NCH Software

Express Mix MultiTrack Recorder

This user guide has been created for use with

Express Mix MultiTrack Recorder Version 12.xx

Technical Support

If you have difficulties using Express Mix MultiTrack Recorder please read the applicable topic before requesting support. If your problem is not covered in this user guide please view the up-to-date Express Mix MultiTrack Recorder Online Technical Support at

www.nch.com.au/mixpad/support.html.

If that does not solve your problem, you can contact us using the technical support contacts listed on that page.

Software Suggestions

If you have any suggestions for improvements to Express Mix MultiTrack Recorder, or suggestions for other related software that you might need, please post it on our Suggestions page at www.nch.com.au/suggestions/index.html.

Many of our software projects have been undertaken after suggestions from users like you. You get a free upgrade if we follow your suggestion.

Express Mix MultiTrack Recorder

Contents

Audio Cleanup

NCH Sound Library

Settings

Screen References

Overview

MixPad is a multi-track mixing software package designed for professional audio production. Simply load existing audio files, or record new files into MixPad, adjust the volumes, pans, fades, add effects and mix your audio clips down to create a single high-quality audio file. It is the digital equivalent of using a multi-track recorder and mixing desk.

Features:

- Pan and volume envelope automation.
- Intuitive, simple graphical interface.
- Independent volume, pan and effects for each track.
- High accuracy for precise mixing.
- Solo and mute feature on each track.
- Add multiple chained effects to each track.

MixPad is just one component of the NCH Software suite of audio software. If you have not done so already, why not visit www.nch.com.au to download and trial our other professional audio software packages.

Getting Started with MixPad

A MixPad project consists of one or more tracks. Each track has its own set of controls on the left side of the screen for controlling the overall volume, pan and other playback features of the track.

Each track can contain audio clips. Audio clips are the pieces of audio that you wish to mix together into one file. Whenever you add an mp3 file, or make your own recording, an audio clip will be added to the currently selected track.

When you have adjusted all the volumes and fades exactly how you want them, you can output your project to a .wav, .mp3, .gsm or one of many other formats by simply selecting Export Project as Audio File from the File menu and then choosing the format you wish to save in. Alternatively, if you are not finished mixing and want to come back to it later, you can save your project to be reopened for later.

Choosing Your Audio Settings

Setting up the Project Audio Options

Select 'Audio Options' from the Recording tab.

To set audio options that are specific to an individual track, select the track and click either the 'Audio Options' button under the 'Recording' tab or click the 'wrench icon' button to the right in the track's control panel.

Track

The **Select track's settings** displays the track you are editing by it's Track Index number. You can switch to any other of the tracks listed from the drop down list. If you make a change to the any of the options below, choose **Apply Settings** by clicking **Apply** or click OK to apply those settings to the track. If you make a change then click **Apply**, Choose **Apply Settings to all tracks** it will apply the change to all tracks.

Recording options

Here, you will be able to set the recording device you wish to use for this track.

To set up recording using an audio device, select the **Record using an audio device** option and select the device you want to record. You can specify whether to record as mono or stereo and if the device can record multi-channel you can specify which two channels will be recorded.

You can specify whether to record as mono or stereo and if the device can record multi-channel you can specify which two channels will be recorded.

MIDI Playback options

This section allows you to choose a MIDI device for playback.

Loading an Audio Clip into the MixPad Work Area

Use one of the following methods to load a clip into a track. Note that MixPad may relocate the clip to the next track if the clip being loaded will overlap other clips.

Record a track using MixPad

For details on recording a track in MixPad see the Recording a Clip topic in this manual.

Load an existing clip

To open an existing audio file and add it as a MixPad clip, tap the Clip --> Add Clip command at the menu.

Positioning, Selecting and Copying Clips

When you need to select only one clip, just touch on the title of the clip you want to select.

The easiest way to change the position of a clip is by simply touch down on the caption of the clip window and dragging it to a new location. You can drag a clip to a new location within its current track or to a new track altogether. If the clip is very small then you may find it useful to zoom in closer. This will give you a larger caption bar to work with.

If you need to move a clip by very small amounts, you can zoom in very close so you can see the precise start point of your clip.

For accurate and consistent positioning of clips, you may find it useful to use the grid lines. When the grid lines are turned on, clips will snap to the nearest grid line as you drag them. You can turn the grid lines off altogether by touching the toggle button in the tool panel on the bottom left of the MixPad screen.

You can copy and paste clips to different positions and tracks. Simply select the clip you wish to copy, press Copy command at the menu, select the new location of the clip and press Paste to paste it.

To delete a clip from the work area, just touch Clip-->Delete.

Envelope Fade Points (Automation)

Working with fade points

Automation allows you to vary the volume and pan of each track over time. By default, MixPad displays the automation for the volume of a track. You can swap between Pan, Volume and No automation by using the drop down list on each track's control panel (left of the track). You can add fade points to this envelope by placing the project cursor at the location you would like to have a fade point, and choose 'Add Fade Point' from the 'Edit' menu.

To edit a fade point tap on the point you wish to edit and drag it to a new location.

To remove a fade point, select the fadepoint, then select Edit --> Delete Fade Point from the menu.

Editing Clips

Basic Editing in MixPad

MixPad comes with basic editing tools for arranging your projects. The included editing functions are listed below.

Selected Region

Many of the edit functions below apply to a selected region. To select only one clip, just tap on the title of the clip you want to select. To select a region, touch and hold the clip for one second and drag a selection box around the region.

Undo/Redo/History

To Undo is to restore the file to its state before the previous edit function. This is useful if you want to try an edit or just make a mistake. To undo your last action tap Edit --> Undo. Redo is the reverse of Undo.

History lists all available undo/redos and you can choose one from the list to recover.

Cut

To 'cut' is to delete the selected region but to keep a copy on the clipboard so it can be 'pasted' somewhere else. This is useful when moving parts of the audio around in the file.

To cut, select a region and then tap Edit --> Cut.

Copy

To 'copy' is to make a copy of a selected region to the clipboard so you can paste it in another location. This is useful if you want to duplicate a part of the audio and insert (or mix it) in another file.

Select the region and then tap Edit --> Copy.

Paste

Paste can only be used after you have used the Cut or Copy functions (above) to take a selected region to the clipboard.

The paste function replaces the currently selected region (or inserts at the cursor on the selected track if there is no selection). To replace a selection tap Edit --> Paste. To insert tap on the position within a track and tap Edit --> Paste.

Delete

To delete the clip or selected region tap Edit --> Delete. This is similar to the cut function but a copy is not taken to the clipboard.

Trim

To 'trim' is to cut off the beginning or the end of the file. This is useful when you have just recorded a file but there is silence or noise before the start or after the end.

To trim just select the clip you want to trim, select a region which you want to keep, and then tap Edit --> Trim Region, it will trim all the other parts of selected clips and keep the selected region part.

Silence Region

Silence Region can make the selected region silent. Select the clips you want to edit, then select a region. Then tap the Edit --> Silence Region to make the selected region part of clips silence.

Split Clip

Use this option if you want to quickly split the current clip into two small clips. To do this, tap on the position where you want to split the clip and tap the Edit --> Split Clip.

Takes

Use this command if you want to change a multi-take clip. This command is only available when you select one multi-take clip. You can select which take to use or expand the multi-take clip to separate clips

Rename Clip

Use this command if you want to change name of a clip. This command is only available when you select one clip. Select one clip and then tap the Clip-->Rename Clip

Working with Tracks

MixPad allows you to work on number of tracks. By default, there are two tracks in the project window, but you can add more tracks or delete unneeded.

Adding and removing tracks

To add a new track, tap Track --> Add Track on the menu.

To insert a new track, tap Track --> Insert Track.

To delete a track, select it and then tap Track --> Delete Track.

Select and move tracks

You can move a track by tapping Track --> Move track up or Track --> Move track down command at the menu.

Scroll tracks

You can scroll the visible tracks up and down by using two fingers to drag up and down.

Recording on multiple tracks

To record to many tracks at the same time, tap the Rec button in the control panel of the tracks you want to record to. Tapping the Rec button puts that track into record standby mode. To begin recording all tracks in record mode, tap the main record button in the controls at the bottom of the project window.

Muting and soloing tracks

When you are listening to a MixPad project you may wish to only listen to a few tracks at any one time. For instance, you may wish to listen to a quiet piano piece without hearing a loud drum track. The mute button **M,** which is found in the track control panel, allows you to silence a track. Any track on mute will not play when you play the project. Conversely you can set a track to solo by tapping the **S** button, also found in the track control panel. When one or more tracks are set to solo, only tracks set to solo will be played.

Pan

You can use the pan slider, located in each track control panel, to pan the audio of a track so that it comes out the left channel only, the right channel only, or anywhere in between. The effect of the pan slider on the track will combine with any pan specific fade points you have added to individual clips.

Adding an effect chain

MixPad allows you to build a live effect chain on each track. A live effect chain means that any effect you apply will be applied during playback, which eliminates the need to wait for your audio to render with the effect. To create or edit an effect, simply tap the Fx button in the track control panel. MixPad will present you with a window showing the list of effects currently applied to the selected track. Tap Add to add a new effect, or select an existing effect and tap Edit to edit its properties. Select an effect and tap Remove to remove the effect from the chain.

DB Display

DB display is the colorful meter show the DB level when you recording clips. Every track has a DB display, and you can notice there is a peak red line on db display when you recording. This helps you to know the exact db level, and you can clear the peak db red line just simple tap on db display.

Working with Your Project

A MixPad project is your entire mix. It is all your tracks and all your clips.

Playing

To play a MixPad project move the play cursor to the position you would like to start from and tap the play button at the bottom left side of the screen. Tsp the fast-forward and rewind buttons to search through your audio, or tap the Go to Start or Go to End buttons to quickly jump to the start or end of the selected track.

Saving and Loading

You can save your project as a .mpdp file with a data folder(same name as project file with a .ProjectData postfix) and then load it for use again later. When you save a MixPad project, all audio files are saved into the project's data folder. This means that if you want to move your project to a different computer or location, the things you need to save are the project file and associated data folder.

Timeline Modes

The timeline can be viewed in either minutes and seconds or in bars and beats. To toggle between the two, tap the timeline mode button located above the play controls and just right of the loop button.

Loop Play Mode

To loop a section of your MixPad project over and over, turn loop play mode on by tapping the loop mode button, located on the Controls tab toolbar (also found just above the Play button in the lower left corner of the MixPad window). Next, select the portion of your project you want to loop by tapping and dragging in the timeline area beneath the tracks. When you tap play and the cursor position is before the end of region, MixPad will play from current cursor position and once reaches the end of selected region, the selected region will loop as long as the loop play mode button is activated.

View Grid Lines

Using grid lines will help line up loops and pieces across several tracks. To turn on grid lines, tap the grid lines button to the right of the timeline mode button, just above the play controls. Note that if you change the BPM, you will have to reset the grid lines by turning them off and back on again.

Metronome

Playing a metronome during playback will help you align tracks with the master tempo of the mix. To toggle the metronome on and off tap the metronome button, located below the tracks in the lower left corner of the project window. If you want to use your custom metronome file you have to enable the metronome in the Metronome tab of Settings dialog. See the

Settings ~ Metronome section of this manual for more information on metronome options.

Recording a Clip

Tap the **Rec** button on the track's control panel to place the track into record-standby mode - MixPad will not start recording yet. Check that the audio level meter, located at the bottom of the track control panel, is registering an input. If there is no meter reading here, your audio input devices have not been set up properly.

Finally, press the main record button (found at the bottom with the other project control buttons), and MixPad will start recording.

Punch and Roll

Select a region then choose **Punch and Roll** to record exactly on the selected region.

To start punch and roll recording, select a region and tap the **Record** button from the Recording Tab and select **Punch and Roll**. You can also tap the arrow next to the Record button and select **Punch and Roll**.

There are different record modes in the **Punch and Roll** recording. You can access this by going to the MixPad Settings -> Generals tab then tap the **Punch and Roll Settings…** button

- The Auto mode chooses between "Fixed" and "Flexible" smartly based on the type of audio we are recording over.
- The Fixed mode automatically stops based on the duration of the selected region. This mode is recommended for music audio types.
- The Flexible mode continues recording from the selected region and stops until the user decides to. This mode is recommended for voice audio types like podcasts, etc.

MIDI Playback

MixPad supports the playback of MIDI files. Unlike most digital audio workstations, MixPad allows you to add a MIDI clip to any track - it does not distinguish between a MIDI track and an audio track. This allows you greater flexibility in your mixes.

You add a MIDI clip to your project the same way you add a normal audio clip - either by dragging and dropping or from one of the Load options under the 'Clip' menu.

It is important to note that a MIDI clip is not an audio file. It is just a sequence of instructions which can be sent to a synthesizer which can then convert the instructions into audible notes. As such, if you wish to playback your MIDI clips you will need to select an output device to play back to. If you have an external MIDI synthesizer then this will be your best option because the latency is likely to be very low and you will keep good synchronization between your MIDI clips and audio clips. If you do not have an external MIDI device then you can use the built in Microsoft synthesizer. Please note that the built in synthesizer has a significant lag which can make synchronization difficult.

MIDI files can contain more than one MIDI track. If you load a track into MixPad that contains multiple tracks, MixPad will give you the option to merge all the tracks into one clip, or separate them out in individual clips.

MIDI volume control is a little bit different from audio. Currently, if you want to control one MIDI clip volume, you have to use VST Instrument. Or you can use the master volume control on the right bottom of MixPad , which controls all the MIDI device volume and audio volume. You can't adjust one single MIDI clip volume directly while playing through MIDI device.

MIDI Editing

MixPad allows you to edit MIDI files using a special MIDI editor window. To edit a MIDI clip, simply long press that clip and select 'Edit Clip with MIDI Editor' from popup menu. This will open the clip in the MIDI editor.

From this editor you can do the following:

- Move MIDI notes around - grab the note with the mouse and drag.
- Resize the length of MIDI notes - drag the edges of the note.
- Delete notes - simply select a note by tapping on it and press Delete.
- Property Panel - This contains the 'Notes' and 'Channels' tab. You can show or hide the 'Property Panel' by tapping the 'Property Panel' on the View Tab.
- Add new notes - Drag in the editor where you would like the new note created. The new note's velocity and channel is determined by the settings in the Context Panel's 'Create Notes'.
- Change the properties of notes - simply select the note you would like to change. In the Property Panel you can change the selected notes' property.
- Filter -- you can show/hide notes by channels. If you do not want to see all note on channel 0, then just tap the eye icon for channel 0.
- Copy & Paste -- After selecting a note by tapping on it, tap Copy to save the selection to the clipboard. Then tap on Paste to add those copied notes on to the rollout. They will be added where the cursor on the timeline is located.
- Quantize and Humanize -- Quantize can let you align notes' start, length etc to divisions. And Humanize is the reverse of Quantize which introduces random errors for notes' start position, length etc.
- Events List -- Events list window lists all events including program change, controller change events etc.
- Play MIDI notes in loop - turn on loop mode and drag the timeline to select the region you want to loop in.

Beat Maker

Overview The Beat Maker is a utility which you can use to create beat tracks for your project. It comes with a range of sample drum kits and drum patterns to help get you started.

Getting Started

To get started you will need to download or create a kit. If you do not have any kits installed then MixPad will ask if you want to download a pack. If this is your first time using the Beat Maker then this is highly recommended.

Creating a Pattern

All Beat Maker compositions are built from patterns you add to the project area (the composer). To add a new empty pattern simply double tap on an empty bar in the composer.

Adding a Beat to a Pattern

You can then add beats to the pattern by tapping on one of the bar division lines in the pattern. You can delete the beat by long pressing. Press play to hear the results of your first pattern. The higher a beat is placed in the pattern the louder it will be.

Resizing a Pattern

You can resize a pattern by dragging the small handle on the top right hand corner of the pattern.

Changing a Pattern's Divisions

You can set how many divisions per bar a pattern has by tapping on the up and down arrows at the upper left corner of the pattern.

Joining Patterns

Patterns which are on the same track can be joined by tapping the Join button in the toolbar.

Splitting Patterns

Patterns can be split relative to the cursor by tapping the Split button in the toolbar.

Grouping Patterns

Patterns can be grouped and ungrouped using the Group and Ungroup buttons in the toolbar. When patterns are grouped, then tapping on one pattern will also select any other patterns which it has been grouped with.

Linked Patterns

When you copy and paste a pattern, then those patterns become linked. If you add a beat to one of the patterns then a beat will be added to the other pattern too. This is useful for when you have a repeating theme and you want to make an adjustment that should be replicated throughout the entire composition. You can unlink patterns so that they can be individually edited by clicking the Unlink button in the toolbar.

Locked Patterns

You can lock and unlock patterns by tapping on the padlock icon in the top left hand corner of the pattern. When a pattern is locked, you can't add, edit or delete beats from the pattern. You can however, delete or move the pattern. This helps with preventing accidental changes to patterns you are already happy with.

Using the Pattern Library

MixPad has a pattern library from which you can take frequently used patterns and add them to your composition. To do so, double tap an item in the pattern library and it will be added to your composition. You can add a new pattern to the pattern library by selecting one or more patterns from your composition and then clicking the Add to Library button on the toolbar. MixPad will ask you to give your pattern a name and a category.

Composition Properties

In the top right hand panel, you can adjust the properties of the composition you are working on.

Looping Region

The green bar above the timeline shows the looping region of the composition. You can change this by moving the small handles on either end of the bar or by editing the numbers in the composition properites panel. Looping can be activated and deactivated by clicking the loop button in the playback controls at the bottom of the window.

Changing a Track Sound

Each track is assigned a sound from the active kit. You can change which sound is assigned to the track using the pulldown menu in the track control panel at the left of the track.

Mute and Solo

Tracks can be muted by tapping the M button and soloed by clicking the S button in the track control panel.

Track Volume

You can adjust the track volume by using the slider in the track control panel. This volume will affect the patterns in the corresponding track.

Editing and Creating Kits

MixPad comes with some sample drum kits, but you are able to create your own kits from scratch. To get started click on the Edit Kits button in the toolbar.

Creating a New Kit

To start a new kit tap on the New Kit button in the Kit Editor window. Then you will need to assign sounds to the kit. Each row of the table represents one component of the kit. Using the pulldown list on the left, select the type of component you are adding to the kit. Then select the file path to the audio file which contains the sound. In the 3rd column you can enter a display name for the component.

Saving and Loading Compositions

Once you are happy with your composition you should save it so that you can easily come back and make any changes at a later stage.

Exporting Your Composition

When you are finished composing your beat track then you can tap the Export button in the toolbar and MixPad will render your composition and insert it into your project where you will be able to use it like any other clip you import into MixPad.

Exporting

When you are happy with your mix, you will probably want to export it to an audio file, such as .wav or .mp3, so that it can be played by normal players. To do this in MixPad tap the Export Mix icon on the Home tab. You will see a dialog box called Export Options. Here, you can use the Browse button to select where, and in which format, you would like your mixed project to be saved. Once you have selected the location and file type, you can tap the Settings button which will allow you to define the settings MixPad will use when exporting.

Cloud Support

Dropbox MixPad allows you to directly upload and download your project from your Dropbox account. To upload your project tap Save Project to Dropbox on the toolbar. This will save all files necessary for your MixPad project to be opened on another computer. Tap Download from Dropbox to download a MixPad project already stored in Dropbox.

Effects

MixPad allows you to build a live effect chain on each track. A live effect chain means that any effect you apply will be applied during playback, which eliminates the need to wait for your audio to render with the effect. To create or edit an effect, simply tap the Fx button in the track control panel. MixPad will present you with a window showing the list of effects currently applied to the selected track. Tap Add Effect to add a new effect, the new added effect will be highlighted means it is selected and the checkbox before the effect indicates that it is enabled. Select an existing effect, you can see and edit all the effect's properties. Select an effect and tap Remove Effect to remove the effect from the chain. Untick the checkbox before the effect you can disable the effect but still keep it.

Once you setup the effect chain. you can save the effect chain to a file by tapping the save effects chain and load it into another track through load effects chain.

Effects

Amplify

To 'amplify' is to increase the loudness or volume. The volume is entered as a percentage: 100 being no change, 50 being -6dB softer or 200 being +6dB louder.

Chorus

The chorus sound effect is used to make one voice or instrument sound like 3 voices or instruments by playing the original with variably-delayed and slightly pitch-changed copies of the original.

Note: Chorus is a very useful way to make a mono source sound stereo. You should convert your file to stereo before applying the chorus effect.

Dynamic Range Compressor

A Dynamic Range Compressor limits the volume levels of a sound recording so that it stays within a certain loudness range. An example of where it is used is in TV broadcasting, where it ensures that the volume levels of ads are perceived as being louder than the television program itself, without any change in the actual broadcast volume.

It also has a use for recording audio from one medium to another, where the two mediums are not capable of handling the same range of volume levels (e.g., A CD can handle a much greater range than a cassette tape).

The "Threshold" setting works by detecting when the sound recording volume exceeds a defined decibel level. It then gradually attenuates the sound to bring it down below the dB level, and does it in such a way that the listener will not be aware the attenuation is occurring.

The "Ratio" setting limits the amount the volume level of the recording increases at any one time. If, for example, you wanted the volume levels of a recording to only increase by at most 1/4 of the amount they would normally increase, then this would correspond to a Ratio of 4:1. So if the recording volume level increased by 8dB, then you would only hear a 2dB volume increase.

The "Limit" setting defines at what maximum decibel level the sound recording will be allowed to rise up to. So if, for example, the Limit was set to 0dB, then you will never hear the volume level of the recording get louder than 0dB. The Limit setting has similarities to the Threshold setting, but the main difference is that the Threshold does allow sounds to go above the defined decibel level (for a short time), whereas the Limit does not.

You will find that the minimum Limit volume you can set is the same as the maximum Threshold value. This basically means that, in any situation, the sound will start to attenuate at the threshold level, but will never be heard louder than the limit.

Distortion

While normally we do everything to reduce distortion, sometimes you want to add it. It is popular for use with guitars. The distortion is measured between 0.0 (off) and 1.0 (clipping). The level where it kicks in can also be specified in dB. For a more consistent sound, you should apply Dynamic Range Compression first before you add distortion.

Echo

An echo is a repeat of the sound after a short time. It can sound like the person or instrument is in a large stadium or is shouting between two mountains. Specify the duration and amplitude of the echo. The duration is the length of time after which the sound repeats - usually this is between 1 and 4000ms. The amplitude can be between 1 - 99% (99 being a very loud echo).

Flanger

A Flanger sound effect is created by mixing a slightly delayed signal that is slowly modulated over time with the original. You specify the starting delay time (default 5ms), the frequency of modulation in times per second (default 0.5Hz which is 2 seconds) the depth of modulation (default 50%) and the wet dry gain (100% for wet, 0% for dry).

High-Pass Filter

A high-pass filter (sometimes called a low cut filter) removes all low frequencies below a specified Hz. This is useful if you want to make your recording sound 'clearer' or less 'muddy'. It is very usual to use a high-pass filter of about 300Hz on all voice recordings to improve intelligibility.

Reverb

Reverb is many small randomized reflections of a sound that come after a set time. It is most noticeable in when someone is speaking in a room, hall, etc. When you record in a studio, there is usually very little reverb which can make the recording sound flat. Adding reverb to your tracks can help to make the recording feel more 'live.' The reverb level is the amplitude - 99 is very wet, 0 is dry. The time can be between 100 and 800ms - 200ms sounds like a small room or 800ms sound like a large hall.

Bookmarks

Bookmarks

Bookmarks are locations within your project that you might frequently want to return to. For example, you can use bookmarks to denote different parts of a song you are creating eg. intro, verse, chorus etc. You can also use bookmarks to create working regions.

The easiest way to use bookmarks in MixPad is via the Bookmark Manager dialog. You can access this dialog by selecting View --> Bookmark Manager from the menu. You can also see the bookmarks on the project timeline.

Tap the **Add** button in the Bookmark Manager dialog to create a new bookmark at the current cursor location in your project. You can also fine tune the location of the bookmark by changing the parameters in the dialog box. You can specify the location of the bookmark by entering the time in hours/minutes/seconds or using musical timing notation in the form of bars and beats. You can also specify a color for the bookmark.

To edit a bookmark simply highlight the bookmark and tap the **Edit** button in the Bookmark Manager dialog. Alternatively you can drag the bookmark on the project timeline to a new location.

When you want to return to a bookmark you have created, open the Bookmark Manager dialog, and a list of bookmarks will be shown. Tap the bookmark you want to return to.

To navigate between the bookmarks, you select the bookmark you want to return to in the Bookmark Manager dialog.

If you have more than one bookmark in your project you can quickly select the area between the two bookmarks as a region. Creating a region in this way will allow you to create a section for looping playback and looping recording. Looping recording is a feature which allows you to record multiple takes without having to first stop and restart the recording. To select a region highlight the bookmark which will be the start of the region and tap the **Select Work Region** button in the Bookmark Manager dialog. Remove the work region by tapping the **Clear Work Region** button.

Audio Unit Plugins

Audio Units Effects Plugin and Instrument Plugin Support

MixPad supports Audio Units (AU) plugins. You can use these plugins to enhance MixPad's capabilities. You can use effect plugins to add an effect to one of your tracks, or you can use instrument plugins to synthesize playback of MIDI files. Other plugins may provide some type of visual feedback of the audio signal.

To add a AU effect to one of your tracks, tap the **Fx** button in the track control panel on the left of the track. In the left hand panel of the dialog that opens you should see a sub-heading **Audio Units** . Any valid AU plugins you have installed will be listed underneath this sub-heading. Apple supplied audio units will be always listed underneath.

Playback and Monitoring of MIDI with Audio Unit Instruments

An AU instrument is a special type of plugin which takes MIDI as its input and outputs audio. You can use an AU instrument to playback your MIDI clips and you can also use them to monitor any recordings you make using MIDI hardware. To add MIDI playback or monitoring to a track, open the track options dialog by tapping on the small spanner icon on the track controls. In the dialog that opens you should see a **MIDI Playback Options** option group.

Select the radio button which says "Play MIDI using Audio Unit instrument" and then select which instrument you would like to use.

You can do the same for monitoring MIDI recordings in the **Monitoring Options** option group. Select the radio button which says "Enable MIDI monitoring via Audio Unit Instrument" and select the instrument you would like to use from the associated pull down list.

Once you have selected an AU, you should be able to open your chosen plugins' GUIs via the two small keyboard icons on the control panel of each track. The button on the left opens the AU currently assigned for playback and the button on the right opens the GUI of the AU assigned to monitoring recordings.

Editing - Beat/Tempo Analysis

Overview MixPad can scan music files added to your project and automatically determine the precise tempo and location of the beats.

Tempo Detection

You can configure how the automatic beat analysis is performed by going to the 'Clips' tab of the options dialog and selecting one of the options. By default, MixPad will scan any clips you add to your project in the background and then update the beat information on the clip when the scan is complete. You can also choose to scan the clip before loading or only scan on demand. Full details on these options can be found here.

Analysis will only be successful on source files with a predictable tempo. If MixPad is unable to determine the tempo reliably then it will not show any beat information. However, you can always add the tempo manually using the tempo editor.

If you load a single clip into MixPad, and you have selected the background analysis option, then MixPad will also change the project tempo to match the tempo of the loaded clip.

Tempo Editing

If MixPad is unable to determine the tempo of a clip, or it makes an error, you can manually adjust the tempo using the Tempo Editor. To access the editor, long press on a clip and select 'Edit Tempo and Beat Grid' from the menu.

Details on how to use the Tempo Editor can be found here.

Effects - Amplify

Amplify

To 'amplify' is to increase the loudness or volume of the selected region. To make a part of the recording softer or louder, select it and then use the menu Effects -> Amplify. The volume is entered in percent (100 being no change, 50 being -6dB softer or 200 being +6dB louder).

Effects - Normalize

Normalize

To 'normalize' is to adjust the volume so that the loudest peak is equal to (or a percentage of) the maximum signal that can be used in digital audio. Usually you normalize files to 100% as the last stage in production to make it the loudest possible without distortion. Another reason to normalize is to have multiple tracks sound equally loud, or to have equal average loudness.

The 'Peak' normalization method finds the sample of the greatest magnitude within the file. Normalization is then done with this value as the peak. With the Normalize Peak Level set to 100% (0dB), the whole file will be amplified so that the peak reaches 0dB.

The 'Average Loudness (RMS)' normalization method normalizes according to the file's average loudness, or volume. Multiple files normalized to the same peak level using this method will have equal average loudness. The 'Normalize Peak Level' for this method should be set much lower than for the Peak method, because the average loudness will always be lower than the peak sample.

The 'Peak Loudness (RMS)' normalization method attempts to normalize according to how loud the loudest part of the file will sound. This is the best method to use to make multiple tracks sound equally loud. As with Average Loudness, the 'Normalize Peak Level' for this method should be set lower than for the Peak method, because the peak loudness is lower than the peak sample. The actual algorithm used takes the RMS of each 50ms window in the file, ranks the windows from loudest to quietest, and then takes the 95th percentile of these as the 'peak'. Note that no adjustment is made for humans' differing perception of different frequencies.

The Perceived Loudness (dBA) normalization method uses A-weighted decibels which is an expression of the relative loudness of sounds in air as perceived by the human ear. Normalization is made because the human ear is less sensitive at low audio frequencies, especially below 1000 Hz, than at high audio frequencies.

The Integrated Loudness (EBU) normalization method uses the R-128, an international standard for loudness normalisation and maximum level of audio signals. Several platforms use this such as Spotify (-14dB), YouTube (-14dB), Apple Music (-16dB), Facebook (-18dB), and TV / Radio (-23dB). See EBU R-128 for more details.

Effects - Compressor

Dynamic Range Compressor

A dynamic range compressor limits the volume levels of a sound recording so that it stays within a certain loudness range.

An example of where it is used is in TV broadcasting, where it ensures that the volume levels of ads are perceived as being louder than the television program itself (without any change in the actual broadcast volume).

It also has a use for recording audio from one medium to another, where the two mediums are not capable of handling the same range of volume levels (e.g. A CD can handle a much greater range than a cassette tape).

The Dynamic Range Compressor dialog has two tabs: "Simple" and "Graphic". Changing settings on the Simple tab will also change the graph on the Graphic tab, but not vice versa as the graph allows more control. There is also an "Advanced Compressor Settings" dialog for adjusting more advanced features.

The Simple Tab

The "Simple" tab of the Dynamic Range Compressor dialog contains settings called "Limiter", "Compressor", and "Noise Gate". While these sound like three different things, they are more accurately viewed as three different ways of using the dynamic range compressor.

The "Limiter" defines the maximum decibel level that the sound recording will be allowed to rise up to. So if, for example, the Limiter Threshold was set to -2dB, then you would never hear the volume level of the recording get louder than -2dB. Any signal over the limiter threshold would be clipped, which would probably cause distortion. Note that setting the Limiter Threshold to 0dB effectively turns the limiter off, because 0dB represents the loudest signal possible in a digital recording.

The "Compressor" reduces the volume of any sound which exceeds its "Threshold" setting. When a signal exceeds the threshold, the compressor gradually attenuates the sound to bring it down below the dB level, and does it in such a way that the listener will not be aware the attenuation is occurring. The compressor differs from the limiter in that the compressor does allow sounds to go above its threshold (for a short time), whereas the limiter does not.

The "Ratio" setting defines the ratio of the reduction in volume of sounds which exceed the compressor threshold. For example, if the ratio is 4:1 and the volume exceeds the threshold by 4dB, then the volume will be reduced to only exceed the threshold by 1dB. Note that a ratio of 1:1 means that there will be no change in volume; it effectively turns the compressor off.

The "Noise Gate" works similarly to the Compressor, except that is reduces the volume of sound below its Threshold. This can be useful for reducing or removing softer background noise from a recording.

You will find that the maximum Compressor Threshold you can set is the same as the current Limiter Threshold value. This basically means that, in any situation, the sound will start to attenuate at the Compressor Threshold, but will never be heard louder than the Limiter Threshold. Similarly, the maximum Noise Gate Threshold you can set is the same as the current Compressor Threshold.

The Graphic Tab

The "Graphic" tab of the Dynamic Range Compressor dialog shows a graph which represents the relationship between input and output volumes. The horizontal axis shows input volumes in dB from -60dB to 0dB. The vertical axis shows output volumes on the same scale. The graph will be changed by changes to settings on the Simple tab, but changes to the graph will not be reflected on the Simple tab, because it is possible to represent a wider variety of settings on the graph than is possible in the controls on the Simple tab. When the dynamic range compressor is applied it will use the settings from the Graphic tab.

To change the graph, click and drag the black vertex markers, or click anywhere else to create a new vertex. To remove a vertex, right-click on it.

Advanced Compressor Settings

Clicking on the "Advanced" button in the Dynamic Range Compressor dialog will open the Advanced Compressor Settings dialog. In it are controls for the following properties of the compressor:

-Input Level Sensing - Peak or RMS:

-This controls how the compressor determines the audio level. "Peak" sensing looks at the highest point in the window of audio which it examines. It will almost always give a higher reading than "RMS" sensing, which uses an average, or Root Mean Square of the window to determine the audio level. RMS sensing more closely corresponds to the audio level which a human listener would perceive.

-Compressor Response:

-

-Attack:

-The time (between 0 and 1000 milliseconds) that it will take to apply the gain adjustment. The total gain adjustment required will be gradually introduced over this period.

-Release:

-The time (between 0 and 5000 milliseconds) that it will take to remove the gain adjustment once gain adjustment is no longer needed. This is the opposite of attack.

-WindowLength:

-The length (between 10 and 50 milliseconds) of the window to use when calculating the current audio level. A shorter window responds to level changes more rapidly, but anything less than 50ms will start to respond inconsistently to bass, since 50ms (20Hz) is the wavelength of the lowest human-audible sound.

-LookAhead:

-How far ahead (between 0 and 100 milliseconds) to look at the input level when determining the output gain adjustment. This can cause the compressor to start responding to a change in volume before it happens. If this value is the same as the attack time, then the full gain adjustment could be made by the time the louder signal is reached.

-Side-Chain Equalizer:

-This determines how strongly the compressor should weight different audio frequencies when determining the input level. For example, to compress only when there is a loud bass sound, turn the Bass level up and/or reduce the MidRange and High levels.

-Auto Makeup Gain:

-When this option is selected compressor automatically makes up the gain lost in the compression process. Select this option if you want to amplify the compressor output to the original audio level.

Dynamic Range Compressor Presets

The following presets have been defined for your convenience. A preset will change the settings of the dynamic range compressor, after which you can make further adjustments if necessary. The presets are: -Default:

-Pressing the "Default" button will cause the compressor to have no effect. It sets the output levels to be exactly the same as the input levels, and also resets the advanced settings to their defaults.

-Fast Compressor:

-This compression preset will cause any spikes over -20dB to be rapidly reduced, but will not cause distortion. It uses peak input level sensing and a fast attack, which will reduce the volume of transient sounds (such as a snare drum hit), but may also change their characteristic sound. Compare this with the Smooth Compressor preset below.

-Smooth Compressor:

-This preset reduces the volume more gradually when the signal climbs above -20dB. The slow attack time will mean that transients (such as snare drum hits) will not be changed, or if they are then they will be uniformly reduced, thus their characteristic sound will not be significantly altered.

-Heavy Compressor:

-This preset uses a lot of compression whenever the average volume climbs over -30dB, resulting in a very uniform dynamic range. This can be useful for making the quieter parts of music with a large dynamic range (such as classical music) easier to hear in noisier environments, such as in a car or a restaurant.

-Hard Limit:

-This preset does not allow any sounds to exceed -12dB. This may cause distortion due to clipping in some tracks.

-Soft Limit:

-This limit allows short spikes over -6dB, but will prevent longer durations of audio over this threshold.

-Noise Gate:

-This will remove soft sounds from a track. This can be useful for removing the crackle of a record player during silences, or background noises in a dictation.

Effects - Equalizer

Equalizer

An equalizer changes the frequency response of a signal so it has different tonal qualities.

After you select Effects menu -> Equalizer you will see a dialog containing three different Equalizer representations. Use the tabs at the top to select between the Visual Equalizer, Graphic and Parametric Equalizer views.

Visual Equalizer

Tap on any point to create a new band point. To remove a band point long press on it. To assist you with shaping the Equalizer graph in the way you want, there is a preset list that displays the most common sorts of filters used in the Equalizer graph. You can choose any preset filter from the list and then manipulate the filter to achieve the effect you desire. The list of filters to choose from and how you can shape them are explained below. Note that all fields where a frequency value is entered can have a maximum value of 20000 (Hertz).

Graphic Equalizer

The Graphic Equalizer uses discrete sliders to set the gain or attenuation of a signal at a particular frequency. You can select how many sliders you would like to manipulate by entering a value between 3 and 20 in the box at the top of the display. When you change the number of sliders you would like to utilize, the frequencies are automatically allocated to best span the audible frequency range from 20Hz to 20kHz. Selecting presets allows you to easily configure common filters such as low pass or high pass. Note that when you change the Graphic Equalizer, the Visual and Parametric Equalizer views are not changed, as the changes in the three views are not compatible.

Parametric Equalizer

The Parametric Equalizer is similar to the Graphic Equalizer, but with more control. Here you can adjust the frequency and bandwidth of the individual sliders by tapping on the frequency or Q values below each slider. Frequency must be set between 20Hz and 20,000 Hz. The Q parameter must be set between 0.05 and 20. A higher Q causes the gain or attenuation peak at the frequency to be much sharper, and therefore less likely to impact adjacent frequency content, while a lower Q applies the modification more smoothly across the frequency spectrum.

- Band Pass Filter
- Keeps only those frequencies in the audio between a certain range. - Start Frequency
- The lower cutoff frequency value, in Hertz.
- End Frequency
- The upper cutoff frequency value, in Hertz.
- Slope Length
- The width of the slope extending from the lower and upper cutoff points, in Hertz.
- Amplitude
- The degree that the frequencies outside the cutoff range are suppressed. 6dB means the volume is reduced to one-half, 12dB means the volume is reduced to one-quarter. Maximum value is 60dB.

- Band Stop/Cut Filter
- Keeps all frequencies in the audio except those between a certain range.
- Start Frequency
- The lower stop frequency, in Hertz.
- End Frequency
- The upper stop frequency, in Hertz.
- Slope Length

-The width of the slope extending from the lower and upper stop points, in Hertz.
-Rejection
-The degree that the frequencies inside the stop range are suppressed. 6dB means the volume is reduced to one-half, 12dB means the volume is reduced to one-quarter. Maximum value is 60dB.

-High Pass Filter
-Keeps only those frequencies in the audio above a certain value.
-Pass Frequency
-The point at which all frequencies above are to be kept, in Hertz.
-Slope Length
-The width of the slope extending from the pass frequency, in Hertz.
-Low Pass Filter
-Keeps only those frequencies in the audio below a certain value.
-Pass Frequency
-The point at which all frequencies below are to be kept, in Hertz.
-Slope Length
-The width of the slope extending from the pass frequency, in Hertz.
-Notch Filter
-Attenuates the frequencies in the specified range to very low levels and passes all other frequencies unaltered. There is no slope - frequencies are either attenuated or not.

-Start Frequency
-The lower cutoff frequency value, in Hertz.
-End Frequency
-The upper cutoff frequency value, in Hertz.
-Boost Filter
-Either attenuates or boosts frequencies in the specified range and passes all others unaltered.

-Start Frequency
-The lower boost/cut frequency value, in Hertz.
-End Frequency
-The upper boost/cut frequency value, in Hertz.
-Slope Length
-The width of the slope extending from the lower and upper boost/cut points, in Hertz.
-Amplitude
-The degree that the frequencies inside the boost/cut range are either boosted or cut. 6dB means the volume is boosted to twice the original amount, and 12dB means the volume is is boosted to four times the original amount. 20dB.

-High Pass Shelf Filter
-Attenuates signals of frequencies below the cut frequency and passes all others unaltered.
-Start Frequency
-The lower cut frequency value, in Hertz.
-Slope
-The width of the slope extending from the lower and upper cut points, in Hertz.
-Rejection
-The degree that the frequencies inside the cut range are cut. 6dB means the volume is attenuated to about half the original level, and 12dB means the volume is attenuated to about a quarter of the original level.

-Low Pass Shelf Filter
-Attenuates signals of frequencies above the cut frequency and passes all others unaltered.
-Start Frequency

- The lower cut frequency value, in Hertz.
- Slope
- The width of the slope extending from the lower and upper cut points, in Hertz.
- Rejection
- The degree that the frequencies inside the cut range are cut. 6dB means the volume is attenuated to about half the original level, and 12dB means the volume is attenuated to about a quarter of the original level.

If you are using the equalizer simply to drop lower frequencies, you should always try the High Pass filter first (Effects menu -> High Pass Filter), because it is better and faster for very low frequencies.

Effects - Envelope

Envelope

The 'envelope' is the change in volume of the selected region over time. This can be used to make fine adjustments to the volume over time or even more crude changes like fade in or fade out.

Select the region you want to change the volume over and use the Tools tab -> Envelope. Tap and drag on any point to adjust its volume (long press removes the point). Tap the Set Flat button to reset the volume and remove extra volume points.

Effects - Stereo Pan

Stereo Pan

The stereo pan effect allows you to change how loud the sound is that comes out the left or right speaker. For example if you had a stereo recording with all the sound coming out of only one speaker, you could use the pan effect to "center" the sound yourself. You can also make a centered sound change move one from speaker to the other as the sound file plays.

Select the region you want to change the pan for and choose Effects -> Stereo Pan. Click on a point and move it upwards for an increase in volume on the left speaker, or move it downwards for an increase in volume on the right speaker.

Please note the stereo pan effect only works on stereo files. If your file is not stereo you must first convert it to stereo by choosing Edit -> Convert Channels -> Stereo.

Effects - Echo

Echo

An echo is a repeat of the sound after a short time (usually 400 - 1000ms). It sounds a bit like the person is in a large stadium or is shouting between two mountains.

To add echo select the region and use the menu Effects -> Echo then specify the duration and amplitude of the echo. The duration is the length of time after which the sound repeats - usually this is between 400 and 1000ms. The amplitude can be between 1 - 99% (99 being a very loud echo).

Effects - Pitch Shifter

Pitch Shifter

Pitch Shifter is a sound effect that raises or lowers the pitch of audio signals. You can adjust pitch shifter speed by dragging the slider in the settings.

Effects - Reverb

Reverb

Reverb is many small reflections of the sound that come after a set time. It usually occurs when someone is speaking in a room, hall, etc. More reverb is called wet, no reverb is called dry. When you select the reverb effect, you will see a dialog with two tabs.

Simple

The first tab of the reverb effect allows you to adjust the reverb level and time. The reverb level is the amplitude - 99 is very wet, 0 is dry. The time can be between 100 and 800ms - 200ms sounds like a small room or 800ms a large hall. If you add too much reverb it can sound like the person is in a pipe or in the bathroom.

The Simple tab also includes preset options to choose from, depending on how large the space being simulated is. Click the play button at the bottom of the tab to preview the reverb effect on your audio.

Room Design

The second tab of the reverb effect allows you to specify the dimensions of a room, the position of the source and listener, and the room absorption with preset options for the materials that make up the walls, floor and ceiling of the room. Click the play button at the bottom of the tab to preview the reverb settings on your audio.

Effects - Phaser

Phaser

The phaser sound effect is created by mixing a slightly delayed signal with the original. You can set the delay in ms (default 5ms) and the wet dry gain in percent. 100% is wet. 0% is off/dry.

Effects - Flanger

Flanger

A Flanger sound effect is similar to the phaser except that the delay is slowly modulated over time. You specify the starting delay time (default 5ms), the frequency of modulation in times per second (default 0.5Hz which is 2 seconds) the depth of modulation (default 70%) and the wet dry gain (100% for wet, 0% for dry, default 60%).

Effects - Vibrato

Vibrato

The vibrato sound effect is a pulsating of the pitch at a depth and frequency specified by the user. The higher the Frequency (Hz) set, the more often the pulses will be heard, and the higher the Depth (semitones), the wider the fluctuation in pitch will be.

Effects - Tremolo

Tremolo

The tremolo sound effect is similar to the vibrato effect, except that the amplitude pulsates rather than the pitch. The higher the Frequency (Hz) set, the more often the pulsation will be heard, and the higher the Depth (%), the deeper the fluctuation in volume.

Effects - Doppler

Doppler

The doppler effect simulates the sound of a passing vehicle, which has a high pitch while approaching, shifting to a low pitch when traveling away from the listener. Specify the Velocity (in km/h) of the passing source; a higher velocity will result in a higher starting pitch and lower ending pitch. Adjust the Listener Horizontal and Vertical Positions to indicate the listener's horizontal and/or vertical position to the passing source; play around with the values to achieve different combinations of pitch.

Effects - Wah-Wah

Wah-Wah

As the name suggests, the effect modulates a specified frequency band within the sample, which results in the characteristic "Wah wah" sound. The effect is a bandpass filter with its center frequency (not to be confused with the center frequency parameter, below) alternating between a min frequency and max frequency (specified by the center frequency and depth parameters) and from max frequency to min frequency. The frequency of alternating direction is represented as a triangular wave with a frequency specified by the wah frequency parameter.

Resonance: also known as Q or emphasis, this parameter controls the resonant peak of the bandpass filter. This value determines the sharpness of the wah-wah effect. Higher values produce more resonant/peaky tones.

Depth: this parameter determines the frequency range swept by the bandpass filter. Its range is specified as a percentage of the range (0 to center frequency). If the value of the percentage of the range (0, center frequency) is specified as X, the min and max frequencies are (center frequency - X) and (center frequency + X).

Center Frequency: This parameter is the center frequency of the bandpass filter sweep, and is used to determine the min and max frequencies as mentioned above.

Wah Frequency: This is the frequency of alternating the direction of the sweep, or the frequency of the wah-wah sound. It is the frequency of the triangular wave described above.

Effects - Chorus

Chorus

The chorus sound effect is used to make one voice or one instrument sound like 3 voices or instruments by playing the original with variably delayed and slightly pitch changed copies of the original.

Note: Chorus is a very useful way to make a mono source sound more stereo. You should convert your file to stereo first before using Chorus.

Effects - Distortion

Distortion

While normally we do everything to reduce distortion, sometimes you want to add it. It is popular for use with guitars. The distortion is measured between 0.0 (off) and 1.0 (clipping). You also specify the level where it kicks in in dB.

For a more consistent sound, you should apply Dynamic Range Compression first before you add distortion.

Effects - AM Radio Effect

AM Radio

This simulates an AM Radio. We have made it accurately simulate a 'good' AM radio. To make it worse, apply the effect twice. For a really bad sound, paste mix some soft white noise (use the Tone Generator tool) to simulate bad reception.

Effects - Telephone Effect

Telephone

This simulates the audio down a telephone line. It simulates a 'good' telephone line. To make it worse apply the effect twice and paste mix soft white noise.

Effects - Reverse

Reverse

This effect reverses the selection in the same way playing a record or tape backwards would.

Effects - Fading

Fade In

To fade in use the menu Effects -> Fade In.

Fade Out

To fade out use the menu Effects -> Fade Out.

Fade Out and Trim

The fade out and trim option is a combined function which fades out over the selection then marks the end of the selection as the end of the file. This is frequently used at the end of music tracks.

CrossFade

The CrossFade tool allows you to mix together voice and music in a variety of different ways. You can, for example:

- Fade out a music track while fading in another track,
- fade out a music track and cue in a voice track at full volume (or vice versa), or
- overlay the end of one voice track with the start of another track.

To use the tool, first select the region of audio you want to perform the crossfade on. If you want to crossfade between two files, you must combine the two files together first into one file. Next, go to Effects menu -> CrossFade. A window will appear, showing a graph and a number of data fields.

The graph is divided into two sections, the top section shows the fading in part of the audio, the bottom shows the fading out. The area that the crossfade is to be performed on is highlighted in blue, and surrounded by markers showing the start and end of the crossfade region. There is a one second portion of the waveform on either side of the highlighted section, which is there to provide a better view of the crossfade.

If you hover your mouse over any part of the graph, you can see what parts of the graph correspond to what time in the audio waveform.

The data fields work as follows:

-Start and End Selected Positions

-Tells you the start and end times of the audio you selected in the waveform. **Note:** These times do NOT correspond to the start and end times you see in the graph window! Read on!

-Gap Time

-This says how long the crossfade region will be, in milliseconds. This time may be be modified when the crossfade is performed, if the fade in and fade out times are larger than this value. **Note:** If this time is shorter than the audio you selected then the middle of the selected audio will be lost as a result of the crossfade.

-Fade In Time

-The length of time to fade in the end of the selected audio. For example, if you select 5000ms of audio and a Fade In Time of 1000ms, then the last 1000ms of your selection will fade in over the last 1000ms of the crossfade.

-Fade Out Time

-The length of time to fade out the beginning of the selected audio.

-Fade Type

-The shape of the transition over time of the fade in or out. You can select from Linear, Logarithmic, Expoential or Sinusoidal fade types.

- Linear - Volume change will occur evenly over time.
- Exponential - Volume change will start slowly to begin with and then rapidly become faster towards the end.

- Sinusoidal - Volume change will start slowly, then increase rapidly, then slow down again towards the end.
- Logarithmic - Volume change will start rapidly and slow down towards the end.

So with the above information in mind, the crossfade will work as follows:

1. A Fade-Out buffer will be created with a length of the Gap Time. At the start of the buffer will be the start of your audio selection, fading out over the Fade Out Time.
2. A Fade-In buffer will be created with a length of the Gap Time. At the end of the buffer will be the end of your audio selection, fading in over the Fade In Time.
3. The Fade-In and Fade-Out buffers will be mixed together, and replace your audio selection.

Effects - Speed and Pitch Changing

Simple Speed and Pitch Change

This plays the recording faster or slower which in turn increases or decreases the pitch too. This function is useful to correct slow or fast tapes.

Speed Change

Normal speed changes (i.e. "Simple Speed and Pitch Change" above) changes the pitch in proportion to the speed. If you want to change the speed but keep the pitch the same use this function. Speed can change the duration of the audio. The time duration (in seconds) can also be adjusted using this effect.

Pitch Change

This changes the pitch of the recording without changing the speed (i.e. the converse of the above). Change of semitones can also be adjusted using this effect

Pitch Speed Profile

This allows you to specify how much to change pitch, speed, or pitch and speed at any point in the file, using a graph.

Effects - Reduce Vocals

Reduce Vocals

If you want to reduce the vocals from a music track you can use this effect. This effect will attempt to identify the voice in the left-to-right spectrum of a stereo recording and remove it. The recording must be stereo (from an original stereo source like a CD - simply converting a file to stereo will not work). It will also remove any instruments near the voice in the stereo spectrum.

Perform Only Simple Channel Subtraction - This option works best if the vocals are recorded in the center of the stereo recording with instruments spatially separated. If the stereo is simply a mono copied to the left and right channels, the lack of spatial separation will result in silence rather than a relative reduction in vocals.

Note: it is impossible to remove the vocals perfectly without the original mix track. You will notice some instruments might be removed too and some vocal remain. The effect will also not work on some files which have previously encoded in a highly compressed form like mp3 (because this remove some stereo depth).

Effects - Anonymous Effect

Anonymous Effect

Use this effect to make a voice recording sound anonymous or unrecognizable. Use one of the default presets of the effect or adjust the Pitch, Speed and Distortion amounts to create your own Anonymous voice effect. Use the Pitch Change to make the sound rough or sharp. Use the Speed Change to make the recording faster or slower. Add distortion to increase the white noise in the voice and neutralize background sounds.

Press the Play button to listen to the recording with the effects applied as you make adjustments. Press the Bypass button to listen to the original file without the effect. When you are satisfied with the result, press the Apply button to add the effect to your audio file.

Use the Save Preset… button to create your own Preset based on your preferred settings. Delete a Preset by selecting it in the list, then clicking the Delete Preset button.

Note: This effect can only make the voice unrecognizable to the human ear. It is possible that digitally, the voice file could be returned to its original sound. For a more secure method, use Navigate Speech on the tools tab to convert the recording to text. Then use the Text to Speech tool to have the text recorded to a voice file by your speech engine.

Effects - Voice Change

Voice Change

The Voice Changer allows vocal distortion by changing pitch, semitones, cents, and timbre, by modulating tone, and adding whisper/noise to the voice.

Effects - Pitch Correction Window

Pitch correction is a powerful tool that can be used to make minor adjustments to a voice's pitch. Load a voice clip, select the part of the clip that needs to be adjusted, tap the **Pitch Correction** button on the toolbar to open the pitch correction window.

In the Pitch Correction window, you can see the voice in notes. If needed, add some node points to the pitch line by tapping it, and drag the points to increase or decrease the pitch at the desired places. Remove a node point by long-pressing it. Changes will be applied before playback or when the dialog is being closed.

Use **Automatic Correction** in the toolbar to add fade points at both ends of all detected notes and move the notes to the nearest pitch. **Manual Correction** will also place fade points on both ends of detected notes but it will ask you to provide the key from which the notes will be snapped to.

Audio Cleanup - Noise Reduction

There are two ways of reducing noise. The slow but accurate "Spectral Subtraction" method - usually used where noise is really a problem - and the fast "Multiband Noise Gates" method - usually just automatically on batch voice recording jobs.

Sometimes using both (spectral always must be first) then multiband gates works very well.

- Spectral Subtraction
- - Automatic Method
- This approach will automatically estimate what is noise and what is not. It usually works well on voice and is nice and easy to use; just select the region and apply the effect.
- Manual Method
- To use this you must:

1. Select a short part of 'noise only'. Usually this is from a gap in the audio.
2. Select Effects -> CleanUp -> Noise Reduction -> "Grab noise sample from selected area for spectral subtraction".
3. Select the entire file.
4. Select Effects -> CleanUp -> Noise Reduction -> "Spectral subtraction based on noise sample".

- Multi-Band Noise Gate
- Multi-Band Noise Gate removes all the audio data below a given threshold in the audio file. A good noise threshold for most audio files is usually between -30dB and -20dB.

Noise Gate

A noise gate is a filter which controls the volume of an audio signal. Any part of your audio which is below the Threshold will be attenuated by the amount you specify.

- Threshold
- Audio falling below this threshold will be attenuated.
- Hold
- The period of time (in milliseconds) to wait before applying the attenuation.
- Release
- The period of time (in milliseconds) taken to fully apply the attenuation.
- Attack
- The period of time (in milliseconds) taken to fully remove the attenuation.
- Attenuation
- The amount to attenuate the audio signal when it falls below the threshold.

Noise Removal Wizard

An easy-to-use noise removal wizard can be found on the Tools tab. The wizard guides you through choosing the best parameters based on noise type description, then applies the noise reduction. To use the wizard, tap the Noise Removal button on the Tools tab.

Audio Cleanup - Click/Pop Removal

Auto Click/Pop Removal

This tool allows you to apply a repair of a single click/pop artifact. To use it properly, you must zoom right in to the artifact and select a small region around it. Then select Tools menu -> Auto Click/Pop Removal. The repair will be performed straight away.

Parametric Click/Pop Removal

This tool is designed to remove click and pop sounds from recordings. It is ideal for those who have recorded music onto their computer from LP records and want to repair any defects caused by dust and scratches on the vinyl.

To use the tool, click Tools menu -> Parametric Click/Pop Removal. In the window that appears, you can configure settings for the following fields:

- Click Sensitivity
- This is the degree of aggressiveness (as a percentage) that will be applied by the tool when searching for click and pop artifacts. If you don't know what to enter, you can start by leaving it at 50%. The more a piece of audio is damaged, the higher you may have to set it. Moderately damaged audio can require settings of 60% - 80%. Be careful though - if you set it too high, the tool will start thinking parts of the audio are actually clicks/pops. If you set it too low of course, the tool will think some clicks/pops are part of the audio. Try experimenting to find the right value, and note that the level you apply to one file may be different to the level you apply in another file.

- Maximum Click Length
- This is the maximum length that a click lasts in your audio, in milliseconds. As a general guide, use 450ms if you don't know what to enter. 350ms is appropriate for audio with only small amounts of defects, whereas 550ms or 650ms is appropriate for audio with lots of defects.

Audio Cleanup - High-Pass Filter

High-Pass Filter

A high-pass filter (sometimes called a low cut filter) removes all low frequencies below a specified Hz. This is useful if you want to make your recording sound 'clearer' or less 'muddy'. It is very usual to use a high-pass filter of about 250Hz on all voice recordings to improve intelligibility.

Audio Cleanup - Low-Pass Filter

Low-Pass Filter

A low-pass filter removes all high frequencies above a specified Hz. This is useful if you want to make your recording sound 'clearer'. It is very usual to use a low-pass filter of about 1600Hz on all voice recordings to improve intelligibility.

Audio Cleanup - Band-Pass Filter

Band-Pass Filter

A band-pass filter removes both low and high frequencies below and above a specified Hz respectively. This is useful if you want to highlight a specifc range of frequency to make vocals or instruments more prominent.

Audio Cleanup - Automatic Gain Control

Automatic Gain Control

Automatically adjusts the amplitude to compensate for variations at the input, in order maintain an suitable amplitude at the output. Changes the level of the file up and down over periods of time (time frame: 5 seconds)

De-esser is applied to avoid the AGC increasing "S" sounds and making them too loud (default for voice: 2200Hz).

High-Pass filter is applied to avoid low frequency bass making the AGC reduce volume (default for voice 450Hz).

Other features which can be used with AGC:

- Normalization – changes the level of the whole file equally (time frame: whole file duration)
- Dynamic Range Compressor – instantaneous regulation of volume so softer sounds are increased and louder sounds are decreased (time frame: 100ms)

Audio Cleanup - DC Offset Correction

DC Offset Correction

Often when you record audio using bad electronics the recording has a constant 'DC' level throughout the file. Because the ear cannot hear this you will not notice it until you attempt to edit in other audio when you can hear horrible clicks. If you think this is the problem you can run DC Offset Correction over the entire recording before you begin to edit. Another (and possibly better) way to deal with this problem is to run a high pass filter (say at 50Hz) over the recording.

Audio Cleanup - De-esser

De-esser

Reduce the excessive prominence of sibilant consonants, such as the sounds normally represented in English by "s", "z", "ch", "j" and "sh".

To use the tool, click Tools menu -> De-esser.

NCH Sound Library - NCH Sound Library

The **NCH Sound Library** is a collection of thousands of royalty-free sound effects that can be added to your project.

Once you have opened the library, you'll see the following:

Folder Tree

On the left hand side, each folder represents a category of sounds. Expand a folder to either see its subfolders or a list of sounds it contains.

Sound List

On the right hand side, all the sounds in the currently selected category are listed. This will be empty until a category is selected.

Preview Sound

Select a sound in the list then tap the **Play** button to hear it. When you have finished, tap **Stop** .

Download

Select a sound in the list then tap the **Download** button to download the sound (if it hasn't already been downloaded).

Settings - General

Default Playback Sample Rate

Select the sample rate you would like to use as the default for any new project while playing back using MixPad. If you are aiming for CD quality then choose 44100Hz. Use "Auto" to let Mixpad choose the best sample for you. To change the options for your current project, please use "Menu->File->Project Options...".

Default Recording Sample Rate

Select the sample rate you would like to use as the default for any new project while recording a sound clip.

Match content sample rate

If this feature is active and content is added that has a higher sample rate than the selected recording sample rate, the recording sample rate will be adjusted to the highest of any content.

Offset recordings for latency

If the recordings you are making with MixPad are not perfectly synchronized with each other, then you may have hardware latency issues. You can specify an offset in milliseconds here which will automatically adjust your recordings.

Replace underlay clip(s) when recording with new clip

If you are recording over an underlay clip, the default setting is to replace that clip. To change this, go to MixPad Settings->General Audio Recording.

Replace underlay clip(s) - The new recording will replace the underlay clip(s).

Note: If this option is unchecked, will keep the underlay clip(s). The new recording will be placed on top of the underlay clip(s).

Settings - Messages

Warn when deleting a clip

Check this box if you would like to be warned whenever you try to delete a clip from the project window.

Warn when deleting a track

Check this box if you would like to be warned whenever you try to delete a track from the project window.

Ask to render MIDI files when panning

Check this box if you would like to be prompted to render your MIDI clips when you change the pan on a track which contains MIDI clips.

Warn when moving locked clips

Check this box if you would like to be warned whenever you try to move locked clips.

Tell me when creating a beat clip was successful

Check this box if you would like to be notified when you successfully create a beat clip using the beat maker.

Tell me when one or more Beat Maker tracks have no sound assigned

Check this box if you would like to be notified when one or more Beat Maker tracks have no sound assigned.

Offer to load files in 'mix-tape' formation

Check this box if you want to be offered the option of loading files end to end with a crossfade between them.

When project export is finished

Choose the action you want to happen after successfully exporting a project.

Settings - Metronome

Choose a custom metronome sound

Check this box if you would like to choose your own metronome sound. Otherwise, MixPad will use the default metronome sound.

Remove silence from custom sound

Check this box if you want MixPad to automatically trim silence found at the beginning of your custom beat sound.

Use the metronome volume slider to control how loud the metronome sound is played when it is enabled.

Pre-roll settings

Only play metronome during record pre-roll

Select this box to have the metronome only make a sound while it is in the pre-roll phase. This can help if you want to be counted into a recording but then would prefer time based on other already recorded tracks.

Enable pre-roll

Check this box to have the current pre-roll settings apply to your project

Playing pre-roll (beats):

Configure this value from 0 to 280 to decide how many beats before playback. 0 means start playback immediately.

Recording pre-roll (beats):

Configure this value between 0 and 280 to decide how many beats to wait before recording. 0 means start recording immediately.

Subdivisions:

Configure how many subdivisions are between beats. At the end of each subdivision, sub-beat sound will be heard.

Beats Per Bar:

Configure the beats per bar of the project by selecting one entry from this pull down list.

Settings - Clips

Beat Analysis

MixPad can analyze clips that you load for beat information. Beats are indicated on the clip waveform by a vertical line. There are several options available to you for how this analysis is performed:

1. Automatically (high priority). This option will perform a complete beat analysis before the clip is loaded and displayed on the project window.
2. Automatically (low priority). This option will load the clip immediately without beat information and then perform the beat analysis as a background task. Once the task is complete, the beat information will be added to the clip display.
3. Manually. This option allows you to choose to analyze beat information only when you need it. You can do this by long pressing on a clip and selecting "Automatic beat analysis".

Automatic beat analysis

Automatically analyze beats (background)

Choose this option if you would like all clips loaded to your project to be analyzed for beat information. This analysis will occur in the background and beat information will be displayed once the process is complete.

Manual beat analysis (via the long press menu)

Choose this option if you want to choose when your clips analyzed. You can do this via the long press menu of the clip.

Screen References - Stretch or Shrink Clips

Use this feature to alter the duration of the selected clip without trimming any audio. MixPad will stretch or compress the clip to meet the specified duration. You can also use the *Keep Pitch Constant* option to avoid changing the pitch of the audio.

Screen References - Color Picker Dialog

Color Picker Dialog

The color picker is divided into multiple parts, with several ways available to choose a color:

Visually Adjusting Color

The top half of the dialog consists of 2 colored boxes, one that lets you pick any point in a 256 x 256 square and the other that works as a long slider. Which colors are displayed in these two boxes depends on which radio button is selected (see below), but in all cases the position of the slider will change the colors available in the square box. Which ever point is selected in that square box is the selected color.

Adjusting by RGB or HSV

Underneath the square colored box, there is a series of 6 options (1 each for hue, saturation, value, red, green, and blue) that each contains a radio button, a slider, and an up-down number control. If one of the 3 HSV radio buttons is selected, that value will be represented by the visual slider (see above) while the other two values will be the axes of the square color box. The controls work the same way if one of the RGB radio buttons is selected. You can also adjust these values with the sliders provided next to the labels, as well as the up/down controls next to the sliders. Any adjustment of these values will update the visual controls, and any adjustment of those visual controls will update these values.

Selected Color

The selected color is displayed to the right of the RGB and HSV controls. The box is split diagonally, with the upper left section labeled New (the color you are creating), and the bottom right section labeled Current (the color that already exists, and will continue to be selected if you Cancel the dialog). Beneath that is the Hex Value of the selected color (you can also enter any valid hex value there to update the selected color). There is an eyedropper button to the right of the selected color that will allow you to set the selected color by sampling anywhere on the desktop.

User-defined Color Palette

At the bottom of the dialog you will find a series of 12 colored rectangles alongside a button titled Set to Swatch. These comprise the user-defined color palette, and allow you create and save (these will persist even when ExpressMix is closed) your own custom colors. To set a swatch's color, first select that swatch by clicking it. Then choose your desired color, and finally click the Set to Swatch button. You can then load any of the saved colors by simply clicking on that swatch.

Screen References - Project Settings

Project Playback Sample Rate

Select the sample rate you would like to use for the current project while playing back using MixPad. At the project level, this will override the Mixpad options.

Project Playback Channels

Select the channels would like to use for the current project while playing back using MixPad. If you want to create a surround sound project, please set up the channels as 5.1 or 7.1 surround. At the project level, this will override the Mixpad options.

Project Recording Sample Rate

Select the sample rate you would like to use for the current project while recording a sound clip. At the project level, this will override the Mixpad options.

Screen References - Rename Clip

Enter a new name for the selected clip in the text field. This will only change the name of the file within MixPad; it will not change the file name from the location you loaded the file from.

Screen References - Select Speed

You can control the speed the audio plays back at. Normal speed is at 100%, half speed is 50%, and double speed is 200%.

Screen References - Mix Tape Creator

Mix Tape

Creating a 'mix tape' in MixPad means arranging a series of clips end to end with a small cross-faded overlap between each clip. MixPad can do this for you automatically using this dialog. There are two main ways to create a MixTape.

1. Click the 'Load' button in the main window and select all the audio files you want to have in your mix tape. MixPad will then ask you if you want to load them in the mix tape format. If you agree, then MixPad will show this dialog which allows you several options for how you would like the mix tape to be created. Note that MixPad will only offer the mix tape dialog if the option is turned on in the MixPad options.

2. The second way of creating a mix tape is to select a number of clips which have already been loaded to the project window. Select all the clips you want to include and then click 'Clip->Create Mix Tape' and follow the same instructions as above.

Using the Mix Tape Creator Dialog

There are two main methods for creating a mix tape in the Mix Tape Creator.

- The first way is to use the *Smart MixTape* option. If you select this option then MixPad will automatically analyze the structure of your selected songs and arrange them in the best way it can trying to maximise the compatibility between each song to make each transition as seamless as possible.

- If you want more control over how the clips are arranged and mixed then you will need to choose the *Custom MixTape* option.

Custom MixTape

1. Drag and drop the songs in the list into the order you would like them to appear in the project
2. Select *Align Crossfades on Beats* if you would like MixPad to analyze your songs for beats and mix them so each one transitions to the next with minimal interuption to the tempo.
3. Choose to either allow MixPad to apply the best type of crossfade between each song by selecting *Automatic Crossfade Selection* or specify the parameters of the crossfade yourself by selecting *Crossfade Using Custom Fade Parameters*.

Screen References - Tempo Editor

Tempo Editing

If MixPad can't determine the tempo of a clip, or makes an error and gives the wrong tempo, then you can use this dialog to adjust the result. Any changes you make to the tempo of a clip will be immediately visible on the clip in the main project window. This can help with correctly adjusting the tempo and the offset.

If you know the correct tempo then you can directly enter it into the tempo field. Alternatively, you can use the 'Tap' button which will allow you to tap along to the track while it plays. MixPad will base the new tempo on how you press the button.

The 'Grid Offset' field refers to the location of the first beat in the clip. You can adjust this offset so that the grid correctly aligns with the beats in your clip. The +1/2 and +1/4 buttons will move the grid forward by exactly one half beat or one quarter beat based on the current tempo of the clip. Grid misalignment by 1/2 or 1/4 of a beat is the most common type of error when using automatic beat analysis, and this feature lets you fix it quickly.

Screen References - Auto Duck Settings

Use this feature to apply an auto duck effect to a track. When Auto Duck is on, MixPad detects the volume of control track and applies a volume reduction to the other tracks. For example, it is used to lower the music background track when a vocal track is the main focus; then raise the music background when the vocal track is silent.

Click the **Duck** icon on the track control panel, which opens the options for that track.

- **Auto-Duck Fade this Track** - Automatically lower or adjust the volume level based upon the Master/Control track. This is also known as the
- **Auto-Duck Track is Control** - No ducking effect on to the track. This will act as the control track that the duck tracks respond to
- **Auto-Duck Smart (Music Fade, Voice Control)** - Automatically set all tracks as duck or control based on their classification of Voice or Music
- **Auto-Duck Ignore this Track** - The auto duck effect is off on the track

Suppose Track 1 is a voiceover narration and is selected as a **Auto-Duck Track is Control**. Track 2 is the background music and applied as a **Auto-Duck Fade this Track**. MixPad will detect the volume level of Track 1. Once the volume crosses over the **Threshold**, the volume of Track 2 will lower automatically with a **Fade-In** and **Fade-Out** applied for smooth transitions in volume levels. When Track 1's volume is lower than the **Threshold**, Track 2 will return to its original volume level.

Auto Duck Settings

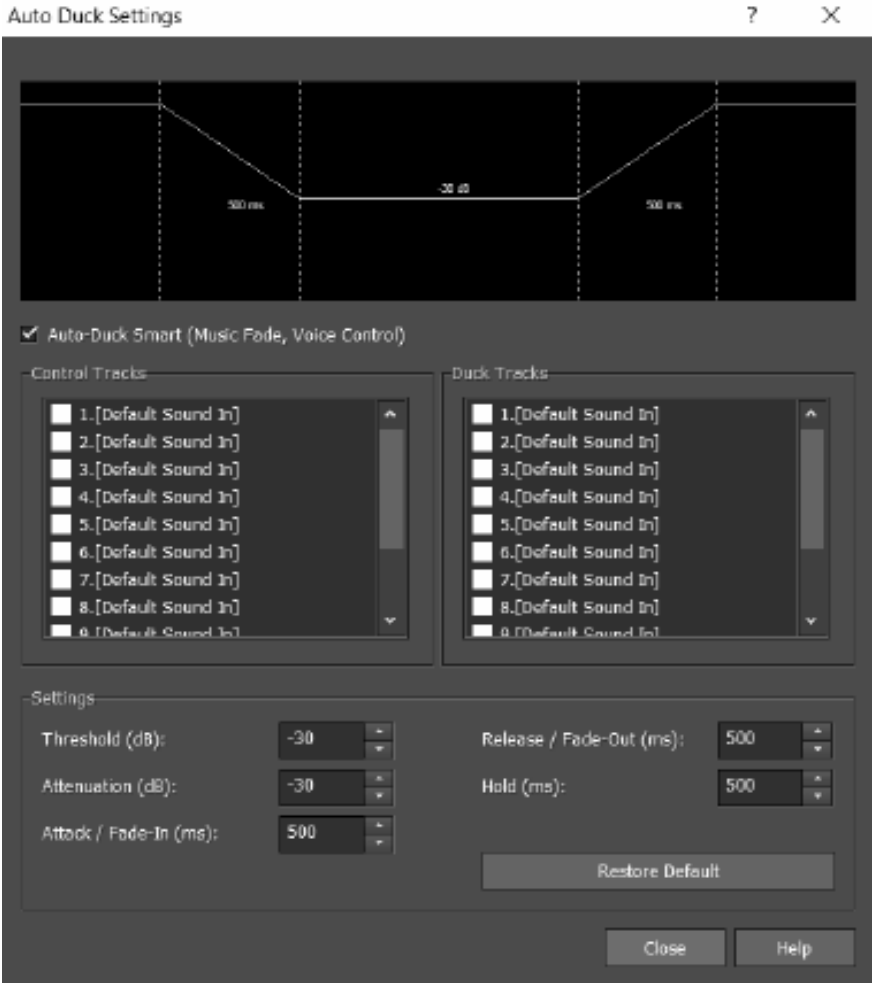

will automatically set a track as duck or control based on its audio classification of Music or Voice.

The **list of the Control Tracks and Duck Tracks** allow you to classify your project tracks as Control or Duck from one dialog instead of adjusting the setting on each individual track.

Settings:

- **Threshold** - The level in the control track that determines when the volume of the duck track should be lowered.
- **Attenuation** - This determines how much to lower the volume of the duck track.
- **Attack / Fade-In** - This determines how quickly to transition from full volume to the duck volume.
- **Release / Fade-Out** - This determines the speed that the volume returns from duck to normal.
- **Hold** - This determines how long the lower volume should be held in the duck track, even if the threshold is no longer exceeded in the control track.

Screen References - Apply Auto Duck

Use this feature to apply an auto duck effect to some tracks. MixPad detects the volume of control track and applies volume reduction to other tracks. It is helpful if you want to lower the music background track when a vocal track is the main focus and then raise the music background when the vocal track is silent.

Suppose Track 1 is selected as a control track and the auto duck effect is applied to Track 2, Mixpad will detect the volume level of Track 1 and once that volume crosses over the threshold configured a dB reduction, a Fade-In and Fade-out will be applied to Track 2 automatically. When Track 1‘s volume is lower than the threshold, Track 2 will return to its original volume level.

- *Threshold (dB)* is the level of the control track at which reduction will be applied on the track being ducked.
- *Reduction (dB)* is the amount of reduction when ducking is fully applied to the track being ducked.
- *Fade in length (ms)* is the length in ms over which the ducking is applied or its fade-in time.
- *Fade out length (ms)* is the length in ms over which the ducking is released or its fade-out time.