Technical description - KA7OEI Eclipse TX and RX equipment

Transmitter description:

Overview

The goal was to *simultaneously* transmit identical signals on as many HF amateur bands as practical, the prime signal being FST4W, one of the modes in the WSJT-X suite of protocols. This mode is unique in that as part of its receive statistics it provides a reading of drift during the *(nearly)* two minute cycle which, when combined with stabilized transmitters and receivers, may be used to divine Doppler shift. As a solar Eclipse results in a nighttime-like ionospheric configuration, being able to measure this shift may provide insight into how these layers behave before, after, and during the eclipse itself.

I also wished to produce two more carriers that could allow additional analysis: The “T1” signal being modulated at 0.5 Hz using BPSK - the transition occurring in sync with the GPS second - and the “T2” signal being an unmodulated carrier. Both of these carriers could be used to provide additional points of measurement - most notably Doppler, - and the “T2” signal trivially so, offering greater frequency and temporal resolution than that of the FST4W signal.

With identical signals being emitted across the HF spectrum, having the multiplicity of signals to analyze individually may also help determine to what extent the ionization is affected by noting both the Doppler shift and amplitude of all of the disparate signals with each amateur band being affected differently by the undulating ionosphere.

The Beacon Blaster and the BPSK modulator/PIC IDer in the lid. This photo was taken prior to the addition of the "T2" carrier circuitry. The combiner and MMIC amplifier may be seen on the far right of the chassis.

Brief description of the transmitted signals and method of generation

All transmit signals are initially generated at about 45 MHz - the transmit “IF”. This signal is then mixed down to each amateur band as needed. There are three signals generated: The FST4W signal, the “T1” carrier and the “T2” carrier. The composite signals on each band are then amplified using a QRP-Labs 10 watt linear amplifier *(one per band)* and then combined to common antenna(s) using series-tuned band-pass filter networks.

The FST4W signal itself is set precisely at 45.000 MHz using a “Beacon Blaster” running a configuration modified as follows:

- The FST4W signal is generated “upside-down” owing to the use of a high-side local oscillator to convert to the individual amateur bands.
- Optionally enabled *(by flipping a switch)* there are two other signals that may be transmitted along with the FST4W - each of these other signals also being produced *(unmodulated)* by the Beacon Blaster:
  - **The “T1” carrier:** At 500 Hz below the FST4W signal is another that is BPSK-modulated by the 1PPS signal from the GPS receiver onboard the Beacon Blaster: The phase is flipped 180 degrees on the GPS second, plus additional delays, yielding a modulation rate of 0.5 Hz. This signal is nominally 10 dB below that of the FST4W signal.
  - **The “T2” carrier:** At 988 Hz below the FST4W signal is a CW *(unmodulated)* signal. This is nominally 15 dB below that of the FST4W signal.
  - Both the BPSK and CW carriers are identified by Morse using amplitude reduction every 600 seconds *(10 minutes)*.
- The BPSK and CW signals, when down-converted, appear ***above*** the FST4W signal in frequency so that they may be recorded/captured using a upper-sideband receiver using standard SSB bandwidth tuned to yield 1500 Hz signals at the center of the WSPR subband.

The precise transmit frequencies *(e.g. my FST4W signal being 92 Hz below the center of the WSPR band, or 8 Hz above the bottom)* were set on a per-band basis in the programming of the local oscillator for each band. The configuration of the BeaconBlaster is discussed in Appendix A, near the end of this document.

Discussion:

The rationale of the additional carriers are as follows:

- **BPSK carrier (“T1” signal):** It ***may*** be possible to determine absolute time-of-flight between the transmit and receive sites by determining the timing of the 180 phase change of the BPSK signal and comparing this with a similar signal *(on a slightly different frequency)* produced at the receive site. If this is possible it would likely be so by analyzing and correlating multiple phase transitions over time to reduce uncertainty.
  - It has been suggested that a better means of this would be to increase the “chipping rate” of such a signal. One idea considered would be a QPSK signal with the “I” channel conveying a 256-bit PN sequence and the “Q” channel carrying the BPSK signal modulated at the 0.5 Hz rate.
    - A 256 Hz chipping rate was considered as it is a convenient 2^n value, and below the *(current)* FCC-mandated maximum HF band “baud” rate of 300 Hz.
    - Because “Spread Spectrum” is not legal on HF, the use of SS-related terms like “chipping” rate and “spread spectrum” would have to be avoided.
  - Generating this 256Hz/0.5Hz QPSK signal was not done due to lack of available time to construct such hardware - plus the lack of a means to “receive” such a signal.
    - Additionally, the bandwidth of such a signal would likely require that it be placed elsewhere, away from the WSPR sub-band, likely too far away to be received using the same SSB bandwidth as one might use for WSPR reception.
    - Without precise filtering *(requiring much more complicated hardware and software)* the sidebands of such a signal would likely be visible over SSB bandwidths *(e.g. main lobe and sidebands extending past 1 kHz)* further complicating deployment.

- **CW carrier ("T2" signal):** This was transmitted to enable the trivial, real-time analysis of absolute frequency and Doppler Shift: One simply plots the frequency of this signal over time. While FST4W can measure Doppler shift, there are several complications:
    - **Low temporal resolution.** The Doppler numbers from FST4W are, at best, available for 2 minute windows - a very significant portion of the duration of "totality" of the eclipse.
    - **Lower absolute frequency resolution.** The measurement of the absolute frequency of an FST4W signal - using WSJT-X tools - is much lower than the Doppler Shift, making precise tracking of Doppler Shift *(sign and magnitude)* more difficult. The direct measurement of the Doppler shift of the "T2" carrier is trivial by comparison.

45 MHz exciter - detailed description

This device generates not only the main FST4W signal, but also the auxiliary "T1" and "T2" carriers which are typically transmitted at a lower signal level than the main. The "T1" carrier is BPSK-modulated at 0.5 Hz with the transitions occurring in sync with the start of the GPS second while the "T2" carrier - transmitted at a lower level than T1 - is unmodulated. These two auxiliary carriers may be used for direct analysis of Doppler and amplitude shifts during transmissions.

The heart of the 45 MHz exciter is a Turn Island Systems "BeaconBlaster-6" which contains an Atmel processor, GPS receiver, a non-precision 10 MHz oscillator and several Si5351-type frequency synthesizers. In normal use, an external 10 MHz reference *(typically a GPSDO)* is applied for precise frequency reference, overriding the less-accurate onboard oscillator, and an external GPS antenna allows the onboard receiver to supply timing and location information for the FST4W transmissions which, in a typical application, produced on a per-band basis by separate synthesizer and amplifier output. This device also has a USB "B" connector which permits access to the configuration files in the form of a mass storage device as well as access for the programming environment *(IDE)*.

Typically, this board is configured so that the *(up to)* six outputs produce digital transmissions - but in this case only output #1 is used to produce FST4W while two other outputs generate CW *(unmodulated)* signals for the "T1" and "T2" carriers. A special configuration was used to cause the FST4W frequency modulation to be spectrally-inverted to accommodate the high-side local oscillator injection. A copy of this configuration file and notes are included in the appendix of this document.

Here is the as-built diagram of the exciter:

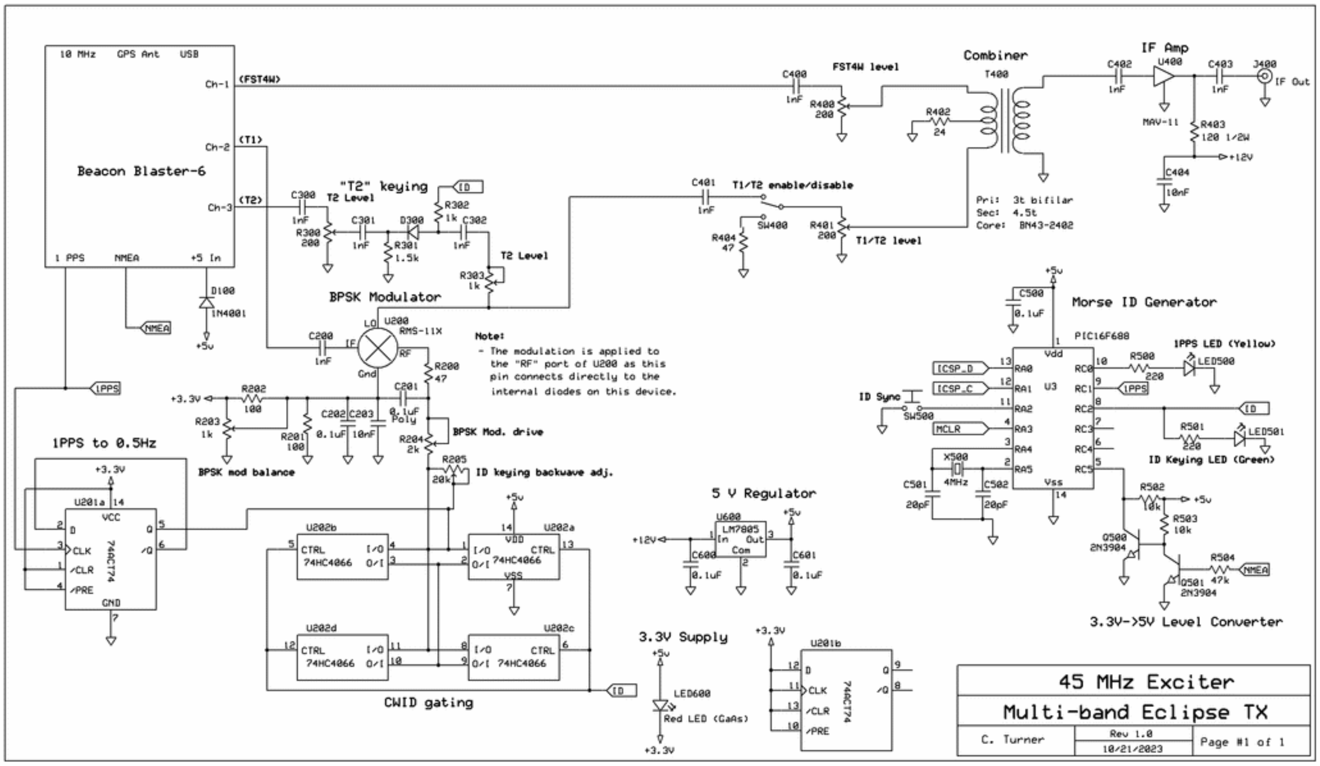

The BeaconBlaster is represented by the block in the upper-left corner and the signals from it are applied to additional circuitry. Using an oscilloscope, the 1pps and NMEA signals from the GPS receiver were identified and brought out to the external circuitry, noting that these signals use 3.3 volt logic levels.

The 1pps output is applied to U201a, half of a 74ACT74 wired as a divide-by-two flip-flop. This device is powered from a 3.3 volt supply created by dropping the 5 volt supply through an old-fashioned red LED with about 1.6-1.8 volt drop so that it can be properly driven by the 3.3 volt logic signal from the GPS receiver. Because U201a is triggered by the rising edge, its output changes state upon the beginning of the 1pps pulse. The other “half” of U201 is unused.

The divided-by-two *(0.5 Hz square wave)* signal from U201 is applied to the “RF” port of U200, a diode-ring doubly-balanced mixer *(DBM)*. Unlike most DBMs where the internal diodes appear on the “IF” port, the diode connections of the RMS-11X appear on the “RF” port - likely due to the industry convention of defining the combination of ports that provide the greatest LO to RF port isolation. Practically any diode ring mixer could be used here provided that the “baseband” modulation is applied to the mixer port with direct connection to the internal switching diodes.

Resistors R201 and R202 powered from the 3.3 volt supply - are used to establish a mid-supply reference for the DBM and it and its “ground” are thus-elevated at DC only. By doing this the 0-3.3 volt logic output from U200a will swing above and below the DBM’s “ground” and permit bi-phase modulation of the RF passing through it. Resistor R200 establishes a 50 ohm-ish impedance at the RF port and C201 provides RF bypassing and a small degree of pulse shaping. R204 sets the amount of current - and thus the amount of attenuation in the “on” state” - of the drive signal from U200a.

To be legal, the T1 and T2 signals must be identified and this is done by reducing the amplitude of these carriers for about 250 msec and then doing a Morse ID of my call at 20 words per minute - the elements of the characters being transmitted at “normal” power: After the ID, another 250msec pause at the low level is followed by a resumption of normal power level. This keying signal is generated by U3, a PIC16F688 microcontroller which can also “see” the 1pps signal from the GPS receiver and this is used to time an ID to occur every 600 seconds: The A/D converter in the PIC is used to detect the 3.3 volt logic level 1pps pulse rather than do a 3.3 to 5 volt level conversion as the timing of the Morse ID need not be precise to the microsecond. Q500/Q501 are used to translate the 3.3 volt logic from the NMEA output to 5 volts and applied to U3: At this time the software does not *(yet?)* use this information to time the ID cycle and the “ID Sync” is used to restart the 600 second timer and trigger an ID manually. Typically, I timed the beginning of the Morse ID to coincide with the end of the FST4W transmission *(e.g. at about 54 seconds on the “odd” minute.)*

Comment: It ***may*** have been possible to configure the Beacon Blaster to do a Morse ID, but since I wanted the carrier to be phase-continuous through the ID to prevent disruption of frequency measurements with very long sample times, I did the keying externally. It would be worth knowing of the BeaconBlaster can, in fact, key the carrier without disrupting the phase.

To reduce the carrier for the purposes of the ID the PIC’s “ID” line is brought low which turns off the gates of U202, a 74HC4066. This device - a quad analog switch - was used as it is agnostic to the different *(3.3 vs 5 volt)* logic levels: R205 is connected “across” this switch to allow a bit of current to flow when this switch is open to set the level of the “backwave”, allowing the “key up” carrier level to be set for the T1 signal - typically about 10 dB below the normal “key down” power level. All four gates of U202 are connected in parallel for lowest resistance.

For keying the T2 carrier for the ID, a simple diode switch *(D300)* is used. This diode is floated on the same mid supply rail as the diode ring mixer and when it is forward biased *(the “ID” line is high)*, it passes RF - but blocks it when it is reverse-biased *(e.g. the “ID” line low)*. Potentiometers R300 and R303 allow adjustment of the level of this carrier relative to that of the T1 carrier, thus providing on/off keying for the Morse ID.

The FST4W and the composite T1/T2 carrier source are combined via the hybrid combiner using T400 and related components which is preceded by R400 and R401 which are used to set the absolute levels of the FST4W and T1/T2 carriers. The output of T400 is amplified by approximately 12dB by U400, a MAV-11 MMIC and made available on J400 to be output to the 45 MHz filter/amplifier module. Switch SW400 allows selection as to whether the T1/T2 carriers should be transmitted or not.

Additional comments:

In retrospect, the use of a 45 MHz IF wasn’t really the ***best*** choice - and not for the reason that you are probably thinking. The possible issue is that the 35th harmonic of 45 MHz lands at 1575 MHz - very near the GPS “L1” frequency. Fortunately, this wasn’t really a problem because the boards were appropriately constructed and the circuits featuring proper bypassing and decoupling. In the field, the GPS “puck” antenna was placed about 6 feet *(3 meters)* away from the transmitter and no obvious issues were noted. The use of 45 MHz was dictated by the ready availability of inexpensive, relatively narrow quartz crystal filters.

Were I to build this again, I would likely do a few things differently - specifically to improve the flexibility of signal level adjustment of the FST4W and the T1 and T2 carriers. Since it worked exactly as it was intended, I have no immediate plans to rebuild it.

45 MHz amplification and filtering:

The composite 45 MHz signal from the box containing the Beacon Blaster goes to the box containing the mixers and synthesizers - but it is first passed through a PCB equipped with four poles of quartz filtering. The nominal bandwidth of this band-pass filter is around 8-10 kHz. This filter was used because of the presence of many other signals *(e.g. harmonics, ingress from the transmitted signals themselves)* could get into the IF signal path and cause a myriad of low-level spurious signals. This board is located in the same enclosure as the downconverter boards.

Circuit Description

This circuit is designed to strongly filter the IF prior to down-conversion to assure that it does not spurious contain frequency components - either from the BeaconBlaster or from ingress of transmit signals on the cable conveying the 45 MHz signal from it. This effort was done to reduce the probability of generating spurious signals in the transmit path.

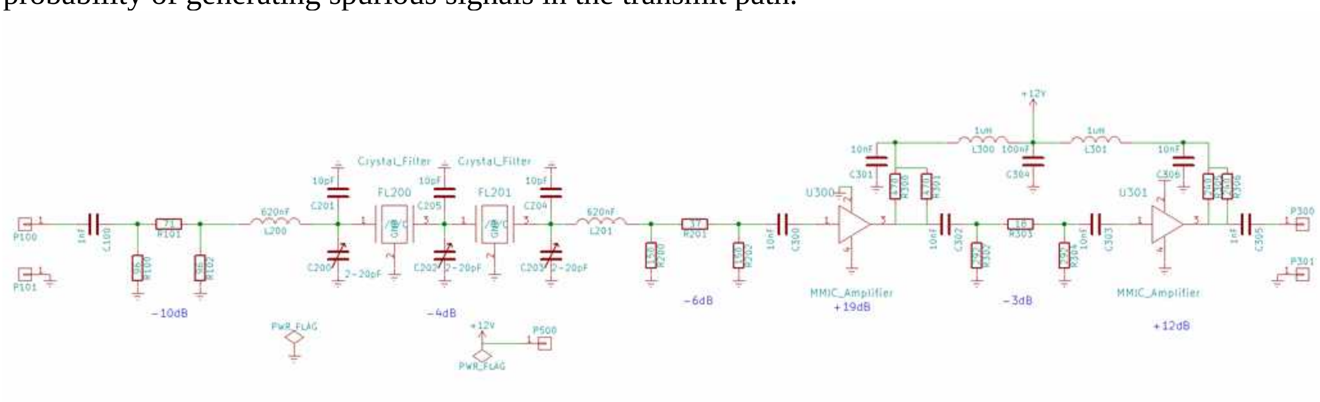

This filter uses an ECS-45K7.5A quartz filter, which comes as a pair of 2-pole, matched quartz filter units to provide a bandwidth of roughly 10 kHz. The input signal is matched/padded via resistors R100-R102 with L200 and C200/C201 to match to the 330 ohm impedance of the crystal filters. Between the two crystals, C202/C205 provide interstage matching while C203/C204/L201 match the output back to 50 ohms where it is a 6 dB resistive pad provides termination.

Adjustment of the crystal filter was done using a NanoVNA. Careful adjustment of the variable capacitors permitted a symmetrical passband response with a nice, flat top *(to better than 1dB)*.   When making such a measurement it is typically necessary to place use resistive attenuators on both the input and output of the VNA *(6-10 dB each)* so that the onboard amplification cannot overload it and cause erroneous readings.

The resulting signal is them amplified by a pair of MMIC amplifiers - the first being an MSA-2086 to overcome the losses of the crystal filter, matching and padding and then by an MAV-11 to boost the signal even further.

The ECS-45K7.5A filters are through-hole, but in retrospect it would have been just fine if I used a ***single*** surface-mount filter unit *(2-pole)* and of wider bandwidth - say 15-30 kHz: The only circuit change would be to rework the values of the L/C network used to match the crystal filter to 50 ohms. The quartz filtering was characterized and it was determined that the propagation delay between the incidence of the 1pps transition and the zero phase crossing was 82.20 microseconds.

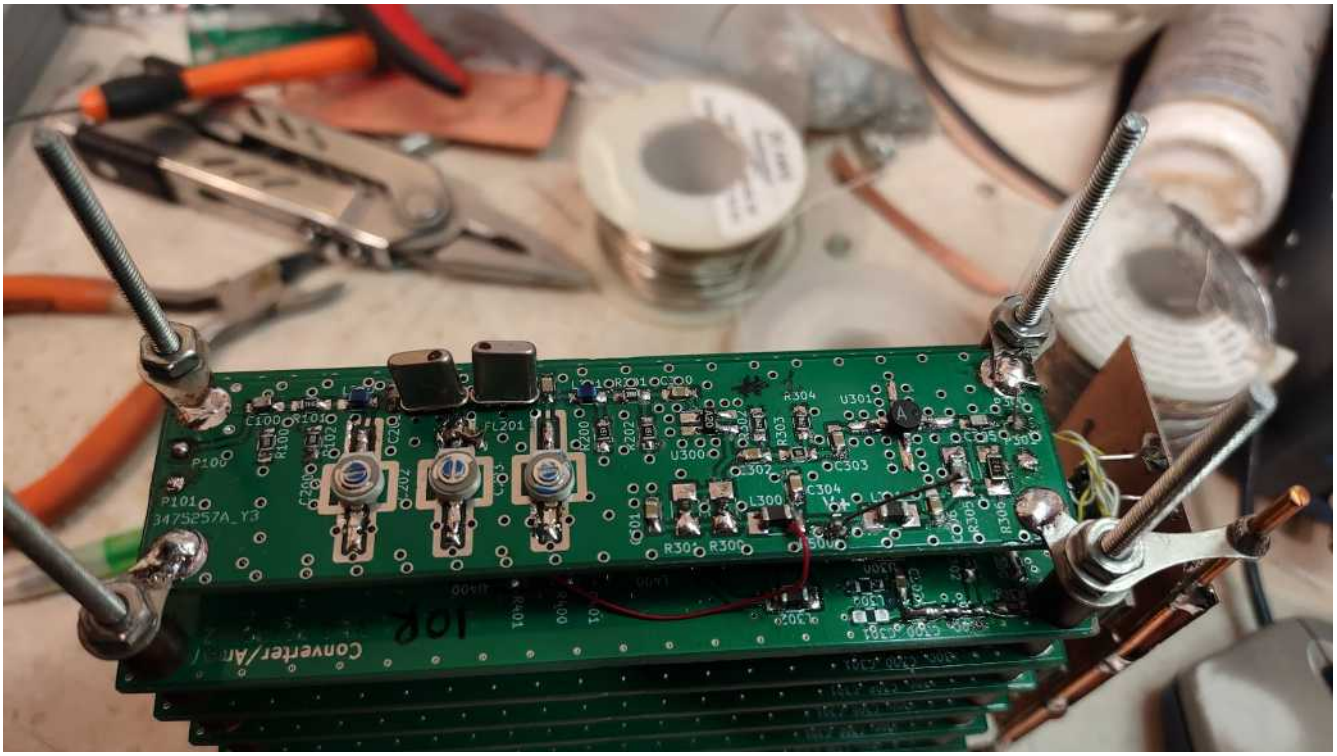

The 45 MHz filter/amplifier board. This is on the "R1/R2" generator, but identical to that in the TX converter chassis.

The output of the filter module is amplified by two MMICs, the total gain being about 25 dB. The output MMIC, a MAV-11, is driven to no more than about +5dBm to preserve linearity of the composite signal.

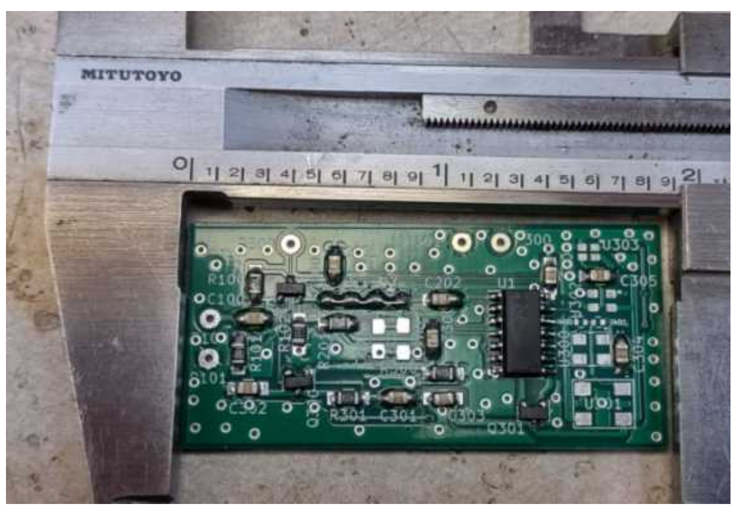

The "clock/switch" board: Four sets of pads accommodate different size oscillators.

Clock switch board:

I decided to standardize on 10 MHz as the single, precise frequency and this source would be used as a reference for all other oscillators from the BeaconBlaster to the downconverter module. Within the downconverter modules are Si5351-based frequency synthesizers, also locked via the 10 MHz source.

While intended to be driven by a precise, external 10 MHz frequency source, I also wished to allow the downconverter to operate without an external signal to permit testing, and for this I designed a “clock switch” board: If the external reference is absent, an onboard TCXO is powered on, but if such a reference is present, the onboard TCXO is turned off and the external signal is conditioned and passed through. This allows the board to ***always*** produce a signal at/near the desired frequency - useful for casual testing.

Circuit description:

The external clock signal is applied via P100 and buffered by Q100, an emitter-follower, which also establishes a mid-voltage bias supply for U1a. Resistor R103 was found to be necessary as the lack of termination provided by it ***could*** result in self-oscillation if no input source is connected - particularly if an unterminated coaxial cable was connected to the input: Since this switch board is frequency agnostic, ***any*** such oscillation would cause the onboard TCXO to be disabled.

Following the buffer are components *(C200, R200, C201, L200)* that may optionally be used to provide a low-Q band-pass filtering of the input signal - possibly useful if the waveform is particularly “ringy” or otherwise nasty. If this band-pass filter is ***not*** needed *(it usually is not)* then R200/C201/L200 are omitted - C200/R200 are replaced with jumpers - and the bias from Q100 gets passed through to U1a.

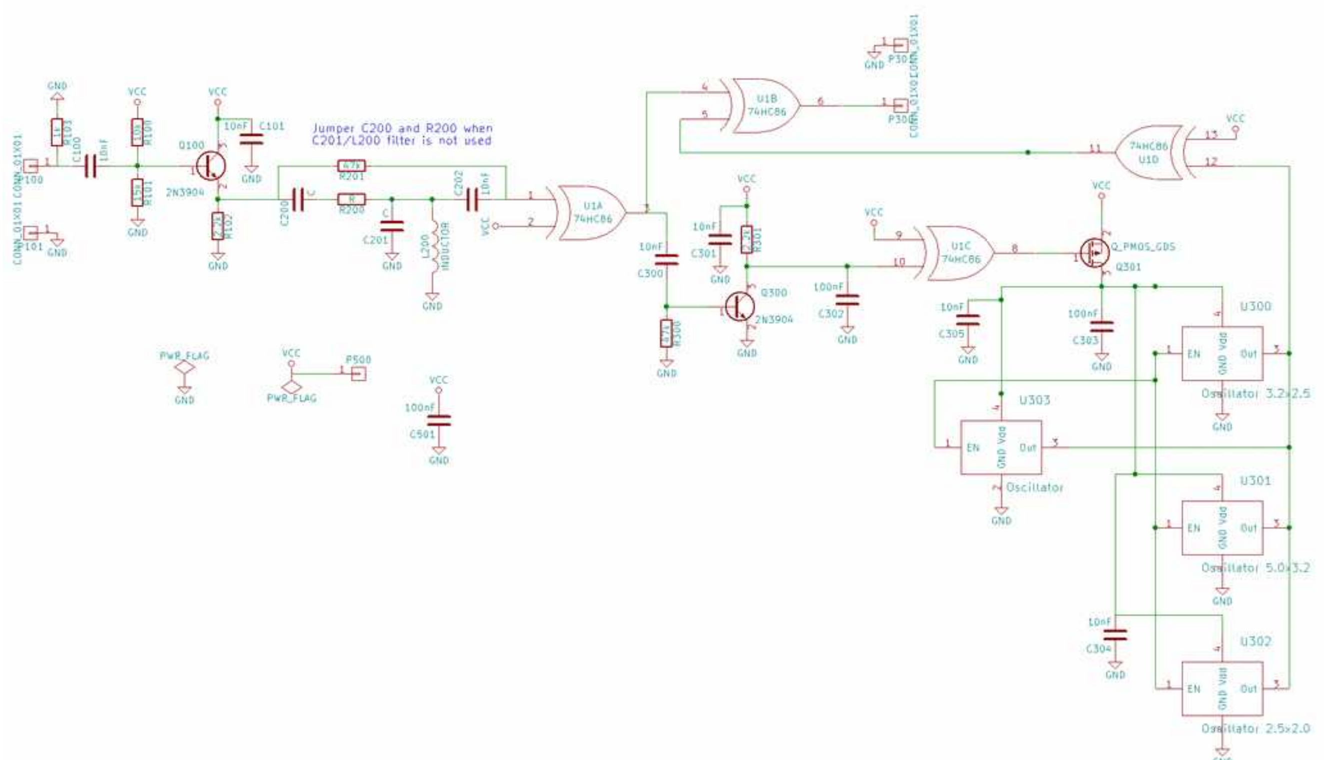

Clock switch unit. A 74HCS86 (Schmidt trigger version) is used for U1.

U1A, an XOR gate, which is preferably a ***74HCS86*** - a version with Schmidt-trigger inputs - squares up the signal and provides a peak-to-peak output commensurate with the supply voltage. The output of this feeds U1b which provides an output signal.

The output of U1A is also sent to Q300 via C300: If there is a signal on the input, Q300 is repeatedly turned on by that signal, pulling its collector - and one of the inputs of U1C low and held there by capacitor C302. U1C, wired as an inverter, thus has a “high” output in this condition and Q301 - a low-threshold P-channel FET - is held off and no power is applied to the onboard TCXO.

If there is ***no*** signal applied to the input, the Q300 is never turned on and its collector goes high, the output of U1C goes low, turning on Q301 and power is applied to the TCXO which is then buffered by U1D and then output by U1B. Because U1D is wired as an inverter, if the TCXO is powered down - which results in no voltage on U1D’s pin 12 - its output is high, allowing the external clock signal to be passed via U1B.

The diagram shows ***four*** TCXOs because the circuit board has four separate footprints, sized to accommodate the most common sizes of these devices.

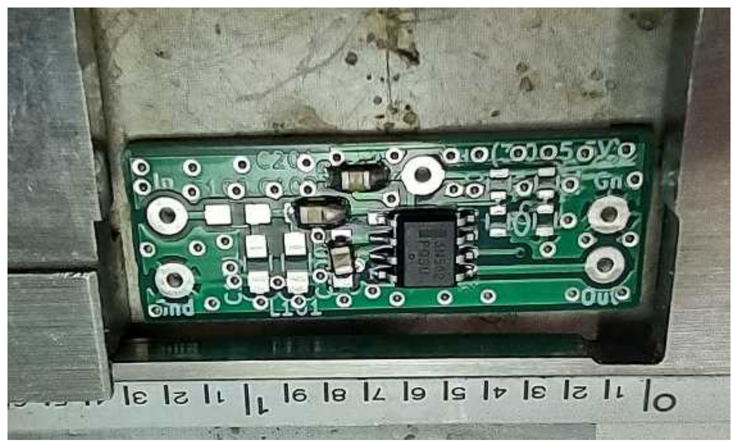

The clock multiplier board

In testing this circuit was found to be capable of operating down to at least 2 volts and was able to pass signals to at least 80 MHz at 3.0 volts - but it was only tested with 10 and 25 MHz TCXOs. It is nominally powered at 5 volts.

It has been suggested that squaring a sine wave to apply to logic devices - or using something like a Schmidt Trigger - is fraught with difficulties like increasing the phase noise, etc. While this may be true, such degradation was not observable with the test equipment on hand *(Tiny SA Ultra, Agilent E4440A, etc.)* Also, what is it that ionospheric propagation does with RF, again?

LO generation - 2.5x multiplier:

The 10 MHz signal is applied to a board that contains an NB3N502DG clock multiplier. Strapped for a multiplication factor of 2.5, this produces a 25 MHz signal from the 10 MHz source for the ProgRock 2 synthesizers. This is based on the NB3N502DG clock multiplier chip and is nearly a direct copy of the data sheet.

The only departure is the inclusion of components L100, C101 and L101 which - along with the appropriate selection of C100 - can form a simple 4-pole band-pass or low-pass filter, depending on which components are fitted: This filter might be useful if the inputted signal is “ringy”. If this filter is not needed, L100 is replaced with a jumper an C101 and L101 are not fitted - which is how I built it for this project.

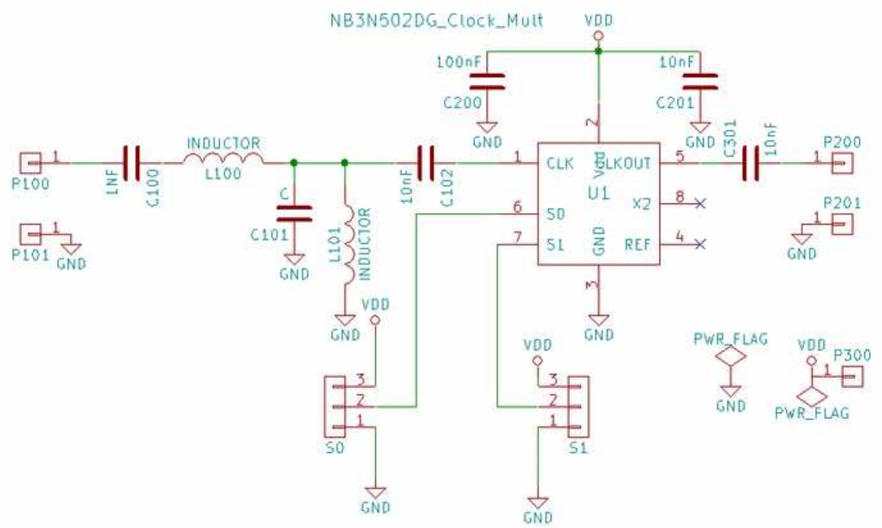

The NB3N502DG is capable of a number of different multiplication factors, selected by jumpers “S0” and “S1”. In the table below, “0” means grounded, “1” means connected to Vcc and “M” means floating:

S0	S1	Multiplier
0 (or M)	0	2x
1	0	5x
0 (or M)	M	3x
1	M	3.33x
0 (or M)	1	4x
1	1	2.5x

For the Eclipse equipment, I needed to produce 25 MHz from 10 MHz, so both S0 and S1 were jumpered to Vcc.

Comment:

In testing this board, I have also determined that it could be used to provide clocking for a KiwiSDR - and even a Rasperry/Flydog SDR from a 10 MHz source as follows:

- Use one multiplier board in “2x” mode to produce 20 MHz. *(Alternatively, a Bodnar GPSDO may be able to produce both 10 and 20 MHz).*

- Use another multiplier board in “3.33x” mode to produce 66.66 MHz to feed to the KiwiSDR(s).
- If a clock is needed for a Raspberry or Flydog SDR, yet another board could be used in “2x” mode to produce 133.32 MHz. An input frequency of 66.66 MHz is beyond the specifications of the NB3N502DG, but in 2x mode I found that it would reliably work with an input signal beyond 80 MHz *(e.g. 160 MHz output).*

This circuit seems to work well between 3.0 and 5.5 volts: The lower voltage limit of 3.0 volts should be carefully observed as it definitely does ***not*** work well below this - and the output drive seems to drop more than you might expect, possibly being marginal below 3.5 volts for applications requiring a full voltage swing.

LO generation - Si5351 modules:

Because of limited time and, to an extent, laziness, I used three QRP Labs “ProgRock 2” boards *(link: **http://qrp-labs.com/progrock2.html**)* to generate the needed eight local oscillator signals. These boards - costing $18/each at the time of writing - are very small *(they will fit inside an HC-6 crystal can)* and can be directly programmed via USB.

The ProgRock 2, front and back

While specified to be able to operate at up to 12.0 volts, I “discovered” that doing so would cause the regulators to overheat and destroy them in a matter of minutes, so they were, instead, operated from a local 5 volt regulated source.

As the ProgRock 2 boards have an onboard 0.25ppm, 25 MHz TCXO, these had to be removed to allow the desired precision of an external clock. This was *(carefully!)* done using hot-air rework gear and the 25 MHz clock from the 2.5x multiplier was applied to the “output” pad of the original TCXO, via a blocking capacitor already on the ProgRock 2 board. The 25 MHz clock signals were applied to all three ProgRock 2 units in parallel using twisted wire-wrap wire.

Additionally, the configuration of the ProgRock 2 does not allow one to directly specify the input source frequency in the calculations. Practically speaking, I could have probably applied 10 MHz to the ProgRocks and applied a 2.5x multiplier to the frequency, instead, but I chose the “safe” path and used 25 MHz: Testing the “10 MHz” configuration is on my “to do” list.

As the ProgRock 2 boards' Si5351 operate at 3.3 volts, I placed a 100 ohm resistor in series with the output of the 2.5x multiplier board *(which was being operated at 5 volts)*. Originally the multiplier was operated at 3.3 volts, but I discovered that the '5351's sometimes behaved erratically, depending on temperature: Even though the clock signal being applied to them looked OK, it took boosting the swing of the signal by powering the multiplier board from 5 volts *(with the series resistor to limit current)* to make everything unconditionally stable.

The major disadvantage of using the ProgRock 2 in this application is that the divisor polynomial applied to the Si5351 is not known and this means that ***exact*** calculation of the output frequency cannot be known. This fact required that I make individual measurements of the component frequencies to determine - with mHz precision - the actual, on-air frequencies.

The three Progrock 2 Si5351 modules mounted to the stack of converter boards. Just barely visible is the clock multiplier board, seen edge-on on the top side of the board. The "clock switch" board - to the right of the multiplier - is not visible.

This measurement involved the use of a Schlumberger Si4031 - which produces ***exact*** ratiometric output frequencies relative to its *(10 MHz)* reference to a resolution of 100 Hz. By mixing the signal to be measured with the output of the '4031 set to a slight offset *(usually a few hundred Hz)* the resulting audio frequency can be measured with precision using the "Spectrum Lab" program by DL2YHF. With this - and enough time - it is possible to make sub-mHz frequency measurements. It's worth noting that this technique also requires that one measure the sample rate of the sound card in the computer with the same precision and take its error into account as well.

Frequency conversion

Custom boards were designed and built to convert the 45 MHz IF signal to the desired HF frequency. The circuitry is rather straightforward - but care must be taken to preserve linearity and signal cleanliness.

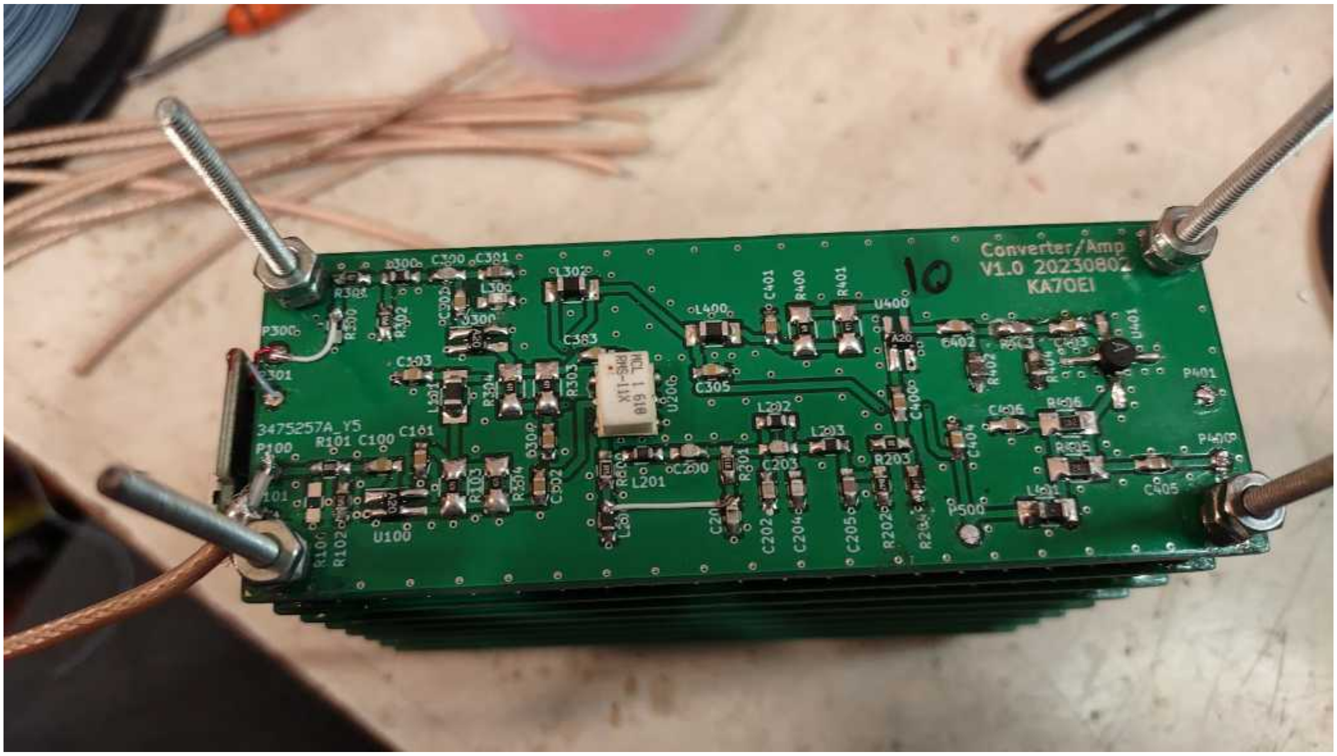

TX Converter module - this one for 10 meters.

The LO signal is applied and roughly filtered with a low-Q L/C network *(mainly to reduce the possibility of non-harmonic energy ingress)* before being amplified by a MMIC to approximately +8dBm. This signal is then fed to an RMS-11X doubly-balanced mixer.

Meanwhile, the 45 MHz signal from the filter network is applied via a resistive divider/attenuator with a 200 ohm input and nominal 50 ohm output: The 200 ohm inputs of the eight converter boards are simply connected in parallel, presenting a 25 ohm load to the output MMIC on the filter board. The output of the divider/attenuator is applied to a MMIC and amplified by about 16 dB and then sent to the mixer.

The output of the mixer is terminated with an L/C/R diplexer that passes the desired frequency *(the HF frequency to which signals are being converted)* but terminates into 50 ohms all other frequencies. From here, a low-pass filter strips off harmonics and local oscillator leakage before being amplified again by about 16 dB.

After this, there is either an attenuator *(for the lower bands)* or a high-pass filter: It was "discovered" that on some of the higher bands *(>=15 meters)* some spurious signals *(about -40dBc)* were appearing - but the application of a 3-pole high-pass filter *(L-C-L)* knocked those down by 10s of dB. Practically speaking, the transmit filter would have removed these, anyway.

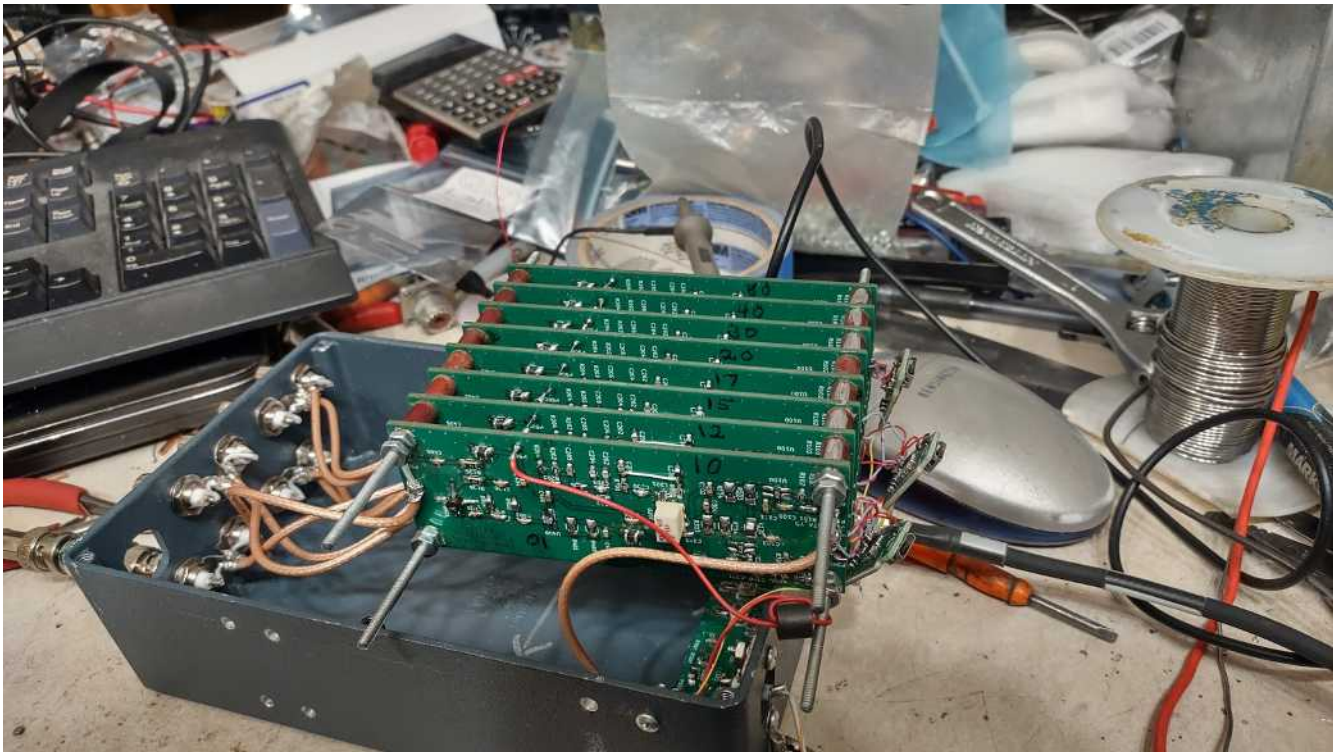

Stack-up of TX converters, with Progrock 2 modules hanging off the side for testing

The final amplification is via a MAV-11 MMIC which can cleanly *(with IMD components <30dB)* produce more than +10dBm - more than enough to drive the output RF amplifiers.

Comments:

- The stack-up of eight converter boards and 45 MHz filter/amp board draws about 2 amps at 13 volts, requiring a fan to keep everything cool: That's what I get for using MMICs everywhere!
- The lower-level MMICs are equivalent to the Avantek MSA-2086 - but in a slightly different package. These devices are no longer manufactured, but I used them because I have *(or had!)* a rather large quantity of them.
- The MAV-11 *(a.k.a. MSA-1105)* were used, also because I have a lot of them. One advantage of this device and the '2086 *(as opposed to more modern devices like the Mini-Circuits ERA series and the GALI-74)* is that they start to run out of steam north of 1-2 GHz, making it a bit easier to maintain unconditional stability when they are surrounded by reactive components *(e.g. filters)* - plus, they were free! *(Did I mention that I have a lot of them?)*
- The eight boards were mounted on 6-32 threaded rod with spacers to permit a compact arrangement. The ground, input signal and V+ connections are "daisy-chained" as these connections on the individual boards are inline with each other with this board arrangement.

Detailed circuit description

The frequency converter takes the 45 MHz IF signal *(or about any frequency, really)* and converts it - using an external local oscillator - to the desired amateur band.   There are ***two*** versions of this circuit: The first is built as depicted below, capable of about +10dBm output - enough to drive the QRP-Labs 10 watt amplifier - and another version that is designed to produce low-level signals that are injected into the receive signal path which will be described later.

The use of a high IF frequency - and “high-side” local oscillators - has the distinct advantage of offering “cleaner” frequency conversion: Using the ***difference*** of the local oscillator signal and the 45 MHz IF *(calculated using the sum of the 45 MHz and the desired HF frequency)* the other mixing product - the ***sum*** of the two - is forced to occur at a much higher frequency than HF, the worst case being 80 meters where this image lands around 93 MHz, making it very easy to filter out - and even easier for higher bands.

This same technique has long been used in double and triple-conversion superhetrodyne receivers and transmitters - and for the very same reasons.

This board does ***not*** include the local oscillator source itself - this being supplied externally by *(in my case)* Si5351-based QRP-Labs “ProgRock 2” synthesizer modules.

The diagram of this module is as follows:

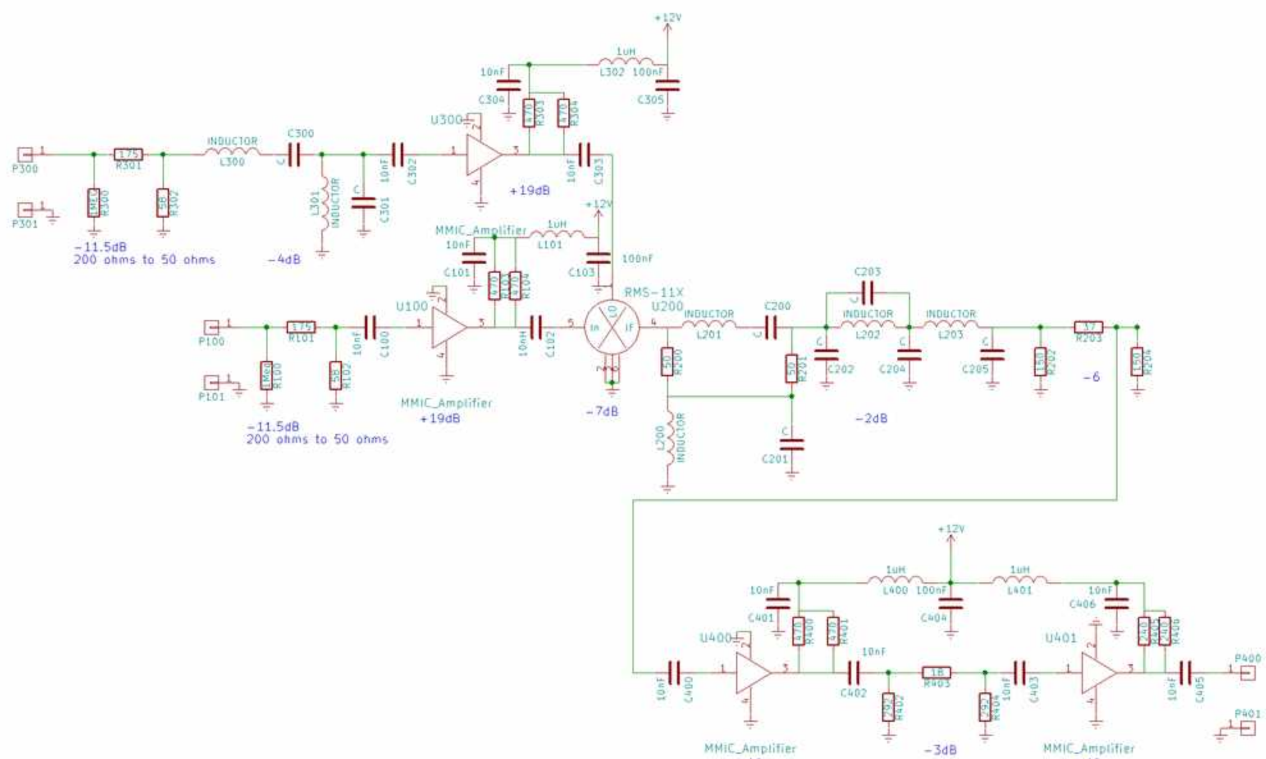

The local oscillator signal is applied via P300 through a 200 to 50 ohm resistive pad *(attenuation of 11.5 dB).* 200 ohms was chosen as this lighter loading *(compared to 50 ohms)* offers somewhat better performance - and less stress on the ‘5351. Following this is L300, C300, L301 and C301 which form a broad band-pass filter. The reason for this filter is to minimize the ingress of non-harmonic signals from the outputs of the other converter modules - or even via the power supply leads. Because this filter is quite broad, the same component values are used for the LO of 80-20 meters and a different set is used for the LO of 17-10 meters.

U300, an MSA-2086 MMIC boosts the output from the Si5351 and happens to have a saturated output level of about the right level to optimally drive U200, the RMS-11X diode ring mixer. It is unimportant that the output of the MMIC is driven into distortion *(e.g. more of a square wave than a sine)* as this is actually preferable for optimal performance *(provided that it has only harmonically-related content)* - a fact mentioned in the manufacturer’s own literature. *(See MCL document AN00-011 - **link**.)*

Meanwhile, the 45 MHz IF signal is applied via P100 and it, too, has a 200 to 50 ohm resistive divider - but this time it’s to allow direct paralleling of all eight converter modules across the output of the 45 MHz filter module, presenting to it a 25-ish ohm impedance: The MMIC really doesn’t care! U100 boosts the level of the IF signal - the optimal level being somewhere around -10 and 0 dBm to get the maximum output level from the mixer with an acceptable level of intermodulation distortion.

On the output of the mixer we find components L200, L201, C200, C201, R200 and R201: L201 and C200 form a low-Q band-pass filter which permits the desired frequency to go through while L200 and C201, connected in parallel, are “invisible” at the desired output frequency, but with R200 and R201 cause all other frequencies to be terminated at 50 ohms. This “diplexer” is very important in optimal performance of a diode ring mixer as it minimizes intermodulation distortion by preventing reflections of these “other” frequencies band into the mixer where they can *(ahem)* “re-mix”.

Immediately following this is a low-pass filter consisting of C202, C203, C204, C205, L202 and L203. For lower frequencies *(below 15 meters)* this is configured as a Chebychev low-pass filter *(omitting C203)* to remove local oscillator leakage and the images, but on 15 meters and above a Cauer configuration is used with L202 and C203 forming a notch at the 45 MHz IF and increasing the steepness. Following this another resistive pad: In retrospect I should have probably put the pad between the diplexer and the low-pass filter, but that’s probably bordering on OCD!

Again, we used an MSA-2086 MMIC to boost the signal and following this is another pad - or not: Again, a bit of OCD happened here and I typically jumpered the position of R403 and omitted R402 and R404 - but in the case of the 15 meter and higher converters, R402 and R404 were fitted with inductors and R403 with a capacitor to form a 3-pole high-pass filter to knock down low-level mixing products *(typically at 5 MHz and below)* even farther than they already were.

The final amplification *(U401)* is our friend the MAV-11, capable of a P1dB power level of +17 dBm, but perfectly happy to produce clean *(IMD <30dB)* output at even +13dBm - more than enough to drive the QRP-Labs 10 watt amplifiers.

Special attention was given to power supply decoupling of the various sections: 1uH series inductors and bypass capacitors are used on each MMIC amplifier’s supply line to minimize the amount of RF that might appear on the power supply bus. Failure to do this would surely result in spurious signals and possible amplifier instability!

This board is intended to operate from 11 to 15 volts - more gain and power output being possible at the higher end of this range.

The following is a table of various components used on a per-band basis. Standard values of inductors and capacitors were used, requiring a slight amount of "fudging" - but since the "Q" of these circuits is quite low, they are very forgiving in terms of component tolerance.

Band	**L200**	**L201**	**C200**	**C201**	**L202**	**L203**	**C202**	**C203**	**C204**	**C205**	**L300**	**L301**	**C300**	**C301**	**R402/4***	**R403***
80	3.9	1.2	510	1500	2.7	2.7	910	N/A	1500	910	330	82	25	100	N/A	N/A
40	2.0	620	250	820	1.4	1.4	470	N/A	750	470	330	82	25	100	N/A	N/A
30	1.5	470	180	560	1.0	1.0	333	N/A	560	330	330	82	25	100	N/A	N/A
20	1.0	470	120	390	680	680	240	N/A	430	240	330	82	25	100	470nH	200pF
17	820	240	100	330	560	560	200	N/A	330	200	330	68	22	100	360nH	150pF
15	680	220	82	270	330	470	75	33	200	100	330	68	22	100	330nH	120pF
12	560	180	72	220	330	470	82	33	180	100	330	68	22	100	270nH	120pF
10	470	180	68	200	300	430	82	33	180	120	330	68	22	100	240nH	100pF

All capacitor values are pF. Inductor values <=10 are in uH while those >10 are in nH.
Capacitor/inductor values at R402-R404 comprise a high-pass network.

Receive signal path version:

As mentioned above there's a version of this same circuit that is used to generate much lower-level signals that are injected into the receive signal path to generate the so-called "R1" and "R2" carriers. This circuit's operation and configuration is described in Appendix C.

This version uses the same circuit board as the transmit version, but a number of amplifier sections are omitted and jumpered as much lower signal levels are required. Specifically, these changes are:

- **Removal of IF amplifier.** U100, C100, R103, R103 and R104 are omitted and a jumper is placed between R102 and C102.
- **Removal of post-mixer amplification.** C400-C405, R400-R406, U400, U401, L400 and L401 are omitted and a jumper wire is placed between R204 and P400/C405.

Power amplifiers

Rather than design my own RF power amplifier, I chose, instead, to buy nine *(eight, plus a spare)* of the "10 watt linear amplifiers" from QRP-Labs - See: **http://qrp-labs.com/linear.html** These amplifier kits are relatively cheap *(they are now $26/each)* and they have nominally 26 dB of gain across the 80-10 meter frequency range meaning that +10dBm of drive would theoretically yield about 4 watts - very close to the nominal target of 5 watts of RF output per band.

After building one or two, it took me about 75 minutes to assemble and test each amplifier and barring an assembly error, they all worked as they should While the gain is specified to be "26 dB +/- 1dB" I suspect that my samples exceeded this slightly - but since a 2 dB variation implies nearly 40% variance, I wasn't too surprised to see that my nominally "+10dBm" output from the converter modules yielded about 18 watts on 80 and 40 meters, dropping to about 2 watts on 12 and 10 meters.

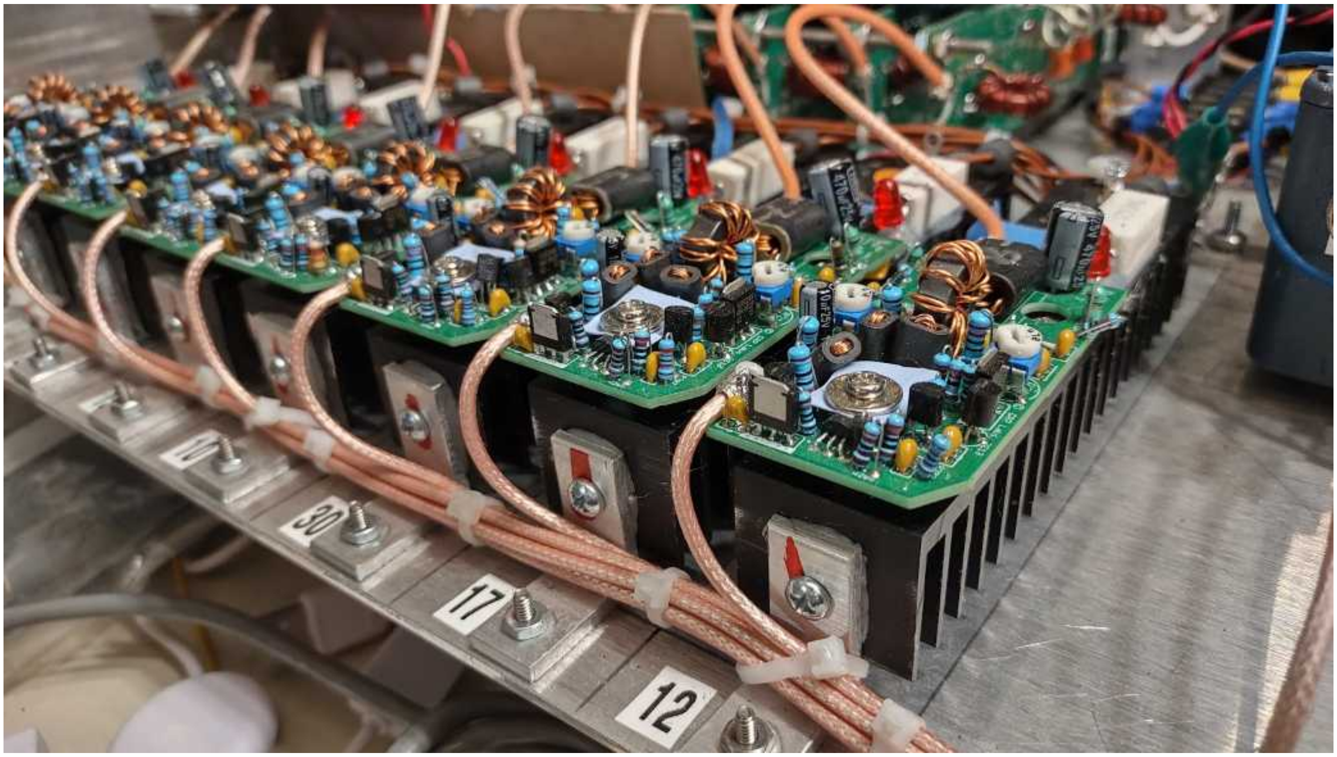

The QRP-Labs 10 watt linear amplifiers. Below the PC boards can be seen the method of mounting these amplifiers to the chassis, described in detail below.

While the gain is not directly adjustable, by adjusting the bias applied to the final transistors of these amplifiers *(IRF510)* the effective gain could be adjusted by about 6 dB without any obvious effect on linearity: For the lower bands I set the bias no lower than 50mA per PA transistor *(100mA total)* and for 12 and 10 meters I set the bias to 200mA per device - with the manual-recommended 125mA/device for the bands in between. Ultimately, I observed about 13 watts for 80 and 40 and about 4 watts on 12 and 10 meters.

Because these amplifiers are very "compact", they can be a bit awkward to mount - and there is no obvious way to do so. To mount these amplifiers I cut 1/2" wide tabs from some angle aluminum and drilled holes for mounting screws as can be seen. In the heat sink of the amplifier itself - on each end - I drilled a hole and tapped it with 8-32 threads to receive a mounting screw and then, when assembling, I put some "JB Weld" epoxy between the tab and the heat sink and tightened the screw firmly to prevent the heat sink from pivoting on the screws' axes.

I would recommend to anyone building these amplifiers to burn them in for several days before "trusting" them. During "burn in" I had two random component failures:

- One failure was that of one of the monolithic bypass capacitors across the main V+ rail *(C212)*. This amplifier was fine during initial testing, but it tripped the power supply overcurrent *(despite my inclusion of a 5 amp PTC fuse)* when I did first testing with all of them mounted and integrated. Ultimately I used a 5 amp power supply *(set to about 3 volts)* and a FLIR thermal camera to figure out exactly which component failed as the short was too low resistance to measure and locate with a standard *(non-Kelvin)* ohmmeter.
- Another failure was a random base-to-collector failure of Q201 and/or shorting of Q202. This resulted in an uncontrolled bias voltage on the driver transistors, causing them to experience sudden, rapid self-disassembly: I had to replace Q201-Q204 to restore operation.

As mentioned before, I included 5 amp self-resetting PTC fuses to prevent a catastrophic failure of one amplifier *(e.g. short of output transistor or bypass capacitor)* from causing the wire to immolate - something that could cause a cascade failure of a wire bundle: These PTC fuses - tucked between th board and heat sink - also would open if the heat sink got extremely hot *(>80C or so, depending on ambient temperature)* potentially protecting the amplifier from thermal destruction. Across each of these PTC fuses is wired a red LED and series 1.5k resistor so that should a fuse "blow" due to overcurrent - or from an amplifier heat sink getting too hot - it would glow.

Also included each amplifier is a 0.1 ohm, 5 watt resistor *(the white, rectangular things in the background)* that are used to monitor the current of each, individual amplifier - particularly useful for in-situ adjustment of the bias of the output transistors. Typically this was done by removing RF drive *(interrupting the 45 MHz IF or unplugging the amplifier's input cable from the downconverter assembly)* and using the lowest voltage scale on the voltmeter.

Finally, I noted that there is no DC continuity between the heat sink and the circuit ground, which is concerning as tabs of the final transistors are bolted to it using gray thermal pads as insulators, resulting in a few 10s of PF of coupling, potentially making the heat sink slightly "live" with harmonic-rich RF. Using an extra hole in the supplied heat sink, I installed a lug on the heat sink and soldered it edge of the circuit board near the output transformer to provide an RF ground.

Transmit antenna filter/combiner

To allow multiple transmitters to exist on the same feedline I constructed a relatively high-Q band-pass filter network for each transmitter. I had designed/used similar filters for the individual bands on the Northern Utah WebSDR *(and built a similar device for KFS)* - but in that case I could trade selectivity for insertion loss, allowing the filtering/separation of all bands 160-10 on a single signal path.

For transmitting - particularly when the goal is to achieve 1dB or less insertion loss - the challenge is greater. For most of the non-WARC HF bands, this is pretty easy: 80, 40, 20 and 10 meters are an octave apart and a filter sufficient to obtain the necessary isolation is not all that difficult. The inclusion of 15 meters into the mix makes it a bit more difficult - but appropriately sizing the inductors and choosing low-loss *(typically silver-mica)* capacitors makes it possible.

Throwing the WARC bands into the mix is where things get difficult: 30 meters isn't too much different than the separation of 15 meters from 20 or 10 meters - but 17 meters and, especially 12 meters, pose a challenge.

To an extent, this challenge is averted by the fact that a typical wire antenna *(an end-fed half-wave or a full-wave loop)* will happily match to something resembling 50 ohms on 80, 40, 20, 15 and 10 meters, but not so much on any of the WARC bands which means that it would make little sense to combine those into the mix but, rather, feed them separately.

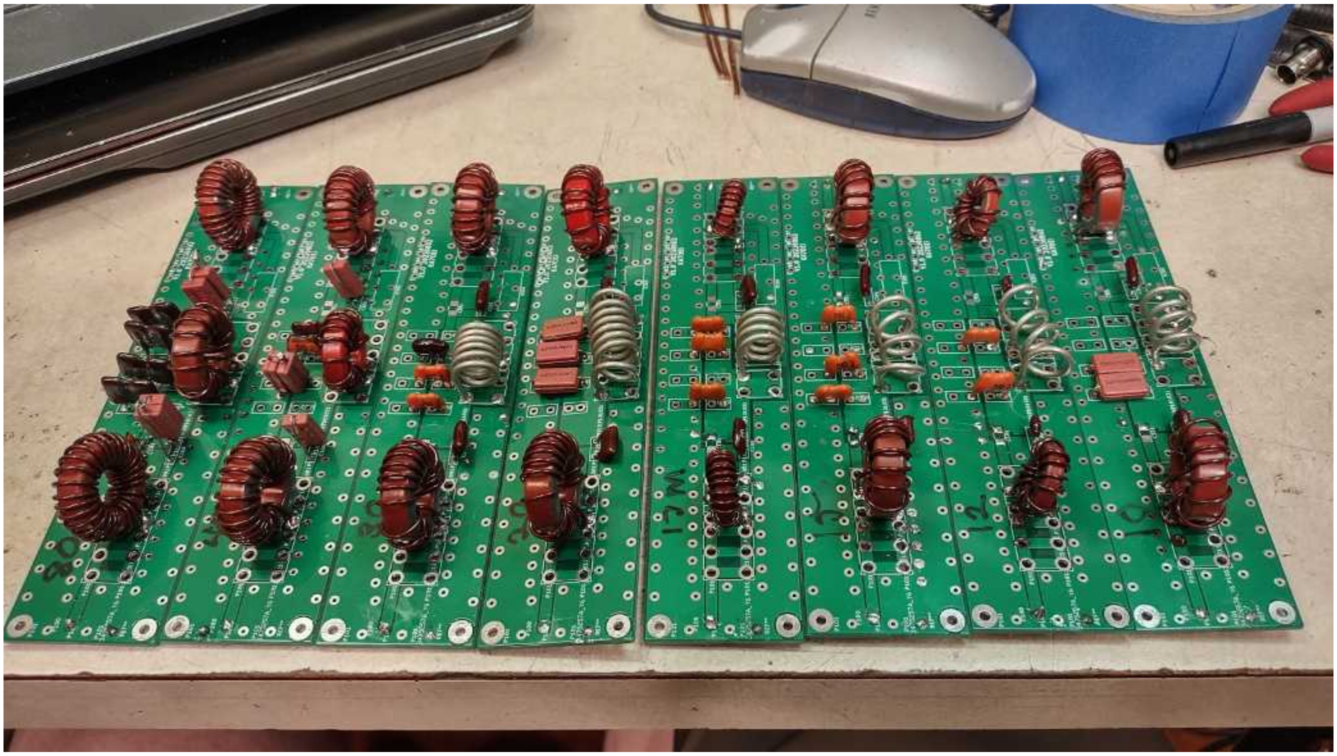

The individual transmit/combiner filter modules

These band-pass filters consist of “series-parallel-series” L/C networks, which provide a band-pass response narrow enough to suppress harmonics down to FCC requirements and sufficient isolation between transmitters in the “non-WARC” case.

Ultimately, the transmitter was equipped with four antenna outputs: One for a combined 80, 40, 20, 15 and 10 meters, one for 30 meters, another for 17 meters and yet another for 12 meters. In testing I found that I could combine the “80/40/20/15/10” output with the 30 meter output with little degradation and even adding the 17 meter output was tolerable - but as expected adding the 12 meter output to the mix significantly degraded both 15 and 10 meters.

During testing at home I ran it for several days with all outputs combined - noting negligible production of IMD products between the outputs - and fed it to my 10-band vertical antenna. Despite the fact that the VSWR on the 80 meter frequency was about 10:1 *(the vertical is tuned to the upper end of the 80 meters)* I received reports on all bands - and the amplifiers seemed to tolerate the abuse just fine.

Circuit description

This filter provides harmonic filtering of the RF amplifiers as well as allowing multiplexing of several bands on a single antenna, the diagram being as follows:

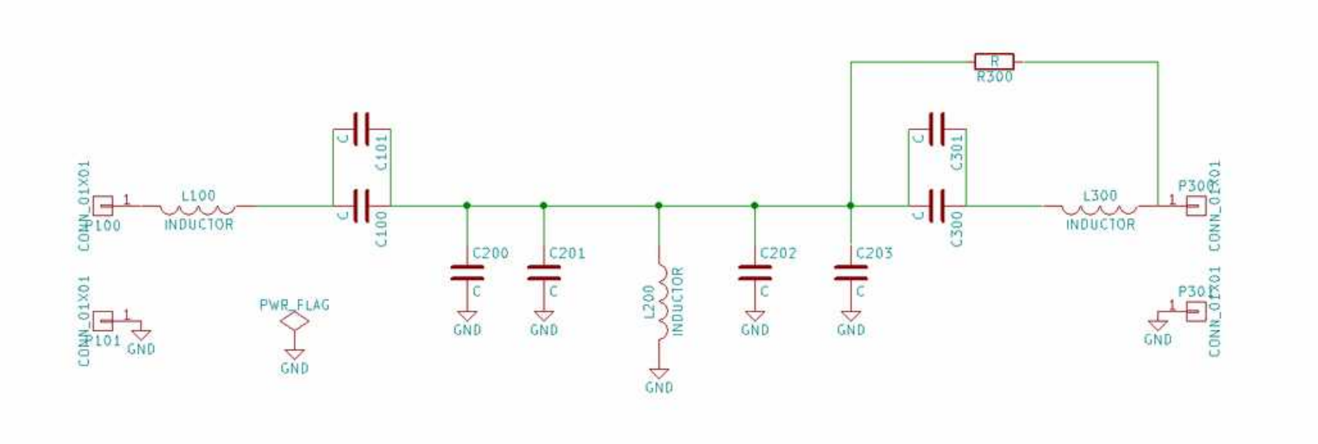

This filter is symmetrical - that is, the L/C values for C100/C101/L100 are the same as those for C300/C301/L300 - except for R300. This filter's termination impedance is low only near the passband, but extremely high everywhere else, and if it had been found that the amplifier was unstable due to this termination, this resistor - the value to be determined - would have been fitted to stabilize the amplifier and a commensurate decrease in filter performance would have been tolerated. As it turned out, the QRP-Labs 10 watt amplifiers seem not to care about their termination impedance, so it was not needed.

For the constructed unit, all capacitors are silver-mica with precise trimming of C100/C101 and/or C300/C301 being done with silver-mica or NP0/C0G. L100 and L300 were wound on rather large Mix-2 toroids - T80 and T94 - using #17 wire *(because it was on-hand)*. For the 80 and 40 meters, L200 was wound using the same toroids, but for 30 meters and higher, air-core coils of 12 AWG tin-plated wire was used. Tuning was performed using a NanoVNA. At first, tuning is quite tricky as all adjustments interact violently, but one soon gets the "feel": It's recommended that the lowest-frequency filters be adjusted first as they are a bit easier to tune. Approximate component values for these filters are as follows:

Band	L100/300	C100/C300	L200	C200
160	11uH	717pF	3.8uH	2228pF
80	5.4uH *25t T94-2*	365pF *2x165pf+33pF*	1.8uH *14t T94-2*	1065pF *4x240pF*
40	2.8uH *17t T94-2*	170pF *1x165pF*	836nH *11t T80-2*	568pF *2x240pF+82pF*
30	2.7uH *16t T94-2*	93pF *100pF actual*	208nH *6t #12 0.5" ID, 0.75" L*	1217pF *2x240pF+750pF*
20	2.4uH *15t T94-2*	52pF *47pF actual*	257nH *8t #12 0.5" ID, 1.25" L*	497pF *3x165pF*
17	1.64uH *14T T80-2*	47pF *47pF actual*	106nH *4t #12 0.5" ID, 0.6" L*	730pF *2x 240pF*
15	1.64uH *10T T94-2*	33pF *33pF actual*	83nH *4t #12 0.5" ID, 1.0" L*	720pF *3x240pF*

12	1.1uH *11t T80-2*	37pF *39pF actual*	84nH *4t #12 0.5” ID, 1.0” L*	478pF *2x240pF*
10	1.62uH *10t T80-2*	20pF *20pF actual*	98nH *4T #12 0.5” ID, 0.7” L*	330pF *2x165pF*

Note: *Inductor specifications are approximate: Physical adjustment is required during filter alignment to fine-tune inductance values. For the higher bands, the board traces may be a significant portion of L200's inductance values. The values for 160 meters were calculated and simulated, but the filter was not constructed.*

The schematic shows multiple, parallel capacitors and only the lowest-number capacitor designator is depicted in the table above. In particular, for the parallel L/C network consisting of L200 and C200, there are actually four positions for capacitors *(C200-C203)* and the intent is to install ***multiple*** lower-value capacitors at this position to minimize losses of the capacitors in addition to being able to “dial in” a precise value.

As mentioned earlier, the above filters are sufficient to allow the combining of 80, 40, 20, 15 and 10 meters on the same antenna with minimal interaction. The original intent was to have 30, 17 and 12 meters each on their own antennas - which could have been simple dipoles or portable vertical or loaded dipole antennas, but it turned out to be convenient to connect 30 and 12 meters together *(as the frequency spacing permitted such)* and feed the combined signal into an 80-10 meter end-fed halfwave: Despite the lack of a particularly good match at either band *(likely around 3:1-4:1)*, the durability of the amplifiers and the relatively low loss of the coaxial cable resulted in good performance at both bands. The remaining band - 17 meters - was radiated via a portable dipole, a “JPC-7” with its feedpoint set on a tripod at a height of 10 feet *(about 3 meters)* - also with excellent results.

As with the converter boards, the filter boards were “stacked” using lengths of threaded rod and lots of nuts and split washers. The same right-angle aluminum pieces used to mount the amplifiers were used to mount the four corners of this assembly to the chassis.

Final integration:

The eight power amplifiers, the stack of transmit combiner/filters and the converter board were assembled onto a large aluminum chassis *(it used to be part of a rack shelf, I think)* - the final result pictured above. As the 45 MHz exciter *(the part with the Beacon Blaster)* was considered to be a separate unit *(plus the lack of space)* it is kept separate as it considered to be the component that is most likely to be modified.

In the foreground on the far left is the converter chassis: Because it, by itself, consumes about 2 amps, it requires its own fan for cooling. To the right of that is a terminal strip where power is distributed to the various pieces - particularly the individual amplifiers, and right of that is the stack-up of the individual band/combiner filters: These are bolted together using 6-32 “all-thread”, nuts and lock washers, using right-angle tabs to attach to the chassis.

Also along the near edge - just left of center and on the far right - are the individual band output: BNC “tee” connectors can be used to combine the various bands together for testing or operation as needed - or not!

Along the back edge can be seen the eight QRP-Labs 10 watt amplifiers. As the 80 and 40 meter amplifiers run with the most power, a fan was installed above them to move a little bit of air past the heat sinks. Those with sharp eyes might notice that there is an air dam between the filter bank and the amplifier cut from a box of crackers to better-direct air between the chassis and heat sinks.

The complete multi-band chassis - sans the 45 MHz exciter. In the background are the eight QRP-Labs 10 watt amplifiers: A ninth amplifier was built and tested as a ready spare.

Practically speaking this fan wasn't strictly necessary - but it's a good idea to keep things pretty cool to improve reliability - particularly during a "one time" event in the middle of nowhere!

Appendix A: Configuration of and notes related to the BeaconBlaster 6

The configuration of the Turn Island Systems BeaconBlaster 6 *("BB6")* is via a text file found on the USB-accessible mass storage device named "commands.txt". More details about the operation of the BB6 may be found in the manual found on the Turn Island Systems web site: **https://turnislandsystems.com/downloads/**

As mentioned earlier, my configuration is not standard because:

- With "high side" local oscillators used for conversion of the 45 MHz IF down to the HF bands, the FST4W modulation had to be "backwards" - that is, Tone 0 is the highest and Tone 3 is the lowest frequency.
- I needed to generate two CW *(unmodulated)* carriers for use as the "T1" and "T2 signals which would be modulated as required externally.

The following are the contents of this configuration file:

```
# Delete the "Testing Only" commands for proper operation.

# Testing Only: this allows the BB-6 to run on the internal crystal if the ref clk is absent
#CLK 0

# Testing Only :this allows the BB-6 to free-run without GPS input.  GPS is still required to get the 1pps clock
#GPSSET 0

# The BB-6 normally monitors the amplifier power supply.  No power supply here...
# Ignore loss of amplifier voltage
VAMP 0

# Set your callsign, power level, and grid
# callsign is mandatory.  Power defaults to 30.  Grid will be calculated from GPS if not provided here.
ID KA7OEI 37 DM39

# There are four PLLs in two Si5351s: 0,1 (used for channels 1-3), and PLL # 2,3 (channels 4-6)
# PLL [PLL #] (A] [B] [C]
# set PLL #0 to 850 MHz (10 MHz ref clk)
PLL 0 85 0 1

# Configure Chan 1 for 44,999,999.707 Hz, inverted FSK
# FREQ [channel] [PLL #] [A] [B] [C] [FSK 0] [FSK 1] [FSK 2] [FSK 3]
FREQ 1 0 18 803920 904408 15 10 5 0
# setting channel 1 for transmission every timeslot
```

```
sched 1 1 1


# Formula:  Divisor = A + B/C
# FREQ <chan#> 1 <A> <B> <C>
# A should be EVEN
# B and C are 20 bit unsigned, maximum 1048575
# Useful calculator:  johndcook.com/rational_approximation.html

# set PLL #1 to 899.99 MHz
PLL 1 89 999 1000
# set Chan 2 to use PLL 1 and divide by 20 (= 44.9995 MHz)
FREQ 2 1 20 0 1
# set chan 2 to "fixed" mode
FIX 2 1


# set Chan 2 to use PLL 1 and divide by 20 + offset (= 44.999012 MHz)
FREQ 3 1 20 211 972827
# set chan 2 to "fixed" mode
FIX 3 1

# end
```

Comments:

- There is a rare bug in which leaving the grid square blank *(allowing the computer to calculate it from GPS data)* will result in transmitting an unmodulated FST4W carrier. Rather than risk transmitting a dead carrier, I was prepared to modify the “commands.txt” file in the field to reflect the current location. This bug is being investigated.
- As mentioned, the 4-level FSK signal comprising the FST4W signal had to be spectrally inverted so the frequency of each tone is defined in the configuration.
- The “T1” and “T2” carriers were both generated from PLL1 of the Si5351 *(which was set to 899.99 MHz)* using the polynomials defined in the Si5351 data sheet and the document “Manually Generating sn Si5351 Register Map for 10-MSOP and 20-QFN Devices” found here: **https://www.skyworksinc.com/-/media/Skyworks/SL/documents/public/application-notes/AN619.pdf**
    - Specifically, section 4.1.2 of the document where it talks about programming the “Multisynth”. All that one really needs to know are the following rules:
        - The VCO is divided by the number resulting from: a + b/c to yield the output frequency where:
            - a is an *even* integer from between 8 and 2048 *(inclusive)*
            - b and c are 20 bit integers

        - As noted in the comments in the above configuration file, the values for "a", "b" and "c" are included on the command line.
        - As the precision of the above fractional divisor easily exceeds the number of decimal places of a standard pocket calculator, I used the Windows calculator, which goes out to many more decimal places to help calculated the decimal value required to get as close to the desired frequency as possible.
        - Since the "b/c" portion is an integer fraction, I used the iterative calculator found at **http:johndcook.com/rational_approximation.html** to determine the best-fit approximation of the decimal number, inputting a value of 1048575 as the largest denominator.
            - The fraction produced by the above tool was then converted back to decimal and used to calculate the "exact" frequency produced by the Si5351. For Channel 2 *(T1)* this value was exact, but that for Channel 3 *(T2)* had a slight residual of about 12 mHz.
        - Comment: Since there was another Si5351 on board, I could have probably programmed it to produce the "T2" channel as it would have had a PLL all to itself! Since it is enough to ***know*** the frequency precisely - but not necessarily set it to a nice, even value, the method used was sufficient.
- The BeaconBlaster has a "bug" related to its file system that, as of the time of this writing, is under investigation. While I'm confident that the issue will be solved, I will explain it here just in case.
    - The symptom is that the unit will not transmit when expected, due to the loss of the "commands.txt" file. This is most likely to happen when powering up the unit, but I experienced this issue *during* operation as well.
    - It seems that there is the possibility of a subtle corruption of the file system on the mass storage device and this can manifest itself in two ways - likely related to the same problem:
        - **The "commands.txt" file will disappear.** It should go without saying that this file should be backed up, anyway, but one should be prepared to - in the field - restore this file at any time.
        - **The mass storage device will fail to enumerate or *everything* on it will disappear.** If this happens:
            - The unit must be completely power-cycled *(note that if the USB is connected to a computer, it, too, must be disconnected along with the normal power supply for the unit as either can power the processor.)*
            - When powered back up, the mass storage device should enumerate once more, but the files may still be missing and/or there may be little or no space left on the device.
            - Using Windows Explorer, right-click on the BB6's storage drive and select "tools" and "fix" *(or repair)* the device. This should take only a few seconds and it will likely create a directory with a lot of files with the ".CHK" suffix. Delete this directory *(and thus the .CHK files)* in its entirety and power-cycle the unit again when done. This should "fix" the device and allow you to load your "commands.txt" file once again and restore it to proper operation.

Appendix B - Frequency reference information and observations of the Leo Bodnar GPS unit during the eclipse.

While most of those out in the field used a GPSDO - specifically one of those produced by Leo Bodnar *(see: **https://www.leobodnar.com/**)* - typically the “Mini Precision GPS Reference Clock” - I used a Rubidium-based atomic frequency reference - mainly because I had it. Specifically, I used an Efratom LPRO-101 which outputs a 10.0 MHz signal with stability and accuracy easily exceeding 10E-10 within about 10 minutes after power-up. As a back-up, I also had another Rubidium referenced based on the FEI FE-5680 which - although slightly less accurate, was also known to attain at least 10E-10.

Also on hand was a Leo Bodnar “Precision GPS Reference Clock” which can produce up to two outputs - but if just one output is enabled, the “disabled” output produces a 1pps pulse. While this wasn’t used as a frequency reference, it - like the FE-5680 - was available as a back-up.

As the LPRO-101 worked flawlessly, we didn’t need to use any other device, but this gave us the opportunity to simultaneously compare it, the FE-5680 and the Bodnar using a 4-channel digital oscilloscope to observe discrepancy between the three signal sources. As expected, the two Rubidium standards’ 10 MHz waveforms moved very slowly with respect to each other, taking somewhere around 15-20 minutes to “slide” one cycle. In comparison, the Bodnar would *usually* move seemingly randomly - albeit quite slowly - over the same time period, sometimes drifting slightly high and low in frequency.

While this was an unscientific endeavor, the purpose of this exercise was to see if the Bodnar *(or more specifically, its GPS signal)* was affected by the presumed, rapid and *(geographically)* “local” changes of the ionosphere in terms of propagation. While it is unknown if this was due to the eclipse, it was observed that during the “darkest moments” *(+/- 20 minutes or so of maximum annularity)* that compared to the two “stable” Rubidium waveforms, the 10 MHz output of the Bodnar was observed to “hunt” at a very observable rate: The output seemed to slide *(wobble)* 10s of nanoseconds up and down in frequency over periods of just a few 10s of seconds.

This “wobble” behavior was not observed in the hour or so after the eclipse, so it cannot be known if this was normal behavior due to certain satellite geometries, the number of satellites, or if there was, in fact, some sort of GPS timing anomaly introduced by the *(presumed)* rapidly-changing ionospheric propagation through which the satellites’ signals had to pass that introduced timing uncertainty during the depth of the eclipse. Finally, it is unknown if these changes would have been enough to “noise up” the frequency-versus-time data of the transmitters and receivers that ***had*** been using the Bodnar for frequency reference enough to be visible *(and distinguishable!)* in the recorded data.

As solar eclipses are quite rare, it is rather difficult to study and quantify this *(presumed)* phenomenon, but considerations should be made to do so during future events.

Appendix C: Receive site signal generator

In addition to the transmitter gear, similar equipment was constructed for use a the receive site - namely the Northern Utah WebSDR. The purpose of this gear is to inject signals into the antenna signal path called the "R1" and "R2" signals. Based on the sound that one hears when tuning the "R1" signal, we dubbed it the "clicky box".

- **"R1" signal:** This is analogous to the "T1" signal in that it is BPSK-modulated at 0.5 Hz from a GPS reference. As the signal path is the same as for the transmitter, this signal, too, is delayed by 82.20 microseconds from GPS due to propagation delay.
- **"R2" signal:** This is an unmodulated *(CW)* carrier. When compared with the transmitted "T2" signal, continuous, direct measurements of frequency and Doppler shift may be made.
- As neither of these signals are transmitted "over the air", there is no Morse ID.

Rather than using a "Beacon Blaster" to generate the 45 MHz signal, a QRP-Labs ProgRock 2 *(modified for external clock input)* using the same circuit as in the transmitter for the BPSK modulation. As with the transmitter, the signal is passed through a 45 MHz filter/amplifier module to assure spectral purity of the IF signal and down-converted in the same manner to the individual HF bands.

The same converter boards are used - but modified by omitting all amplifiers other than U300 - that used to amplify the local oscillator - the reason being that the signals needed are many orders of magnitude weaker: Only about -90 dBm or so is actually needed on the antenna line.

The individual converter modules are "combined" using a row of 200 ohm trimmer potentiometers, "high" side connected to a common bus and applied to a 16 dB gain bipolar amplifier to boost the composite *(80-10 meter)* signals. This composite signal is then inserted into the signal path of the RX-888 and the KiwiSDRs via a 20 dB directional coupler.

Finally, the R1 and R2 signals are measured to determine the equivalent signal power applied to the main antenna feedline *(from the TCI-530 omnidirectional antenna)* so that a direct comparison may be made with the R1/R2 and received T1/T2 signals.

Generation of "R1" and "R2" carriers

At the receive site, analogs of the "T1" and "T2" carriers *(the "R1" and "R2")* are generated and injected into the antenna signal path at ***known*** levels, equivalent to signals in the -90 to -100 dBm range at the receive antenna. Using the same - but modified - converter boards as the TX generator, these signals are generated at 45 MHz.

Rather than using a BeaconBlaster, a QRP-Labs ProgRock 2 board - modified to accept an external 25 MHz reference *(derived from an external 10 MHz source and multiplied by 2.5)* is used to generate the 45 MHz signals.

The "R1/R2" Generator, sitting atop a 10 MHz GPSDO.

As with the T2 carrier, the R2 carrier is unmodulated, and neither the R1 or R2 need to have a Morse ID, simplifying the circuitry: All that is required is that the R1 signal be BPSK-modulated in the same manner as the T2 carrier.

The R1/R2 unit has its own GPS receiver - an inexpensive Ublox unit similar to that on the BeaconBlaster - to supply the 1pps. The circuitry to generate the BPSK signal is comprised of the components with designation in the “200s” *(e.g. U200, U201)* - but not including U202 or U205 which is used to do Morse keying. *(Note: The the converter boards are modified versions of the transmit converter boards - the differences noted in the previous section describing the TX converter boards.)*

The outputs of each of the converter boards is applied to the “high” side of a 200 ohm potentiometer - the other side being grounded - and the wipers connected to a common bus via a 220 ohm resistor. While a bit of a “kludgy” and lossy resistive combiner, it is effective and any IMD products produced were below the level of detection.

The composite output was then applied to a simple linear amplifier based on a 2N5109 - the diagram on the right. Practically speaking, I could have used a 2N3904 or similar transistors and accept a roll-off of a couple of dB at 30 MHz. To a degree, the gain can be adjusted by varying the value of R2 - with the 680 ohms shown here as being near the high end of what is practical, yielding about 16dB of gain. It is possible to optimize input matching to 50 ohms by careful selection of R5, but that was not important in this application.

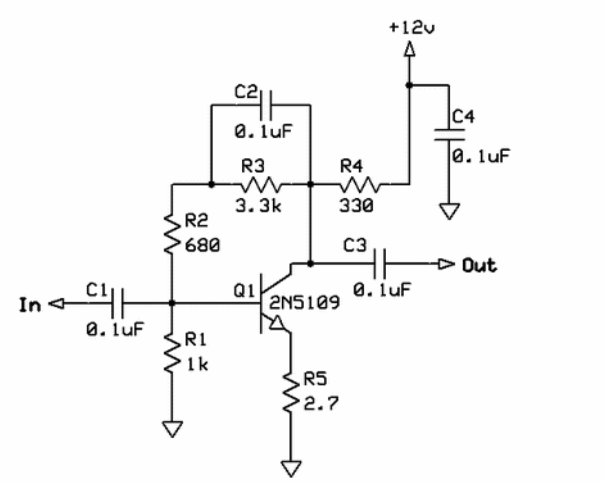

General Purpose RF Amplifier

While the 2N5109 is no longer made, a reasonable substitute - in SMD format - is the BFU590 which has generally similar RF performance albeit with somewhat lower voltage and current ratings - but it would have been more than adequate in this and many other applications.

Combining into the receive signal path

The output from the amplifier was combined into the receive signal path using a Mini-Circuits ZFDC-20-3 - a 20 dB directional coupler. Operated this device “backwards” *(the labeling assumes that one is sampling a signal from the “through” line, not inserting a signal)* it has only about 1 dB of insertion loss from MF through VHF.

At the Northern Utah WebSDR there are two main signal branches on all of the antenna filtering/protection: The “narrowband” branch which feeds band-specific pass filters and the “wideband” branch which has only AM broadcast band reject filtering and low-pass filtering . This branch is what feeds the KiwiSDRs *(and the RX-888 Mk2)* via a “shelving” filter *(described here: https://ka7oei.blogspot.com/2020/08/revisiting-limited-attenuation-high.html)* to maximize the usefulness of direct-sample SDRs. The “R1” and “R2” signals were inserted *only* into the signal path that feeds the KiwiSDRs and the RX-888 that are connected to the TCI-530 Omnidirectional antenna.

* * *

[END]

This document written by Clint Turner, KA7OEI, October 2023

V20231030