

WINRAD

by Alberto I2PHD

What is Winrad ?

Winrad is a software program that implements a so-called Software Defined Radio (SDR), meant to run under Windows XP, Windows 2000, or Windows 98SE. In a nutshell, it accepts a chunk of up to 192 kHz coming from a complex mixer in form of two signals I and Q fed to the PC sound card, does a fine tuning inside that segment with a point-and-click technique, demodulates (LSB/USB/CW) what has been tuned and optionally applies a series of filters to the results of the demodulation.

It implements only a subset of the functionalities offered by the excellent program Linrad by Leif Åsbrink SM5BSZ, which, as its name implies, is meant to run under some versions of Linux. Winrad is not a porting of Linrad to Windows. It has been rewritten from scratch, using the visual and programming tools offered by the C++Builder compiler from Borland, and the Win32 threading model, which has been preferred over pthreads. Only some of the Linrad functionalities have been implemented, with a special consideration to user friendliness, ease of installation and ease of use. Winrad must hence be considered "work in progress", and it is presently distributed as a time locked beta version, which will cease functioning after a given calendar date, so to not leave old and obsolete versions lying around. Before the expiration date of a given beta version, a new version will be released so to not leave anybody with a non functioning program.

Installation

To install Winrad, just execute the Winrad setup program, and answer to a couple of questions (destination directory and if you want an icon created for you). After the installation, just double-click on the new Winrad icon to start the program. Winrad will then present you a list of the sound cards it has found on the PC, and will ask you to specify which one you want to use for input and which one for output. You can specify the same card for both, or use two different cards. You don't have to worry about different sampling rates in input and output, as the program will do, if needed, an upsampling of the output audio signal, just before sending it to the sound card for reproduction.

Important: this program requires a display resolution equal or greater than 1024x768, with 1280x1024 as ideal. If your display is set for a lower resolution, Winrad will not run.

Use

Winrad has been designed with the goal to make its use the more straightforward possible. I have often said that a good program should not require any user-guide manual to be correctly used by anybody. Winrad does not quite stand up to this requirement, hence the need for the document you are reading just now, but I strived to make it as simple as possible.

Winrad when started takes over the entire screen, there is no the customary window with the title bar, system icons, etc. You can minimize the program to carry out other Windows tasks, leaving it running in the background, by using the Minimize button.

The first step to use it is of course to connect the two I and Q signals from your hardware to the left and the right channels of the sound card. You have also the possibility to feed Winrad with just one signal, albeit with reduced performance (half of the useful bandwidth, no unwanted band rejection). This mode of use can be useful when you just want to use the filtering capabilities of the program, e.g. after a standard Rx that provides itself the sideband selection function. Winrad needs to be informed of this, and you can do that by using the ShowOptions menu as shown in Figure 1 :

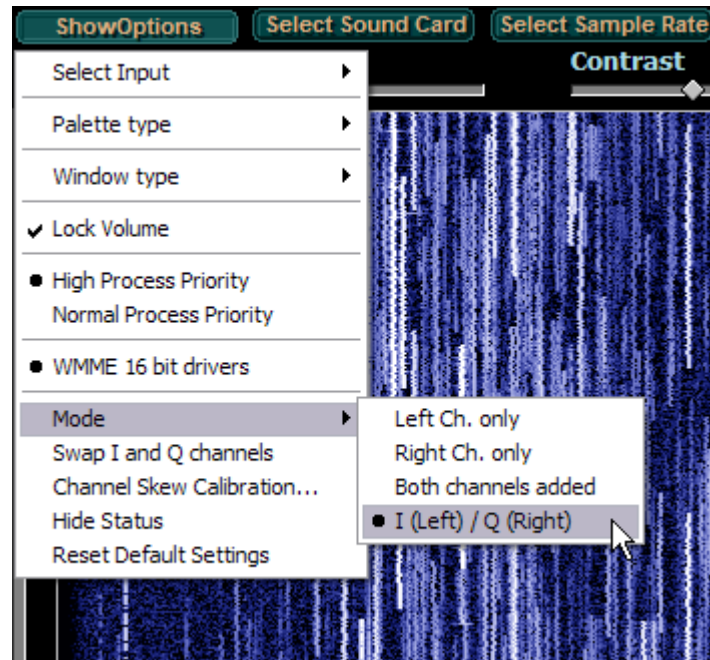


Figure 1 - The ShowOptions menu

Once done that, you choose whether the input is coming from a radio hardware, or from a previously recorded Wav file :

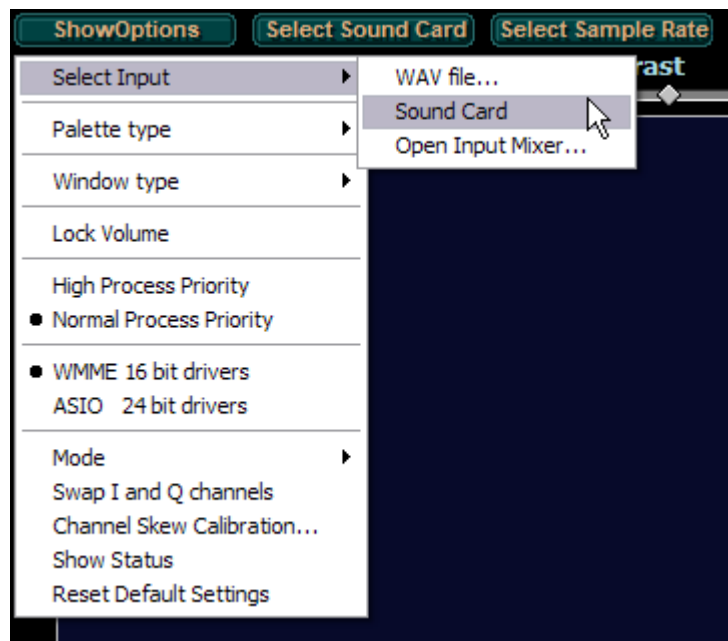


Figure 2 - Input selection menu

If a Wav file is not open, the default behavior is to read from the sound card. The "Input Mixer" choice is used to open the Windows mixer Recording Panel

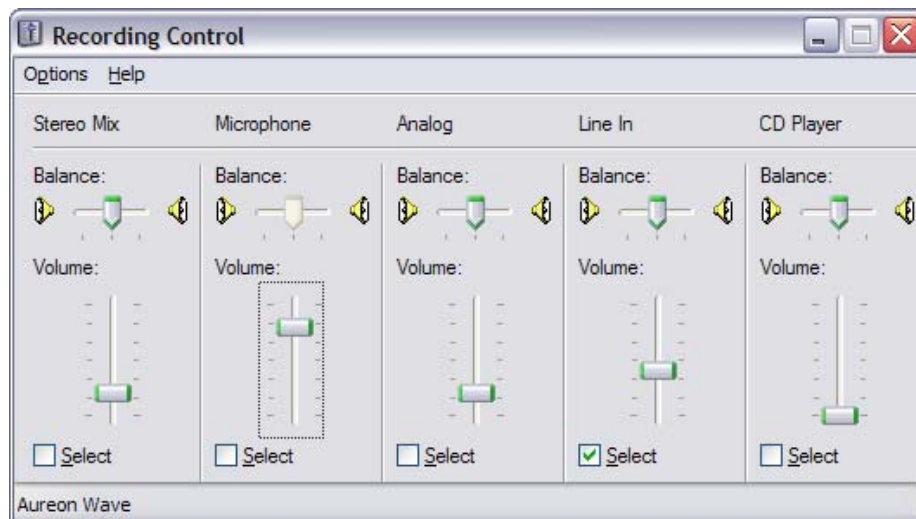


Figure 3 - The Windows mixer panel

In Figure 3 the controls refer to a specific audio card. You may find some minor differences depending on the card you have on your PC. Just a note here : some semi-professional sound cards, notably the Delta 44 from M-Audio, do not implement in their drivers the functionalities needed for this panel. If you select "Input Mixer" when such a card is selected as input, you will receive an error message from Windows. If you use a Delta 44 sound card, please set its Control Panel as follows, where the important settings are the DMA Buffer Size, the Multi Track Driver Devices and the Rate Locked:

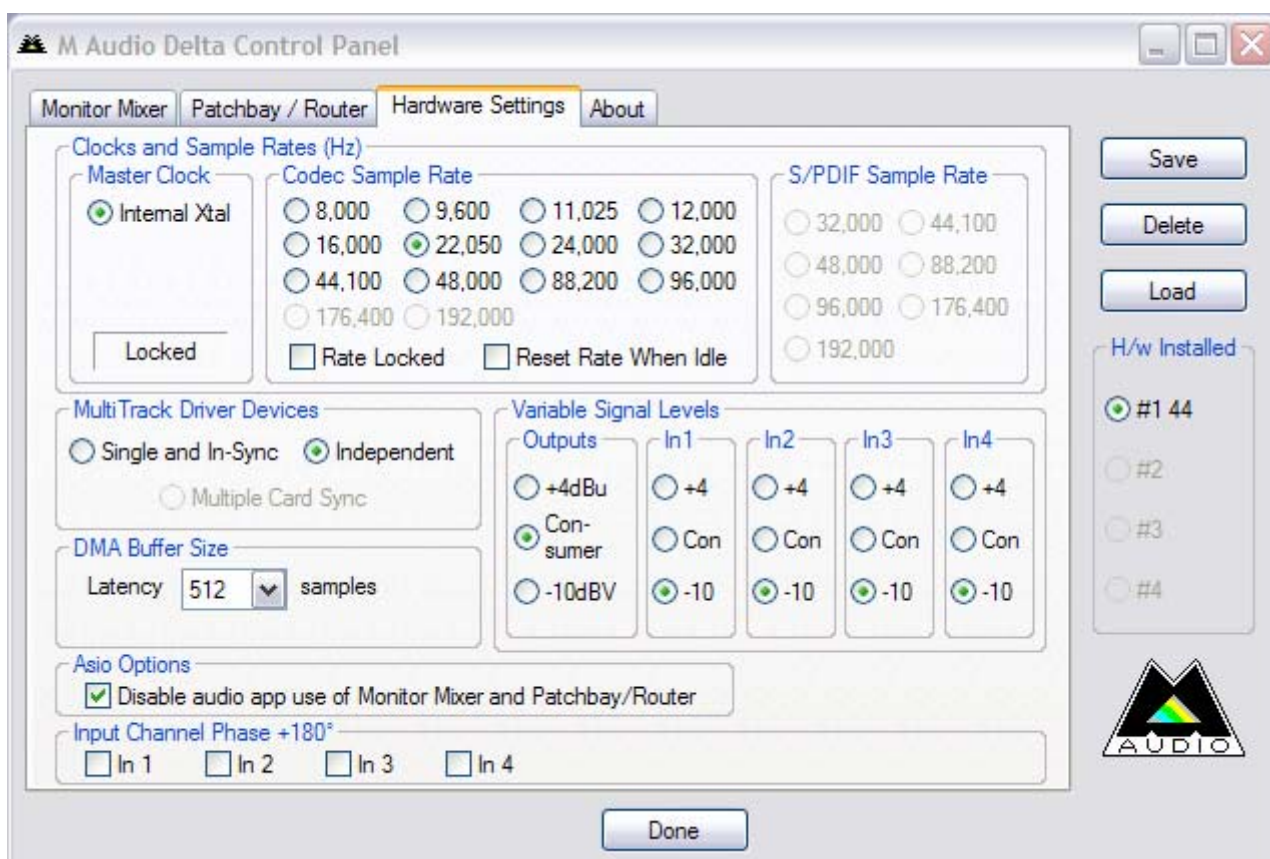


Figure 4 - How to set the Delta 44 Control Panel

Other choices present in that menu are the selection whether to use the WMME drivers for your sound card or the ASIO drivers (if installed), the Palette type and the window type. This latter choice does not refer to the Windows version, but rather controls which kind of windowing is done to the signal in the FFT process. If you don't know what all of this means, leave the default set by the installation program (\sin^5), which is the optimum for almost all the cases. There are no performance penalties in the choice of a window over another, as they are all precomputed. The "Swap I and Q channels" choice is useful for hardware that outputs the I signal on the right channel

and the Q signal on the left. The “Channel Skew Calibration” choice activates a routine that allows to compensate for unbalances in the two channels of the sound card. More information is shown on the panel that opens when selecting that choice.

Pressing the “Select Sound Card” button you can choose which card to use, if you have more than one. Figure 5 is what I see on my PC.

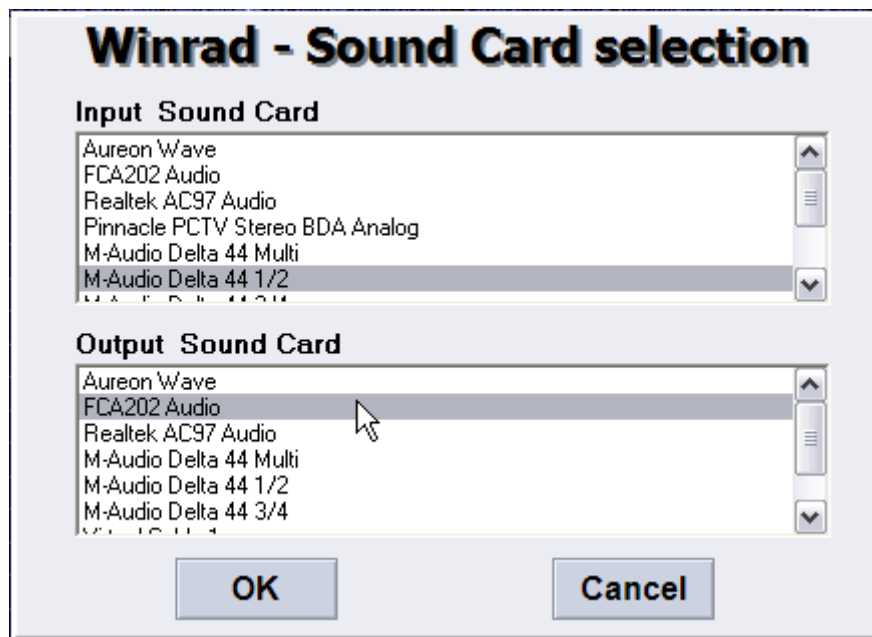


Figure 5 - The sound card selection panel

Then you have the possibility to specify whether you want the volume “locked”. What does this mean? When the volume is set to a too high level such as to cause saturation in the output sound card DAC, then the setting of the volume slider is lowered automatically. This is not always desirable, e.g. when there are strong atmospheric crashes caused by a nearby thunderstorm. You would end up with a too low volume level. In such cases it is useful to “lock” the volume slider, which means that the automatic adjustment is not performed. Of course you are still free to adjust it manually.

A further choice is that of the Process Priority of Winrad. There is no a firm rule for this. Try both settings if you are experiencing audio dropouts, to see which one works better for your setup.

Another thing you have to select before starting is the sampling rate to use, and this can be done with the “Select Sample Rate” button, as shown in the following picture.



Figure 6 - Selection of the sample rate

When using the ASIO drivers, for some rates you may receive an error message stating that the current ASIO driver does not support that rate. It happened to me with the old driver for the Delta 44 card. The latest driver seems to have expanded the range of the supported rates. The output sampling rate is automatically chosen by the program to be 8 or 11.025 kHz, depending on the input rate. If you have just one card which does not support different input and output rates and your operating system is Windows 98SE, or you don't rely too much in the resampling quality of Windows XP, then you can force Winrad to use the same sampling rate for output as for input. The internal processing is still done at 8 or 11.025 kHz, but a final upsampling is done just before sending the audio data to the card for reproduction.

When you press the Start button, the program starts to acquire audio data, processes them and sends them to the DAC of the audio card. The upper window is divided into two panes, the waterfall and the spectrum. It is possible to choose how to assign the available space between the two panes by using the right mouse button. Just place the mouse pointer on the frequency scale. The pointer will change shape, indicating that you can drag left and right the portion of the spectrum that is displayed (keeping the **left button** down) or you can drag up and down (keeping the **right button** down) the frequency scale which divides the spectrum from the waterfall. Much more complicated to explain than to do...

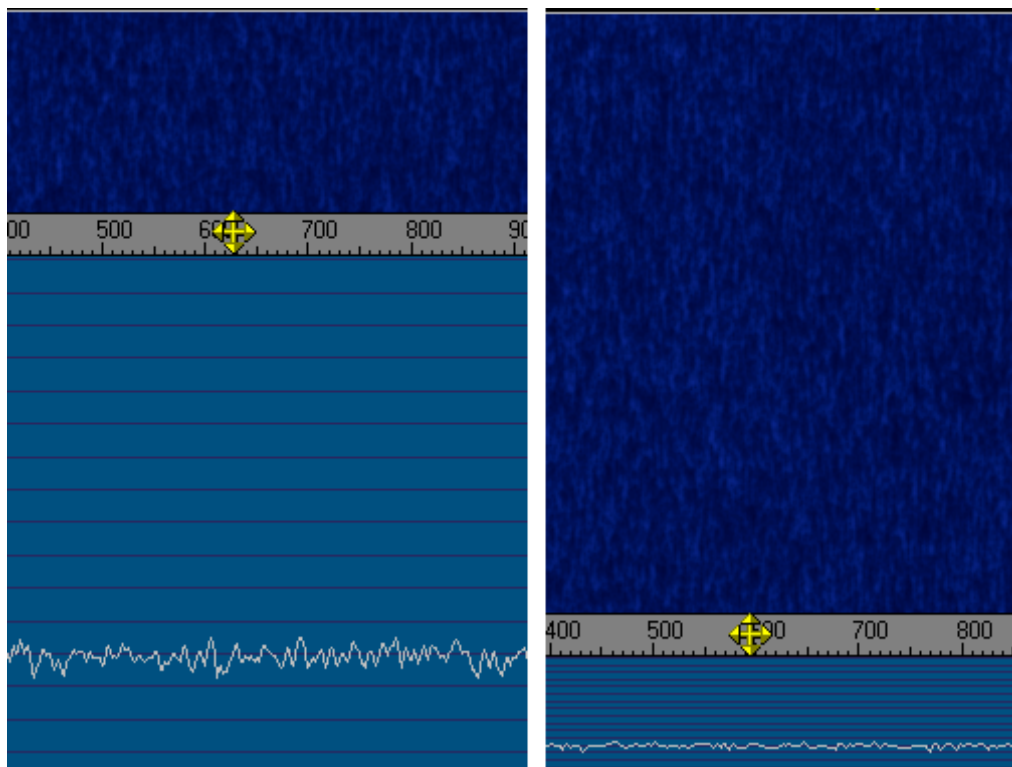


Figure 7 - Two different ways to divide the space available in the upper window

Note that the effect of the drag (horizontally or vertically) will be apparent only when Winrad is started, i.e. when the waterfall is flowing. Similarly two sliders appear at both sides of the upper window only when the program is started. The left one controls the gain level and the right one the base level of the spectrum display, much as you would find on a hardware spectrum analyzer.

Ok, now the program is running, and you have to choose the decoding mode. Some purists will object to the use of the word "decoding", as SSB signals are not encoded in any way, they are just translated in the spectrum. But please let me use this word to signify the process to bring them back to baseband. You choose the mode simply pressing one of the three buttons :



Figure 8 - Choosing the decoding mode

Currently only the USB, LSB and CW modes are implemented. Time will tell if other modes should be added. No matter what your choice is, you will see in the upper window a light blue area, which corresponds to the limits of the passband currently set, and shown graphically on the lower window.



Figure 9 - How the passband is shown in both windows

Note that the shape in the lower window is not an approximation of the passband. It is the *real* passband, dynamically computed and graphed while the program is running. Even the ripples due to the Gibbs phenomenon are visible. As you can see in Figure 9, it has an out-of-band rejection of more than 160 dB, thanks to the 1537 taps equivalent FIR filter used for the computation of the fast convolution.

Changing the limits of the passband is quite easy. Just place the mouse cursor on the upper or lower limit, the cursor shape will change indicating that you can now drag that border, as shown in Figure 10.




Figure 10 - Dragging the upper passband limit with the mouse

To drag that border, simply press the left mouse button and move the mouse keeping the button pressed. Similarly for the lower border. You can do this operation only when the program is processing input, i.e. when the Start button has been pressed.

To tune the desired signal while in SSB mode, click in the upper window on the frequency of the (suppressed) carrier of the SSB signal. This frequency is shown with a gray line and reported numerically above the window. Alternatively, you can drag the light blue window itself, by placing the mouse cursor inside it. The cursor shape will change and dragging is now possible. Fine tuning is possible with the mouse wheel, in steps of 10Hz.

The upper and the lower frequency scales

The upper scale indicates the frequency of the internal NCO (Numerically Controlled Oscillator) used to bring to zero IF the signal, summed to the frequency of the external LO. Hence, for SSB modes, it corresponds to the (suppressed) carrier frequency. For CW it corresponds to the CW signal frequency. The NCO frequency is the frequency of the internal software oscillator used to bring to baseband the received signals, and its value, as said above, is summed to the value of the external LO so that to indicate the tuned frequency in the display marked Tune :



The image shows a digital display with two frequency values. On the left, '07.030.000' is displayed in a blue, glowing font, followed by the word 'Tune' in a smaller, light blue font. On the right, 'LO 07.050.000' is displayed in the same blue, glowing font. The background is black.

Both those values can be changed using the mouse. If you hover the mouse on one digit of either displays, that digit is marked as being changeable. You can change it either by rotating the mouse wheel, or by keeping pressed the left mouse button (continuous auto increase) or the right mouse button (continuous auto decrease)



A close-up of the frequency display showing '07.032.000 Tune'. A mouse cursor is positioned over the digit '2' in the third position after the decimal point. The digit '2' is highlighted with a small yellow rectangle, indicating it is the active digit for modification.

When Winrad is used with a companion DLL meant to control an external hardware, like the SDR-14, SDR-IQ or the Winrad-RX, changing the LO value with the mouse has the effect to change the frequency where the hardware is tuned to.

The lower scale indicates the offset from the frequency indicated on the upper window, so for SSB it shows the audio frequencies of the signal received. The scale reversal when in LSB mode is automatically taken care of by the program. When in CW mode, the lower scale indicates the pitch of the CW note, which is user-settable. In CW mode a faint gray line appears in the center of the passband, in correspondence with the pitch set. When in this mode, the tuning can be refined by clicking on the lower window (spectrum or waterfall, doesn't matter). The frequency clicked on will be made the new carrier frequency, and hence positioned under the gray line, in the center of the passband.

To change the desired pitch frequency (default 550 Hz), position the mouse cursor in the lower window on the desired pitch, then press the right mouse button, while keeping pressed the Ctrl key on the keyboard.

Filters

When needed, a few filters can be applied to the signal, to make the reception better. One of the audio filters that it is possible to insert when receiving is the Denoiser, or Noise Reduction Filter, that is activated with the **N Red** button as shown in Fig. 11 on the next page.



Figure 11 - The control buttons panel




Figure 12 - The adjustment of the denoiser level



Figure 13 - The adjustment of the Noise Blanker level

Its aggressiveness can be controlled by a slider that appears when right-clicking on that button. See Fig 12. This denoiser is especially effective in presence of the so-called AWGN (Additive White Gaussian Noise), in other words the common hiss. Its effectiveness with impulsive noise is limited. For this type of noise the **Noise Blanker** filter is more indicated. A word of caution here. The Noise Blanker, if inserted in absence of impulsive noise, can make the reception worse... so use it only when needed, starting with its Strength set to the minimum, and slowly raising it up to the point where the reception becomes possible. On one of my Web sites there is a wav file that you can download to test the working of the Noise Blanker.

Get it from http://sundry.i2phd.com/CQ_EME.wav

It is a large file, slightly more than 10MB. When processing it, set first the Mode to I/Q, the LO frequency on the upper window to zero, and the Tune frequency to -12343Hz, CW mode. Try both with and without the Noise Blanker. Insert also the Denoiser and the CW Peak filters. Without the filters, it will be impossible to copy the call sign of the caller. With the filters active it will become an armchair copy. The file has been produced using a similar file downloaded from the Flexradio Web site.

The CW Peak filter, just mentioned, is a combination of a dynamic expander and a peaking filter, and can be activated only when in CW mode. The peaking filter has the shape shown in Fig 14 below, simulated with Matlab.




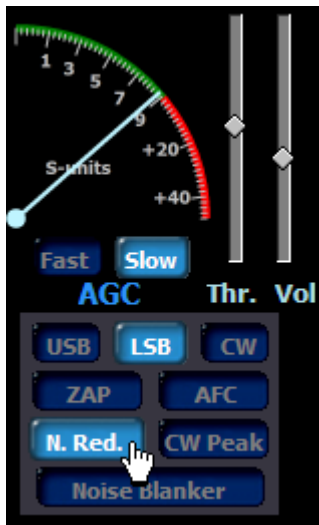
Figure 14 - The shape of the CW Peak filter

It is implemented as a simple IIR resonator. You can control its “pointedness” with the slider that is shown when you right-click on the CW Peak button. It’s the digital equivalent of a Q-multiplier in the analog domain. Differently from the analog Q-multiplier it cannot be pushed into oscillation, as this is prevented by the software, but if you push it to the extreme, the peak will become so narrow to not leave the CW manipulation pass... dots and dashes will sound as a semi-continuous tone, the filter is ringing, as any IIR filter has the rights to do...

To make the best use of this filter, pay attention to the fact that the CW signal is centered in the passband in the lower window, in other word it is positioned on the gray line indicating the CW pitch. You can use the ZAP button (described below) to do this automatically.

The other controls

Figure 15 - The control buttons



In Fig. 15 at the left, the complete set of the control buttons is shown. Starting from the top, the Level meter indicates, in S-units, the strength of the signal, to which it is logarithmically related. The full scale corresponds to the saturation level of the sound card ADC. Right to it, the slider AGC Thr. sets the knee of the AGC slope. Above this knee, the slope is kept almost horizontal, so to have a constant audio level. Below the knee, the slope is almost linear, in the sense that the differences in signal strength are no more compensated. Pay attention to the fact that the slider has a reverse behavior : when at its minimum, the knee is high, and viceversa. This has been done to make its use more intuitive. Try to use it when listening just at the weak background noise, and it will become clear.

The Vol slider is just a volume control. Its background will turn red when the level has been set too high, and there are audio peaks greater than the full scale value of the sound card DAC (which will be clipped).

Fast and Slow just indicate the release time constant of the AGC. The attack time is fixed and quite fast. A very short hang has been implemented, which goes almost unnoticed.

The ZAP button can be used only when in CW mode, and it searches for the strongest peak in the passband in the lower window, and adjust the NCO frequency so to center that peak, which will now be set at the CW pitch frequency. It doesn't stick, it has a momentary action. When you want to use it again, you simply press it again. It can also be activated by pressing the key Z on the keyboard. The AFC button, also active only in CW mode, tries to keep the signal peak centered in the passband, continuously adjusting the LO frequency. In this release of Winrad, the speed of this “frequency hunting” is fixed, but probably in one of the next releases it will be made adjustable.

At the bottom of the figure there are the two averaging controls, AvgSP1 for Spectrum 1 (the upper window) and AvgSP2 for Spectrum 2 (the lower window). The averaging can be done either on the spectrum only, or on both spectrum and waterfall. See later on for the relative controls.



Figure 16 - The waterfall Gain and Contrast controls

Both windows have two controls, Gain and Contrast, with which you can optimally set the mapping between the waterfall brightness and the intensity of the signal. If properly adjusted, the waterfall can show traces of signals so weak to be inaudible. If you know my other program Spectran, the behavior is similar.



Figure 17 - A few controls below the upper window

With the speed control you adjust the refreshing speed of the spectrum and the waterfall displays. The /10 button divides the current speed by ten, so to allow for very low refreshing rates. The button marked F (which means Full) unlocks the spectrum display speed from that of the waterfall. If activated, the spectrum display will run always at full speed, while the waterfall speed is set by the controls. The button marked Rev (Reverse) makes the waterfall flowing from bottom to top (never heard of water falling up instead than down, but some prefer this behavior). And finally the button WF Avg applies the averaging not only to the spectrum, but also to the waterfall.

The two arrow buttons increase or decrease the resolution of the display. The Hz value indicated is the width of the FFT bin, which always corresponds to a single pixel. Making the resolution finer increases the CPU load, and makes worse the time resolution, so adjust it wisely.



Figure 18 - The controls under the lower window

Under the lower window there is the speed control, similar in function to that of the upper window, the left and right arrows to control the resolution, and a button marked F (which means Full band), that controls whether the spectrum displayed must be windowed by the passband limits or not. Look at the two different cases :




Figure 19 - The F button is **not pressed**



Figure 20 - The F button **is pressed**

Of course, the F button affects only the display, not the audio, which will always be affected by the passband limits set. The button WF Avg is similar in function to the similar button below the upper window, but now it relates to the lower window waterfall.

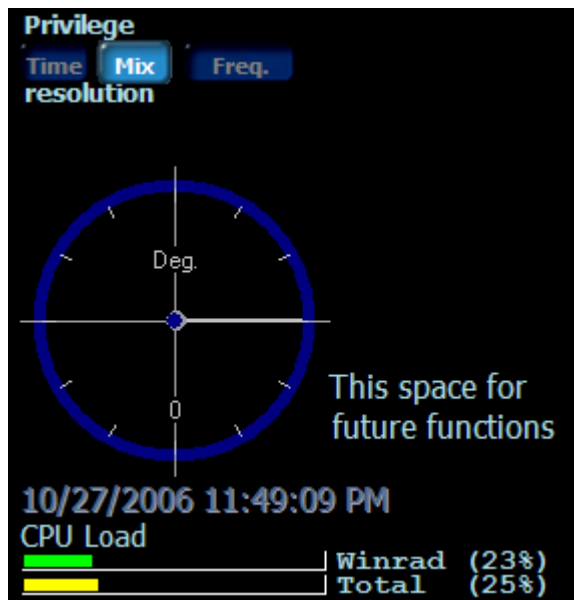


Figure 21 - Some other controls

Using the upper three buttons in [Fig. 21](#) it is possible to tailor the spectrum and the waterfall displays in both windows so to give more importance to the frequency resolution or to the time resolution, or to compromise between them. Make your choice depending on the characteristics of the signal. For normal speed CW signals, you may certainly want to leave it set to Time.

The round phase display is for functions to come.

The two bars at the bottom indicate the current CPU load, both that imputable to Winrad, and the overall load caused by the sum of all the programs running, Winrad included. This display is available only when Winrad runs under Windows 2003 or XP, but not Win98

nor Windows 2000.

Some caveats

Now a few words about the sound card sampling frequency, and the consequences of errors in its value.

If this world were perfect, a sampling rate of 48 kHz would be exact, and so would be 8 kHz. If this were the reality, then I could produce one output buffer at 8 kHz out of 6 input buffers, and everything would work ok. Unfortunately this is not how things are... the sound card sampling rates have errors, and the consequence is that, if I downsample exactly by a factor of 6 going from 48 to 8 kHz, I would have either buffer starvation or buffer overflow depending on the real ratio being lower or greater than 6.

To remedy to this problem, the only solution is to monitor the number of buffers ready for output and to change dynamically the decimation ratio so to keep this number constant. This value is of course a discrete, integer, number, not a continuous variable, and this complicates a bit the computation of the optimal decimation ratio. That notwithstanding, when real sound cards are involved, the algorithm converges in an acceptable lapse of time, and then the decimation ratio becomes a rather stable number.

With using e.g. the SDR14 hardware the input sound card is replaced by the USB port, which is where the audio come from. Under Windows the USB port is not designed for a constant, steady flow of data, but rather as a block device, where large amount of data are pumped a chunk at a time, with some unavoidable jerkiness. This perturbs that computation, and the result is that the computed optimal value cycles around a mean value, and it is possible to observe this effect in the lower window. Here you see the frequency changing slightly, as a variation of the decimation ratio involves a different final frequency for the signal.

The value of the decimation ratio is shown in the Status panel, that can be opened by a choice in the ShowOptions menu.

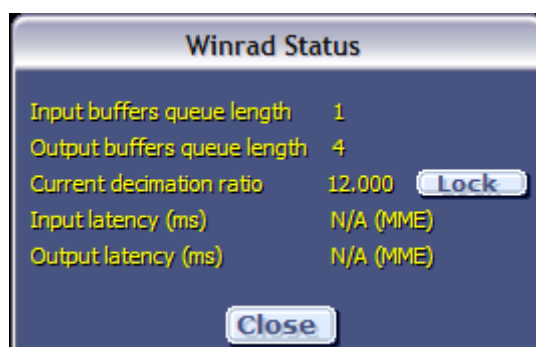


Figure 22 - The Status panel

Winrad tries to keep the value of the output buffer queue length to 4, by adjusting the decimation ratio. The button marked "Lock" can be used to lock temporarily the current value of the ratio. This will stop the wandering of the output frequency, but of course it will eventually result either in the queue length value going to zero (causing clicks in the audio) or increasing up to the maximum allowed, causing some automatic discard of buffers to avoid overflow.

If you see that Winrad has problems in reaching a stable value for the decimation ratio, there are a couple possible remedies. One is to "Force" the output sampling rate to be equal to the input, using the same card both for input and output (if you use two different cards, there is no guarantees that the two rates are equal). The other is to use the ASIO drivers for the card, if they exist for that specific model of sound card. Using the ASIO drivers, the input and output interrupts on the card are coupled, so they are guaranteed to happen simultaneously, exactly at the same time.

Well, this more or less ends this synthetic User Guide, which I hope will make you able to use Winrad in an acceptable way. I would like to thank Jeffrey Pawlan WA6KBL, whose invaluable advice has been of great help during the development of the program, and thanks also to the beta testers who gave suggestions on how to improve Winrad.

If you live in the US, and would like to have guidance on the use of Winrad from a person who speaks your own language, you may contact Jeffrey at jpawlan@pawlan.com .

If you are in an European Time Zone and prefer to contact me, my email address is i2phd@weaksignals.com .

Winrad is freely downloadable from <http://www.weaksignals.com>

Enjoy the program and please give feedback, positive or negative, about its use by joining the Winrad discussion reflector, which you can do by visiting <http://join.winrad.org>. Thanks.

Alberto I2PHD
i2phd@weaksignals.com

ASIO is a trademark and software of Steinberg Media Technologies GmbH

Appendix A : SDR-14/IQ support

Starting with V1.0 Winrad supports external hardware through specific DLLs. The first such DLL written is for support of the SDR-14 and the SDR-IQ (see <http://www.rfspace.com>).

To activate this support you must download from my Web site also the SDR14-IQ.zip file, which contains :

SDR14X.dll
ExtIO_SDR14.dll
Register.bat

Please place these three files in the same directory of Winrad (important), and execute Register.bat which will register with Windows the base OCX for supporting the SDR-14 or the SDR-IQ.

You can power on the SDR-14 either before or after launching Winrad, it doesn't matter. If all is well, when you press ShowOptions | Select Input you will see that among the possible inputs, also the SDR-14 (or the SDR-IQ) is listed :

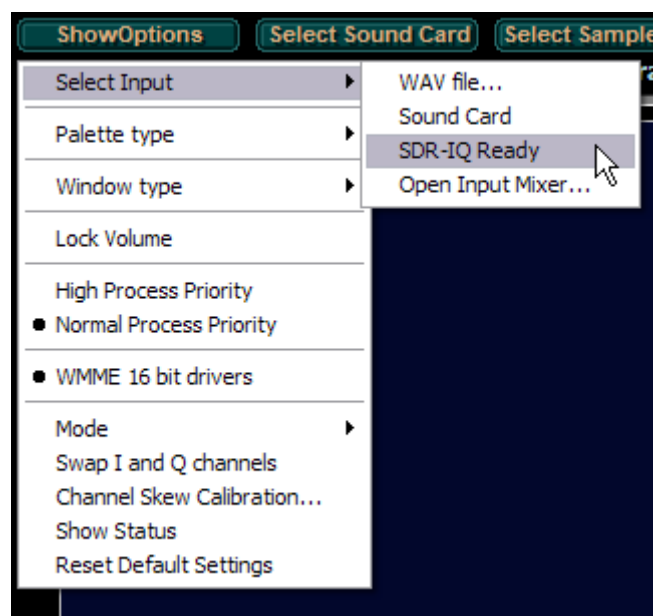


Figure 23 - How to select the SDR-14/IQ as input device

Clicking on that choice will open the SDR-14/IQ control panel :

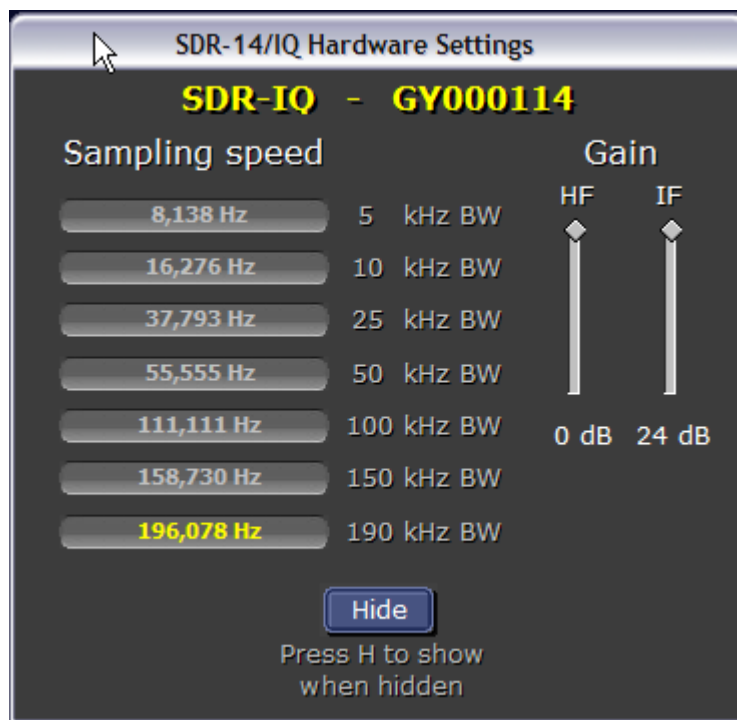


Figure 24 - The SDR-14/IQ Control Panel

Here you can select the Sampling Rate, and the Input connector of the SDR-14, and set the HF and the IF Gain of the device. Refer to the SDR-14/IQ manual for explanations.

Note that when using the SDR-14 as input device, the ASIO drivers choice of Winrad is disabled, and only the WMME drivers can be used.

To change the frequency the SDR-14/IQ is tuned to, change the value of the LO frequency display on the upper part of Winrad :



When you change it, also the Tune display will change accordingly, to indicate the correct reception frequency
