Tutorial 03
Pitch and Time
Using changes in pitch and time to create new sounds

written by Stephan Schütze

Overview
This tutorial will cover two of the simplest and yet most effective ways of manipulating sound. These techniques originated with film sound designers creating sound effects on tape machines by speeding up and slowing down the rate at which the tape ran through the machine. Current digital editors make altering pitch and time much easier but there are still limitations.

This tutorial will use Audacity a free to download audio editing program which can be downloaded at the following location:

http://audacity.sourceforge.net/

If you do not have Audacity installed download and install it now.

Part 1. Basic Method
Open a sound file for the tutorial.
I am using **Paving_stone_move01.wav** from the materials/earth section of the library. Click on the link to download this sound if you don't have it already.

You can use any sound file for this tutorial but the results will obviously differ. Open the file in Audacity in the same way you would open any file on a PC. Regularly saving project work avoids loosing material should there be either software or hardware issues.

Paving_stone_move01 is a recording of a small concrete paving stone (about 50cm square) being dragged across another paving stone. This tutorial will go through the steps to alter one sound to create a new sound that may otherwise be too difficult to record normally.

Select the sound file buy holding down the **Ctrl** key and pressing **A** (this is *select all* **Edit/Select All**)
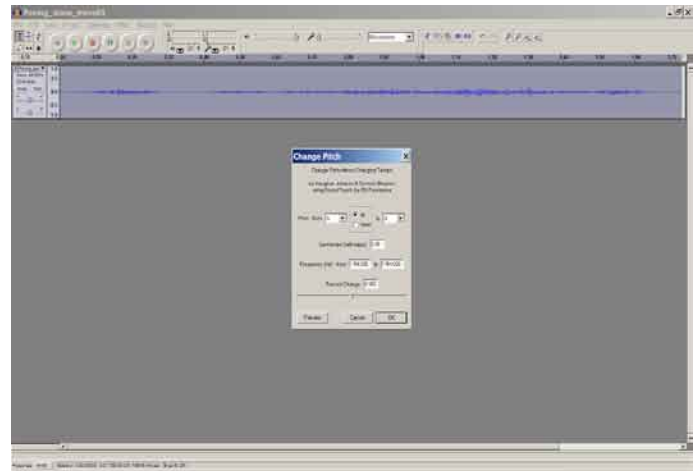Go to **Effect/Change Pitch**. The following window should appear:

Figure 1. Change Pitch window
This function allows for the pitch of a sound to be altered by various methods.

The window in Figure 1. allows the pitch values of sound files to be altered. The **Preview** button allows you to hear any changes made to the sound in this window before committing the changes. Pressing **Ok** will alter the sound file with any changes made. If no changes are to be made press **Cancel**. Clicking **Preview** now should play the sound as it is in the edit window because there have not yet been any changes made.

There are several different ways to change the pitch using this window. The musical note letter values change the pitch from its current value of "A" up or down to a new musical note value. This changes the pitch by one semi tone for each note value up or down. Altering the frequency values will also alter the pitch up or down, but by smaller increments. This is very useful for achieving more exact results; to frequency match different sounds to mix together later. The third method is the simplest, using the slider to change the overall percentage value in relation to pitch. Negative values will make a sound lower in pitch and positive values higher.

You are going to make the paving stone sound like a much bigger stone. In the real world as things get bigger they tend to have a lower pitched sound. This isn't always the case, but it's a good general rule. Because of this I you should select a percentage value on the slider of around -82%, this value does not need to be too exact for this tutorial. Once the slider has been moved to its new position press the **Preview** button to have a listen to what the new sound will be. Once you click **Ok,** this change will be committed to the original sound file. Tap the **Spacebar** to listen to the new sound. It should sound much lower in pitch than the original sound.

Time Out

I want to pause here for a second and explain something about the process just used. The original sound has been altered by changing its pitch and only its pitch. This is a process that has only recently become possible with the invention of digital sound processing software. Before computer based software, it was possible to alter the pitch of sounds on tape by slowing the tape down, but this also resulted in the sounds becoming longer. The tape is slowed down. It moves through the tape machine slowly, taking more time, but being lowered in pitch. This is the same as playing a 45rpm record at 72rpm, the record sounds much lower in pitch, but it also takes longer to play. The reverse is also possible. Speeding up a tape or record will raise its pitch. This is the classic "Chipmunks" effect.

Digital sound processors allow pitch to be altered without changing speed and in the same manner speed to be altered without changing pitch. Audacity's **Change Pitch** function only alters pitch, not speed. I have found that drastic alterations of only one characteristic can often result in artefacting. This is when severe changes to a sound file slightly corrupt the sound information and results in extra (usually unwanted) sounds mixed with the original. Badly compressed mp3 files often have this problem. The process of slowing a tape is a natural process, as a guitar string gets longer it gets lower in pitch and the same thing happens with tape sounds. When a software package tries to alter one characteristic without altering the other, it may need to add or subtract digital information.

The software may not always correctly calculate how to add or subtract this information and this is what causes the artefacts. Newer software can often do a very good job, but the more dramatically you alter the sound, the more the software needs to guess, and it can often get things wrong.

The reason why I mention this point is because I sometimes like to have a sound lengthen when I pitch it down. Often in real life a large low pitched sound takes longer to generate anyway. A small bird tweet is a high pitched, quick short sound; where as an elephant or whale make long low sounds that generate more energy. To turn a high pitched bird tweet into a sound for a giant dinosaur, it often makes sense for the sound to become longer as well as lower in pitch.

In the case of the stone paver it needs to sound like a giant block of stone from a mummy's tomb. It would certainly take a lot of effort and probably time to move such a big chunk of rock, so the sound needs to be longer.

Time stretching

Now the pitch is appropriately low, but the sound needs to be longer. The first process only altered the pitch. It is possible to alter just the time by using the **Effect/Change Tempo** function, this will alter just the time characteristic of the sound. Another option is to change both simultaneously.

Undo the pitch change from the previous process. Go to **Edit/Undo** and click to undo the **Change Pitch** process (**Ctrl Z** will also undo the last step). The sound file should now be in the original state it was when opened. This next process will change the pitch as well as changing the time in the same manner as with using an old tape reel. Use **Effect/Change Speed**.
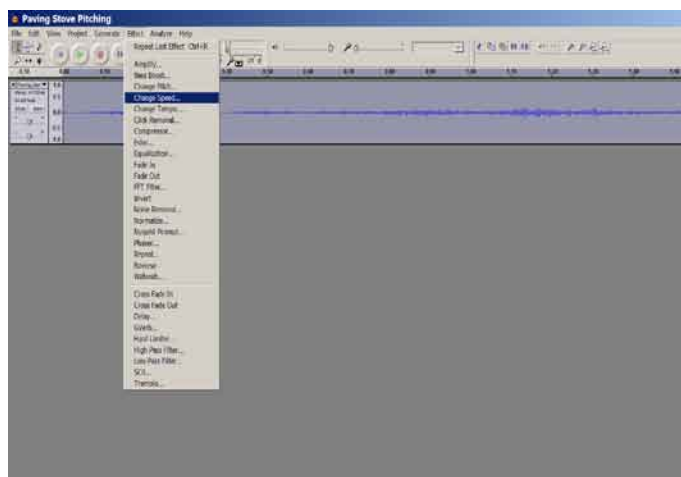

Figure 2. Accessing the Change Speed window
This function allows for the speed at which a sound plays to be altered

This window has fewer options than the change pitch window because like speeding up/slowing down a tape or record the one action will affect both pitch and length. Before changing anything note that the original sound file is about 1.7 seconds long. Move the percentage slider down to about -84% and click on **OK**. Once the software has processed the file you will notice that the sound file is now about 10.5 seconds long. At the very least, it has made the sound much longer. Press **Play** now and you should here what sounds much more like a giant slab of stone being dragged or pushed across a stone floor.
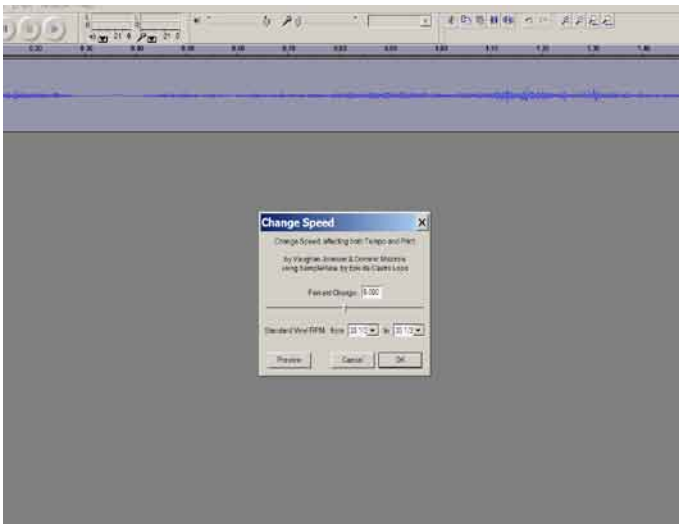
Figure 3. Change Speed window
This function allows for the speed at which a sound plays to be altered.
This functions in the same way altering the speed of a tape deck would.

Just one simple process has turned a regular household paving stone into the door to the mummy's tomb!

The next process will enhance the sound's dynamic range and maybe sound more like what you might hear in a movie. EQ or frequency equalisation. This is boosting certain frequencies of the sound to make them stand out more, just like turning up the bass on a home stereo system. Go to **Effect/Equalization** and this window should appear. You are going to work with the little graph window part of this tool.
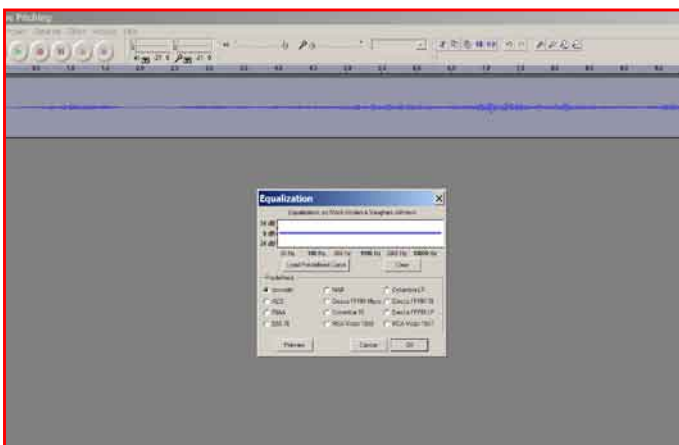

Figure 4. Equalization window
Altering the amplitude of specific frequency groups within a sound file can alter the overall feel of the sound.

The light blue line represents all the frequencies present in a sound file. Clicking on the line creates a node that can be dragged up and down. The point where the node is created is the range of frequencies that will be altered. Dragging the node upwards will increase the amplitude of those frequencies, dragging the node down will decrease them.
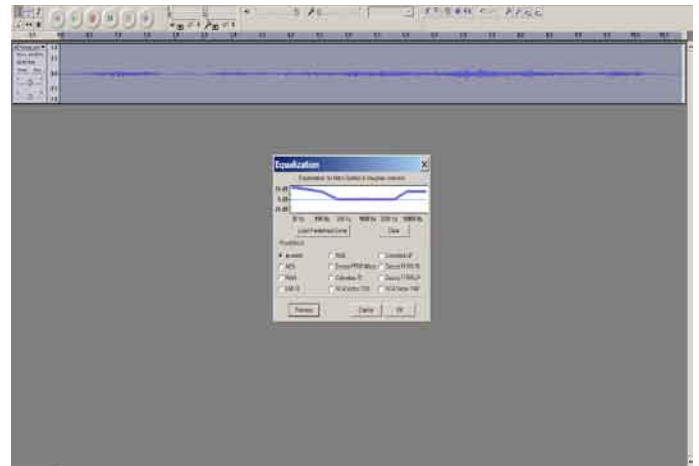
Figure 5. Altering frequency levels
The line indicates frequencies progressing from the lowest at the left edge to the highest on the right . This sound will have a boost to frequencies at both the low and high ends with the middle frequencies remaining flat.

In Figure 5. the frequencies have been raised considerably at the lower end to give the sound more weight. A lot of movie sounds are very heavy on the bass end. Be careful though, if you increase only the bass frequencies a sound can become very muffled and unclear. That's why the high may need to be raised as well. This helps to make the sound more crisp. The "sand" element of sound in our sound file will stand out more by increasing the high end frequencies and make the sound cleaner. Experiment adding nodes anywhere on this line and drag them around. Listen to how the sound changes by using the **Preview** button. This is the best way to understand how EQ tools work. Getting the absolute best out of EQ-ing can take a lot of time and practise and often comes down to personal preferences.

The final step is to "normalize" the sound. This can be a good habit to get into with sounds that are being created to be used as part of a project with other sounds. When normalized, the amplitude of the sound is set to a level defined by the editor. Often this is maximum amplitude. This tends to be the 0 dB point in your editing tools. Anything above this point means your sound will red line and is prone to distortion. Normalizing alters the level of the entire sound file proportionally, it increases the amplitude of the loudest part of the sound file to 0 dB while keeping the ratio between all amplitude levels the same. This keeps the overall ratio of dynamic range within the sound file the same. The entire sound will be increased consistently. The advantage of normalizing is that the loudest points of all the sound files in a project will be equal. This avoids sounds being inconsistent in amplitude with each other. Normalizing will set the sounds at the pre determined level this may be an increase or a decrease depending on the level that is defined. Go to **Effect/Nor-malize.**
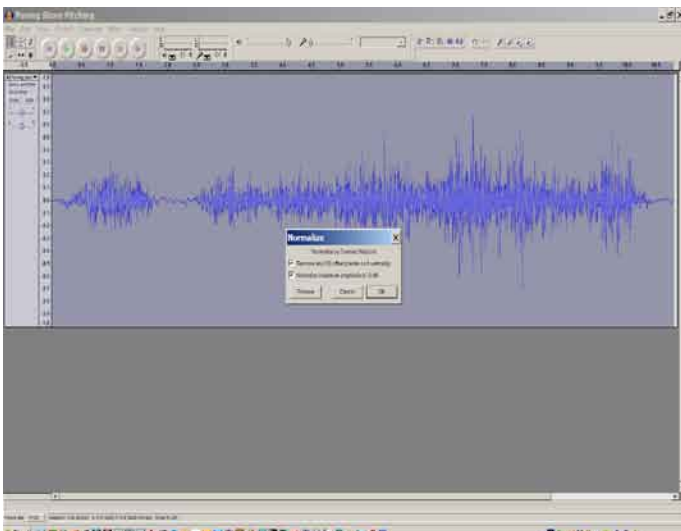
Figure 5. Normalization window
Normalizing will maximize the volume of a sound file while keeping its dynamic levels relative.

The window has two tabs. One refers to normalizing to -3DB. This ensures that the sound file will not distort by going over 0DB. Some programs will allow you to set the normalization range as an overall percentage with 100% being 0DB. Each editing program will have its own way of dealing with various processes.

Leave both tabs ticked and click OK. The sound file has now been pitch and time shifted with some EQ and normalization added. It's a good time to save your work again.

Part 2. Quick and easy monsters

Continuing on from part 1 this section will alter an animal sound to create a basic monster sound effect. I used the sound **Lorrikeets09** This is the sound of an Australian Rainbow Lorikeet. This is a small brightly coloured bird that has a range of different sound it produces. The sound in question is a series of agitated chirps from a lorikeet when it thought I was getting too close to its breakfast bowl. Open the file, it should look like this.
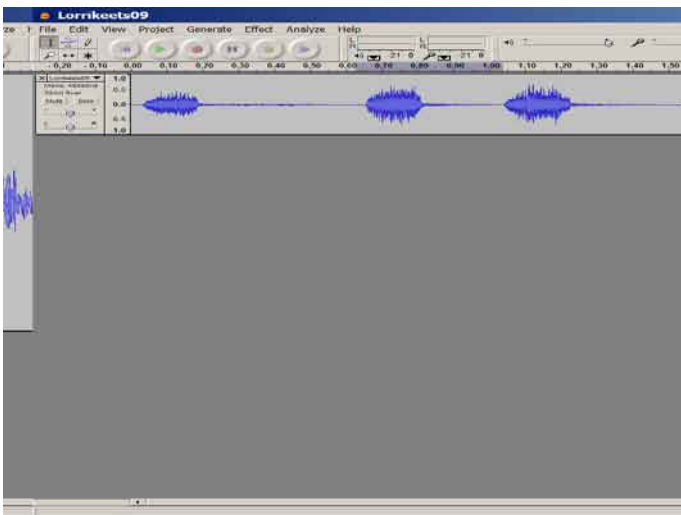

Figure 6. Sound file of Lorikeets Chirping
This is a recording of a single bird making three chirps

If you listen to the sound you can hear that it's a series of chirps that are about three and a half seconds long in total. The whole sequence is not needed, so isolate one of the chirps. Drag select an area of the file that contains a single chirp. The area should turn a darker blue.
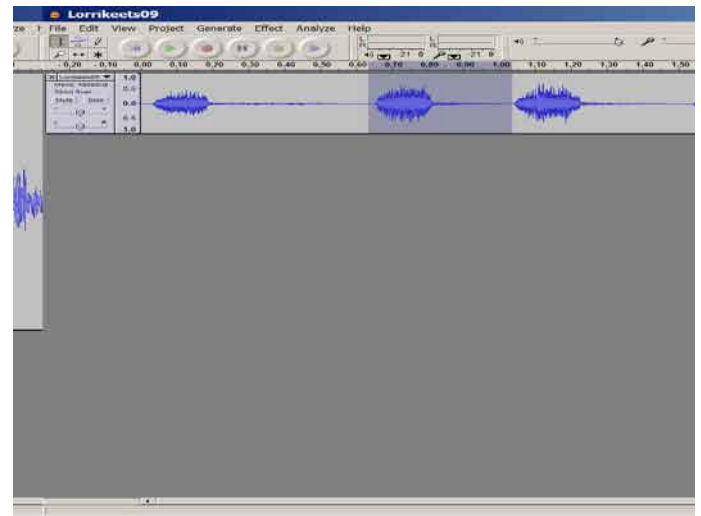

Figure 7. Drag Select highlight
By dragging an area within a sound file it is poosible to select an isolated section. This section can be copied or have effects added to it without effecting the rmaining sound file.

This is the same process used to select the end of the file in the previous section. Go to **Edit/Trim**. This function will delete everything but the selected area. There is now a single chirp sound to work with. Using the time shift tool you can drag the sound file back to the left hand edge of the work space. This is not necessary but I prefer to work this way.
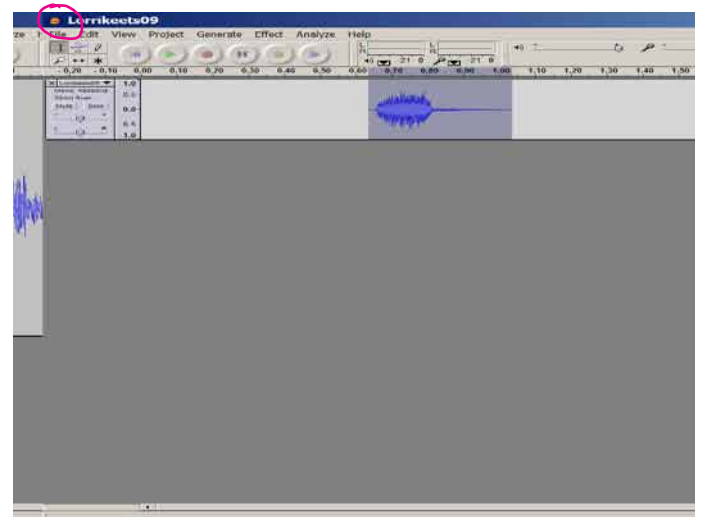

Figure 8. Time Shift Tool
The time shift tool allows the user to drag the selected area back and forth in the time line. This can be used to order multiple sounds or to simply arrange a single sound in a prefered positon to edit.

Bring up the change speed window and move the slider to about the -75% mark. Click **Ok**. The sound has changed from about 0.4 of a second to 1.7 seconds. It has also pitched down considerably and sounds more like some giant winged dinosaur. You can repeat this process with all of the individual chirps in the original file. Isolating them and then pitching them down. To finish off each sound add a fade out over the last 0.25 of a second and then normalize them. You can experiment with altering the speed change value slightly for each one. That way you can create the effect of a herd of prehistoric beasts. For further practice in this area simply select any bird or animal noise from the library and play around with its speed or pitch value. With a little exploring and persistence I suspect you will start to find many of the monster sounds heard in films and TV.

Section 3   Pitching up

For this last section on pitch shifting you are going to simply reverse things and take a low pitched sound and make it a higher pitched sound. Start with **fist_on_metal_echoe02**   This is a sound of me bashing a large metal surface with my hand and recording it. Because the metal was part of a playground that was half underground there is a long echo in the sound as well. This is useful because the reverse process to sections 1 and 2, raising the pitch by speeding up the sound, it will also be shorten the sound.
Select the sound and use the **Effect/Change Speed** window. Change the speed 380%. This time the sound file will alter from about 3.5 seconds to only about 0.7 of a second. The sound has also changed from sounding like a large heavy sheet of metal to sounding more like a tin can. Clean up the sound by deleting a small amount off the end of the file that is very rattly sounding. Add a **Fade Out** on the last 0.25 of a second. This will give the sound a good tight punchy metal sound. Unfortunately Audacity doesn't have a reverb function, a small amount of reverb added to the sound would add a musical quality to the sound. The end result I think could work really well in some kind of industrial or noise piece of music. In this way you can create new sounds to use as sound effects or even new sounds to use as samples to create music.