# **AM Transmitter**

Gord Rabjohn February 2024

## Introduction

I have designed a low-power AM band transmitter optimised for legal operation. This is to provide audio programing throughout my house for my AM radios. My goal is to get as much radiated power out as is legal. Not that anyone would notice if I exceeded the legal limit, but I wanted to set a challenge for myself.

I wrote this article in a logical order to make it easy to follow, but the path I actually took was winding, with dead ends and late-in-the-day revelations. I was both aided and abetted by measurement errors, lapses in logic, serendipity, misfortune, foolish optimism, desperation, oversights, laziness, and sometimes good old stupidity. But I learned and re-learned so much, and had a lot of fun doing it. I think the resulting transmitter vindicates my design approach, but it is not an easy design approach to get right.

## The Law

FCC “Part 15” regulations, and Canadian regulation RSS-210 (entitled: “Licence-Exempt Radio Apparatus: Category I Equipment”, Annex B, issued by Innovation, Science and Economic Development Canada) allow unlicensed transmission on the AM band. This allows us to make small transmitters that make our old AM radios actually do something if all the commercial AM transmitters vanish (which has not happened as of February 2024, but seems like a real possibility). However, there are two limits to the unlicensed AM band transmitters (unless you are able to accurately measure field strength): the DC power into the final stage of the transmitter must be is under 100mw and the antenna mut be under 3m or 10 feet long.

*“…the total input power to the final radio frequency stage shall not exceed 100 mW, and the total length of transmission line, antenna and ground lead (if used) shall not exceed 3m…”*

The 100mW restriction is actually not the big problem; many ham radio operators operate in “QRP” (low power) mode, and with the right antenna, under favorable conditions at the right frequency, with a good receiver, and some luck, 100mW can propagate for hundreds of kilometres. Though QRP officially refers to powers under 5W, many operate at much lower powers. There is a 1000 miles-per-watt award available (from QRP ARCI). However, hams use large efficient antennas, in bands with less interference than the AM band, and in bands where efficient antennas are possible on reasonable pieces of real estate. This gives them a decided advantage.

Note that the 100mW power is defined as DC power into the final amplification stage, not RF power. This means that efficiency is very important; RF losses cannot legally be made-up by increasing the DC power into the amplifier. Any inefficiencies (and there are many: finite amplifier efficiency and finite inductor “Q” being the main ones) come out of the 100mW budget.

Of course, your may or may not want to adhere to the law. There are many tube-based AM transmitters out there, and I guarantee that most consume more than 100mW. Does it really matter? You decide. Likely, many are not radiating more than 100mW, so arguably *sorta* meet the intent of the law. And, if you really are radiating more power than that, who is going to know? The AM band is an almost uninhabited wasteland these days, so maybe the law is outdated. However, my goal was to try to see what could be done while staying within the bounds of the law.

## The Antenna

The 100mW restriction is not the most important impediment to an effective transmitter. The 3 metre antenna length restriction at 1MHz, on the other hand, is a very severe limitation. A 3m antenna can theoretically be an efficient radiator, but because it is so short relative to a wavelength (which is 300 metres in the middle of the AM band, so such an antenna is about one hundredth of a wavelength), the impedance will be very high, very difficult to match to. We usually think that a high "Q" is good, but in this case, we want the antenna to lose power to radiation, so this high "Q" is not helpful. A short wire (not a loop) antenna will look like a small capacitance with either a tiny series resistance or a huge parallel resistance. (At one frequency, the two representations are equivalent). One can parallel-resonate out the capacitance (using an inductor), leaving a huge resistance. This resistance is caused by the losses in the metal and insulation, and also the energy lost to radiation (we encourage this energy loss). To optimise transmitted power, we must a) minimise the undesired antenna losses by using low resistance metal and good insulators, b) get as much RF voltage across that antenna resistance as possible, and c) minimizing other transmitter losses, especially those in the resonating inductor. Traditional thought is to use a matching network to match the antenna impedance to the impedance of the final stage of the transmitter, but this is a very difficult network if matching up from a low voltage, low impedance solid state amplifier. Why? If the "Q" of the matching elements is lower than the "Q" of the antenna, then they burn more power. Any way you slice it, to radiate power efficiently, you will end up with a high voltage at the antenna, so you will need an extreme impedance transformation from your solid-state amplifier. With a short antenna like this, you will need a substantial RF voltage on it (hundreds of volts at 100mW) to get it to radiate.

If you have a choice, operating at the high frequency end (shorter wavelength end) of the AM band will give you an easier antenna impedance to match to, because the 3m length is a larger fraction of a wavelength at shorter wavelengths.

(We should also mention the loop antenna. The long-wire antenna I discussed above terminates in an open circuit; a loop antenna terminates in a short. A loop antenna looks like a small inductance with either a small resistance in series or a large resistance in parallel. (At one frequency, the two representations are equivalent) I have been told by someone much smarter than me that a loop antenna is less efficient than a long-wire antenna. A long wire antenna runs at high voltages, so dielectric losses dominate. However, we have numerous very high-quality dielectrics available. A loop antenna runs at high currents, so metal losses dominate. Even the best metal still has measurable loss. For this reason, long distance electrical power transmission is done at high voltage because dielectric losses tend to be lower than metal losses. In any case, I did not investigate the use of a loop antenna, but I probably should. One could resonate the inductance out with a big series capacitor, leaving the small resistance, which might interface well with a low voltage solid-state amplifier. Ferrite bars could be used to increase the effectiveness of the loop antenna (they work well as receiving antennas, so should work equally well in a transmitter) All of this is for another day. )

So, I should just strive to get as much RF voltage into the antenna as possible. The antenna capacitance gets in the way, making it more difficult to develop a high RF voltage, so I parallel-resonate it out. I made some early attempts at measuring the real part of antenna impedance which indicated a resistance of somewhere around 200kohms. About 100V RMS develops 50mW in 200kohms (that's an implicit assumption of 50% efficiency). You could do this with a transistor, but in my opinion, this screams out for a vacuum tube! Even without any matching network, a decent high RF voltage is available from a tube. I made the somewhat arbitrary decision to design a transmitter based on a tube amplifier operating at 200V DC on the plate (which should give me as much as 140V RMS), at 0.5mA plate current, which is exactly at the 100mW limit.

(Another side thought: I elected to parallel-resonate the antenna with an inductor and drive with a high impedance. You could also series-resonate the antenna with an inductor and drive with a low impedance. My thought was that the voltage transformation was going to make that more difficult; the circuit must transform a low voltage up to several hundred volts of RF voltage. This is what you would do if you were to drive it with a solid-state amplifier. I tried this briefly, and results were poor, but others are most welcome to try.)

To keep within the law, we need to run at 100mW (DC input to the final). For example, 0.5mA at 200V or 0.4mA at 250V or 0.33mA at 300V. This encourages me to reduce losses as much as possible; I can't just pump in more power to compensate for the losses. The 0.5mA current calls for a very small tube. My first choice was the diminutive 6AK5 because it is a pentode (giving it a high output impedance that will not load the output tuned circuit), but I write much more on that later.

I considered several different modulation modes. The traditional approach would be plate modulation, but at these low currents an enormous plate inductor would be required. Or, more straight-forwardly, just a high voltage amplifier with a resistive load would suffice, since such low plate current is required. However, I ignored conventional wisdom, and elected to use cathode modulation. With the aid of low voltage op-amp circuitry, it is possible to precisely control the cathode current in sympathy with the audio input. The RF is imposed on the grid. A parallel L-C tank on the plate cleans everything up.

## The First Test Bed

My first test bed included the tube, the 1MHz drive oscillator, a 200V DC power supply for the plate, a separate screen grid supply, and a plate tank (parallel L-C). I initially used a high-quality ferrite variable "loopstick" for the "L", and a mica fixed "C", but after doing some measurements, I was disappointed in its Q. So, I removed these. For the "C" I used the best variable capacitor I could find, a ceramic insulated air dielectric Hammarlund with 110pF max. I knew that I would need to experiment with the "L", so I used 2 binding posts that would allow me to easily connect different inductors. I also included a 1000:1 capacitive divider at the output that allowed me to monitor the voltage swing at the plate without loading it.

Here is how I selected the values for the plate capacitor and inductor (the "Tank"). Ideally, the inductor would exactly resonate with the antenna capacitance. Some variability will be required, for fine tuning and to accommodate different antennas: this could be done by using a variable inductor (but high Q variable inductors are rare) or adding a variable capacitor across the antenna. With an estimated antenna capacitance of 40pF, plus maybe 10pF for tuning, the inductor would need to be 500uH. Such an inductor will be quite close to resonance at 1MHz. You can look at this 2 ways: being close to resonance tends to elevate the apparent inductance, OR being close to resonance means that there is significant capacitance, which would add to the 50pF mentioned above. So, the inductor will have to be lower than 500uH. You don't want it too much lower because the lower it gets, the more current it carries, and this will increase its resistive loss. Also, a lower inductor would require more capacitance, making the tuning even narrower. I ended up using values in the range of 200uH - 400uH at 1MHz. For most of the experiments, I used two large encapsulated 100uH fixed inductors in series because I had them, they are stable, and they have a fair "Q". These provided approximately the same loss as I expected from an antenna (at 1MHz) (I was wrong), and at 0.5mA cathode drive, something like 100-200V p-p of RF swing was present at the plate. I later found that three 100uH chokes in series was even better, offered more output power. Four was too much at 1MHz. I made some quick tests at 1.544MHz; at that frequency even 200uH was too much, so I used 100uH.  My highest "Q" coil is one that I wound with coarse Litz wire for a crystal radio project. It offered the very best "Q" but is not shielded and is very sensitive to its environment (including my hands), making experiments unreproducible. However, it was the best coil for use in the final transmitter.

The cathode current was set with a variable resistor in the cathode. This resistor needs to vary between about 1K and 500K for a complete curve, so I used a multi-turn pot in series with a resistor substitution box. The screen grid was supplied with an external lab supply.

I did most of my work at 1MHz, because I have a 1MHz crystal, and it's a nice round number. I probably should have used a higher frequency, just because antenna impedances will generally be more favorable at that end of the band. For the 1MHz (fixed) RF source, I decided that a crystal (as opposed to a VFO) would be the most reliable approach. A crystal gives excellent frequency stability, and it turns out that stability is an issue with the high Q of the antenna. I had started with a JFET "Analog" oscillator with a resonant tank at the output, but I had issues with reliable starting, and pulling. More on that later.

I assumed class A operation of the amplifier at first, but I soon decided that driving the grid heavily into overdrive (deep into class C) made the system more efficient and predictable.

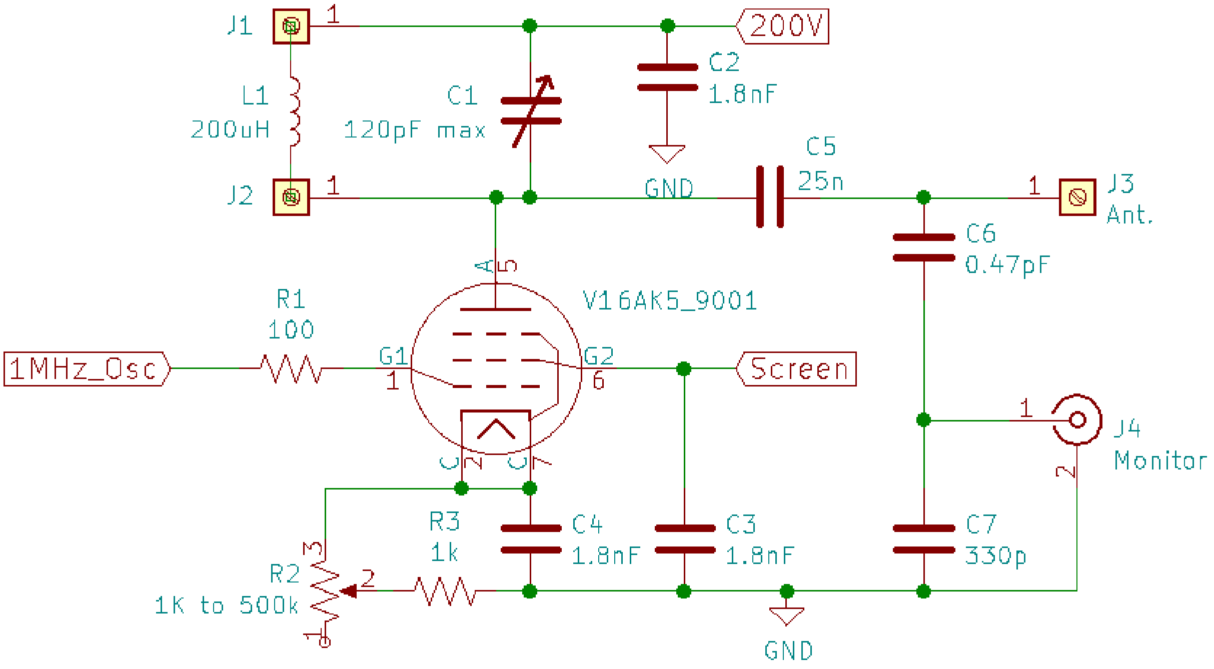

This monitor port was connected to an RF voltmeter and attenuator (with a cable with about 100pF of capacitance). Since I wanted to measure linearity over a 30dB range of powers, I was concerned about the linearity of my RF voltmeter (a HP 400EL). Rather than rely on it being linear over 3 decades of power, I used an RF attenuator (a JFW 50BR_008, and later an old GenRad step attenuator) cascaded to the input of the voltmeter. The ports were both padded with 50 ohms to make sure the mismatch did not affect the attenuator's accuracy. The attenuator was adjusted in 1dB steps until the RF voltage read on the voltmeter was within 0.5dB of "-60dB" ("0" on the most sensitive scale of the 400EL). The -0.5 to +0.5dB residual was also recorded.  Since the voltmeter was always operating at almost the same power, its linearity was unimportant. The attenuator had been proven to be linear to within 0.1dB (with a DC test). So, a sweep consisted of:

Set DC cathode current to 0.5mA
Peak the plate capacitor for maximum RF voltage swing
Set DC cathode current somewhere between 0.01mA and >1mA
Adjust the attenuator in 1dB steps until the RF voltmeter reads close to 0dB
Record the current, attenuator setting, and residual (delta from 0dB)
Repeat 3,4,5 for cathode currents from 0.01mA up to over 1mA in 2 or 3dB steps.
I paid special attention to output swing at 0.5mA current (as a proxy for efficiency) and at saturation.

I did this sweep to evaluate tubes, tuning, predstortion, effect of plate load, you name it. I did dozens of these sweeps, manually.

## Tube Selection

Perhaps the most interesting work was selecting the right tube. I needed a tube that operates well at low current, and I wanted a sharp cut-off to minimise drive requirements and ensure that the tube could be cut off. As I mentioned, I thought a pentode would be best because of its high output impedance.  I measured the static plate resistance of various small pentodes. The 9001 pentode (at Ip=0.5mA, Vp=200V, 150V screen grid voltage) is at least 5Mohm, increasing at higher plate voltages. The 6AK5 is not as good. However, either of these tubes will not significantly de-Q the plate resonant circuit. A 200uH inductor with a Q of 500 (which would be extremely good) at 1MHz has a real part of about one tenth of the plate resistance, so plate resistance can be ignored.

As I said, my first choice was the 6AK5, a very common miniature VHF pentode, and I wired up my prototype for that tube. As long as pin 2 and 7 are connected together, there are many tubes that are plug-in compatible (such as 6AG5, 6AH6, 6AS6, 6AU6, 6BC5, 6BH6, 6CB6, 6CE5, 6CF6, 6DC6, 6DE6, 6DK6, 6EW6, 6JK6), and I tried several. I started with the 6AK5 and 9001, and also looked at the 6AG5, 6AH6, 6AS6, 6AU6, 6BC5, 6CB6, and 6DK6. I believed that a remote cutoff pentode would not work well, and the one I tried, the 6BZ6, proved me correct. In each case, I measured RF output swing vs. cathode current over a wide range of currents (I plotted a linearity graph for each one). I was looking for the most output power (both at high drive and at 0.5mA drive) and the best linearity. One tube stood out: it was what *I thought* was a 6AK5: it had the best efficiency, and oddly the screen grid voltage was entirely unimportant. Looking at the tube closely, I realised that it was not a 6AK5, rather it was a 6HA5 (with a deceptive label): a modern high performance VHF triode that is pin-compatible with the 6AK5. It offered 2dB better output power (at 0.5mA cathode current) than any of the pentodes with good linearity. I assume that since the screen grid current is wasted, the efficiency of pentodes seem lower.

With my sudden and entirely serendipitous discovery that a triode might be a better tube for me, I tested several other triodes. The diminutive 6CW4 Nuvistor came very close. I also tested a 6C4 (which is half of a 12AU7) (terrible, a medium mu triode) and half of a 12AT7 and 12AX7, as well as half of a subminiature 6112. The 6HA5 seemed to beat them all.

I researched the 6HA5/EC900 on the web. It seems to be a pretty common tube, I have 4 in my personal stock, and the OVRC tube stock has more. There is not much written about it except for a Japanese web site from https://radiomann.sakura.ne.jp/HomePageVT/TV_Tuner3.html that describes a range of "frame grid" triodes that were used in TV tuners. Frame grids are wound on a ridged metal frame, allowing finer grid wire to be used, with the frame helping to maintain the mechanical strength and shape of the grid. The 6HA5 (and similar tubes like the 6HM5, 6HQ5 that share the same base, and several (like the 6FQ5, 6GK5) with different base connections) has a shield between the grid and the plate, reminiscent of the beam forming electrode in a beam tetrode. The purpose of this shield is to better isolate the plate from the grid (Lowering plate-to-grid capacitance), allowing higher stable gain. Indeed, the grid to plate capacitance is on the order of 0.3pF, considerably lower than any other triode I was able to find (ignoring the very high voltage 6BK4). The grid wires of the 6HA5 are apparently 8 microns in diameter, about a tenth the diameter of a human hair!

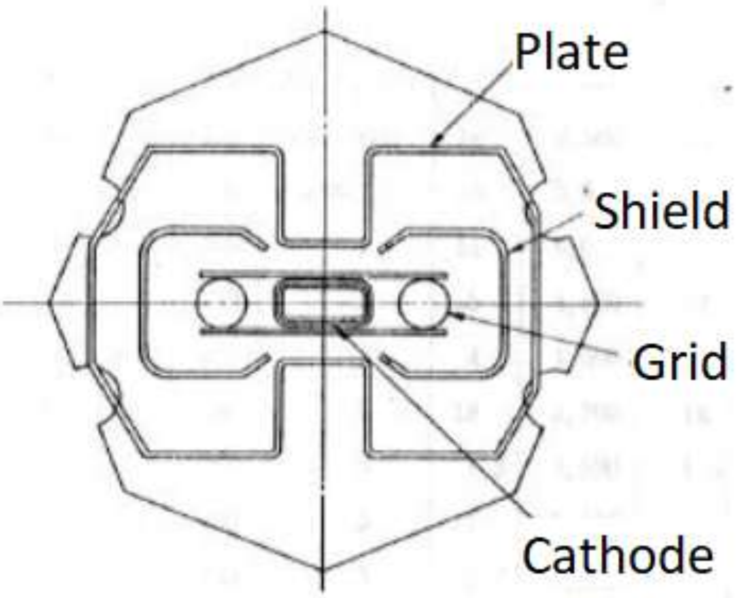

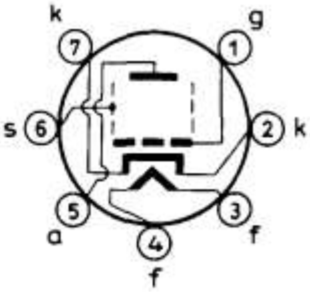

I include some plots (linear-linear) of some of the tubes. Do not examine the too carefully as there are issues with my technique in some cases, but they illustrate the differences between tubes, and give a general view of the shape. In these graphs, the x-axis is DC cathode current in mA, and the y-axis is the "gain" as defined by RF output voltage divided by DC cathode current. (with rather arbitrary scaling)

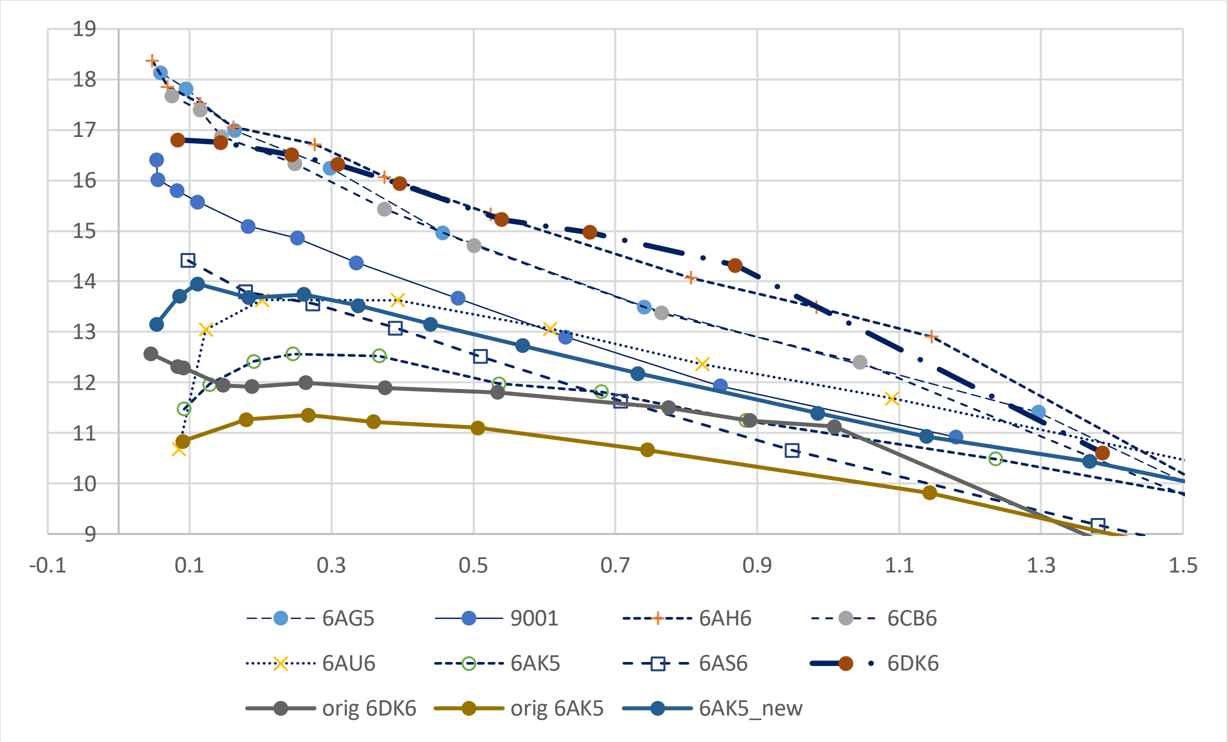

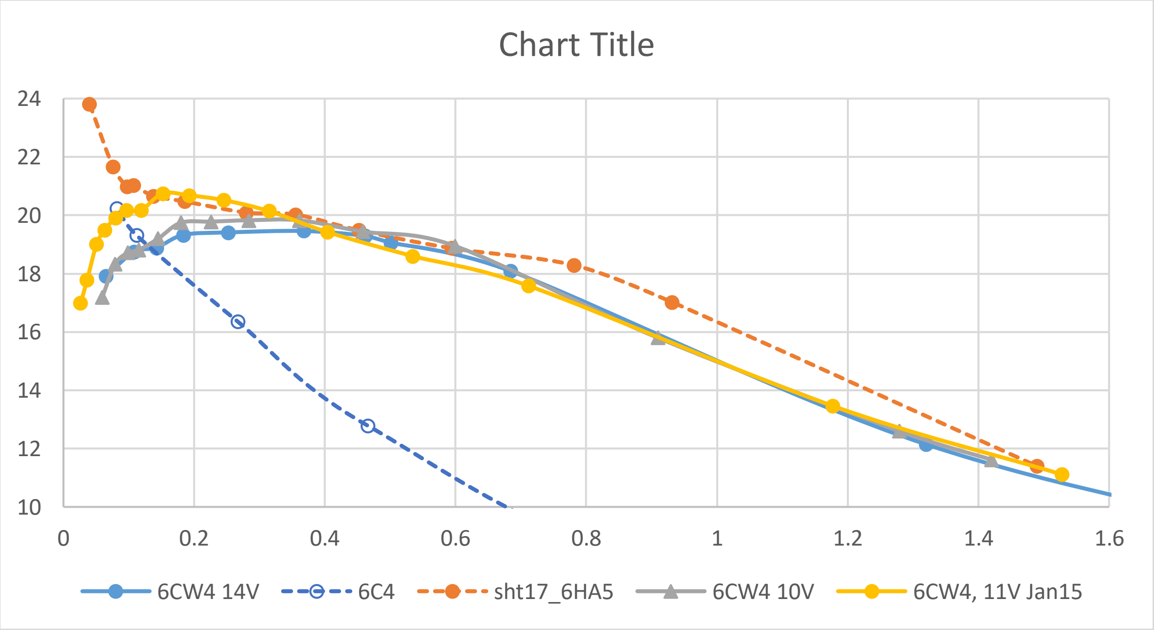

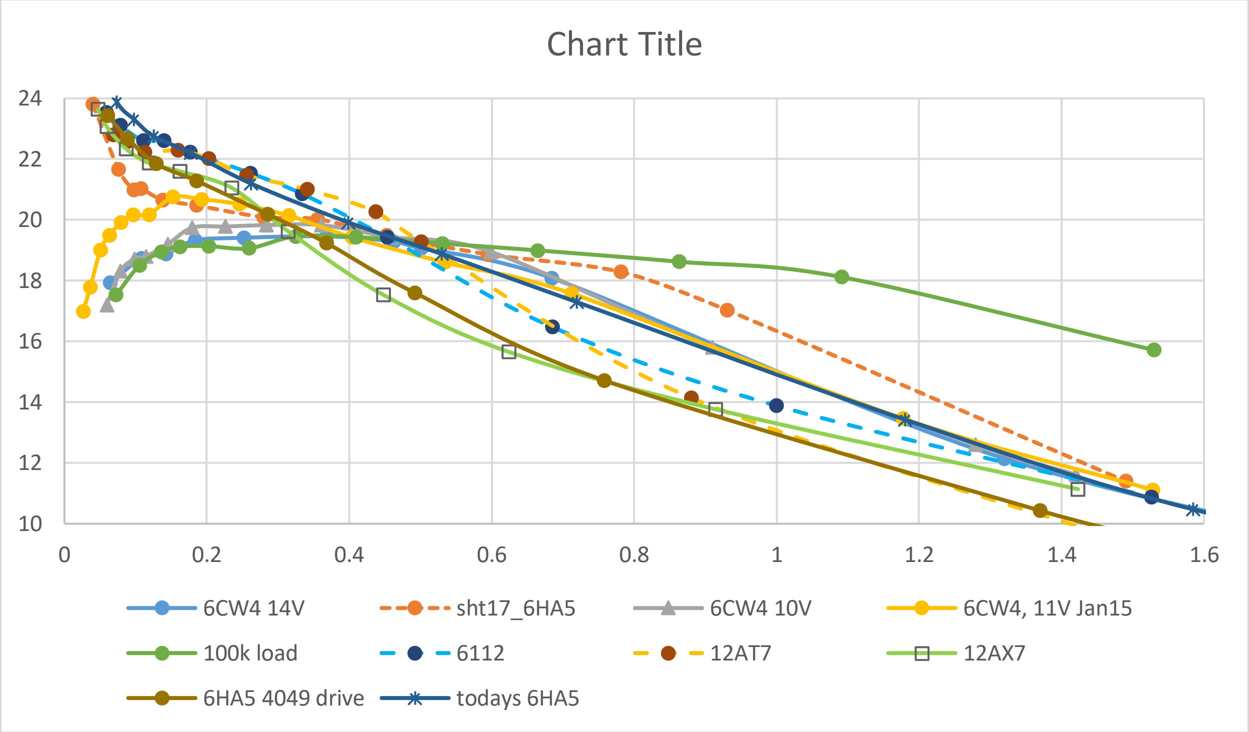

## Experiments Leading Up to the Final Design

I discovered that the shape of the resulting curve depends on the plate load. The capacitance required for peak voltage depends slightly on the cathode current. If you peak plate voltage at low current, you get a different shaped curve than if you peak it at 0.5mA or over 1mA. The reason for the difference is because the plate impedance depends on the cathode current.

I played with grid drive. My first crystal oscillator circuit was “linear” using a JFET oscillator followed by a tuned buffer. I found that the changing impedance of the grid loading the oscillator (especially with the triodes) caused the output tank to become mistuned. (Referring to Terman, it is well known that the grid impedance in a class C amplifier varies as the plate load is changed, and can even become negative) This caused hysteresis and discontinuities in the curves. The DC voltage on the grid (to ground) has no effect, since the cathode circuit keeps the tube operating at the same current and therefore regulates the grid to cathode voltage.

I also played with the grid drive and the shape of the grid waveform. Initially, the waveform was a 15V p-p sine wave. The DC cathode voltage approximately followed the voltage at the peak of the sine wave. This makes sense, as the cathode voltage will adjust itself until the cathode current is right. With 15V p-p grid swing, the tube is only conducting current at the top of that sine wave, so the cathode voltage will be around the same voltage as the peaks of the wave. The tube “wanted” as much grid drive as I could give it, so it was likely conducting only at the tip-top of the grid voltage swing. This suggested to me that maybe driving the grid with pulses (say 10% duty cycle) might work well, and indeed it did. So, I build an oscillator based on a CMOS 4049 inverting buffer, and fed that square wave into a pulser. I discovered that pulses of more than 8V were required, which lead to a dual supply arrangement. The duty cycle is set by the series C13 (150pF), and is approximately 10%.

One other advantage of the 4049 oscillator is that harmonics of a crystal can be used. I was able to use the 2nd and 3rd harmonic of a 500kHz crystal and several harmonics of a 250kHz crystal. (without any circuit modifications, as long as the plate tank was tuned to the harmonic) The selectivity of the plate resonant frequency is adequate for suppression of all but the desired harmonic.

I mentioned the technique I used to measure linearity. Certainly, a limit to linearity will be clipping; when the waveform on the plate approaches 400V p-p (with a 200V DC plate supply), no more swing is possible (since the plate runs out of voltage on the negative excursions), and the amplifier will saturate. The current drive that makes this happen will be determined by the impedance presented to the plate; higher Q of the plate load means higher impedance, and therefore less current will be required to reach saturation. But what about below saturation? As long as the plate impedance remains constant and the duty cycle of the tube remains constant, the swing at the plate should be exactly proportional to the cathode current (once leakage is subtracted out: there will be some level of leakage around the tube and thru the grid-plate capacitance. When calculating linearity, this must be subtracted out.) It turns out that it is approximately proportional, but it’s not perfect. There is a slow slope in the curve (which gets well hidden in log (dB) graphs) that I cannot explain!

## A Description of the Final Design

Please refer to the schematic. It shows several blocks: central to the operation of the circuit is the tube itself, which is similar to the testbed. The tube grid receives RF drive from the crystal oscillator. The cathode current is controlled by the modulator, which includes the audio amplifier and a predistorter. The tube is monitored with a circuit that can detect the RF output or measure cathode current. The monitor also includes circuitry that detects clipping. A power supply energises all of these circuits.

V1 is the 6HA5 tube configured as a class C amplifier. The cathode is decoupled for RF by C5, which should be small enough to not affect audio frequencies. The plate bias is applied through L1, which is mounted with binding posts for the purposes of experimentation. C1 (plus the antenna capacitance) resonates with L1 at the operating frequency. Note that C1 is “live”, care must be taken to ensure you do not receive a shock from it or its shaft. (It would have been safer to place it on the other side of C2, but I wanted to make the path of the resonant circuit short.) C2 is simply a DC blocking capacitor. C3 and C4 (with an external instrument cable in parallel with C4) form a 1000:1 capacitive divider. This divider will provide minimum perturbation (degradation of Q) to the RF signal at the plate. The circuitry surrounding the tube is on copper clad board so that a ground plane is always close.

The grid sees pulses of fixed amplitude from the crystal oscillator. The cathode current is regulated by the cathode modulator. The cathode voltage will adjust itself until the average cathode current is at the set value. Higher cathode current causes the cathode voltage to become more negative, so the grid-cathode voltage increases at the peaks of the pulses.

Using the best possible inductor at L1 is important (C1 is also important, but good air-variable capacitors are fairly easy to come by). Any loss in the inductor is wasted RF power, and with such a high Q antenna, you can easily loose most of your RF power in the inductor. The best inductor I have is a large basket-wound air-core litz wire inductor I had wound for a crystal radio project. Its inner diameter is 4.5” and made with litz wire of about 0.046” diameter. 46 turns yields 234uH, with a Q too high for me to accurately measure (but over 200, probably over 400. The LCR meter displayed a negative Q which is impossible; negative Q means gain rather than a loss), but most certainly better than the any other inductor in my collection. It is perfect for 1MHz operation. For 1.544MHz, operation, I tap it down to 132uH (32 turns). Its Q is still too high for me to measure accurately. I used large Cambion 100uH encapsulated inductors for precision work, because the open basket inductor is very sensitive to changes (moving hands, for example) in the local environment. The 100uH inductors are reported to have a Q of 75 minimum at 2.5MHz. Measured Q is 330 at 1MHz, but my LCR meter is optimistic, so it is certainly lower than that.

The crystal oscillator uses a very traditional circuit based on a CMOS inverter, U11C. R11 sets the bias into the linear range. C11 and C12 provide the correct load for the crystal. R12 limits crystal current and adds phase shift. The output of the oscillator is "sharpened up" with U11B. The edges of the 1MHz square wave are differentiated with C13 and R13, resulting in 100ns pulses into U11F. D11 prevents the input of U11F from going below ground. This train of pulses is applied to the grid of the tube through a resistor (good practice for stability). I experimented with duty cycle, and 10% is approximately a sweet spot. Note that this circuit operates between the -Vss and +Vdd line, so it is operating at about 15V (the Zener diode D44 limits the negative voltage to keep it within specs for U11). The output swings up to about Vcc (about 8.5V). Since the duty cycle is about 10%, the peak cathode current is about 10X the average current. This forces the grid positive with respect to the cathode at high currents, so cathode voltage will actually be slightly lower then Vcc at high drive.

The modulator uses a standard op-amp current source configuration. The current through Q21's collector is determined by the voltage drop across R25, which is determined by the voltage on pin 3 of U21. A 2K current monitoring resistor means the current will be set to 0.5mA per volt on pin 3. R25's value needs to be high enough to make the predistorter useful, but not too high that it imposes head room limitations. Pin 3 is driven by the "Carrier" potentiometer (that sets the static carrier level (generally 0.5mA)) and the audio from blocking cap C22. The audio from the external source is amplified by U21B, where R22 and R23 set the gain to 10. The "Mod'n" potentiometer sets the modulation depth. A regulated voltage for the "Carrier" potentiometer is generate by U22, a 3-terminal regulator, where R26 and R27 set the output voltage to about 3V.

The modulator includes a predistorter that consists of potentiometer PreD_Thr, Q22 and potentiometer PreD. At higher levels of cathode current (and RF swing), the output starts to saturate such that more cathode current does not provide proportionately more swing. "Saturation", when the tube is running at maximum swing and the swing cannot be made any larger, cannot be fixed, but the "softening" of the curve can be. The simplest circuit would be a diode (in place of the E-C junction of Q22) and the PreD resistor. At low drive, the voltage across R25 is not enough to turn on the diode. At higher drive, the diode comes on and shunts R25 with the PreD resistor. Now, the effective resistance for the current sensing resistor is lower, so the current is higher than it would have been without the predistorter. This has the effect of propping up the output power as the tube starts to saturate. This more flexible circuit replaces the fixed diode with Q22 and PreD_Thr (Predistorter Threshold), which acts like a diode with variable turn-on voltage. This allows the voltage at which the predistortion starts to have an effect to be varied. The circuit works and measurements show that it can extend the linearity curves by several dB. In my opinion, the usefulness of this circuit is limited; it is difficult to set up properly, and it does not produce a clear-cut improvement in audio fidelity. These three parts can be safely eliminated.

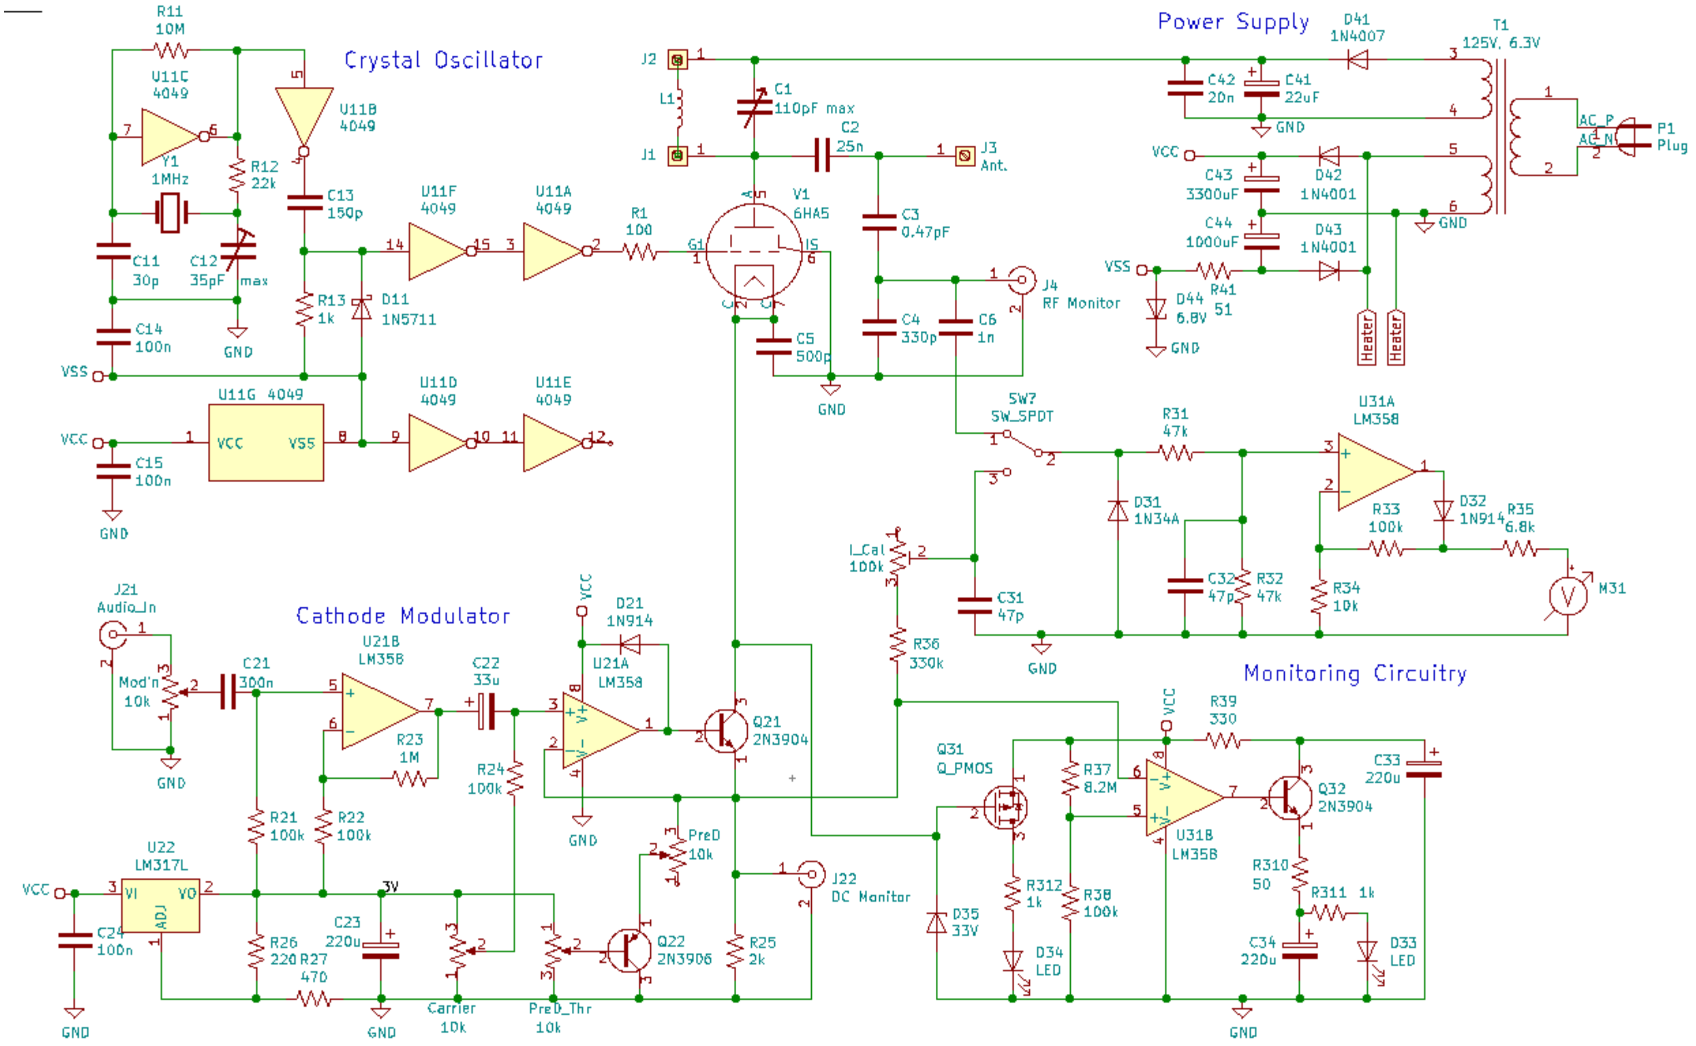
Crystal Oscillator
Cathode Modulator
Power Supply
Monitoring Circuitry
R11 10M
U11C 4049
U11B 4049
Y1 1MHz
R12 22k
C11 30p
C12 35pF max
C13 150p
C14 100n
R13 1k
D11 1N5711
U11F 4049
U11A 4049
U11D 4049
U11E 4049
U11G 4049
VCC
VSS
C15 100n
GND
R1 100
V1 6HA5
J1
J2
L1
C1 110pF max
C2 25n
J3 Ant.
C3 0.47pF
C4 330p
C5 500p
C6 1n
J4 RF Monitor
SW7 SW_SPDT
D41 1N4007
T1 125V, 6.3V
C42 20n
C41 22uF
P1 Plug
AC_P
AC_N
C43 3300uF
D42 1N4001
C44 1000uF
D43 1N4001
R41 51
D44 6.8V
Heater
Heater
R31 47k
U31A LM358
D31 1N34A
C32 47p
R32 47k
R33 100k
D32 1N914
R35 6.8k
R34 10k
M31
I_Cal 100k
C31 47p
R36 330k
J21 Audio_In
Mod'n 10k
C21 300n
U21B LM358
R23 1M
C22 33u
U21A LM358
D21 1N914
Q21 2N3904
R24 100k
R21 100k
R22 100k
U22 LM317L
VI
VO
ADJ
C24 100n
R26 220
R27 470
C23 220u
3V
Carrier 10k
PreD_Thr 10k
Q22 2N3906
PreD 10k
R25 2k
J22 DC Monitor
Q31 Q_PMOS
R37 8.2M
R39 330
U31B LM358
Q32 2N3904
C33 220u
D35 33V
R312 1k
D34 LED
R38 100k
R310 50
R311 1k
C34 220u
D33 LED

There are 3 parts to the monitor circuitry. The first monitor circuit causes a LED to illuminate at the peaks of the envelope. Q31 is a p-channel MOS device that turns on when the cathode is appreciably lower than Vcc (which, as stated above, happens in high drive conditions). When it is on, it illuminates LED D34. Honestly, it activates only at the highest levels of overdrive. Diode D35 prevents the cathode voltage from rising to a voltage that might damage Q31 at very low cathode currents. The second monitor circuit causes a LED to illuminate at the troughs of the carrier, when the carrier disappears. U31B is used as a comparator, which turns on if the voltage across the current sensing resistor R25 drops below a threshold set by R37 and R38. The output of U31B turns on the emitter follower Q32, which illuminates LED D33. C34 is a pulse stretcher that makes the illuminated LED easier to see. R39 and C33 prevent current spikes from the LED from affecting other parts of the device. The third monitor circuit is a detector that rectifies the RF from the monitor port and indicates its relative strength on a meter. This is useful for tuning. This circuit is also used to indicate cathode current. When the switch is in position 1, the diode received RF from the monitor port and causes a rectified voltage to appear. R31, R32 provide the DC return for the diode, and C32 filters out the RF. This is applied to an op-amp DC amplifier with a gain of 10 (set by R33 and R34). The output goes to a meter. Resistor R35 is selected so the meter reads full scale when the tube is running at maximum RF voltage swing. This circuit can also be switched (position 2) so that the meter reads the cathode current by sensing the voltage across R25. R36 and I_cal are adjusted for an accurate meter reading.

The power supply is straight-forward, with 3 half wave rectifiers supplying approximately 200V (actually about 192V) for the plate, 8.5V for Vcc and -6.8V for Vss. The tube filament is supplied from the 6.3V winding of the transformer.


DC Monitor
Power Supply
RF Monitor
Cathode Modulator
Monitoring Circuitry
C1
V1
Audio Input
Crystal Oscillator
LEDs


car.
mod.

## Tuning the Transmitter

If frequency is important, adjust C12 while monitoring the output port with a frequency counter. Generally, I find that the crystal is close enough without the need for fine-tuning the frequency. (as long as the crystal is controlling the frequency, which may not always be the case, so it is wise to verify the frequency of operation)

Connect the antenna. Start with predistortion off (PreD at maximum resistance, the base of Q22 at maximum voltage), and audio at minimum. Adjust Carrier (cathode current) for a current of 0.5mA cathode current (by measuring the voltage at J22 or using the meter). Connect an oscilloscope to J4 with a 1m BNC cable (The capacitance of this cable is included in the design). Adjust C1 for a peak in output power (as measured on an RF voltmeter or an oscilloscope on J4). You would like to see approximately 0.2V p-p on an oscilloscope on J4 (indicating about 200V p-p on the plate). If it is less than 0.1V p-p, the antenna may be loading the circuit too much. If it is much more than 0.2V p-p, the antenna may be loading it too lightly. Monitor the RF voltage at J4 while sweeping the Carrier (cathode current) from 0 to 2mA. The RF output should vary, following the cathode current. It may saturate above 1mA. Turn off the meter. Apply your audio input and increase the Mod'n control. The signal as monitored on the oscilloscope should start to show audio. The modulation can be increased until distortion is heard, or sharp zeros are seen in the envelope.

A good way to see what is happening is to inject a little audio, say a sine wave from a generator, at a low index of modulation. Connect J4 to your oscilloscope. Tune up the plate tank C1 for maximum swing. Then vary the carrier level. You will see that at low carrier levels, the sine wave distorts because the carrier hits zero at some points. At high carrier levels, the sine wave distorts because the amplifier is running out of head room; with a 200V supply you cannot expect more than about 400V p-p (0.4V p-p at J4) on the plate no matter how hard you drive it. (In fact, it's more like 340V p-p, because the tube plate needs a certain minimum voltage) Between these extremes, the amplitude of the audio sine wave in the envelope should be fairly constant, though I do notice that the amplitude starts to drop slightly at high currents, even before clipping (that's what predistiortion should correct). If you have a dual trace scope, you can display the waveform from J22 at the same time, and that makes it easier to see when the sine wave is distorting. Set the carrier to be half way between these extremes. When you apply "real" audio, it should remain in the range where the sine wave amplitude was constant.

In my case, with a high Q litz wire basket wound tank coil, a 10 foot antenna, 1MHz, running at 0.5mA quiescent current was too high. I ended up running at closer to 0.15mA quiescent current, otherwise I was clipping. At 1.544MHz, (smaller inductor) I could run at 0.5mA quiescent current and get fair quality audio.

Below, I show screen shots of a 400Hz sine wave constant amplitude audio with various levels of carrier. The first shot shows the bottoms of the 400Hz modulation flattening out when the carrier of the transmitter essentially turns off at the lowest extremes of the audio. The next 3 shots show a faithful sine wave, though you will note that the 3rd shot shows some reduction in the sine wave amplitude, indicating the soft saturation. The last 2 curves show severe distortion as the amplifier saturates and runs out of power. The vertical scale is approximately 50V per division (actually, a bit less than that), so you can see that at saturation, the waveform is approaching 400V p-p (likely about 340V p-p).

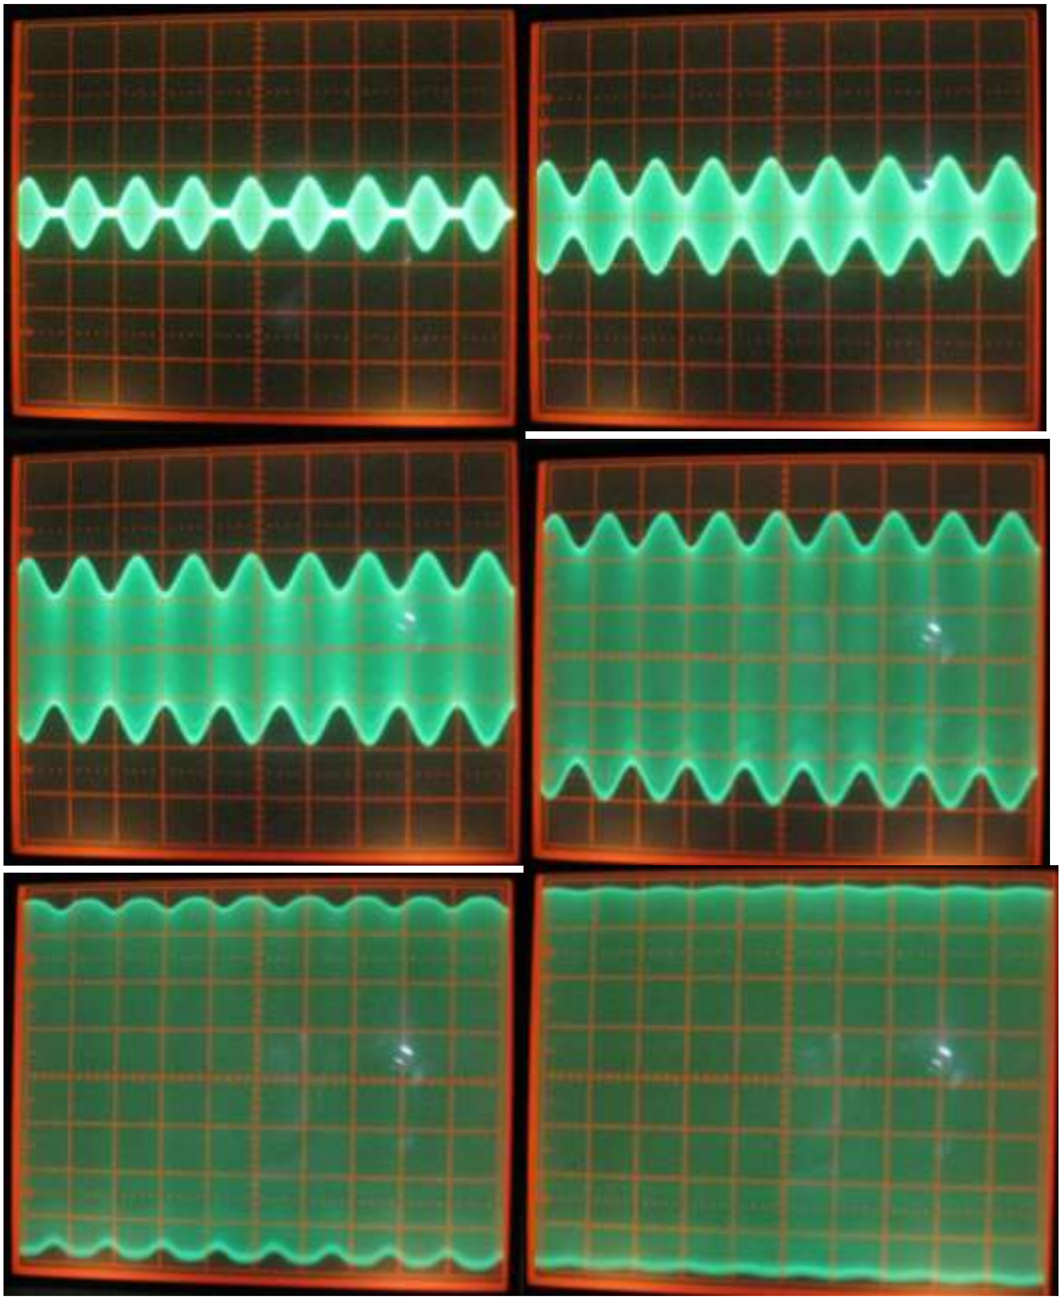

Below I show triangle wave modulation, this time with the audio triangle wave on a separate trace, arbitrarily scaled to overlap onto the RF envelope. The audio is the brighter trace. The envelope should follow the audio exactly. Once again, the first trace shows the bottom tips of the triangle wave being clipped, the next 3 looking pretty good, and the last 2 show saturation. With the audio triangle wave superimposed over the envelope, the distortion becomes very apparent in the last 2 traces.

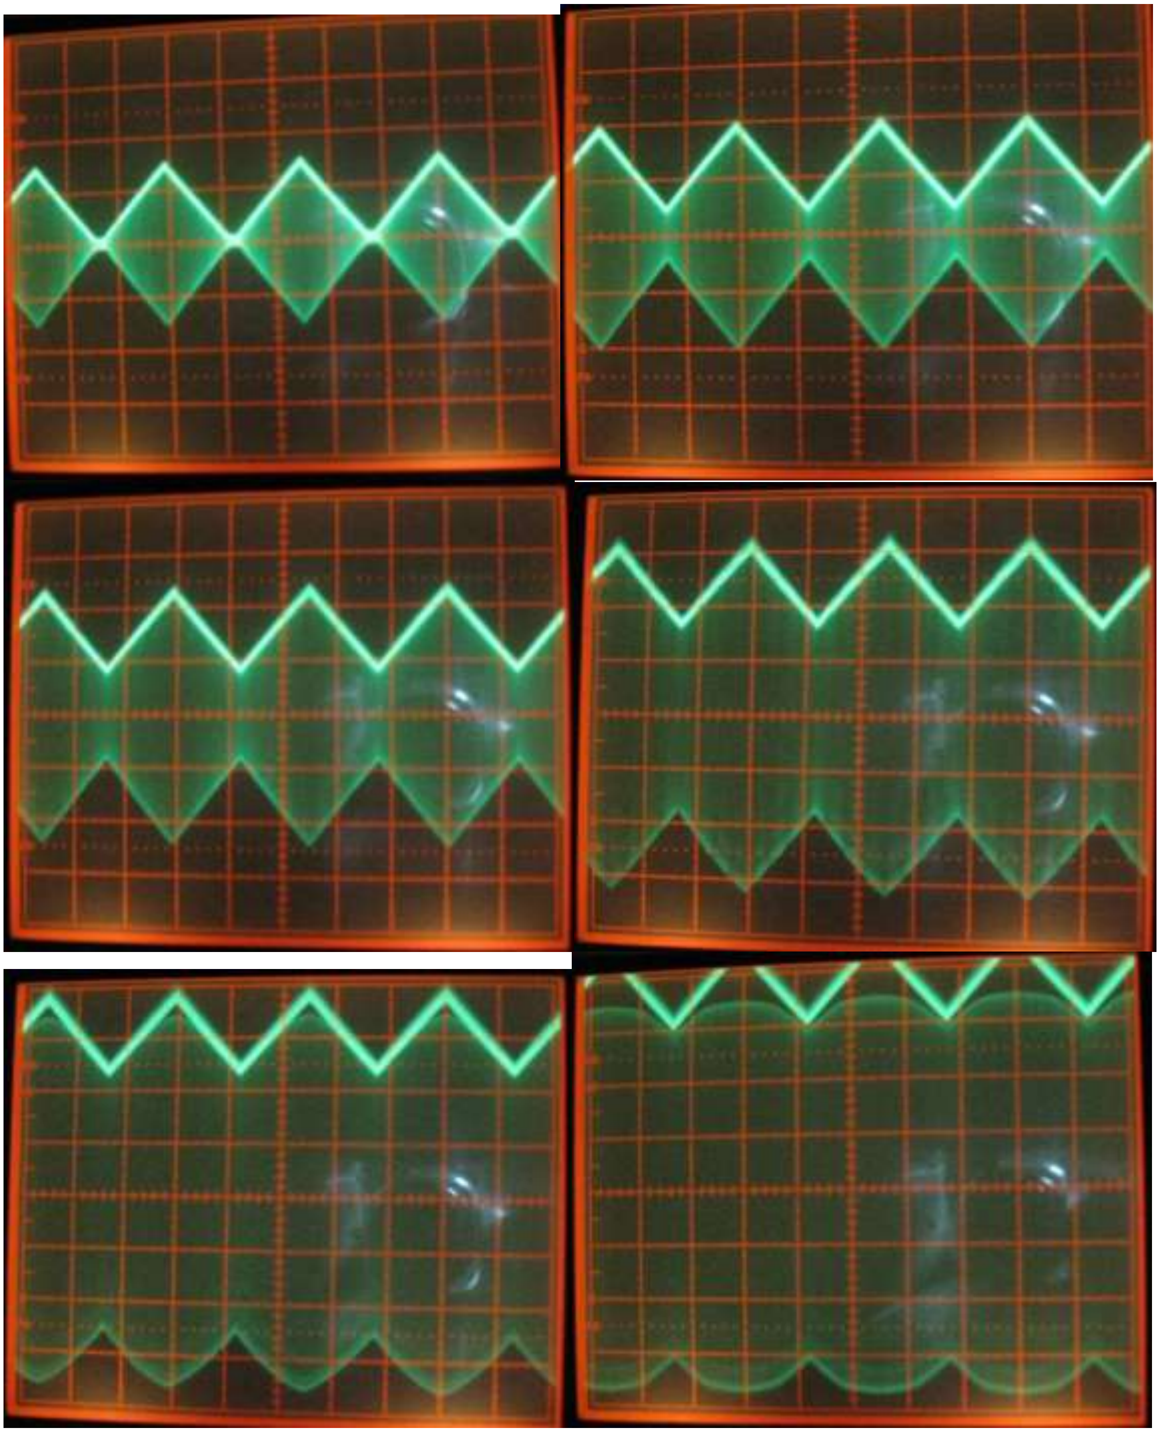

Finally, I show actual music being broadcast. On the left trace with the larger excursion, there are bright spots where the audio clips. The right trace has adequate head-room. It is more difficult to interpret these traces than those with ideal sine waves.

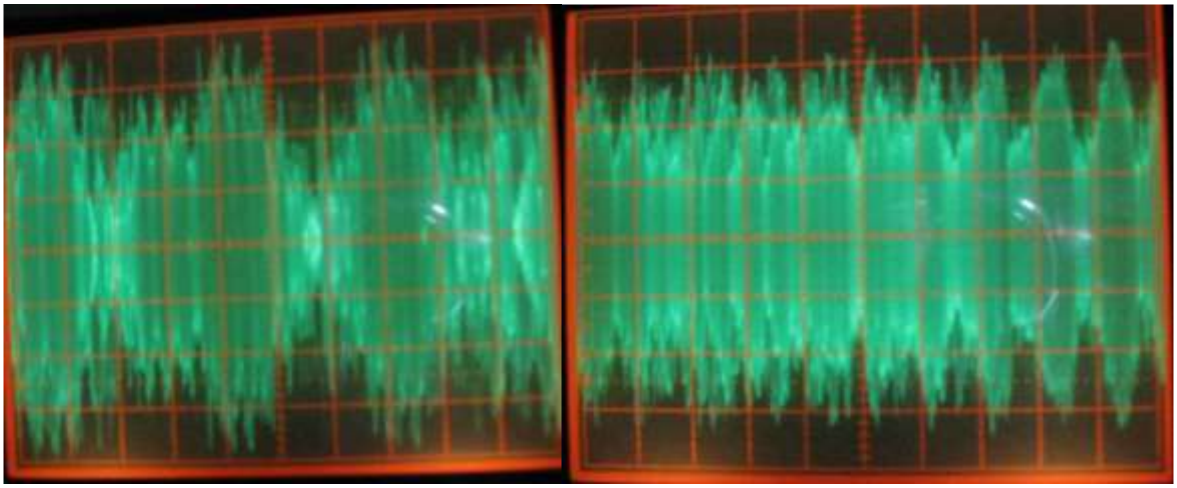

Trapezoid Technique

Another good way to quickly monitor the output is with a trapezoid pattern on your oscilloscope in ‘X-Y mode’. The x-axis input is connected to the audio program (or a function generator) which is also connected to J21, and the Y axis is connected to J4. The shape of the resulting trapezoid can be used to diagnose the modulation of the transmitter. (I refer the reader to any 1950-1980 ARRL handbook for a more detailed description of this technique) In the following photographs, the audio signal is a triangle wave. The first shot shows a moderate level of modulation. The sides of the trapezoid are straight, indicating a linear relationship between the audio signal and the RF envelope. The second photo shows excessive modulation causing the left end of the trapezoid to come to a point. This point indicates that the carrier is actually disappearing during the negative most extreme of the audio signal, and distortion will result. Decreasing the amplitude of the audio, or increasing the carrier level will correct this. The third photo shows excessive modulation causing RF clipping; the edge of the trapezoid curves at the right side. At the highest extremes of the audio excursion, the tube is running out of plate voltage; decreasing the carrier level or decreasing the audio amplitude will correct this.

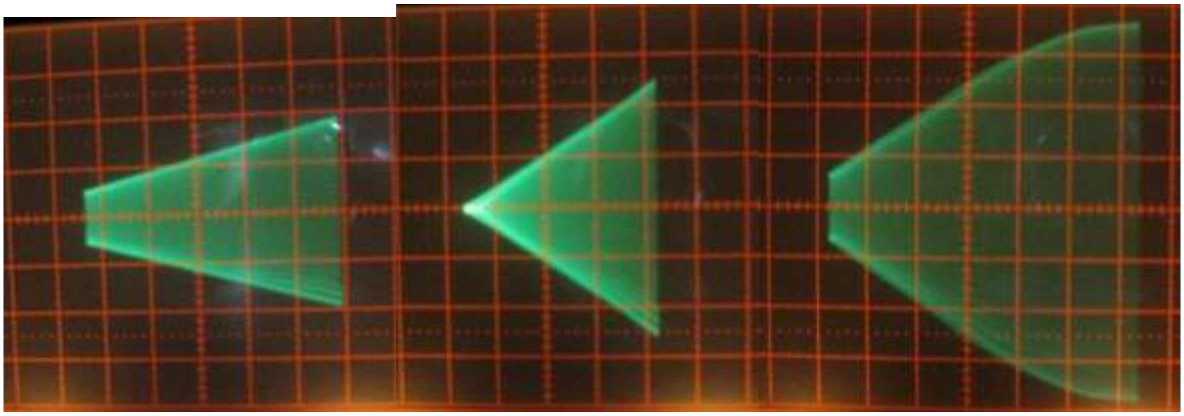

Some Measured Results

Fundamentally, the transmitter should translate cathode current (which is a result of the audio input and the carrier setting) linearly to an RF swing at the antenna. So, the most fundamental test we can perform on the transmitter is plotting the RF voltage at the antenna vs. DC cathode current. An extension of these tests would include the frequency response of the audio signal. I made many measurements of output RF voltage swing vs. DC cathode current, and plotted the results in a variety of ways. You would like the RF output swing to be exactly proportional to the cathode current, so the ratio of output swing to cathode current (which I call "gain") should be constant. I knew going into this that it is not possible because there will be leakage (RF swing without cathode current) and clipping (RF swing maxes out because of B+)

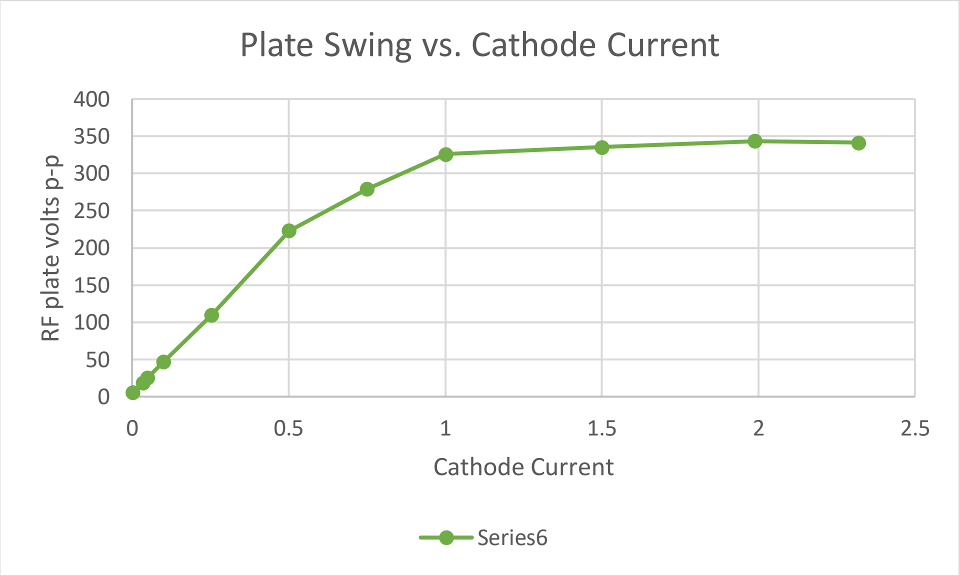

The above graph shows measured plate RF swing (p-p volts) vs. average cathode current. This is taken at 1.544MHz, with an antenna, with the high Q basket-wound coil. Note that at zero current, the RF swing is small (about 2Vp-p), indicating that feedthrough is small. As current increases, the curve starts to soften and eventually saturates (clips), because of supply voltage limitations. With a 200V plate supply, one cannot expect more than 400V p-p of RF swing. Plate swing was measured directly (thru the 1000:1 divider) with the HP RF voltmeter, but I do not trust the accuracy of the voltmeter over a wide range of powers. Most other tests are done with a well calibrated variable step attenuator to keep the power to the RF voltmeter constant.

There are better ways to display this information. Below are plots of gain vs. drive level. The first is linear to linear, with cathode current on the X-axis and gain (V of output swing per mA of cathode current, with some arbitrary scaling) on the Y-axis. A more conventional way to present this data is by looking at the output swing in dB on the X-axis and the gain in dB on the Y-axis.

These plots were taken at 1.544MHz with antenna connected, with the high Q inductor at L1 with various settings of the predistorter. Note that the predistorter can make the curve flatter (gain more constant), but it is a temperamental setting. However, it works well if you get the settings right.

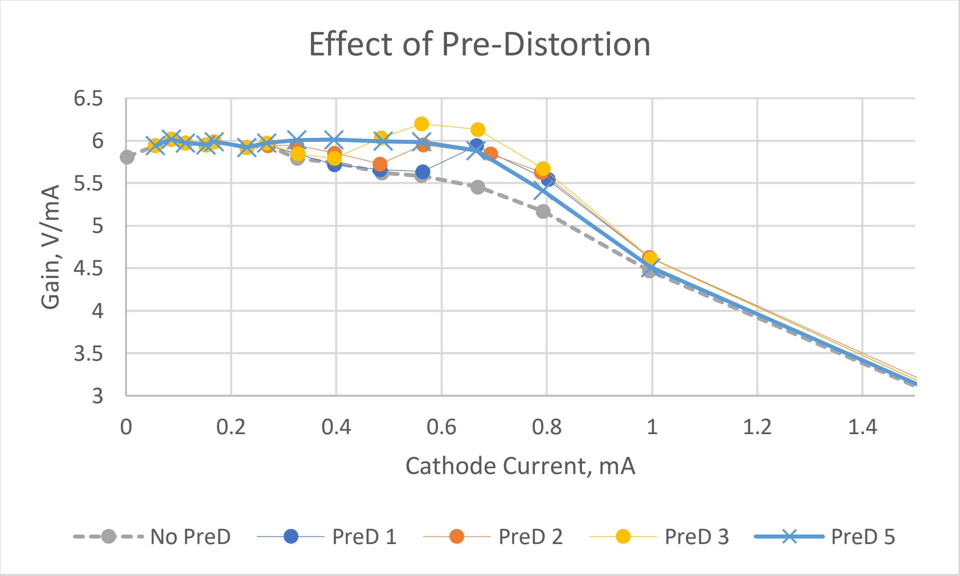

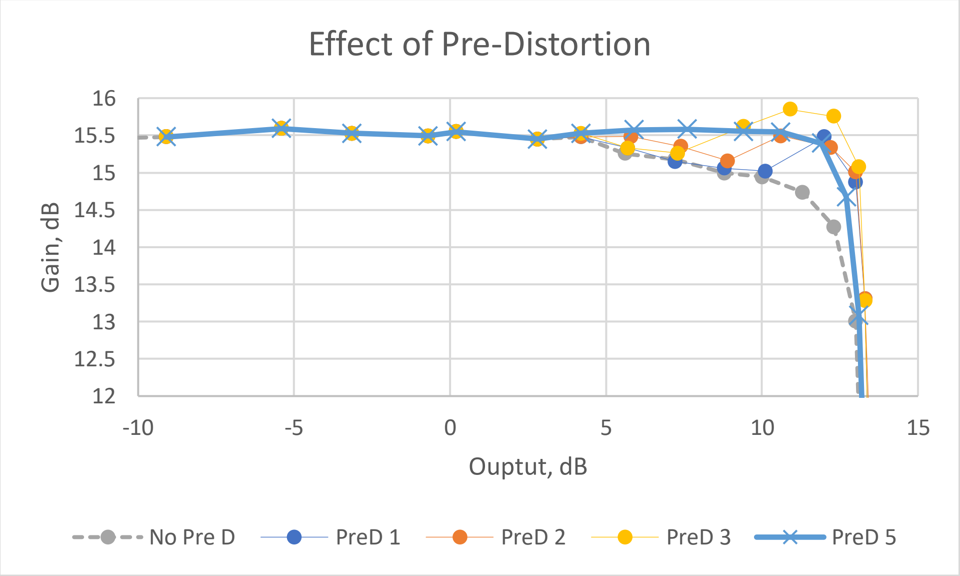

## Observations and Lessons Learned

I'm always a bit reluctant to describe my "Lessons learned", because some of them look painfully obvious in retrospect. So, bear with me as I demonstrate my naivety!

Does it work?

Yes, absolutely. I have a 1MHz crystal and a 1.554MHz crystal, and both can be used with my 10 foot antenna. At 1MHz, the antenna impedance too high (I could remedy this with a higher plate voltage) (This is done with my high Q litz wire coil). The 0.5mA DC quiescent cathode current (required to get 100mW power) is too high for linear operation. At 0.5mA, I have over 300V p-p at the plate (already starting to clip, it would have been almost 500V p-p if clipping could be eliminated). At 0.15mA carrier current, it's about right (150V p-p), but of course that is only 30mW of DC power, so I am not taking advantage of the maximum legal power. At 1.554MHz however, the antenna impedance is lower, so the transmitter really "wants" to operate at a higher current, and 0.4-0.5mA is about right. (This is done with the high Q coil, but tapped down to a lower inductance) At the 0.5mA limit for legal operation, the quiescent voltage swing (0.5mA) is about 220V p-p . So, you can develop substantial RF voltage with only 100mW of power!

Modulating this 0.5mA with music does what you would predict, you receive music on your AM receiver. Modulation depth needs to be moderate; as soon as I get close to full modulation, the audio sounds distorted. I am not certain whether this is caused by clipping, or the slight slope on the gain curves that I commented on before, or something else.

The “clipping” lights come on when clipping is quite severe. The audio gain needs to be well backed off from the point that the lights flash. The lower clipping light will come on earlier when operating at a lower current; it might be better to modify the design so that the threshold for the clipping light tracks quiescent current, as shown in the schematic below (which I have not tried).

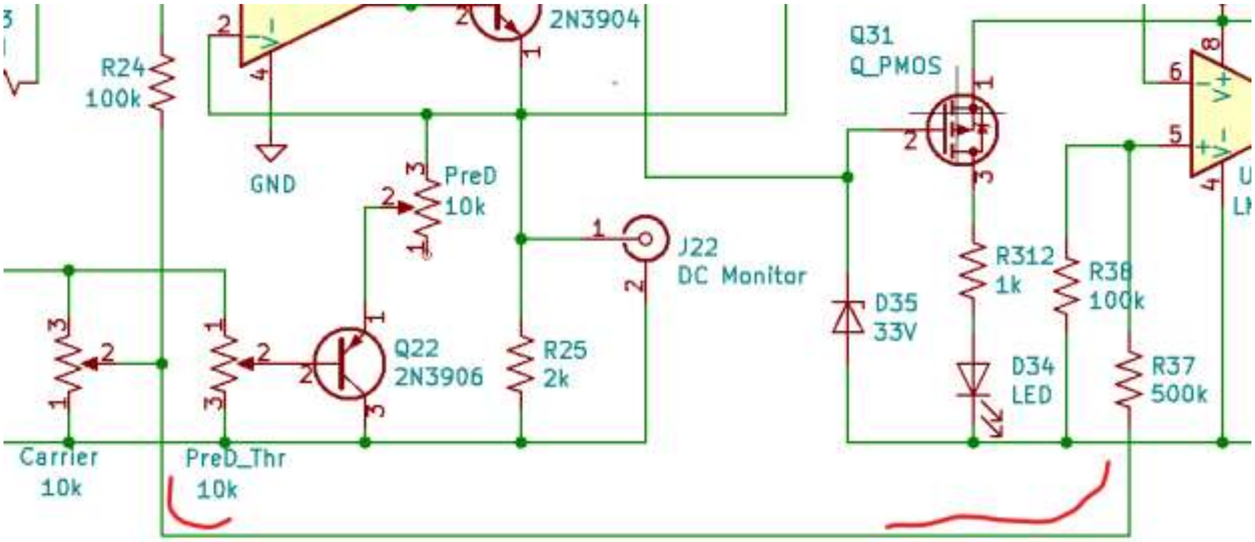

I did a listening test around the block, when operating at 0.5mA cathode current, 200V supply, high Q inductor in the tank, 1.544MHz, 10 foot basement horizontal antenna (so, just barely above ground level), noon, with music playing. I could get 2 houses away with a useful signal and 4 houses away before the signal was too noisy to be useful. That is about 40m for a good signal (with pockets of good signal strength beyond that), and about 60m to where the signal was disappearing into the noise. There was a barely discernable signal sporadically at 200m. The receiver was a run-of-the-mill 8-transistor pocket transistor radio, nothing remarkable. Daytime is the best time to conduct this test as there are fewer LEDs, dimmers and TVs causing interference, and at night I periodically had interference from other distant stations.

An inflexible design.

One disadvantage of this approach is that the transmitter is designed for a specific antenna radiation resistance. It can accommodate different antenna capacitances, but there is no impedance matching, so it is optimum for only one antenna radiation resistance. If the antenna radiation resistance is too high, then the transmitter will saturate prematurely. If the radiation resistance is too low, then the swing at the plate will not hit saturation (at the legal power limit), so efficiency suffers. Another unintended consequence is that the design will change if you want to run the transmitter at the low end or the high end of the AM band, because the antenna impedance changes across the band. It would help to make the power supply variable: higher voltage & lower current could be used for higher impedance (low band) loads, or lower voltage and higher current for lower impedances (high end of band) . I wanted to avoid an antenna tuner, because they can be lossy (especially because the impedances are so high, inductors will generally be working near their resonant frequency, not good for losses), but that may be the only other answer. Of course, the other solutions are to ignore the 100mW limit (frequently tacitly done) and/or de Q the antenna (also frequently tacitly done) .

## The severe limit imposed by a 3m antenna

Such a short antenna looks like a high “Q” capacitor, and you are trying to force as much power as you can into that capacitor’s loss. You can do this with “Brute Force”, just charge and discharge the antenna capacitance into an amplifier’s output impedance, but this will not be efficient. If you stay within the 100mw budget, 99% of the power is lost to the capacitor, only 1% is radiated (a “Q” of 100). (This is de-Qin the antenna) Or, you can do what I did and match out the capacitance, and dump the amplifier’s power into the antenna’s loss. BUT, since the antenna Q is so high, this is a very narrow band match, and the system can be easily detuned by a) a hand near the antenna, or antenna positioning b) a small change in frequency (making crystal or synthesizer frequency control a necessity), or c) small changes in component value. This is a fundamental limitation; you cannot get around this regardless of the match you use at the antenna. Either you sacrifice efficiency (effectively de-Q the antenna) or you live with a narrow temperamental match. (This becomes much less onerous as the antenna length increases or frequency increases) It also means that the audio frequency response will be limited by the narrow tuning. With a Q of (say) 100, the audio response of the transmitter will roll-off significantly by 5kHz. Broadcasters do not have these restrictions because their antennas are much longer and easier to match to. They still pay attention to losses, however, because of the enormous power they use.

## About supply voltage

A disadvantage of this architecture is that there is no variable matching (other than resonating out the capacitance) between the transmitter and the amplifier. Therefore, the transmitter (in particular the power supply) must be designed FOR the antenna. Not a great scenario unless you have intimate knowledge of the antenna, but, as I said, matching the antenna seems to result, ultimately, in a loosing proposition because of losses.

The optimum supply voltage will depend on the antenna used, and matching thereto. In the optimum scenario, the plate swing, which is seen at the antenna, at 100% modulation, will just be at the point of clipping. (so, with our 200VDC supply, that would be an RF voltage approaching 400v p-p) If you assume that the Q of the plate inductor has been optimized such that most of the available power goes into the antenna radiation resistance (which, if you use really good inductors, seems to be possible), then the swing will depend directly on that radiation resistance.

## That darn slope

Looking at the gain vs. cathode current curve, there is always a slight downward slope, most noticeable on linear plots. In the plots, the gain is simply RF output swing (volts) over DC cathode current (mA). Plotted in dB, (the way I always plotted power amplifier compression curves during my professional life), the slope is almost invisible, and it has what looks like an excellent curve. However, there is a slope, and I suspect that it is causing the audible distortion.

I do not know the cause of this slope. At the plate, there should be current pulses that are a constant width (by virtue of the CMOS pulses driving the gate) with a current proportional to the well-regulated cathode current. (as long as DC grid current is negligible) Constant width pulses from a constant impedance source should result in an RF swing exactly proportional to the pulse current. Errors could be caused by errors in the DC current (which could be caused by errors in the constant current source, say if it runs out of head room, or could be caused by grid current subtracting from the cathode current), or variations in the pulse width, or variations in the plate impedance of the tube. At saturation, one of these effects must occur.

## A Question of Fidelity

I admit that I was never quite fully satisfied with the audio fidelity. Obviously, an AM signal as received by any conventional radio will lack high frequency response, as IF amplifiers generally limit audio bandwidth. The high Q antenna match will also limit treble. However, I always felt that I should get a bit more high-frequency response. The bottleneck (aside from the output tank) may be the capacitor on the cathode. I originally had a capacitor that was 6X larger and reduced it to improve fidelity. In its current configuration, if you ignore the effects of the narrowband tuning at the plate, the frequency response is flat from under 20Hz to over 14kHz. However, a decent Q plate inductor drops the high frequency audio response, as low a 2kHz at the 3dB point.

I find that I need to keep the modulation depth very moderate (probably under 70%), otherwise the audio “sounds” slightly distorted. I find that the base “thumps” on the source material are especially distorted. Looking at the linearity curves for the tubes (most noticeable in linear scale), there is generally a slope to the gain vs. output power curve. This could result in distortion, but I cannot quantify how much distortion that curve would produce. Also, I do not know the root cause of the slope in the curves; I tried many things (grid drive levels and waveform, screen grid voltage (when I was using pentodes) degeneration resistor in the cathode, different tubes (really, all exhibited this slope over most of their operating range), predistortion (which helps a bit), but the slight slope persisted. I do not know if this slope is important, or if it is the cause of the distortion.

Predistortion should allow wider dynamic range. Predistortion works as it should, but like everything else in this transmitter, its set-up is critical. The threshold at which it should kick in depends on when clipping starts, and is therefore highly dependent on plate impedance. As you can see from the graphs on page 20, it can work quite well, but the audio quality does not seem to be noticeably better with predistortion turned on.

If you are very observant, you noticed the hysteresis visible in the trapezoidal curves. There edges of the trapezoid are marked by 2 lines, one representing the upward slope of the triangle wave, the other representing the downward slope. This effect gets worse at higher audio frequencies: below it is shown with a 2kHz triangle wave.

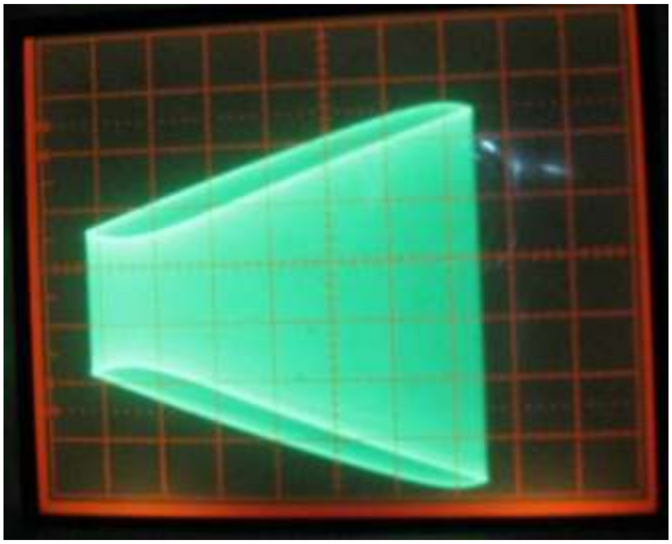

What causes this? The answer is that it is caused by the ringing in the plate tank circuit. It behaves like a delay. If you compare the triangle wave with the envelope, it becomes clear that the RF envelope is lagging the triangle audio excitation. The left waveform (2kHz triangle wave) shows that the peak of the envelope happens later than the peak of the audio signal. This is with just the high Q tank on the plate. If we de-Q the tank with a 10K parallel resistor, the second photo is the result, where the delay has pretty much disappeared.

It can also be seen on square wave audio drive: the three shots show a 2kHz square wave drive, first with low loss in the plate tank, second with a 47K shunt resistor, and third with a 10k shunt resistor. DE-Qing the tank clearly makes it respond faster.

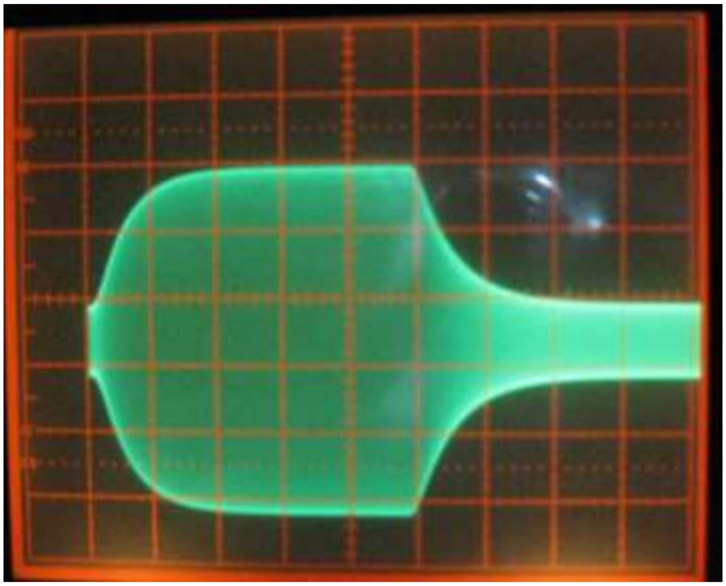

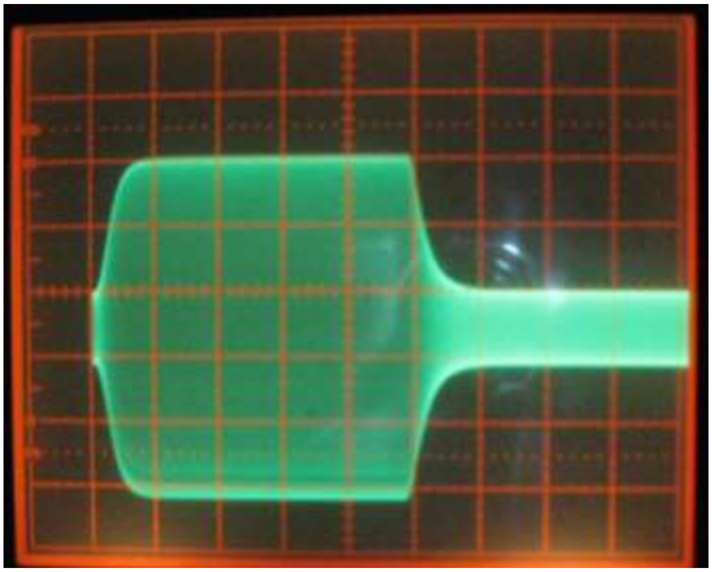

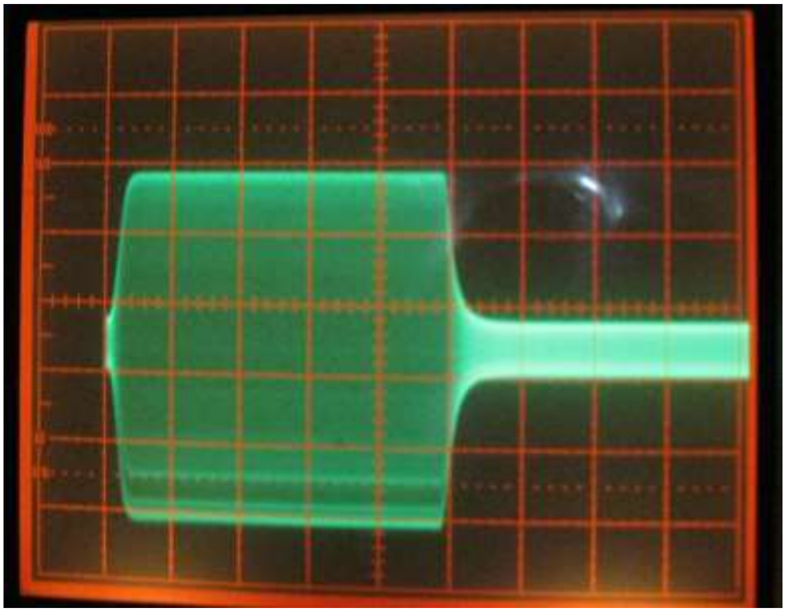

## Measuring Antenna Characteristics

I decided to carefully measure the impedance of my antenna, a 3 metre length of hook-up wire held with ceramic insulators on my basement ceiling. The antenna is kept as far away from everything as possible, because anything could add loss, and any movement causes variability. The high impedance of the antenna makes conventional instruments ineffective, so I developed my own technique, centered around the transmitter. I connect the antenna to the transmitter and tune the plate capacitance to resonance and note the voltage at resonance. I remove the antenna, and replace it with high Q mica capacitors, trying different values until voltage peak happens at the same tuning of the plate capacitance. I then try various R's (small quarter watt resistors) in parallel with the "C" until the RF voltage is the same as what I observe with the antenna connected. So, now I have a parallel R-C that behaves just like the antenna. That R-C is the equivalent circuit of my antenna at the test frequency. I discovered that my antenna looks roughly like a 30pF capacitor in parallel with a 560,000 ohm resistor at 1MHz. At 1.544MHz, it looks like 30pF with 180,000 ohm resistor in parallel. That's my antenna model.

Armed with that information, we can calculate how much RF voltage we would need to generate. The regulations specify 100mW DC into the final amplifier stage. Amplifiers are not perfect, and there will be losses elsewhere, so I assumed that we would get maybe 50mW at the antenna (3dB of loss). (that may seem like a lot of loss, but consider even just the tank inductor that matches out the capacitance. If its "Q" is similar to the "Q" of the antenna, which is about 106 at 1MHz, it will immediately burn half (3dB) of the power delivered to the antenna structure.) That works out to be 167V RMS (470V p-p) at 1MHz. Now, this is at average. Under AM modulation, half the time the voltage will be less, half the time it will be more, as much as double at 100% modulation. Clearly, we will need more plate voltage at 1MHz if we expect to be at the legal limit (and we already proved that). The situation is better at 1.544MHz. We would need 94V RMS (270V p-p) to get 50mW into the antenna, which seems more reasonable (we already proved that as well).

## Measurement of the Effective Output Resistance of the Transmitter

The output resistance of the transmitter is important. Ideally, we would like it to be much higher than that of the antenna. If it is not, then some of the RF power generated by the transmitter, maybe most of the power, will be dissipated in the transmitter (likely the choke) not the antenna. This could be a problem with ANY transmitter, not just this architecture.

I measure the transmitter impedance by looking at the output swing (using the 1000:1 port) with different resistors replacing the antenna. (Choke is still there, of course. We need DC) The effective resistance (seen by the plate) will be the parallel value of the known resistance and the unknown internal resistance. I then plot inverse resistance (the paralleled value) against the inverse voltage, and extrapolate to 0,0. (0 implies infinite resistance) (because we know that the swing would be infinite if there were no losses, that is, if the effective resistance was infinite) I adjust the unknown internal resistance until the curve extrapolates down to 0,0. In other words, the internal resistance is the resistance that, when added to the known resistances, results in a curve that extrapolates to 0,0.

With the plate load of three 100uH inductors in series (300uH), at 1MHz, (always tuned up resonant) the effective transmitter resistance is 140,000 ohms. In this case, I used quarter-watt resistors from 100k up to 4.7M. They had negligible capacitance (tuning did not change), and any inductance they might have will be lost when compared to their resistance.

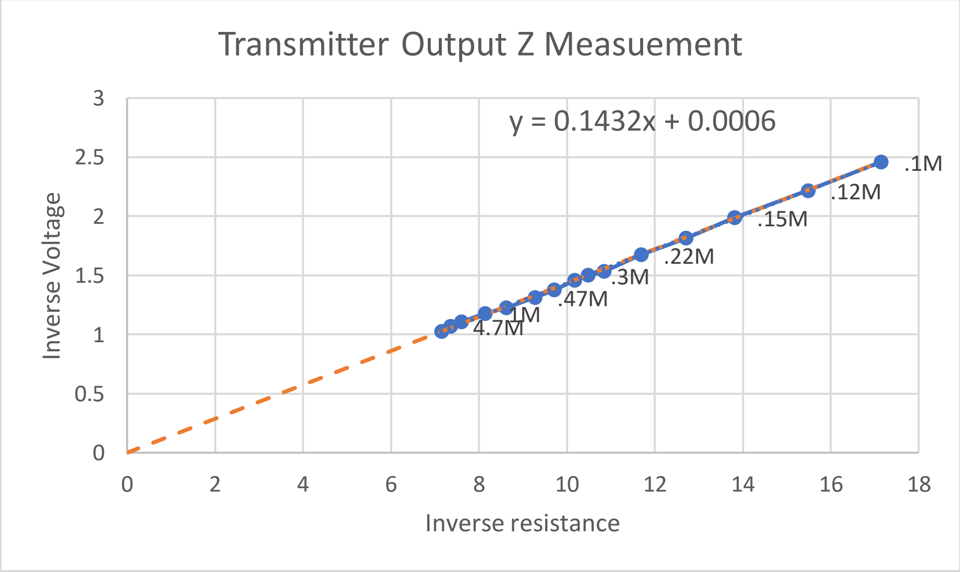

I measured the “Q” of the inductor (300uH) with my HP LCR meter; it said that the “Q” is about 200 at 1MHz. Honestly, I don’t believe it, because I measured the “Q” of my basket wound litz wire inductor at 1MHz, and it was “negative 160”. Of course, Q can not be negative for a real passive inductor, so my LCR meter seems to be optimistic for Q. We can safety say that the 200 value is a maximum, and it’s probably lower. If the inductor Q was responsible for the entire measured output resistance of 140000 ohms, then the inductor Q would be about 74, which I think is too low. That 140000 ohms will probably be dominated by the Q of the inductor, but the Q of everything else in the plate circuit, as well as the plate resistance of the tube will also contribute.

Assuming class “A” amplifier theory, the best efficiency we can expect (with perfect inductors) is 50% at saturated power. Furthermore, efficiency when modulated is even worse, say 25% at 50% modulation. In fact, we might do a little better by driving the tube hard, pushing it to class B or maybe even class C operation.

## Overloading nearby radios

At one point, I was trying to diagnose terrible audio distortion and I wasted a lot of time debugging the transmitter. Long story short, I was overloading the (high quality) AM radio that I was using for a monitor, causing it to distort. Any radio I tried (except for my “Miller” crystal radio) near the transmitter was distorted simply because the signal was simply too strong.

## Poor Man’s Synthesizer

You could easily make a “poor man’s synthesizer”. Run the crystal oscillator at 3.58MHz (use a common TV colour burst crystal), but add a CMOS 4017 configured as a divider (to divide by 3 to 10) by connecting an output to the reset pin, then use one of the harmonics. By selecting different combinations of divide ratios and harmonics, you can get 597, 716, 796, 895, 1023, 1074, 1193, 1343, 1432, 1534, and 1591 kHz , all as stable as a crystal.

## Conclusions

This was a great education for me, but I do not feel that the result is 100% successful.

On one hand, the transmitter works as intended, and it covers my house with lots of margin, while staying completely within the law.

However, it is fiddly to set-up, largely because of my goal to keep the transmitter efficient, and keep input power under 100mW. I can confidently say that *any* efficient design (that does not throw lots of power away) will be fiddly with this short antenna. However, further careful work, either with low-loss matching networks, or a variable plate supply voltage, would make this a more adaptable transmitter.

I believe that there are issues with cathode modulation that are affecting large signal fidelity. Fidelity is fine for voice, but there is noticeable distortion on music unless a low modulation index is used. I believe the cause of this is the slope on the gain vs. cathode current curve, but I do not know the cause of this slope, or how to fix it. Predistortion is an option, but it is of limited benefit, and adds complication to an already “fiddly” design.

If you want an AM transmitter to make programming available for your AM radios, I actually suggest that this may not be the design for you; there are simpler designs that apparently have better performance out there, but violate the letter of the law. However, it is a good platform for understanding more about transmitters and inherent limitations of short antennas.