Computer Methods and Programs in Biomedicine 61 (2000) 125–150
Computer Methods and Programs in Biomedicine
www.elsevier.com/locate/cmpb

MASTER: a Windows program for recording multiple auditory steady-state responses

M.S. John, T.W. Picton *

Rotman Research Institute, Baycrest Centre for Geriatric Care, 3560 Bathurst Street, Toronto M6A 2E1, Canada

Received 11 March 1999; received in revised form 25 May 1999; accepted 25 May 1999

Abstract

MASTER is a Windows-based data acquisition system designed to assess human hearing by recording auditory steady-state responses. The system simultaneously generates multiple amplitude-modulated and/or frequency-modulated auditory stimuli, acquires electrophysiological responses to these stimuli, displays these responses in the frequency-domain, and determines whether or not the responses are significantly larger than background electroencephalographic activity. The operator can print out the results, store the data on disk for more extensive analysis by other programs, review stored data, and combine results. The system design follows clear principles concerning the generation of acoustic signals, the acquisition of artifact-free data, the analysis of electrophysiological responses in the frequency-domain, and the objective detection of signals in noise. The instrument uses a popular programming language (LabVIEW) and a commercial data acquisition board (AT-MIO-16E-10), both of which are available from National Instruments. © 2000 Elsevier Science Ireland Ltd. All rights reserved.

Keywords: Auditory evoked potentials; Frequency analysis; Objective audiometry; Steady-state responses

1. Introduction

The multiple auditory steady-state response (MASTER) technique provides a rapid and objective assessment of hearing. The technique is based upon the statistical evaluation of the electrophysiological responses evoked by multiple auditory tones presented simultaneously. These auditory steady-state responses can be recorded from the human scalp intermixed with the other activity in the electroencephalogram (EEG). A combination of averaging and frequency-analysis can distinguish the responses from the background EEG activity. Typically, eight continuous tones are presented (four to each ear) and each tone is sinusoidally modulated at a unique frequency. The detection of the interwoven responses becomes possible after the electrophysiological data are transformed into the frequency domain. The response to each tone is then identified at the specific frequency at which the tone was modulated. The technique thus evaluates the responsiveness of the human auditory system to several different tonal frequencies in the same time it would take

* Corresponding author. Tel.: +1-416-785-2500, ext. 3505; fax: +1-416-785-2862.

E-mail address: picton@psych.toronto.ca (T.W. Picton)

PII: S0169-2607(99)00035-8

to record one response if each stimulus was presented separately.

Electrophysiological responses to sounds are useful when patients are unable or unwilling to give accurate or reliable behavioral responses [1]. This occurs in newborn infants, young children, comatose or anesthetized patients, patients who have difficulty communicating because of neurological or psychiatric disorders, and patients who may be feigning hearing loss. The auditory steady state responses have several advantages over other electrophysiological procedures for assessing audiological responsiveness. For example, data can be obtained from either awake or sleeping subjects since arousal state has minimal effect on the steady-state responses to stimuli presented at rates of 70 Hz or faster [2,3]. Responses can be rapidly obtained by presenting four simultaneous stimuli in each ear (a total of eight stimuli), provided the carrier frequencies (f_c) are separated by at least one half-octave and the modulation frequencies (f_m) are separated by at least 1.3 Hz [3,4]. Moreover, because the stimuli are sinusoidally modulated tones, the stimulus energy contains less spectral 'splatter' than transient stimuli such as tonepips.

This article discusses the general principles underlying the MASTER system, which includes both software and hardware, and which can be used with a PC running Windows95 to collect and analyze auditory steady-state responses. While the system was primarily designed to facilitate research with the auditory steady-state responses it may also be used by clinicians and audiologists to assess hearing in patients who require objective audiometry.

2. Background

Human auditory steady-state responses were first recorded by Galambos et al. [5], using stimuli presented at rates near 40 Hz. Subsequent studies showed that steady-state responses may be recorded over a wide range of other stimulus rates [6,7]. Since the 40-Hz responses fluctuate with the level of arousal [8], faster rates (>70 Hz) may be preferable for audiometric purposes because the responses at these rates are little affected by arousal [2,3]. Combining both amplitude-modulation (AM) and frequency-modulation (FM) produces larger auditory steady-state responses than either type of modulation alone [2]. Steady-state responses can be statistically distinguished from the background electroencephalographic activity in which they normally occur by means of several frequency-domain techniques [9–13]. The idea of presenting several stimuli at once was first investigated over three decades ago by Regan and coworkers [14,15], who showed that steady-state responses elicited by several simultaneously presented visual stimuli could be recorded and analyzed independently as long as each stimulus was presented at a different rate. More recently, Lins and Picton. [3] extended these findings into the auditory modality and recorded responses to multiple amplitude-modulated tones.

One of the most important applications for the auditory steady-state responses is in assessing hearing in newborn infants. Otoacoustic emissions can provide a rapid screening test for hearing impairment but cannot determine the severity of the hearing loss. Auditory steady-state responses perform well in this context [16–19], and protocols are being developed to monitor the treatment of the detected hearing impairment [20]. Accordingly, we decided to develop a system which would use standard hardware and a flexible programming system, thus allowing easy modifications of a general steady-state response program with evolving techniques.

After reviewing various data acquisition platforms, we chose to create our programs using LabVIEW from National Instruments. LabVIEW™ is a graphical programming language which lets the user create 'virtual instruments' (VIs) that look and act like real instruments [21,22] and which is tailored to work with a large number of data acquisition cards. LabVIEW is attractive because it contains a vast number of functions which enable system control and data acquisition. Other researchers have used it to build systems for data acquisition and analysis in neurophysiology [23,24], cardiovascular research [25], clinical audiology [26], neurosurgical monitoring [27], and intraoperative monitoring [28].

We also chose LabVIEW because it operates in a Windows95 environment which supports communication between programs (e.g. for evaluating data in spreadsheet-programs), and because it allows the program to be easily modified for different protocols through the interface of the Windows operating system. For the data acquisition board we chose a National Instruments AT-MIO-16E-10 board, which provides two channels of digital–analog (DA) conversion (12-bit resolution), and 16 channels of analog–digital (AD) conversion (also 12-bit resolution). The present system uses only one of the AD channels because we have not yet finished developing techniques to evaluate multiple channel recordings. Prior to the development of the MASTER system, we used a more rudimentary program written in C, the major drawbacks of which were the lack of real-time processing and display, the need for custom-built interface boards, and the lack of user-friendliness due to the non-Windows environment.

3. Overview of the MASTER technique

A steady-state evoked response is a repetitive evoked potential "whose constituent discrete frequency components remain constant in amplitude and time over an infinitely prolonged time period" [29]. In actual practice, an auditory steady-state response may be elicited by presenting an auditory stimulus at a rate that is sufficiently rapid to cause the brain's response to any one sound stimulus to overlap the response to the preceding stimulus. This technique results in a response containing frequencies at the rate of stimulation and/or at harmonics of this rate. The sinusoidal modulation of the amplitude or frequency of a continuous carrier frequency may also be used to evoke steady-state responses [6,30,31]. Because the modulation frequency evokes the responses these have also been called envelope-following responses [32].

The MASTER technique presents multiple tones, each of which is modulated at a unique frequency. The compressive rectification occurring in the ear causes each tone to evoke a response at the frequency by which it is modulated [7]. Because the scalp-recorded activity contains the multiple superimposed responses to the various stimuli, it is difficult to distinguish the individual responses in the time domain. However, if the data are converted into the frequency domain using a fast Fourier transform (FFT), the amplitude and phase of the response to each stimulus can be measured at the specific frequency at which the tone was modulated.

As well as identifying responses at particular stimulus-modulation frequencies, MASTER also evaluates a noise estimate that is derived from neighboring frequencies in the spectrum at which no stimulation occurred. If there were no response, the power at the modulating frequency would be within the range of the power at these neighboring frequencies. An F-ratio can estimate the probability that the amplitude at a stimulus-modulation frequency (signal) is within the distribution of amplitudes at the neighboring frequencies (noise) [11–13,18]. When this probability is less than 0.05, the response is considered significantly different from noise, and the subject is considered to have heard that tone that corresponded to that frequency of modulation.

In general as more data is collected the signal-to-noise ratio increases. Several approaches are used to utilize the incoming data in order to increase the signal to noise ratio of the recording. Averaging together repeated recordings or 'sweeps' reduces the level of activity in the recording that is not time-locked to the stimuli. Selectively rejecting recordings wherein the noise level is particularly high (usually because of noncerebral potentials or 'artifacts') prior to averaging increases the efficiency of the averaging process [33]. Artifact rejection protocols are set up in MASTER to operate only on sections of the recording sweep or 'epochs.' Increasing the duration of the activity submitted to the FFT increases the frequency-resolution of the analysis. Provided the stimulus frequencies are precisely locked to the recording frequencies, increasing the sweep duration thus reduces the noise level, by distributing the power across more FFT bins, without affecting the amplitude of the response which is represented within a single FFT bin [4].

When converting time series data into the frequency domain the amount of data submitted to the FFT becomes an issue. The specific frequencies available from an FFT operation are integer multiples of the $\mathbf{1/(\mathit{Nt})}$ resolution of the FFT, where N is the number of time points and t is the time between each data point. During data collection a series of data sections or 'epochs' are collected. In the experiments that will be discussed in this paper the digitization of the electrophysiological signals occurred at 1000 Hz and the buffer had a length of 1024 points. Therefore, each epoch lasted 1.024 s. Because submitting single epochs to the FFT analysis would produce a frequency resolution of only 0.98 Hz, the individual epochs are linked together to form longer data windows termed 'sweeps'. By using 16 epochs in each sweep, data windows of 16 384 time points were created, producing a frequency resolution of $1/(16 \times 1.024 \times 0.001)$ or 0.061 Hz in the amplitude spectrum which spanned from 0 to 500 Hz (the Nyquist frequency).

Not all the data that were collected were submitted to the FFT routine. Artifact rejection was accomplished by rejecting any epoch containing a voltage value that exceeded ± 50 μV. If an epoch was rejected, the next epoch which was not rejected was used in its place when linking the epochs of data into sweeps. This procedure does not violate assumptions of stationarity because each of the stimuli which evoked the responses were constructed so that each epoch contained an integer multiple of the modulation and carrier frequencies and started with identical phases. More detailed discussions of the processes of stimulus generation and artifact rejection occur later in this paper (see Section 5.2 and Section 5.5).

In general, as more data are collected the signal-to-noise ratio is improved. We have shown that the signal-to-noise ratio can be reduced with similar efficiency by either averaging the data in the time domain or using the data to increase the duration of the sweep which will be submitted to FFT analysis [4]. We use both techniques. A sweep length is selected that will provide sufficient frequency resolution for the statistical analyses, and sweep recordings are averaged together to reduce the noise. When enough data for an entire sweep is obtained, an FFT analysis computes the spectrum. When the next sweep is collected it is added to the prior sweep in the time domain and the result is again submitted to FFT analysis. This is more efficient than combining the spectra, which would require vector averaging to consider phase as well as amplitude. The responses are then evaluated using the previously described F-ratio technique. When the F-ratio of a response at a particular modulation frequency is significant at $P < 0.05$ then the subject is considered to have heard the carrier frequency. This technique therefore offers a rapid and objective method to test hearing.

4. System description

4.1. Hardware and peripherals

In addition to the digital signal processing that occurs in the MASTER system, recording the auditory steady state responses requires analog signal conditioning for both the auditory and the electroencphalographic signals. The current MASTER system is relatively inexpensive because it relies upon other instruments, which already exist in most laboratories and audiology clinics, to amplify and filter these signals (Fig. 1). In addition to the software, the system includes a connection box which contains a series of input and output connectors and several resistor bridges for voltage-attenuation. For example, the voltage signals (± 10 V) from the DA outputs of the board are attenuated by 10 at two of the terminals so that they may be connected to the 'tape input' of a clinical audiometer. This audiometer enables the operator to adjust the intensity levels of the stimuli and to select the transducer (earphone, free field speaker, bone conduction vibrator, etc.). For the data presented in this paper, a Grason-Stadler Model 16 Audiometer was used. Additionally, the electrical activity recorded from the scalp (usually in the range of ± 20 μV) must be amplified before reaching the AD input of the data acquisition board. For the data presented in this paper, we used a small, battery-operated EEG amplifier

(Model P55, Grass Instruments) with a gain of 10 000, a high-pass filter of 1 Hz and a low-pass filter of 300 Hz (− 3 dB points, 6 dB/octave).

4.2. Main screen

When a user starts the MASTER program the Main screen appears (Fig. 1). The left side of the screen contains buttons for establishing the experimental protocol and ensuring proper system performance, and the right side concerns options for the collection and analysis of data.

Although the figures in this article appear in black and white, the actual screens use a simple color scheme to designate the functions of various components. For example, three shades of blue are used: a light blue to indicate information which is important to the operation of the screen; a dark blue for buttons that evoke various operations within the program; and a middle blue for the background. White lettering is used to provide descriptive labels for the parameters displayed in the screens. Three highlight colours are also used: yellow to highlight spectral data values at the modulation frequencies, and to indicate controls that can manipulate the processing of the recorded data; red to display cautionary information (such as illegal values); and green to signify that given parameter value is acceptable.

4.3. Load Protocol screen

Pressing the 'Load Protocol' button invokes a screen for defining experimental parameters (Fig. 2). This screen contains controls that allow the user to generate various types of stimuli (top section of screen) and configure AD and DA operations (bottom section). Control values can be modified by mouse-clicking on the arrows that are adjacent to each control value, by entering the values directly from the computer's keyboard, or by loading a set of stored values from disk. Since there are numerous parameters for any given experiment, the user will commonly load values stored in text files that have '.pa1' or '.pa2' filename extensions, e.g. 'default.pa1.' To create new protocols the user can adjust the parameters and then store them in a new file by pressing the 'write file' button. Long file names are supported under LabVIEW. Both reading and writing from

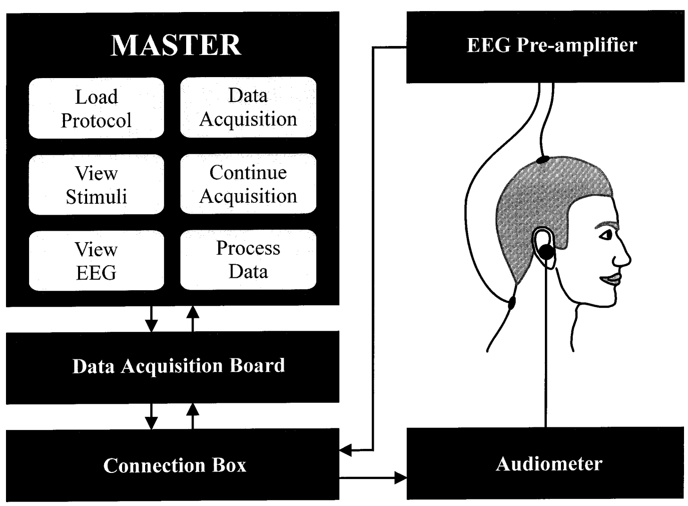

Fig. 1. MASTER system. The system includes software running on a PC which operates a data-aquisition board. The DA outputs of this board are sent via a connection box to an audiometer which acts to control the intensity and the presentation of the stimuli to the subject. A pre-amplifier is used to amplify the subject's electrophysiological responses to a level appropriate to the AD input. The program begins with a Main screen which shows the various options which are available to the user. Choosing any button invokes another screen dedicated to the selected option. When the specific task of that screen is completed, the user is returned to the Main screen.

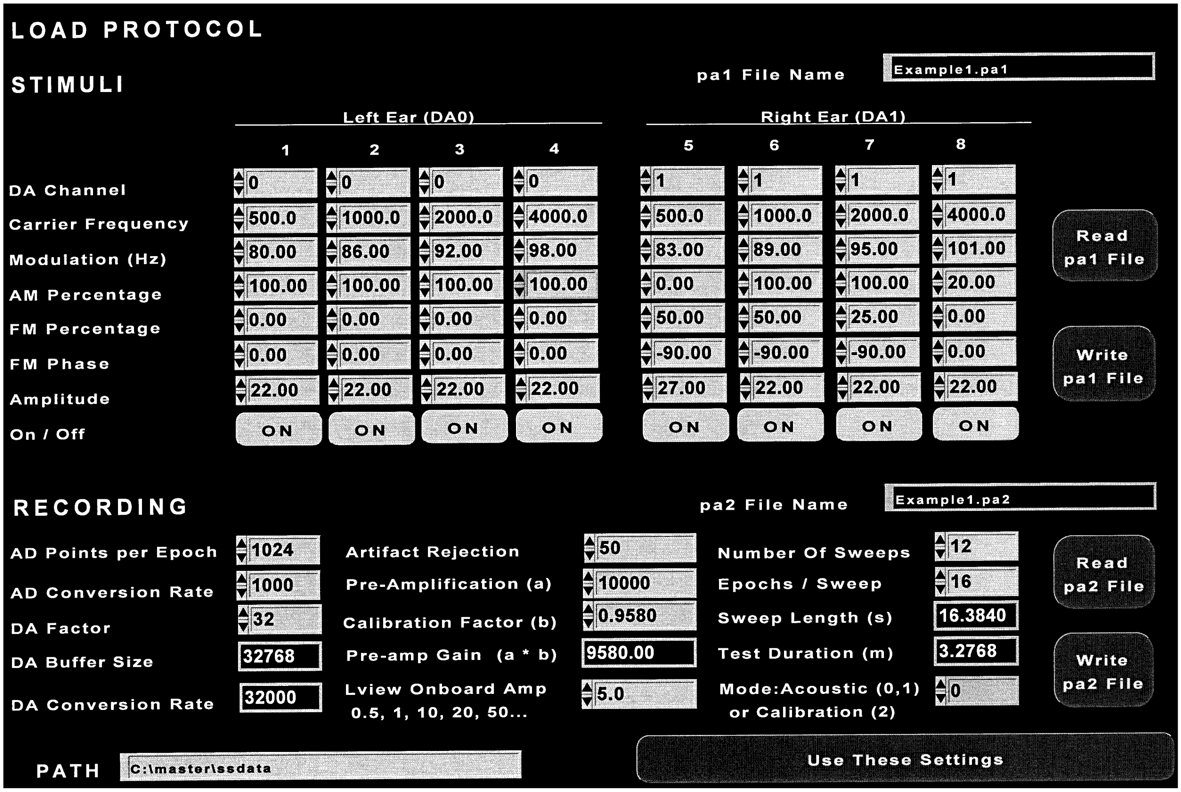

Fig. 2. Load Protocol screen. The top of this screen is used to define up to eight stimuli which may be either amplitude-modulated, frequency-modulated or both. The stimuli defined for the left ear are typical of what would be used to record responses. The stimuli defined for the right ear are only to illustrate important aspects of the stimuli, which are displayed in Fig. 3. The lower part of this screen enables the user to change parameters related to recording the responses.

disk use standard Windows dialogue boxes for file operations (e.g., 'Open File'). Since pa1 and pa2 files are text files, the user can also modify the parameters for an experiment using a standard text editor. However, it is generally safer to edit parameters from within the MASTER program to ensure that one is not choosing 'illegal' values. Many of the parameters must conform to certain principles so that processes like the synchronous timing of the AD and DA conversions occur successfully.

As is indicated by the controls located in the top section of the screen, the MASTER system is currently configured to present a maximum of eight stimuli. The most common testing protocol utilizes four stimuli in each ear since previous research has shown that using a greater number of stimuli can decrease the amplitude of the responses [4]. This typical setup is indicated by the titles on the screen which indicate which stimuli will normally be presented to the left and right ear. However, the software can actually present up to eight stimuli in one ear if none are presented to the other ear. For each stimulus the user defines the following options:

- *DA channel:* determines the DA channel to which the stimuli are sent (0 is for the DA0 output which usually goes to the left ear, and 1 is for DA1 output which goes to the right ear).
- *Carrier frequency:* usually between 0.5 and 8 kHz, determines the center frequency of the stimulus.
- *Modulation frequency:* determines the frequency of the modulation envelope, between 80 and 200 Hz for the rapid auditory steady-state responses, or near 40 Hz for the middle latency steady-state responses.

- *AM percentage:* determines the amount of amplitude-modulation expressed as a percentage (0–100%) according to the formula $\mathbf{100(a_{max} - a_{min})/(a_{max} + a_{min})}$ where $\mathbf{a_{max}}$ is the maximum amplitude of the signal and $\mathbf{a_{min}}$ the minimum amplitude. When $\mathbf{a_{min}}$ is zero the amount of modulation is 100%.
- *FM percentage:* determines the amount of frequency modulation expressed as a percentage (0–100%) according to the formula: $\mathbf{100(f_{max} - f_{min})/f_c}$, where $\mathbf{f_{max}}$ is the maximum frequency of the signal, $\mathbf{f_{min}}$ the minimum frequency, and $\mathbf{f_c}$ the carrier frequency. For example, when a carrier frequency of 1000 Hz is frequency-modulated at 20% the frequency moves between 900 and 1100 Hz, i.e. $\pm 10\%$. The other term used to describe FM is $\mathbf{\Delta f}$ which is the excursion between the middle frequency (the carrier) and either $\mathbf{f_{max}}$ or $\mathbf{f_{min}}$. In this example, $\mathbf{\Delta f}$ is 100 Hz or 10% of the carrier frequency.
- *FM phase:* determines the phase (in degrees) of the frequency modulation relative to the amplitude modulation so that changes in the frequency of the stimuli can occur either in-phase or out-of-phase with those of the AM. Since the maximum frequency will occur at 180° sine phase and the maximum amplitude will occur at 90°, a phase setting of $-90°$ will ensure that the tone that is both amplitude- and frequency-modulated will reach maximum frequency at the same time as it reaches maximum amplitude (see Fig. 3).
- *Amplitude:* determines the amplitude of the stimulus expressed as a percentage of the maximum voltage allowed by the output buffer (± 10 V). Since the stimuli within each ear are added together, when four stimuli are used, each of the four stimuli should be defined to use no more than 25% of the total range.
- *On/Off:* determines if the stimulus will be presented to the subject. This control allows the operator to select a subset of the stimuli for presentation.

The lower half of the screen enables the user to set the experimental parameters for the data acquisition and output buffers. The left column permits the user to modify the AD and DA conversion rates and the buffer sizes. After a sampling rate is chosen, the user chooses the size of the AD buffer. For optimal computational performance (i.e. using the fast Fourier transform rather than a slower discrete Fourier transform) the number of points in the AD buffer should be set equal to a power of two. The duration of a buffer is a function of the number of points in the buffer and the rate at which the buffer is filled or emptied. In the example shown in Fig. 2, where the number of points in the buffer is set to 1024 and the input conversion rate is set to 1000 Hz, it will take 1.024 s for the buffer to become filled. Each time the data buffer is filled, one epoch of the recording is completed. When choosing AD conversion rates, the operator must consider the range of frequencies that are present in the incoming data. In order to prevent aliasing, the AD conversion rate should be greater than twice the highest frequency in the input. Thus, when the low-pass filter of the pre-amplifier is set to 300 Hz, the AD conversion rate should be at least 600 Hz.

It is essential that the AD and DA buffers have an identical duration (in this example 1.024 s) so that the input and output buffers remain synchronized during the recording. In order to make the duration of these two buffers identical, the parameters for the DA operations are calculated as a function of the AD parameters. Accordingly, the operator enters a ‘DA factor.’ In the example shown in Fig. 2, the operator has chosen a DA factor of 32, giving a DA buffer size of 32 768 points and a DA rate of 32 000 Hz. If one wishes to view the spectrum of the stimuli using the View Stimuli screen, it is helpful to make the DA factor a power of two so that the FFT algorithm can be used to rapidly calculate the spectrum (provided the size of the AD buffer is also a power of 2). The screen uses the convention that numbers derived from other inputs (e.g. the DA buffer size which is derived from the AD buffer size and the DA factor) are shown in light blue on a dark blue background.

Normally, we use a maximum carrier frequency of 8000 Hz in our audiometric testing procedures and a DA rate of 32 kHz, since this will present the stimulus without significant distortion. Increasing the D/A rate above 32 kHz is not permitted because of the speed limitations of the AT-MIO-16E-10 board. The screen provides a warning in the form of a flashing red indicator if the numbers entered are not possible given the clock resolution and/or the processing speed of the board. These red indicators occur if the DA conversion exceeds 32 kHz or if the rate of either the AD or DA conversion is not an integer fraction of the board's 20-MHz clock rate. The operator can then change the entered values until the red light is turned off. (Because the MASTER program may in the future be used with different boards having different clocks, it has not been programmed to detect discrepancies based on this particular board's clock rate and adapt the input automatically).

The central column contains controls which enable the user to modify parameters which affect the acquisition and scaling of the electrophysiological data. The artifact rejection level deter-

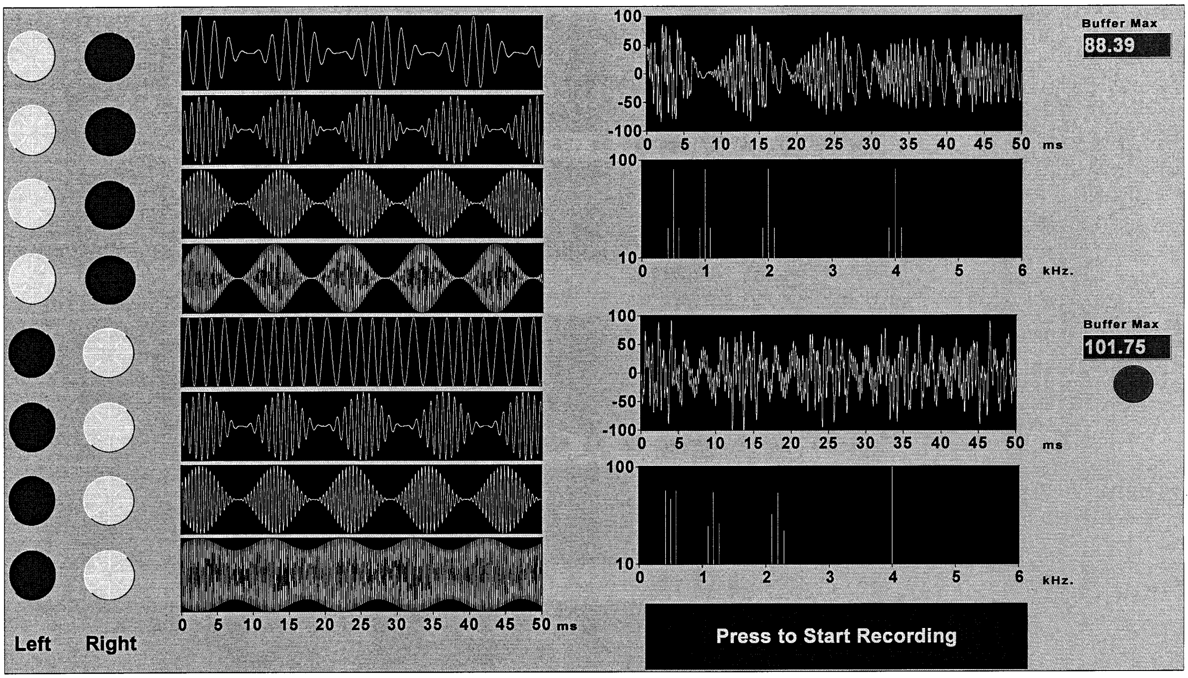

Fig. 3. View Stimulus screen. Up to eight stimuli can be displayed here and summed together to create the stimulus that will be presented to a subject. Note that only the first 50 ms of each stimulus are shown. The top right hand graph shows a digital representation of the stimulus that was created by adding the first four stimuli in the left hand column, each of which has been amplitude-modulated at 100% as is indicated by the parameters shown in Fig. 2. Under this waveform is the amplitude spectrum of the stimulus showing four peaks in the spectrum corresponding to each of the carrier frequencies together with flanking side-bands related to the modulation. In order to show the spectra more clearly, the x-axis of the graph has been changed so that extends only to 6 kHz rather than the usual 10 kHz. The four stimuli presented to the right ear demonstrate what occurs when combining AM and FM. The first stimulus is only frequency-modulated. The phase is set at −90° so that the peak frequency occurs at the same time as the peak amplitude would have occurred if the stimulus were amplitude-modulated (this stimulus can be compared to the first stimulus in the left ear which also has a carrier frequency of 500 Hz and which is modulated a little more slowly). The second and third stimuli in the right ear are both AM and FM. Because the largest amplitude occurs at the highest frequency, the spectrum is not centered on the carrier frequency but is shifted upward. This shift is very clear for the second stimulus which is 50% FM, and is a little smaller for the third stimulus which is only 25% FM. The fourth stimulus shows a 20% AM stimulus. The amount of modulation is sufficiently small that the side-bands are below the lower limits of the y-axis. The 'buffer max' indicates that the right ear stimulus exceeds the buffer. Even though the sum of the amplitudes $(27 + 22 + 22 + 22)$ is less than 100, the peak amplitudes of the AM stimuli (the second and third in particular) have been increased to maintain the RMS intensity. Unless the stimuli are changed the DA converter will clip.

mines the value which must be surpassed in order for a recording epoch to be rejected. The amplification factor and calibration factor are entered so that the data can be correctly scaled and displayed in the proper units. The AT-MIO-16E-10 data acquisition board has its own onboard amplification which can be adjusted to a factor of 0.5, 1, 5, 10, 20, 50 or 100. The range of the AD conversion after the onboard amplifier is ± 5 V. The optimal configuration occurs when the pre-amplification and onboard amplification are set so that the full 12-bit range of the input buffer is used to represent the incoming analog signal. When the pre-amplification is set at 10 000, typical EEG signals which are ± 20 μV (peak to peak) are amplified to ± 200 mV. Using an onboard gain of 5 increases the incoming signal to ± 1 V which reproduces the signal well without clipping. The clipping occurs at ± 5 V after the onboard amplification, equivalent in this case to ± 100 μV in the EEG.

The two top controls in the third column enable the user to determine the amount of data that will be collected and the number of data epochs which will be used to compose each data sweep. By increasing the number of epochs that are concatenated together to form a sweep, the operator will increase the frequency-resolution of the FFT. The spectral analysis will utilize the FFT if the result of multiplying the number of epochs per sweep and the number of points in each epoch produces a number of points which is a power of two (e.g., 16 384). If the number of points is not a power of 2, the program will run and the amplitude and phase of the responses will still be calculated correctly but the program will use a discrete Fourier transform, rather than zero padding the data and using the FFT. In order to maintain a reasonable speed of processing, we therefore recommend that both the number of points in an epoch and the number of epochs in a sweep be set to powers of 2.

Since the currently defined AD epoch will last 1.024 s and each sweep has been set to contain 16 epochs, each sweep will last 16.384 s. Further, since the user has indicated that 12 sweeps will be collected, 192 epochs will be collected, causing the recording session to last about 3.3 min. If the user wishes to increase the frequency-resolution of the amplitude spectrum by using a sweep containing 32 rather than 16 epochs, the program would automatically adjust all the parameters and internal computations in accordance with this change.

The last control of the third column determines the type of stimuli that will be used during the recording period. The user selects whether the stimuli are presented according to one of three modes. In Mode 0 the stimuli have a constant root-mean-square amplitude, while in Mode 1 the stimuli will have a constant peak-to-peak amplitude. Mode 2 presents stimuli consisting of only the modulating frequencies (without any carrier frequencies) which are used to calibrate the EEG amplifier (see Section 5.2 and Section 5.6).

Like the stimulus parameters, the recording parameters may also be saved in and loaded from text files. The files for the recording parameters have the filename extension 'pa2'. The use of separate pa1 and pa2 files allows users to mix and match different parameters for stimuli and recording. For example, one might wish to use the same stimuli with two different pre-amplifiers, or with different AD conversion rates. Lastly, at the bottom of the screen is a control which allows the operator to enter a path that determines the directory where the parameter files are stored and where the raw data, and results obtained during the recording, will be stored.

4.4. View Stimulus screen

The acoustic stimuli used in the MASTER program are carrier frequencies that may be either amplitude-modulated, frequency-modulated, or both. Fig. 3 illustrates the stimuli that were generated according to the control values displayed in the experimental parameter screen in Fig. 2. In the left ear, there are four individual carrier frequencies (f_c), each of which is amplitude-modulated at a unique modulation frequency (f_m). These four are added together to create the stimulus that will be presented to the left ear. The time course of this stimulus (for the first 50 ms) can be seen in the upper graph of the right-hand column. Below that is the amplitude spectrum which indicates that the stimulus is comprised of the original four carrier frequencies, each with two sidebands at $f_c \pm f_m$. To the right of these plots is an indicator showing the results of a buffer check. The output buffers are defined so that they can output values ranging from ± 100 (which produces the full ± 10-V range of the output buffer). Since the maximum value of the stimulus which is to be presented to the left ear is 88.39, this stimulus can be successfully represented in the buffer. However, the stimulus constructed for the right ear exceeds the maximum. Since all values over 100 would be clipped at 100, the right ear stimulus would be significantly distorted if used in an actual experiment. Because peak-to-peak amplitudes of the stimuli may be altered to maintain a constant RMS in Mode 0, this buffer check is non-trivial (illustrated in Fig. 3).

In the figure the left ear stimuli are typical of what might be used to evoke auditory steady-state responses. The stimuli defined for the right ear were created for demonstration purposes only and would not be used in an actual experiment. The first stimulus is a carrier frequency that is frequency-modulated at 50%. The second stimulus is a carrier with 100% amplitude modulation and 50% frequency modulation. The phase of the frequency modulation has been set to –90° in order to align the maximum frequency of the stimulus with the maximum amplitude. In the third stimulus 100% AM and 25% FM (phase still –90°) were used. The fourth stimulus shows only 20% AM and no FM.

4.5. View EEG screen

Before beginning a recording session a user will normally choose to view the incoming electrophysiological data. Pressing the View EEG button, located in the bottom left corner of the main menu, will invoke a simple screen which displays the subject's EEG. On this screen, the user can change the x- and y-scales, or the A/D rate, to observe the quality of the incoming EEG signal For instance, the user may increase the sampling rate to 8000 Hz to ensure that a slow frequency is not due to the aliasing of a high frequency.

4.6. Data Acquisition screen

Fig. 4 shows the screen used during data acquisition. The data which are displayed in the figure were obtained from a recording session which presented stimuli to both the left and right ears. In the upper left section of the screen is a chart labeled 'EEG' which displays the incoming electrophysiological signal. Its y-axis spans ± 100 μV and its x-axis automatically adjusts to show the all the data points of the current epoch (with the number of points rounded down to the nearest 100). If a user wishes to display fewer time points of each sweep, the x-axis values are configurable, and may be changed by clicking on the value with the mouse and entering a new value from the keyboard. The horizontal lines that appear at ± 60 μV on this chart show the artifact-rejection limits (these are red on the computer screen). If the current epoch contains values which exceed these limits, it is rejected, i.e. not averaged and not stored in the raw data file.

In the upper right of the screen is a graph of the amplitude spectrum of the average response. The amplitude ($\boldsymbol{a}$) is measured as the base-to-peak amplitude of the cosine wave, $\boldsymbol{a}\cos(2\pi f)$, provided by the FFT. The data submitted to the FFT is the average response of the subject and is obtained by adding each recorded sweep to a running sum and dividing this by the number of sweeps collected. Therefore this display is updated each time enough artifact-free data epochs have been collected to add an additional sweep to the average. The x-axis normally spans from 0 to 200 Hz, while the y-axis is automatically scaled according to the highest amplitude value. The axes may also be configured by the user by clicking the mouse on a value on the axis and inputting a new value from the keyboard. Additionally, because the lower frequencies usually have higher amplitudes than higher frequencies, the user can click on the button labeled 'Auto' to change it to '$A > 70$'. In this mode, the scaling will set the maximum of the y-axis to the maximum amplitude of frequencies above 70 Hz. The X-scaling button can also be used to change the frequency-range of the spectrum. On the computer screen the amplitudes of the spectrum which occur at

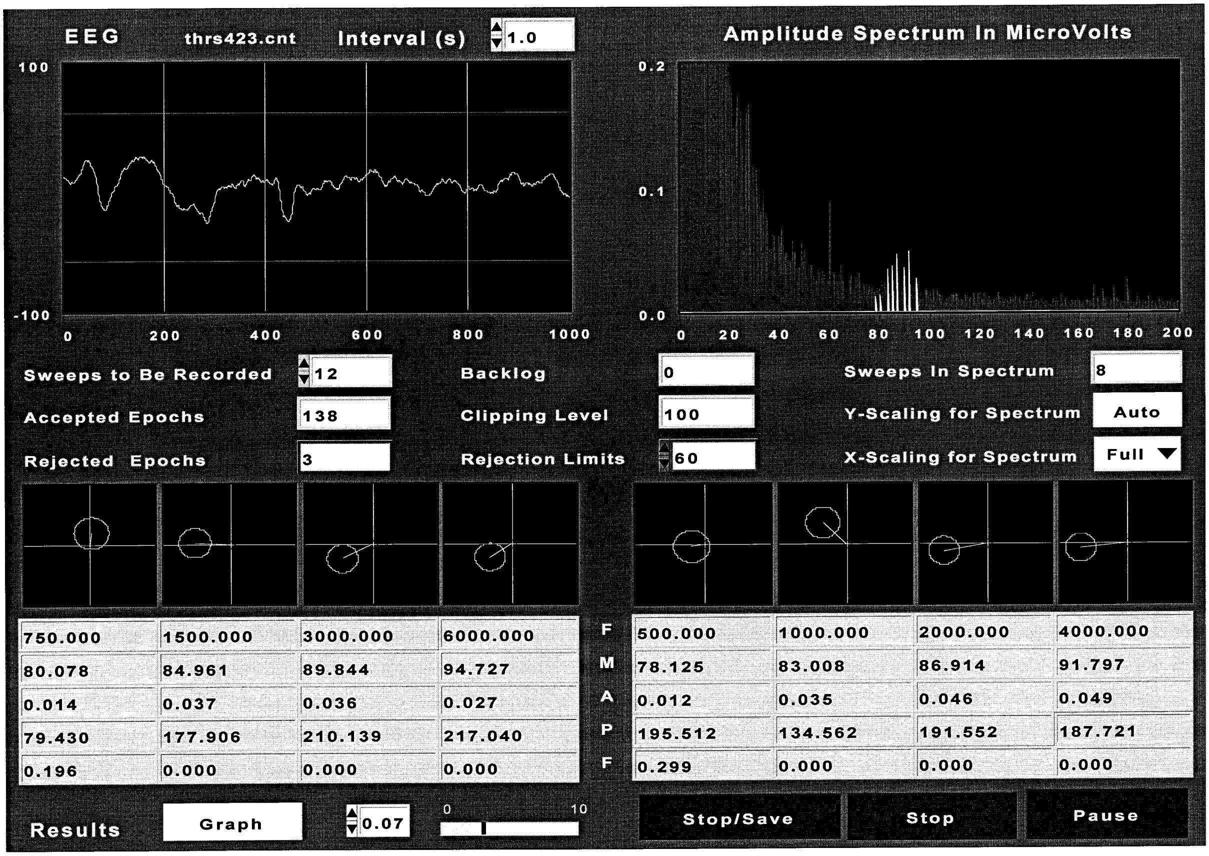

Fig. 4. Data Acquisition screen. The incoming EEG is monitored in the upper left graph. This figure was obtained by reading data from a disk-file, the name of which is presented at the top of the graph. The spectrum of the average responses is shown in the upper right graph. The frequencies in the spectrum corresponding to the modulating frequencies are highlighted. For this figure the y-axis of the spectrum was changed amplitude so that both the responses and the residual EEG spectrum are easily visible. The normal options are to set the y-axis automatically to the maximum value (which would make the responses small) or to the maximum value above 70 Hz (which would clip the lower frequencies). Another option would be just to display the spectrum above 80 Hz (X-scaling button). A line noise peak can be seen at 60 Hz (in order to avoid phase distortion we chose not to use any analog notch-filtering). Responses can be recognized between the frequencies 80 and 100 Hz. There are also smaller peaks recognizable at the second harmonics (between 160 and 200 Hz). In the lower half of the screen are shown polar plots of the responses to each of the stimuli. Responses to six of the eight stimuli being presented to the subject have become significant after only about 2 min of recording. The significance of the response can be seen monitored by seeing whether the circle incorporates the intersection of the axes, or by checking the probability of the F-value (F). The responses to 500 and 750 Hz have not reached significance.

frequencies corresponding to the modulation frequencies are highlighted in yellow. In Fig. 4, six of the eight peaks corresponding to the modulation frequencies are clearly larger than the background activity.

Below the displays of the EEG and amplitude spectrum of the average response are several indicators which allow the operator to monitor and in some cases control the data acquisition:

- *Sweeps to be recorded:* indicates the number of sweeps required for a complete recording. This is a control as well as an indicator. During data collection the operator can thus increase or decrease the number of sweeps according to how well the responses are being recognized.
- *Accepted epochs:* indicates the number of epochs that did not exceed the rejection threshold.

- *Rejected epochs:* signifies the number of epochs that exceeded the rejection criteria. Rejected epochs do not enter into the averaging and are not stored with the raw data.
- *Backlog:* shows the amount of data left over in the software buffer after it has been read. If this number intermittently deviates from zero, or becomes very large for only a moment, then the experimental protocol may be near the memory and processing capability of the computer (see Section 5.1 for details).
- *Clipping level:* indicates the level at which the AD converter clips the EEG signal. This information is useful for setting the artifact-rejection limits (which should be less than the clipping level), and for checking that the amplifier settings are correct.
- *Rejection limits:* indicates the current threshold for rejection. This value can be changed online during the experiment to adjust to different states of the subject.
- *Sweeps in spectrum:* indicates the number of sweeps in the current average response, which has been submitted to the FFT routine to give the displayed spectrum.
- *Y-scaling:* This button determines whether the y-scale of the spectrum display depends on the whole spectrum or just that part of the spectrum greater than 70Hz.
- *X-scaling:* This button determines whether the x-scale of the spectrum display starts from zero, from 80 Hz or from 160 Hz.

The lower half of the figure displays the evaluations of each responses both graphically and numerically. The left four graphical polar plots represent the responses produced by stimuli presented to the left ear. Each plot shows the magnitude of the response as a vector which extends from the origin outward. The phase of the response (relative to the phase of the modulating frequency) is represented as the angle between the line and the x-axis. The 95% confidence interval for the noise is represented by a circle centered at the end of the vector. As long as the circle does not include the origin, a response lies outside of the 95% noise estimate. In other words, there is only a 5% chance that there is no real response and that the measurement occurred by chance.

This display is based on that used in the Hotelling's T^2-test [9,10], which assesses the variability of the average response by making multiple measurements of the response in subsections of the recording. In the T^2 technique the circle indicates the confidence limits of the mean measurement. The F-test is different because it measures confidence limits of the activity recorded at frequencies other than the signal (i.e. a noise estimate). The diagram might therefore more properly show the circle around the origin and the vector of the response either exceeding or not exceeding the radius of the circle, for significant or non-significant responses, respectively. However, because the T^2- and the F-tests yield equivalent results [34], the display uses the plotting convention of the T^2-test, since this makes it visually easier to see the significance of the response. In this record, the first four graphs indicate that a significant response has occurred for carrier frequencies at 1500, 3000, and 6000 Hz. while the 750-Hz tone (far left) does not produce a response that is significantly different from background noise levels. The second four responses, due to stimuli presented in the right ear indicate that the 500-Hz tone failed to produce a significant response. The failure of the 500- and 750-Hz tones to evoke a significant response is not unexpected, because in some subjects about 6 min of recording is required for these frequencies to produce significant responses. The scaling for these plots can be changed using either the sliding control or the digital control located at the bottom of the screen.

The responses are also described numerically in the eight columns of numbers located below the polar plots. The left four columns show the current results for the left ear stimuli, while the right four contain data for the right ear. The results for each response are displayed using five values: carrier frequency, modulation frequency, amplitude of the response, phase of the response, and F-ratio level of significance for that frequency. The first two values of each column describe the stimuli and therefore remain constant, the last three values change as the recording progresses.

The user may choose to change the polar graphs to a simpler representation of whether or not the response is significant. Depressing the button labeled 'Graph' under the tabulated results replaces each of the polar plots with a colored circle. A green circle means that the probability that the response occurred by chance is less than 5%. An orange circle means that the probability is 5–10%, while a red circle indicates that it is more than 10%. The ability to change the display online demonstrates an attractive feature of LabVIEW which does not exist in other programming languages. LabVIEW offers the ability to toggle the state of any component of a screen from visible to invisible. This is a powerful feature because, entirely different screens can be created for users of different skill levels (e.g. scientist, clinician, student), with each of the screens still using the same source code.

In the lower right corner of the screen are three buttons which can be used to halt the operation of the program. The button labeled 'Stop' is used to halt the entry of new data into the recording, while the input and output buffers continue their operation. This can be used if the data acquisition needs to be halted in the middle of a recording session, as might occur if the user needs to reattach an electrode to the subject's scalp because it fell off before the end of the testing procedure. When the problem has been fixed, the user can then choose 'Continue Acquisition' from the Main screen (Fig. 1), and the subsequently collected data will be appended to the data that were already collected. The button labeled 'Stop/Save' will stop all input and output operations, clear the input and output buffers, and save any data that has been collected prior to that point. This can be used if the user feels that enough data have been collected and wants to continue on to another recording.

Two features of this screen are only available during the playback of data from disk and are not available during data collection due to the synchronization requirements of the input and output buffers. The 'Pause' button is used to halt the program only when reviewing data from a recording that had been previously stored to disk. Additionally, a control labeled 'Interval' which is located above the chart of the incoming EEG data is used during data review. This control determines how long each data epoch is displayed on the screen before the next epoch is read from disk.

4.7. Process Data screen

After the data have been acquired and stored, the operator may wish to combine these data sets across subjects or across replications within one subject. The Process Data screen allows the operator to work on data sets which have already been recorded. Three different subscreens are available through this screen: 'Review', 'Combine' and 'Log Results.'

The 'Review' features are used to examine data that have already been recorded, and to print out hardcopy figures. The figures are organized in a similar manner to the way the data are presented on the Recording Screen, except that graph of the incoming EEG is omitted and the spectrum of the average data is expanded to fill its place. The screen also allows the operator to choose the number of responses that are graphed, and to enter comments about the data.

The 'Combine' routines support either the averaging or subtracting of previously recorded data sets (Fig. 5). A user may wish to average data because, for instance, the results from a short recording session may not show a significant response even though the subject might hear the stimulus. This is usually due to the fact that the noise has not been sufficiently reduced by averaging and still obscures the responses. The results from several recording sessions must therefore be combined in order to demonstrate a significant response. As discussed earlier, computing the average of spectral results is accomplished by averaging the time series waveforms from the different files, and then re-computing the amplitude spectrum.

Sweep-weighted averaging is commonly used when combining different data sets obtained from a single subject using the same stimuli. Weighting is used because the average waveform from each recording session may have been generated from different numbers of sweeps. For example, the user might have increased the number of sweeps in one of the sessions in order to gain slightly more data. In this situation, each sub-average should be weighted during averaging by the number of sweeps in that sub-average. For two sets of data, X_1 and X_2, based on N_1 and N_2 sweeps, respectively, each time point ($i = 1$ to the size of the sweep) in the sweep-weighted average is

$$(X_{1i}N_1 + X_{2i}N_2)/(N_1 + N_2)$$

and this is then stored as the average of $N_1 + N_2$ sweeps. These calculations are performed using double precision in order to attenuate the effects of round-off error if the sum becomes large.

Alternatively, the operator may opt for non-weighted averaging when combining data across subjects in order to weight each subject equally. Merely because more data may have been collected for a particular subject does not mean that this set of data should contribute more to the average. For non-weighted averaging, the formula for combining two averages is simply

$$(X_{1i} + X_{2i})/2$$

and this is stored as the average of two 'sweeps'. The convention for the number of sweeps is arbitrary. By using the number of subjects for the number of sweeps, files can easily be distinguished as a weighted average rather than a non-weighted average. Furthermore, using the number of subjects allows one later to average together non-weighted averages using sweep-weighted averaging. For example, the average of the first eight subjects can be combined with the average of the next two subjects using weighted averaging to obtain the average of ten subjects with each subject equally weighted.

Subtracting responses may be necessary when evaluating the effects of masking. The derived response technique, for example, sequentially subtracts responses recorded using masking noise at

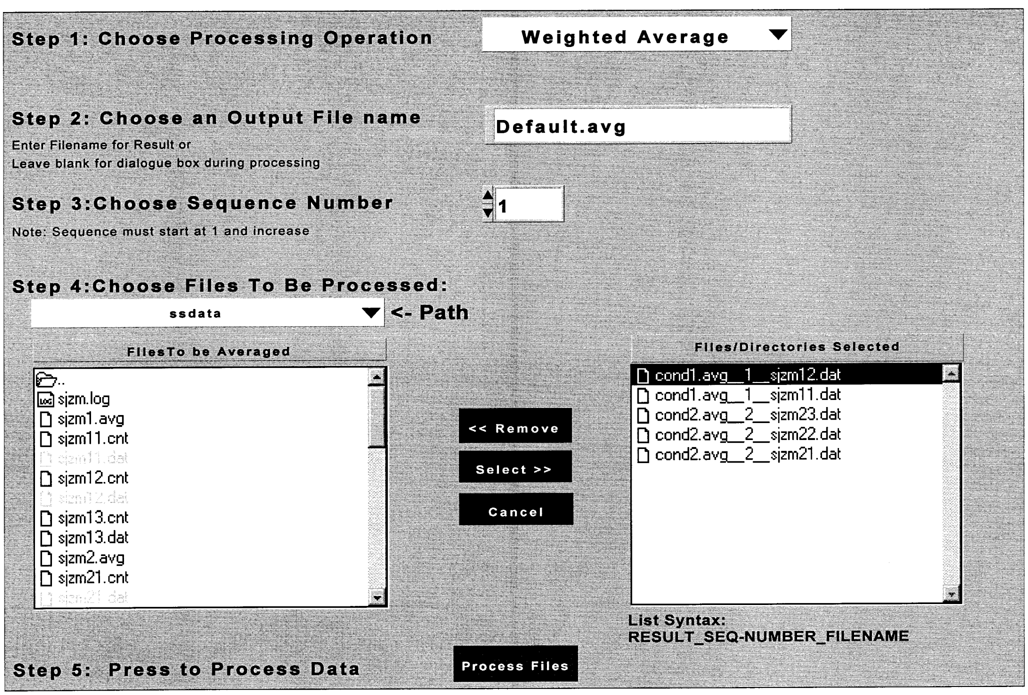

Fig. 5. Process Data screen. This screen illustrates how stored data can be combined. The operator selects the type of processing to be done. Then he or she enters the name of the file where the processed data should be stored and gives the processing a sequence number. The files to be processed are selected from those present in the path and entered in the listing on the lower right. Once everything is set up, the operator presses the 'Process Files' button. The screen shows the setup to perform a weighted average of the data in two files and to put the average data in Cond1.avg, and then to perform a weighted average of three files and to put the average data in Cond2.avg.

decreasing high-pass cut-off frequencies to determine the responses to the frequencies between the cut-off frequencies [35,36]. Subtraction must be done using unweighted calculations. Furthermore, there is no division following the subtraction. When the diminuend is $\boldsymbol{X}\mathbf{1}$ and the subtrahend is $\boldsymbol{X}\mathbf{2}$, the result of the subtraction (the difference) is

$$X_{1i} - X_{2i}$$

and this is stored as the average of one 'sweep'. Again, the convention is arbitrary.

The screen controlling the routines for averaging and subtracting are organized so that the operator inputs the required information step by step. The first step is to choose the type of processing to be used (weighted averaging, non-weighted averaging, or subtraction). The second step is to enter the name of a file for storing the results of the processing. The third step is to give the processing a sequence number. This number should be 1 if only one set of files is to be processed. The fourth step is to select the files for processing. For averaging any number of files can be selected, but for subtraction only two files can be selected (first the minuend and second the subtrahend). The selected files are then listed using the syntax 'results file–sequence number–file to be processed'. Multiple repetitions of the process can be set up by circling through steps 2–4 and adjusting the sequence number each time. Fig. 5 shows the screen that will lead to two separate weighted averages, the first based on two data sets and the second based on three.

The Log Results routines allow users to export a summary of the experimental results to a file that can be used in another software application (e.g. for statistical analysis). Different methods for sorting the data can be chosen and data from different subjects and recording sessions may be logged together into a single ASCII file. The user may sort the data for a subject or a set of subjects in order of carrier frequency, modulation frequency, stimulus number, or ear of presentation. Log files therefore provide a convenient method of sorting data for report generation and statistical analysis of the results.

5. Computational methods and theory

The MASTER techniques were developed according to clear principles concerning the generation of stimuli, data collection, artifact rejection, and response analysis. Violation of any of these principles would cause the system to work either inefficiently, or not at all. This section describes some details of the programming that embodies these principles. Real data are used to illustrate the different principles.

5.1. Synchronizing the AD and DA conversion

AD and DA conversion must be perfectly synchronized. The basic unit within the MASTER technique is the epoch. An epoch of data is created every time the input buffer is filled. If the input and output buffers are exactly the same duration and they are started together, then it follows that they will both end at the same time. During continuous recording the buffers will continue to start and end at the same times, thereby ensuring that each recorded epoch is evoked by exactly the same stimuli. Every epoch should therefore contain responses with identical phase and amplitude which can be added together, removed due to artifact rejection, re-ordered, or re-sampled without changing the responses.

The MASTER system performs synchronized data input and output using a subroutine based upon a VI that is provided with LabVIEW version 5.0 ('Simul AI/AO buffered Trigger (E-series MIO.vi)'). This VI synchronizes the start of signal acquisition and stimulus presentation using a digital pulse that is sensed by the PFI0 pin (programmable function input) which acts as a start trigger for both the DA and AD buffer conversion operations. The original VI required that the trigger be externally generated. Because we desired the program to start on its own, MASTER causes the board itself to generate a pulse on the GPCTR0_OUT pin (general purpose counter) which we have hardwired to PFI0. After the AD and DA buffers have been initialized and after the DA buffer has been filled with values which will produce the acoustic stimuli, simultaneous AD and DA conversion is initiated using the internally generated pulse. Although LabVIEW provides users with routines which use software timers (that rely on the CPU clock) for initiating the AD and DA operations, these were not used because they produce variable delays of more than 1 ms between the two buffers depending upon system performance demands. This modified VI permits MASTER to acquire and generate data simultaneously with no drifts in timing since both AD and DA operations are started by the same pulse and then controlled by the same hardware clock.

The MASTER software uses a circular buffering technique for both DA and AD buffers. This means that after the last point in the buffer is sent out or filled, the next value will be read or written to the first point in the buffer. During operation, while the output regenerates the same buffer of data, the input data are constantly changing, which requires that the data are transferred from the buffer to the computer before new data are written in their place. The 'Simul AI/AO buffered Trigger (E-series MIO.vi)' also contains a 'scan backlog' indicator which shows the number of data samples that remain in the buffer after the transfer subroutine has executed. If the backlog value increases steadily, then the data are not being read fast enough to keep up with the acquisition, and the newly acquired data may overwrite unread data, which would cause the program to halt and display an appropriate error message. This important feature can indicate that the computer speed is not fast enough or the amount of RAM insufficient for the protocol that the user has designed.

As each AD conversion is completed, the value is transferred into a first-in–first-out (FIFO) buffer. Because the FIFO can collect up to 512 values before any information is lost the computer can ignore the incoming information for a time equal to the product of the FIFO size and sampling rate. LabVIEW software routines manage data transfer between the FIFO on the board and the direct memory access (DMA) buffer which has been defined in the computer's memory. While data are transferred from the FIFO into a part of the DMA buffer, the CPU is removing data from another part of the DMA buffer and processing it (storing it on disk or averaging it with the previously recorded data). This setup obviates the need for the traditional use of a double buffer. However, the acquisition parameters and DMA buffer size must be chosen so that the LabVIEW routines are able to transfer data out of the DMA buffer fast enough to keep up with the rate at which the FIFO is transferring new data into the buffer. If not, unread data in the FIFO will be overwritten, and the user would be alerted by the MASTER program that an error has occurred. The DMA buffer must be large enough to inhibit the FIFO from overwriting itself, while not being so large as to affect the RAM needs of the CPU. The input DMA buffer has been set at 16 384 values which is 16 times the values that are read using inputs of 1024 points each, since this is the default number of points used in data acquisition in the MASTER program.

The timing of both the AD and DA conversions is determined by counters that wait through a specified number of ticks of the clock on the AT-MIO-16E-10 board. Since this is a 20-MHz clock, the time between successive AD and DA conversions (i.e. 1/conversion rate) must be an integer multiple of 50 ns. The timing will drift during the experimental session if this 50-ns requirement is not met, and will cause the buffers to become desynchronized. The program therefore permits AD rates of such as 500 or 1000 Hz but does not allow 3000 Hz (Fig. 2). The DA rate also has to be an integer multiple of 50 ns. For example, using 1000 Hz for the AD timing and a DA factor of 32 leads to a DA timing of 31250 ns (1/32 kHz), which is 625 clock ticks. If the DA factor were set to 30 rather than 32, the required timing for the DA conversion (30 kHz) is not an integer value of the clock ticks. The DA timing would be rounded off to an integer and the DA and AD conversion would become unsynchronized. If the user enters a value in these controls that produces illegal AD or DA values then a red warning indicator appears on the screen next to the illegal value(s).

5.2. Stimulus generation

The parameters entered in the Load Protocol screen (Fig. 2) define the stimuli. The processes by which these stimuli are generated rely on a set of formulae that are reviewed in the following paragraphs. More intensive discussions of these formulae are available in Hartmann [37] and Stanley [38]

Dividing the AD buffer size by the AD rate gives the epoch time (s), e.g. 1.024 s. The DA conversion time (t) is equal to the reciprocal of the DA rate. Prior to constructing the stimuli, the stimulus parameters provided by the operator are altered slightly based upon s to prevent acoustic artifacts. During data acquisition the values in the DA buffer are continuously converted into analog signals. The DA buffer operates in a circular fashion, which means that after the last value is output from the buffer, the conversion process begins again from the beginning of the buffer. In order for the transition from the last point in the buffer to the first point to occur without acoustic artifact an integer number of cycles for all modulation frequencies and carrier frequencies must occur within the length of the buffer. Therefore, the frequencies provided by the operator for both carrier and modulation signals are adjusted to give an integer number of cycles within the buffer using the equation:

$$f = (\mathrm{int}(sf'))/s$$

where s is the buffer duration in seconds, f is the actual frequency to be used, f' is the frequency input by the user and **int** is a function returning the closest integer. For example, a modulation frequency of 85 Hz is adjusted to 84.961 Hz when s is equal to 1.024 s. As well as enabling the DA sweeps to be linked together without acoustic artifact, this also ensures that the frequency of the response is exactly equal to one of the discrete frequencies measured by the FFT.

The other parameters defined in the Load Protocol screen are then also adjusted in order to simplify the calculations. First, the modulation indices m_a and m_f are calculated from the entered percentage modulation values entered for amplitude modulation m'_a and frequency modulation m'_f (rows 4 and 5, respectively, in Fig. 2):

$$m_a = m'_a/100$$

$$m_f = (m'_f/100)(f_c/(2f_m))$$

Note that the index for frequency modulation varies with the both the carrier frequency and the modulation frequency. This is necessary for the calculations, which depend on manipulating the phase of the carrier signal at the modulation frequency. The use of a percentage of the carrier frequency to define the amount of frequency modulation is not common and is open to ambiguity: is the numerator the total amount of frequency change from maximum to minimum or just the change from the carrier to the either the maximum or minimum frequency (Δf)? In keeping with previous papers in this field [2,31], we define the percentage frequency modulation as 100 times the difference between the maximum frequency and the minimum frequency divided by the carrier frequency. Thus, a 1000-Hz signal that is 20% frequency modulated will range between 900 and 1100 Hz.

The amplitude value is adjusted so that a value of 100 equals one-half of the range of the ± 10-V DA buffer:

$$a = 10'a/100$$

where a' is the amplitude input by the user. The LabVIEW DA conversion routines convert these voltages into the actual levels used in the converter, which has a 12-bit resolution (4098 levels or 72 dB) and a minimum level difference of 48 mV.

A frequency modulation term is calculated to adjust the phase and thus the frequency of the signal:

$$P = m_f \sin(2\pi f_m t i + \theta\pi/180)$$

where f_m is the modulation frequency, t is the time per address in the sweep in seconds, and θ is the phase difference between the frequency- and amplitude-modulations expressed in degrees.

The MASTER program provides three different modes of stimulation. Mode 0 calculates stimuli which maintain their root-mean-square (RMS) pressure constant across the different amounts of amplitude-modulation [39,40]. Mode 1 calculates stimuli with a constant peak intensity. Mode 2 (for calibration) uses just the modulating frequency. The stimulus (s) is thus calculated according to one of the following formulae:

Mode 0 (constant RMS): s(i)

$$= a(1 + m_a \sin(2\pi f_m ti) \sin(2\pi f_c ti + P)) \times /(1 + m_a^2/2)^{1/2}$$

Mode 1 (constant peak): s(i)

$$= a(1 + m_a \sin(2\pi f_m ti) \sin(2\pi f_c ti + P))/(1 + m_a)$$

Mode 2 (calibration): $s(i) = a(\sin(2\pi f_m ti))$

The MASTER technique combines all the individual stimuli for one ear into a single waveform. The traveling wave separates a complex stimulus into its component frequencies with the higher frequencies activating the more proximal portions of the basilar membrane. Since, the cochlea allocates the different carrier frequencies to different regions of the basilar membrane, each carrier frequency then evokes a separate response. The response occurs at the modulation-frequency specific to that particular carrier tone [7].

For each ear, the AT-MIO16-E10 board uses a 12-bit DA buffer which extends over a range of ± 10 V. Because the individual stimuli are summed together to obtain the final stimulus, the combined stimulus contains amplitudes which vary with the algebraic sum of the specific peaks and troughs of the individual stimuli. This combined stimulus will almost always contain amplitudes that are larger than the individual stimuli. The MASTER system defines the amplitude of its stimuli in terms of a 'percentage of buffer'. In order to keep the combined stimulus within the range of the DA converter the amplitudes of the individual stimuli should not in general exceed $100/N$, where N is the number of stimuli within one ear. When N is 4 the individual stimuli are then represented with only 10-bit accuracy, and the combined stimulus is represented with the full 12-bit accuracy. Individual stimuli can be represented with a resolution as low as 6 bits and still evoke reliable responses, but at lower resolutions the acoustic waveforms are significantly distorted.

In order to ensure that the amplitude of the combined stimulus does not exceed the maximum value that the buffer can represent, the program checks the maximum amplitude and produces a warning if the maximum is exceeded (see Fig. 4, right ear stimuli). The buffer check routines become non-trivial when using Mode 0, because although the user might define, for example, four stimuli, each of which has an amplitude that is 25% of the buffer, the amplitude correction factor (for constant RMS) may modify the signals so that they have considerably larger peak values. The summation of these new stimuli may then greatly exceed the buffer maximum, causing significant distortion when these values are clipped at the maximum value during DA conversion.

5.3. Choice of stimuli: AM and AM/FM stimuli

MASTER allows the user to present stimuli using either AM or FM, or both. The advantage of using both AM and FM has been explored by Cohen et al. [2] who demonstrated that larger responses (gains varying between 1.28 and 2.17 for different frequencies) could be obtained using both AM/FM stimuli over either AM or FM alone. In their study, Cohen et al. suggested aligning the maximum frequency of the FM with the maximum amplitude of AM. This choice is supported by studies which have shown that evoked potentials respond preferentially to transitions from slower to faster frequencies [41,42]. We have extended Cohen's finding from single stimulus recordings to the multiple stimulus case. Four stimuli were presented at 50 dB SPL to the left ear (0.5, 1, 2, and 4 kHz) and four others to the right ear (0.75, 1.5, 3, and 6 kHz). The stimuli were modulated at frequencies that varied from 85 to 95 in 2.5-Hz increments with the higher carrier frequencies in one ear having higher modulation frequencies. The stimuli were modulated by either AM only (modulation depth 100%) or combined AM and FM (modulation depths of 100 and 25%, respectively). The effect of recording time on the number of significant responses was investigated in eight subjects using 16.384-s sweeps and averaging 12 of these sweeps (just over 3 min of data). The difference between AM and AM/FM stimuli can be seen in the bottom of Fig. 6. The responses to the AM/FM stimuli were significantly larger in

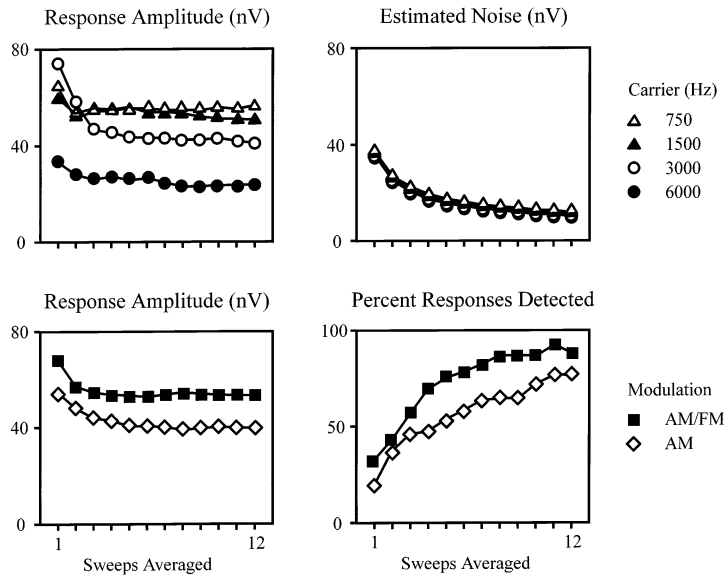

Fig. 6. Combined AM/FM stimuli. The upper half of the figure shows the change in the responses and the estimated noise levels as the number of sweeps averaged increases from one to 12. For these graphs, the results for the AM and AM/FM stimuli have been combined to show the effect of carrier frequency. For simplicity only the data from the right ear are used. The recorded amplitude drops slightly over the first few sweeps because unaveraged EEG noise contributes to the measurements in these sweeps. The response to the 6000-Hz carrier is smaller than the other responses. The upper right of the figure shows the effect of averaging on the estimated noise (background EEG). The noise decreases by the square-root of the number of trials. In order for the F-test to be significant, the response has to be about three times the estimated background noise. Comparing the upper two graphs shows that the response to the 6000-Hz carrier is not significantly different from noise even after 12 sweeps have been averaged. The lower left half of the figure shows the comparison between AM and AM/FM stimuli, with the data collapsed across carrier frequency. The lower right half of figure shows the percentage of responses that are judged significant as the averaging proceeds. The AM/FM stimuli clearly become significantly faster than the AM stimuli. The increase in amplitude caused by the AM/FM modulation results in between seven and eight responses rapidly becoming significant within the 3-min recording period. These data were collected from eight subjects.

amplitude (average gain of 1.35) than the responses to the AM stimuli (main effect of modulation type $F = 18.8$; df 1,7; $P < 0.01$; no significant interaction with carrier frequency). The time required for testing depends upon how rapidly a significant signal-to-noise ratio can be attained. Since the responses to AM/FM stimuli are significantly larger than the response to AM stimuli, the testing time can be shortened by using AM/FM stimuli. By the sixth sweep, 75% of the AM/FM responses (on average six of the eight) are significantly different from noise. The responses to the AM stimuli do not reach this level of significance until after 12 sweeps have been collected.

The two carrier frequencies that did not produce reliable responses in the 3-min period were 0.5 and 6 kHz. These responses generally display smaller amplitudes than for the other carrier frequencies [4,7] and often require twice that amount of time before they become significant. The different amplitudes of the responses for different carrier frequencies are shown for the right ear stimuli in the upper part of Fig. 6. There is only a slight reduction in the estimated noise (calculated as the mean amplitude in the 60 frequency bins above and the 60 frequency bins below the response frequency) with increasing carrier frequency (and increasing modulation frequency).

5.4. Detection of steady-state responses

The F-test enables the program to automatically provide the probability of whether a signal is present at the frequency of stimulation. This test derives historically from Fisher [43], who described a test used by Schuster [44] to investigate hidden periodicities in meteorological phenomena. This approach has recently been used to examine whether a response is present in recordings of the otoacoustic emissions [11]. The assumption is that the activity in the adjacent frequency-bins is random with a mean of zero and equal variance in the real and imaginary dimensions. The ratio of signal power (the sum of the squares of the response in the two orthogonal dimensions) to the sum of the powers in N adjacent bins is distributed as F with 2 and $2N-1$ degrees of freedom (Wei, [45], p. 259, equation 12.1.8 noting that n equals $2N+2$). Since we are not subtracting out the mean, we used 2 and $2N$ degrees of freedom (cf. Lins et al. [18] and Zurek [11]). The F-ratio may be computed using:

$$\frac{(\boldsymbol{x}_{\mathbf{s}}^2+\boldsymbol{y}_{\mathbf{s}}^2)/2}{\left(\sum_{i=1}^{J}\boldsymbol{x}_{\mathbf{i}}^2+\sum_{i=1}^{J}\boldsymbol{y}_{\mathbf{i}}^2\right)/(2\boldsymbol{J})}$$

in which $\boldsymbol{x}$ is the sine term of each polar vector and $\boldsymbol{y}$ is the cosine term, and $\boldsymbol{J}$ is the number of vectors which are used in the summation operation. We used $\boldsymbol{J}$ of 120 (60 measurements on either side of the signal). Because the sum of the squares of these sine and cosine terms equals the square of the amplitude, the equation reduces to:

$$120\boldsymbol{a}_{\mathbf{s}}^2/\sum_{i=1}^{120}\boldsymbol{a}_{\mathbf{n}i}^2$$

where $\boldsymbol{a}_{\mathbf{s}}$ is the amplitude at the signal being tested and $\boldsymbol{a}_{\mathbf{n}i}$ is the amplitude noise at each of the $\boldsymbol{i}$ adjacent frequencies.

When 120 bins of the FFT are used in the estimate of the noise, the ratio may be referenced to the F-ratio value for 2 and 239 degrees of freedom. The 120 values for the noise estimate are those which are adjacent to the modulating frequency (60 on either side), omitting those frequencies at which another response might occur. Valdes et al. [13] found that the F-test performs as well as other techniques such as Hotelling's T^2-test [9,10] or the Rayleigh test which only looks at phase [46].

The F-ratio procedures assumes that the noise is equally distributed across the frequencies. The noise would then be equally distributed across the FFT-bins. The amplitude of the EEG decreases (approximately exponentially) with increasing frequency (as can be seen in the amplitude spectrum in Fig. 2). This could lead to an over-estimation of the noise levels and a more conservative test. In order to assess this, we empirically checked the performance of the F-ratio on real data from the experiment described in the preceding section. Using one recording (12 sweeps) from each of eight subjects we analyzed 400 frequencies to obtain the number of false positives (those showing amplitudes larger than expected from the distribution of the adjacent 120 frequencies at $P<0.05$). The samples were taken by taking consecutive frequencies starting at 72.99 Hz and increasing to 97.96 Hz, in 0.061-Hz increments (and omitting the frequencies containing responses). To evaluate the responses at the selected frequencies, the value at each frequency was compared to the 120 adjacent frequencies (60 above and 60 below). The data were also examined in the same way that they are evaluated online during a testing session: the first sweep (16.384 s of the record) was submitted to an FFT analysis and the number of false positives were calculated, then the average of the first and second sweep was analyzed, and so on until all 12 sweeps of the record had been averaged and evaluated. Across all subjects, and for all the sweep sub-averages in the data set, the average number of false positives was 4.33% with a maximum of 6.25% and a minimum of 2.75%. The statistic is therefore slightly conservative when used with real data. There was no trend for the number of false positives to increase or decrease as the recording duration was increased. This means that the testing procedure was reliable even for very short recording sessions: false positives were not more common after only one sweep was analyzed than after 12 sweeps. Rapid recordings are therefore not prone to increased false positives.

The F-ratio technique was also evaluated using a Kolmogorov–Smirnov test which checks whether a sample of results comes from a particular distribution [47]. We used this test to compare the cumulative distribution of the probabilities from the F-test against the expected monotonic increase from 0 to 1.0. None of the cumulative distributions were significantly different ($0.14<P<0.89$) from the expected straight line.

A statistical problem not addressed in the present program derives from the number of statistical tests performed as the recording proceeds [44]. It is possible for the operator to stop the test at any time that a desired significance level is reached. If the operator makes this decision every time the current sweep average is analyzed, the significance level resulting from the F-test is incorrect because multiple tests rather than a single test have been performed. In our laboratory we avoid the problem of repeated measurements by always collecting a set amount of data and only evaluating the response significance at the end of the recording period. However, we are beginning to work on algorithms which will allow the test to be stopped once the required level of significance has been attained. These algorithms will take the factor of repeated measures into consideration.

5.5. Artifact rejection

When data windows or 'sweeps' of various lengths are submitted to the FFT routine, they produce a frequency resolution of $\mathbf{1/(Nt)}$, where $\boldsymbol{N}$ is the number of sample data points and $\boldsymbol{t}$ is the time between each sample. The longer the sweep the better the frequency resolution. For our analyses we have used sweeps varying in duration from about 12 to 65 s. When a subject produces high-amplitude artifacts from non-cerebral generators such as movement or muscle, it has become routine in the analysis of evoked potentials to exclude these data from the analysis. However, if an entire sweep was rejected due to an artifact that lasted only 300 ms, the testing procedure could become very long. If we build each sweep by linking together shorter epochs, artifact rejection does not consume too much time. Accordingly, when a data epochs contains artifact, it is not appended to the ongoing sweep, and the next artifact-free epoch is inserted instead.

The concatenation of temporally discontinuous epochs of data causes the recorded data to be nonstationary. However, this characteristic does not affect the spectral energy at the frequencies of stimulation since these frequencies are integer multiples of the epoch rate (the reciprocal of the epoch's duration). However, the spectral energy in the recorded data which was not integer multiples of the epoch rate will be dispersed by the nonstationarity into frequencies $\boldsymbol{Nf_e - f_x}$ where $\boldsymbol{f_x}$ is the frequency of interest, $\boldsymbol{f_e}$ is the epoch rate and $\boldsymbol{N}$ is integer. This is illustrated in the top half of Fig. 7.

In order to examine how significant these nonstationarities might be, we considered both modeled and real data using a 500-Hz AD rate, a buffer duration of 1024 ms and a total sweep duration of 16.384 s. Nonstationarities were modeled by reversing the order of the epochs in the averaged sweep to give a 'jumbled' sweep (causing a jump to occur between each epoch and its neighbor). The top two lines of Fig. 7 illustrate the effects of nonstationarity on a signal that is not an integer multiple of the epoch rate (0.977 Hz). The top line shows only a portion of the time data to illustrate the nonstationarities, one of which occurs at 1024 ms. The second line shows a portion of the spectra for the normal sweep and the jumbled sweep. We also modeled the effects when a signal at an integer multiple of the epoch rate occurs in random noise. The signal remains the same but the noise spectrum is changed (at all the frequencies that are not multiples of the epoch rate). However, there is no change in the estimate of the noise amplitude over a range of frequencies. In the case of the actual EEG data, there is a greater energy at low frequencies than at high. The nonstationarity could then cause a redistribution of the lower noise frequencies toward the higher. This could then increase the background noise levels estimated from the activity adjacent to the frequencies of the stimuli. In order to assess this we combined a synthetic response (at 9.8 Hz) with real EEG background noise (third line of the figure). There is indeed a small increase in the background activity near the response.

At the frequencies most commonly used for recording the auditory steady-state responses, the falloff in EEG amplitude with increasing frequency is small. The bottom row of Fig. 7 shows the amplitude spectrum of an original recording using four simultaneous stimuli at modulation rates of 80.1, 85.0, 89.8 and 94.7 Hz. The amplitudes of the four responses were 1.30, 0.92, 0.78 and 0.43 μV, with F-ratio probabilities less than

0.0001 for the first three responses and 0.0323 for the fourth response. While the spectrum of the jumbled sweep demonstrated minor changes in its pattern, the amplitudes at the frequencies of stimulation remain unchanged. Despite the change in the pattern of the spectrum, there were no changes in the F-ratio probabilities before the third decimal place for any of the eight subjects examined. Therefore, while there was a small redistribution of the noise energy, this is not significant for the procedures we are using to detect responses (at least in the case of these relatively high modulation frequencies). The incorporation of a windowing procedure for each epoch as it is concatenated into the sweep could reduce this effect. However, the effect seems too small to warrant the use of such windowing and the amplitude recalibration it would entail.

5.6. Calibration procedures

All data acquisition instruments must have calibration procedures so that the operator can check the accuracy and reliability of the measurements. In order to perform AD calibration, the user selects Mode 2 in the Recording Protocol screen. The program will then present 'stimuli' which are only the modulation frequencies of the stimuli that are normally used. During calibration one of the DA channels is connected to the input of a

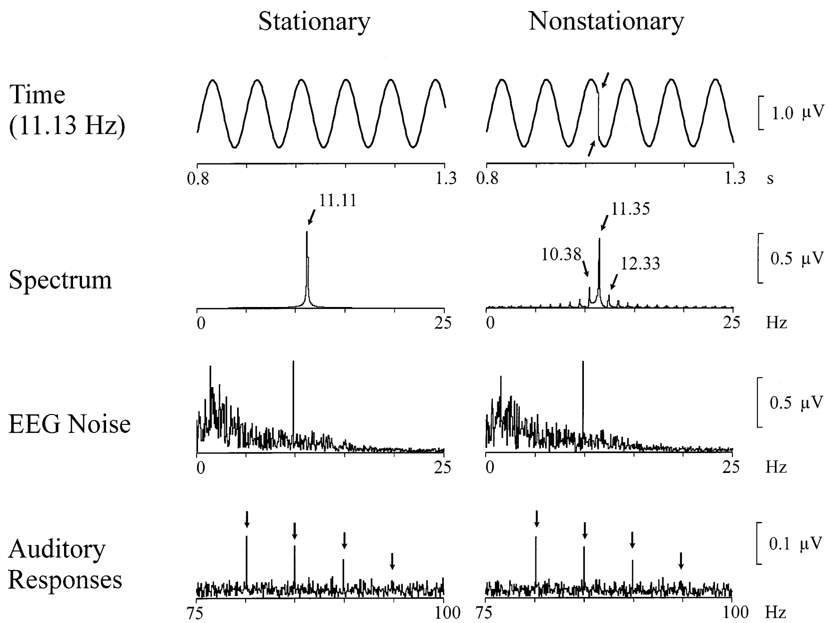

Fig. 7. Effects of nonstationarities in the data. The upper half of this figure shows the effects of inserting nonstationarities into a simple sine wave. The sine wave had a frequency of 11.13 Hz, which does not give an integer number of cycles in the recording epoch of 1.024 ms. The nonstationarities were induced by jumbling the order of the 16 epochs which composed a full sweep. There is then a sudden jump in phase at the transitions from one epoch to the next. An example is shown at the upper right at 1.024 ms into the sweep (two arrows). The spectrum of the jumbled signal shows energy at $Nf_e - f_x$ where N is integer, f_e is the epoch rate (0.977 Hz) and f_x is the frequency of the sine wave. The maximum peak occurs at 11.35 Hz which is when $N = 23$ (given the proviso that the measured frequency has to be an integer multiple of the frequency resolution for the whole sweep of 0.061 Hz). The third line of the figure shows what happens when EEG noise is jumbled. A modeled signal at exactly ten cycles per epoch (9.8 Hz) was mixed with the EEG. Jumbling the epochs caused no change in this signal. The nonstationarities change the spectrum of the EEG background. As well as random changes in the microstructure of the spectrum, there is some spread of energy from the lower frequencies to the high. The bottom line shows what happens when a recording of four simultaneous responses in the 75–100-Hz frequency-range is jumbled. There is no change in the responses (arrows). The nonstationarities change the microstructure of the background EEG spectrum. Because the spectrum is relatively flat in this frequency region, however, there is no significant change in the general estimate of the background noise across frequencies.

million-to-one attenuator of the connection box whose output is sent to the EEG pre-amplifier. This setup enables the user to quantify both the amplitude and phase effects that the electronic circuitry would produce on the recorded steady-state responses. If the signals are at an amplitude of 25%, the amplitude presented from the DA output is ± 2.5 V and the responses should show an amplitude of 2.5 μV in the amplitude spectrum display. The calibration factor can be adjusted until this is true. We usually start the calibration procedure with the calibration factor equal to 1.0 since this allows an easier estimation of the calibration factor. If the calibration procedure is performed using different modulation frequencies, the recorded activity will show changes across these frequencies related to the analog filters in the pre-amplifier. Both the amplitude and phase of the signal may change with frequency. At present we do not change the calibration factor with the frequency of the response, and we therefore recommend that the filter settings should not be too close to the modulation frequencies or show too high a slope. As well as checking the pre-amplifier, the calibration procedure also allows the user to ensure that the DA and AD buffers are synchronized over time, since no change in phase or amplitude of the spectral peaks at the modulation frequencies should occur during calibration. In practice, whenever a new paradigm is developed, it is generally a good idea to calibrate the system.

Calibration of the acoustic stimuli is also necessary. The analog signals are sent from the DA port through a ten-to-one attenuator in the connection box to the tape input of the audiometer. This instrument accepts an external input with a ± 1 V range, which it then amplifies, attenuates to the chosen decibel level and directs to the selected transducer (e.g. EarTone 3A ear-inserts or TDH headphones). Although the MASTER system enables four stimuli to be played in each ear, the DA calibration procedure should be performed on single stimuli. Using a sound pressure level meter (e.g. Brüel and Kjaer model 2230) and the appropriate coupler (e.g. a DB0138 2-cm^2 coupler in the case of insert earphones) the sound pressure level (SPL) of each single stimulus is measured. Across-frequency adjustments in intensity can be made by changing the amplification of the tape input (as monitored on the audiometer's Vu-meter). Frequency-specific adjustments can be made (with due consideration to the resolution of the 12-bit DA converter) by changing the amplitudes of the stimulus on the parameters screen.

6. Data storage

At the termination of a recording session the results are automatically saved to a disk file having a '.dat' filename extension. Additionally, the user is given the option of saving the raw data as well (using a '.cnt' filename extension to denote the continuous recording). The raw data for a 3-min recording at an AD rate of 1000 Hz requires about 1.6 Mb ($180 \times 1000 \times 9$ bytes). Files that store results of data processing such as averaging ('.avg'), subtraction ('.sub'), or summarization ('.log') are also labeled by a unique file name extension. Unless the user specifically chooses a binary format, all data files are written in ASCII format which can be easily viewed using a word-processor or imported into analysis programs such as MATLAB™ or Excel™. All files contain a header, which displays the information defined in the pa1 and pa2 files, so that the experimental protocol used during data collection can be immediately ascertained. Software for reviewing raw data files is provided with the MASTER system. This is exactly the same as the software for online recording except that the 'Record' button in the Main Control Menu enables the user to load data from disk rather than collecting it from the hardware buffers.

7. Performance issues

Since MASTER runs under a Windows operating environment, users are able to benefit from many Windows features. For example, whole screen images can be sent to a network printer or copied and pasted into a graphics program. Alternatively, a user can copy any of the specific components of a screen by clicking the right mouse button while the cursor is above the object and then choosing 'copy'. The graphics object can then be individually pasted into a graphics program. Several programs may be executed concurrently while MASTER is run. The user may therefore run customized analyses on the data files as soon as they are collected. Running extensive analyses during actual data collection is, however, not recommended due to memory demands. While communication with other programs is possible using the ActiveX™ or dynamic data exchange (DDE) functionality of LabVIEW™ this feature is not used at the present time by the MASTER program since raw data can be stored to disk and MASTER already provides extensive real time analysis of the data. Unlike other Windows-based programs, the screen displays for MASTER may not be easily resized. MASTER automatically runs on a full screen setting (800 × 600 pixels resolution) so that the only other screen component is usually the Windows Taskbar. While the current version of LabVIEW™ supports resizing for different screen resolutions, we are aware of complaints that this feature has shown problems on some systems. Accordingly, MASTER warns the user if the computer is set with incorrect display settings.

Our system runs easily on a Pentium 166 with 64 Mb RAM. Attempts to run the system on computers with only 32 Mb RAM have occasionally run into problems. Because of the inclusion of the 'Backlog' indicator on the main data collection screen, users are warned if the experimental protocol begins to tax the resources of the system.

8. Conclusion

The MASTER system is a Windows based system for the real-time collection and analysis of auditory steady-state responses. We have designed it to be flexible, intuitive, and user-friendly. The system will easily allow new researchers to study the auditory steady-state responses without having to set up and test their own programs. Further information, demonstration software, and a tutorial are available at http://www.rotman-baycrest.ont.ca/users/sasha_j/master.html.

Acknowledgements

This research was funded by the Medical Research Council of Canada. The authors would also like to thank James Knowles and the Baycrest Foundation for their support. Richard Mowrey, Adrian Kellett and Martin Rivers provided programming assistance during the development of earlier versions of the program. Hilmi Dajani, Peter Picton, and the National Instruments Technical Support Staff provided advice about LabVIEW programming.

References

[1] R.A. Dobie, Objective response detection, Ear Hear. 14 (1993) 31–35.
[2] L.T. Cohen, F.W. Rickards, G.M. Clark, A comparison of steady-state evoked potentials to modulated tones in awake and sleeping humans, J. Acoustic Soc. Am. 90 (1991) 2467–2479.
[3] O.G. Lins, T.W. Picton, Auditory steady-state responses to multiple simultaneous stimuli, Electroencephalogr. Clin. Neurophysiol. 96 (1995) 420–432.
[4] M.S. John, O.G. Lins, B.L. Boucher, T.W. Picton, Multiple auditory steady-state responses (MASTER): stimulus and recording parameters, Audiology 37 (1998) 59–82.
[5] R. Galambos, S. Makeig, P.J. Talmachoff, A 40-Hz auditory potential recorded from the human scalp, Proc. Natl. Acad. Sci. USA 78 (1981) 2643–2647.
[6] F.W. Rickards, G.M. Clark, Steady state evoked potentials to amplitude-modulated tones, in: R.H. Nodar, C. Barber (Eds.), Evoked Potentials II, Butterworth, Boston, 1984, pp. 163–168.
[7] O.G. Lins, P.E. Picton, T.W. Picton, S.C. Champagne, A. Durieux-Smith, Auditory steady-state responses to tones amplitude-modulated at 80-110 Hz, J. Acoustic Soc. Am. 97 (1995) 3051–3063.
[8] R.D. Linden, K.B. Campbell, G. Hamel, T.W. Picton, Human auditory steady state evoked potentials during sleep, Ear Hear. 6 (1985) 167–174.
[9] T.W. Picton, J. Vajsar, R. Rodriguez, K.B. Campbell, Reliability estimates for steady-state evoked potentials, Electroencephalogr. Clin. Neurophysiol. 68 (1987) 119–131.
[10] J.D. Victor, J. Mast, A new statistic for steady-state evoked potentials, Electroencephalogr. Clin. Neurophysiol. 78 (1991) 378–388 Published erratum appears in Electroencephalogr. Clin. Neurophysiol. 83 (1992) 270.
[11] P.M. Zurek, Detectability of transient and sinusoidal otoacoustic emissions, Ear Hear. 13 (5) (1992) 307–310.
[12] R.A. Dobie, M.J. Wilson, A comparison of t test, F test, and coherence methods of detecting steady-state auditory-

evoked potentials, distortion-product otoacoustic emissions, or other sinusoids, J. Acoustic Soc. Am. 100 (1996) 2236–2246.

[13] J.L. Valdes, M.A. Bobes, M.C. Perez-Abalo, M. Perera, J.A. Carballo, P. Valdes-Sosa, Comparison of statistical indicators for the automatic detection of 80 Hz auditory steady state responses, Ear Hear. 18 (1997) 420–429.

[14] D. Regan, J.R. Heron, Clinical investigation of lesions of the visual pathway: a new objective technique, J. Neurol. Neurosurg. Psychiatry 32 (1969) 479–483.

[15] D. Regan, J.R. Heron, R.F. Cartwright, Simultaneous analysis of potentials evoked by focal illumination of two retinal areas by flicker and pattern stimuli; application to migraine and to visual field investigations, Electroencephalogr. Clin. Neurophysiol. 29 (1970) 105–106.

[16] F.W. Rickards, L.E. Tan, L.T. Cohen, O.J. Wilson, J.H. Drew, G.M. Clark, Auditory steady-state evoked potential in newborns, Br. J. Audiol. 28 (1994) 327–337.

[17] G. Rance, F.W. Rickards, L.T. Cohen, S. De Vidi, G.M. Clark, The automated prediction of hearing thresholds in sleeping subjects using auditory steady-state evoked potentials, Ear Hear. 16 (1995) 499–507.

[18] O.G. Lins, T.W. Picton, B. Boucher, A. Durieux-Smith, S.C. Champagne, L.M. Moran, M.C. Perez-Abalo, V. Martin, G. Savio, Frequency-specific audiometry using steady-state responses, Ear Hear. 17 (1996) 81–96.

[19] G. Rance, R.C. Dowell, F.W. Rickards, D.E. Beer, G.M. Clark, Steady-state evoked potential and behavioral hearing threshold in a group of children with absent click-ABR, Ear Hear. 19 (1998) 48–61.

[20] T.W. Picton, A. Durieux-Smith, S.C. Champagne, J. Whittingham, L.M. Moran, C. Giguère, Y. Beauregard, Objective evaluation of aided thresholds using auditory steady-state responses, J. Am. Acad. Audiol. 9 (1998) 315–331.

[21] J.P. Poindessault, C. Beauquin, F. Gaillard, Stimulation, data acquisition, spikes detection and time/rate analysis with a graphical programming system: an application to vision studies, J. Neurosci. Methods 59 (1995) 225–235.

[22] G.W. Johnson, LabView Graphical Programming: Practical Applications in Instrumentation and Control, 2nd ed, McGraw-Hill, New York, 1997.

[23] D. Budai, L.J. Kehl, G.I. Poliac, G.L. Wilcox, An iconographic program for computer-controlled whole-cell voltage clamp experiments, J. Neurosci. Methods 48 (1993) 65–74.

[24] M.A Nordstrom, E.A. Mapletoft, T.S. Miles, Spike-train acquisition, analysis and real-time experimental control using a graphical programming language (LabView), J. Neurosci. Methods 62 (1995) 93–102.

[25] C. Davis, A. Mazzolini, D. Murphy, A new fibre optic sensor for respiratory monitoring, Australasian Phys. Eng. Sci. Med. 20 (1997) 214–219.

[26] H. Kunov, P.B. Madsen, Y. Sokolov, Single system combines five audiometric devices, Hearing Journal 50 (3) (1997) 32–34.

[27] T.F. Bardt, A.W. Unterberg, K.L. Kiening, G.H. Schneider, W.R. Lanksch, Multimodal cerebral monitoring in comatose head-injured patients, Acta Neurochirurg. 140 (1998) 357–365.

[28] E.J. Heyer, A. Wald, A. Mencke, Intraoperative data acquisition for the study of cerebral dysfunction following cardiopulmonary bypass, J. Clin. Monit. 11 (1995) 305–310.

[29] D. Regan, Human Brain Electrophysiology: Evoked Potentials and Evoked Magnetic Fields in Science and Medicine, Elsevier, Amsterdam, 1989, p. 35.

[30] S. Kuwada, R. Batra, V.L. Maher, Scalp potentials of normal and hearing-impaired subjects in response to sinusoidally amplitude-modulated tones, Hearing Res. 21 (1986) 179–192.

[31] T.W. Picton, C.R. Skinner, S.C. Champagne, A.J.C. Kellett, A.C. Maiste, Potentials evoked by the sinusoidal modulation of the amplitude or frequency of a tone, J. Acoustic Soc. Am. 82 (1987) 165–178.

[32] W.F. Dolphin, M.E. Chertoff, R. Burkard, Comparison of the envelope following response in the Mongolian gerbil using two-tone and sinusoidally amplitude-modulated tones, J. Acoustic Soc. Am. 96 (1994) 2225–2234.

[33] T.W. Picton, R.D. Linden, G. Hamel, J.T. Maru, Aspects of averaging, Semin. Hear. 4 (1983) 327–341.

[34] R.A. Dobie, M.J. Wilson, Objective response detection in the frequency domain, Electroencephalogr. Clin. Neurophysiol. 88 (1993) 516–524.

[35] M. Don, J.J. Eggermont, D.E. Brackmann, Reconstruction of the audiogram using brain stem responses and high- pass noise masking, Ann. Otol. Rhinol. Laryngol. Suppl. 57 (1979) 1–20.

[36] D.R. Stapells, T.W. Picton, A. Durieux-Smith, Electrophysiologic measures of frequency-specific auditory function, in: J.T. Jacobson (Ed.), Principles and Applications of Auditory Evoked Potentials, Allyn and Bacon, New York, 1993, pp. 251–283.

[37] W.M. Hartmann, Signals, Sound and Sensation. American Institute of Physics, Woodbury, New York, 1997, pp. 393–467.

[38] W.D. Stanley, Electronic Communication Systems, Reston Publishing, Reston, VA, 1982, pp. 141–213.

[39] N.F. Viemeister, Temporal modulation transfer functions based upon modulation thresholds, J. Acoustic Soc. Am. 66 (1979) 1364–1380.

[40] S. Sheft, W.A. Yost, Temporal integration in amplitude modulation detection, J. Acoustic Soc. Am. 88 (1990) 796–805.

[41] A. Maiste, T. Picton, Human auditory evoked potentials to frequency-modulated tones, Ear Hear. 10 (1989) 153–160.

[42] M. Kohn, K. Lifshitz, D. Litchfield, Averaged evoked potentials and frequency modulation, Electroencephalogr. Clin. Neurophysiol. 45 (1978) 236–243.

[43] R.A. Fisher, Tests of significance in harmonic analysis, Proc. R. Soc. (London) Ser. A 125 (1929) 54–59.

[44] A. Schuster, On the investigation of hidden periodicities with application to a supposed 26 day period of meteorological phenomena, Terrestrial Magnetism Atmospheric Electricity 3 (1898) 13–41.
[45] W.W.S. Wei, Estimation of the spectrum, in: Time Series Analysis: Univariate and Multivariate Methods, Addison-Wesley, Redwood City, CA, 1990, pp. 256–287.
[46] B. Lütkenhöner, Theoretical considerations on the detection of evoked responses by means of the Rayleigh test, Acta Otolaryngol. Suppl. 491 (1991) 52–60.
[47] W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, Numerical recipes, in: C: The Art of Scientific Computing, 2nd ed, Cambridge University Press, Cambridge, 1992, pp. 623–626.