A direct conversion HF SSB receiver (Rev 1)

By Juan H la Grange, ZS6SZ

1. Introduction

This is the description of a direct conversion, phasing method SSB receiver, covering the range 1.8 to 30 MHz, specifically for radio amateur use. It is published for use by private enthusiasts and hobbyists.

It must be stated up front that it uses and builds on the ideas of many other people. They will be recognised where applicable and the internet will provide even more references.

The author started the work more than 40 years ago but then left it until he found time again now in his retirement. Components at hand were used where possible, thus no guarantee is given that those will be available for purchase or even that they may be optimal.

2. Operating principles

As mentioned, the phasing or second method of SSB is used, which employs 90° phase shift in the local oscillator (LO) and AF stages to filter out the unwanted sideband. The LO produces two signals which have a 90° phase difference and feeds these into two balanced mixers, following an RF stage. After the mixers there are thus two audio channels, with the one being phase shifted by 90° again or both being shifted, one by +45° and one by -45°. These are then combined for the audio output.

The balanced mixers filter out the local oscillator carrier, although this has little relevance for the receiver.

Direct conversion means there is no IF, conversion is directly from RF to AF in the balanced mixers, with most of the filtering and amplification done at AF. The sharp bandpass filtering, active or passive, that can be done at AF eliminates the need for crystal filters.

The LO runs at the frequency being tuned, which means that measuring and displaying this frequency is relatively easy – there is no IF difference. Using digital circuits in the LO means that a 90° phase shift is quite easy to achieve.

3. Block diagram

The simplified block diagram is as follows:

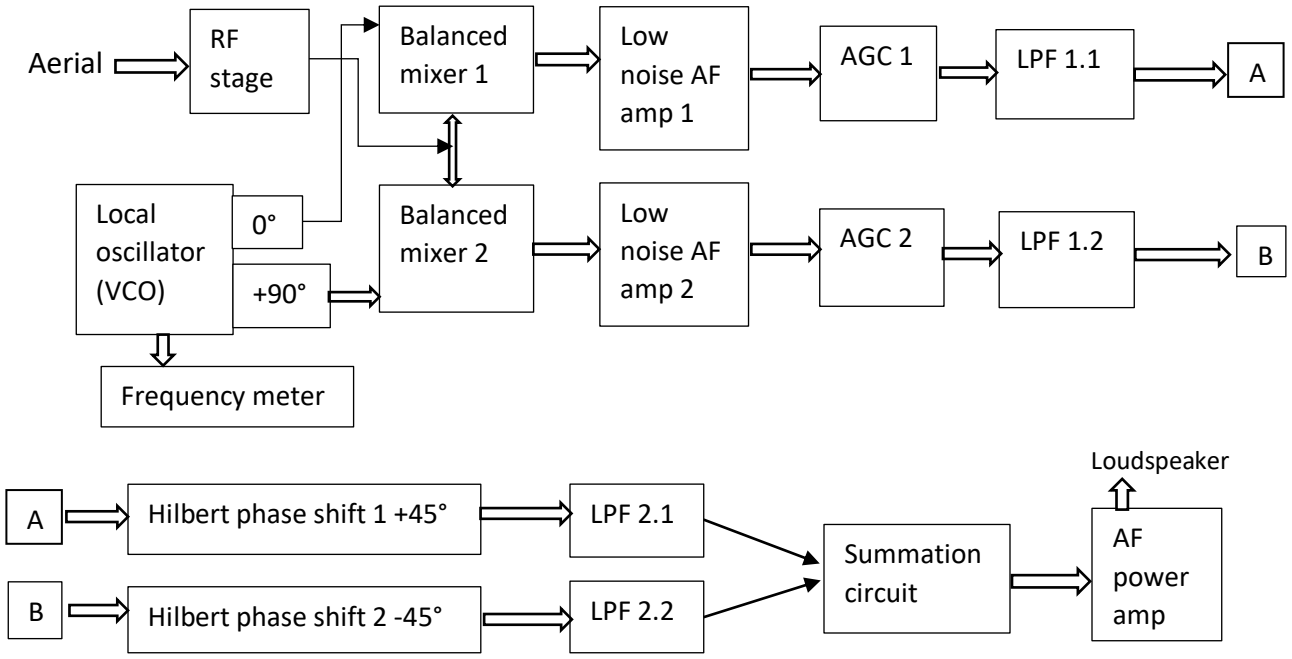

A design alternative, with far fewer components, would have been to have the phase shift and summation stages much earlier and then only have one channel of filtering, AGC and amplification. However, the choice of a digital Hilbert phase shift on a PIC 16F1619 microchip with "limited" capabilities and thus a relatively low sampling rate, forced the above design. (The selected PIC device can clock at 32 MHz and the ADC has a minimum 7 mV step, which is quite amazing considering its low cost. But much more enhanced chips are available, of course at a price.)

The reason for the two LPF2s, and not just one after the summation circuit, is that an intended conversion of the receiver to a transceiver will require two of these.

Tuning is in two steps (and with two knobs) in that the desired frequency must first be tuned to via the Voltage Controlled Oscillator (VCO) and the frequency meter. Then the RF stage is tuned to that frequency.

Sideband selection, not shown above, takes place between LFP1 and the Hilbert phase shift.

To contain EMI, all logic circuits are bypassed by means of a ferrite bead in series with and next to the +5V pin and a 10 nF capacitor from that pin to the ground pin [c.f. Application Manual for Power Supply Noise Suppression and Decoupling for Digital ICs, Murata C39E.pdf, July 20, 2010]. Ferrite beads are put on all circuit board power supply leads. The circuit diagrams below do not include the power supply connections for logic circuits, to reduce clutter.

Construction was done on copper strip boards and no layouts are provided here as correctly updated versions are not available. (All the designs changed as the receiver was being put together and problems showed up, which changed the board layouts. All the initial paper design, breadboarding and testing was obviously not enough, and much learning took place in the process. Including the mistakes made along the way, most of which will rather not be mentioned here, these school fees turned out to be quite expensive and the process of learning humbling but fascinating.)

Circuit blocks, according to the block diagram above, are screened from each other with copper-clad PC boards in the receiver as per the photographs. (Up to the AGC stage the two channels are split physically to reduce cross-coupling. After the AGC stage the two channels are constructed on one board per set of blocks.) Twisted pairs are used to convey the signals from one block to another at RF. The boards are mounted on the aluminium chassis or on the PCB screens by means of stand-off plastic posts.

4. RF stage (Rev 1)

A double-tuned RF stage is employed for the following reasons:

- The intent is to convert the receiver into a transceiver, which will require the tuned circuits to filter out harmonics prior to transmission.
- The local oscillator output is a square wave containing many harmonics, which may lead to unintended signals being detected if there is no selectivity beforehand.

The circuits are as follows:

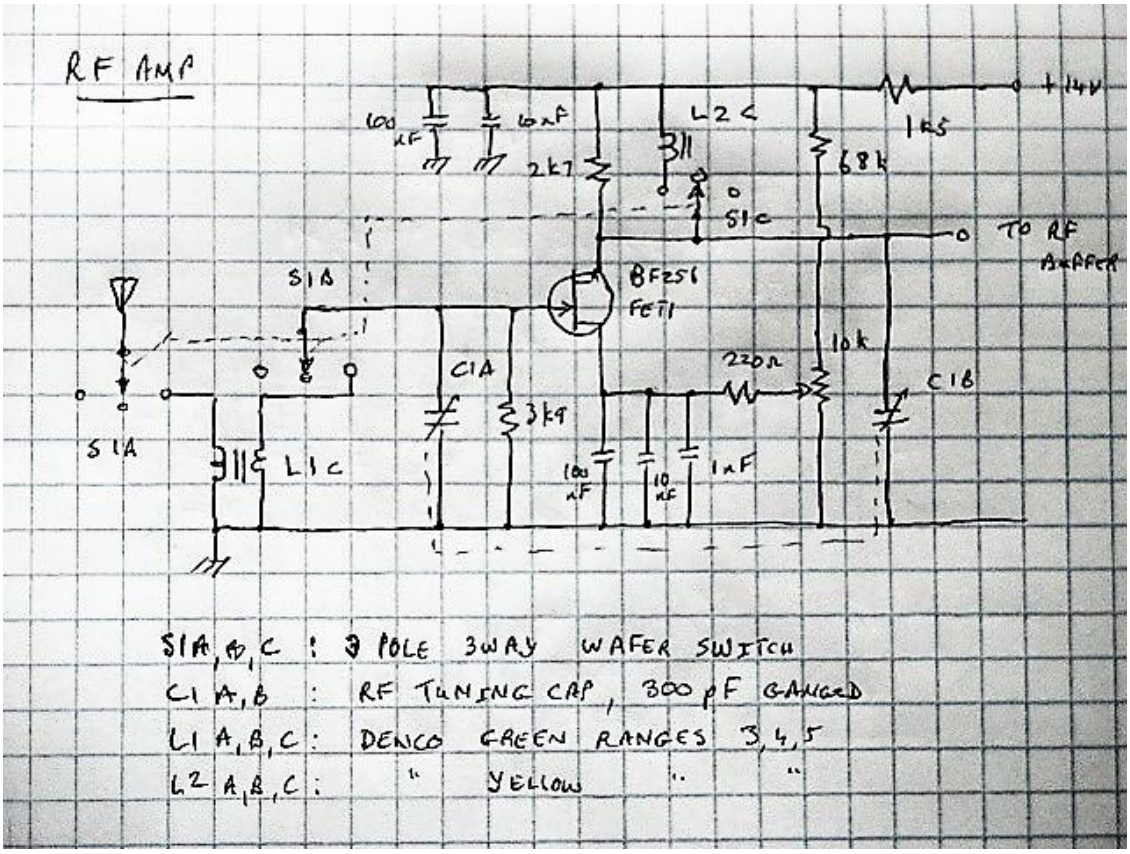

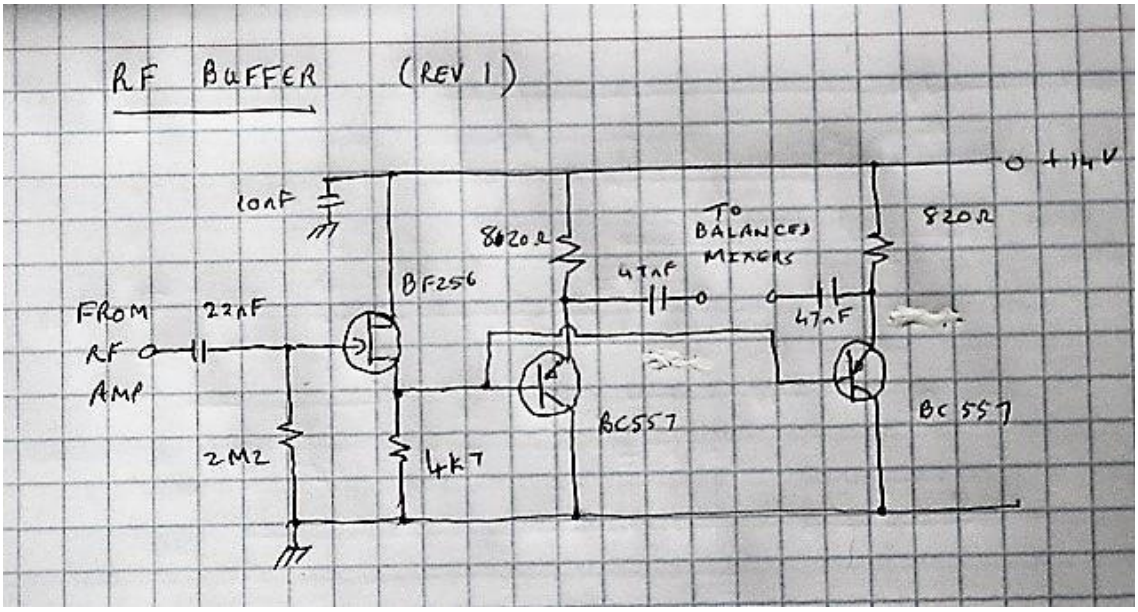

The first stage is a conventional FET amplifier followed by an RF buffer with a FET and two emitter followers to drive the balanced mixer with adequate power, low impedance and some isolation. Maximum overall gain is about 20 dB.

Initially a FET cascode RF amplifier was built but it oscillated. Similarly, a BC547 common emitter transistor buffer was used but it also oscillated. A previous-revision FET buffer without the two emitter followers had an estimated 200-ohm output impedance, which meant significant signal loss when coupled to the two 50-ohm balanced mixers in parallel. Impedance matching through transformers was considered but, in the end, the active components provide more power.

Both tuned circuits are shunted by resistors to decrease the Q and thus make RF tuning easier. In addition, it improves stability as oscillations were (still) experienced without them.

Similarly, there is no attempt made to match the FET1 circuit output impedance to the second tuned circuit.

Denco coils are used, ranges 3, 4 and 5, which were at hand. Their specs are available on the internet [www.g4dmp.co.uk/dencocoil.pdf] and there is also information on winding them yourself.

The three RF tuning bands are:

1. 1.8 – 5.3 MHz
2. 5 – 15 MHz
3. 10.5 – 30 MHz

5. Balanced mixer

The balanced mixer comes from a design by Tasic Sinisa-Tasa YU1LM [www.yu1lm.qrpradio.com/Renaissance of HF DC RX YU1LM.pdf]. It uses self-wound wide-band transformers and Schottky diodes. The impedance is low and although the author has no means to calculate or measure the impedance, it is estimated to be lower than 50 ohm.

The circuit is as follows:

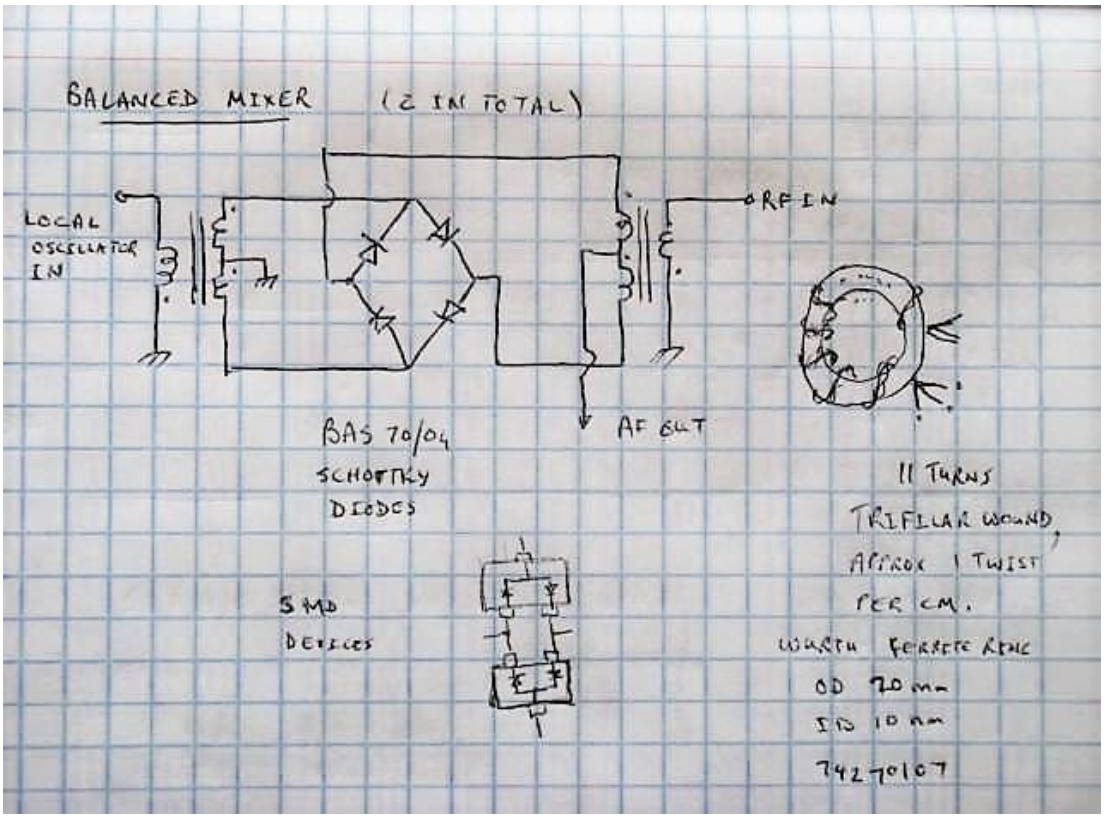

6. Local oscillators

The basic oscillator of the receiver, VCO1, comes from an idea in Wireless World of May 1972, [Cheap, Stable Local Oscillator, p218], which uses a PAL delay line to obtain a very stable oscillator which can be tuned over more than an octave, in this case 2.5 to 5.5 MHz.

This reference frequency is then multiplied up and down through a second oscillator, VCO2, which is phase locked to VCO1.

Of course, any oscillator can be used even one that is purely restricted to the amateur bands.

6.1 VCO 1

VCO1 is the core concept that drove the development of this receiver, being a very stable, tuneable oscillator. (If the author did not get this right, the rest of the receiver would not have followed, and he would have taken up Gardening or Home Maintenance instead.)

The circuit is as follows:

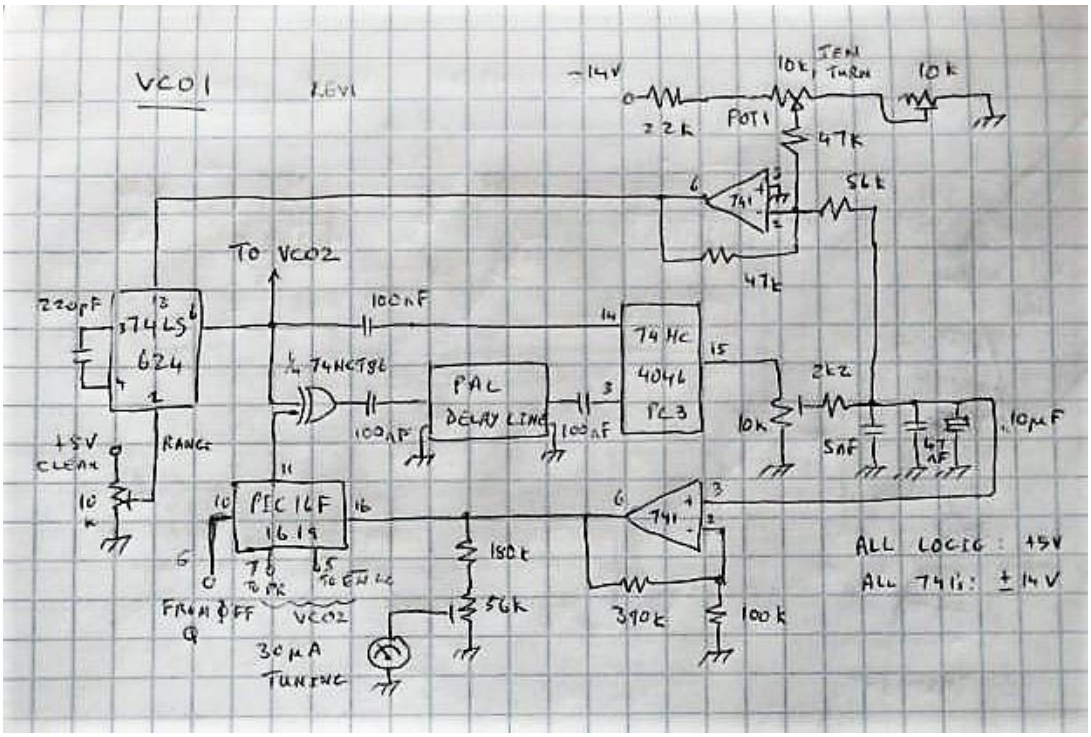

The 74LS624 VCO produces a good square wave over the frequency range. The only part used of the 74HC4046 is Phase Comparator (PC) 3, which is an excellent 360° phase detector not requiring a 50% duty cycle square wave as input. The 74LS624 is used because it has less jitter than the 74HC4046 VCO. The output of PC3 goes to the low pass loop filter and is then added to the variable tuning voltage from pot 1, a ten-turn pot, feeding back to the VCO.

The VCO output also goes via an exclusive-or gate to the PAL delay line. The delayed output goes to the other input of PC3 via a capacitor, so that a lower level output from the delay line can be accepted without amplification.

The basic oscillator tunes in theory to a roughly 16 kHz range, the inverse of the 64-microsecond delay, after which it will jump to the next 16 kHz range. In practise this range is more like 13 kHz with a 6 kHz dead band between tuning ranges. This leads to some frequencies not being tuneable at all, such as 3.8 MHz. When the frequency jumps to the next range, it jumps to the top of the new range if it is jumping upwards from the top of the old range and to the bottom if jumping down from the bottom of the range. This makes contiguous tuning somewhat inconvenient.

To solve the problem of unreachable frequencies, it was decided to flip the phase in the delay line leg by 180° when the top or bottom of the tuning range is approached. In this way all frequencies can be tuned. It is implemented by means of

a PIC 16F1619 chip which measures the tuning voltage via its ADC and then outputs a one or a zero to the exclusive-or gate to achieve the phase reversal.

Although VCO1 can now tune to all frequencies, the above leads to somewhat quirky and unpredictable tuning. Most of the time iterating between higher and lower ranges is required to tune to a specific frequency.

The front panel of the receiver has an analogue meter which displays the tuning voltage as an aid to tuning, taking some of the guesswork away. It shows where the phase flips and jumps start and the linear tuning area.

The oscillator of the 74LS624 is supplied from the Clean +5V supply, as well as the Range pre-set pot.

The software for the PIC was developed using the Microchip MPLABX-IDE development tool. It was loaded onto the chip by means of a Curiosity board.

6.2 VCO 2

VCO2 is a Phase Locked Loop (PLL), locked onto VCO1, which can tune to the following bands by successively dividing the VCO frequency by two; and produces two signals 90° out of phase with each other:

1. 1.25 – 2.75 MHz
2. 2.5 – 5.5 MHz
3. 5 – 11 MHz
4. 10 – 22 MHz
5. 20 – 30 MHz

The 90° phase shift also comes at the cost of dividing the primary frequency by two. Thus, for a 30 MHz max output, the VCO2 oscillator must run at 60 MHz.

The 74LS624 VCO is specified to operate up to 20 MHz but it operates well up to 60 MHz. It was compared to a 74S124 VCO and no difference was seen. It has the distinct advantage that two complementary (180° phase shifted) outputs, y and z, are available. These are fed into the clock inputs of the two 74S113 flip-flops which result in a 90° phase difference in their outputs.

Beyond 60 MHz the VCO still works but the PLL loses lock. High speed Schottky logic is used in the critical stages of the divider chain to minimise propagation delays which otherwise cause loop instability.

Similarly, the divide-by-two portion of the 74LS90 is used because its propagation delay is less than that of other LS or HC flip-flops.

Common emitter buffers are added to drive the balanced mixers and minimise the loading on the flip-flop. Care was taken to have minimum conductor lengths and a balanced layout for the output.

The oscillator of the 74LS624 is supplied from the Clean +5V supply, as well as the 10k Range pre-set pot.

The circuit diagram of VCO2 is as follows:

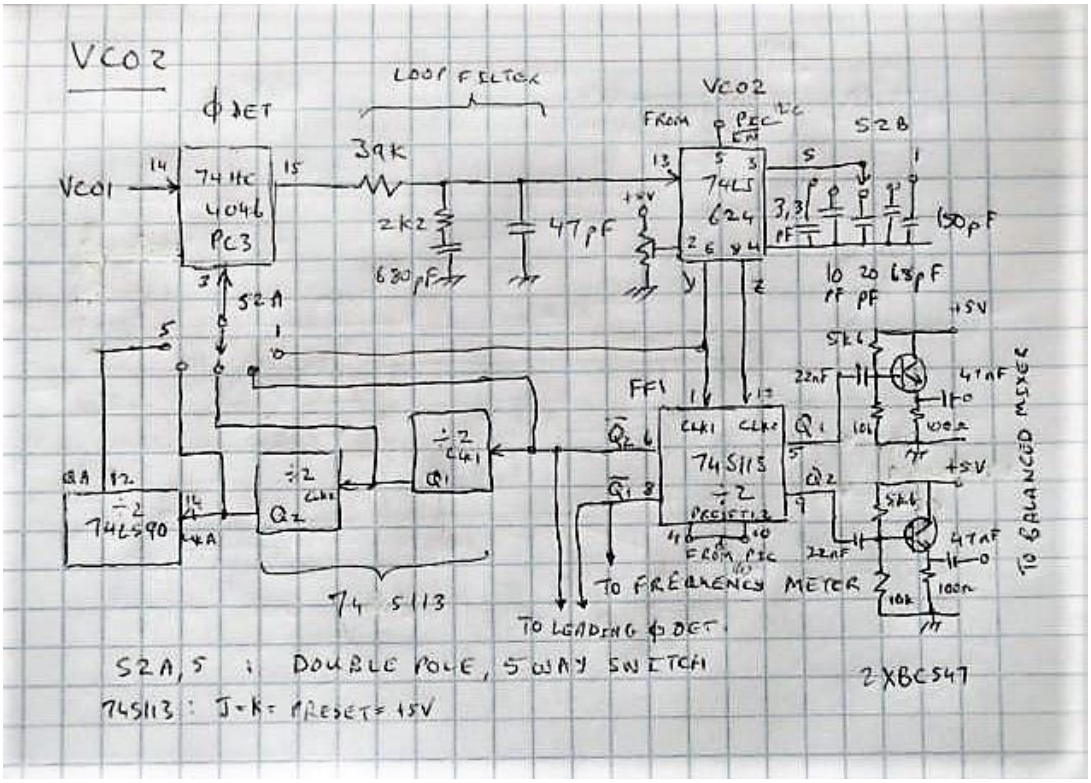

It was found that the relative phase output of VCO2 was not always consistent, sometimes Q1 would lead Q2 by 90° and sometimes the other way around. This of course leads to USB and LSB selection swopping around, which is not acceptable. The cause of this inconsistency is unknown and is seen mostly upon powering up and when VCO2 is band-switched. The solution was a 74HCT109 J-notK flip-flop which measures the relative phase, informing the PIC device of VCO1, which has plenty of spare processing capacity, when a swop occurs. The routine then issues a Preset command to the 74S113 flip-flop 1 and a Disable command to VCO2 74LS624 for a short period, which puts things right again.

This circuit is as follows. When the signal applied to the J and notK inputs of the 74HCT109 leads the signal applied to the clock input, Q is a one and vice versa.

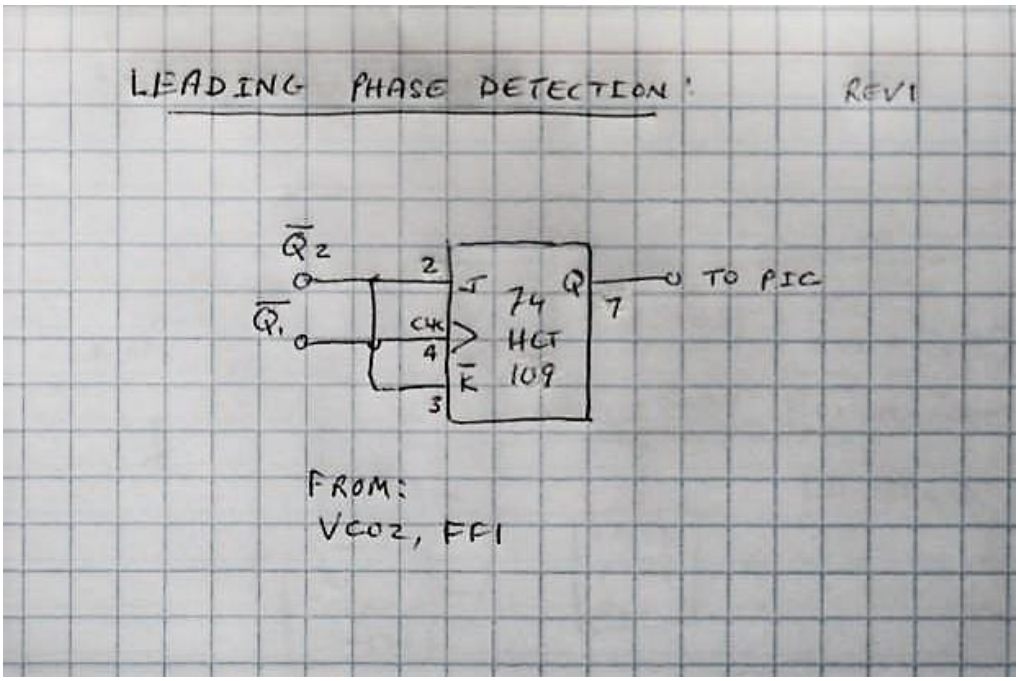

7. Low noise AF amp (Rev 1)

Hum has been mentioned as the bane of direct conversion receivers [W4JBM, Power supply modification, www.hamuniverse.com/w4jbm/12voltps.htm]. So, hum must be minimised especially in this stage and in the power supply.

The first stage is a common base amp using the low noise BC550 transistor, with active decoupling to reduce the hum. It is based on the design of Tasic Sinisa-Tasa YU1LM [www.yu1lm.qrpradio.com/Renaissance of HF DC RX YU1LM.pdf]. The gain is about 40 dB. The input impedance is low to match the balanced mixer.

Since it was found that the AGC stage introduces significant hum for small signals, for reasons unknown, it was decided to add another amplification stage before AGC in the form of a common source FET. Overall gain is now about 56 dB.

There is no attempt to filter out the RF. The 56-ohm resistor at the input is simply there to properly terminate the balanced mixer for RF.

The circuit diagram is:

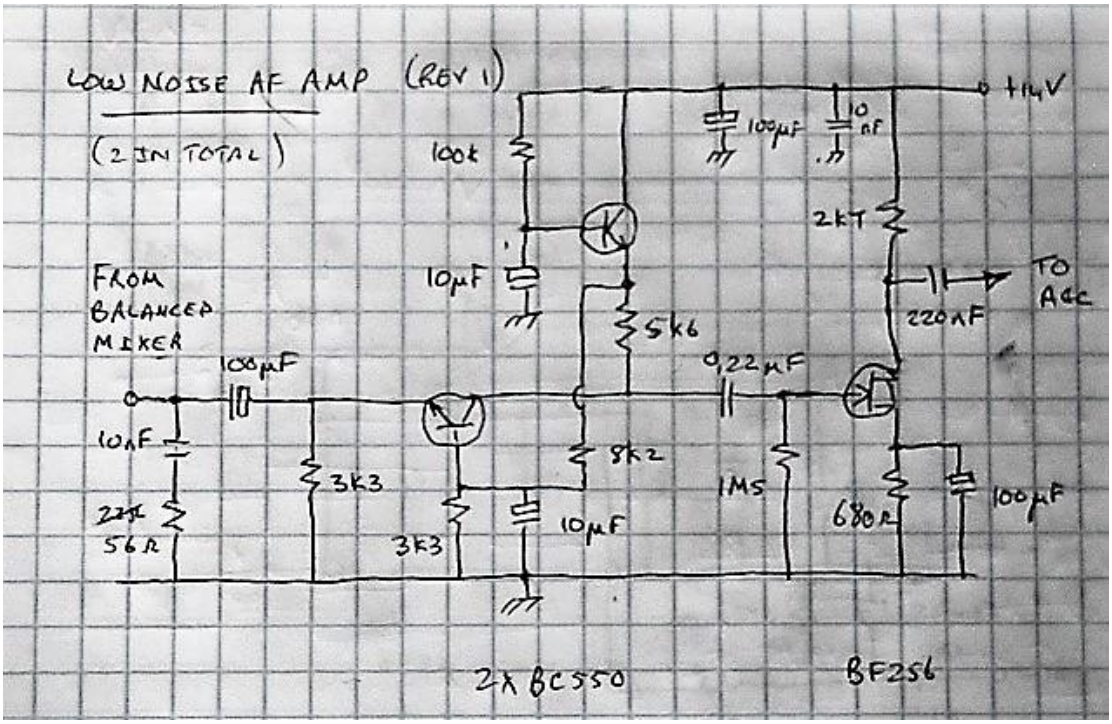

8. AGC

The basic design of the Automatic Gain Control (AGC) stage comes from the Engineers Garage website [https://www.engineersgarage.com/tutorials/circuit-design-automatic-gain-control] and other websites, where more details can be found.

The circuit diagram is as follows:

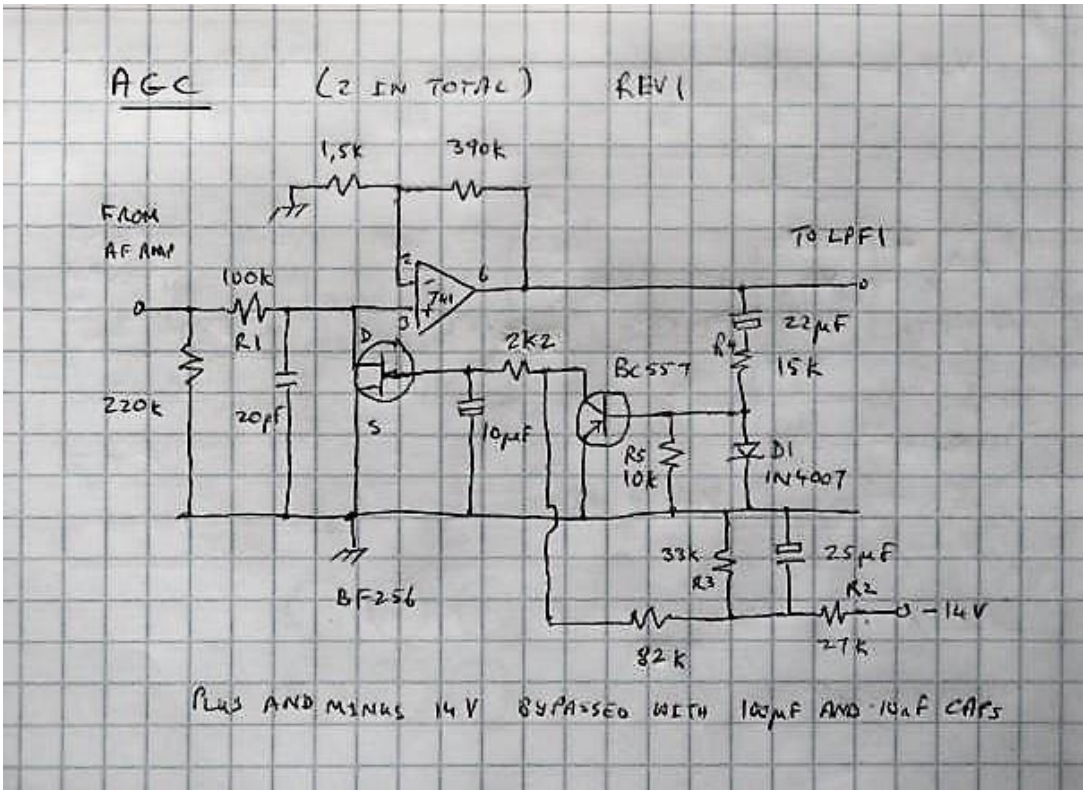

The principle is that the FET and the resistor R1 form a voltage divider, based on the amplitude of the output signal. Since the input signal to the op amp is very small, the FET operates as a resistor and linearity is excellent.

Diode D1 protects the PNP transistor from excessive reverse bias. The divider network R2 and R3 protects the FET from over voltages during transients. (Both these were found the expensive and painful way by blowing transistors.)

The output signal voltage level is set by the R4 and R5 voltage divider to about 1 Vrms maximum. The signal remains roughly at this level for the rest of the receiver up to the AF amp.

Maximum gain and the dynamic range is some 48 dB.

Since quite a bit of RF reach this stage, it was decided to filter it out by adding the small capacitor C1 to the input of the op amp.

9. Side band selection

Sideband selection takes place just before the Hilbert phase shift stage.

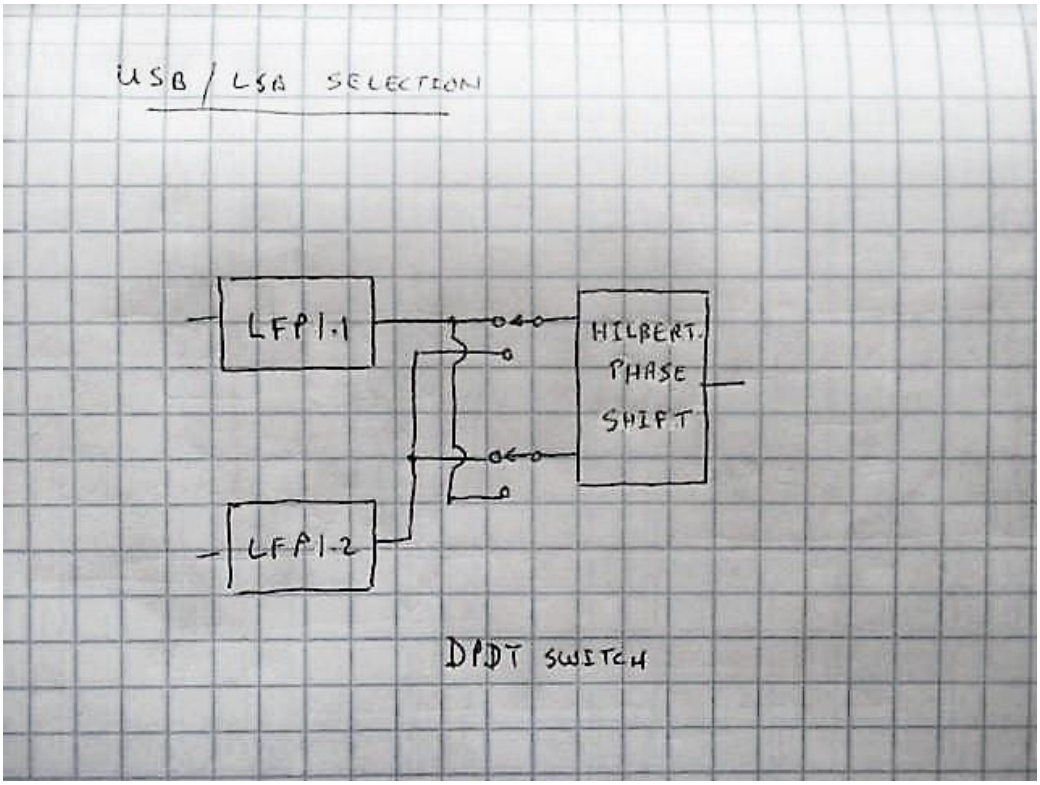

10. Hilbert phase shift

It was decided to use a digital Hilbert filter to obtain the audio 90° phase shift. Many of the ideas and all of the constants come from the Iowa Hills website [http://www.iowahills.com/].

The implementation is via a PIC 16F1619 microchip per channel clocked at it its maximum speed of 32 MHz or 8 MIPS.

Since the phase shift is not perfect across the pass band, it is better to use two Hilbert filters, one in each leg, one at plus 45° and one at minus 45°. The 90° difference across them balances out the imperfections and is very good in practise.

The microchip contains both an Analog to Digital Converter (ADC) and a Digital to Analog Converter (DAC), easing implementation. Sampling is done at 9 kHz or every 111 microseconds. The chip also contains an essential Maths Accelerator which provides integer multiplication and accumulation at 2 MIPS.

The chosen filter has 29 taps which means that 29 samples are used in the calculations. The samples are stored in a circular buffer with the newest overriding the oldest.

Every 111 microseconds the program initiates the ADC process and waits for the result. It uses one byte and calculates the output according to the formula

$Y(z) = z0*k0 + z1*k1 + \ldots. + z28*k28$
Where z0 is the current sample value, z1 is the value of the last sample, etc. The constants k0 to k28 are obtained from the named website and normalised.

Since the constants from the website filter simulation program are decimal figures less than one (if the gain is one), they are normalised to integers so that the chip can efficiently process them. Both sides of the equation above can be multiplied with a factor without any effect, which is what the normalisation does. An off-line normalisation spreadsheet converts the decimal figures to 16-bit integers which are stored in the program. It also selects the byte to be written to the DAC.

The program multiplies the 8-bit sample values with the 16-bit constants to produce a 32-bit value which is accumulated. The selected byte is then written to the DAC. The output of the DAC is a stepped sine wave for a continuous sine wave fed into the ADC.

The 9 kHz sampling of the Hilbert filter is onerous on the filtering. Much “easier” filtering would have been possible if a higher sampling frequency, such as 33 kHz was chosen. The 9 kHz rate is a design trade-off between using more taps in the filter (e.g. using 100 taps provides a much smoother and more accurate phase shift across the pass band than using 29 taps) and processing power of the chip. Using 29 taps results in the CPU being fully utilised for most of the 111 microseconds available at its maximum clock rate. A much more powerful (and expensive), 16-bit chip would have to be used for 33kHz sampling, but the selected chip provides adequate quality. Or of course analogue phase shift could have been used.

Since the ADC only works with positive voltages, the AC input signal needs to be offset so that it is only positive. This is done by means of a level-shifting common emitter stage DC coupled to the ADC.

The voltage references for each ADC are obtained via pre-set pots. This allows the outputs of the two channels to be optimised and balanced, as a change in the ADC reference voltage is similar to a change in gain.

Using one crystal clock oscillator for both PIC devices ensures that they operate exactly in step and no spurious phase shift is introduced.

The circuit diagram is as follows and the C code and other relevant files (from MPLABX-IDE) are attached.

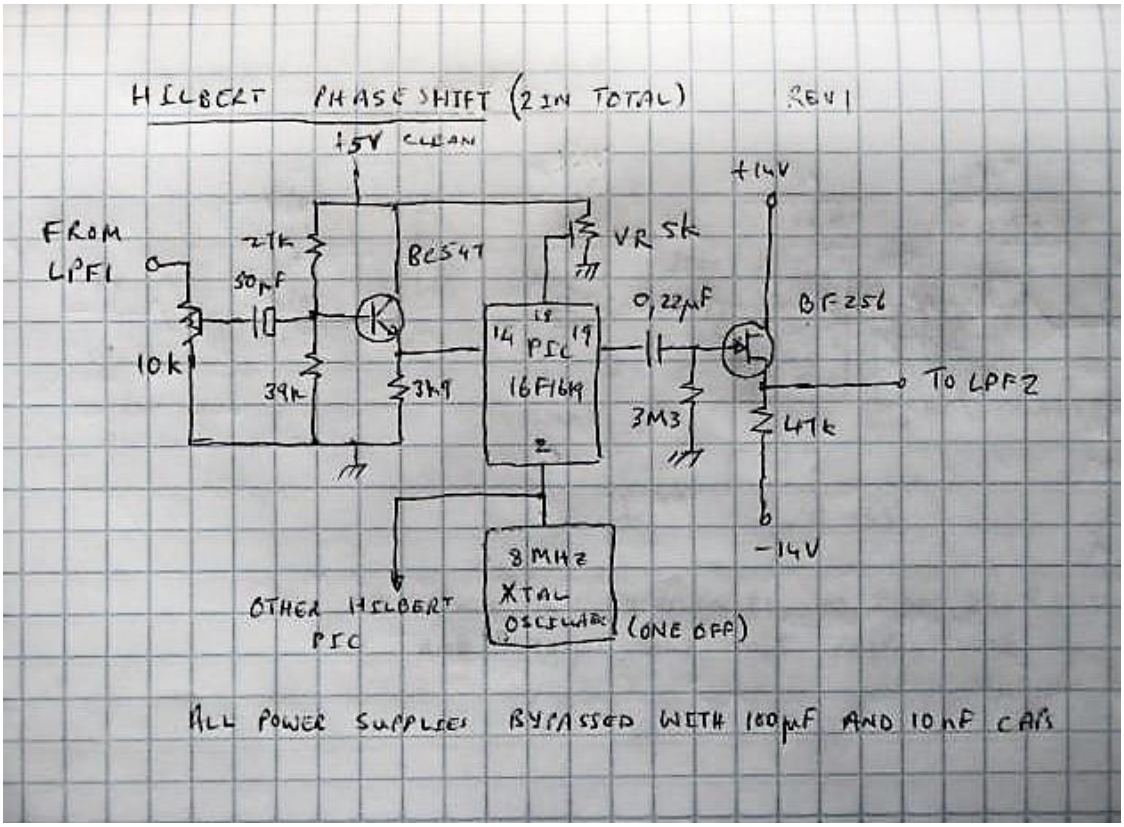

11. Low pass filters

Most of the selectivity of direct conversion receivers comes from audio filtering. In this case, the filters also provide other essential functions.

The first Low Pass Filter (LPF) is an anti-alias filter which should prevent frequencies of more than half of the sampling rate, i.e. 4.5 kHz, reaching the ADC. The maximum sensitivity of the ADC is just below 10 mV per step when 8 bits are used, so ideally frequencies of higher than 4.5 kHz should be attenuated to below 10 mV which will block them from entering the Hilbert filter. Since the maximum output of the AGC stage is about 1 Vrms or 2.8 Vp-p, reducing 2.8 V to 10 mV requires an attenuation of about 48 dB.

Assuming that the filter bandwidth would be the standard SSB value of 2.7 kHz and looking at the characteristics of (Butterworth) filters, this means an 8-pole filter is required.

The second LPF needs to remove the unwanted components from the stepped output of the DAC. The spectrum of such a stepped output contains, in addition to the desired audio, also the sampling frequency and the sampling frequency plus and minus the audio signals. Thus, for a 4 kHz signal, which could conceivably still get through the first

LPF and ADC, the unwanted components would be at 5, 9 and 13 kHz. Again, a sharp filter is needed to remove especially the 5 kHz component.

So, the decision was to implement identical 8-pole filters for both LPFs. Initially L-C filters were designed but their performance was atrocious due to wide inductor tolerances. Active filters were then designed using the TI Application Report SLOA049B and 0.5% and 1% resistors. Performance is good, they look as follows:

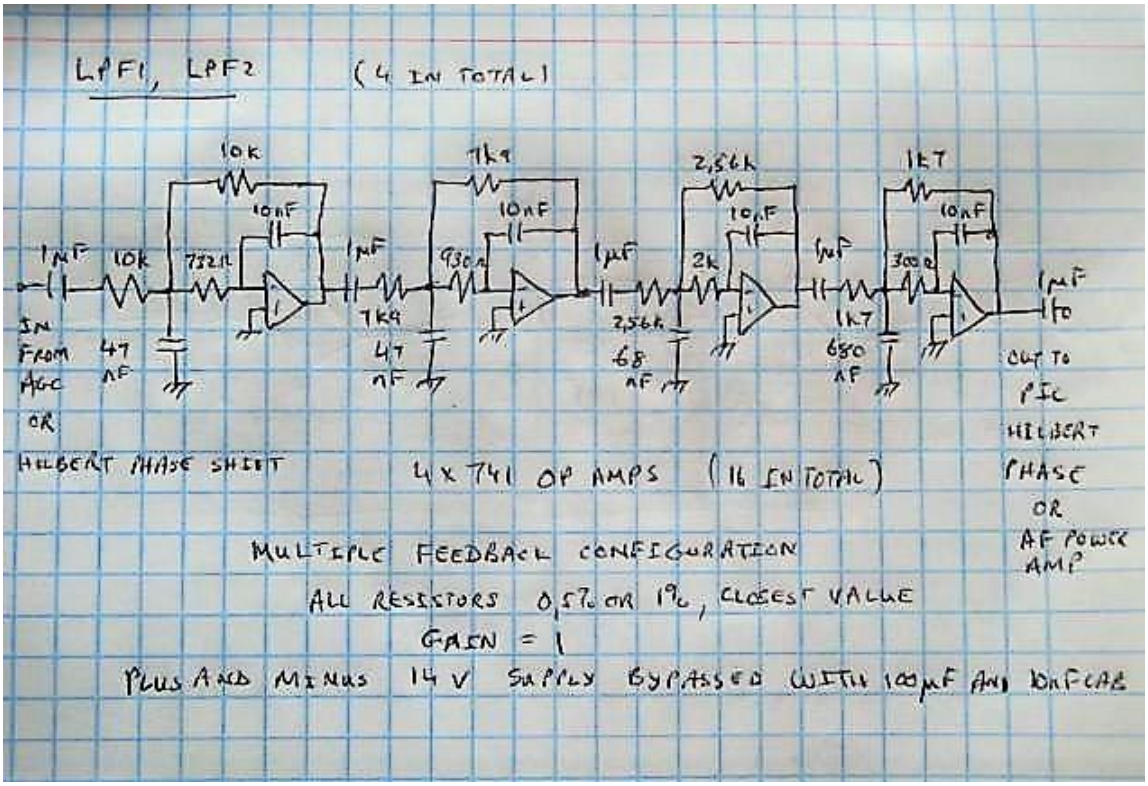

The output of LPF1 is adjustable via a pre-set pot to balance the two channels. The input to LPF2 is via a high impedance FET buffer to not load the DAC. The design gain of both LPF1 and LPF2 is one.

12. Summation circuit and AF power amp

A standard AF power amplifier, the TDA7297 is used to drive an 8-ohm loudspeaker at a maximum of 8 W, or to drive headphones. It is preceded by a volume control.

The power amp is preceded by a summation stage, which is simply an op amp. The two desired sideband audio signals arrive at the summation stage in phase from the two channels and are added while the unwanted sideband signals arrive 180° out of phase and cancel each other out, a thing of beauty to behold.

To balance the loading on the power supply, the AF power amp is driven from the -14 V supply.

The circuit diagram is as follows:

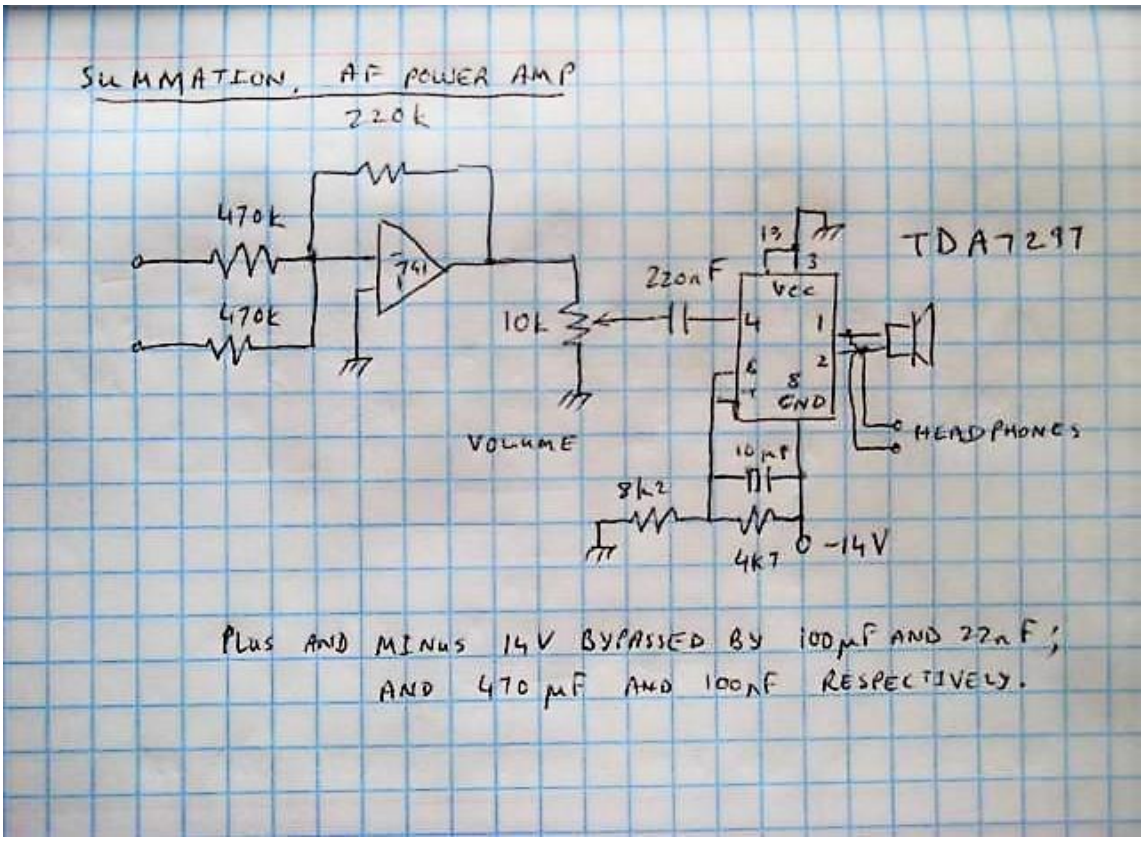

13. Frequency meter

The frequency meter is based on the design by John Main [www.best-microcontroller-projects.com/pic-frequency-counter.html]. It is implemented on the PIC 16F1619 device.

The design is such that Timer3 of the PIC counts the number of incoming pulses over the period determined by Timer1, which is 1 second. The result is displayed on an LCD device. The basic frequency for the timers is determined by a 4 MHz crystal oscillator.

The circuit diagram is as follows and the C code and other relevant files are attached.

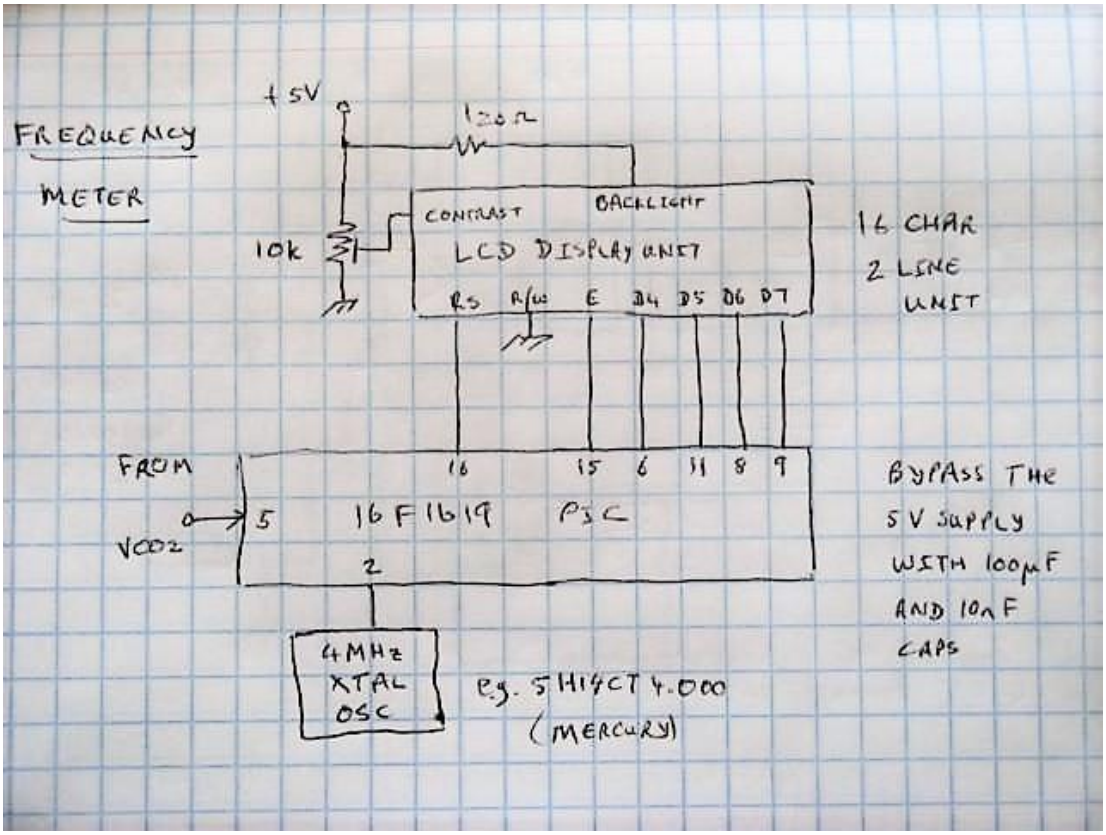

14. Power supply

The linear power supply provides the following voltages:

- 5V for the logic circuits
- 5V “clean”
- + 14V
- - 14V
- 12 V for the changeover relays (not used in the receiver)

To reduce hum, the diodes in the bridge rectifier are bypassed as per W4JBM, [Power supply modification, www.hamuniverse.com/w4jbm/12voltps.htm]. The Rules of Thumb from John Errington [www.skillbank.co.uk/psu/thumb.htm] were used for the design.

The reasons for the use of plus and minus 14 V (and not say 15V or 12V) are that some components like the FETs are rated at 30 V which means that plus and minus 15 V will overstress the FETs; while 14 V allows more power output from the AF power amp than 12 V.

The 5V regulators are preceded by an interim regulator to take the voltage down from 20 V to 13 V. This spreads the heat load and reduces the size of heatsink required.

The circuit is as follows:

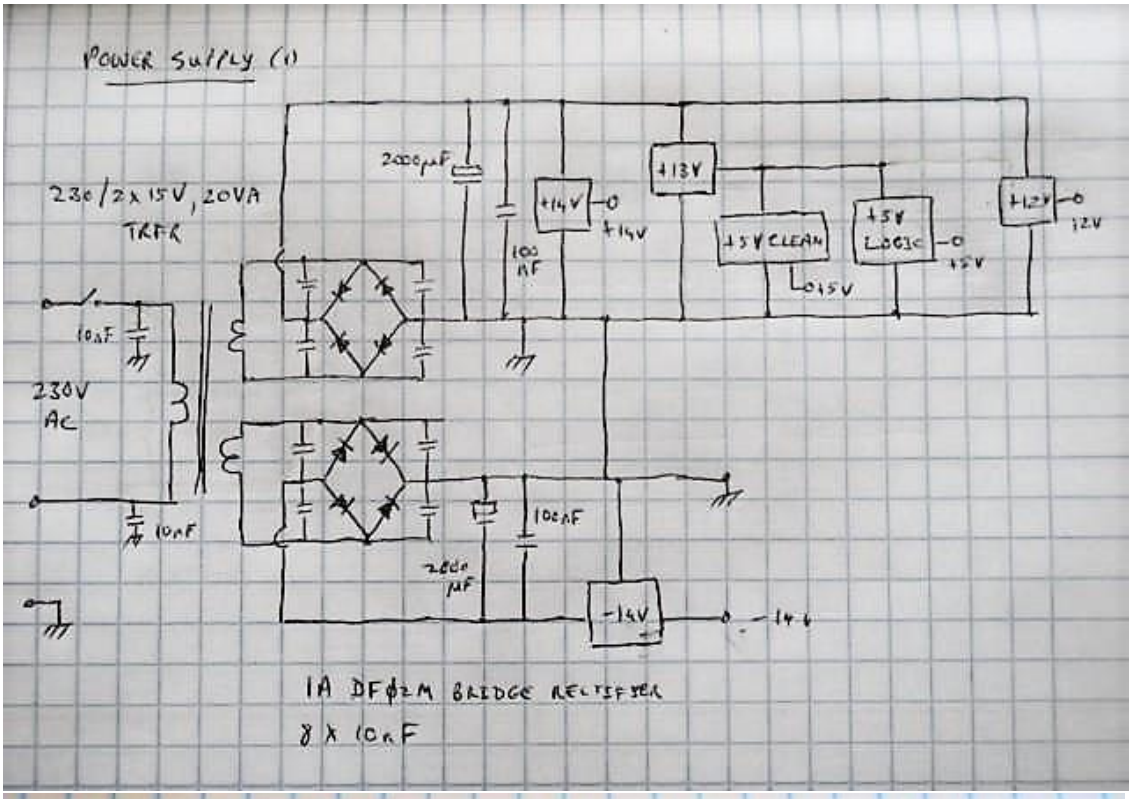

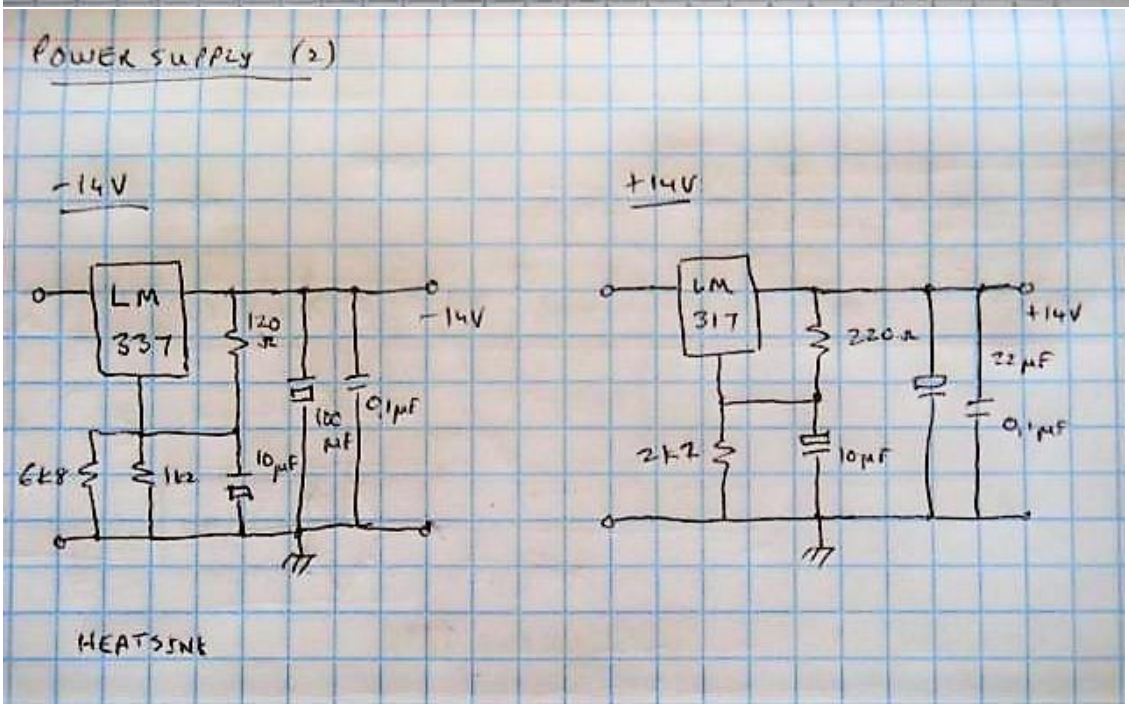

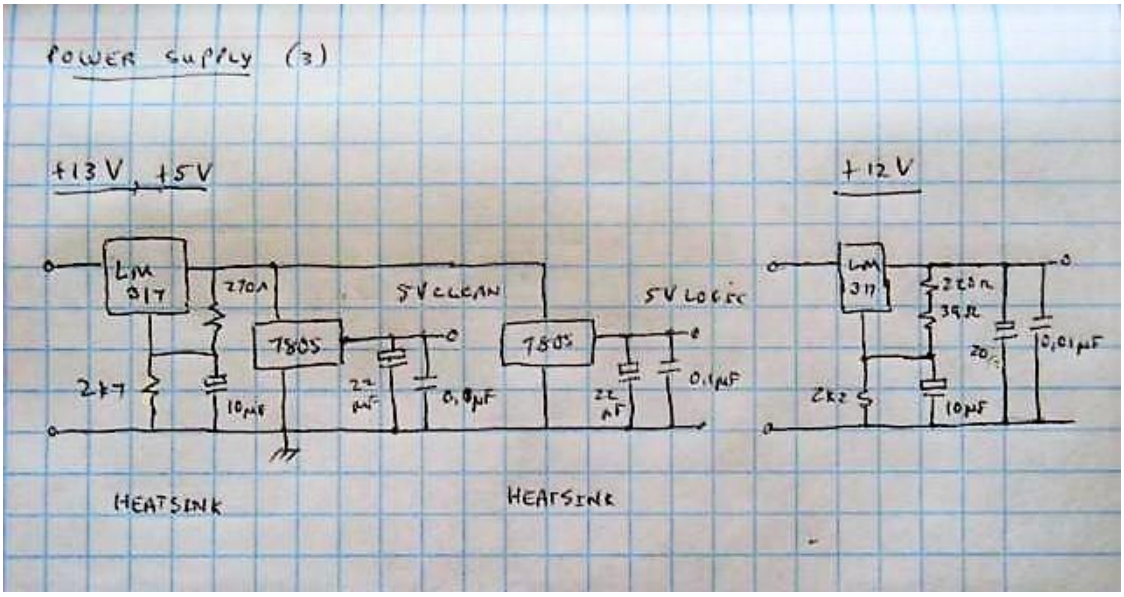

15.Setting up and performance

15.1 Setting up

Setting up is relatively simple. The tuned circuits in the RF stage need to be aligned and this is done band by band using a signal generator and an oscilloscope.

Then the two audio channels need to be balanced. Firstly, the two LPF1s are balanced by feeding them the same signal and then adjusting the pre-set pots. Make sure the highest amplitude input signal to the Hilbert phase shifter is less than about 3Vp-p so that the level shifter and ADC operate in their linear zone.

Secondly the DAC outputs are balanced by adjusting the ADC voltage reference pre-set pots while measuring at the output of the LPF2s for each sideband. At the same time, the quality of the waveform produced by the DAC is checked, as it depends somewhat on the reference voltage.

Side band selection can be checked by e.g. tuning to an amateur SSB transmission in the 40m band, where LSB is used. If swopped, reverse the connections from the LFP1s to the selection switch.

The Range pre-set post of the 74LS624s can be set to zero volt as a start and then adjusted if there are range or loop stability issues.

15.2 Performance

- Unwanted sideband suppression performance: The unwanted sideband signal is about 20 dB below the desired one at 1 kHz.

- Noise performance: The in-band noise is about 40 dB below the signal, for a decent signal.
- Audio and VCO2 phase shift performance: The performance of the Hilbert phase shift was measured through Lissajous figures to be good. The VCO2 phase shift was not measured but probably shows up in the unwanted sideband suppression.
- Jitter performance: The VCO2 jitter gets progressively worse as the frequency increases. It shows up as a frequency modulation of the AF signal. Up to 15 MHz the audio quality is very acceptable.
- Local oscillator stability: The stability of VCO1 is a function of the stability of the PAL delay line, which is excellent, the stability of the 74LS624 and the stability of the reference voltage, which causes a drift of maybe 300 Hz in the first half hour after switching on and then stabilizes.

16.Conclusion

The next step will be to convert the receiver into a transceiver by installing several changeover relays and adding an RF power amp.

The question can be asked why even bother with such a complicated design as above, when Software Defined Radio (SDR) or digital signal processing (DSP) chips are available these days. The author is afraid the answer is similar as to why people climb mountains. The challenge to design and build this receiver has been there for many years and it was important to “do it himself”. At the same time, the performance of this receiver is quite acceptable. (Of course, SDR caters for all other types of modulation as well. DSP adds additional features such as noise blanking.)

The receiver is actually “SDR or DSP ready.” The wide-band signal just before the AGC stages can be taken to an SDR receiver or DSP chip, although the bandwidth of the RF stage needs to be checked and perhaps increased.

When the receiver was being tested, since the VCO tuning is “quirky”, it was very useful to have a WebSDR display up when tuning. Active transmissions were easily identified, and the receiver could be quickly tuned to that frequency.

17.Photographs