

# SURROUND SOUND IN LINUX

## 1. Introduction

Getting applications in Linux to play surround sound can involve some work, especially when you want the highest quality. Don't think you're done just because you can hear sound, surround or otherwise. It's possible you still need to do some fine-tuning and calibration to get the most out of it. In this article, I will show you what you need to do get surround sound working properly, in both games and other applications. When you're done, there's no reason to boot Windows to watch movies anymore. In fact, the sound quality is possibly higher than you were used to in windows. Read on to find out why.

Note that I'm not talking about duplicating the audio of the front speakers to the rear speakers. Two channel sources are to be played back on two speakers. Anybody who tells you otherwise doesn't know what stereo-imaging and soundstages are. What I mean by surround sound, is playing four or more channels of source audio, like surround sound in games and DVD movies.

What you first need to do, get Alsa to work. For most people, this is no issue, and is already done.

Note: there is actually a mistake in this article. The 10 dB boost with which the LFE channel is supposed to be played back is not included. I'm working on a correction.

---

## 2. Setting up a custom Alsa device

Most applications output in 5.1. If you only have a 4.0 soundcard, like me, you have to make a custom sound device for the applications that don't provide downmixing themselves, or if you don't trust applications to downmix properly (like me...). I prefer to let every application which plays Dolby Digital 5.1, also output in 5.1, and control myself how that's routed to 4.0. This way, I'm sure it's done properly. As they say: "if you want something done right, you have to do it yourself."

Alsa has predefined output devices, like *surround40* or *surround51*. When you instruct a DVD player to output to *surround51* on a 4.0 soundcard, the center and LFE channel will not be played. You want to control how downmixing to 4.0 is done, so you're going to define a custom sound device in */etc/asound.conf*:

```
pcm.51to40
{
    type route
    slave.pcm surround51
    slave.channels 6

    # Front and rear, at 50% of original signal strength
    ttable.0.0 0.5
    ttable.1.1 0.5
    ttable.2.2 0.5
    ttable.3.3 0.5

    # Center channel routing (routed to front-left and front-right),
    # 6dB gaindrop (gain half of main channels) per channel
    ttable.4.0 0.25
    ttable.4.1 0.25

    # LFE channel routing (routed to front-left and front-right),
    # 6dB gaindrop (gain half of main channels) per channel
```

```
# TODO: actually this is wrong. This article awaits an update.
ttable.5.0 0.25
ttable.5.1 0.25
}
```

You then have to configure each application to use 51to40 as 5.1 device. Much to my surprise, with one DVD program, this was not possible. I believe it was Ogle or Video Lan Client, but I don't know for sure. For Xine (xine-lib 1.1.4), I have to specify "*plug:51to40*" as output device.

You may be wondering why I don't trust applications to downmix themselves. There are multiple reasons: one is that people often think the gain drop should be -3 dB (or 0.707 gain) to play one channel over two speakers, but that should be -6 dB. You can test this very easily with *speaker-test* (comes with *alsa-utils*), testing with non-directional sound (low frequencies). Run "`speaker-test -c 6 -D 51to40 -t sine -f 40`" to play a 40 Hz sine on every channel. Center, left and right will sound equally loud, whereas center would sound louder when you apply only -3 dB gain drop.

Another reason is that the application may decide to apply dynamic range compression to the downmixed result, to avoid reaching outside the maximum sampling value of a 16 bit value, which can occur when summing signals in the digital domain. I absolutely don't want dynamic range compression and this way I can be sure it's not applied. To avoid clipping distortion, I decrease all the channels in gain somewhat. The gain of the main channels should be 0.5, and the center and LFE 0.25. The reason for that is, with all channels at maximum signal strength, the summed total value is exactly 1.0, so you don't get clipping.

By decreasing the gain on all the channels like I do, you do need a sound system that has some room to spare on the volume control and not have too much noise,

otherwise you may not be able to get enough usable volume out of it. A better solution than decreasing the gain, is getting a soundcard which can output in 5.1, and then use those separate signals to feed an amplifier. And of course, when feeding a digital amplifier with s/pdif from a DVD or other AC3 source, this issue is avoided altogether (as is the need for this entire article).

When you instruct software to downmix dolby 5.1 to 2.0 matrix encoded (there usually is a special option for that), dynamic range compression is always used. The Dolby Digital standard demands it, and because two channels have a lot less combined dynamic range than six, this is also desirable. However, if you really want all the dynamic range of the surround encoding in two channels, you can use this output configuration:

```
pcm.51to20
{
    type route
    slave.pcm surround51
    slave.channels 6

    # Front and rear, at 33% of original signal strength
    ttable.0.0 0.33
    ttable.1.1 0.33
    ttable.2.0 0.33
    ttable.3.1 0.33

    # Center channel routing (routed to front-left and front-right),
    # 6dB gaindrop (gain half of main channels) per channel
    ttable.4.0 0.16
    ttable.4.1 0.16

    # LFE channel routing (routed to front-left and front-right),
    # 6dB gaindrop (gain half of main channels) per channel
    ttable.5.0 0.16
    ttable.5.1 0.16
}
```

}

The dynamic range compression requirement also indicates why it's bad for the "standard" to become 7.1 when you only have 5.1 equipment (like most people do). It's very possible you lose dynamic range, and therefore realism, if you downmix it to 5.1. I don't know if 7.1 audio formats demand it (like Dolby Digital 5.1 demands it for downmixing to 2.0), but it's possible. Perhaps 7.1 is better (I don't think so, but that's a different story altogether), but it's annoying how the industry keeps changing the standards. For once, I would like to keep using my equipment for more than two years, before it's outdated.

Yet another reason to configure the mixing yourself, is to make sure nothing is forgotten. You may find it hard to believe, but for a long time, the Windows Player PowerDVD neglected to mix in the LFE channel with the fronts in stereo or 4-stereo mode, even if it was enabled. This is not as obvious as a missing center channel (because that would cause the dialogue not to be audible), and therefore you won't notice it as easily, but you have to agree that it's highly undesirable. When the developers of PowerDVD make stupid mistakes like that, it's also possible it doesn't handle the clipping distortion in the correct manner (and it doesn't give you an option to configure it). I never saw any DVD program that lets you configure that. All very good reasons to dump proprietary players, and use something of which you can be sure it sounds the way it is supposed to :).

A final word of caution: if you have a sound system which can't handle the bass from the LFE channel, or it won't go low enough to make it audible, you'd better leave it out (set gain to 0). If you don't, you may cause damage to your speakers. Even hi-fi full range speakers can have difficulty coping with the LFE signal sometimes. If your speakers tend to have a high excursion for bass signals, it may not be wise to mix in the LFE signal. At some point, you might watch "Black Hawk

Down", and suddenly find your speakers making break-up noises at the "Fucking Irene" scene, because of the intense 7 Hz infrasonic tone...

---

## 3. Setting up OpenAL for games

For surround sound to work properly in games, you need to configure OpenAL properly. Newer versions of OpenAL work differently, so there are separate sections for old and new versions. I don't know exactly which version started using the new scheme, but I'm guessing pre 1.0 is old and post 1.0 is new.

### 3.1. Configuration for old OpenAL versions

The OpenAL configuration file is `/etc/openalrc`, or `~/.openalrc` to override system wide settings for a local user. For a four speaker setup, it should contain:

```
(define speaker-num 4)
(define devices '(alsa))
(define alsa-device "surround40")
```

The apparent syntax error on line two is no error, since the file is LISP based.

If you have more speakers, it's possible you need to change the *speaker-num* variable and the *alsa-device* (*surround50*, *surround70*, etc). However, you might want to experiment with this setting, as not all games support it. Unreal Tournament 2003 for example reverts back to stereo when you use anything other than four channel output. I don't know if this is Unreal Tournament 2003's fault, or OpenAL's (I haven't tested any other game...). I have seen a remark somewhere that older OpenAL versions only supports two or four channels, so it would seem to be an OpenAL limitation.

## 3.2. Configuration for new OpenAL versions

The config file is either `/etc/openal/alsoft.conf` or `~/.alsoftrc`, depending on if you want to use system wide or user specific settings. The equivalent of the config file above for newer OpenAL versions would be:

```
format = AL_FORMAT_QUAD16

[alsa]
device = surround40
```

You can of course specify different output schemes, if you want more channels for example. If you need more information about that, you can look at the example config file in `/usr/share/doc/openal-(something)` or at the one I use for version 1.7.411.

Whereas with older OpenAL versions, you had to use 4 speaker output for games that didn't support otherwise, with the newer version, OpenAL does the downmixing for you. So where Unreal Tournament 2003 didn't work with 5 channel output on older versions of OpenAL, it does on newer.

## 3.3. .1 output

I must say I don't really know what to do with the .1 output. I do know that for older versions of OpenAL you shouldn't use it (for example, use `surround50` as opposed to `surround51`), but the newer one does seem to understand it. However, in my opinion, it has no added value for OpenAL (unless apps like dvd players start using OpenAL). Computers always drive a receiver or speaker set of some sort and those are the ones that should take care of bass routing. The computer, *which is a source*, should only output .1 data if the source has .1 data (a DVD with AC3 sound for instance). When there is no .1 data, the .1 output of

the computer should be silent. The reason is very complicated to explain, but it has to do with that people often think that the .1 output on a DVD equals subwoofer output, but it most certainly does not.

If the OpenAL developers did their work properly, the .1 output on your computer should be silent with games (because they use mere stereo or mono PCM data for the sound effects) and have information with AC3 5.1 soundtracks.

Perhaps I'll do some tests in the future to determine OpenAL's bass behavior, but for now, it doesn't really seem necessary.

### **3.4. Pointing to the correct OpenAL library**

If by now surround sound still doesn't work, it's possible the game in question is shipped with it's own OpenAL library, instead of using the system's version, and may not support surround sound. This is the case for Unreal Tournament 2003 for example. Find the openal library file in the game's directory, rename it and symlink the system's version (most likely called */usr/lib/libopenal.so.something*) in it's place. Repeat this procedure when you upgrade the game, because the file will most likely be replaced when upgrading.

---

## **4. Calibrating the mixer**

Mixer calibration can be a bit confusing. For the Sound Blaster Live!, there are separate controls to control the level of the rear speakers. One is for stereo mode, the other for surround mode. Added to this, is that the master control only controls the front channels. The control called "surround" is for controlling the level of the rear speakers in surround mode, when using *surround40* for example. The control

called "wave surround" is for controlling the level of the rear speakers in stereo mode, and will contain a copy of the fronts. Because stereo should be stereo, I recommend turning that control all the way down.

When you're done calibrating, save the state with `"alsactl store"` (you can restore with `"alsactl restore"`). I'd also advise you to disable saving the mixer state upon shutdown of your computer. Where this is configured, is distribution dependent, but in Gentoo, it's configured in `/etc/conf.d/alsasound`. Debian based distributions, like Ubuntu, probably have some file for that in `/etc/default/`. I also suggest you configure applications to use software volume controls instead of hardware, or you will easily mess up your carefully made calibration.

---

## 5. Disabling dynamic range compression for Dolby Digital AC3 playback

Dolby Digital AC3 (which is used on DVD) has meta data in it's stream to allow the audio decoder to compress the dynamic range of the audio in a controlled manner. The default for most software (and hardware) is to have this enabled. If you care about the audio quality, I recommend disabling this. Every program that can deal with AC3 (mplayer, xine, etc) has an option for this. A side note here, is that the Windows programs PowerDVD and WinDVD don't support the highest dynamic range. Even when you set it to "highest", it actually uses the "normal" setting. That is, the last time I used Windows (for playing DVD's), which is quite a while ago.

A funny hypothesis I have formed, and seems to hold up, is that this default setting, of both hardware and software AC3 decoders, is actually the reason why

people often say DTS sounds better, while in fact, they sound just as good. Disable dynamic range compression, and notice how the AC3 tracks of your DVD's suddenly sound a lot better. Just think about it, would Dolby create a standard that makes the audio sound so much worse than the original master? In fact, Dolby's technical documents show they picked the bitrate so that there is no perceivable quality difference with the master. By the way, that's also a reason to prefer Dolby Laboratories over Digital Theatre Systems, they are much more open with information, and they don't sabotage open source development of decoders and encoders for their product.

I will probably write a dedicated article about the subject. Should you in the mean time want to know more about my reasoning behind this hypothesis, or supply me with information why DTS really is better, you can contact me.

---

## 6. Testing

You can use this file to test your surround setup. It's a six channel wav file which identifies which channel is what, by playing "front left", "front right", etc. At some point you hear "front left", "front right", etc played back simultaneously. That's the LFE channel.

It's a special kind of wav file, so it won't work in just any player. I use the command line player *aplay*. You can use the parameter *-D* to specify the sound device, like `"aplay -D 51to40 chan-id.wav"`.

You can also use the *speaker-test* utility (mentioned above) for testing. If you supply it the *-t wav* option, it plays back a recorded voice identifying each channel. However, on my system, it calls the LFE channel "rear center", which is

obviously wrong, even though my configuration works fine. I submitted a bugreport about that to the Alsa bug tracker.

Source : <http://www.halfgaar.net/surround-sound-in-linux>