

This project represents my first foray into radio frequency design, although it's far from my first electronics project. Somewhat counter-intuitively, this has resulted in it being more complex than an equivalent receiver from a more experienced designer – because I was not sure of what would work, I adopted a “belt-and-braces” approach whenever I could. For example, the receiver has two AGC circuits; not because I think this is the best arrangement, but because I was not sure which would work best. Extra complexity was also added by a lack of top-down design and “noodling” before solder melting commenced. Many of the functional blocks were built and tested as stand-alone circuits, and then required additional interfacing circuitry to get them to play nicely together (the gain staging in the IF amplifier is an example).

The RF circuits are constructed Manhattan style on copper clad boards, and enclosed in diecast aluminium boxes to form modules. Connection to the modules is mostly, for reasons of economy, by RCA phono connector. Legend has it that these have very acceptable RF performance into the UHF, although it turns out that they lack screening – for this reason some of the more sensitive connections are made by BNC. Each module accepts 15V dc which is regulated to 12V locally, and DC (including CV) connections are all filtered through a “poor man’s feedthrough filter” consisting of an axial inductor and ceramic capacitor mounted at the connector. These filters, as well as the decoupling used liberally throughout, are not shown on the diagrams for reasons of clarity.

The AF amplifier sections are built on stripboard and (at the time of writing) are all al-fresco. The various digital circuits are also on stripboard, with the microcontroller and FPGA being on commercially made breakout boards. The 7-segment display is also a commercial module.

The receiver itself is a conventional single conversion superheterodyne design with an IF at 10.7MHz. To the best of my knowledge, this isn’t regarded as an ideal IF for shortwave receivers, but my hand was forced by the crystal filter I had available. This is very much a mono-band receiver for 20m, having a tuned RF amplifier and a LO of very restricted range – both of these blocks would need to be changed to make it suitable for another band. It is also very much just a receiver. No provision has been made for re-using the elements (mixers, amplifiers, filters) in a transmit section. This, again, is a result of it being my first radio project, and quite complicated enough without having to design-in a future transmitter at the same time. In the final build it is intended to take a feed from the two oscillators so that a separate transmitter can be constructed to operate in synchronicity.

The receiver features an RF front end amplifier, two stages of IF amplification and further amplification at AF, the IF and AF stages being equipped with AGC. All three stage gains are adjustable by a control voltage, since I had some temptation to make the entire thing remotely controllable by software. The two oscillators are also voltage-controlled (using varactors) and are tuned under command of a microcontroller – which allows tricks such as frequency presets, as well as holding a steady tuning frequency. The control loops for the VCOs are completed by a pair of frequency meters implemented on an FPGA. Although I experimented with other (cheaper, more elegant) ways of measuring frequency, in

the end the FPGA (which had lain dormant in a box for several years, waiting for a project) with its combination of speed and density, allowed a brute-force approach to the problem which would have been impractical with discrete logic ICs. The FPGA configuration was entered as a schematic which, if nothing else, illustrates vividly why HDLs are the right way of doing it. The MCU is an STM8 and is in control of everything frequency-related in the system.

Construction was done with no specialist RF testgear, although I am lucky enough to own a scope fast enough to at least look at the signals. Therefore, all frequency responses, etc, are taken on faith and have not been measured.

1. RF amplifier

The RF amplifier is tuned. Ideally, of course, it should have a flat response across the band with a steep skirt each end – however the Q is low enough that it is broadly flat from 14 to 14.35MHz. Combined with the 5-pole roofing filter, the image frequency at 35MHz should be on the order of 90dB down vs. the wanted signal at 14MHz (numbers from simulation, so to be taken lightly salted). The amplifier gives about 26dB of gain - not as much as expected, but worth having - and simulation (again) suggests an input-referred noise of 1.5nV per root Hz.

The amplifier itself is best described as a common base amplifier with feedback. The input impedance is low (a bit over 50 Ohms in this case) and the transistor acts to direct the input current through the collector load (a tuned circuit) but at a higher source impedance – allowing a high Q and high voltage gain at resonance. The gain is adjusted by simply damping the Q of the resonant tank by turning on the PIN diode. This obviously harms the relative rejection of out-of-band signals but, assuming it's only done in response to strong inputs, this should not be too detrimental. The fast SS9018 transistor proved prone to parasitic oscillations, which had to be quenched by a ferrite bead on the emitter.

Borrowing from Hayward & Kopski's popular TIA amplifier, the output is buffered by a pair of emitter followers to decouple from the tank from the load.

2. Local oscillator

The LO tunes from 24.7 to 25.05MHz, with a little to spare each side. When setting up (by adjusting the variable capacitor), the oscillator span is set such that the required frequency range is towards the top of the span. At this point the control voltage is highest and the varactors are operating in the flattest part of their logarithmic response, which makes for the most accurate tuning and lowest susceptibility to control voltage noise.

The oscillator itself is a textbook Clapp type, built around the 2N3819 JFET. The Zener holds the source at a positive potential, allowing the ALC to pinch off the JFET by pulling its gate towards 0V. The

automatic level control (ALC) is tapped from the circuit output, post buffer, and operates to bias the JFET off once the output waveform is sufficient to turn on the BC337.

The output is buffered by a cascode amplifier (one of a number of RF amplifier topologies I experimented with during this project) which adds 6dB or so of gain, and then further buffered by an emitter follower to the output. The technique of adding an RF choke into the emitter load seems to improve linearity a little for the same idle current (by helping sink current during the negative half cycle).

The inductor is air-cored (in an attempt to get maximum Q) and the (salvaged) trimmer capacitor I believe to be about 5pF. The so-called varactors are simply a pair of common 1N4002 rectifiers, which have adequate capacitance swing for this application.

3. First mixer

The mixer itself is a Minicircuits SBL-1, with its LO port is driven from the oscillator output via a series resonant circuit by way of filtering. Because the oscillator was by this point completed (with a single output), and I had no desire to re-visit it, it was decided to pass the LO signal through the mixer module on its way to the frequency meter input. To this end a buffer amplifier with +10dB gain was added; for the sake of simplicity, this was constructed using an AD8055 high-speed opamp.

The 9-kHz wide crystal filter has a terminating impedance on the order of 800 Ohms, so a 1:4 transformer was used to match this to the 50 Ohm output of the SBL-1. Originally a buffer amplifier had been used between mixer and filter but, while it worked, I judged it must be adding unnecessary noise and that a transformer would be a quieter solution. To fine tune the terminations, a trimmer resistor and capacitor are paralleled at each end. The filter was found to be fairly insensitive to the driving impedance, but far more so to the capacitance at the output end. The termination was adjusted by using a jerry-rigged 14MHz oscillator (similar in design to the LO and BFO) at the RF input of the mixer, whose frequency was swept by a slow triangle wave. By synchronising the oscilloscope trigger to the triangle generator, the envelope of the IF waveform could be observed as the frequency swept up and down through the filter passband.

4. IF Amplifier

The IF amplifier is in two stages, each giving a maximum of about 35dB gain - as other amplifiers, the gain is voltage controlled. The IF Amp assembly also contains an envelope detector which is used for signal metering, AGC and AM detection.

The amplifiers are the same basic design as the RF front end, augmented by a second input transistor to form a differential input pair. This gives the stage a high input impedance and boosts the gain, as well as (we hope) reducing the second harmonic distortion. The collector load is again a series tuned tank, but the additional input transistor presents a problem when the gain-regulating PIN diode is brought into the picture. In order to keep the two transistors' collector currents roughly balanced, a second (1N4148) diode is added, to draw approximately the same current as the BAP64 PIN diode. This is an inelegant solution, and probably not very effective, and so requires further work. The tail current of the input pair

limits the maximum input signal to about 50mV peak, regardless of the gain; this caused some problems when trying to arrange the AGC (see following section). This block uses a Sziklai pair follower, which should provide better performance than the Darlington in the RF amplifier.

The two amplifier sections are identical save for the additional coupling network at the final buffer, which also runs a higher bias current in order to deliver 500mV peak into the second mixer. The two sections had to be de-tuned slightly from each other to prevent the entire assembly oscillating.

The output is tapped off to drive an envelope detector based around a pair of Schottky diodes and a dual op-amp. The arrangement provides some positive feedback in the bias voltage on the first diode – the one that does the rectification. The second diode carries the same average current, such that the drain of the BS170 MOSFET is two diode drops below the output DC voltage. Half of this voltage is then fed to the first op-amp, which forces the junction of the two diodes to equal this voltage, thus forming the positive feedback loop. Some additional bias may be added via the preset resistor, and circuit output is in the form of current from the BS170's source.

The output from the envelope detector drives two auxiliary outputs: Directly to an AM output and via a peak detector to the signal metering circuits. Note that the signal metering (which is via a 10-segment LED bargraph) is tapped from after the IF amplifier, and so depends upon its gain setting, rather than directly reflecting the power incident at the antenna input as would be normal.

The envelope detector also drives the AGC circuit, with the next op-amp in the chain setting the threshold and ratio of the compressing action (its output swinging high when gain reduction is required). Also input to this circuit is the manual gain control voltage, which is combined with the AGC signal in a highest-takes-precedence manner via two diodes. The resultant gain control signal is applied to the input of a logarithmic current mirror (with a little diode-shaping to improve the control law) which, in turn, sinks current from the PIN diodes in the amplifiers.

Ideally, the total gain of the amplifier chain should be reduced by lowering the gain of the second stage before the first – a scheme which preserves the greatest signal-to-noise ratio. However, reducing the second stage gain to unity would leave the first stage (still operating at +35dB) attempting to drive its input at up to 500mV peak – far in excess of the 50mV maximum swing mentioned above. To keep the second-stage input drive below 50mV, its gain cannot drop below 20dB. The ZVP4105 P-MOSFET is so arranged as to cut off the control current to the second stage at a preset level, transferring all additional current (and therefore gain control) to the first stage. The overall gain range is therefore about 50dB.

A final 2-transistor sub-circuit detects when the AGC signal has exceeded the manually programmed gain and illuminates an LED to indicate gain reduction.

5. Beat Frequency Oscillator

The BFO is identical, save for some component values, to the LO.

6. Second mixer

Likewise, the second mixer is similar to the first. In place of the crystal filter, a simple LC lowpass filter removes RF mixer products from the signal before it passes to the AF stages.

7. Audio amplifier

The audio stages can be divided into three blocks: The preamplifier, the filter and the power amplifier.

The preamplifier is built around an NE5534 op-amp and has a voltage-controlled gain up to about 30dB with AGC. The gain control element is a 2N5457, with the gain being reduced as the gate is brought down towards ground. The diode/resistor network linearises the control law. The control voltage, normally driven from the gain pot, is pulled low by the BC337 as the AGC becomes active – the flow of current through the 820R resistor causing the gain reduction LED to illuminate. The maximum AGC threshold is set by the B-E threshold of the BC337, and can be reduced by applying an external bias voltage.

Following the preamplifier is a switchable phone bandwidth filter. Since the crystal filter is too wide (at 9kHz) for normal SSB phone signals, a further filter at AF is required. The effect of this is that the lower passband edge is provided by the crystal while the upper edge is provided by the active filter at AF; this gives that interesting property that the lower edge (ie, the highpass frequency) of the phone passband can be adjusted independently of the upper edge, by adjusting the BFO. Of course, it also has the disadvantage that large signals in the 6kHz between the top of the desired phone signal and the upper crystal cutoff frequency can still overload the IF stages, despite being “inaudible” to the operator. The filter is an 8th order Butterworth - while simple to build, with hindsight a sharper cutoff would have been desirable. The full or limited bandwidth signal is selected to drive the power amplifier by a J112 JFET switch.

The audio power amplifier is built around an LM301 op-amp with a discrete power stage, the loudspeaker being capacitively coupled to the rails. The preamplifier AGC threshold is set so as to keep the power amp out of clipping, since it will only deliver around 2V peak into a 4 Ohm speaker. To minimise crosstalk the power stage is driven directly from the main 15V supply, not the locally regulated 12V.

8. Digital section

Much of this section is mundane, and is presented in block form on the schematic. The microcontroller is in charge of tuning the receiver, which it does with the help of an FPGA. The MCU runs a software servo loop to control the two oscillators while the FPGA implements the frequency counters and various I/O blocks. Communication between the two is by SPI, with the MCU as master. As the various FPGA functions are essentially independent, a bus of 3 address lines is used to select the desired functional block.

At the centre of the frequency servo mechanism is a PCM1753 24-bit audio DAC, which provides ample resolution for the control voltages required by the two VCOs. The (inherently noisy) output from the DAC is amplified and heavily filtered before leaving the digital section (there being more filtering in each VCO block).

In an attempt to avoid contaminating the analogue ground, the signal from the oscillators is coupled in via transformers and large common mode chokes. Having been shifted to the centre point of the 3.3V digital power domain, the oscillator signals are fed into the FPGA via LVDS pin pairs (these are ideal for the job).

9. FPGA overview

The Altera Cyclone II FPGA hosts several independent circuit blocks, which are described below.

9.1 Frequency meters

These high-speed 24-bit counters represent a brute-force approach to RF frequency measurement. Each frequency meter consists of a “first” counter which is preloaded via SPI while the block is selected. As soon as the block is deselected the first counter commences counting, at a fixed 50MHz rate, towards the preload value. While the first counter is counting, a gate is opened allowing a second counter to count at the unknown VCO frequency – when the first counter reaches its preload value the gate closes and both counters stop. The block remains idle until it is selected again, and while a new preload value is clocked into the first counter, the value accumulated on the second counter is read out. Both counters are automatically cleared as the block is deselected.

9.2 I2S generator

The FPGA is also a convenient place to generate the I2S (inter integrated-circuit sound) bus signals for the DAC. This block consists of a chain of serial-to-parallel latched shift registers which de-serialise incoming SPI data and a complementary chain of parallel-to-serial registers to re-serialise it at the correct I2S rate. Some synchronisation logic allows the MCU to write an update at any time, with the new value only becoming available at the end of the I2S word, to avoid glitches. The various I2S clocks are also generated, and the clock rate can be selected (one of two values) by a “spare” output from the GPIO section. This allows the MCU (and thus the operator) to toggle the DAC clock frequency and thereby move any spurs generated by it – despite my best efforts, such spurs do still exist!

9.3 General purpose input/output

This section handles the operator interface. There are 16 input lines total for 8 buttons and 4 encoders, which are multiplexed to 8 lines by the input buffer hardware. The FPGA handles demultiplexing, debouncing the buttons and decoding the rotary encoders (this saves precious program memory in the MCU). The result is presented as a 16 bit word and serialised over SPI. The incoming SPI word is deserialised and latched to outputs – buffer hardware drives LEDs.

9.4 SPI pass-through

The final addressable block is simply a pass-through that echoes the SPI clock and data to the 8-digit display module. Thus, the MCU can access all its “peripherals” on one bus.

10. Frequency servo

Each VCO has its own software servo loop (also frequency meter and DAC channel). The VCO set-point frequencies are calculated from the desired tuning frequency and BFO offset, meaning that the BFO frequency can be adjusted while the receiver tuning does not change.

Running each loop is a constant sequence of starting a measurement on the frequency meter, waiting for that measurement to complete, then adjusting the DAC and starting again. Because high-resolution and fast response are both required, the frequency meter gate times are varied adaptively (high-resolution measurement requires a long gate time, but this precludes fast operation). With a known set-point frequency, software makes a fast but low-resolution measurement of the VCO using the shortest gate time (about 3ms). If the measured frequency is within the error bounds for that gate time, the gate duration is doubled and the process repeated. At each step the error between the measured and set-point frequencies is used to alter the DAC setting. The longest gate time available is 400ms, for a 2.5Hz resolution. If the frequency error becomes too great then a shorter gate time is selected, in the opposite of the process described above.

Because the VCOs do not respond linearly to control voltage, a servo loop calibration is performed on start up. This consists of sweeping the control voltage in 16 steps across its entire range; at each step the difference in frequency obtained is recorded. Later, when the servo is active, the inverse of this value can be used to determine how many DAC counts are required to deliver a required frequency change, at any particular operating point in the DAC range.

Although the servo system works quite quickly, it is nowhere near as fast as a PLL, and is not fast enough to give a pleasant feel to manual tuning. In order to achieve a responsive tuning action, the servo system is disabled while the tuning control is in motion; this gives two distinct control regimes: “Servo locked” and “free tuning”. In the latter mode the frequency meters are still active, but serve solely to update the display – as a digital readout of the VCO frequency. The VCO control voltage is adjusted up or down according to the motion of the 2400ppr encoder, the rate of increment being adjusted according to the calibration table described above, so as to give a linear tuning scale. When the tuning control stops moving, the last measured frequency is stored as the servo set-point and the servo is switched on, keeping the receiver on the frequency to which it was tuned.