Digital Ears User Manual

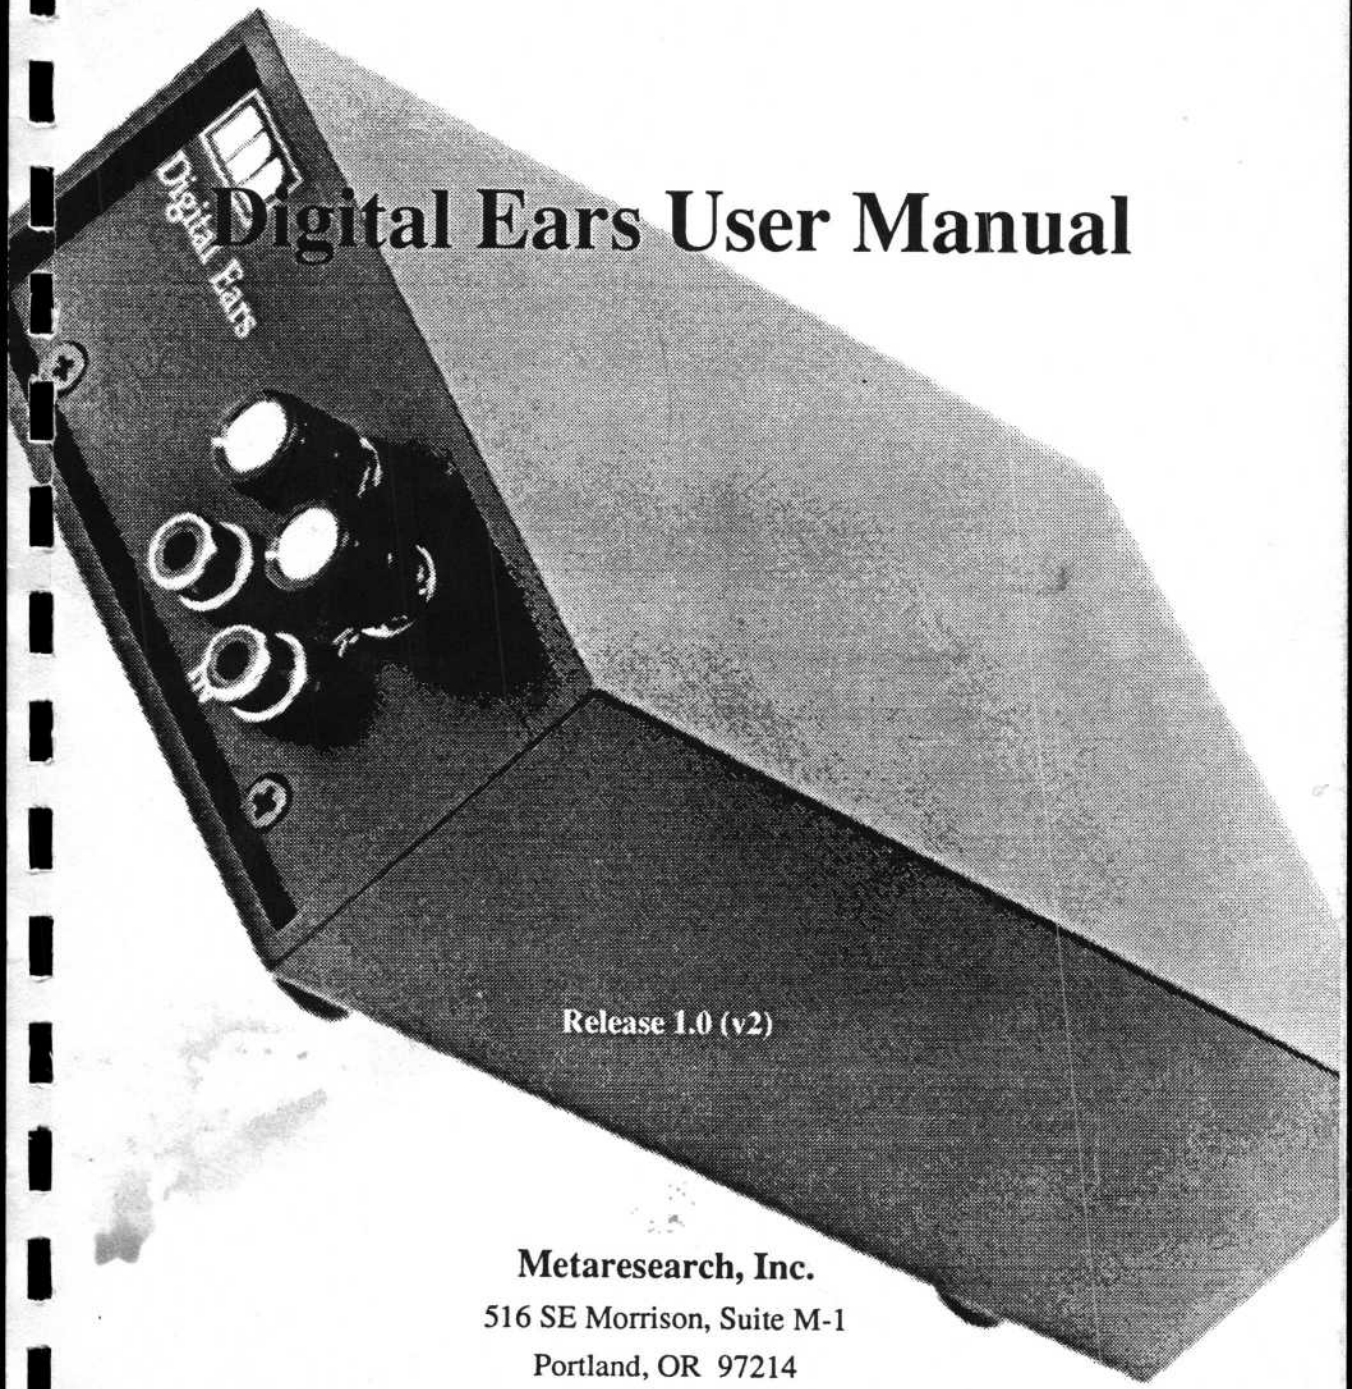

Release 1.0 (v2)

Metaresearch, Inc.
516 SE Morrison, Suite M-1
Portland, OR 97214
(503) 238-5728

Digital Ears
User Manual
Version 1.0

Metaresearch, Inc.
516 SE Morrison, Suite M-1
Portland, OR 97214
(503) 238-5728

Legal Notices

Copyright

Trademarks

Digital Ears, *DE-1*, *SoundWorks* and *DEController* are trademarks of Metaresearch, Inc. NeXT is a trademark of NeXT, Inc. The NeXT logo is a registered trademark of NeXT, Inc.

Limited Warranty on Metaresearch products

30-DAY MONEY BACK EXCHANGE POLICY ON METARESEARCH HARDWARE PRODUCTS:

If for any reason you are not satisfied with a Metaresearch hardware product, Metaresearch will refund you the purchase price, provided that you return that product, along with all accompanying hardware, software, documentation and other materials, to Metaresearch within 30 days of receiving the product. This offer does not apply to Metaresearch products that have been modified, which have been subject to unusual physical or electrical stress, or which have had their serial numbers removed or altered.

ALL IMPLIED WARRANTIES ON THE MEDIA AND MANUALS, INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, ARE LIMITED IN DURATION TO ONE YEAR FROM THE DATE OF THE ORIGINAL RETAIL PURCHASE OF THIS PRODUCT.

Even though Metaresearch has taken all possible steps to test the software and review the documentation, we make no warranty or representation, either express or implied, with respect to software, its quality, performance, merchantability, or fitness for a particular purpose. This software is sold "as is", and as the purchaser, you are assuming the entire risk as to its quality and performance. In no event will Metaresearch, Inc. be liable for direct, indirect, special, incidental, or consequential damages resulting from any defect in the software or its documentation, even if advised of the possibility of such damages. In particular, Metaresearch, Inc. shall have no liability for any programs or data stored in or used with Metaresearch products, including the costs of recovering such programs or data.

The warranty and remedies set forth above are exclusive and in lieu of all others, oral, or written, express or implied. No Metaresearch dealer, agent, or employee is authorized to make any modification, extension, or addition to this warranty. This warranty gives you specific legal rights, and you may have additional rights which may vary from state to state.

Warning

Never use the Digital Ears without its cover. Always turn the NeXT Computer off before connecting or disconnecting the Digital Ears.

Use of controls or adjustments, or operating the Digital Ears in any manner not specified in this manual voids the manufacturers' warranty.

Authors

Juliette Guilbert, Chris Lydgate, and Andy Laird.

Illustration and Design

Marianne Colgrove and Juliette Guilbert

Support

Kim Orumchian

Contents

NeXT software and the Digital Ears

Laboratory Data Acquisition

Appendices

Welcome

Welcome to Digital Ears. Digital Ears brings high-quality sound recording and high-resolution data acquisition to the NeXT Computer. With Digital Ears and the NeXT Computer, high-quality sound is about to assume a new and important role in personal computing.

Thank you for choosing Digital Ears. We are always open to suggestions for additions to both our hardware and software. If there is a capability you would like to have access to, please give us a call.

Sincerely,

Jean Delord
President
Metaresearch, Inc.

About this manual

Getting Started in a Hurry provides summarized installation and recording instructions for experienced NeXT users who want to begin to use their equipment immediately.

The **Introduction** explains what the Digital Ears is and how it works. It also provides a brief description of the software that is available for use with the Digital Ears.

Using the Digital Ears gives detailed, step-by-step instructions for connecting the Digital Ears to the NeXT Computer and setting up equipment for sound recording.

NeXT Software and the Digital Ears describes the utility programs for recording and manipulating soundfiles that are included with the NeXT system software. This section also discusses MonsterScope, an oscilloscope application bundled with the NeXT software.

Laboratory Data Acquisition explains how to use the Digital Ears as a laboratory-grade analog-to-digital converter. It also explains how to change the Digital Ears' settings to suit your particular requirements.

The **Appendices** include a discussion of some potential DE-1 applications, a glossary, and a description of soundfile formats. The appendices also state some technical specifications. For more detailed information, please contact Metaresearch.

Getting started in a hurry with SoundWorks

Always remember to turn the NeXT Computer off before connecting or disconnecting the Digital Ears.

1 **TURN THE NEXT COMPUTER OFF.**

2 Connect one end of the Digital Ears cable to the DSP port on the back of the Digital Ears, and the other end to the DSP port on the back of the NeXT cube.

3 Start up the NeXT Computer.

4 Connect a source device, such as a tape deck or a CD player, to the inputs of the Digital Ears.

5 Mount the Digital Ears 1.0 optical disk in the NeXT drive.

6 Copy the SoundWorks application folder to an Apps directory.

7 Double click on the SoundWorks icon to open it.

8 Set the input recording level:

- Turn on the VU meters by selecting *VU* on the button under the disabled meters.
- Turn the gain adjustment knobs on the front of the Digital Ears counterclockwise as far as they will go.
- Slowly turn the gain adjustment knobs clockwise until the bars of the VU meters start to light up.
- Play part of the sound passage you wish to record.
- Adjust the knobs until the VU meters read roughly between -2 and 0 dB.

9 Choose the *New* option from the File menu.

10 Hit the *Record* button. The NeXT will prepare for recording, and when everything is ready to go, the *Pause* button will light up.

11 Hit the *Pause* button to begin recording.

12 Hit the *Stop* button to interrupt recording. The *Stop* button also interrupts playback.

13 Press *Play* to play the currently active sound.

14 Select *Save, Save As... or Save To...* from the File menu to save the sound.

Introduction

What is the Digital Ears?

The Digital Ears is a high-speed, high-precision analog-to-digital converter with input filters. Its compact-disk quality input mirrors the NeXT Computer's compact-disk quality output, and fulfills the NeXT Computer's potential as a powerful environment for sound recording and data acquisition.

The Digital Ears takes line-level audio signals and converts them to digital information in true CD format (16-bit, 44.1 kHz stereo). It then transmits this information to the digital signal processor (DSP) on the NeXT Computer.

What does all of this mean?

The real world is largely analog - that is, most phenomena that we can observe are continuous. Consider the arc of an arrow shot from a bow. It traces out a smooth curve until it plunges to the earth - it does not descend jerkily. In the same way, when you record a violin solo

through a microphone, the voltage signal it generates consists of peaks and valleys, not of staircases.

Computers need digital information

But the world of computers is digital - no matter how sophisticated, a computer must represent signals in discrete numbers. This is why you can't just plug a receiver into a computer and expect to get any meaningful results. You must first find some way of converting continuous signals from the real world into discrete digital information that the computer can manipulate. Hence, the Digital Ears.

A waveform can be used to represent a sound's loudness and pitch:

Analog representation of signal

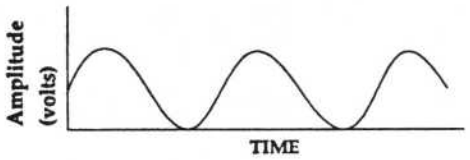

The sound's amplitude, or relative loudness, is shown by the height of the waveform. Its pitch is represented by frequency, or number of peaks per second. Frequency is expressed in Hertz: one cycle per second is a frequency of one Hertz. The lower the frequency, the lower the pitch of the sound.

The Digital Ears changes analog signals to digital signals

The Digital Ears takes an analog signal like the one represented by the waveform above - let us say the signal that represents a violin concerto - and samples that signal at 44.1 kHz per channel. This means that it determines the signal's amplitude at a certain instant, and produces a digital value, a number, that corresponds to that amplitude. Digital Ears repeats this process, which is called sampling, 44,100 times a second.

Digital representation of signal

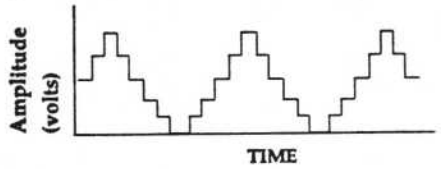

Once Digital Ears has converted the anaolog signal to a discrete binary number, the value is transmitted to the NeXT's DSP chip. Software that accesses the DSP can then process the data for analysis, display, or storage.

Each sample is a 16-bit binary number, which means that the number of possible values is 2^{16}, or 65536. This resolution is the same as that of a compact disk recording. This means that recordings made with the Digital Ears will be faithful reproductions of the original.

What software can be used with Digital Ears?

Bundled with DE-1

These applications are included with the Digital Ears under release 1.0.

- **SoundWorks:** a user-friendly digital sound processor, which can record and play sounds and prepare soundfiles for inclusion in other applications.
- **DEController:** a C object which facilitates the development of custom applications that use the Digital Ears. DEController includes methods for recording, playing, opening and saving sound/datafiles as well as methods for monitoring the DSP port for incoming data.

Bundled with the NeXT Computer

These applications are included with the NeXT system software.

- **MonsterScope:** a powerful three-channel 44.1 kHz oscilloscope and spectrum analyzer that lets you view data as it comes into the Digital Ears.
- **UNIX sound utilities:** the utility programs for manipulating soundfiles that are included with the NeXT system software. UNIX utilies are all documented in the on-line UNIX programmer's manual.
- **Sound Kit and Music Kit:** these kits provide software tools for playing, editing, synthesizing and displaying sound. Also, they include usable code for controlling the NeXT's DSP chip.

Available from third parties

MediaStation and FrameMaker are two examples of third party applications that use Digital Ears; there are many more in development.

- **Media Station:** a software package by Imagine, Inc. which provides the means to combine and process images, sounds and text to create multimedia applications and presentations.

- **FrameMaker:** an application from Frame Technology Corporation which combines sophisticated word processing, graphics and layout capabilities with the ability to incorporate sound into documents.

Anatomy of the Digital Ears

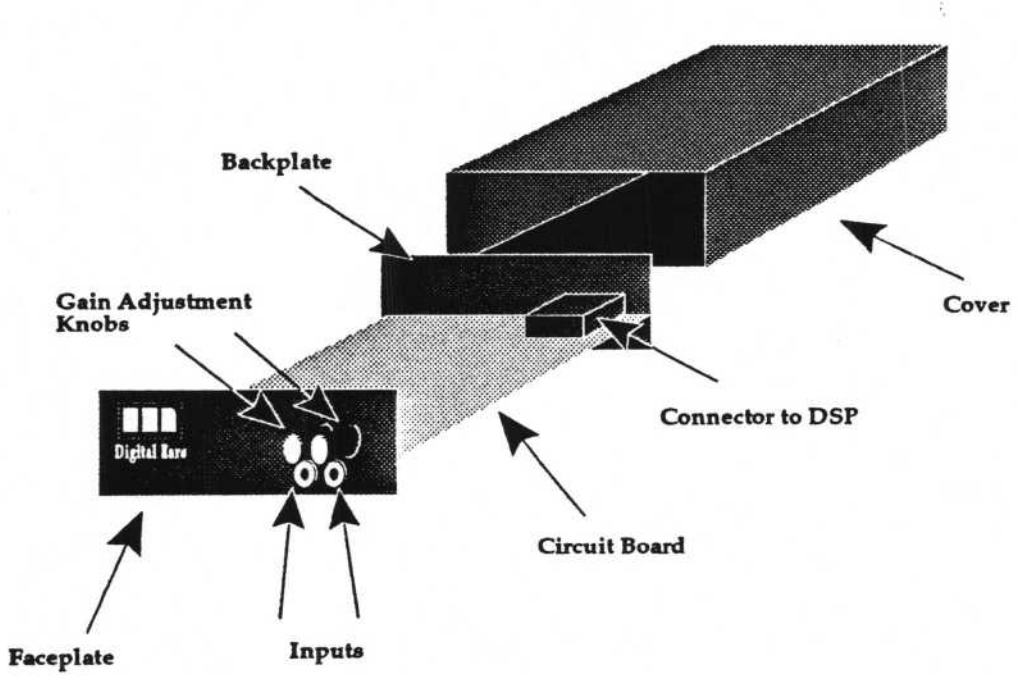

Figure 1. Inside the Digital Ears

This exploded view shows the basic components of the Digital Ears.

Using the Digital Ears

What you need

- A NeXT Computer with a hard disk and at least 8 megabytes of RAM.
- A Digital Ears with connector cable (included).
- A line-level audio source and a cable to connect it to the Digital Ears.
- NeXT system software release 1.0.

We also recommend the supplemental hardware listed below. While not essential, it will help you get the most out of the Digital Ears.

- A receiver, or pre-amp/power amp combination for sound playback.
- A pair of high quality speakers.
- Several spare RCA male-to-male patch cables.
- Additional hard disks.

A printed copy of the NeXT technical documentation, especially the sections concerning the Sound class, recording, Interface Builder and the SoundKit, can be very helpful.

Precautions

- **ALWAYS** turn the NeXT Computer off before connecting or disconnecting the Digital Ears.
- **NEVER** remove the cover from the Digital Ears while it is connected to the NeXT Computer.
- **NEVER** enter speaker (amplified) output from an audio power amplifier into the Digital Ears. The Digital Ears is designed to take a standard line-level input only.

Installing the Digital Ears

1 **Turn the NeXT Computer off.** (THIS IS IMPORTANT!)

2 Connect one end of the Digital Ears cable to the DSP port on the back of the Digital Ears.

3 Connect the other end of the cable to the DSP port on the back of the NeXT cube.

4 Start up the NeXT Computer.

The Digital Ears is now active and ready for use.

Setting up your equipment

Your Digital Ears set-up will depend on what you want to do with it. If you want to use the Digital Ears as a laboratory data-acquisition device, please refer to the **Laboratory Data Acquisition** section.

There are several different ways to install and use the Digital Ears. If you're recording sounds, you'll probably want to listen to them at the same time. The set-up shown in Figure 2 (see next page) lets you record from several different sources while listening to what you record; the set-up shown in Figure 3 is intended to illustrate one of the potential uses for the Digital Ears: real time voice recognition.

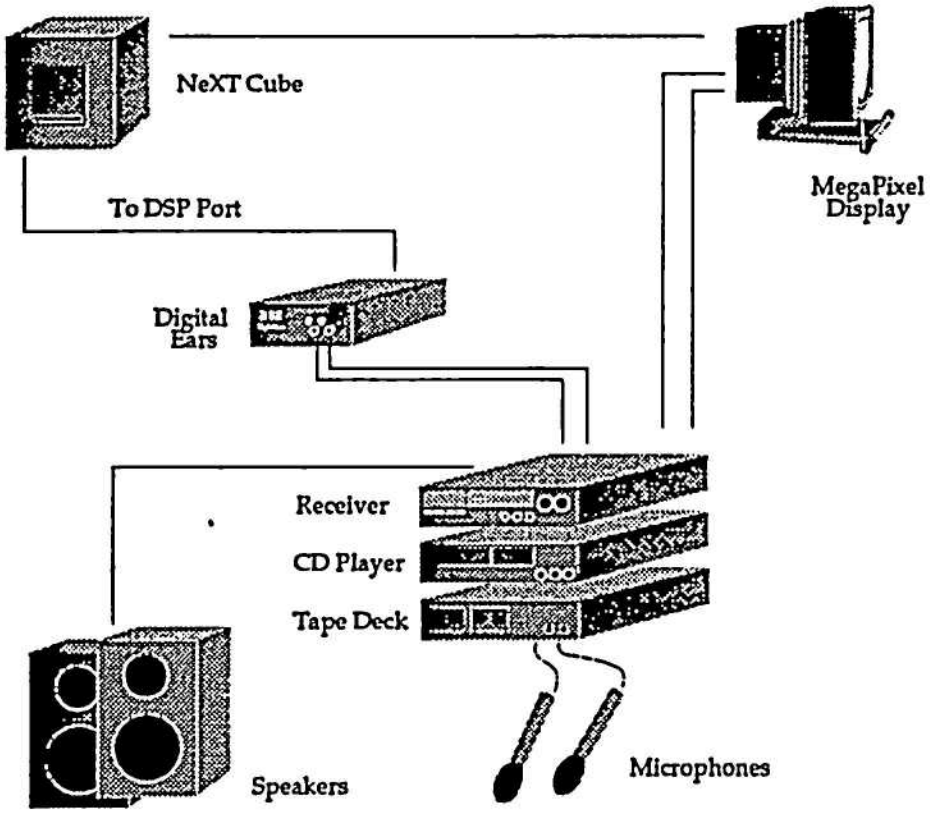

Figure 2. Sound configuration.

This configuration lets you listen to sounds while you record them.

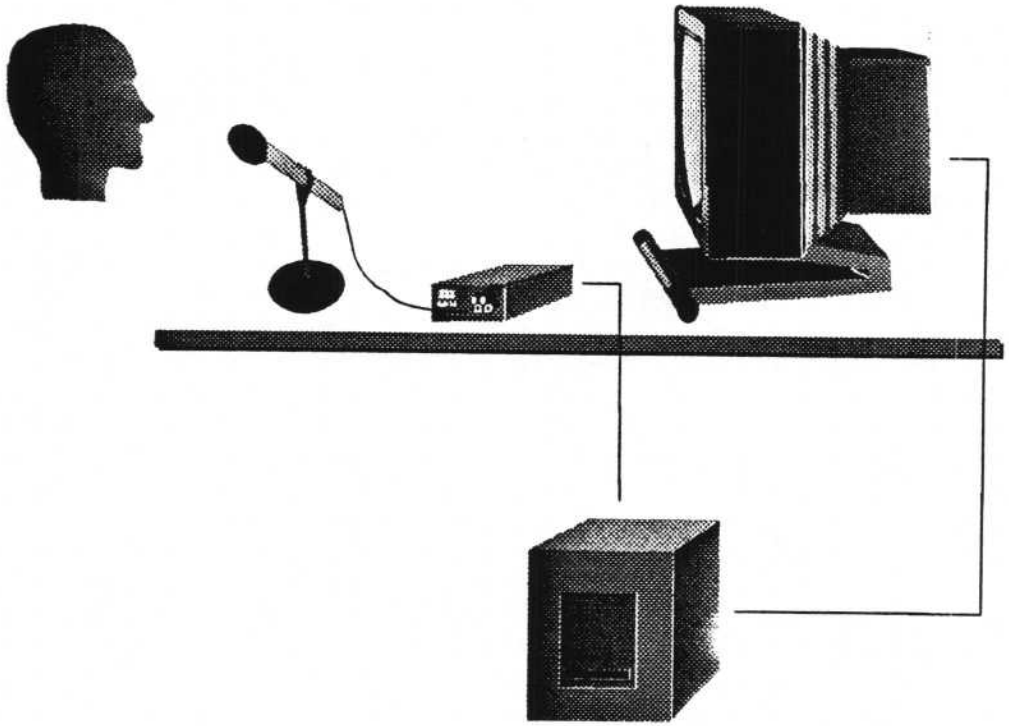

Figure 3. Voice recognition

This configuration illustrates one of the Digital Ears' potential application areas.

Most receivers have a tape monitor button; by connecting the NeXT Computer to the Tape In (PLAY) jack and the Digital Ears to the Tape Out (RECORD) jack on your receiver, you can monitor levels the same way you would with a tape deck. The only drawback to this approach is that there is some degradation of signal quality as the signal is routed through the various components and cables.

To set up your equipment for sound recording and playback:

See Figure 2 for an illustration of this set-up.

1 Connect the Digital Ears to the NeXT (see **Installing the Digital Ears**, above).

2 Connect the **Tape Out** (Record) jack on the back of the receiver to the **Input** jacks of the Digital Ears.

3 Connect the **Tape In** (Play) jacks on the back of the receiver to the **Line Out** jacks on the back of the Megapixel display.

4 Connect the sources' outputs to their corresponding inputs on the receiver.

5 Make sure that the speakers are connected to the receiver.

6 Play some music on the system to test the connections.

You can set input levels for recording using either SoundWorks or MonsterScope.

Setting input levels using SoundWorks:

1 Make sure the Digital Ears is connected to the NeXT Computer.

2 Start up SoundWorks.

3 Turn the level adjustment knobs on the front of the Digital Ears counterclockwise as far as they will go.

4 Slowly turn the level adjustment knobs clockwise until the bars of the VU meters start to light up. If the meters do not respond, the most likely causes are:

- You haven't turned the knobs far enough.
- The Digital Ears is not properly connected to the NeXT.
- The source device is not properly connected to the Digital Ears.
- The source device is not sending any signal.
- The signal is too weak.

5 Play part of the sound passage that you wish to record. For best results, play the loudest part of the passage.

6 Adjust the knobs until the VU meters read roughly between -2 and 0 dB.

The system is now ready to record.

Things to remember

Whenever you use the Digital Ears, please remember the following facts:

Storage requirements

High-quality sound takes up a lot of storage space on disk. The following chart shows how much storage space is required for recordings of various lengths made at various sampling rates:

Length of Recording	*Sampling Rate in kHz*		
	CODEC	22.05 stereo	44.1 stereo/ 88.2 mono
1 sec	8 kbytes	88.2 kbytes	176.4 kbytes
5 sec	40 kbytes	441 kbytes	882 kbytes
30 sec	240 kbytes	2.65 mbytes	5.29 mbytes
1 min	480 kbytes	5.29 mbytes	10.58 mbytes
5 min	2.4 mbytes	26.46 mbytes	52.92 mbytes

When you're recording, be sure you have plenty of free space on your disk. **We recommend that you have at least twice as much space as your soundfile will require.** Disks that are too full can develop unpleasant side effects.

Input levels

Before recording with the Digital Ears, you must first adjust the input levels. This is a delicate but essential operation, since the quality of the recording depends on the audio input level. (See **Setting input levels using SoundWorks** (p.11), or **Setting input levels with Monsterscope**, in *NeXT Software and the Digital Ears*.)

Audio sources

The Digital Ears will accept any source of music or sound, provided that its output signal is within a range of ±8 volts. Signals within this range are called "line-level" signals. Signals that are too strong, such as signals that have been amplified by a power amplifier, will be subject to distortion and could ultimately corrupt the system. Signals that are too weak, such as those produced by some microphones, must be passed through a pre-amplifier before they are fed into the Digital Ears.

Optimal signal size for input to the Digital Ears is in the ±0.5 to ±5 volt range. Signals below or beyond these limits will be recorded, but the signal-to-noise ratio may be low.

NeXT software and the Digital Ears

Overview

NeXT system software Release 1.0 includes several tools for working with sound. These are: UNIX sound utility programs for manipulating soundfiles; MonsterScope, an oscilloscope/spectrum analyzer; and the Sound and Music Kits, which provide software tools for playing, editing, synthesizing and displaying sounds. UNIX utilities and MonsterScope are discussed in this section; for information regarding Sound Kit and Music Kit, please refer to the on-line NeXT documentation.

UNIX sound utilities

If you aren't familiar with UNIX, you may want to skip this section.

To use any of these utilities, you should first start up a UNIX shell window, using either the Shell or Terminal applications included in the NeXT system software. Remember that these utilities are part of NeXT's software release, and they are all documented in the on-line UNIX programmer's manual. For more detailed information, double click on the UNIX Manual icon in the Digital Library window, and type the name of the command you want to learn about.

Recording sounds with *sndrecord*

sndrecord can make recordings from the Digital Ears or from a microphone plugged into the MegaPixel display. It has a number of powerful options for developers; if you just want to make sound recordings, however, you may find SoundWorks easier to use.

If it records from the microphone, sndrecord sets the sampling rate at 8012 Hz, and stores the resulting data in 8-bit mu-law format, mono.

If it records from the DSP port (i.e. the Digital Ears), sndrecord sets the sampling rate at 44.1 kHz, and stores the resulting data in 16-bit linear format, stereo.

sndrecord syntax

```
sndrecord [-options] filename.snd  <return>
```

The options are summarized as follows:

-d Record from the DSP port. This is the correct option for recording from the Digital Ears. The recording will be stored in output file **filename.snd.**

-f DSP-PROGRAM Run the specified dsp-program on the DSP, and store the resulting output in the output file.

-w Warn if sndrecord detects an output overflow.

-s SIZE Record SIZE bytes. Unfortunately, this is the only way you can specify the length of a recording before the recording actually starts. If you're reading data from the microphone, a 1-second recording will require 8012 bytes. If you're reading data from the Digital Ears, a 1-second recording will require 176,400 bytes. Unlike the others, this option starts recording as soon as you enter the command. (see below)

After you enter the sndrecord command, a message should appear telling you to hit a carriage return to start recording, followed by a second carriage return to stop recording. If you use the **-s** option you will not get this message - sndrecord will start recording as soon as you enter the command, and stop when it has recorded the number of bytes you specified.

When the recording is finished, sndrecord will report the number of bytes recorded, and your prompt will reappear. Once a soundfile has been created, it can be copied, moved, or modified like any UNIX file.

Remember to end your output files with ".snd" so that the NeXT will recognize them as soundfiles.

sndrecord examples

`sndrecord test.snd`
Record from microphone. Length of recording depends on when you type second carriage return.

`sndrecord -d test.snd:`
Record from Digital Ears. Length of recording depends on when you type second carriage return.

`sndrecord -d -s 882000 test.snd:`
Record 882,000 bytes (5 seconds) from Digital Ears as soon as the command is entered.

`sndrecord -s 80120 test.snd:`
Record 80,120 bytes (10 seconds) from microphone as soon as the command is entered.

If you encounter problems, the most likely cause is that another application which uses the DSP port is currently running (such as MonsterScope or SoundWorks). If one of these programs is running, quit from it, and try sndrecord again.

sndplay

sndplay simply plays the soundfiles you give as arguments. The specified files must be in standard NeXT soundfile format. If you give more than one soundfile as argument, sndplay will play the soundfiles in sequence.

sndplay syntax

`sndplay [file1.snd file2.snd]` `<return>`

sndplay examples

`sndplay splash.snd:`
Play soundfile splash.snd.

`sndplay splish.snd splash.snd splosh.snd:`
Play soundfiles in sequence.

sndinfo

sndinfo will read a soundfile's header, and report the following attributes:

- size in bytes
- format: 8-bit or 16-bit, mu-Law or Linear;
- sampling rate: 8012Hz, 22.5kHz, or 44.1kHz;
- number of channels;
- any comments stored in the header.

`sndinfo splash.snd`
Report information about splash.snd.

UNIX shell command release 1.0 user notes

- Do not use sndrecord as a background process.
- Do not use sndrecord when MonsterScope, SoundWorks, or any other application has control of the DSP port.

MonsterScope

Overview

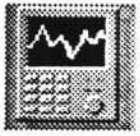

Figure 4. MonsterScope's icon.

MonsterScope is a combination oscilloscope/spectrum analyzer application which supports three input channels: DSP A, DSP B, and CODEC microphone. It may be used to view both waveforms and frequency spectra of data as it enters the NeXT DSP port. It can also be used as an alternative to SoundWorks for setting input levels. MonsterScope is bundled with the NeXT system software.

Running MonsterScope

To run MonsterScope you will need:

- The NeXT Computer booted from a SCSI disk.
- A Digital Ears connected to the NeXT Computer.
- MonsterScope, an application which should be located in NeXT Developer/Demos directory.

Find MonsterScope with the Workspace Manager. Double click on MonsterScope's icon to launch the application. In a few moments, the Oscilloscope window should appear on the screen.

There are two buttons in the upper left-hand corner of the Oscilloscope window labeled DSP A and DSP B. These buttons activate the left and right channels, respectively, of the DSP port.

Beneath the DSP buttons are several groups of buttons. The Inputs group contains buttons which control the AC/DC or GND (ground) mode of MonsterScope. You should set both channels to AC mode.

This initial configuration is shown in Figure 5 on the following page.

Setting Input Levels with MonsterScope

For a more detailed discussion of the issues involved in setting input levels, please refer to **Setting input levels,** in *Using the Digital Ears.*

To set the input recording levels for the Digital Ears using MonsterScope:

1 Make sure the Digital Ears is connected to the NeXT Computer.

2 Start up MonsterScope.

3 Turn the gain adjustment knobs on the front of the Digital Ears counterclockwise as far as they will go.

4 Connect a source device, such as a tape deck or a CD player, to the inputs of the Digital Ears, and make sure it is sending a signal to the Digital Ears.

5 Press the DSP A button to activate the left channel, and make sure the AC button is on. Slowly turn the gain adjustment knob clockwise until you see a signal emerge. If you do not see a signal appear, the most likely causes are:

- You haven't turned the knob far enough.
- The Digital Ears is not properly connected to the NeXT.
- The source device is not properly connected to the Digital Ears.
- The source device is not sending any signal.
- The signal is too weak.

6 As the signal gets stronger, you should see peaks of the signal become flattened at a fixed level – this is clipping (see Figure 6). Don't worry if you cannot make the signal large enough to observe clipping, but in order to make a recording, you will need some signal.

7 Deactivate the DSP A button and repeat steps 5 and 6 for DSP B.

8 To ensure balance between left and right channels, look at both signals simultaneously on MonsterScope's display. Press the "Identify" button

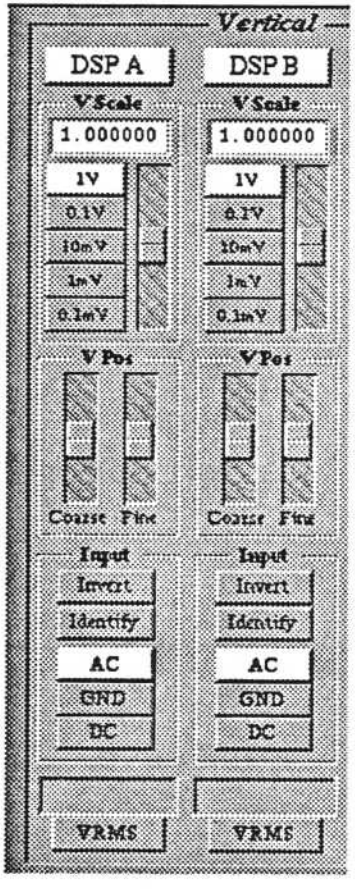

Figure 5. MonsterScope settings.

These are the correct settings for using MonsterScope with the Digital Ears.

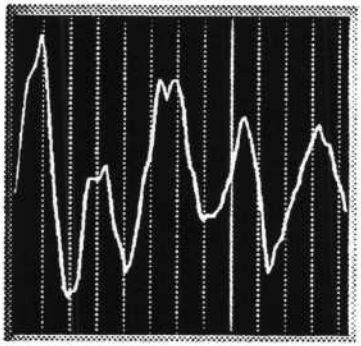

Figure 6a. Normal signal.

Signal is completely visible within scope display.

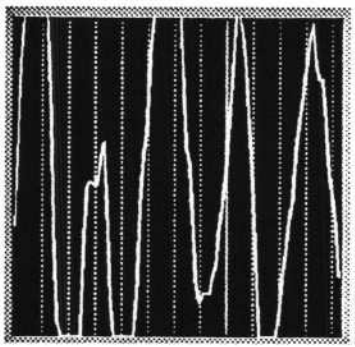

Figure 6b. Clipped signal.

Signal is clipped when Digital Ears' input levels are set too high.

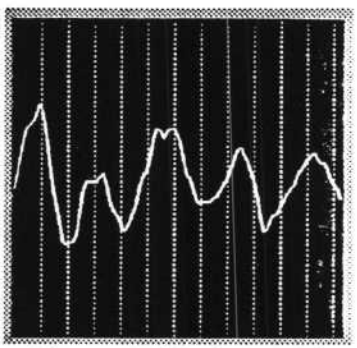

Figure 6c. Ideal signal for sound recording.

Signal is contained within the middle third of the scope display.

for one channel to distinguish it from the other. Are the signals the same size? You may wish to make some final adjustments to set the balance. If balance changes are required, it is usually best to decrease the larger signal to match the smaller one.

MonsterScope 1.0 User Notes

- The controls are slower to respond when the DSP channels are active. If you are changing settings it is best to set do it with both DSP A and DSP B inactive.
- If you want to use SoundWorks, or any other program that uses the DSP, make sure to turn off the DSPA and DSPB control buttons first.

Laboratory Data Acquisition

Overview

Laboratory data acquisition differs from sound recording in several ways. First, laboratory users usually want to take DC measurements, which involves changing the Digital Ears' factory settings. Second, soundfiles on the NeXT Computer have a special format that includes a header. When saving laboratory data, this header is unnecessary but can be useful if adapted for lab data.

Setting up your equipment

1 Disconnect the Digital Ears from the NeXT Computer and change the factory settings as needed. (See **Changing the Factory Settings** below)

2 Connect the Digital Ears to the NeXT Computer.

3 Connect the signal sources to the RCA inputs on the Digital Ears.

You are ready to begin recording data. Make sure the signals to be measured are between ±8 volts – we recommend you use MonsterScope to look at the data coming in to the Digital Ears to make sure that levels are appropriate. Remember that you cannot look at data using MonsterScope and SoundWorks at the same time, because both applications require access to the DSP driver.

Recording laboratory data

As in recording sound, there are several options for recording laboratory data. You can use SoundWorks or the UNIX sound utility sndrecord to create sound and/or data files.

Using SoundWorks to record laboratory data

1 Start up SoundWorks by clicking on its icon in the Workspace Manager.

2 Click the Record button to begin recording.

3 Remember that if you have changed the Digital Ears factory setting to mono, rather than stereo, the Digital Ears will sample the left channel at 88.2 kHz. Since stereo format is Left-Right-Left-Right, however, this presents no problem – mono input will be stored Left-Left-Left-Left.

Take care when using the Digital Ears in mono mode.

Using *sndrecord* to record laboratory data

1 Start up the Shell or Terminal applications to get a UNIX window.

2 Enter the sndrecord command with the appropriate flags (see **Recording with sndrecord** in the *UNIX Sound Utilities* section).

Examples of sndrecord:

A useful convention is to end the names of all data files with a ".dat".

`sndrecord -d test.dat:`
Record from Digital Ears. Length of recording depends on when you type second carriage return.

`sndrecord -d -s 882000 test.dat:`
Record 882,000 bytes (5 seconds) from Digital Ears as soon as the command is entered.

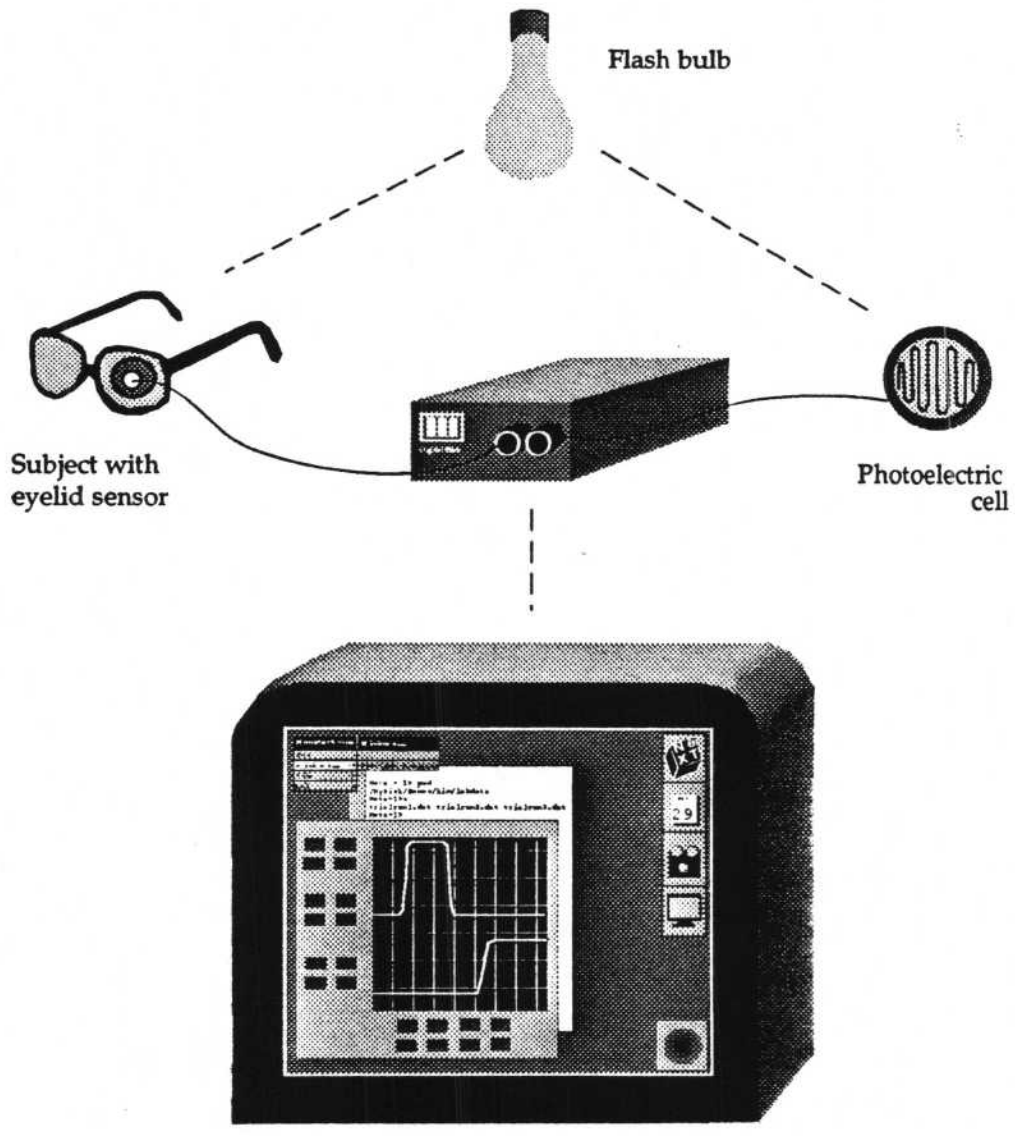

Figure 7. Sample lab application

3 Strip off the soundfile header if you do not require it. We suggest you consider retaining the header, however, since it preserves information about the file's format, and thus prevents the proliferation of enigmatic binary data files of unkown format. See *Appendix C*, **Reading the Soundfile Header.**

Changing the factory settings

Change the factory settings for laboratory data acquisition or special situations.

The Digital Ears' factory settings should be correct for most sound recording needs, but for laboratory data acquisition or for special situations, you may want to change them.

To change these settings you will have to manually reset jumpers on the Digital Ears' circuit board.

Always disconnect the Digital Ears from the NeXT Computer and from any inputs before removing the cover.

To change a jumper setting:

1 Turn off the NeXT if it is linked to the Digital Ears.

2 Disconnect any inputs to the Digital Ears.

3 Remove the two silver Philips-head screws on the bottom of the Digital Ears case.

4 Slide out the circuit board by pushing on the front faceplate and place it on a hard flat surface.

5 Locate the appropriate jumper in Figure 5, and set it according to the table in Figure 6.

Default settings

The Digital Ears' default factory settings are:

- AC mode, both channels
- Stereo mode
- Input filter frequency cut-off at 20 kHz, both channels

AC / DC mode selection

The Digital Ears factory setting is AC mode in both channels. To change AC/DC mode, set jumpers JP1 and JP2 according to the table in Figure 12.

Stereo / mono mode selection

The Digital Ears factory setting is stereo. To change stereo/mono mode, set JP3 according to the table in Figure 12. When the Digital Ears is in mono mode, the left channel becomes the mono input, and the sampling rate becomes 88.2 kHz.

If you change this mode, you will probably want to change the input filter frequency cutoff mode (see next page).

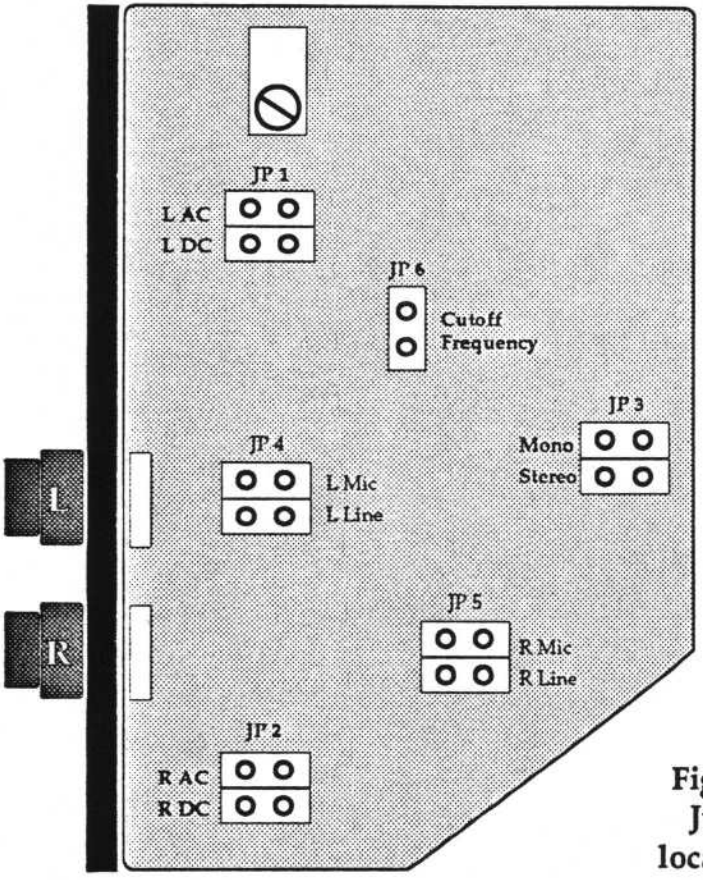

Figure 5. Jumper locations.

Input filter frequency cutoff mode selection

The Digital Ears factory setting is 20 kHz, both channels. The Digital Ears has low-pass input filtering to protect users against playback distortion. Even in situations where there is no audible need for filtering, high frequency components of signals can produce a kind of distortion known as aliasing. Since the Digital Ears can run in mono or stereo, there is a provision for changing this cutoff frequency of the input filter. The two choices are a cutoff of 20 kHz for stereo, and a cutoff of 35 kHz for mono.

To change the cutoff frequency, set JP6 according to the table in figure 12. Leaving JP6 open means the cutoff is 20 kHz. Closing JP6 changes the cutoff to 35 kHz.

As a technical matter, it is always advisable to ensure that your signal has no strong frequency components near the sampling rate. These inaudible components are the cause of aliasing — the result being glitches or chirps in your playback. One source of aliasing distortion is FM stereo tuner output, which may have a 38 kHz stereo "pilot" tone you cannot hear.

If you have signal conditioning equipment, you may want to filter your signals before sending them to the Digital Ears.

FRONT FACEPLATE

Jumper	Setting	Result
JP1 (Left) or JP2 (Right)		AC mode
		DC mode
JP3	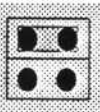	Mono
		Stereo
JP6		Frequency filter=35 kHz
		Frequency filter=20 kHz

Figure 6. Jumper settings.

Appendices

Appendix A: What can the Digital Ears do?

The Digital Ears is designed to function in two distinct capacities: as an audio input device, and as a laboratory-grade A/D converter.

The Digital Ears can serve as a compact disk quality audio input device.

Together with the SoundWorks™ software, the Digital Ears can save high-quality sound recordings on the NeXT Computer. Also, since soundfiles created by the Digital Ears software can be incorporated into a user interface using Interface Builder™, the Digital Ears can be used as a platform for software that incorporates high-quality sound input. Possible examples include:

- **Digital Recording Studios:** Since the Digital Ears can take line-level audio signals from CD players, mixing consoles, tuners, and tape decks, it can be used for tasks involving live recording, sound editing, digital mixing and digital mastering.

- **Audio help and error messages:** Instead of having traditional text files for on-line help, programs could include soundfiles that lead a user through a sequence of actions with recorded voices, music and sound effects. Error beeps could become error messages spoken by the computer.

- **Sound as a conceptual aid:** Sounds corresponding to specific user responses could reinforce or discourage certain actions. A CAD program could groan when a structural member is placed in the wrong position; a circuit board layout program could play a different sound whenever a different component is placed on the board.

- **Sound as a presentation tool:** Presentation technology holds great promise for education as well as business, and high quality sound is an important part of any presentation. A software demonstration could include a recorded narrative, complete with background music, to explain a slide-show. With this approach, teachers and managers can make more effective demos more quickly than before.

- **Music Education:** There are many possibilities for software that can record, play, display, edit, and filter sound. A music laboratory could compare the structure, harmonic content, and other aspects of recordings stored on optical disk, or even modify famous masterpieces.

- **Voice Recognition:** Because the Digital Ears sends high quality, high-resolution data directly to the DSP port, software that reads information from the DSP port could perform real time voice recognition.

- **Language Laboratories:** Software could play a phrase for a student, record the response, and critique the student's pronunciation.

- **Medical and Clinical applications:** Software for the analysis of speech and hearing as a diagnostic tool.

The Digital Ears can serve as a laboratory-grade, high-resolution A/D converter.

The Digital Ears can also act either as a two-channel, 44.1 kHz 16-bit A/D converter, or as a single-channel 88.2 kHz, 16 bit A/D converter. In both cases, the Digital Ears can function in either AC or DC mode. The Digital Ears' data acquisition capabilities have applications in harmonic/spectral analysis of signals, and other speed and resolution-intensive data acquisition tasks, in a number of fields:

- **Electrical Engineering**
- **Physics**
- **Signal Processing**
- **Biomedicine**
- **Normal Mode Analysis**
- **Harmonic and Signal Analysis**
- **Sampling theory**
- **Waveform analysis**

Appendix B: Technical specifications

Frequency Response (Stereo)	DC - 20kHz (-3db)
Sampling Rate (Stereo)	44.1 kHz
Sampling Rate (Mono)	88.2 kHz
Signal to Noise Ratio	
Input gain stage	98 db
A/D stage	85 db
Total noise + harmonic distortion relative to full-scale	
1 KHz sine wave	0.005%
A/D Input Voltage Range	±5V
Max Input Voltage Range	±8V
Input Impedance	100 kOhms
Gain Stage Noise Equivalent	25 nV/√Hz @ 10Hz
Linearity Error (DC mode)	±4 LSB (4 parts in 65,000)

Appendix C: Timing diagram

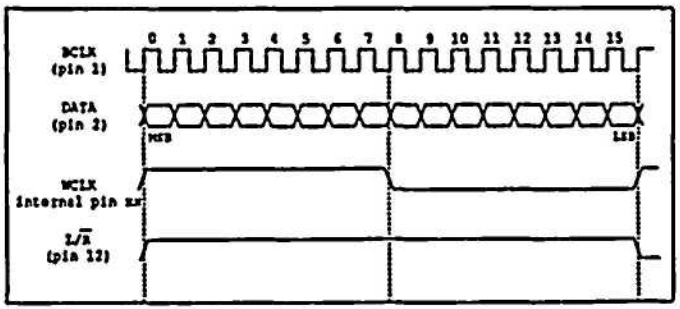

Appendix D: Reading a soundfile header

NeXT soundfiles all contain an SNDSoundStruct header that contains the information about the soundfile. Here is a fragment of C code that reads the information from a soundfile header, and then strips off the header to create a datafile.

```
/* Reading a soundfile header */
SNDSoundStruct *s;
int fd, fn, n;
int start;
int sorcery;
int size;
int format;
int rate;
int numchannels;
char comment[256];
char buffer[2048];

fd = open("soundfile.snd",O_RDONLY,0);
SNDReadHeader(fd,&s);

start = s->dataLocation;  //pointer to sound
data.
sorcery = s->magic;  //magic soundfile iden-
tifier.
size = s->dataSize;  //number of bytes of
data.
format = s->dataFormat;   //format of data.
rate = s->samplingRate;   //sampling rate of
data.
int numchannels = s->channelCount;  //number
of
                                    //channels
                                    //that data
                                    //represents.
strcpy(comment,s->info);   //comment string.

fn = creat("datafile",0666);
lseek(fd,start,0);
while((n = read(fd, buffer, 2048))>0;)
      write(fn,buffer,2048);

//now datafile contains the data of
soundfile.snd
```

Glossary

A/D. Analog-to-digital conversion. The process of converting a continuous signal into discrete digital information. See Appendix A, How Does the Digital Ears Work? for more information.

aliasing. A kind of distortion that results from sampling a signal whose frequency is more than half the sampling rate. It is manifested by a harsh, or grainy sound in the recording.

boot, boot up. To turn on a computer and load its operating system.

byte. A unit of computer memory, equivalent to 8 bits. Memory and disk space are usually measured in bytes, kilobytes(1,000 bytes) and megabytes (1,000,000 bytes).

CD. Compact Disk.

clipping. The distortion that results from an input signal whose amplitude exceeds the input tolerance of the Digital Ears. This tolerance is ±8 volts. Ideal recording range is from ±0.5 volts to ±5 volts.

CODEC. A coding/decoding scheme based on a standardized logarithmic scale, called a mu-law scale. CODEC technology is useful for compressing dynamic range. For example, the NeXT CODEC microphone has 8-bit accuracy, but 12- to 16-bit dynamic range because of an internal CODEC processor.

D/A. Digital-to-analog conversion. The complement of A/D conversion, that is, the conversion of discrete digital information to a continuous analog signal.

data acquisition. The process of measuring, digitizing, and sending laboratory data to a computer for subsequent analysis, testing, or display.

down-sampling. The conversion of data from a higher sampling rate to a lower sampling rate. For example, down-sampling from 44.1 kHz to 22.05 kHz involves deleting every other sample from the data, and doubling the length of time that the remaining samples represent. So sampled data (s0,s1,s2,s3,s4...) becomes(s0,s2,s4...)

This means the down-sampled data will have half the resolution of the original, and occupy half as much space.

DSP. The NeXT Computer's Digital Signal Processor. This processor is specifically designed for digital signal analysis. "The DSP port" refers to the 15-pin connector port on the NeXT cube that allows you to connect peripherals like the Digital Ears to the DSP. The DSP port acts as a high-speed serial port. The NeXT DSP is a Motorola 56001.

cut-off frequency filter. A filter which eliminates all frequencies higher than the cut-off frequency. For analog-to-digital conversion, the cut-off frequency should not exceed one-half the sampling rate.

Hertz. A unit of frequency, equal to one cycle per second.

k, kbyte. A kilobyte, equal to 1,000 bytes.

kHz. Kilohertz, equal to 1,000 Hertz.

line-level signal. A signal in the range typically produced by audio equipment before amplification, usually in the range of ±5 volts.

meg, mb, mbyte. A megabyte, equal to 1,000 kilobytes, or 1,000,000 bytes.

mHz. MegaHertz, equal to 1,000 kiloHertz, or 1,000,000 Hertz.

mu-law. A standardized logarithmic scale useful in compressing dynamic range. See CODEC.

signal. The term given to a continuous analog voltage that varies in time that represents a sound.

soundfile. A file on the NeXT Computer which stores binary numbers that represents a sound. When we talk about playing a soundfile, we mean generating a sound from the information stored in a soundfile. When we talk about recording a soundfile, we mean taking the information that represents a sound and writing it to a file on disk.

DERecorder saves soundfiles in 16-bit binary linear format, written Left-Right-Left-Right in stereo mode, and Left-Left-Left-Left in mono mode.

Soundfiles are like any other files on the NeXT Computer: they can be moved, deleted, opened, or edited. In the Workspace Manager, double-clicking on a soundfile's icon will make the NeXT Computer play that soundfile. Soundfile names always carry a ".snd" suffix, e.g. "splash.snd".

sampling rate. The number of times per second that the Digital Ears evaluates the amplitude of the input signal, and generates a number that corresponds to that amplitude. Sampling rate is often described in frequency, i.e. 44.1 kHz means 44,100 times per second. The higher the sampling rate, the more accurate the recording, and the more storage space the recording requires.

It is an axiom of sampling technology that signals that are greater than half the sampling rate will not be accurately represented by analog-to-digital conversion. The Digital Ears samples both channels at 44.1 kHz, which means the upper limit of the frequencies it can represent is about 20 kHz, which is also roughly the upper limit of the frequencies that the human ear can detect. This is why 44.1 kHz is sufficient for for high-quality sound recordings. The Digital Ears can also sample one channel at 88.2 kHz.

up-sampling. The conversion of data sampled at a lower sampling rate to a higher sampling rate. The up-sampled data will occupy twice as much space as the original data. Up-sampling is not supported in Release 1.0. See *down-sampling*.