



# Data Modes

This month Mike Richards G4WNC, continues his look at SDR radio systems, with some tips on how to make the most of your SDR software.

Welcome to *Data Modes (DM)* and – as you'll have seen from earlier articles in this series – there are lots of different software programs for software defined radio (SDR). Additionally, the good news is that the vast majority are free! However using SDR software to monitor radio signals requires a different approach to a traditional communications receiver.

The first culture shock is the loss of the ubiquitous tuning knob. While most systems provided a simulated tuning knob – it's not usually the best way to get around the bands. If you are desperate for a conventional tuning knob there are a few 'tuning knobs' available that connect, linked to your PC via a USB lead, but you will have to check compatibility with your set-up.

However, it's important to explore the new opportunities presented by the much more flexible interface of SDR software. By far the most common tuning technique is known as 'point and click'. This technique is where you use the mouse pointer with the spectrum display and simply click on the part of the display you want to monitor.

The technique is a particularly quick way to get close to where you want to be and you can usually complete the fine-tuning by 'click-dragging' the tuning point with the mouse or by rolling the mouse wheel. However, there are a few settings that need to be fine-tuned to help you get the best from the interface.

### Spectrum Display


There's usually a spectrum display, that dominates most SDR interfaces and usually set shows the entire bandwidth of the IQ signals coming from the SDR hardware. You may well find that in its default set-up the display is very animated and erratic, making it difficult to make sense of what you're seeing.

The solution to this rapidly moving display, is to set the display to show the signals, averaged over time. When averaged, each point on the display will


depend on previous readings, which has the effect of smoothing out the movement and making observation very much easier. The critical point is, of course, how much averaging to use – too much and short duration signals will disappear – too little and the display will still flicker.

is very difficult for me to recommend a specific setting as it depends on the software you're using and the speed of your PC. A good starting point can be found by tuning to the Morse (c.w.) segment of a busy band and experiment with the averaging until you find a setting that keeps c.w. signals visible, but smooths-out the flicker. I've shown an example of the difference averaging can make in Fig. 1.

The screen shots shown in Fig. 1, show the spectrum plots of a WinRadio Excalibur Pro with no averaging and then with 0.4s averaging. As you can see, the averaged plot is very much cleaner and the signal traces are more



Raw display - no averaging



0.4 second averaging

Fig. 1: Averaged and non-averaged spectrum displays, note with the averaged display (bottom display), it's much smoother to read.

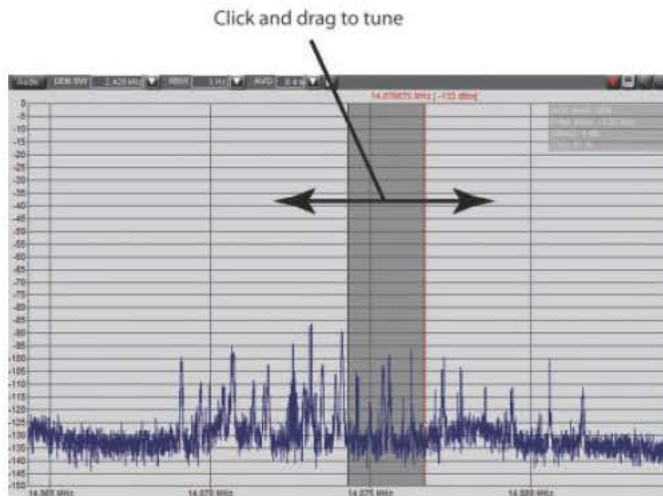


Fig 2: Fine tuning a spectrum display.

obvious. If your software uses the number of samples to set the averaging to three or four, a good starting point.

Most SDR software will show the demodulator bandwidth as a highlighted, or grey bar, superimposed on the spectrum display. This can be dragged from side to side to complete the fine-tuning, see **Fig. 2**. On most systems you can also use the mouse wheel to scroll around. If you happen to have a graphics tablet handy this can also be used for tuning and I find a tablet quicker and more precise than a mouse.

The receive bandwidth bar is often adjustable simply by clicking and dragging on one edge of the bandwidth display, see **Fig. 3**. On some systems this has to be done in the audio rather than the r.f. spectrum display. This feature is particularly powerful, allowing you to trim the receive bandwidth carefully to minimise interference. But I'll cover filters in more detail in a later article.

### Waterfall Display

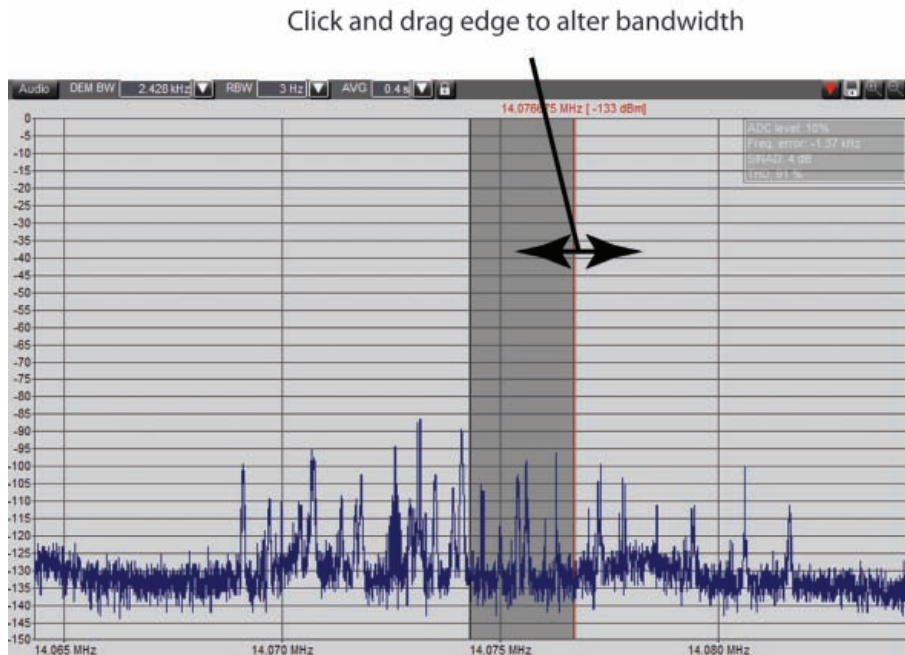
The waterfall display found on many SDR programs, is one of the most useful features, as it provides a record of a signal's spectrum over time. Instead of showing the spectrum as a varying line graph, the waterfall shows different levels in the spectrum as points of varying brightness – the stronger the signal, the brighter the trace. I've shown an example of a spectrum display and the corresponding waterfall display in **Fig. 4**.

The real power of the waterfall display comes when you allow the display to build up for a few minutes. Any intermittent signals show as dashed lines and you will note that different types of transmissions have their own characteristic patterns. This is another subject I'll be looking at in more detail later.


The waterfall display is also very good for click-tuning as it's much easier to click on a stable trace on a waterfall than a signal that's bobbing about within the spectrum display! This is especially true if the station has just finished a transmission as the trace will still be visible in the waterfall but no longer show in the spectrum display.

### Display Resolution

In addition to the controls I've already discussed, there are a couple of other settings associated with Fast Fourier Transform (FFT) displays that often cause confusion. The first setting is the display resolution, which is set by a combination of sample rate and the FFT



**Fig. 3:** Adjusting the demodulator bandwidth.



**Fig. 4:** Spectrum and waterfall displays compared, the spectrum above the waterfall display.

size. Let's quickly re-cap on the FFT process to see how sample size and FFT size relate.

The FFT technique is a well established and very useful mathematical algorithm that can process a stream of digital data to separate-out the frequency components. In very simple terms, it does this by generating a number of memory stores (known as bins or points) each one of which, represents a narrow band of frequencies. As the data is processed, each bin is allocated a number that corresponds to the energy level of its particular frequency band.

The output of these bins is then used to drive the spectrum and waterfall

displays. The bandwidth of each bin is determined by the sample rate of the signal and the number of FFT bins with the simple formula:  $\text{Bandwidth (in Hz)} = \text{sample rate} / \text{number of bins (or points)}$ .

As always, there are compromises to be considered. The larger FFT sizes take longer to fill, so there's a time lag

between the signal being received, and it showing on the display. The optimum setting will be specific to your set-up and preferences –but feel free to experiment!

The method of adjusting the display resolution varies between systems and whilst the *Excalibur Pro* and *HSDR* show the bandwidth in Hz via a simple drop-down menu, many SDR software packages only offer FFT size selections with typical options being: 512, 1024, 2048, 4096, 8192, etc.

Here's an example to show you how to work out the display resolution. If the incoming sample rate is 96kHz and the FFT point size is 4096 then each of the FFT bins and therefore the display






Fig. 5: 1Hz resolution analysis of a PSK31 signal using the Excalibur Pro receiver and its software.

resolution will be  $96,000/4096 = 23.4\text{Hz}$  wide. Whilst that resolution would be okay for general use, finer resolution is necessary if you want to take a close look at narrow data signals such as PSK31 that are only 31Hz wide.

In the case of trying to sample a PSK signal, you'd probably need to take the FFT size up to 32,768, as that would provide a more useful resolution of 2.7Hz. I've shown an example of a PSK31 signal analysed with 1Hz bandwidth in Fig. 5, where you can see the waterfall display. The display is very revealing and provides a powerful analysis tool that can be used to check the transmission quality and identify the transmission mode but I'll cover this later.

### Further Confusion

The further source of confusion concerns FFT windows. Why do we need to worry about these windows? For the FFT to work accurately and drop all the signals neatly into the correct bins each signal in the FFT sample has to be periodic within each data sample. That just means each signal should start and finish at 0. As we're dealing with a wide variety of unknown radio signals, clearly that will not be the case.

If the signals aren't periodic within each measurement sample, then the allocation to bins becomes 'smeared'

and a single frequency signal will appear in more than one bin. This effect is known more technically as leakage. If left unchecked, this will spoil the resolution of our display.

The solution is to intervene and force all signals to zero at the beginning and end of each measurement sample. Of course, there are lots of ways in which the signal can be reduced.

Signals could be either abruptly stepped down to zero, or they can be gently reduced over a period of time. As you can probably imagine, there are countless techniques that could be used but a number of standard shapes have been developed over the years and these shapes are known as

windows and are generally named after the people that developed them, hence the window types called **Blackman**, **Hanning**, **Hann**, **Nuttall**, etc.

I've shown two common examples for such windows in Fig. 6. As always, selecting the window type involves compromises but for general radio use Hanning is probably the best starting point as this provides good resolution with minimal leakage. If you want to achieve accurate measurement of a single tone signal, then a flat-top or rectangular window is likely to give the best result.

That's it for this month, next time I'll be taking a look at SDR filters. Cheerio for now!

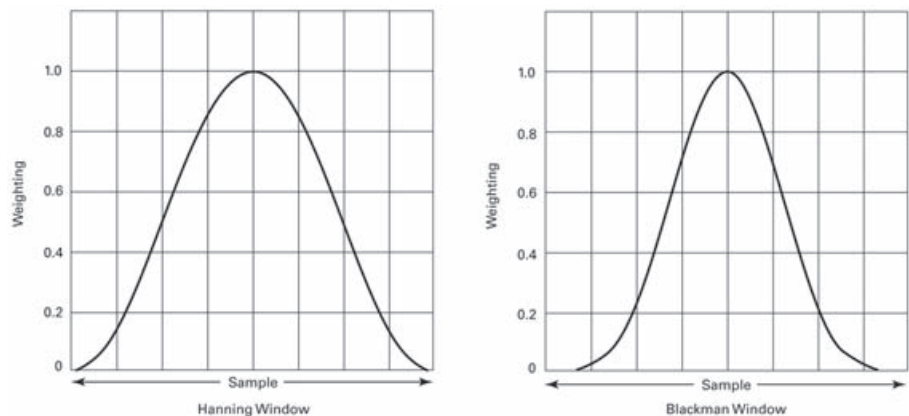


Fig. 6: The Hanning window shape to the left of the Blackman FFT window shape on the right.