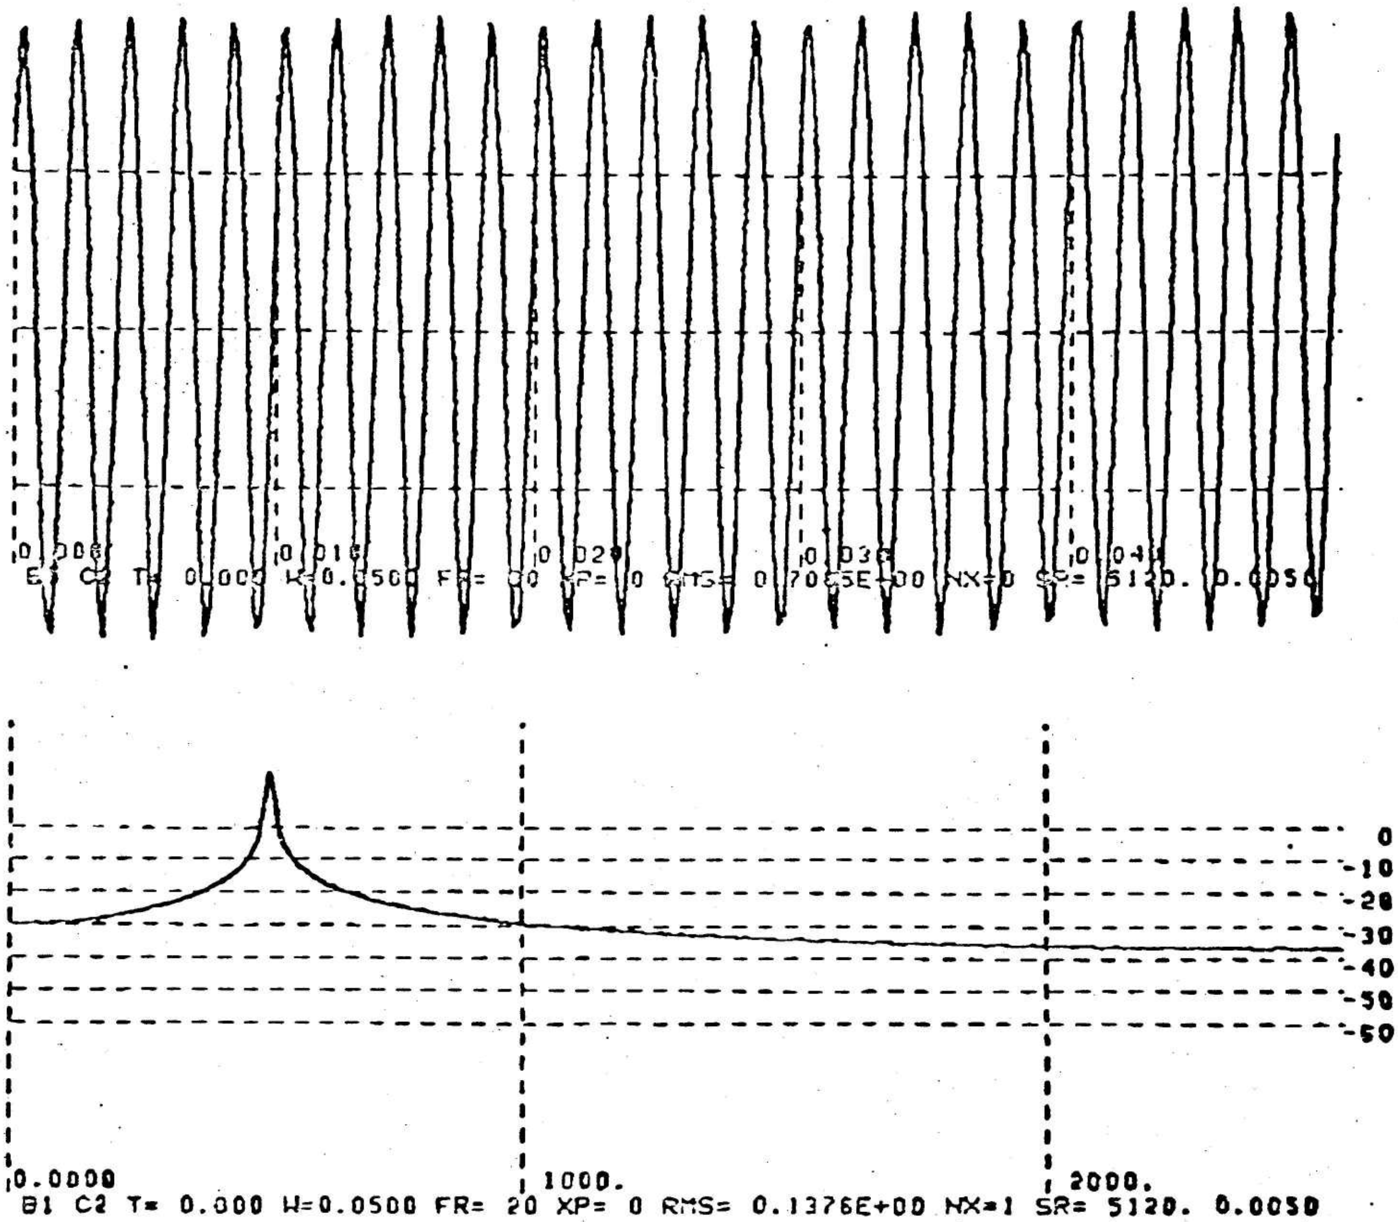

502 Hz test signal

Figure 5.2

first case. The fundamental frequency is certainly correct, but the impulse-like shape has been smeared, and at frequencies far from the fundamental the difference between the two spectra is dramatic, being over 20 dB.

For many forms of frequency analysis, this limitation is minimized by multiplying the time window by a shaping function, such as a Hanning or a Hamming window. This attempts to smooth discontinuities near the boundaries of the window.

It was felt, and initial trials indicated, that this form of weighting intruded on the transient nature of the signals we wanted to analyze.

Fortunately, signals with such sharp transitious are not realistic for testing an audio component. Studies have indicated that virtually no sound likely to be recorded has an envelope of less than about 10 ms rise time; such sharp sounds are produced by instruments like the piano, and a plucked guitar string. [12]

Figure 5.3 and 5.4 show the same two test signals, at 500 and 502 Hz, with a 20 ms exponential attack characteristic; figure 5.5 shows the two spectra overlaid. The lower graph has an expanded scale, from 250 Hz to 750 Hz. It can be seen that the two spectra are much closer to a uniform

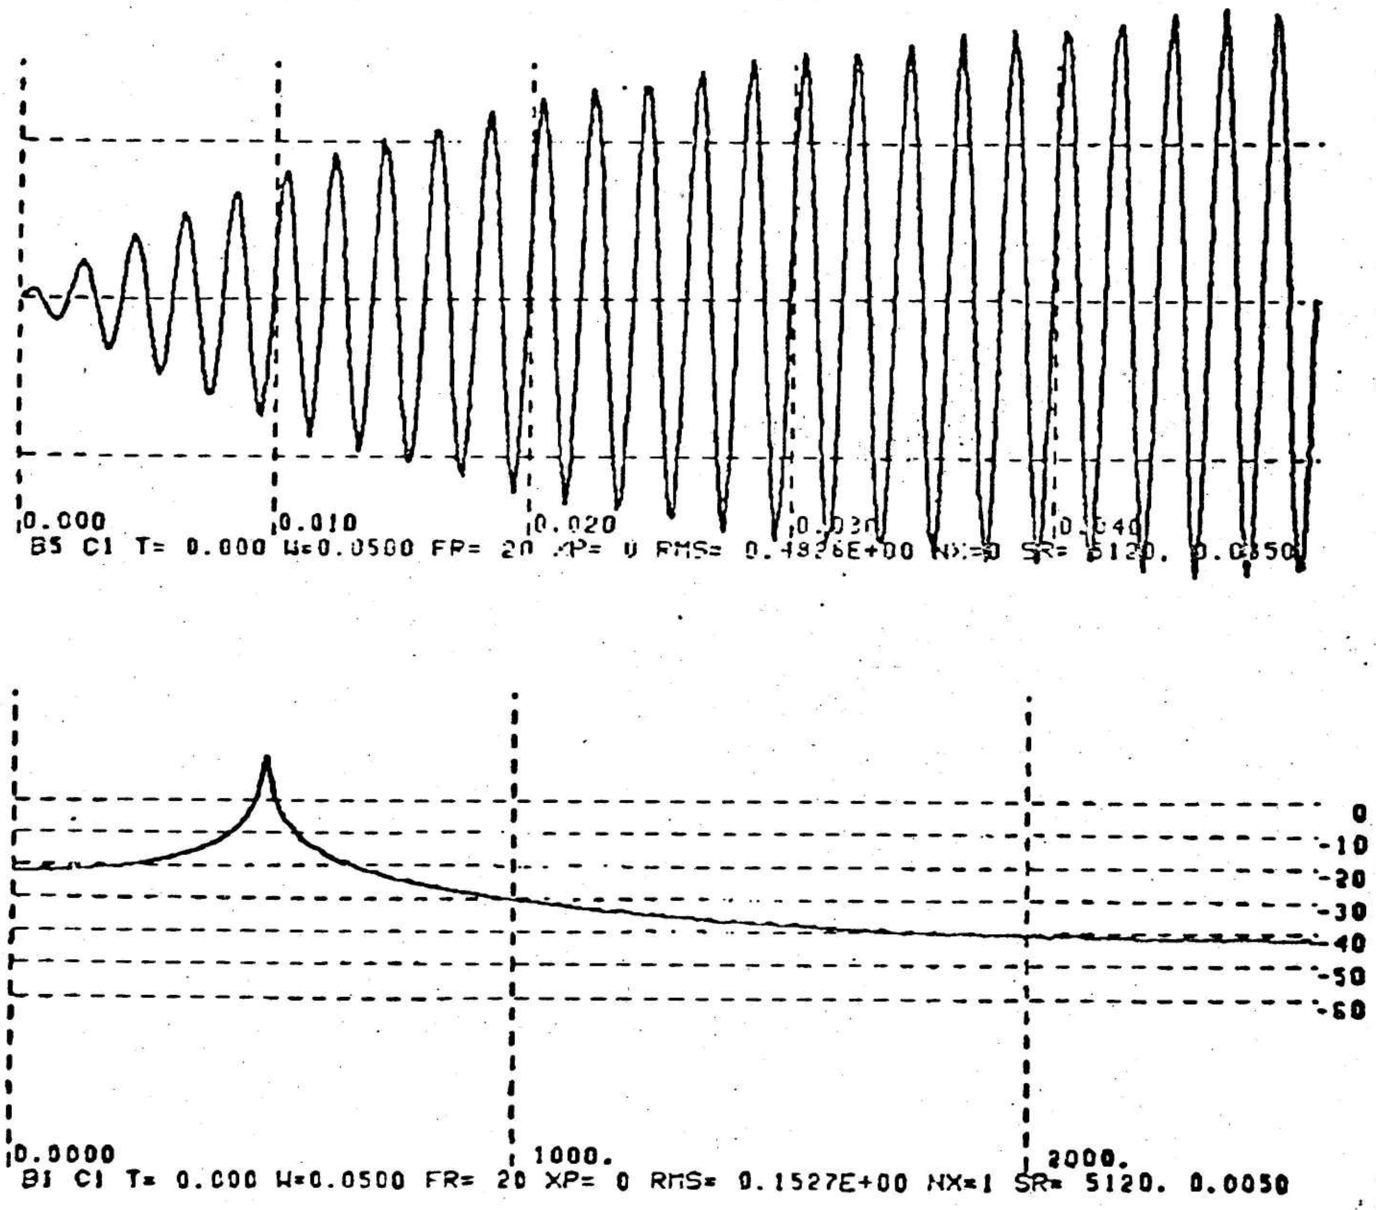

500 Hz test signal
20 ms. time constant exponential envelope

Figure 5.3

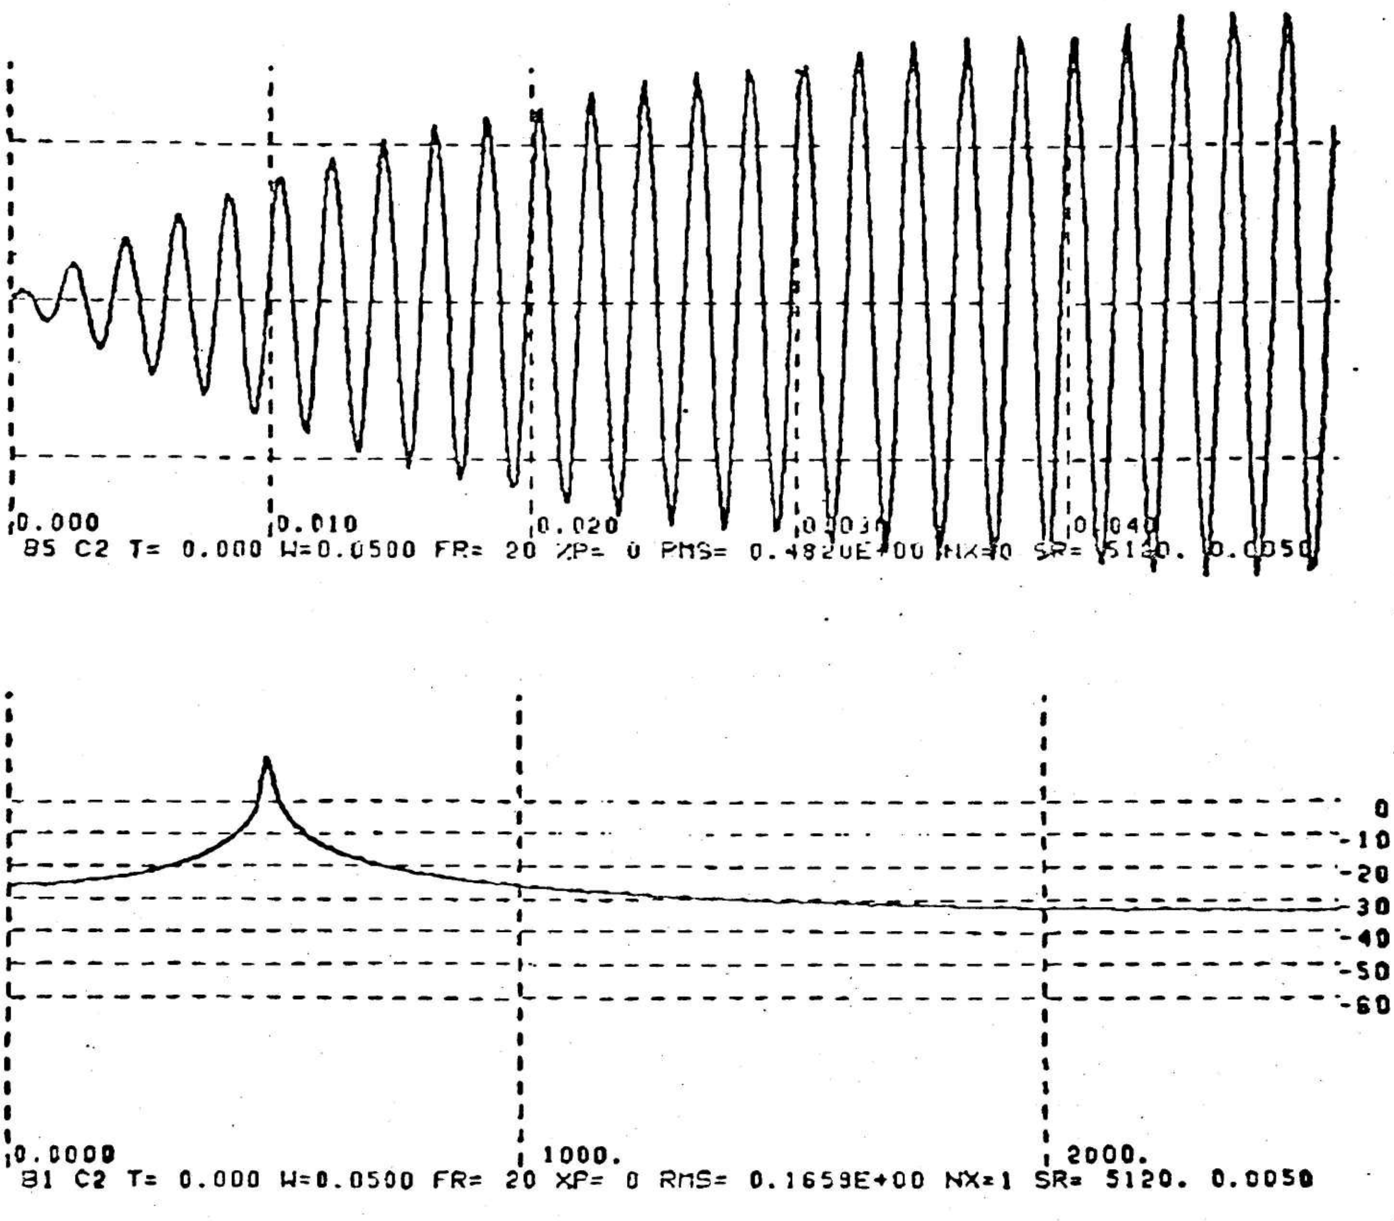

502 Hz test signal
20 ms. time constant exponential envelope

Figure 5.4

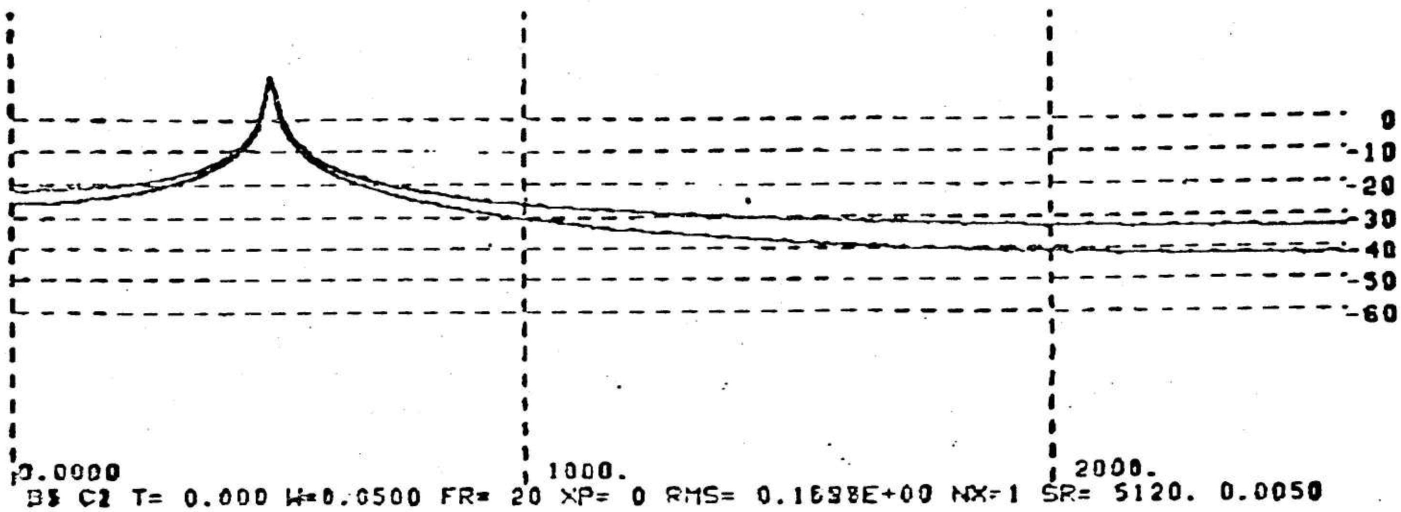

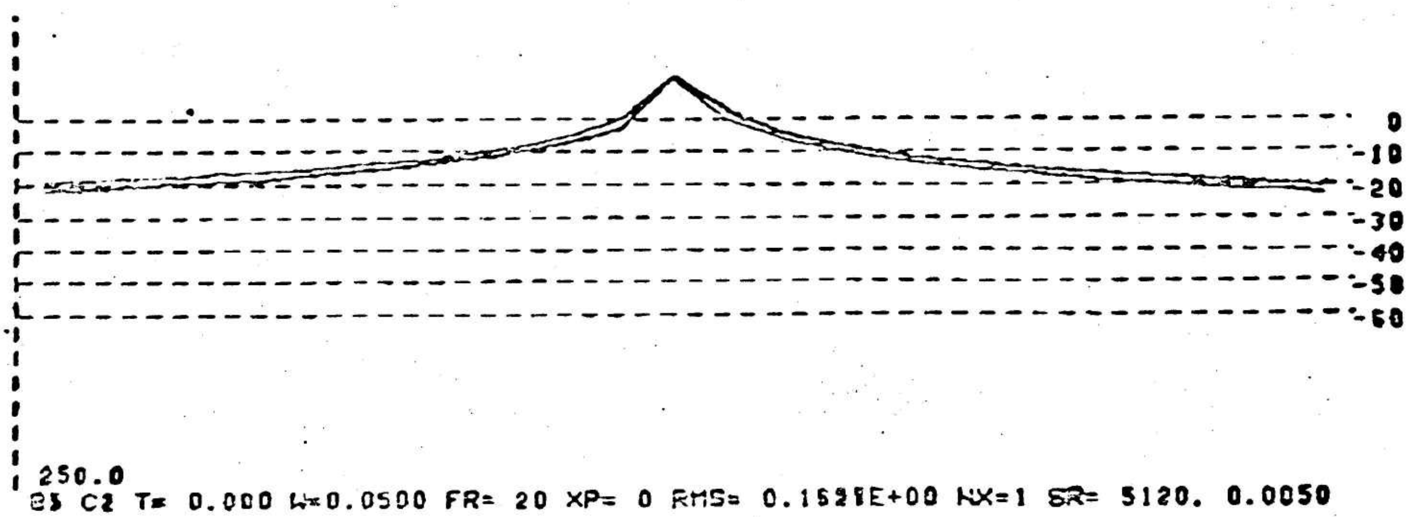

comparison of the spectra of figures 5.3 and 5.4

Figure 5.5

shape than in the previous case. (A 20 ms attack time was chosen for the above illustration because one of the actual test signals was a 1 kHz sine wave with a 10 ms attack time and, except for a scale factor, these two cases are equivalent).

Because of the problems illustrated above, virtually all of the actual tests were done using input signals weighted with a 10 ms time constant exponential. Figure 5.5 gives us a measure of the error we can expect in studying the differences between the dynamic filter's input and output, especially in the region around the fundamental frequency. Specifically, differences on the order of 3 or 4 dB must be regarded as being below the resolution of this experiment.

The Experimental Results

Figures 5.6 - 5.15 illustrate the results of several experimental runs. They can be divided into two sets.

The input signal for the first set is shown in figure 5.6. It is a 1 kHz sinusoid with a 10 ms exponential envelope. Figures 5.7 and 5.8 show the output of the filter with an attack time of .5 ms and 2.5 ms respectively.

The input signal for the second set is shown in figure 5.9. It is a 2.5 kHz sinusoid with a 10 ms exponential envelope. Figures 5.10, 5.12, and 5.14 show the output of the filter with an attack time of .5 ms, 2.5 ms, and 10 ms respectively. Figures 5.11, 5.13, and 5.14 show input spectra overlaid with output spectra, on normal and expanded scales, for the above three cases.

In all cases the quiescent cutoff point of the filter was set at 500 Hz, and the control portion of the filter moved the cutoff point to one-half octave above the test signal frequency.

The attack times for the filter were chosen as representative of current implementations.

Faster rise times for the test signals were found to be unsatisfactory, for reasons previously given, as well as unrealistic of any actual operating conditions. It was felt that slower rise times would yield limited data because of the short time interval the analyzer was constrained to look at.

Simple frequency test signals above 2.5 kHz were not considered realistic, either, since virtually all musical energy above that point is composed of harmonics.

All tests were done at one-tenth rate; time constants were multiplied by ten and frequencies were divided by ten. The scale factor was subsequently corrected in the graphic displays.

All tests were done at an effective 20 kHz sampling rate and all time windows were of 25.6 ms. effective duration, starting at the same point as the test signal.

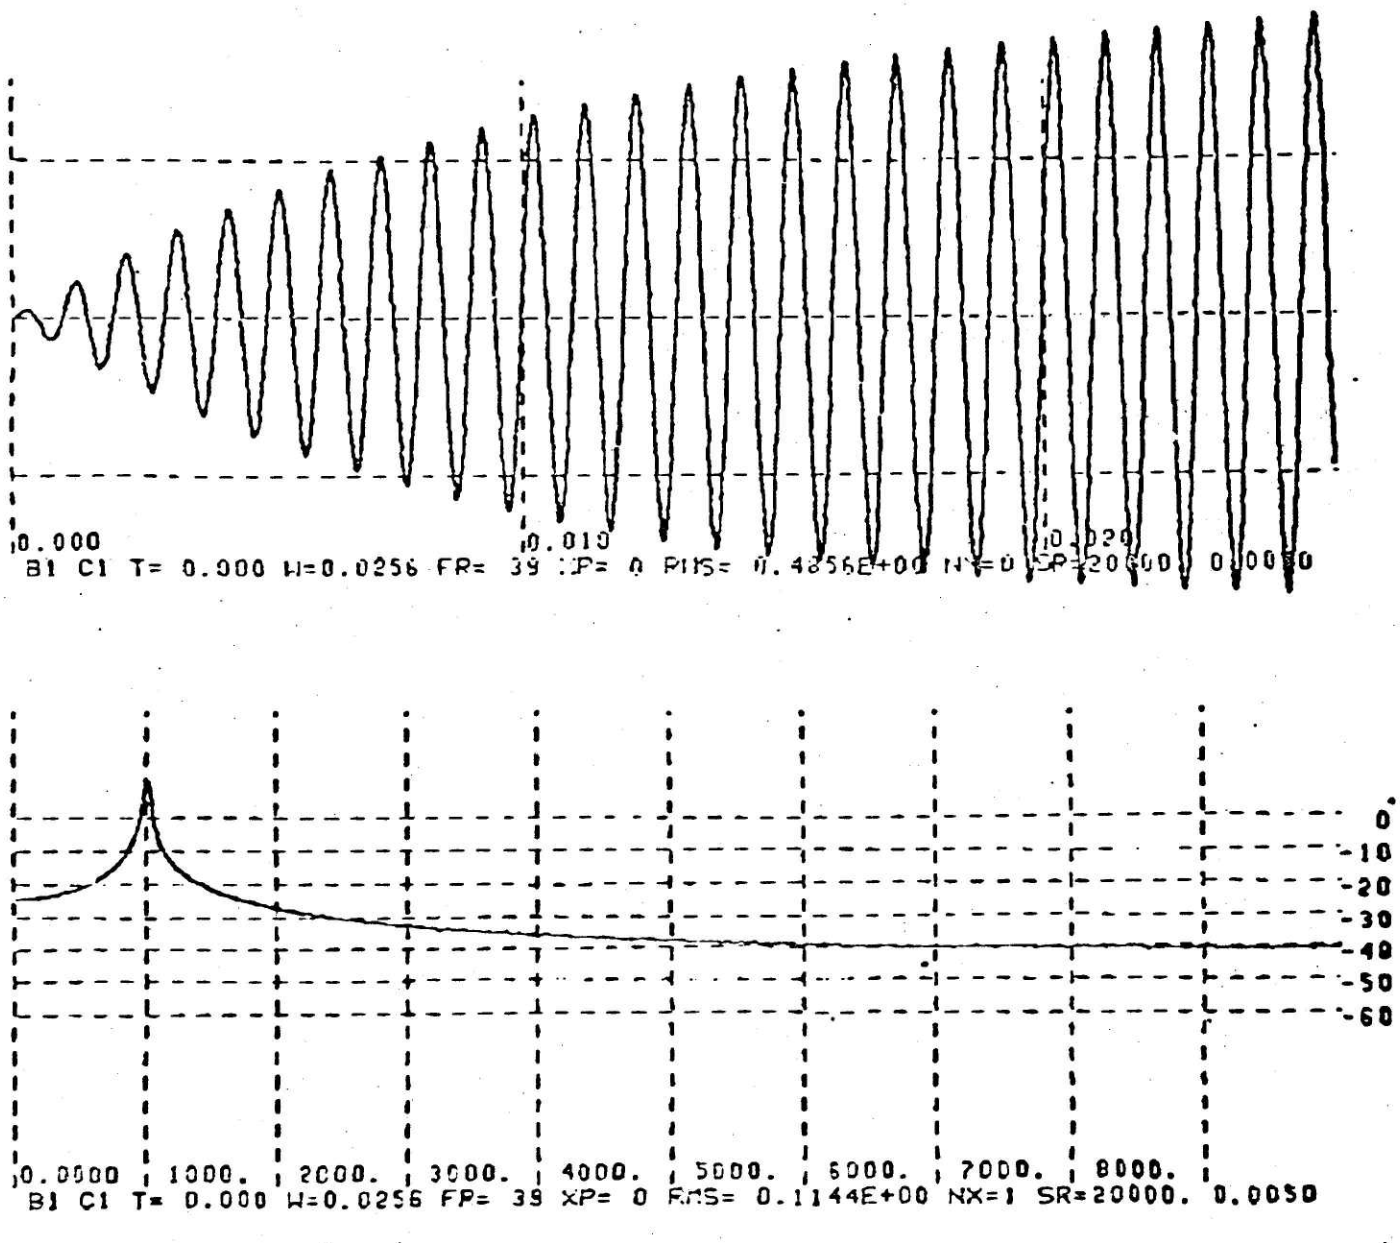

input test signal: 1 KHz with a
10 ms. time constant exponential envelope

Figure 5.6

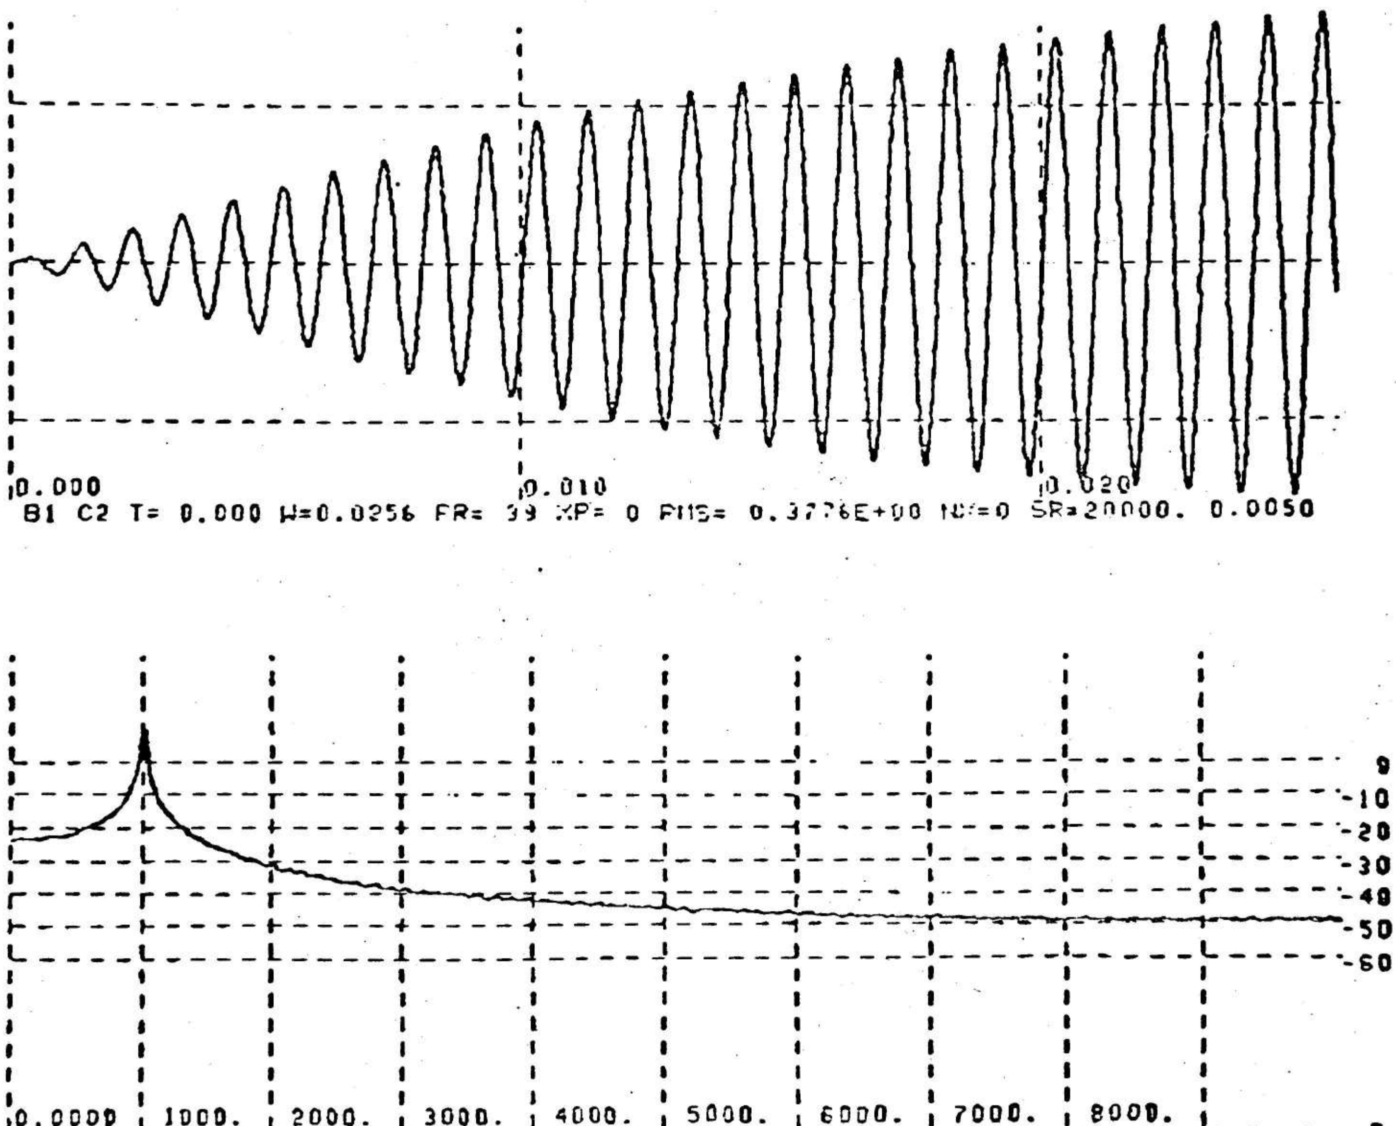

filter output: 0.5ms.
attack time

Figure 5.7

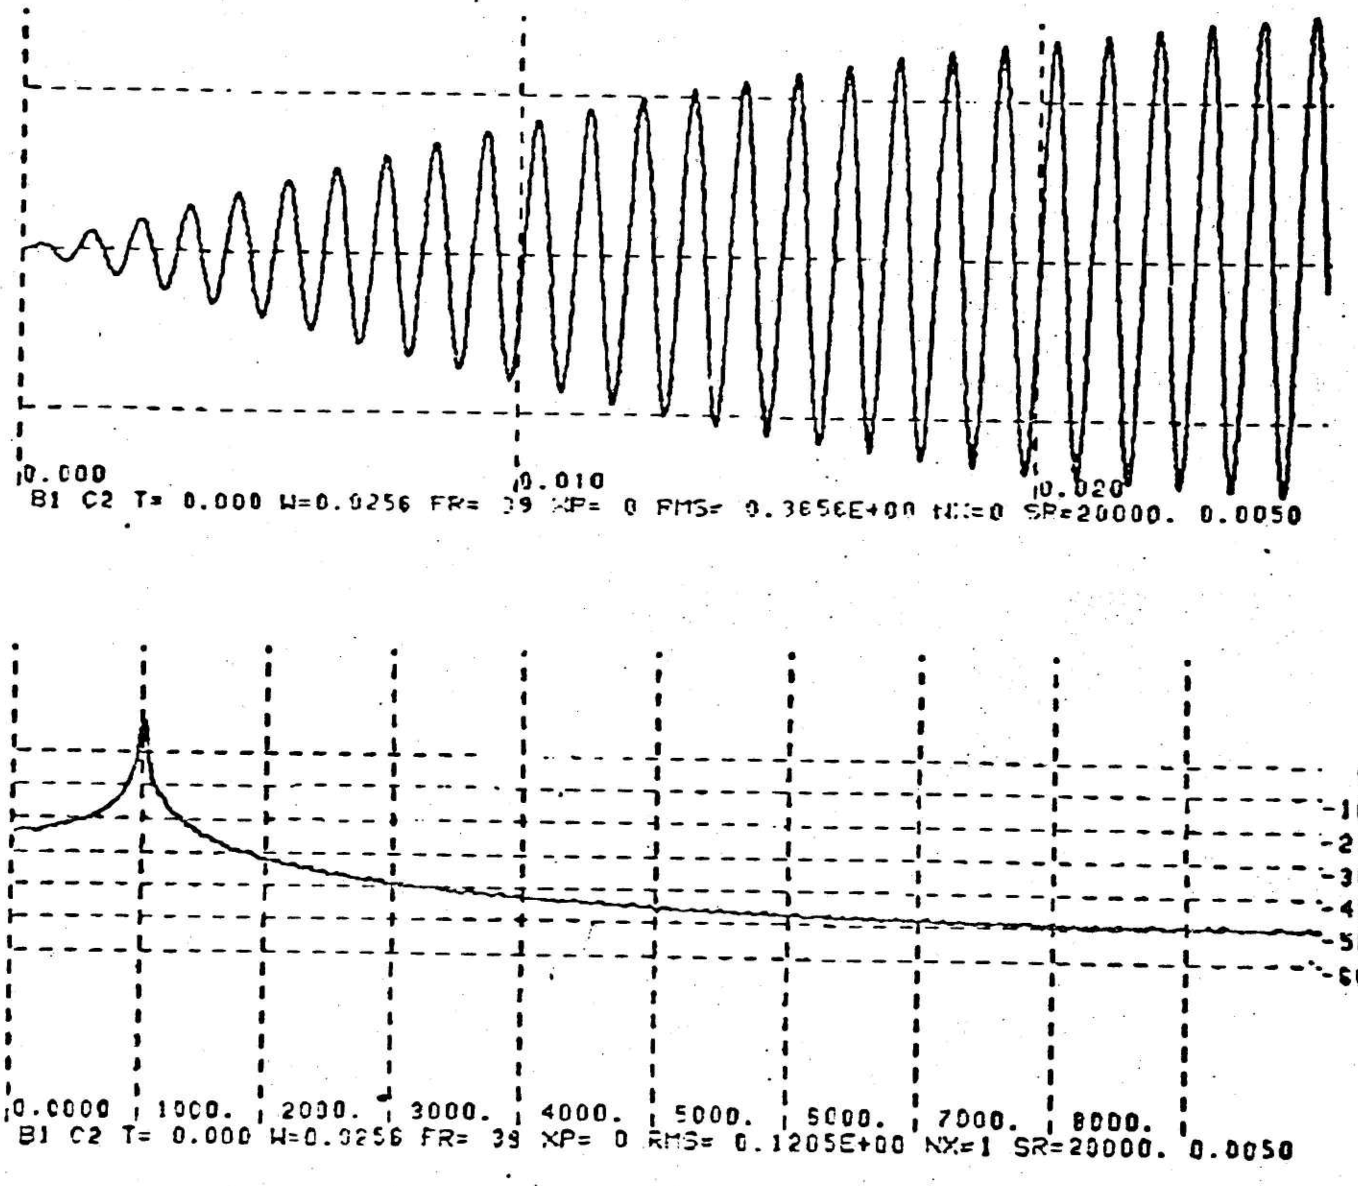

filter output: 2.5 ms.
attack time

Figure 5.8

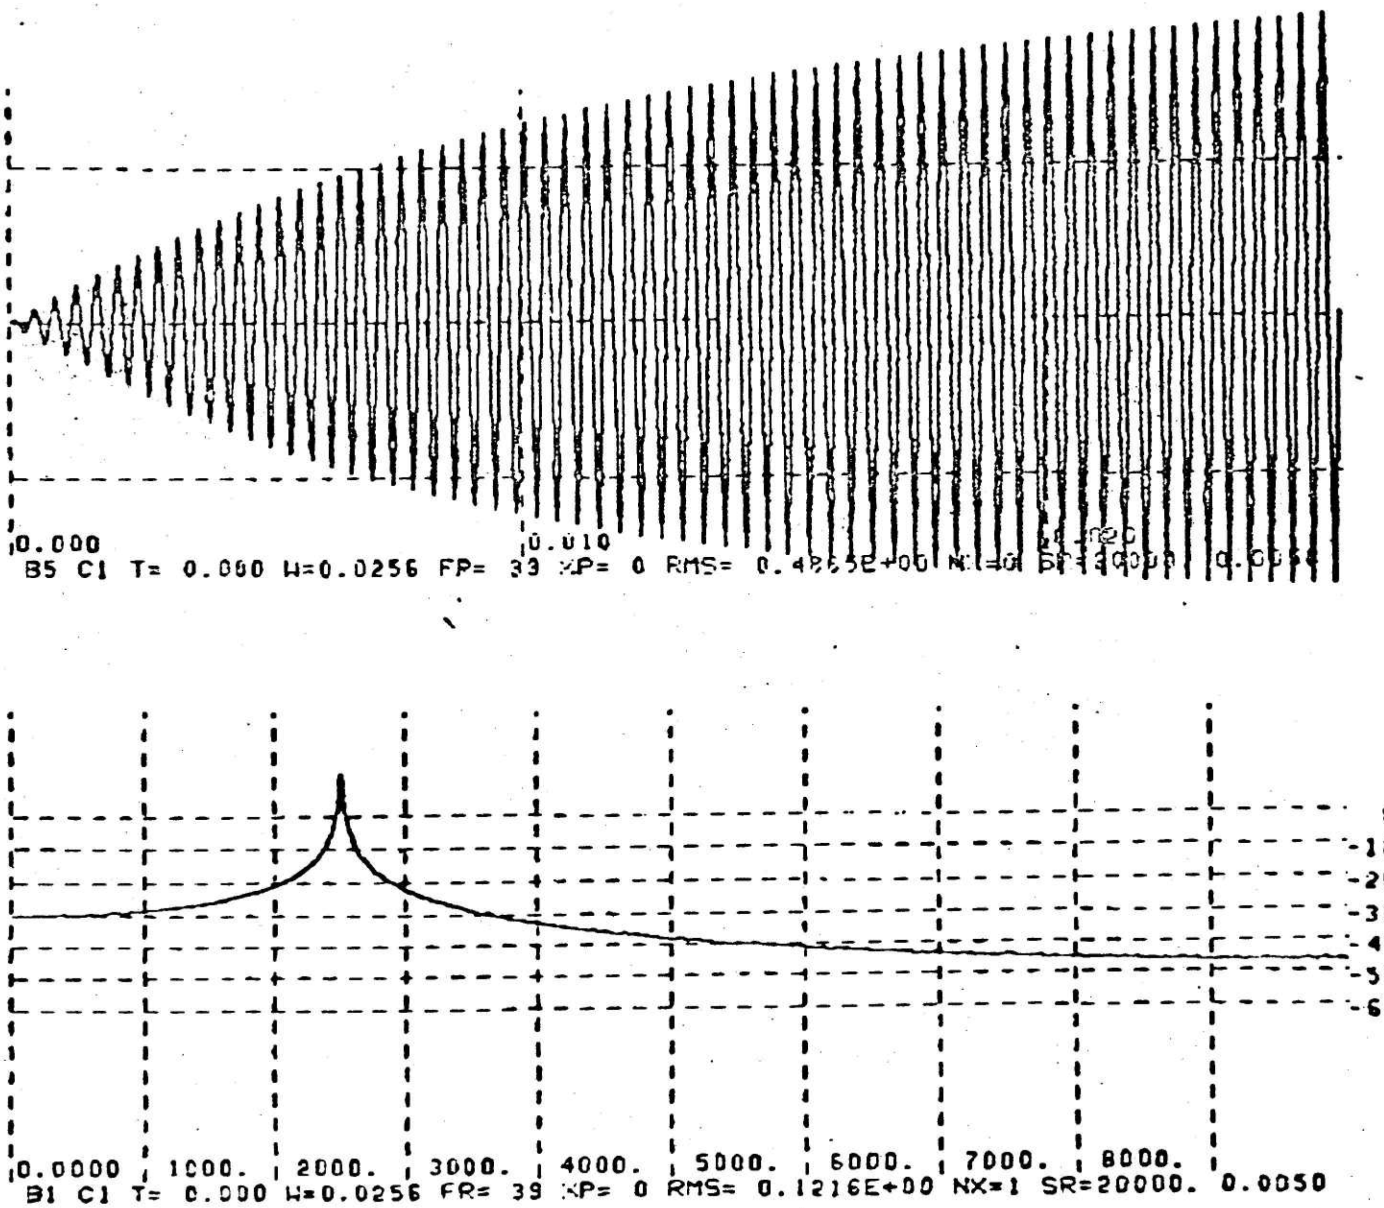

input test signal: 2.5 kHz
10 ms. time constant exponential envelope

Figure 5.9

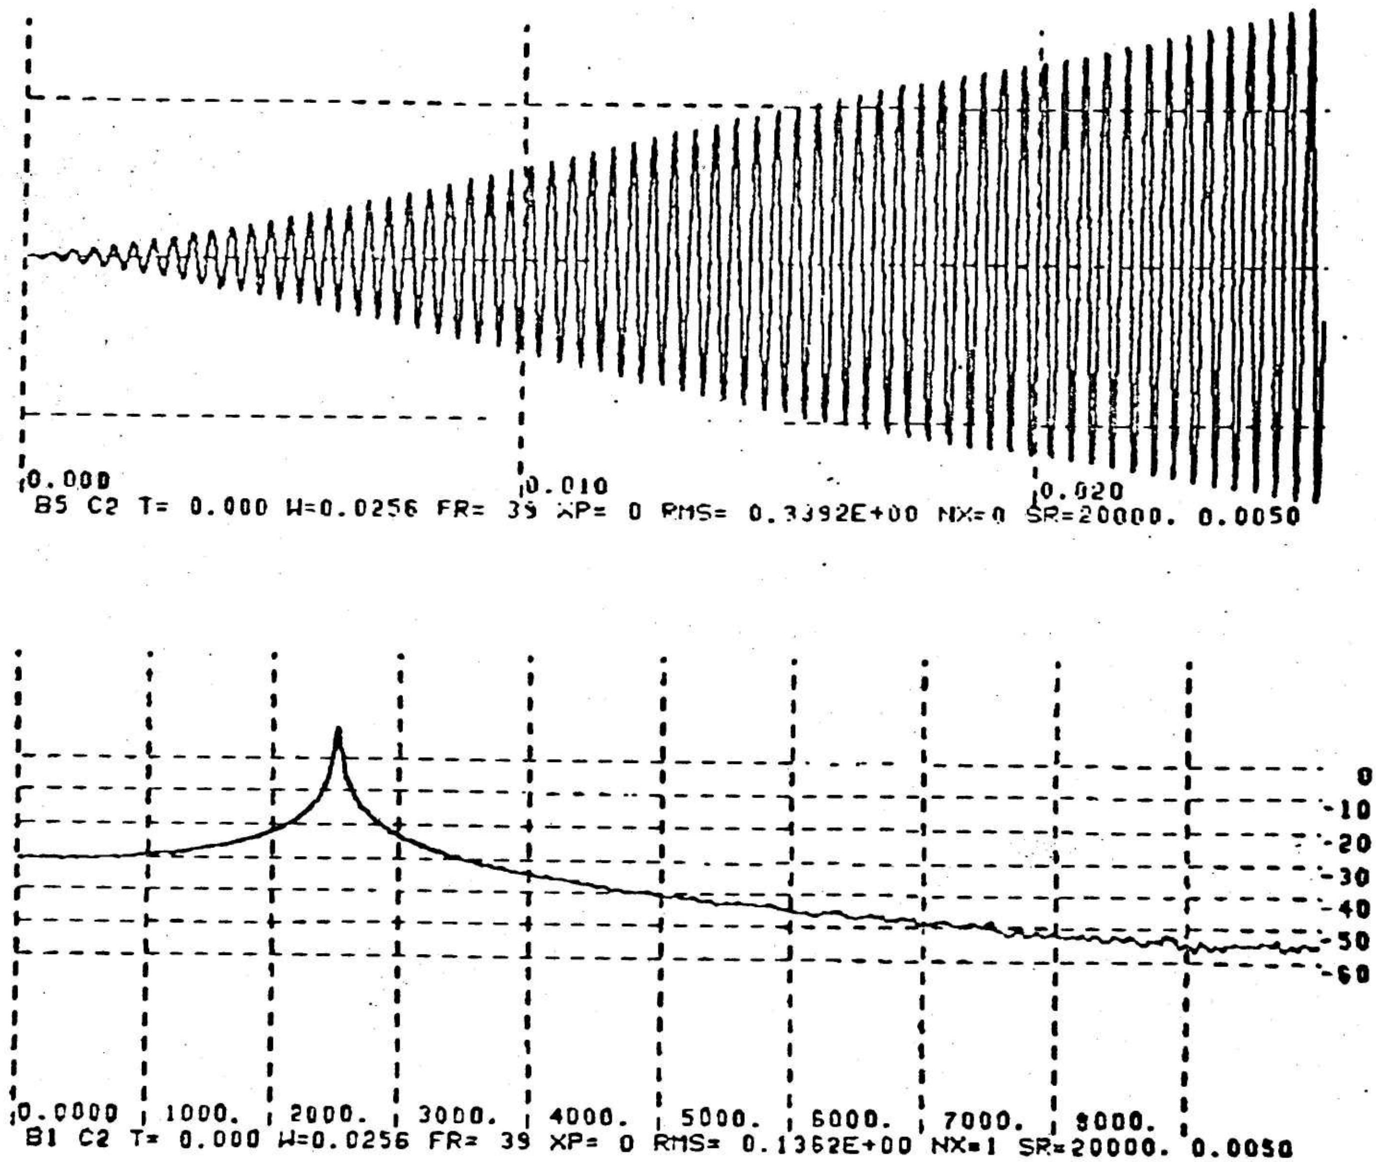

filter output: 0.5 ms. attack time

Figure 5.10

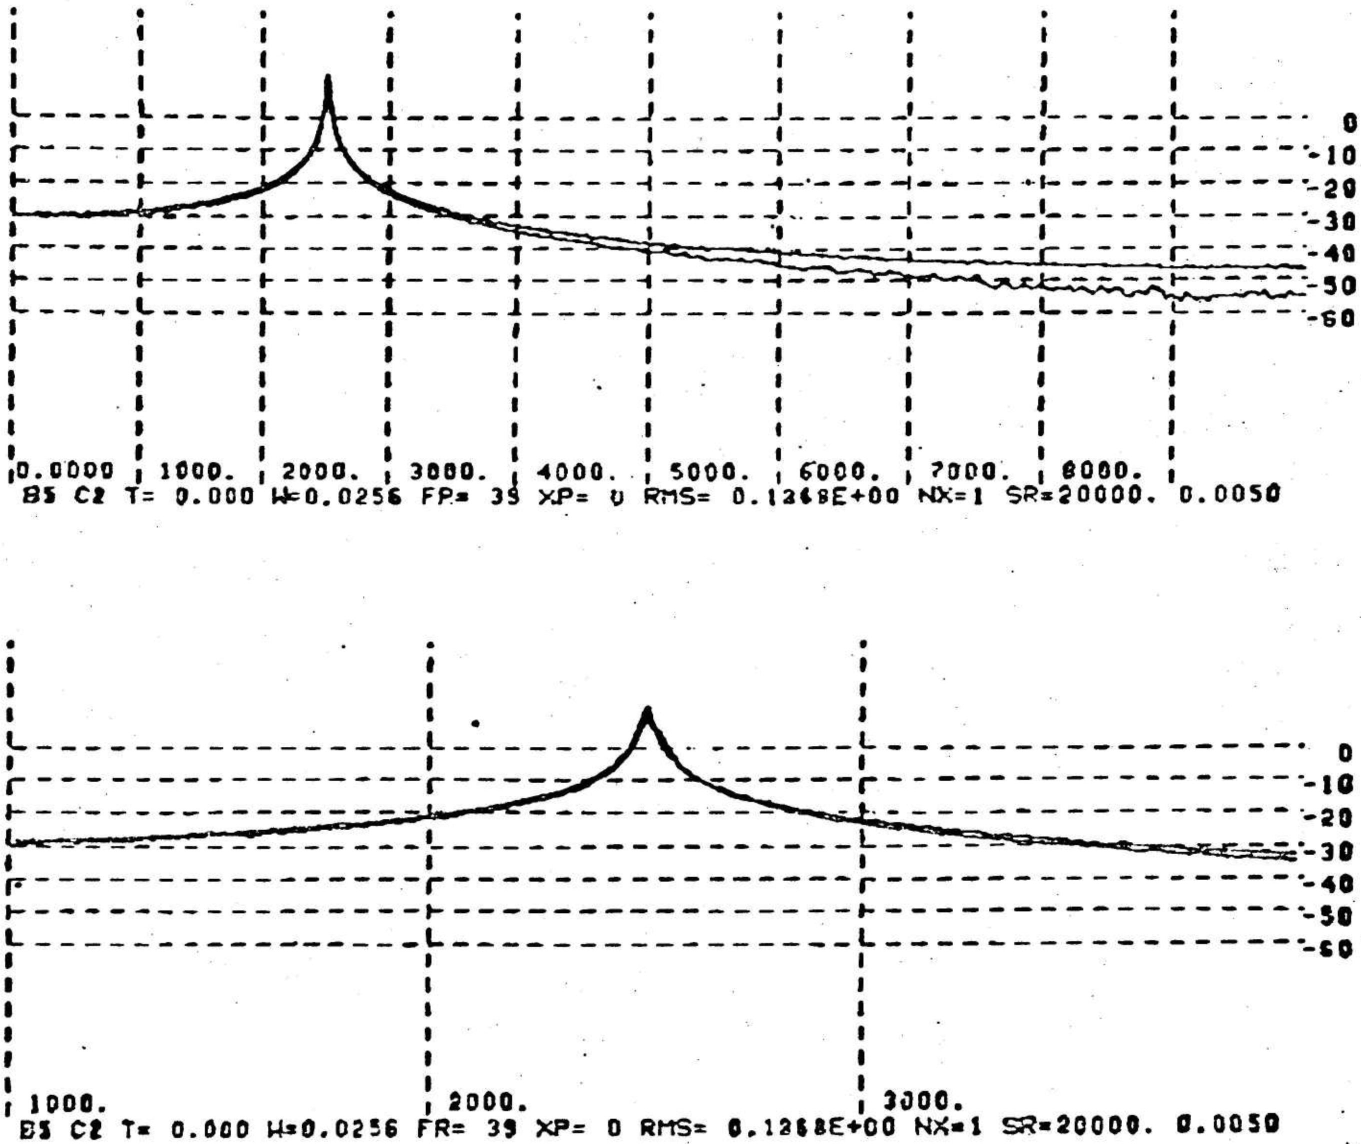

filter input and output spectra: 0.5 ms. attack time

Figure 5.11

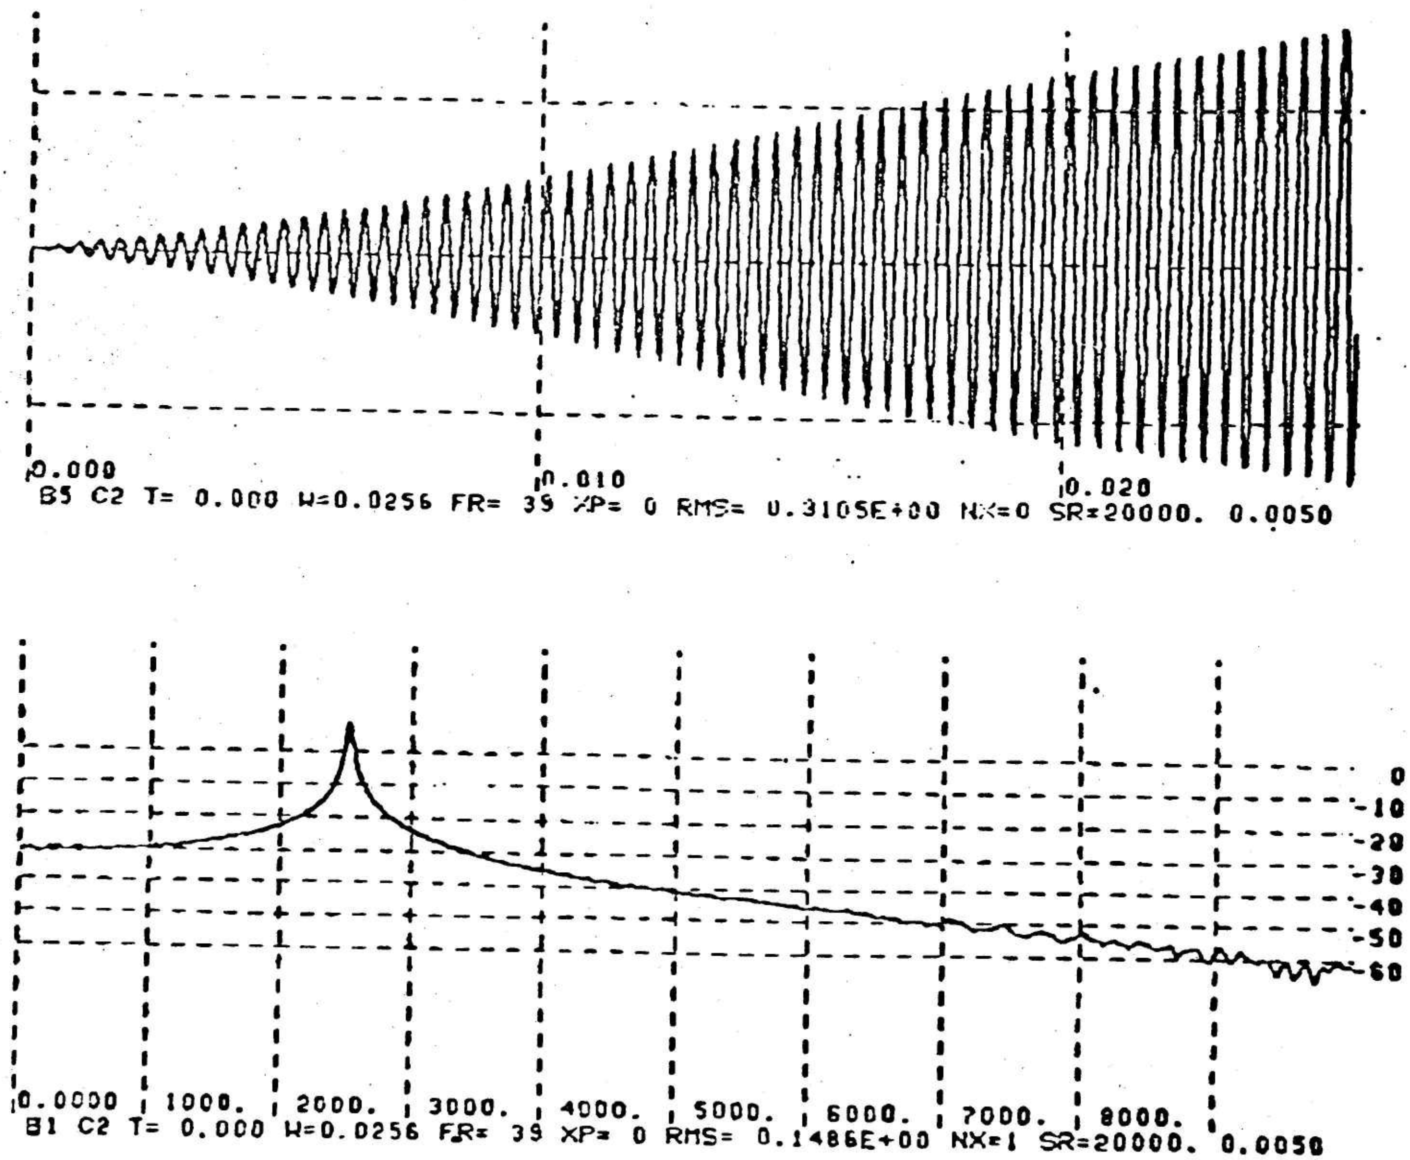

filter output: 2.5 ms. attack time

Figure 5.12

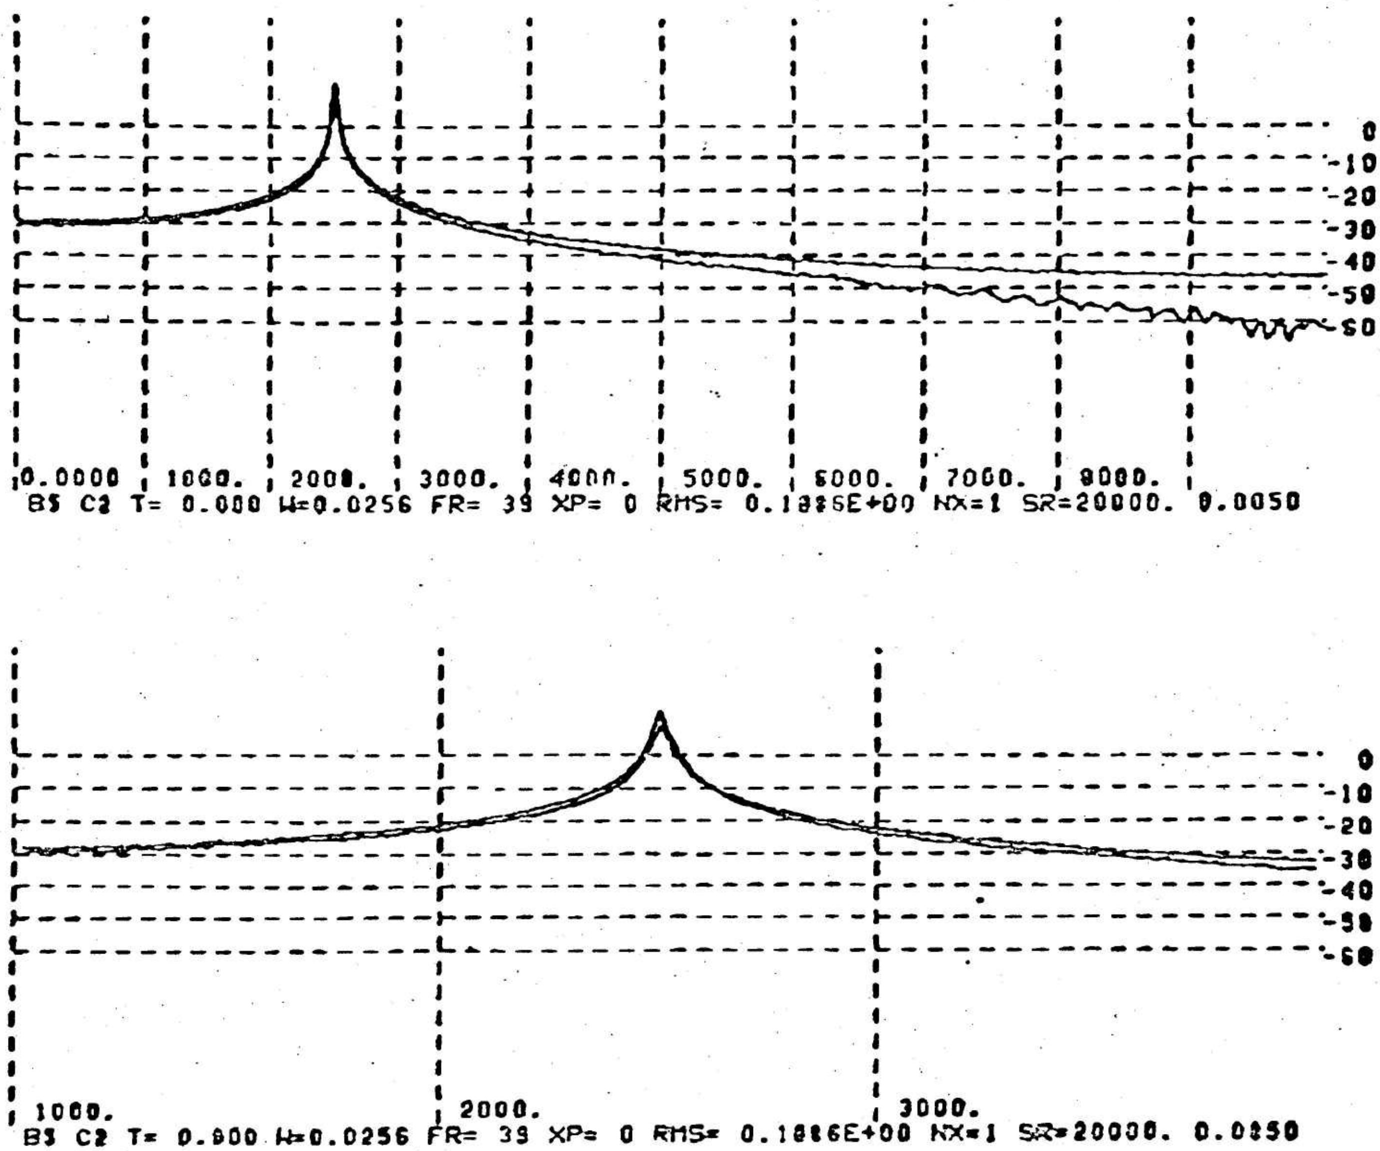

filter input and output spectra: 2.5 ms. attack time

Figure 5.13

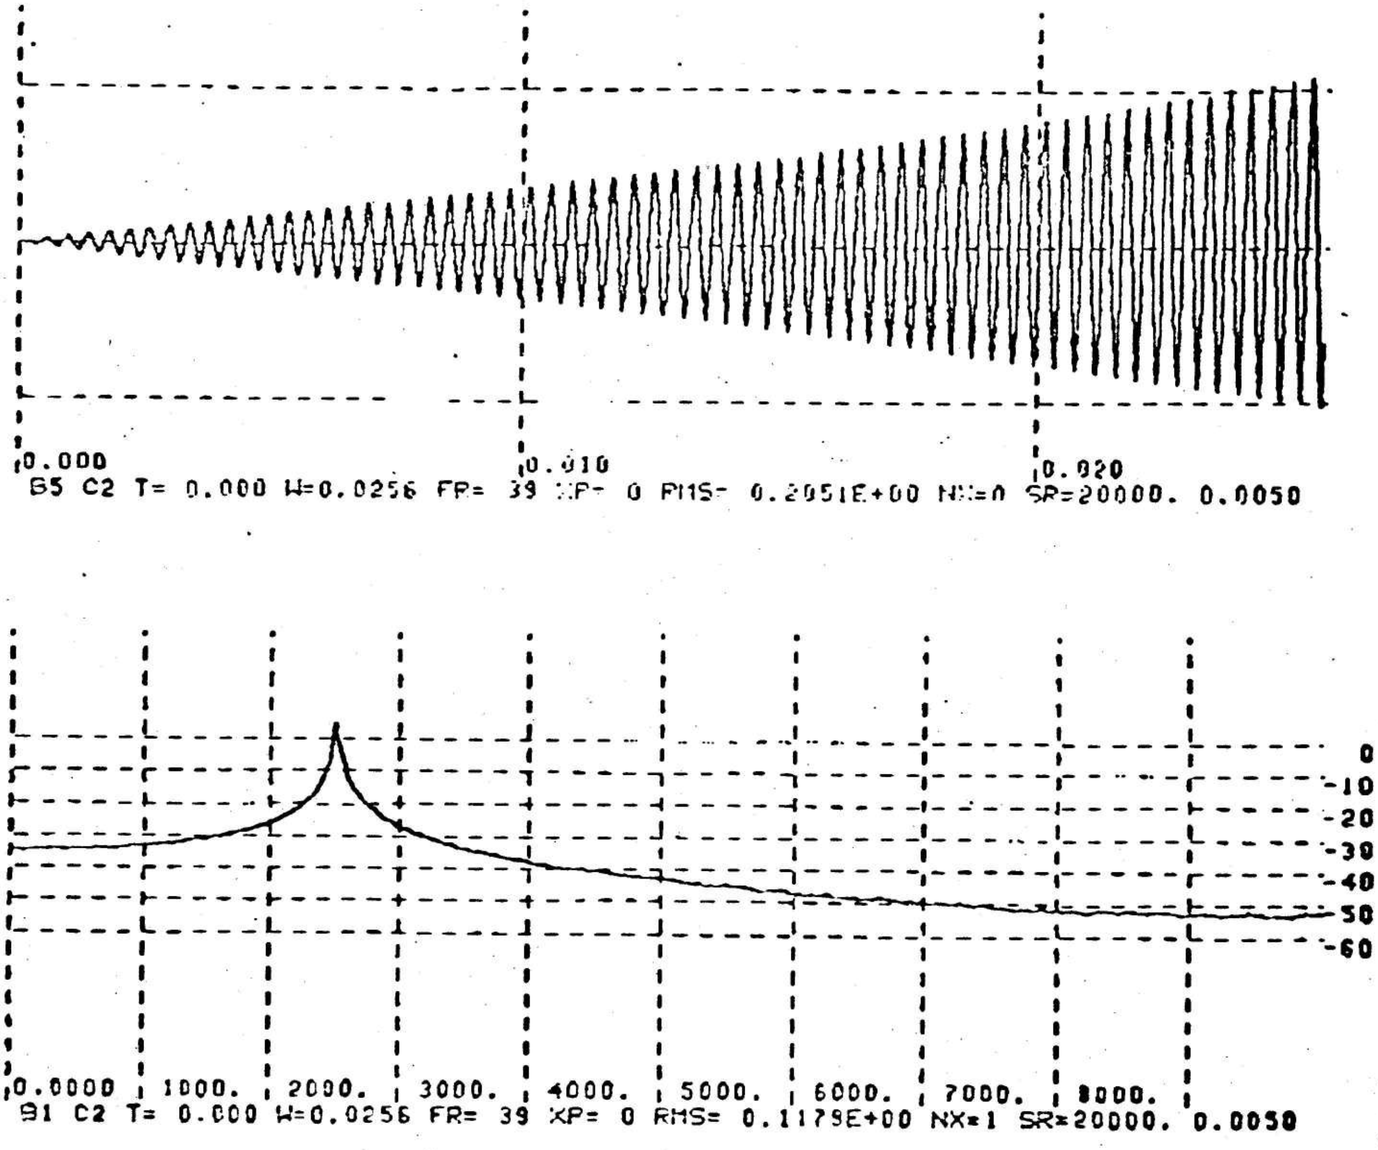

filter output: 10 ms. attack time

Figure 5.14

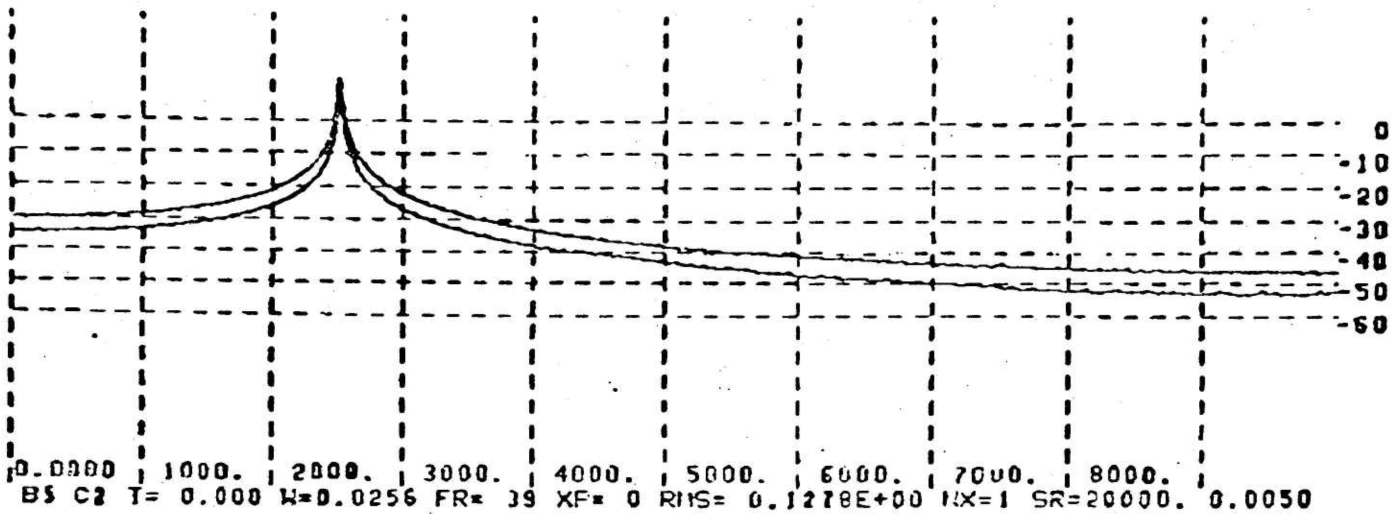

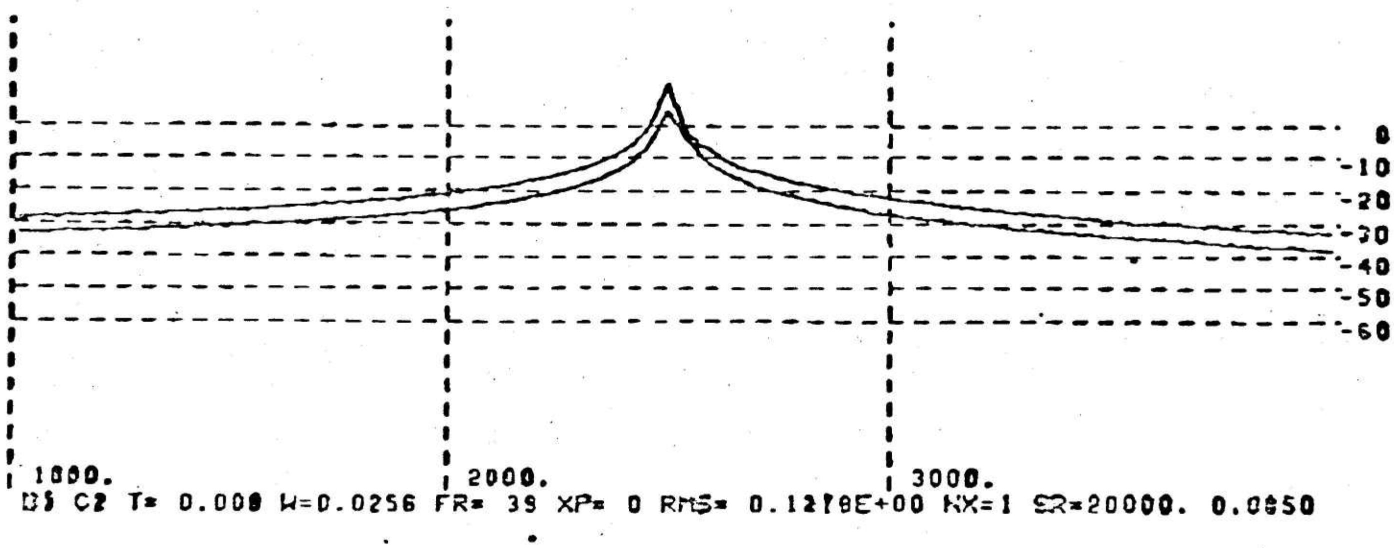

filter input and output spectra: 10 ms. attack time

Figure 5.15

VI DISCUSSION OF EXPERIMENTAL RESULTS AND CONCLUSIONS

For the 1 kHz test signal, the graphic results speak for themselves. There are no major surprises. As expected, the high end of the spectrum shows some rolloff due to the dynamic filter's dulling of the input transient. Both output spectra show some very slight roughness over their entire range, but it is of such small magnitude that it must be considered inconclusive, as it is probably below the resolution of the system. Close examination of the spectra in the vicinity of the fundamental frequency reveals virtually no differences in shape, though the input, as would be expected, is of slightly larger magnitude.

In the case of the 2.5 kHz test signal, the results are almost the same. Again the decrease in high frequency energy is apparent. The two spectra from the faster attack filters show a peculiar roughness in the high end of the frequency range, but it is probable that this represents a system artifact rather than something arising from the dynamic filter's action. Such roughness was not reproduceable in any detail when the same run was attemped using time and frequency parameters that were scaled. In the neighborhood of the fundamental frequency, the spectra of the first two outputs, those with the fast time constants, are virtually identical to the input spectrum. Differences are very small, certainly below the resolution limits that we have set. In the case of the filter with the 10 ms. attack time, the differences in the input and output spectra are larger in the area of the fundamental. However, this is primarily due to amplitude differences rather than to any discernable differences in spectral content.

One interesting result, unrelated to our seach for identifiable non-linear components, was the fact that a fast attack time constant in the dynamic filter will not guarantee good tracking of program envelopes, even though they may have a much slower time constant. This can be seen in figures 5.10 and 5.12. Even though the two filter configurations have attack times of 0.5 ms. and 2.5 ms., respectively, the output envelope lags considerably behind the input envelope, which has a 10 ms. time constant. This implies that even an implementation with a very fast attack time will dull transients, perhaps considerably more than one might expect. The degree to which this is audible is another question, and one with which we are not concerned here.

We must conclude that, at least in the simple dynamic filter that we implemented, any addition of spectral components due to the filter's action is below the resolution of our analysis, and probably negligable.

We would suggest that any further research in this area include:

1) analysis of the filter's action over larger windows of time, and the study of longer program signal attack times.
2) analysis of more complex filters, such as the Ives two pole system or the Phase Linear series of bandpass filters.
3) analysis of the effects of more complex test signals, employing two or more frequencies, to more closely approximate the filter in actual use.
4) rigorous phychoacoustic tests to determine the actual audibility of dynamic filtering, using various attack time constants as parameters.

VII REFERENCES

[1] Putnam, R. S., "Record Scratch Noise Reduction", S.M. Thesis, MIT, Jan. 1975.

[2] Shnidman, D. A., "A Non-linear Filter for the Reduction of Record Noise", S.M. Thesis, MIT, May 1959.

[3] Dolby, R. M., "An Audio Noise Reduction System", Journal of the Audio Engineering Society, vol. 15, no. 4, Oct. 1967, pp. 383-388.

[4] Scott, H. H., "Dynamic Noise Suppressor", Electronics, pp. 96-101, Dec. 1947.

[5] Burwen, R. S., "A Dynamic Noise Filter", Journal of the Audio Engineering Society, vol. 19, pp. 115-120 Feb. 1971

[6] Burwen, R.S., "A Dynamic Noise Filter for Mastering", Audio, pp. 29-34, June 1972.

[7] Ives, F. H., "A Spectral Filter Audio Noise Reduction System", S.M. Thesis, MIT, March 1972.

[8] Ives, F. H., "A Noise-Reduction System: Dynamic Spectral Filtering", Journal of the Audio Engineering Society, vol. 20, no. 7, Sept. 1972, pp. 558-561.

[9] Orban, R., "A Program-Controlled Noise Filter", Journal of the Audio Engineering Society, vol. 22, no. 1, Jan. 1974, pp. 2-9

[10] Carver, R. W., "An Autocorrelator Noise Reduction System", Audio Engineering Society Preprint, May 1975.

[11] Henke, W. F., <u>MITSYN - An Interactive Dialogue Language for Time Signal Processing</u>, MIT Research Laboratory of Electronics, 1975.

[12] Olson, H. F., <u>Music, Physics, and Engineering</u>, Dover Publications, Inc., N.Y., 1967