

10m SSB receiver

Written by Hans Summers

Wednesday, 10 September 2014 03:26 - Last Updated Tuesday, 17 November 2015 07:46

[A 10m band Weaver receiver](#)

[1. Receiver measurements](#)

[2. Experiments and modifications to the Weaver receiver: moving the subcarrier to 2,048Hz](#)

[3. Experiment to see the effect of increasing the coupling capacitor](#)

[4. Experiment to adjust the matching of the audio LPF](#)

[5. Experiment for amplitude-matching the I and Q paths](#)

[6. DC-coupling of the LPF output to the audio mixer](#)

[7. Experimentally changing the summer after the audio mixer](#)

[Conclusions](#)

A 10m band Weaver receiver

Practically the only HF receiver I can find online is a an 80m transceiver article, [December 1984 CQ-DL article by Burkhard Kainka DK7JD](#)

, "80-m-SSB-Transceiver nach der dritten Methode" (

[click to read](#)

). This 3-page article is all in German, which regrettably I cannot speak at all, however the circuit diagram (schematic) is a universal language, and it is possible to gain enough understanding from the circuit to reproduce this project, without understanding the German language.

So, my 10m receiver design is somewhat based on the [DK7JD project](#) from 30 years ago. But I'm using a Quadrature Sampling Detector (QSD) front end, just like the SDR's use. There are various other differences too, but the main inspiration is based on DK7JD's article.

The circuit diagram and photographs of my receiver are below. Click to enlarge any of them. Note the presence of the red numbers on the circuit diagram. For later convenience, in the experiments and modifications sections below, the numbers refer to the experiment section and show where on the circuit diagram the described modifications took place.

{gallery}weaver/photos{/gallery}

The photograph in the middle shows my rough construction of the receiver. The photograph on the right has labelling of each of the circuit blocks so you can see where everything is. For the

10m SSB receiver

Written by Hans Summers

Wednesday, 10 September 2014 03:26 - Last Updated Tuesday, 17 November 2015 07:46

Variable Frequency Oscillator VFO), I am using an [Ultimate3 QRSS/WSPR transmitter kit](#) set to 112.5MHz (four times the reception frequency, for generating a quadrature LO). The

[U3](#)

's official version (as at September 2014) doesn't yet go up above 30MHz but this particular one doesn't have an AD9850 DDS module, it has my new

[OCXO/Synthesiser module \(click here\)](#)

in it instead. This is the purple-coloured module on the left of the photo. I am using another

[Ultimate3 QRSS/WSPR transmitter kit](#)

, this time with the conventional AD9850 DDS module, as my 10m signal generator for testing the receiver performance. This is the one with the LCD screen facing us at the top of the photo. Usually the VFO is the hardest part of a receiver design and even more so, when you want your VFO at 4x the reception frequency. But now with this

[new OCXO/Synthesiser module](#)

such a thing is almost trivially easy.

Some words about the rest of the receiver. The VFO signal is set to 112.5MHz for test purposes and it is fed into a dual D-type 74AC74 flip flop which generates a quadrature square-wave output. A 74AC74 has a much higher frequency capability than a 74HC74 and that is needed here for operation at 112.5MHz. There are four squarewave 28.125MHz outputs, they are at 0, 90, 180 and 270-degrees phase shift with respect to each other. These four outputs operate the individual switches in a 74HC4066. Better devices than the 74HC4066 can be used, but this was conveniently through-hole (not surface mount as some of the better switches) and was available in the junk box.

The signal generator [U3](#) is feeding a classic double tuned resonator band pass input filter, via a 3.9K resistor to reduce the amplitude a bit (prevent overloading the receiver). This BPF tuned very nicely (using the two trimmer capacitors) on the spectrum analyser and it was fantastic to actually see the effect of tuning - the insertion loss, double hump response, and finally the nice reasonably-narrow band-pass filter when all was aligned. In practice it can be aligned too just by tuning to a signal on 10m and adjusting the trimmers for highest amplitude audio output.

The input to the QSD must also be biased to 2.5V mid-rail reference. In this circuit I generate the 2.5V reference from a potential divider boosted by a unity gain op-amp configuration, using a nearby junk-box ICL7612 op-amp. I use the same 2.5V as a later ground reference for all the op-amps in the audio chain, which in my circuit are powered from the same 5V supply as the 74HC logic chips. Ideally, you would want to power the op-amps with a higher supply voltage e.g. 12V single-rail, and arrange for a separate mid-rail for the op-amps at 6V, say. This results in wider dynamic range because there is more head room before the op-amps are overloaded and clip at their output.

10m SSB receiver

Written by Hans Summers

Wednesday, 10 September 2014 03:26 - Last Updated Tuesday, 17 November 2015 07:46

In viewing the circuit be careful to note there are points labeled Gnd (0V), +2.5V (this 2.5V reference) and +5V (the main supply voltage).

After the QSD circuit I use 10uF coupling capacitors to a pair of op-amps (actually one low noise NE5332 dual op-amp chip), that act as pre-amplifiers and produce the conventional I and Q audio outputs. These could be fed to an SDR input e.g. using a soundcard on a PC. However here we're interested in the Weaver demodulator so we don't just stop here.

From here onwards the circuit diagram is somewhat close to the [DK7JD circuit](#), but without the TX/RX switching functions. Mine is wired only for receive. I use a pair of low pass filter the same as his, each path (I and Q) uses a pair of op-amps (one NE5532 chip). In the [DK7JD circuit](#), he precedes the audio sub-carrier mixer with one more pair of amplification stages (two op-amps). It is possible he did this because his circuit is a transceiver and he needed more gain on transmit. I built that same amplification stage and it didn't work, it oscillated somewhere in the region of 600kHz. I couldn't work out why so I just left out that stage altogether. It's still there in my photo and labeled "Not used".

Next comes the audio sub-carrier mixer which I am also calling the Weaver demodulator. Here the magic comes about, the unwanted signals are cancelled out and the wanted ones are re-inforced, and what comes out the far end is the desired audio channel. A miracle to behold, indeed. Another 74HC4066 quad-switch is used for the audio mixer. Again it's driven by a quadrature generator, exactly the same 74AC74 circuit as we use at the RF front end. However this time it is clocked at audio, rather than RF from the 112.5MHz VFO. To generate a suitable audio sub-carrier my circuit differs from the [DK7JD circuit](#). I used a 32.768kHz watch crystal oscillator, which is fed into a 74HC4040 counter/divider chip. In my diagram the two diodes at the Q0 and Q2 output are wired as a

[discrete component AND-gate \(see my page on this technique\)](#)

. This actually turns the 74HC4040 into a divide-by-5 circuit. The resulting frequency fed to the 74AC74 quadrature generator is at 6.554kHz and results in a quadrature subcarrier output at 1.638kHz. That's close enough to 1.5kHz for now. One could also divide by 6 to get 1.365kHz, nearby as well.

Finally some op-amps (another NE5532) sum the signals and right at the end, just like in the [DK7JD circuit](#), there's one final amplifier that also appears to have some low-pass function. This can be

10m SSB receiver

Written by Hans Summers


Wednesday, 10 September 2014 03:26 - Last Updated Tuesday, 17 November 2015 07:46

observed in the measurements too. I feed the audio output into the sound card of a PC where I use the Argo DSP program to measure its amplitude. This program can measure the amplitude and frequency accurately and it is easy to see where all the spectral outputs are, both the wanted signal and unwanted, in the range 0 - 3kHz.


1. Receiver measurements

First a look at the measured response of the audio Low Pass Filter response. The actual measured LO frequency was 28.12232MHz and so the Signal Generator is tuned up from here, in small steps, to 3kHz above.

The chart below shows the measured output of the I and Q channels at the output of the pre-amplifiers. You can see they look really nice. A nice flat output from 100Hz to over 2.5kHz. Furthermore the I and Q channels look pretty equal, which is important for balance and the accurate cancellation of the Weaver reflected "sidebands".



Now at the same time I measured the output of the I and Q channels AFTER the low pass filter. I actually had a junkbox 4-position switch, which is still visible at the bottom right of the photo, and which was used to select between I and Q signals, and before and after the LPF. The chart below shows the audio output level at the output of the low pass filter:




Now if I subtract one from the other, I will have the low pass filter response. It has some gain

10m SSB receiver


Written by Hans Summers

Wednesday, 10 September 2014 03:26 - Last Updated Tuesday, 17 November 2015 07:46

built in. It looks reasonably as I would expect it to be. So it is possible to conclude that up to this point, things are working as they should, which is always comforting. It is a reminder that building things stage by stage, and testing thoroughly at intervals, is a worthwhile development approach. It builds confidence and makes it easier (less difficult?) to identify where any faults are.



You can plot this with a logarithmic x-axis too if you want, and then it looks like below. This is arguably the right way to draw a low pass filter response:



Now a look at the audio output at the end of the receiver circuit chain. This is AFTER the removal of the strangely ~600kHz oscillating gain stage, which appears unnecessary and I couldn't understand why it was oscillating and ruining everything. Here I tuned the signal generator across the LO frequency and measured the audio output using Argo. The y-axis is showing the actual recorded audio level on Argo and this is relative, dependent on the gain of the receiver and some attenuation placed between the receiver output and the sound card. The absolute level is not important, it's the relative level that is important.

Note that a peculiarity of Weaver is that the LO frequency is at the audio subcarrier frequency. So to put it another way, if we were building a receiver using a superhet or conventional phasing method direct conversion architecture, for upper sideband (USB) the 0Hz of the baseband audio would correspond to the frequency of the local oscillator (LO). But in a Weaver receiver, this is NOT so. The LO frequency is set in the MIDDLE of the desired audio channel. This is, in this case, 1.638kHz above the LO frequency of a usual superhet or phasing receiver. The frequency

10m SSB receiver


Written by Hans Summers

Wednesday, 10 September 2014 03:26 - Last Updated Tuesday, 17 November 2015 07:46

of the LO is set to 1.638kHz above the usually recognised "reception frequency".

Notice the deep null at the audio subcarrier frequency, 1.638kHz. Notice another null at 0Hz due to capacitive coupling of stages, including the quadrature sampling detector stage. Finally, notice the response at below 0Hz, this corresponds to the "unwanted sideband" in a conventional receiver (superhet or direct conversion). This is STILL present in the Weaver receiver but now it is determined by the steepness of the audio LPF rolloff, NOT the accuracy of the phase shift and amplitude between the I and Q paths.

Finally, note that the audio output curve is assymmetric on either side of the subcarrier at 1.638kHz. I believe there are two reasons for this. Firstly the output of the QSD after the pre-amps shows a drop in amplitude after 2.5kHz. This may be due to the presence of the 220pF capacitors in the feedback loop, though it seems not likely according to my intuition, though I have not attempted to model this circuit. The capacitors in the QSD themselves will show a gradual audio rolloff but not at anywhere near as low a frequency as 2.5kHz and not suddenly like this. The second reason could be the final audio output stage which appears to have a LPF characteristic, but again this needs to be modelled.



Another point to mention is that this receiver is configured for USB. You can easily change it to LSB just by reversing paths, in several different places. I did actually try it in the audio mixer quadrature generator. Swapping the D-inputs of the two flip flops from the Q and /Q outputs to the /Q and Q outputs, had the effect of reversing the received sideband. In this case, the usual "receiver frequency" is the 0Hz at the TOP of the audio spectrum, since the audio spectrum is reversed. Now the actual Weaver receiver LO frequency must be set to 1.638kHz BELOW the "received frequency". As the signal generator is tuned down in frequency, the audio output frequency tunes upwards, as expected. I tried this swap, and it works well, as it should.

In subsequent charts, I'm going to draw the "unwanted classical sideband" response, where the LO is tuned below the wanted reception frequency, as this reflected line back into the main audio response from 0 to 3kHz. This is to consider the fact that if you are listening to a wanted


10m SSB receiver

Written by Hans Summers

Wednesday, 10 September 2014 03:26 - Last Updated Tuesday, 17 November 2015 07:46


signal say at 500Hz, you will also hear at that same frequency an unwanted signal from the "unwanted sideband" (in the classical receiver sense), at -30dB (or whatever) down on the wanted signal. This chart shows this "reflected unwanted sideband response", and this is how I am going to draw subsequent charts in this page. It doesn't matter how it is shown, as long as it is understood, what it means.

The other thing I'm going to do from this section onwards, is to normalise the audio output level such that the 0dB reference point is measured at 1,500Hz audio. This will simply make it easier to compare the response curves, as I go through making various experimental changes to the receiver. 1,500Hz is arbitrarily chosen because this is the typical centre frequency of a WSPR transmission (though need not be). The absolute level is arbitrary, it is just a matter of convenience for making relative comparisons.



Finally for the initial measurements of the 10m Weaver receiver, I show here a RED line which is the "reflected unwanted Weaver sideband" level. This is the result of inaccuracies in the circuit cancellation. Put another way: if I am receiving a QRSS signal at 1,000Hz audio, and the audio subcarrier is at 1.638kHz, then I will also receive weekly a signal at 2.276kHz audio, reflected across the 1.638kHz subcarrier, onto my wanted 1,000Hz signal. This is the level of the unwanted, within-channel folded-back sideband that is characteristic of the Weaver method.

So in general, the reflected "classical unwanted sideband" line drawn in blue going downwards at the left of the chart, and the Red "reflected unwanted Weaver sideband" are showing the levels of spurious response of this Weaver method receiver.



Ok so this 10m receiver really works, Weaver really works, and it is FASCINATING MAGIC, I can tell you. It was surprisingly easy to build too, and no trouble at 10m HF frequency (28MHz). Now for some experiments on the receiver, to determine if it can be improved and to learn a bit more about it.

2. Experiments and modifications to the Weaver receiver: moving the subcarrier to 2,048Hz

Now the response of the receiver looks as good as might be hoped (see previous section), however the original intention of this design was to see if Weaver could be a valid receiver architecture to build a receiver for QRSS and WSPR. QRSS typically may set the audio to 1,000Hz and a typical receiver centre frequency for WSPR is conventionally 1,500Hz. This means the 200Hz WSPR sub-band centre, so reception really needs to be 1,400-1,600Hz. If you look at the original measurement charts (section above) you can see that at around 1,500Hz audio the signal level is already dropping by 10dB because of the proximity to the 1.638kHz subcarrier frequency. The "subcarrier hole" may not matter in a human speech pattern, but it certainly matters if it falls on top of, or close to, our desired narrow QRSS/WSPR sub-bands!

So we can easily move the subcarrier away from the WSPR audio sub-band. To do so, instead of dividing by 5 in the 74HC4040 divider, just divide by 4 instead. The 32.768kHz signal divided by four is 8,192Hz and the resulting subcarrier is at a quarter of that, i.e. 2.048kHz. Dividing by 4 is a very easy matter, we just take out that AND-gate made from two diodes and the 100K resistor. Instead we connect the reset pin (pin 11) to ground, so that no reset ever occurs. We don't need it now. Then take the 8,192Hz signal from the Q1 output of the 74HC4040, which is pin 6. Very easy modification. As I mentioned before, the modification locations are indicated on the circuit diagram by the large red numbers. Here we are in experiment section 2, and you can see the red "2" near the 74HC4040 reset circuit.


The resulting response is shown below. As I mentioned before, I'm now normalising the audio level such that 0dB is the measured level at 1,500Hz audio. Now notice that we have basically shifted the response curve 410Hz to the right. The response is now quite flat from 1,000Hz to 1,500Hz audio and a little bit either side of it. This would make it suitable for receiving QRSS or WSPR in sub-bands at these audio frequency positions. An additional benefit to moving everything to the right, is that the "classical unwanted sideband" response is also considerably reduced. This is because we are further from the audio subcarrier, but we haven't changed the audio low pass filter, and so being further away means we are further down the curve of the audio LPF. Nice. At 1,500Hz audio for example, the "classical unwanted sideband" is going to

10m SSB receiver

Written by Hans Summers

Wednesday, 10 September 2014 03:26 - Last Updated Tuesday, 17 November 2015 07:46


be at -55dB (from approximately extrapolating the line below), quite a respectable attenuation. Again the red lines show the "folded back Weaver sideband" from within the SSB channel, the other side of the audio subcarrier. The level here is about -28dB relative to the wanted signal.




3. Experiment to see the effect of increasing the coupling capacitor

Now look at the circuit diagram again and notice the 100nF (0.1uF) capacitor between the output of the LPF block and the input of the audio mixer (74HC4066 inputs). 0.1uF is what was used in the [DK7JD circuit](#). But it intuitively feels rather low for audio, don't you think? I speculated that increasing the value of this coupling capacitor might reduce the size of the "hole" at the subcarrier frequency. This would be the expected result, from the circuit. It's near this subcarrier frequency that the audio frequency actually coming from the RF QSD circuit, and coming through the pre-amps and audio LPF, is approaching zero.

Well it was as predicted, and the result is shown in the chart below. The "Weaver hole" now becomes very narrow indeed.



Here's the same curve, drawn on top of the response curve from the experiment 2, described in the previous section. Note that the rest of the response curve isn't altered, it's only the width of the "Weaver hole" which is narrowed. This seems like a nice and worthwhile modification with a noticeable improvement to the original circuit.



4. Experiment to adjust the matching of the audio LPF


Now we should try to see if we can reduce the level of the unwanted within-channel "folded unwanted Weaver sideband" that is caused by mismatch between the I and Q paths.

My first attempt at this is to alter the gain of the Q path in the LPF. See the red "4" in the circuit diagram near the bottom-middle of the page? I replaced the fixed 220K resistor with a variable 100K multi-turn preset potentiometer in series with a 180K resistor, the resulting combination is therefore variable over the range 180K to 280K. Adjustment of this trimmer pot seemed not at all critical. It did seem possible to look at the folded sideband image level and reduce it by adjustment of the pot. It was difficult to know where the pot should be set. I put it at the best minimum I could judge and re-measured everything.

Another change was also made at the same time, which was to insert a 47nF coupling capacitor between the audio pre-amps and the low pass filter. One in each path, with a 100K resistor to the 2.5V reference at the LPF input side of the capacitor. My mistaken thinking, in doing this, was something to do with trying to get a faster roll-off near to 0Hz (the far left of the graph). In fact this is nonsense, since 0Hz audio at that point corresponds to near the "Weaver hole" (now at 2,048Hz). So all I did with this modification is widen the "Weaver hole" width again. Pointless, and that can be seen on the graph below, so let's not mention it again. I took it straight out again right after.

The graph below shows the result of adjusting the Q path of the LPF with the trimmer. It certainly seems that the unwanted in-band Weaver folded sideband amplitude has been reduced by 10dB approximately. A good result! Furthermore the response curve looks flatter along the top, within a couple of dB from 800Hz to over 2,000Hz where the hole is.


I'm not entirely clear on what the adjustment is here. Ideally there should be one adjustment for amplitude balance and one for phase balance. I think that this one probably adjusts both.



5. Experiment for amplitude-matching the I and Q paths

In this measurement run I tried adjusting the amplitude of the Q path by moving the 100K trimmer potentiometer to replace the 47K resistor in the I path which controls the pre-amp gain, see the "5" near the upper centre of the circuit diagram. I also removed the 100nF/100K coupling between the pre-amp output and the audio LPF input, because it was pointless; so you'll also see a pair of "5"'s next to that area. The important change is being able to adjust the relative gain of the pre-amps in each of the I and Q paths, which means any amplitude imbalance can be removed. The adjustment was made at 1,500Hz audio. The adjustment was VERY sensitive, even using the multi-turn potentiometer. It was not difficult to find the optimum adjustment, where moving the potentiometer in either direction increased the amplitude of the unwanted signal.

The result is shown in the chart below. The amplitude of the unwanted "folded image" is less than the un-adjusted case but not as good as the adjustment performed in the previous section 4. This adjustment now is only to the amplitude. That one also probably was making some difference to the phase imbalance. So really, both need to be considered, ideally.



6. DC-coupling of the LPF output to the audio mixer


I wondered if the Null, or "hole" at 2,048Hz could be narrowed even further by DC coupling the LPF output to the audio mixer (74HC4066) input. So those 10uF capacitors in the I and Q paths were shorted out (they were originally 0.1uF, remember). This was the only change for Experiment 6. I did actually try removing the 10uF coupling capacitors between the QSD and the pre-amps, but nothing worked at all like that - I imagine small offsets are driving the pre-amp op-amps into saturation. So that was taken out.

10m SSB receiver

Written by Hans Summers

Wednesday, 10 September 2014 03:26 - Last Updated Tuesday, 17 November 2015 07:46

The results are shown below. DC-coupling worked fine between the LPF output and the audio Weaver demodulator mixer input. I think you can say the hole was slightly narrowed but only slightly. Certainly it wasn't possible to measure such a deep null. It's definitely an improvement and it gets rid of some unnecessary components, so we shall keep it!



7. Experimentally changing the summer after the audio mixer

I couldn't really understand why the [DK7JD circuit](#) has an 11K resistor right before the 74HC4066 audio mixer. Not having found a suitable pair of 11K resistors in the junk box, I used 10K + 1K in series. I wondered if the 11K takes into account the relatively high ON-resistance of the original 4066 CMOS-series chip which was probably used by DK7JD back in 1984. So that the total would become 12K approximately, and given the switch was closed only half the time, that would deliver half the energy through the switch in each cycle, so become approximately equal to the 24K resistors on the other side of the switch. Just speculating, here. However that doesn't really make sense either to me, because that resistor would also be in series with the gain-controlling resistors at the negative input to the op-amps, which already have a 24K resistor.


So, I decided just to try modifying this circuit and see if it helps. I removed the 11K resistors. I inserted a 24K resistor in series with each positive op-amp input, on the right hand side of the switch. This looks to me now like a properly balanced configuration for the op-amp.

The results are shown below. Actually it does seem to reduce the amplitude of the unwanted in-band "folded Weaver sideband", the improvement depends on what audio frequency we are looking at but is more than 10dB at around 1,500Hz. Given that there is improvement across the band, it seems this modification is worthwhile too.

10m SSB receiver

Written by Hans Summers

Wednesday, 10 September 2014 03:26 - Last Updated Tuesday, 17 November 2015 07:46



Conclusions

There would be other things to experiment with, to carry on this project. Optimising the adjustment, by independently adjusting phase balance and amplitude balance, is one thing. There are many more aspects of receiver performance to test too, other than just the unwanted sideband rejection I've concentrated on here. Though I don't think there would be any big surprises. The switching (QSD) front-end is used in a lot of SDR's and is high performance, and well understood. The front end primarily determines the overall radio performance in many aspects.

I think that this Weaver receiver experiment was very educational for me and fascinating too. For the considered purpose of receiving QRSS and WSPR signals I do NOT think it is something I will pursue. The reflected-within-the-audio-passband nature of the unwanted sideband signal is of little consequence for SSB, but for QRSS and WSPR, it makes no difference whether the interference comes from within the 3kHz audio band, or somewhere in the classical "unwanted opposite sideband". It's there, interfering with us, all the same. The unwanted signal rejection isn't better than what could be done with simpler circuits, as far as QRSS/WSPR are concerned. So the additional circuit complexity isn't warranted.

On the other hand, I was encouraged to see Weaver actually working, having read about it on and off over some years, and always wanted to try it. I think it may be a technique with potential, and I do intend to come back to it when my objective is more SSB-related rather than the narrowband weak-signal modes. For SSB I think Weaver might be usable to build a good receiver or even transceiver.