Amadeus II v3.0 reference guide

Martin Hairer

January 24, 2001

Contents

Introduction

Thank you for choosing **Amadeus II**. This shareware is a powerful tool that allows easily to record, play, analyze and manipulate sounds. It provides several quite professional features like direct-to-disk sound manipulation and 24Bit sound handling. It also allows to export files using various compression algorithms, including Mp3 and QDesign Music 2.

This reference guide describes briefly the main functions of **Amadeus II**. It is by no means supposed to be exhaustive. The functionalities described here apply to **Amadeus II v3.0**. Some functions may not be present on older versions.

I do not represent any software development company or whatsoever. I am just a graduate student in physics who likes programming. This program was developed entirely during my free time and is not part of any bigger project. Because of these reasons, no exhaustive testing could have been performed. Since some functionalities are system-dependent or even machine-dependent, it may happen that **Amadeus II** presents some dysfunctionalities on your personal configuration. If so, please feel free to contact me at Martin.Hairer@math.unige.ch. For obvious reasons, I always assume you are in possession of the latest available version of **Amadeus II**. You can check the number of the latest version (and of course also download it) at

http://mpej.unige.ch/~hairer/Amad2.html

Just to avoid getting in any trouble, I have to mention that I am not responsable for any damage that may be caused by this or any other version of **Amadeus II** to your machine.

Installation notice

This release is made of two distinct components.

The **Amadeus Application** folder contains everything that is needed to make the program work. It is important not to remove the application file from this folder. At startup, the program looks for the external files in every folder located at a higher level than the application file in the file hierarchy. This means that if you put the program on the desktop, it will scan your whole disk at startup, which may take a considerable amount of time.

The **Lame** folder contains the shared library used for Mp3 encoding. In order to be able to use these functions, the file "Lame SharedLib" (or an alias) needs to be located either in your "Extensions" folder or in the "Amadeus Application" folder.

1 What's new?

Here is a list of the major new functions in version 3.0.

- A "Window" menu appears.
- The **QuickTime Export...**, **QuickTime Import...** and **Mp3 Export...** functions disapear as those, but all their functionalities are now incorporated in the **Open...** and **Save As...** functions. This hopefully helps to make the user interface more consistent.
- Sonograms can now be exported in many different graphics formats that were not available before. Furthermore, they can be computed at resuolutions higher that the screen resolution (up to 288dpi).
- Windowing functions for spectra and sonograms are available. This especially helps to make sonograms of much higher quality.
- Raw data can be exported and imported.

- Arbitrary sampling rates are supported, even beyond the 64kHz limit.
- A waterfall display is implemented.
- A new multiband real-time spectrum analyzer is implemented.
- Waveform statistics can be computed.
- Average spectra can be computed.
- It is now easy to make smooth transitions between several pieces of sound.
- Automatic generation of markers.

There are also several bug fixes and many minor aesthetic changes.

2   The documents

There are five types of documents handled by **Amadeus II**.

2.1   The sound windows

They are the main "working place". The sound is depicted as a wave showing the pressure as a function of time. If you create a new window, it is gray, meaning that no sound is present. In order to select a part of the sound, just click at the beginning of the select and drag the mouse until the end of the selection. The "shift" key allows you to extend the current selection. If the sound has several channels, all of them will be selected. In order to select only one channel, hold down the "option" ("alt" on some keyboards) key.

There are two small white strips above and below the one(s) containing the sound itself. The lower strip shows a time scale. The upper one contains the marks of the current sound (if some are present). To create a mark, just click into that strip and drag the mouse to the appropriate location.

On the top of the window, the **Show Whole Sound** function provides an alternative way to navigate in a huge document. This function is only enabled if the sound is large enough. The visible part of the sound is depicted by a white rectangle surrounded by a blue frame. You can change it by either dragging it or extending the sides of the frame.

There are four small icons at the left bottom of the window. The first one allows to change the `time scale` at which the sound is to be drawn. If you click onto the second one, a dialog showing the characteristics of the sound opens. If the sound is currently stored on a hard disk, you can try to load it into RAM. The third icon allows to set the volume level at which the sound is to be played. The last icon allows to set if the sound has to be looped at playback or not. If the sound is stored on a hard disk, it may not be a good idea too loop a short sequence, since this causes the reading device to jump very quickly forth and back. If the hard disk does not "keep up" with the sound playback, your machine may crash.

The small icon at the top right of the window allows to change the `amplitude scale` of the displayed sound.

2.2   The "Sonogram" windows

A "sonogram" is a graphic representation of a sound showing the frequency as a function of time. The amplitudes of the different frequencies are shown as colors. For example, a pure sound will be represented by a straight horizontal line, because there is only one frequency present. Sonograms are a quite useful tool to compare sounds or to find out which notes are present in an accord for example.

A sonogram window is composed of three parts: A tool bar, the sonogram itself and the sound from which it was produced.

The tool bar. The "colors" frame contains a pop-up menu that allows to change the color palette and a slider that allows to change the `color scale` of the sonogram. The `color scale` determines which color has to be attributed to which amplitude. If you glide the slider to the right, the color scale becomes more sensitive, *i.e.* small amplitudes will become visible.

The second frame contains two icons. The first one opens a dialog that allows to set the different color palettes. The second one opens a dialog that allows to change the different parameters of the sonogram. Here is a short description of them. **Max. Frequ.** is the maximal frequency shown in the sonogram. If you lower it, you will have a better vertical resolution but you will cut off the high frequencies. **Sizes FFT** is the number of points considered in one Fourier transform. If you increase this value, you increase the quality of the sonogram, but the computation gets slower. **Sizes Picture** is the vertical size of the sonogram. **Frequ. Scale** allows to choose between a linear and a logarithmic frequency scale. **Ampl. Scale** allows to choose between a linear and a logarithmic amplitude scale. In the case of a logarithmic amplitude scale, **Gain** allows to choose the difference in dB of the amplitudes corresponding to the first and the last color in the color scale. **Emph** allows to preemphasis the spectra in order to magnify high frequencies, which have usually lower amplitudes. It is also possible to choose a windowing function in this dialog box (see Section 4.6.1 for more details about windowing functions) and to change the resolution of the sonogram. Changing the resolution has no effect on the picture drawn on the screen, but it causes the sonogram to be recomputed in a higher quality when it is exported.

The sonogram. If you click into this part, a small window appears, showing the frequency corresponding to the location of the mouse, as well as the note that is closest to that frequency.

You can copy the sonogram into the system scrap by choosing **Copy** in the **Edit** menu.

The sound. This part of the `Sonogram window` behaves exactly like the `Sound window`, but you cannot select anything. The three small icons in particular have the same behavior.

2.3 The "Spectrum" windows

A "spectrum" is a graphic representation of the amplitude as a function of the frequency at one point of a sound. A spectrum is in fact a "vertical cut" of a sonogram. The meaning of the functions present in the spectrum windows should be quite clear from what has been said about the sonogram windows.

2.4 The "Animated Spectrum" windows

An animated spectrum contains a sound, together with a sequence of time-equidistant spectra. You can play the sound by clicking onto the **Play** button located at the left of the window. You can also hit the space bar instead. While the sound is being payed, the window shows the spectrum corresponding to the current position. You can also go to a specific position by using the slider at the bottom of the window (use it only when the sound is not currently playing). The "Preemphasis" button allows to emphasize high frequencies. The "Scale" field allows to change the horizontal and vertical scales in order to fit your personal needs.

If you choose **Copy** in the **Edit** menu while an animated spectrum is the frontmost window, the program copies the sound corresponding to that window into the current scrap.

2.5   The "3D Spectrum" windows

An 3D spectrum is much like an animated spectrum, except that the different spectra are drawn all at the same time in order to present as much information as possible at a glance. The functions in the window are essentially the same as for the animated spectrum, except that you can change the color palette used to colorize the spectra. The available color palettes are the same as for the sonograms.

If you click into the window, a small window shows up containing the frequency corresponding to the location of the mouse. If you want to get the frequency of a particular peak, you have to click at the *bottom* of the peak to get the right value.

If you choose **Copy** in the **Edit** menu while a 3D spectrum is the frontmost window, the program copies the picture into the scrap.

3   The floating palettes

There are three floating palettes which can be shown/hidden using the **Windows** menu. Here is a short description of their contents.

The commands palette. There are six icons appearing in this palette. Each of them is linked to a menu item, *i.e.* clicking on it has the same effect than choosing the corresponding menu item.

The four icons symbolizing a tape recorder are linked respectively to the **Play**, **Record...**, **Pause** and **Resume** items of the **Sound** menu.

The icon symbolizing a spectrum is linked to the **Spectrum...** item of the **Analyze** menu.

The icon symbolizing a sonogram is linked to the **Sonogram...** item of the **Analyze** menu.

The memory palette. The progress bar shows the fraction of the heap used by **Amadeus II**. **Total memory** is the total amount of the heap allocated to **Amadeus II** by the system. This amount can be changed by selecting the application program in the Finder and selecting the **Get info...** item in the **File** menu. **Free memory** is the part of the heap that is not used by **Amadeus II**.

The selection palette. It contains the coordinates of the selection of the sound contained in the frontmost window. The first time is the length of the selection, the second one is the start of the selection and the last one the end of the selection. The units can be changed in the **Preferences...** dialog of the **Edit** menu.

If you click into this palette, it has the same effect as selecting the **Set Selection...** item of the **Selection** menu.

4   Menu commands

4.1   The "File" menu.

"New". Creates a new sound file. The characteristics of this file can be set in the **Preferences...** dialog of the **Edit** menu.

"Open...". Allows to open a file (sound or sonogram) previously saved on disk. Fully recognized sound formats are `AIFF`, compressed `AIFC` and System 7/8 `sfil`. The recognized sound compression formats are those provided by the SoundManager, which varies from one system to another (a-Law compression is only recognized under MacOS 8.5 and later). If the document format is different from the formats above, **Amadeus II** will try to open the file using QuickTime. The recognised file formats may differ depending on

your version of QuickTime. If you are using an old version of QuickTime, it may happen that **Amadeus II** is not able to open Mp3 files.

If NavigationServices are present (MacOS 8.5 and later), they are used.

If there is not enough memory left to load the sound, **Amadeus II** opens it anyway and handles it direct-to-disk. This is very useful to open huge sounds.

"Close". Closes the frontmost window. If it contains an unsaved document, asks the user if he wants to save it.

"Save". Saves the content of the frontmost window on the hard disk. If there is no file associated to that window yet, the behavior of **Save** is the same than that of **Save as...**.

"Save as...". Shows the standard saving dialog. If NavigationServices are present (MacOS 8.5 and later), they are used. If the document to be saved is a sound, you can in principle choose between seven formats (Audio International File Format (AIFF), WAVE, Mp3, μ-Law, System Sound, Movie and AVI sequence). For every format, the "Settings..." button allows to choose several parameters, for example a compression algorithm. Except for the Mp3 format, those parameters are provided by QuickTime, so they may vary slightly depending on the version of QuickTime you currently use.

If QuickTime is not installed, only AIFF, System Sound and Mp3 are available. If you compress a sound, be conscious that some quality may be lost, particularly if the compression rate is high, so test the results of the different laws before you use them.

> **Important!** In order to use the Mp3 exporting function, the Lame shared library must be installed! You have to place the file `Lame SharedLib` located in the `Lame` folder into the extensions folder of your system folder or into the `Amadeus Application` folder.

The LAME encoding engine is a free Mp3 encoding engine originally developed by Mike Cheng

http://www.uq.net.au/~zzmcheng .

It is licensed under the LGPL (GNU Library General Public License). For a copy of the LGPL, see the file `COPYING` also located in the `Lame` folder. The Lame source code can be obtained at

http://www.sulaco.org/mp3 .

If the frontmost window is a sonogram, "Save As..." allows to save it either in a file format owned by **Amadeus II** (such files can then be opened again), or in one of the many different graphics formats provided by QuickTime.

"Import Raw Data...". This allows you to import raw sound samples. This may be useful if the sound you want to open is not recognised by **Amadeus II** or if you want to "cheat" on the sampling rate of a sound.

"Export Raw Data...". This item allows you to save the raw sound data of the current sound window into a file. The resulting file does not contain any information about sampling rates, sound quality, markers, and so on. This function is only intended for low-level communication between applications.

"Print...". This allows you to print sonograms, spectrums, and so on. Only the sounds themselves can not be printed for obvious reasons.

"Page setup...". Allows to change the setup for the printer.

"Quit". Terminates the program after closing all documents.

4.2 The "Edit" menu.

"Undo". Undoes the last action performed in the frontmost window. For `Sound windows`, multiple undoing is supported up to a level that can be set in the **Preferences...** dialog.

"Redo". Annihilates the effect of **Undo**.

"Cut". Equivalent to **Copy** followed by **Clear**.

"Copy". Copies the selection of the frontmost window into the scrap. If the window is a `Sound window`, the selection is copied into the current internal scrap. The current scrap can be chosen in the **Scrap** submenu. If you switch to another application, the current scrap is placed into the system clipboard, provided that you checked the "export scrap" box in the **Preferences...** dialog and that there is enough RAM left. If you do not intend to communicate data between other applications this way, it is recommended not to use the "export scrap" option.

"Paste". If no sound is selected, it inserts the content of the current scrap at the insertion point of the frontmost `Sound window`. If a sound is selected, the selection is deleted first. If the quality of the scrap doesn't fit the quality of the frontmost `Sound window`, the content of the scrap is first converted as to fit that quality. Notice that the content of the scrap is not affected by this operation. If you switch from another application to **Amadeus II** and the system clipboard contains some sound data, it is placed into the current scrap.

"Paste special". Mixes the content of the current scrap to the sound contained in the frontmost `Sound window`. It will start exactly at the beginning of the selection, not depending of the size of the selection. If the qualities don't fit, the content of the scrap is automatically converted.

"Clear". Deletes the current selection.

"Select all". If the frontmost window is a `Sound window`, extends the selection of the sound to the whole sound.

"Current scrap". Allows to change the current scrap.

"Copy to New File". Opens a new file with the content of the current selection. It doesn't copy the selection to the scrap, so the content of the current scrap does not get lost.

"Preferences...". Opens a dialog that allows to change the behavior of **Amadeus II**. In particular, you can set the default quality of a new sound and the units at which a time has to be displayed.

"Reload effects". Searches for external filters present in any subdirectory of the directory containing the application file. This is mainly useful for development purposes when you change an external filter and want to test it without quitting **Amadeus II**.

4.3 The "Selection" menu.

This menu contains numerous functions that allow to gain full control over the extent of the selection, as well as some functions to handle marks.

"Set Selection...". Opens a dialog box that allows to set the selection precisely. The numerical values have to be entered in milliseconds. If marks are present, the **Mark...** buttons allow to put the position of a mark into the corresponding text field.

"Mark Selection...". Allows to put one mark at the beginning of the current selection and one at the end. If no sound is selected, it puts a mark at the current insertion point.

"Jump to". The pop-up menu attached to this item allows you to choose the selection, the whole sound, a mark or the time between two marks. The program then changes the `time scale` and the position of the scrollbar of the frontmost `Sound window` in a way that the chosen part of the sound fills exactly the width of the window.

"Extend to Start". Extends the selection to the beginning of the sound.

"Extend to End". Extends the selection to the end of the sound.

"Extend to Next Mark". Extends the selection to the location of the first mark that comes after the end of the selection. If there is none, it has the same effect as **Extend to End**. This function can also be accessed by pressing command-right arrow.

"Extend to Previous Mark". Extends the selection to the location of the last mark that comes before the beginning of the selection. If there is none, it has the same effect as **Extend to Start**. This function can also be accessed by pressing command-left arrow.

"Save Selection As...". Behaves like **Save As...**, except that it creates a sound that contains only the current selection.

"Split According to Marks...". Splits the sound into pieces according to the marks. The "Start color" indicates the color of the marks that specify the beginning of a piece, the "End color" indicates the color of the marks that specify the end of a piece. If you choose the "Consider all marks" option, the function behaves as if all the marks had the same color.

The name of each piece of sound is equal to the name of the mark located at its beginning. Once you have selected "OK", a standard file saving dialog appears, in which you have to enter the name and location of the folder that will contain the various pieces.

"Join Files...". This is the complementary function to the **Split According to Marks...** function. It allows to concatenate several AIFF files into one big file. The dialog box allows to choose the quality of the resulting file. If the chosen files are compressed, **Amadeus II** automatically decompresseses them before concatenating.

"Clear All Marks". Deletes all the marks of the frontmost sound.

"Generate Marks...". Allows to generate equidistant marks on your sound. This is especially useful in conjunction with the **Split According to Marks...** function if you want to cut a very long sound into pieces of reasonable size.

4.4 The "Sound" menu.

This menu regroups mainly the items involved in the acquisition and the reproduction of sound.

"Play". Plays the sound contained in the frontmost `Sound Window`. If the selection is longer than a few milliseconds, only the selection is played, otherwise the whole sound is played. This menu item can be accessed by pressing the space bar.

"Play From Insertion.". Similar effect than **Play**, but the sound is played starting from the beginning of the selection until the end of the sound is reached. This menu item can be accessed by pressing the tabulation key.

"Play Between Marks...". opens a dialog box in which you can choose two marks of the current sound. The sound between the two selected marks will then be played.

"Record...". Opens a window that allows to record a sound from any input device recognized by the Sound-Manager. The first progress bar indicates the length of the recorded sound (If it is filled, the maximal length has been reached). The maximal length of the recorded sound can be set in the **Preferences...** dialog box. If you shorten it, you use less disk space when this window is open. The second progress bar indicates the volume level measured by the active input device. The "Peak" indicator turns red if a too high level has been reached at least once during the recording.

It may happen that your sound input device is not able to sample sound at the quality you request. In that case, **Amadeus II** will do the resampling automatically in real-time while recording.

"Record to New File...". This command behaves exactly like the **Record...** command, except that it creates a new sound document containing the recorded sequence. The quality of this document is the default quality that can be set in the preferences. If you intend to record a sound of a certain length, it is highly recommended to use this function rather than the **Record...** function, since the sound will not be duplicated when you hit "OK" and so everything is much faster.

The **Max. CPU** checkbox allows to freeze all the other applications while recording. This is useful to minimize the risk of having cuts in the recorded sound, especially if your hard drive is not very fast. The **Mark Sound** checkbox allows to generate marks containing the system time at the beginning and at the end of the recording.

"Jump to Play Position.". When a sound is currently being played, this function selects 250 milliseconds of sound about 0.2 seconds before the actual playback position and then calls the **Jump to selection** function. This can be quite useful to detect the position of a crack in a sound for example.

"Stop". Stops playing the frontmost sound. If the **Recording** window is open, stops the recording.

"Pause". If the **Recording** window is open, pauses the recording.

"Resume". If the **Recording** window is open, resumes the recording.

"Characteristics...". Opens a dialog allowing to change the characteristics of the frontmost sound. This operation can also be undone.

Be aware of the fact that 24Bit sound can not be recorded by the actual versions of the SoundManager, so this is mainly useful to preserve high quality even if you apply many effects.

If you work on large sounds in the direct-to-disk mode, it may happen that **Amadeus II** fragments the sound in order to gain speed. If a sound is too fragmented, it may happen that the contrary effect is achieved, so you can defragment a sound. Defragmentation of a sound is automatically performed when you save it.

"Playback Pitch...". Allows to play the sound at a rate different to the recording rate. This function may not produce satisfactory results if the recording rate of the frontmost sound times the playback factor is too high (more than 64kHz). This is related to the internal description of the sound header by the SoundManager and to obvious hardware limitations.

4.5 The "Effects" menu.

"Echo...". Allows to apply an echo to the current selection. If you check the **Go further** box and put a value of s seconds in the text field, the echo of the selection will be prolongated by s seconds. Nevertheless, *no* echo will be applied to the s seconds following the selection.

"Amplify...". Allows to amplify the selection by a given factor. If the "fading time" is non-zero, a smooth transition is made between the amplified and the non-amplified sound.

"Filter...". Allows to apply a frequency filter to the current selection. This item is similar to the graphic equalizer of a hifi chain.

"Reduce background". Allows to reduce the background noise of a sound. Do not expect miracles from this function, lost information can simply not be restored! Nevertheless, it works quite fine sometimes, especially if the sound is not too complicated, for example only a voice speaking. In order to get optimal results, you'll probably have to play around with the **Expert...** function.

Here is a short description of the algorithm hidden behind this function. The "noise profile", as it can be set in the **Expert...** menu, determines a threshold in Fourier space under which frequencies are completely suppressed. In order to get an idea of typical values for this threshold in your case, you may make a spectrum analysis of a pice of the sound and look at the typical amplitudes in regions outside the peaks of significant data. In short, the larger the "Amax" parameter, the more noise will be suppressed. The larger the "ΔA" parameter, the less high-frequency noise will be suppressed.

The remainder of the **Effects** menu shows the external filters detected by **Amadeus II**. For the moment being, the standard release of **Amadeus II** contains eight filters, whoch are shortly described below.

"Extend Selection". This allows to extend the selection to the nearest points where the waveform crosses the origin. The **Settings...** function allows to choose which type of crossings (upwards, downwards or both) are recognised as such.

"Fadings". This effect allows to make fade-ins or fade-outs of various types. To change the type of the fading, open the **Settings...** dialog and select one of the four different types. The **Cross-Fade** function allows you to make a fading between two given amplitudes (and not only between 0 and 100%).

The **Transition to the Left** and **Transition to the Right** functions allow you to easily create a transition between two pieces of sound. This is how it works. Let's say you want to produce a file containing the songs "Song A" and "Song B". Copy them one behinder the other into your file and place a marker at the junction of both songs. Then select say 1 second of "Song B" starting from the marker and select **Transition to the Left**. This will first create a **Fade In** of 1 second at the beginning of "Song B" and a **Fade Out** of 1 second at the end of "Song A". Then it will merge these two parts into one, creating a nice transition.

"Normalize...". This allows to normalize the amplitude of the selected sound with respect to the maximal possible amplitude.

"Repair". This effect allows to suppress a crack in a sound. In order to achieve this, first search for the crack. This can be done for example with the help of the **Jump to play position** function. Then select the crack and a very short piece of sound before and after it. Make the selection as short as possible (about 40-50 ms maximum)! Then you can apply the filter. Try different settings to get optimal results.

The actual version of the **Repair** function allows to search automatically for cracks.

"Reverse". This effect simply makes a time-reversal of the selected sound.

"Sample Filters". This effect allows you to damp either high or low frequencies of your sound.

"Set Pitch...". This effect allows to physically change the pitch of a portion of the sound.

"Simple generators". This effect allows to generate simple sounds like sine waves, silence or white noise. White noise is particularly useful if you want to test the effect of a filter. (Apply it to 5 seconds of white noise and then compute the average spectrum of the noise. This will reproduce exactly the response of your filter, since pure white noise has a flat spectrum on average.)

4.6   The "Analyze" menu.

"Spectrum". Makes a spectral analysis of the selected sound and stores the result in a `Spectrum window`. It is possible to apply a windowing function to the spectrum. It is also possible to compute the average spectrum of the whole selection. The "Overlap" field contains the length (in points) of the ovelap between two successive FFTs. Notice that the "Number of points" contains the number of points in the spectrum analysis. The actual number of sound samples needed to compute this spectrum is in fact two times bigger.

"Animated Spectrum...". Makes an animated spectrum out of the selected portion of sound. You have to select at least 10'000 samples (which corresponds approximately to 0.5 seconds of a sound sampled at 22'050 kHz). The modal box that shows up when you select this function allows you to select the number of spectra per second that are to be computed and the maximal frequency that shows up. If the original sound is a stereo sound, it also allows you to show both channels.
This function may use a lot of memory since the spectra are always stored in RAM.

"3D Spectrum...". Makes a 3D spectrum out of the selected portion of sound. You have to select at least 10'000 samples (which corresponds approximately to 0.5 seconds of a sound sampled at 22'050 kHz). The modal box that shows up when you select this function allows you to select the total number of spectra and the maximal frequency that shows up.
This function may use a lot of memory since the spectra are always stored in RAM.

"Sonogram...". Creates a sonogram from the current selection.

"Waveform Statistics...". Opens a window containing statistics on the current selection. If the sound is stereophonic, the left column corresponds to the left channel and the right column to the right channel. The meaning of the various numbers is the following.

- The "Minimum/Maximum Sample Value" is the minimal/maximal value the sample take in the current selection. These values are normalized in such a way that the clipping values are ± 1.
- The "Peak Amplitude" is the amplitude of the difference between the maximum and the minimum sample value. This amplitude is given in dB with respect to its maximal value (which is 2).
- "Possibly Clipped Samples" indicate the number of samples that take the extremal values ± 1.
- "DC Offset" indicates the average vertical offset of the waveform. It is given in % of the maximal possible value.
- "Minimum/Maximum/Average RMS Power" gives the minimum/maximum/average value of the root mean square power in the selection. These values are given in dB with respect to their maximal values (which are atteined for a square wave with maximal amplitude). The size of the window used to compute the RMS power can be adjusted.

"Real-time spectrum...". Opens a window which shows a real-time spectral analysis of the sound entering into the current sound input device. There are many options to affect the display of the spectrum. The "preemphasis" option allows to amplify artificially the high frequencies.
This function seems to me a great pedagogical tool to show how a note is composed of a fundamental frequency and its harmonics for example.

"Oscilloscope...". Opens a window that shows in real-time the sound entering into the current sound input device.

"Multiband RTA...". Opens a window with a real-time spectrum analyzer on a "log-log" scale. It behae sessentially like the spectrum analyzer on a hifi chain, except that amplitudes are shown in decibels relative

to the mean amplitude of the sound. The consequence is that the frequency profile of a given sound will not depend on its overall amplitude.

"Waterfall...". Opens a window with a real-time waterfall display. The meaning of the picture is exactly the same as the meaning of a sonogram, except that the time-axis is the vertical one an the frequency axis is the horizontal one. The "Settings..." button allows to change the speed of the display. It also allows to dump the result of the spectrum analysis into a file. This file will simply be a text file, so that it can be imported easily into Excel or some other data analysis program like Matlab.

4.6.1   Windowing functions

Windowing functions are available for all spectrum analysis functions. Applying a window to a wavefrom before computing its spectrum is especially useful for sonograms. There it allows to get much sharper and more regular pictures. The "Kaiser" windowing function allows you to specify the "broadness" (or variance) of the windowing function. The smaller the number, the closer the result will be to what you get without windowing. As a general rule, if you increase the number of points you use for the FFT, yo should decrease this value for optimal results.

For a more detailed overview on windowing functions, you can visit the page

> http://www.mathworks.com/access/helpdesk/help/toolbox/signal/hamming.shtml ,

and the other pages on the same site.

4.7   The "Windows" menu.

This menu contains the list of currently open windows and floating palettes. The windows are grouped according to their type. This menu is disabled when a modeless dialog box is currently active.

5   Some Tricks

Here is a collection of some tricks which I found useful when using **Amadeus II**.

- If you want to record a long piece of sound (say more than five minutes), use the "Record to New File..." function. The "Record..." function will make a copy of your sound when you click on "OK", which may take a considerable amount of time.

- In order to manipulate a very long sound, first cut it into pieces of convenient size using the "Generate Marks..." and then the "Split According to Marks..." functions. Work on those pieces separately, then glue them together using the "Join Files..." function.

- When you save files in an intermediate stage of your work (which is highly recommended), always save them as uncompressed AIFF files. This way, you will not loose any quality and the saving and opening will be faster than with other file formats.

- Use option-click to select individual channels and command-click to place a marker at the location of the mouse.

- Take a look at the "Useful Shortcuts" section below, it can save considerable amounts of time...

6   Useful shortcuts

Here is the list of "hidden" shortcuts that can be used with **Amadeus II**:

key	action
space bar	Plays the current selection
space bar	When recording, starts/stops the recording
command-space bar	Plays from the insertion point
tabulation	Plays from the insertion point
tabulation	When recording, places a marker at the current point
left arrow	Shifts the content of the window to the left
right arrow	Shifts the content of the window to the right
down arrow	Jumps to the current play position
up arrow	Places a mark at the start of the selection
command-left arrow	Extends the selection to the previous mark
command-right arrow	Extends the selection to the next mark
option-click	Allows to select only one channel
command-click	Places a mark at the clicked point
shift-click	Extends the selection to the clicked point
page up	Jumps to the beginning of the sound
page down	Jumps to the end of the sound
a	Extends the selection by one point to the left
s	Shrinks the selection by one point at the left
d	Shrinks the selection by one point at the right
f	Extends the selection by one point to the right

7   Registering

Some functions are disabled after 15 days in the demo version; everything else can be accessed. If a function seems to be disabled, it means either that it is not accessible in the current state of the program or that it is not implemented yet. Before these 15 days, everything works **exactly** as if the program was registered.

Registering **Amadeus II** enables those disabled functions again. Moreover, you will be put in a mailing list that keeps you informed about every new version. If you want to be removed from that mailing list, just tell me at `Martin.Hairer@math.unige.ch`.

The registration fee for **Amadeus II** is 25$ US, to be paid to Kagi, *not to me directly* (It costs me about 7$ to cash a cheque). Kagi will the provide you with a serial code, which has to be entered in the **Registration...** dialog of the Apple menu. If, for any reason, Kagi processes your payment but does not provide you with a serial code, please send me an email and I will provide you with one. You can also register online. For more information about registering, see the "Read Me (Register)" file in the "Register" folder.