Contents

Welcome to the CD WAV help file. I'm not a very handy help file creator, so don't expect all the fancy stuff you're used to from the '97 style programs. However, I'll do my best to help you out.

Topics

CD WAV – Why did I write it

I started this project shortly after a friend of mine asked if I could put some of his live Rolling Stones' albums from vinyl to a CD. I figured this wouldn't be much of a trouble, so I said yes. Turned out to be more work than I expected.

First, I have to make a decent recording. With a little trial and error, it's doable to find good volume settings and such, so the wave doesn't clip, and is still at an acceptable volume level.

After recording, I ended up with a 400 MB size recording. I was using the shareware version of Coolwave for the recording. The next task was to locate the song split points, and cut the wave file into smaller pieces. This is a harder task than I figured it would. To correctly split it up, I had to cut on 588 sample (or 2352 byte) borders. This was a lot of typing, calculating and waiting for the program to process its temporary files. I figured I could do better.

What's different

Let me start by telling that CD WAV is not a sound editor, even though it has the same "look and feel". The differences are mainly in the way it is supposed to be used. The primary functions of CD WAV are recording WAV files and splitting huge WAV recordings into smaller tracks, whereas sound editors are usually designed to make modifications to the audio data itself.

Now why would I use CD WAV instead of [fill in your favorite sound editor]?

- A sound editor usually works with temporary files. This becomes very time consuming if the files are hundreds of megabytes.
- CD WAV works in "batch mode". Basically, you define the tasks to be completed (split into tracks), and when you're satisfied, you "commit" these. A sound editor would apply each modification immediately.
- CD WAV records directly into a WAV file on your harddisk. This saves both space and time. The Windows Sound recorder program even can't buffer to harddisk at all.
- CD WAV always works in CD sectors. This ensures that tracks can be written to a disc in Disc-At-Once mode without clicks or unwanted pauses between tracks.

Now I know there exist sound editors that can do some of these things, but try to find one that allows all of this in a simple way…

Step by step

Basically, there are five steps in using CD WAV to put your favorite album on a CDR:

1. Set up the analogue recording part. This usually means connecting a tape deck or pick-up to your computer's sound card.
2. Determine good volume settings and other adjustments for your recording.
3. Digitize the recording, typically by putting it on your harddisk.
4. Split the recording into songs, one per WAV file.
5. Burn the CDR from these WAV files.

Steps 1 and 5 are the only steps CDWAV won't do for you, but I can give some tips.

Setting up your audio equipment

Everybody usually has different equipment. Some of you may own expensive profesional mixing and recording material, others will feel the home or portable HiFi set fills their needs. I cannot give complete information on every combination of audio and computer equipment. Most people owning advanced equipment probably know what they're talking about anyway, and probably can teach me things. I'll stick to giving tips on how to get it working with your home set.

Quick setup

My all favorite way to set up a computer as audio recorder is as a tape deck. The audio card usually has both recording and playback connectors, so if your HiFi set has a tape connector (preferably two if you'd want to keep using your regular set) use it. Connect the 'line out' (some soundcards only have 'speaker out' connectors, you'll have to dig into your manual to see if you can use those as line outputs too) to the 'play' or 'input' connectors on the audio set. Connect the recording output of the set to the line input of the soundcard. If you only have one external tape connector on your set, and you want to use both the computer and the tape deck connected to it, you can put them in 'series'. The play output from the tape deck connects to the line-in of the soundcard, and the line out of the soundcard now goes into the input of the audio set's amplifier. This allows you to directly record from the tape deck, and still be able to use everything like you used to. The tape deck and computer can switch positions to allow playing tapes while the computer is turned off, and to simplify recording from other resources than the tape deck.
Note that connecting a pick-up directly to your soundcard will not work, you need a phono pre-amplifier in between. Some, however, come with built-in pre-amplifiers, so you won't need to have one. Usually, the audio set will take care of it, so setting it up as a tape deck solves this issue.

Tuning volume settings

This is one of the most important steps in successful audio digitizing. To get a good signal-to-noise ratio, the volume should be as loud as possible. But to prevent clipping, the volume setting is limited. So the optimal setting is to have the sound just hit the clipping level. This is almost impossible to accomplish. In general, if your volume is above half the max level, you'll lose less than a single bit out of 16, or 6dB S/N of 98 total dB.

A test mode recording

Press the record button (red circle) to show the recording setting dialog. The first radio button gives you an option to run a 'test mode'. This mode will do everything the audio recording would do, except that it will just throw away the sound data after processing, instead of writing it to disk. The other options, like voice-activation and auto-stop, function exactly the same as in a normal recording, so the test mode can also be used to test these settings. After pressing the start button, the soundcard inputs will be enabled and any input will be processed. On screen, you'll see the min and max levels reached during short 1/3 second intervals. Below the green waveform, you'll see three LEDs. The green LED indicates that the signal is above 'silence'. If the audio level raises above the 'loud' level (usually about 90%), the yellow LED will flash. The red LED will indicate that clipping has occurred, and the sound volume should be lowered. The red LED can be switched of again by clicking it. Each (horizontal) pixel in the waveform display will represent 1/3 second of audio.
Note that the test mode recording only helps to calibrate volume levels and such. You cannot detect buffer overruns by this, because the critical part is writing the data to disk.

Actual recording

After connecting your equipment, and calibrating volume setting, you are ready to record your precious vinyl album or tape to the harddisk.

Harddisk recording

Press the record button. Switch the destination to file. You can type a filename, or use the browse button to point-and-shoot one. An estimate on available space is given directly below the filename.
You can use the other options too. The stop criteria are very handy when you won't be present during the recording. The recording can be stopped after a timeout (you probably know how long the recording will be) and/or a silence interval. Setting a silence stop criterion will stop the recording after the input goes silent for at least the given amount of seconds.

Buffer overrun

While doing CD recording, you'll probably came along the buffer underrun, the event causing sleepless nights and tons of coasters in CDR world. In harddisk recording, you can get a buffer overrun instead. While recording to harddisk, CDWAV records into memory buffers, and writes these to harddisk, so they can be reused. Like most wav recording programs, CDWAV has a total buffer of about 4 seconds. If, for some reason, the data is being written to harddisk too slowly, the buffer will overrun, as more data is arriving from your soundcard, but the buffer is full, so the data is lost. The recording will have skips in it. To prevent this, CD WAV can run as a high or even real-time priority process. See the settings section. To prevent overrun, your system needs to keep up a steady 174kB/s transfer to harddisk. A good 386-class machine can easily live up to these standards.

Splitting the file

Next step in building your own live CD is splitting the recording into tracks or songs.

Let me help you

CD WAV has a built-in auto-split function. This will look for silent parts, and add a split point in the middle of that silence. If your recording has silences, this function will do most of the work for you. If the recording has no silence intervals, as is usually the case with real live recordings, you'll have to do it manually. Typically one will first run the auto-split, and add or remove some split points afterwards.

Do it yourself

The bottom part of the screen holds the split list. You can see what split points have been added, and the corresponding time intervals. To add a split point, find the part of the song where the split is. Use the playback controls and the full view to navigate in the recording and roughly locate the spot. Now click the mouse in the bar directly below this point. The zoomed part of the waveform will show this in 'sector' detail. Click in this window to place the split marker. Press the split button to add the split. The full wave and split list will reflect your changes.

Removing a split

To remove a split, select the split to delete in the list at the bottom. Right click on it to popup a small menu. Select the remove option to delete the split.

Cutting out silent parts or songs

Sometime your recording will have excessively long silent parts in it. You can also cut these out. To do this, treat the silent part as a song. Add a split point at the beginning and end of the silent part (or just one if it's at the start or end of the recording).
Each split point has a check box in front of it. By unchecking this box, that part of the audio will not be saved to disk later. Uncheck the box for silent parts and songs you wish to leave out.

Saving

The last part is to save the songs you just recorded. You can modify the filenames to be written by clicking on them like you would in explorer to rename a file. The default name is the name of the recorded wave with the split index as a two-digit number.
Press the save button, and select a location for the files. All splits with a checkmark in front of them will be saved in that location. These files will be ready to burn.

Burn them

This is the part I can't do for you. You need a CD Recorder and appropriate software to do this. All software I know of will allow you to burn the WAV files to a CDR.
If you experience trouble with software claiming not to be able to read my WAV files, please contact the author of CD WAV, and tell him the complaint. None of Adaptec's CDR products had trouble with CD WAV.
Basically, if your system can record audio with CD WAV, your system is capable of running CDR software. I managed to burn flawlessly at 2x from a 486SX-33 (actually an overclocked 25 MHz) machine with a AVA1542 SCSI card and HP 6020 burner.

Common questions

Why is splitting on 2352 byte/588 sample borders so important?

A CD audio disc is divided into sectors. Each sector holds 1/75 seconds of audio, or 588 samples at 44100 samples per second, or 2352 bytes.
If the size of a WAV recording is not a multiple of 588 samples, the recording software will fill the remainder of the sector with zeroes. If you have a continuous recording (live), you'll hear a short click in between two songs, as a result of the padding zeroes. To prevent this, the program always cuts on 588 sample borders, so two adjacent songs will have no clicks in between.

The WAV files start with a terrible click

If you play back the split files, most of them will have a click at the start. The reason is that the data starts where the previous track ended. This results in a sudden jump in output level when playing back a single file. If you merge two adjacent WAV files together, you'll see that the click has vanished. The data will be merged on the CD, so the click will not be present on the resulting CD. Don't use a sound editor to fade the click, this will introduce a click at recording time! (unless you add a fade out to the preceding track)

CD WAV is shareware

Currently CD WAV is being released as shareware. This means that you can use this program and distribute it, free of any charge, provided that:

- You do not make any modifications to the software or any of its components
- The program is not used commercially. This means putting it on a disk with some other programs and selling this disk is also illegal.
- After a 31 day free trial period, you'll either have to register or stop using this product.

How to register

To register, put $10 (or a local equivalent) in an envelope, and send it to:
Mike Looijmans
Adam van Moorselstraat 12
6023 CD Budel-Schoot
The Netherlands

If you include your e-mail address, you'll receive your registration by e-mail. Otherwise, include an additional $1 for postage and handling, and you'll receive it by snail mail.

Why register

You should register to satisfy your conscience. The registered version has no added functionality over the shareware version. $10 is a ridiculously low price for this program, and if you think you just can keep on using it without paying, well, that means you'll go to hell and have to spend eternity with the other lamers.

The author

That's me!

Snail mail:
Mike Looijmans
Adam van Moorselstraat 12
6023 CD Budel-Schoot
The Netherlands

e-mail address (per January '98):
Mike.Looijmans@nym.sc.philips.com

New file

Choosing New will effectively reset the CD WAV program. You can use this to free resources occupied by CD WAV (like the WAV file and memory). You'll lose all changes and splits.

Open file

Open file will popup a display prompting for a WAV file. Select a WAV file from this list.

After selecting a file to open, the file will be processed. CD WAV reads the data, and precalculates the waveform. This will speed up later processing and the waveform display. The WAV file will never be modified by CD WAV, it always writes “fresh” WAV files.

When opening a file you’ll lose your current settings.

Save

Save will “commit” your work. It will write out the tracks as separate WAV files. You can use these files with a CD-R mastering program to write to a recordable CD, and you can also use them with other programs that can handle WAV files.
Note that you don’t need to “save” after a recording, since the WAV is recorded directly to disk.

Save cue sheet

You can save splits as a DAO compatible cue sheet. These are ASCII text files that contain information of the file to be processed and the splits to be added. Cue sheets have two uses:

1. CDWAV can read them, so you can save your splits for backup or to continue later.
1. Jeff Arnold's DAO program can use the cue sheet to directly record the audio data with split points as tracks to a CDR.

Load cue sheet

The load cue sheet option reads a cue sheet. A cue sheet holds the name of the wave file, and the splits to be made for that file. If the wave file is not already active, you'll be asked if you want CD WAV to load the wave file.

Recording settings

Recording destination

Usually you'll want to record an audio file. Some sound editors never ask for a file at this point, but record into a temporary file (or, like the standard Windows sound recorder, into memory) and copy that data to disk if you press the Save button.
CD WAV always records directly into a (valid) WAV file. This saves both disk space and time. You have to type a filename or select one using the browse button next to the input box.
If you wish, you can do a test mode recording. Everything recorded in test mode is just thrown away. This is useful for tuning volume settings to avoid clipping.

Voice activated start

If you check this button, the recording will wait until the input rises above the silence level before starting to record to harddisk. The time display will start running as soon as the actual recording starts. Thanks to buffering, the recording will include a few sectors before the input raised above silence.

Silence level

The automatic start and stop functions need this to determine whether the input is silent or not. If the sample's amplitude remains below this level, it will be considered silence. The number is a positive sample value. Samples range from -32768 to 32767. For example, setting this to 512 will trigger a voice activated start if a sample below -512 or above 512 has been detected. The green LED lights up during a recording.

Loud level

Much like the silence level, the loud level determines what level is considered to be "dangerously" loud. The only use is that the yellow LED will light up during recording if the input raises above this level.

Stop recording after ... minutes

If you are recording from some other media (like a tape or vinyl), you usually know the duration of the recording. Setting this field will stop the recording after the time elapses.

Stop recording after … seconds silence

This is much like a voice activated start. If, after the recording started, the input goes silent for a period longer than this, the recording will stop. This is convenient if you want to record something while you're eating lunch or so.

Auto Split

Splitting up a WAV file into tracks can be a tedious and time-consuming task. CD WAV can help you split it up by detecting pauses between songs. This is usually the first thing to try after opening or recording a WAV file.

How it works

CD WAV uses two criterions to determine where the pauses are: The silence level and the silence time. The silence level determines what music level is considered to be silence. The level is a sample value, ranging from 1 to 32767. Sample values can range from -32768 to 32767, if the absolute value of all samples is below the silence level, it is considered a silent part. If this silent part lasts longer than the silence time criterion, the silent part is considered to be a pause, and a track split will be added in the middle of the silent section.

Settings

This dialog lets you change some of CDWAVs general behaviour. Choose *Preferences* from the *File* menu to access this dialog.

Windows ’97 style flat buttons

The default is to show buttons in “flat” style, which means they pop-up if the mouse moves over them. If you don’t like this setting, switch it off, and the buttons will have the “classic” Word 95 look instead.

Colors

The full view uses eight colors to display the individual tracks. Click on any of the buttons to change the color.

Recording Priority Boost

As mentioned before, in the actual recording description, recording to harddisk is a time-critical operation. CD WAV can set the priority of its own process to a higher level, to allow the program to get more ‘attention’ from the operating system.

Windows divides time into small units, called time slices. Each time slice, it looks which processes (read: programs) like to use the CPU. It will first handle real-time priority tasks, then high priority, normal priority, low priority and, if there’s nothing else to do, it will run idle priority processes.

You can look upon this as some person working in a shop that has three queues (VIPs, common people, and lamers). He/she will only serve people from the second queue if the first queue is empty. Only if no-one is waiting in the second, he (or she) will serve the first person from the third queue. If there are multiple servants (multiple CPUs), they’ll do the same. So if the first (and/or second) queue keeps filling up, no-one from the second (and third) queue will ever be served. This is called “starvation”.

Assigning CDWAV to high-priority during the recording process allows it to “attract” more attention than other programs to do its tasks. By default, a process (program or task) runs as normal priority. If you are running Word during the recording, Word will not get any attention while CD WAV needs to process incoming data. As CD WAV needs to process only about 172kb/s, even a 486 class system should hardly degrade noticable in performance. You can compare it to a CDR session at single speed.

Setting high or even real-time priority will not improve performance. If no other programs are running, even idle priority tasks will get full resources, and run at the same speed a real-time priority process would run. Thus, if your system isn’t capable of the 172k/s transfer, this option will not help.

If you feel the need, you can even go for real-time. This will give CDWAV even more priority than for example the mouse cursor.

The default is high-priority.

Use seperate thread

This has been added in version 1.2. The actual process of retrieving audio data and writing it to disk will take place in a different thread. This provides a much safer mechanism and can help prevent buffer overruns.