

“What’s all this Bandscope Stuff, Anyhow?”

by Roger Rehr, W3SZ

I guess the first question, for those who don’t know the answer, should be, “What’s all this Bandscope Stuff, Anyhow?” A bandscope is a graphical representation of the signals received over a range of frequencies that is significantly larger than the received audio bandwidth. To be really useful, the bandscope should cover a wide range of frequencies. 48kHz would be minimal useful range; I think having 192 kHz is optimal for weaksignal use, as long as there is the ability to narrow the visible bandwidth and “zoom in” when necessary. Also, the bandscope should have two components, with both a realtime spectral display where frequency is on the X axis and signal intensity is on the Y axis, and a waterfall display, where frequency is also on the X axis, but where the Y axis is time, and signal intensity is expressed by varying the color of the signal pixels. Don’t worry now about what all of this looks like. There will be plenty of examples later in this article.

My interest in Amateur Radio is weak signal communications on frequencies above 50 MHz. I am interested in both terrestrial and Earth-Moon-Earth (moonbounce) communications, but this article will be limited to terrestrial work.

There is not a lot of activity on a daily basis on these bands, and antennas and arrays tend to be highly directional. For these reasons, operating during non-contest event time periods is not highly rewarding, as there are not many other stations active during these non-contest periods.

Thus, my operating activities [but not my experimental activities] are limited to contest periods. In order to get the most out of operating a contest, it is necessary to know at all times what is happening on each band in terms of several different parameters. [1] overall activity on that band, [2] whether or not that band is “open” with enhanced forms of propagation, and [3] whether or not there are new stations on that band that have not been previously worked are examples of band characteristics should be known for every band at every point in time.

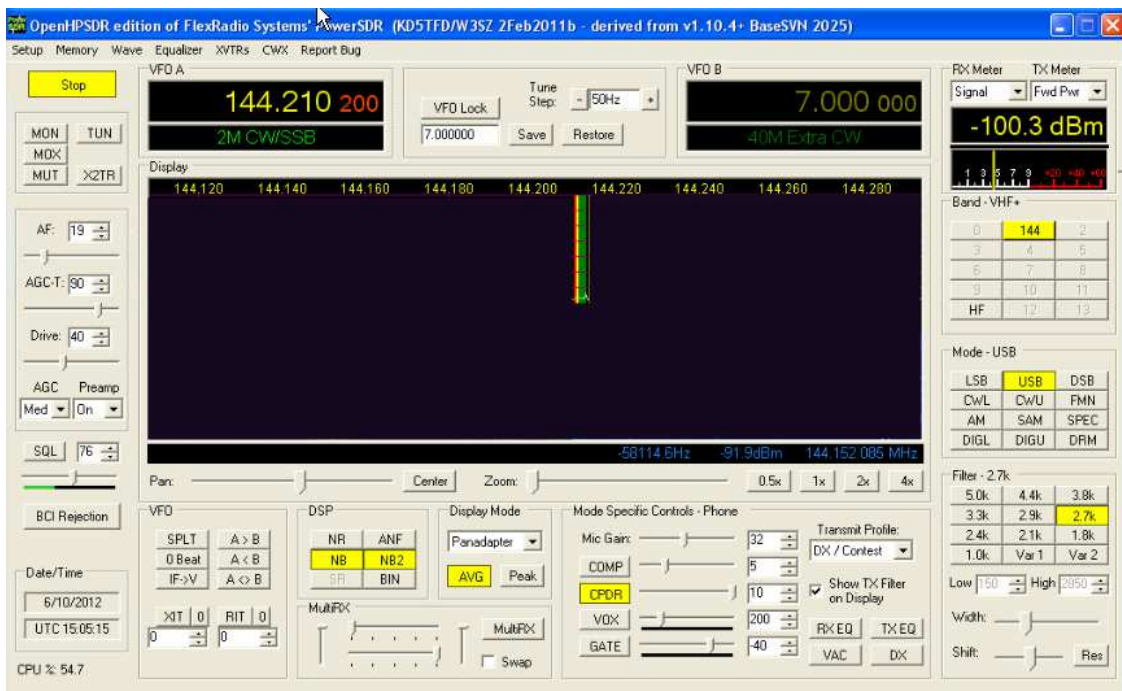
If one is operating a conventional radio, without a bandscope, it is like looking at the world through a narrow tube. At a given time, one hears only a very small, perhaps 2.3 kHz, portion of the spectrum available for use during the contest. The rest of the radio spectrum, in the band on which one is operating, and in all of the other bands as well, remains a black hole, invisible to the station operator.

The “picture” or information that one gets with such a conventional radio without a bandscope is like what one would get with a software defined radio [SDR] with most of the bandscope blacked out. With my radios, I generally have a bandscope of about 192 kHz operating, so by not using such a radio one is seeing only $2.3/192 = \text{ONE PERCENT}$ of the information one would see with an SDR with a 192 kHz bandscope.

This concept of being ignorant of important information when using a conventional radio without a bandscope can be expressed graphically:



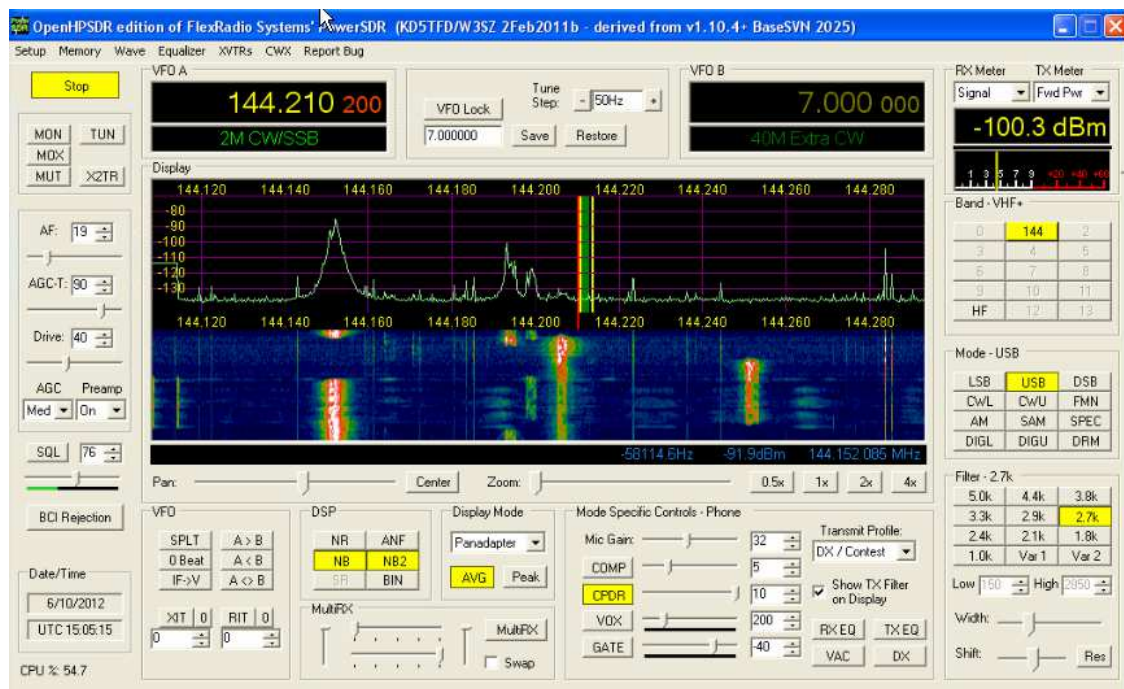
EQUALS:



EQUALS:



I am not saying that the conventional radio used for this illustration is not a good radio. It is a great radio, and it is what I use for 144 MHz terrestrial work both during and outside of contest times. But, I use it with a bandscope, so that I have a full appreciation of what is happening on the 144 MHz band at all times. Instead of the head-in-the-sand picture shown above, I instead see:




Instead of seeing / hearing no signal when I get on the band unless I happen to by good fortune land on top of a station, when I go to a band with a bandscope-equipped radio, I immediately see that there are multiple signals on the band. Furthermore, the bandscope, with its waterfall, gives me an instant picture of what is happening not only over a wide range of frequencies, but also over a wide range of times. You can see that at some times, there was just one signal [or perhaps two] on the air. But the waterfall acts like a time machine, allowing the operator to look back into the past, and see what was happening over a range of times.

This is a major thing. When you tune a radio without a bandscope, you hear only what is on the ONE frequency you are listening to at the ONE time you are listening there. If someone is transmitting on that frequency before or after you tune to that frequency but not while you are tuned to that frequency, you won't hear their transmission. If someone transmits on any other frequency on the band other than the one you are listening to at the moment, you won't hear them or be aware of their presence on the band.

But with a radio that has a bandscope, you can set it up so that you see all the relevant frequencies all the time. And you can adjust the waterfall speed so that you can look back 1, 5, 10, or more minutes into the past and see all the signals that were on the air from the present back to that time 1, 5, 10 or more minutes in the past; whatever the conditions dictate would be most helpful to you.

Having a waterfall is also extremely useful when looking for weak CW signals. It is like being able to hear entire 180 kHz bandwidth, using 262,144 receivers, each with a bandwidth of 0.73 Hz, with no ringing, and no interference from all the frequencies containing only noise. On the bandscope, the very weak signals literally jump out of the noise on the waterfall. If your waterfall is set up properly, you can see signals that are much weaker than you can hear. So if you can't see it on the waterfall, then the signal is so weak that you will never hear it and will never be able to work it. So you don't need to "tune the band" to if there are any signals that you might work. If there is nothing on the waterfall, then there are no signals on the band for you to work. It is that simple. The picture below shows a bandscope with a birdie at 3456.000 MHz, and then a very weak CW signal that is near the noise level, but shows up nicely on the waterfall near 3456.040.



There are other benefits to using the single-band bandscope other than just using it for search-and-pounce station hunting and very weak-signal CW.

For example, you can be calling CQ and running a frequency on two meters, or doing search and pounce, and keep an eye on 144.250 waiting for the rover NN3Q to show up there as he said he would at his next grid-square opportunity. When he pops up, you will see him and you can give a call or take a quick listen to see if it is really him, and not waste time taken from your running stations or doing search-and-pounce looking for him before he really shows up there and is ready to go.

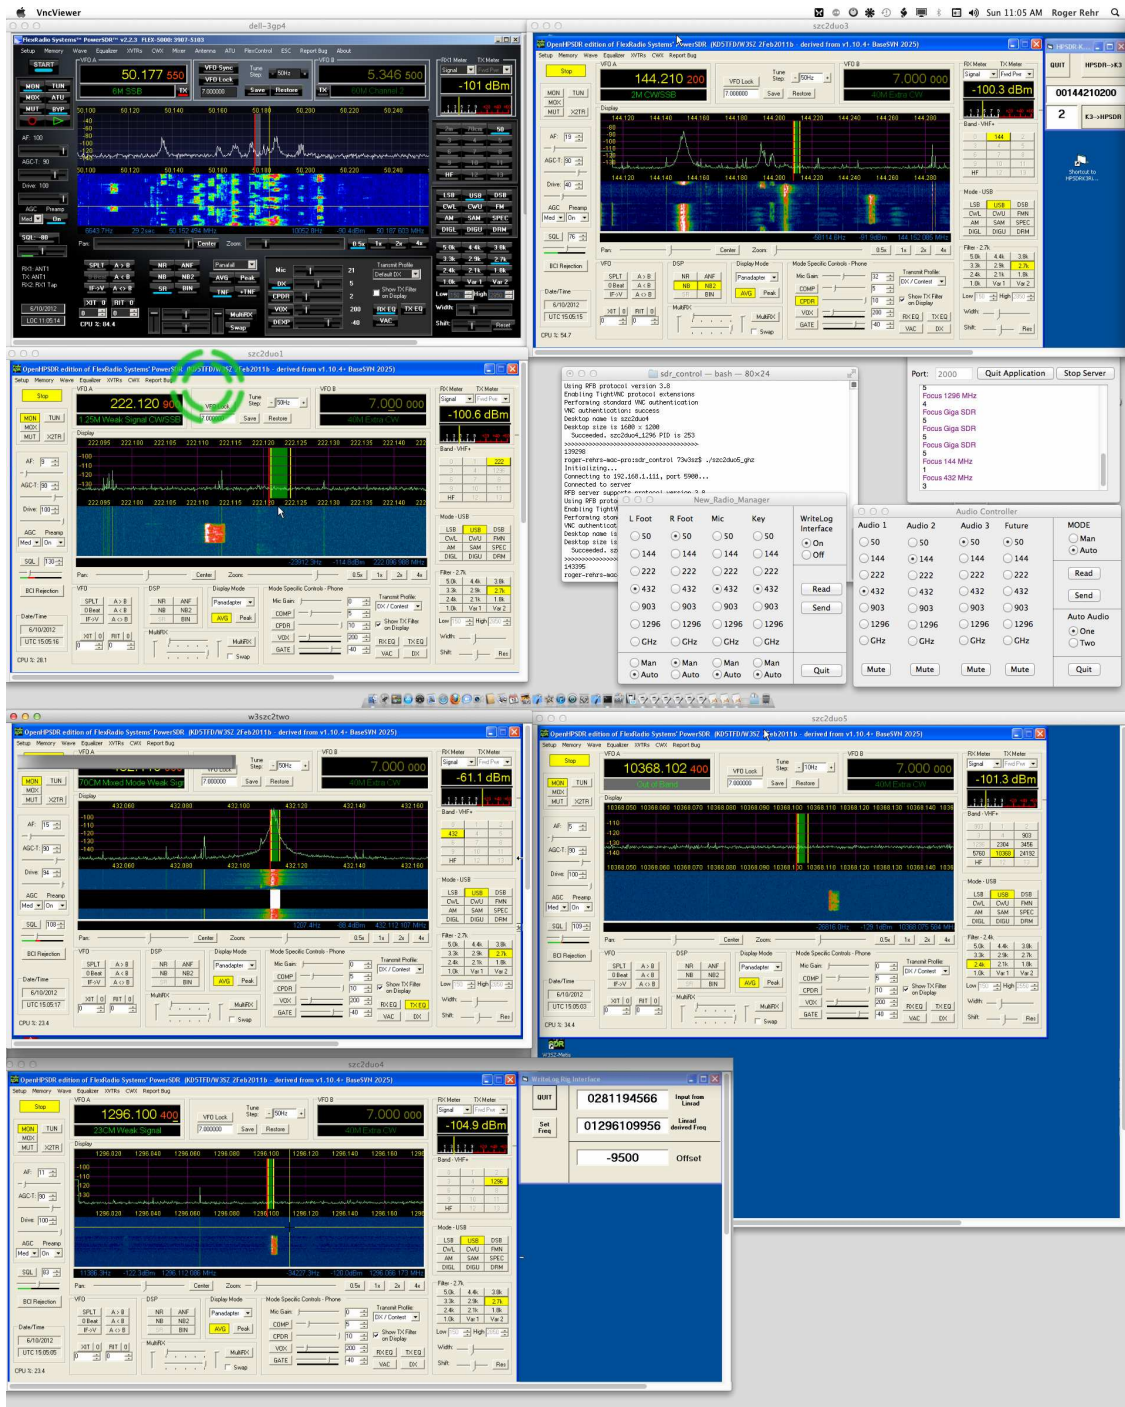
Or, you can be running stations on SSB and notice that a bunch of CW signals are popping up all of a sudden on another frequency on the band. Of course, that likely means that someone has discovered that there is some Aurora or other enhanced propagation starting, and a quick listen to their signal will allow you to figure out what is going on and take advantage of the Aurora well before you might have otherwise stumbled onto it.

Or, it might be that you have worked pretty much everyone on the band. You know that because you have [perhaps with the help of the N1MM bandmap] memorized the bandscope, and you know that every signal appearing is from someone you have already worked. But someone pops up way up at 144.270, likely just “passing through” as he runs the bands. In a split second you can click on his signal and give a call, and if he is a new station, you can work him and even set up a rendezvous to meet in a short while and run the other bands with him as well. Without the bandscope you would have never known that he was there.

If you want to call CQ on the band, you can see the clear spots (“holes” in the band) on your waterfall and jump right to such a clear spot and call CQ. You can put yourself into the “little holes” between stations which become very obvious with the bandscope. Having the history that the bandscope gives you is very helpful in avoiding frequencies where you will cause interference and be interfered with when you start calling CQ.

Also, if you have just switched bands while running the bands with another station and the other station that you are running the bands with is “off frequency” as you land on the new band, you will see him immediately on the bandscope when you jump bands, even though he is off frequency, and you can click on his signal and be on his frequency immediately, and not lose him or waste time trying to find him as you might if you didn’t have the SDR with its spectral display and waterfall.

So having a bandscope is great! What could be better than seeing everything there is to see on a single band? The answer is below:



This is the picture that you should have in front of you at all times during a contest: fulltime bandscope on the lower bands, and a shared bandscope for the microwave bands, so that you can see everything that you need to see to make the right decisions during a contest.

If you were looking at this during a contest, then you would have not only known what was happening on two meters, where you were operating, but you would have been able to see, without disrupting your rhythm on two meters, if there was any activity on 10 GHz, 1296 MHz, 222 MHz, or 432 MHz that might result in additional contacts. You

could potentially quickly get four contacts and multipliers on those bands and still hold your frequency on two meters!

But, the most **IMPORTANT** information that you would have been missing without seeing “all the bands all the time” is the information that six meters was wide open, and you really probably shouldn’t be on any of the just-mentioned bands, but rather on six meters accumulating multipliers and contacts far in excess of what you could be getting anywhere else, and which might not be available later after the opening closed. With always-on bandscopes on the lower four bands, you will know immediately when an opening occurs on any of these bands, and not miss out on it, wasting your time less productively elsewhere. If you are on band X and band Y opens, you will see it immediately if you are paying attention, and you will not miss the opening! One look at the six meter bandscope displayed above and it is obvious to anyone that the band is wide open.

There are other advantages to having multiple bandscopes visible at all times, which you will experience many times every contest if you go this route.

For example, similar to what I described for a single-band bandscope above, if you are jumping bands you can see on the new band what is a clear spot to jump to, and not “land on top of someone” when you jump to a new band, because you have the waterfall “history” of all of the frequencies on the new band right in front of you before you make the jump, and you just have to click on an open spot on the new band’s bandscope and you will be on the clear frequency on the new band, calling CQ or calling the station that you are running with in an instant.

If you lose a station during a microwave contact, you will see immediately if he pops up on any of the frequencies that you were using during your run on ANY of the lower bands, and you will NOT lose the chance to complete the running of the bands, because you are watching all of those frequencies all the time. The likelihood that you will lose a station because of a misunderstanding or because of “bad luck” drops substantially when you are watching every band that he has been on in the recent past, and you know what frequencies he was using on those bands.

There are other advantages to using the bandscope on the microwave bands, besides the ability to see all possible liaison or fallback frequencies simultaneously at all times. On the microwaves, having a bandscope allows you to reduce the number of variables when looking for a station: if he is anywhere within about 190 kHz, you will see him; so you can look for likely signals on the bandscope and rotate the array to try to peak them without having to solve both the frequency and heading variables at the same time. Ordinarily, when you are trying to find a weak station on one of the microwave bands, you need to worry about three variables: frequency, time, and beam heading. With a bandscope, you see all the frequencies all of the time, so that the only variable that remains is beam heading. This makes a huge difference and greatly increases the chances of your finding the weak station, and greatly reduces the time that it takes to do so. Also, with the bandscope waterfall, you can see signals that are too weak to hear or work, so on the microwaves you can find a very weak signal and then peak the heading with the bandscope and turn it into a usable signal and complete the weaksignal contact, when you never would have even been aware of the signal without the bandscope. You do this by finding the station with the bandscope, peaking its signal by optimizing your beam

heading, and then transmitting to the station on his frequency so that all HE has to worry about is beam heading. More often than not, what was a barely visible, non-workable signal becomes a signal that is Q5 at both ends when this procedure is followed. So you can make contacts that would not have been possible without the bandscope.

Furthermore, because you don't need to hunt for the other station as you are running the bands, you can move very quickly from band to band. I have run all the bands from 50 MHz thru 24 GHz in 5 minutes.

Additionally, you can tell if something in your system "breaks" on a band other than the one you are operating on at the time the problem arises, rather than discovering this problem later on when you are running the bands. You can do this by seeing a loss of signals and / or a drop in the noise floor on the bandscope for the band that has developed a problem while you are operating on another band. Having this knowledge in advance reduces the chances of a run up the bands ending prematurely because of your finding yourself in a black hole when you arrive on a band, with no signal heard, and the other station deciding it's not worth trying the higher bands because nothing was heard on the band in question. You can avoid the problem band, or at least forewarn the other station of a potential problem, and then continue up the bands whether or not you are successful on the "problem band".

So the first major point of this talk that I wanted to convince you of is that you gain a big advantage by having a bandscope. Having convinced you of that, the second major point that I wanted you to understand is that you gain a tremendous advantage by having simultaneous bandscopes on multiple bands.

Now that you believe both of these points, I want to introduce the third and most important major point of this paper: the radio must be the bandscope. I know this because over a period of several years, ending several years ago, I tried it the other way. I had bandscopes on the lower four bands and one for the microwaves, using a variety of SoftRocks, an SDR-14, and later some SDR-IQs thrown in too. The delays caused by having to switch focus from the radio to the bandscope and then back to the radio really added up over the course of a contest. Also, although I tried to streamline things by writing software so that I could click on a signal on any bandscope waterfall and then type a computer key to have that frequency sent to my IF radio, I still needed to make sure my station was on the correct band, that is the bandscope's band, and that I was logging on the correct band each time I changed bands.

Although operating this way gave me major advantages over the situation of not having the multiple bandscopes, it was still not ideal.

So I set my goals as follows:

To have constant bandscope monitoring of each of the bands 50, 144, 222, 432, and 1296 MHz, and a shared bandscope for the remaining bands thru 24 GHz.

To be able to select a band by simply clicking on its GUI or typing its frequency into N1MM, and have everything automatically switched so that I could immediately operate on that band, with no need for ANY manual bandswitching of anything, and with N1MM always tracking the frequency and mode in use and logging appropriately.

To have the bandscope BE the radio, so I didn't need to fiddle to bring the band or signal that looks interesting on the bandscope online to transmit or tune in the signal more carefully. After clicking on the signal all I would need to do to make the contact would be to step on the footpedal and speak, or start sending CW.

I have very much enjoyed using a station meeting all of the above goals for the past 2-3 years, and will describe how it does its job. To be more specific about the design criteria that allowed me to achieve the above goals, these criteria are listed below:

The bandscope is the radio - both Tx and Rx.

Always-on bandscopes are used for the lower bands: 50, 144, 222, 432 MHz and for 1296 MHz. A shared bandscope is used for 903 MHz and 2.3, 3, 5, 10, and 24 GHz.

Automatic switching of microphone, footswitch, receive audio, key/keyer to the appropriate radio for the band desired

Ability to manually set the assignment of each of the above I/O devices to a given radio if desired [to supplement or over-ride the automatic assignments for particular situations]

"Extra" footswitch and receive audio channels to facilitate microwave liaison duties

All radios are seen by logging program as one radio [Originally WriteLog, now N1MM].

Full integration of the radios and N1MM.

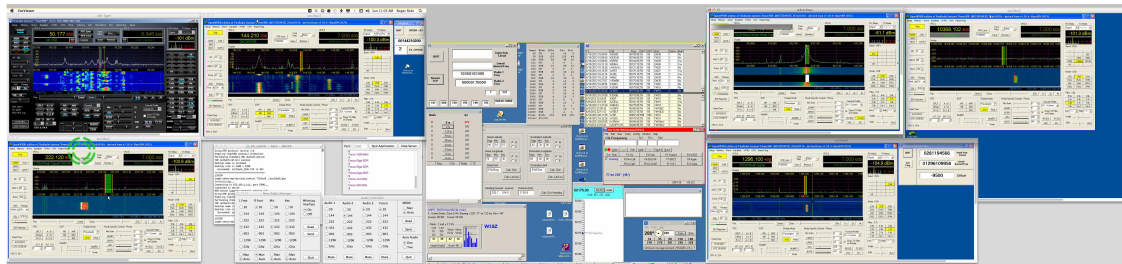
When a band is selected using N1MM, the appropriate radio is selected and all automatic switching performed

When a band is selected by clicking on appropriate radio's GUI, N1MM is set to that radio's frequency and all automatic switching is performed

Voice-Keying and CW Keying from N1MM are implemented and automatically connected with the appropriate radio

Software lockout is provided so that radio cannot switch bands while transmitting, and the radio cannot transmit while switching bands.

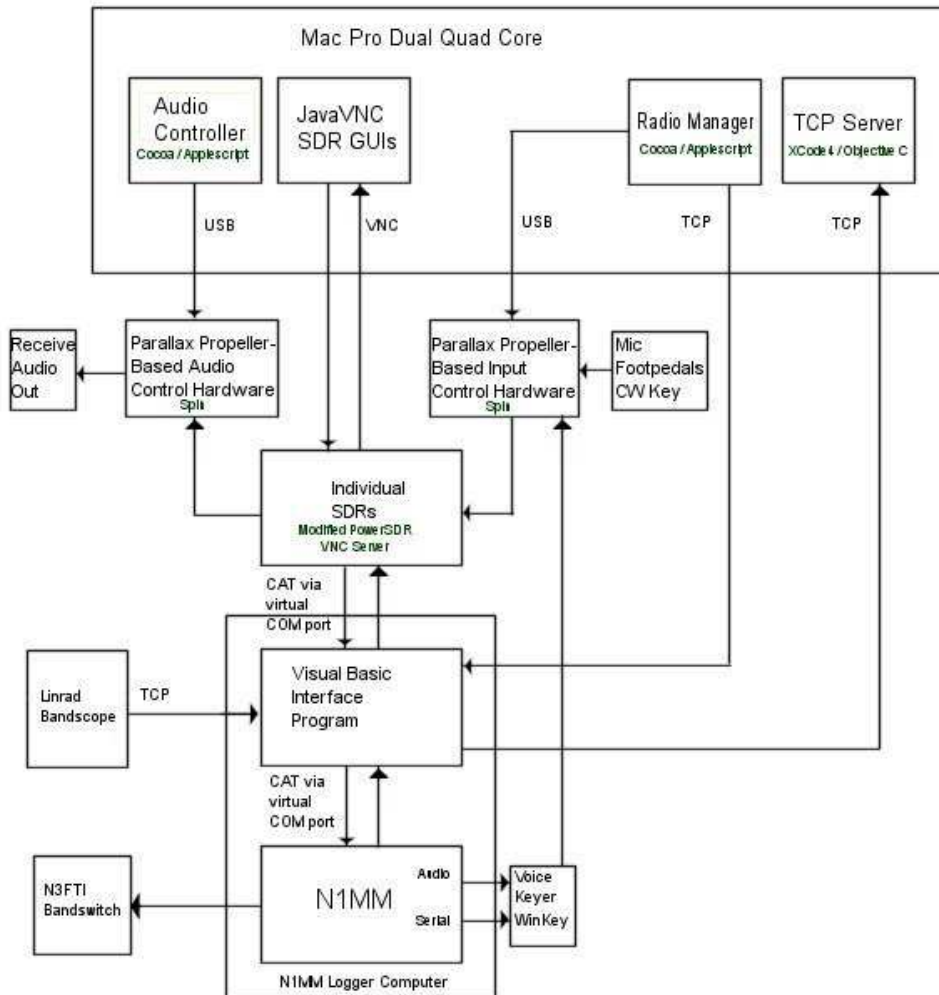
This is what the station looks like to the operator. The picture shows 3 computer screens:



On the left, radios for 50 and 144 MHz are on top, and 222 is below. Below the 144 MHz radio are the GUIs for configuring the automatic switching of the footswitches, Mic, Key, and receive audio. In the center is the N1MM logging computer. On the right

are the radios for 432 MHz and the Microwave bands on top, and 1296 MHz on the bottom.

This block diagram shows the system design (further discussion is on the next page):



A dual quad-core Mac Pro with its associated dual Apple LED Cinema Displays is the control center and aggregator for all of the radio GUIs. Six instances of PowerSDR are controlled from its screen, keyboard, and mouse via individual javaVNC connections to six individual computers, each connected to an SDR running PowerSDR and a VNC server.

The Mac Pro also controls, via 2 USB ports, a homebrew hardware switcher which automatically assigns microphone, footswitch, receive audio, key/keyer to the appropriate radio for the band desired. When a particular PowerSDR GUI is clicked upon, the Mac Pro sets the focus to that GUI (with its associated radio and instance of PowerSDR), and automatically switches the microphone, footswitch, receive audio, and key/keyer to that radio. This is done with two programs that I wrote that run on the Mac Pro, "Radio Manager" and "Audio Controller". Each is written in Cocoa/Applescript, and I have placed the code that I wrote for them on my website.

The Mac Pro also sends, via Radio Manager, band and exact frequency information to the logging computer and N1MM, so that the contact can be logged appropriately. The Mac Pro also receives, via “TCP Server”, band information from N1MM and uses it to the control switching of the microphone, footswitch, receive audio, and key/keyer when band/frequency selection are made from N1MM, and to give focus to the appropriate radio/GUI when frequency selection is made from N1MM. TCP Server is a program that I wrote in Objective C in the XCode environment, that runs on the Mac Pro.

Thus one can at any time select a radio/band and frequency from either the GUIs on the Mac Pro or from N1MM. Any change or selection made at one location will be accepted by the other. For the operator, things function as if there is one 11 band radio with multiple bandscopes coupled to N1MM by CAT control, with logging and all switching functions automatically controlled.

The Radio Manager and Audio Controller Cocoa / AppleScript applications on the Mac Pro allow one to customize the automatic switching of input/output devices, as the needs are different when one is on the low bands vs. when one is on the microwaves (and keeping both the primary and liaison channels open simultaneously). Here are the GUIs for these applications:



When on the microwaves, one footswitch and the microphone can be assigned to the liaison band, and the other footswitch and the key assigned to the microwave band, or vice versa, and a separate receive audio channel dedicated to each band.

The logging computer uses the Windows XP operating system, and it runs N1MM and a Visual Basic application that I wrote which acts as an intermediary between N1MM, the Mac Pro, and the 6 instances of PowerSDR running on individual computers. The logging computer also controls the antenna rotor via N1MM logger and a serial port, and WinKeyer, via N1MM logger and a USB-Serial Port adapter, and runs Win-Grid to serve as a check on the headings supplied by N1MM. Six instances of com0com, a virtual TCP serial port program, and one instance of hub4com, a virtual com port aggregator, provide virtual serial port connections across the network from the N1MM computer to the instances of PowerSDR running on each of the SDR computers.

Six separate computers, each running Windows XP, an instance of PowerSDR, and an instance of VNC Server, are controlled from Mac Pro via javaVNC viewers running on the Mac Pro. Each of these six computers is connected to its own dedicated SDR.


For six meters, a Flex5000 is used as the SDR, and an Elecraft 28/50 MHz transverter is used to get the signal to six meters. For two meters, a K3 is used as the radio, with an HPSDR running as its bandscope and CAT controller, and a DownEast Microwave 28/144 MHz transverter to get the signal up to two meters.

For the remaining bands, HPSDR radios are used, with one HPSDR per band except for 903 MHz/2/3/5/10/24 GHz which share one HPSDR radio. The bands 222, 432, and 1296 MHz each have a dedicated DownEast Microwave transverter attached to the HPSDR radio. The transverters for 903 MHz and 2.3 GHz and higher all share a single HPSDR radio, with switching among the transverters being done automatically as described elsewhere in this article. Except for 10 and 24 GHz, all of these transverters are DownEast Microwave as well. Kuhne transverters are used on 10 and 24 GHz.


In addition to running a VNC Server and an instance of PowerSDR, each HPSDR computer also runs an instance of the com0com virtual serial port program to communicate with the N1MM computer CAT control. Additional Visual Basic programs run on 144 MHz and 1296 MHz to provide specialized functions on those bands that are beyond the scope of this article.

The homebrew SDR Controller is connected via USB to the Mac Pro. This controller uses 2 Parallax Propeller Proto Board USBs to control the automatic switching of the microphone, footswitches, key/keyer, and receive audio to the appropriate radio for the band desired. The two Propeller Boards use their own custom Spin applications to control switching of these devices. The code for these programs which I wrote is also on my website.

Below is a picture of the homebrew SDR controller taken before the top was installed:



Below is a picture of the back panel, which holds all of the external connections:



A homebrew lockout device to prevent simultaneous transmitting and bandswitching also uses a Parallax Propeller Proto Board USB.

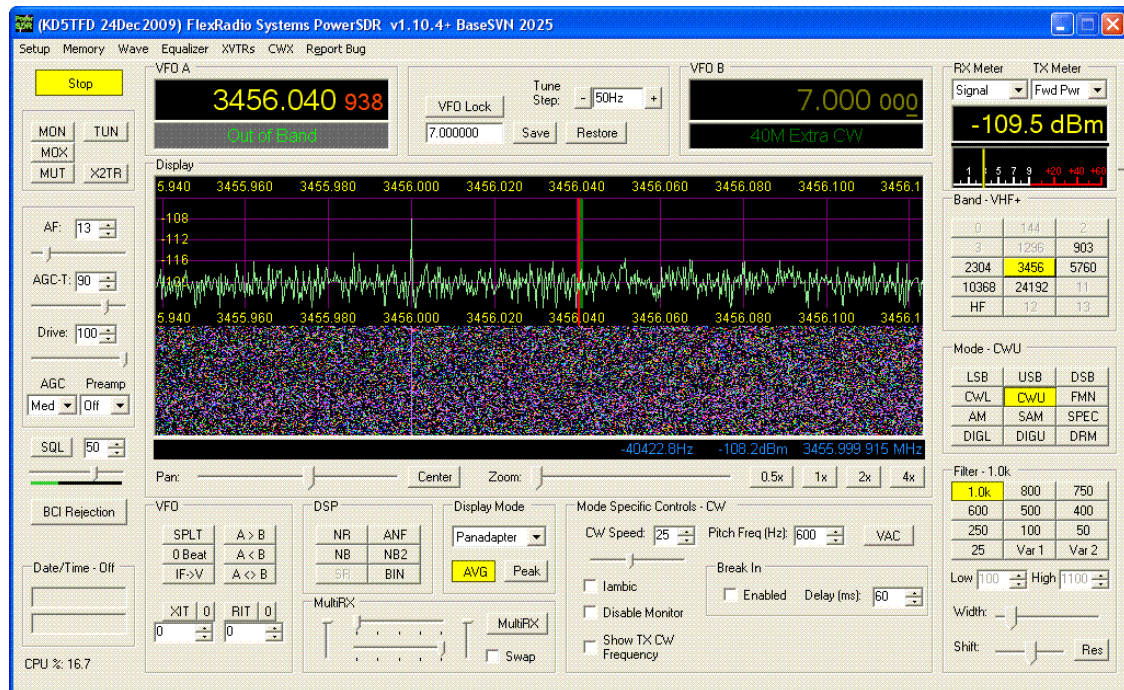
An N3FTI Bandswitch is used for bandswitching transverters and associated hardware for 903 MHz, and for 2.3 GHz and higher bands. This is automatically controlled by the HPSDR / PowerSDR combination used for those bands. Manual bandswitching is NEVER needed.

The choice of bandscopes that I made was deliberate. The first requirement, because “the Bandscope is the Radio”, was that the bandscope radio chosen needed to be a transceiver; a receive-only device was unacceptable. Using receive-only hardware was a “been-there, done-that long ago and not good enough” choice. I wanted to have at least [approximately] 192 kHz bandscope width. That limited my choices at that time to the Flex 5000 and the HPSDR radios. The cost of 6 Flex 5000s was daunting and the less expensive HPSDR radios were a more practical choice. I already had a Flex5000, and so I chose to use that on six meters. When I started out building this station, I wanted a bullet-proof receiver for 144 MHz with wide dynamic range, superb sensitivity, and excellent large signal handling capability. Thus I decided to use the K3 on two meters, but I coupled it to an HPSDR radio used as its bandscope. This combination has worked very well for me. For 222 MHz and up, HPSDR radios were used exclusively.

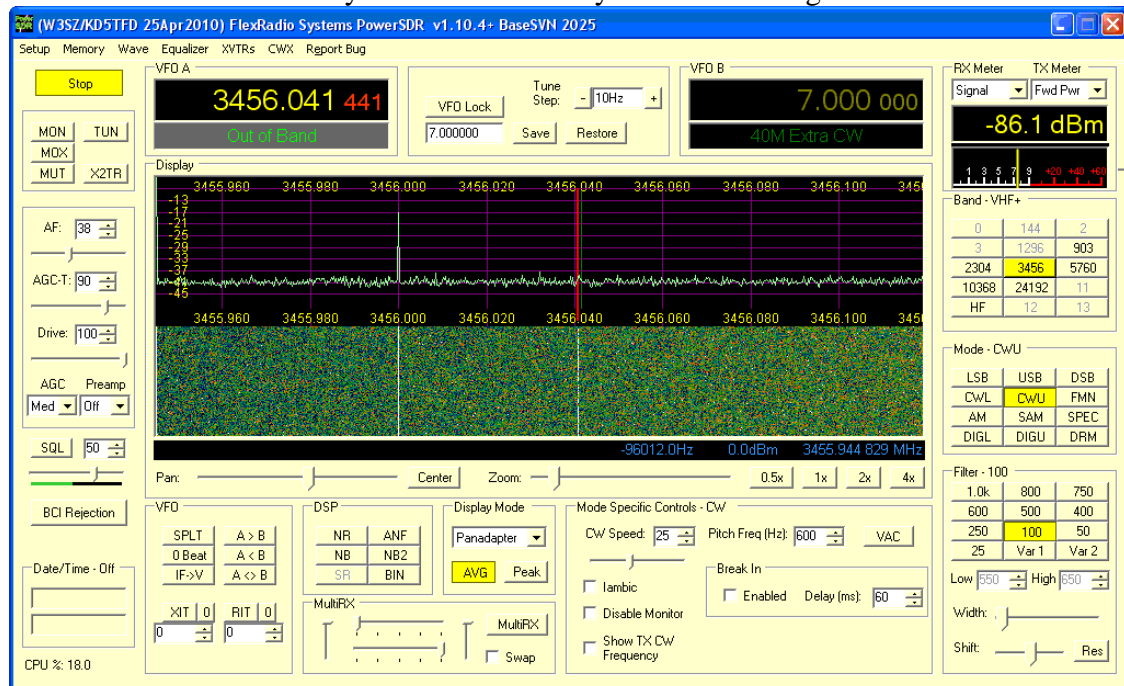
Once I started using the HPSDR radios, it turned out that the FFT size of the stock PowerSDR software [used on both the Flex 5000 and the HPSDR radios] was too small to give optimal bandscope sensitivity. I ran Linrad in parallel with PowerSDR and discovered that I could easily see signals on the Linrad bandscope that were invisible on the PowerSDR bandscope. But once I tuned the PowerSDR radio to the Linrad bandscope frequency, the signals were Q5! So I modified the PowerSDR radio to permit larger FFT sizes, and found that PowerSDR was then able to see the weak signals too. If you don't look for very weak signals running Linrad in parallel with PowerSDR you will never know what you are missing! My homebrew-modified version of PowerSDR allows me to select any FFT size [in powers of 2] between 4096 and 262144.

The series of illustrations below are some screen-grabs that illustrate the need to use the large FFT sizes when doing weak signal work. The large FFT size used is 262144. The same recorded signal was used for all of these screen-grabs. Ignore the internal birdie at 3456.000 MHz.

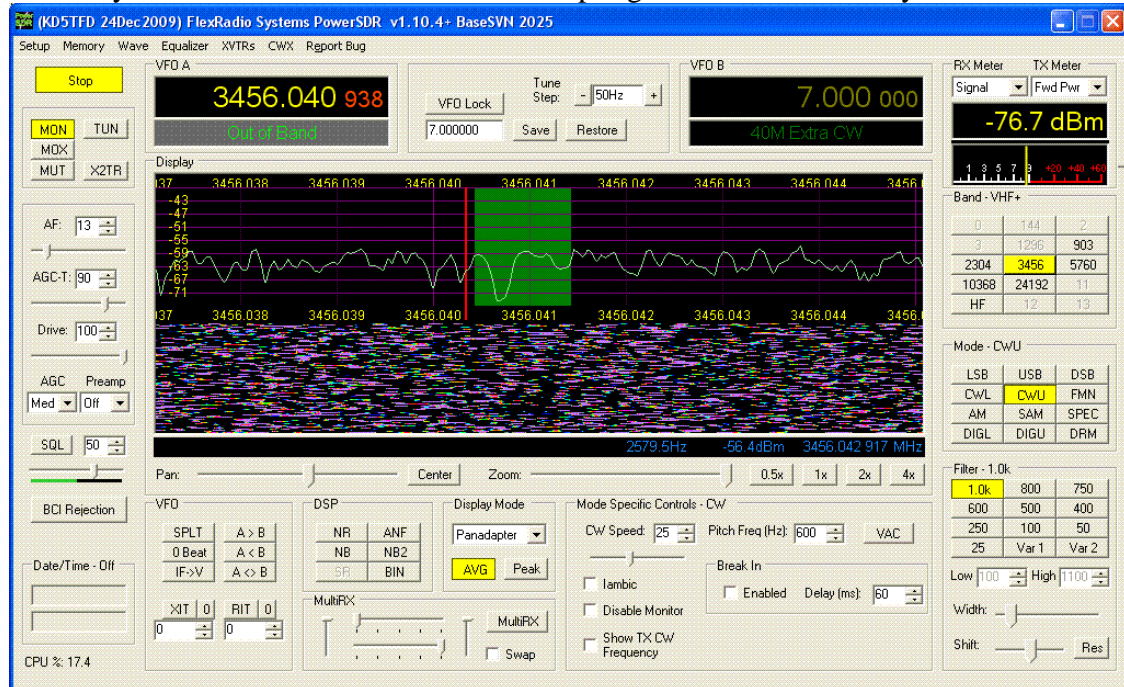
Below is a full width Standard FFT size PowerSDR bandscope. The weak signal at 3456.041 MHz is invisible:



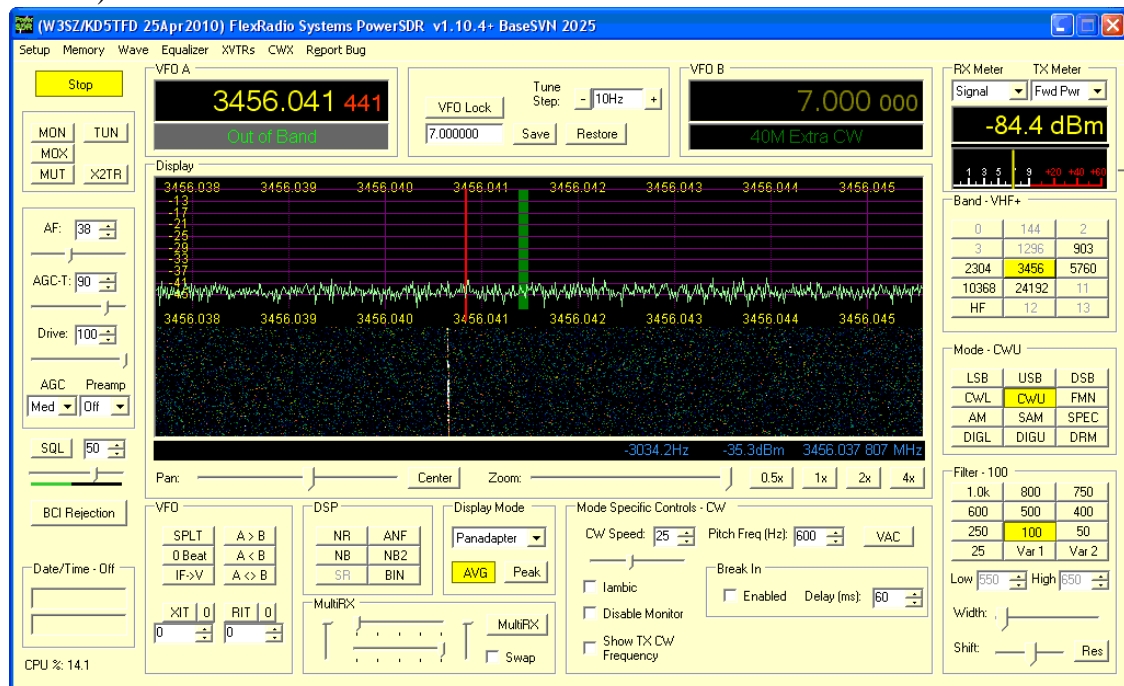
Below is a full width Large FFT size PowerSDR bandscope, same weak signal at 3456.041 MHz. Note that you can now actually see the weak signal on the waterfall:



You can really see what the problem is with the standard FFT size PowerSDR when the bandscope is zoomed in below. The signal is again invisible and the large bin size that is forced by the small number of bins when sampling at 192 kHz is clearly demonstrated:



Below is the same zoomed-in spectrum with the Large FFT size version of PowerSDR. Note the smaller bin size, the superior frequency resolution, and the fact that you can actually see the signal (which is on the waterfall below and just to the left of the vertical red bar)!:



To summarize the main points of this article:

- One bandscope is good.
- Multiple Bandscopes are better.
- The best result is achieved when “The Radio is the Bandscope.”
- Automatic bandswitching and automatic switching of Mic, CW key, footswitches, receive audio, and the logging band is essential.
- Full integration of the radios and the logging software is required.
- A large FFT size is required to maintain bandscope sensitivity.

Additional information on these items, more photos, as well as software sources and schematics for the hardware, are on my webpage at:

<http://www.nitehawk.com/w3sz/osxhpsdrserver.htm>

Other items of interest related to this project that are available on the website include:

Software Sources for the applications mentioned above:

Custom Spin application for the RadioManager Propeller Board:

<http://www.nitehawk.com/w3sz/RadioManager.spin>

Custom Spin application for the AudioManager Propeller Board:

<http://www.nitehawk.com/w3sz/AudioController.spin>

The RadioManager application that runs on the Mac Pro:

http://www.nitehawk.com/w3sz/New_Radio_ManagerAppDelegate.applescript

The AudioController application that runs on the Mac Pro:

http://www.nitehawk.com/w3sz/Audio_ControllerAppDelegate.applescript

The TCP server program that runs on the Mac Pro and receives information from N1MM:

<http://www.nitehawk.com/w3sz/W3SZOSXTCPServer.txt>

The TCP client program that runs on the N1MM computer and communicates with the Mac Pro:

<http://www.nitehawk.com/w3sz/tcpclient.c>

The Visual Basic program that runs on the N1MM computer and communicates with N1MM Logger and the Mac Pro:

<http://www.nitehawk.com/w3sz/N1MMNewHPSDRRigControl.txt>

The web pages describing the HPSDR radios can be found at:

<http://openhpsdr.org/>

Some of the HPSDR components are available for sale by TAPR:

http://www.tapr.org/hpsdr_index.html

Apache Labs, in India, sells Hermes, the new HPSDR transceiver, and some other items:

<http://www.apache-labs.com/>

If you have questions please feel free to contact me via email: w3sz@arrl.net

73,

Roger Rehr W3SZ
July 18, 2012