HELICAL RESONATOR OSCILLATORS
by Harold Johnson W4ZCB

In the quest for the perfect radio, there's not a lot on the list of necessities that is more important than that of a quiet oscillator. There is absolutely no difference to the operator between a radio with a dirty oscillator with a band full of clean signals, and a clean oscillator with a band full of dirty signals. Worse, while a poor AGC system, or a radio with poor gain distribution is a bane to the owner, at least it stays with the receiver. An oscillator with phase noise or spurious signals is troublesome not only to the owner, but is troublesome on transmit to other operators who are unfortunate enough to share the same part of the spectrum. For the present, the *only* thing that spells relief from these problems of phase noise and spurious responses is that old elusive quality Q. This article describes how to get some.

For longer than I care to remember, I have been one leg of a triad designing and building a radio better than one we could buy...at any price. From the beginning, we were aware that we could obtain better selectivity, more sensitivity, more dynamic range and a higher intercept than anything available. We were just as certain that designing and building an oscillator that would do such a radio justice was going to be a daunting task. Indeed, the better the rest of the radio, the better the oscillator needs to be!

Analog variable frequency oscillators have phase noise and are prone to drift. Digital variable frequency oscillators get rid of that noise and drift at the expense of generating their own special kind of noise, spurious radiation. (Random close-in signals seldom more than 60-70 dB below the carrier.) Techniques to stabilize the analog oscillator, and to rid the digital generator of it's spurii are complex to say the least. It's enough to petition the FCC to segment the ham bands into channels and go back to crystal control...And why? Because the crystal has Q. It can stay put, has low phase noise, and doesn't generate more than one signal.

After trying a dozen oscillator circuits, from Vackars, to Colpitts to Hartleys, Seilers to Clapps and even K7HFD's[1] ingenious class C low-noise oscillator, it really dawned on me that it's not so much the circuit (though some are better than others) as it is the components and the way you choose to use them. What was needed was an oscillator with such high tank Q that it would swamp any other characteristics of the circuit.

Hewlett Packard and their pre-synthesis signal generators pointed the way. Their 8640, one of the quietest oscillators that they ever made is built around a gorgeous coaxial cavity that tunes 256 to 512 MHz and is subsequently divided down to frequencies below that in octave bands. I made several new friends on the Internet trying to find a few of those cavities, despite their being roughly twice the size of the rest of the radio. Finding them in short supply, learning of some of their mechanical shortcomings in doing so, and casting about for some other method of satisfying the high-Q requirement, I decided to give the helical resonator a try.

A helical resonator can be described as a transmission line tank, where the transmission line is wound in a helix so as to be of practical size, and encased inside a conductive enclosure. That enclosure may be round or square, and may be open or closed on the ends. The critical parameters of the resonator are the coil pitch, length, diameter, wire size and physical relationship to the outer shield. Zverev[2] has an entire chapter on helical resonators, with equations and graphs defining the optimum geometry for them. Quality factors for a well designed and constructed helical resonator run well above 1000 and far exceed that obtainable with ordinary inductors and capacitors. Hayward[3] also addresses the helical resonator in "Introduction to RF Design", Chapter 4.

A transmission line oscillator is nearly linear in tuning with capacity. A change from 10 to 20 pF in the tuning yields about the same frequency change as a change of from 50-60 pF. This very attractive feature eases the design of a loop filter for a PLL that covers the entire tuning range.

For a cylindrical helical resonator, Zverev has a graph indicating optimum unloaded Q for the device occurring when the diameter of the helix divided by the diameter of the enclosure is 0.53. This relationship is critical and the Q drops significantly with excursions from that ratio of as little as five percent. The length of the helix for best Q is also a function of it's diameter, and should be 1.5 times the diameter. The diameter of the enclosure should run somewhere from 1.6 to 2 times the diameter of the helix, and its length should be equal to the length of the helix plus 1/2 the diameter of the enclosure with the helix centered in that length. The physical length of the helix itself should be a quarter wavelength at the desired upper frequency limit, and in practice, due to fringing and self capacitance of the coil, will run about 5-6 percent less than a full quarter wave. In our case, where the resonator will be loaded by the oscillator circuitry, the foreshortening will be more on the order of 10 percent. On the premise that every equation will halve the readership of an article (S.W. Hawking), I leave the appropriate equations to the end of this article, when it will be too late not to read the rest of it.

.William Carver[4], W7AAZ, and another leg of this triad, designed a synchronous frequency divider (which, along with his very original phase locked loop ideas had better be the subject of another article!) which with three chips, will divide by anything even from 4 to 18. It doesn't take a lot of math to realize that with this power, an oscillator running from 75 to 105 MHz can be divided to cover all amateur bands (including WARC) from 1.8 to 30 MHz. I *needed* a helical resonator covering 75 to 105 MHz! Our transceiver was single conversion without the high IF that compromises general coverage receivers.

Getting started.

A visit to the local hardware store, located copper pipe and pipe caps of 1-1/2 inches diameter. They also had some 1 inch that was too small for my ham-handed machine capabilities, (at least until I really want one of these things running at a GHz!) and some 2 inch that I drooled over until I ran the equations to calculate the difference in Q for the difference in size and decided against it. It turned out that if I hack-sawed it myself, they would sell me less than a 20 foot length of the pipe. Two feet of pipe was enough for many resonator enclosures, and the pipe caps were only a few dollars more. I brought the materials home, ran the equations and found that for 105 MHz I needed a helix for an enclosure of that pipe size with a diameter of 0.8 inches, and length of 1.2 inches, and that the pipe should be 1.9 inches in length. At that length, the pipe caps themselves would be touching before they bottomed out on the pipe, so they needed to be shortened as well.

Chucking the pipe in the lathe, (You *GOTTA* have one of these! Well, you don't *gotta*, a friend with a lathe who owes you a *big* favor will do), turn and square a 1.9 inch length of copper pipe. Then chucking the end caps, turn off a half inch of their collar leaving about a half inch to be used for drilling and tapping mounting bolts. Placing the shortened end caps on the pipe, drill pilot holes radially in both the cap and the pipe, and a hole in one cap that is centered in the bottom of the cap for mounting of the helix support. Remove the caps, drill out the holes in the caps to pass a #8 bolt, and tap the pipe for short #8-32 bolts.

The equations say that for 105 MHz, the helix needs to have 9 turns and be 1.2 inches in length. The diameter of the wire for the helix has a second order effect on the Q of the finished product, and I recommend reading Charles Michaels[5] excellent article "Optimum wire size for RF coils". If you go through his math, you find that #12 copper wire is a good choice for the conductor for this particular helix. Go to the local Home Depot and buy 4 feet of #12 with ground house wiring Romex, and you have 12 feet of copper for your first and future helical resonators. This one is going to use about two feet of wire. Using a 3/4 inch diameter mandrel, grip one end of a piece of #12 in a vise, wrap the other end around the mandrel while walking slowly toward the vise and close wind the copper on the mandrel. If you wind it tight, it will spring to 0.8 inches in diameter. Too much or too little can still be corrected at this point by using a slightly smaller mandrel. (Read drill bit, deep socket of the proper diameter, two sockets in a row, whatever.) Later, spread the turns to make the first 9 turns 1.2 inches in length by slipping a small pencil or drill bit between the turns and rotating the coil around it. Cut the 9 turns with about a half inch left over which will be bent to make connection to the side of the outer copper pipe. Position the helix in the pipe, and bend the last 1/2 inch of the wire so that it extends outward from the helix and just touches the copper pipe when the rest of the helix is centered.

Things not to get caught doing.

I made my first helical resonator oscillator some 3 years ago. constructed along these lines. The helix was supported internally by soldering the ground end to the shell of the enclosure, and supporting the top by connecting to a ceramic standoff that ran through the center of the helix from the cold end to the top. In the interest of shielding this resonator from the AC mains fields, I encased the whole thing in a 2-1/2 inch steel pipe nipple, and machined down a pair of end caps for that. Pieces of UT-141 coax protruded from the end for electrical power, RF out and control voltages. With some trepidation, I took the thing with me to Manchester, England, where G3SBI[6] could measure it's phase noise. Phase noise was terrible. It was enough to get me to drop the project for nearly a year.

By then, when not much else was looking too hopeful, I decided to revisit the helical resonator oscillator. I built a second oscillator very similar to the one residing in England's dump. By now, I had developed a marginal method of measuring phase noise myself, and when I tested my new pride and joy, was horrified to find Colin's observations correct. Due to a terribly misplaced faith in the relative stability of a helix made of #10 wire, 9 turns in length and solidly supported on each end, I found that I had made a fairly good accelerometer. It would "read" the vibrations of my heat pump compressor from 20 feet and was totally unacceptable as an oscillator. You can probably make helical resonator filters with this technique, but don't try to make an oscillator this way. That helix needs to be supported along it's entire length.

The helicoidal way to go

First, you approach all your plastic suppliers and ask them if they carry polystyrene rod. If the supplier does *NOT* ask you which kind of poly you want, hang up and try again till you find one who knows the difference between polystyrene and polystyrene-acrylonitrile. You want the one with the unhyphenated name. The other, despite having an expansion coefficient of half the real stuff has a dissipation factor two orders of magnitude worse!

Go back to your friend with the lathe that owes you a favor, and get him to turn your polystyrene down to 0.8 inches in diameter. (After the copper, which is "sticky" and *very* difficult to machine, he'll bless you.) After that, have him turn threads on the outside of the rod to the pitch of your helix (You know, 9 turns over 1.2 inches is 7.5 turns per inch...call it 8, no sense going overboard on this!)

and to a depth of one half the diameter of the #12 wire. Cut off a length of the rod to the helix length plus 1/4 the diameter of your resonator shell. Drill and thread one end for a #8 bolt and drill the opposite end for the length of the helix to make a tube with perhaps 1/8 inch walls. (A half inch to 5/8 inch drill bit, depending on how good he is on the lathe) The idea is to reduce the strength of the polystyrene where it supports the helix because it has a greater coefficient of expansion than the helix does. If you run the oscillator free running, you might even slit the tube for the length of the helix to provide room for the poly to expand into. In the event you phase or frequency lock the resonator oscillator, this step is not necessary, and I would trade the temperature sensitivity for the mechanical stability of the solid tube.

When you finish, the helix should thread onto the polystyrene support rod fairly snugly. Thread it in from the free end of the helix so that when you screw the rod to the end cap, that end is the free end of the helix. The other end of the helix will be supported by being soldered to the shell of the resonator, and will, with the extra length of the poly rod, position the helix centrally in the length of the pipe when the pipe cap on which it is mounted is fully seated. Again, the free end of the polystyrene support is located at and terminates at the ground end of the helix.

Solder the end of the helix to the pipe wall where it is touching. I find this easiest to do by pre-tinning the wire with a heavy dollop of solder, cutting back the excess on the end so there is just a "collar" of solder on the end of the wire. Install the helix, supported by the poly rod with the pipe cap as a locator, and heating the outside of the pipe with a plumbers acetylene torch. If the pipe is clean and fluxed, the solder will flow off the wire and onto the pipe and make a clean joint. This is where the high current portion of the circuit is! Unless you can make a lower resistance mechanical connection than you can solder, solder it. If you intend to silver plate it, use the lead free tin/silver solder that plumbers use. I have mixed reviews on my silver plating, I have used "Cool Amp" plating powder, and silver cyanide and an electric current to plate these resonators. Both methods have occasionally yielded a lower Q factor after plating than I measured before plating. It makes it look pretty, but I have seen no conclusive results that indicate the superiority of plating over simply clean copper. Hayward claims unloaded Q improvements of 20 percent for a three stage resonator by simply plating the unplated copper helices even with an already plated enclosure. So it may simply mean that my measuring techniques are too coarse. The problem of finding the unloaded Q in a device with Q this high is not a trivial task. Zverev suggests estimation from the observed loaded Q and the insertion loss.

The downhill stretch

After mounting of the helix, it must be tapped to allow coupling to the oscillator. Using a large iron, quickly solder a lead to the cold end of the helix at a point approximately 1/2 turn up from ground. I bring this lead out to a teflon feedthrough pressed into the pipe at a position high enough on the shell to clear the end cap. If necessary, you can notch the end cap to clear the feedthrough. There's a lot of copper here, counting both end caps and the pipe, but it would be a mistake to consider it to be a uni-potential surface. I have had the greatest success in building the oscillator itself on a small piece of circuit board and physically mounting it to the resonator with the radial bolts holding the end cap in place. Almost anything else seems to be inferior, although I have made one oscillator with components remote to the resonator and fed through a reasonably short piece of RG-174 coax. Interestingly, there is not a lot of current flow at the pipe-pipe cap interface. The frequency of oscillation only changes a couple hundred kHz with the complete removal of the end caps although I carefully bolt them down to prevent mechanical movement or other external influences on the helix field.

Helical resonators are slow wave transmission line structures. Like their cousins, the transmission line resonator, they exhibit reentrant modes. A quarter wave resonator will exhibit a 3rd, 5th and other odd overtone modes. These can be triggered inadvertently, and getting rid of them can cause sleepless nights unless you are either an inveterate mathematician or develop a heuristic feel for what the mechanisms are.

Consider a helix of an arbitrary 90 degree length. This length will also exhibit a strong resonance as a helix of 270 degree length. (And others of higher order yet). In searching the cold end of the helix for a tap point of proper impedance for the oscillator, the obvious cold end is zero ohms or pretty close to it. As you move up the helix, the impedance rises, until at the top end of the coil it is extremely high. I have managed to get little RF burns off the top end of a helical resonator oscillator with the active device running at 15 volts! Somewhere between zero and high is the impedance level you're looking for. It should be obvious that on a line of 270 degrees in length, that a given short distance from ground yields a higher impedance than the same physical location does for a line of 90 electrical degrees. If your helical resonator oscillator oscillates on the third overtone, it means that your tap is too close to ground. You have plenty of coupling for overtone operation but not enough to achieve coupling to support fundamental operation. Move upward a small amount and retap and you cannot *make* it oscillate in the overtone mode. If you're actually trying to make an overtone resonator oscillator, simply reverse the above. You can trigger the 5th overtone oscillation in this manner although the tuning range on that overtone is very limited. The proper location for fundamental operation will occur with the tap from 5 to 7 percent from the bottom of the helix. The real goal is to maximize the signal-to-noise of the oscillator with maximum JFet power, with the impedance at the tap point roughly matching the V_{dd} and Id_{ss} limits for the device.

There is, in this particular tapped configuration one other possible pitfall. When you tap the helix to extract oscillator and or tuning connections, you basically have two tuned circuits in the container.

One, a tuned circuit loaded with the oscillator capacity and any tuning C between the tap point and ground, and the other, the rest of the helix above that tap point. One of my early helical resonator oscillators would tune very nicely down in frequency to a point where, if you were observing the waveform on a scope, the signal would actually Morph. The waveform would shrink slightly and then expand and come back as a twice frequency oscillation. That one will drive you up a tree until you have an AHA! in the middle of the night. The triggering mechanism for this occurs, when the tuning capacity and the inductance below the tap resonates, (or gets close to resonating) at twice the frequency of the fundamental mode of the helix. At this point it starts looking like an open circuit to the rest of the helix. The rest thinks it's a helix ungrounded at both ends and the oscillator sees a half wave helix resonance and takes off accordingly. This frequency represents the *LOWER* limit of the fundamental tuning range of the oscillator. No matter what you do, you don't get lower than this. Lower the tap point to get less inductance in the resonant circuit, and it just takes more capacity to get to that point. If you need to go lower in frequency, you make a helix that is calculated for a lower upper frequency limit and settle for the reduction in the top end. All of this seems more arcane than it really is. Just be aware of the possibilities and you quickly settle in on the correct parameters. With this oscillator, compared to the conventional linear transmission line oscillator you at least have one additional variable you can juggle to get the operation you desire, the tap point.

While you can use the oscillator circuit you prefer to drive this resonator, much can be said for the JFET in the Clapp configuration. The device is inherently quiet in it's own right *IF* you keep it out of gate current, it is easy to control the feedback, and AGC is readily applied to hold the drain current constant. In this application, with the drain at ground potential through the tap on the helix, a negative supply is required with an N channel FET, in my case not a consequential problem. If it is, you may either cast about for a P channel FET, or use one of the little DC-DC converters to obtain a few milliamperes of the proper polarity. If you go the latter route, be *certain* to heed the comment below about power supply cleanliness, for these converters are typically *not* quiet!

I initially tapped the helices with two taps, one for the feedback and one for frequency control using a varicap. Refinement brought the realization that the same tap could be used for both, as long as you do not push the power level to that point where the varicap becomes forward biased. Typically, with 10 volts on the oscillator, about +10 dBm may be obtained without encountering this barrier, more than enough to drive a subsequent buffer.

One last reminder of things to watch out for in building a low noise oscillator is the power supply for the oscillator. A perfect oscillator will not be perfect with a noisy source of power. A year or so ago, I ran across an article on the Wenzel Associates[6] web page called "Finessing power supply noise" Basically a shunt AC coupled regulator, it functions by shunting the noise to ground through a small series impedance. With no adjustment at all it is capable of reducing power supply noise by 20 dB.

With a bit of tweaking, that can be raised to a 40 dB improvement contributing directly to oscillator sideband cleanliness. The schematic for this circuit is included here by permission of Mr Charles Wenzel. Frankly, I don't build oscillators or frequency control circuits without it.

These are the experiences of building perhaps 3 dozen helical resonator oscillators. I've become pretty famous locally for eating up all the copper pipe caps the hardware store can keep in stock. What it will give you is an oscillator that *free running* at 100 MHz, when divided down to be the local oscillator for a HF sideband receiver, will copy a sideband round-table for an hour without being retuned. It is readily tuned three quarters of an octave or a shade more. Measured noise sidebands with a divisor of 8 (Each *synchronous* division by a factor of two provides close to a 6 dB improvement in phase noise) are better than -123 dBc/Hz at 1 kHz and drop to -144 dBc/Hz at 5 kHz. At 20 kHz spacing with divide by 8, those numbers measure -156 dBc/Hz and at 50 kHz they exceed my ability to measure them. (-162 dBc/Hz to -165 dBc/Hz depending on how I hold my tongue in my cheek on that particular day.) This is better at *ONE* kHz spacing from the carrier than a current commercial offering recently reviewed in one of our periodicals was at *ANY* spacing.

This paper describes the construction of a helical resonator VHF oscillator with many desirable features suitable for use as the local oscillator in an HF radio. It further explains some of the characteristics of such an oscillator and describes methods of controlling them.

[1] Hayward and DeMaw, *Solid State Design for the Radio Amateur* ARRL 1977. See the oscillator designed by Linley Gumm K7HFD, on Pg 126. Also in Chapter 4 of EMRFD.
[2] Zverev, Anatol I. *Handbook of filter synthesis* John Wiley and Sons. NY. 1967
[3] Wes Hayward W7ZOI, *Introduction to Radio Frequency Design* Prentice-Hall 1982
[4] William N Carver W7AAZ, 690 Mahard Drive, Twin Falls, ID 83301
[5] Charles Michaels W7XC, *Optimum wire size for RF Coils* QEX August 1987.
[6] Colin Horrabin G3SBI, Inventor of the "H"-mode mixer topology, third leg of this triad, and the one to blame for starting all this.
[7] Wenzel Associates, Charles Wenzel, Web Page HTTP://WWW.WENZEL.Com

Figures:

Equations:

Valid *ONLY* for the following condition:

$b/d = 1.5$

Boundary conditions:
$1.0 < b/d < 4.0$
$0.45 < d/D < 0.6$
$0.4 < d_0/Tau < 0.6$ at $b/d = 1.5$
$0.5 < d_0/Tau < 0.7$ at $b/d = 4.0$

Then, the number of turns (From Zverev) is approximately:
$N = [1700/(F_0)(D)(d/D)][\log(D/d)/1-(d/D)^2]^{1/2}$ turns
And the unloaded Q approximated by
$Qu = 60(D)(F_0)^{1/2}$

The characteristic impedance of the resonator is

$$Z_0 = 98000/F_0(D) \text{ ohms.}$$

Where F_0 is the upper resonant frequency, D is the diameter of the outside shell of the resonator, d_0 is the wire diameter of the helix, b is the length of the helix, d is the diameter of the helix, and Tau is the pitch or center to center distance between turns of the helix. All measurements in inches and MHz.

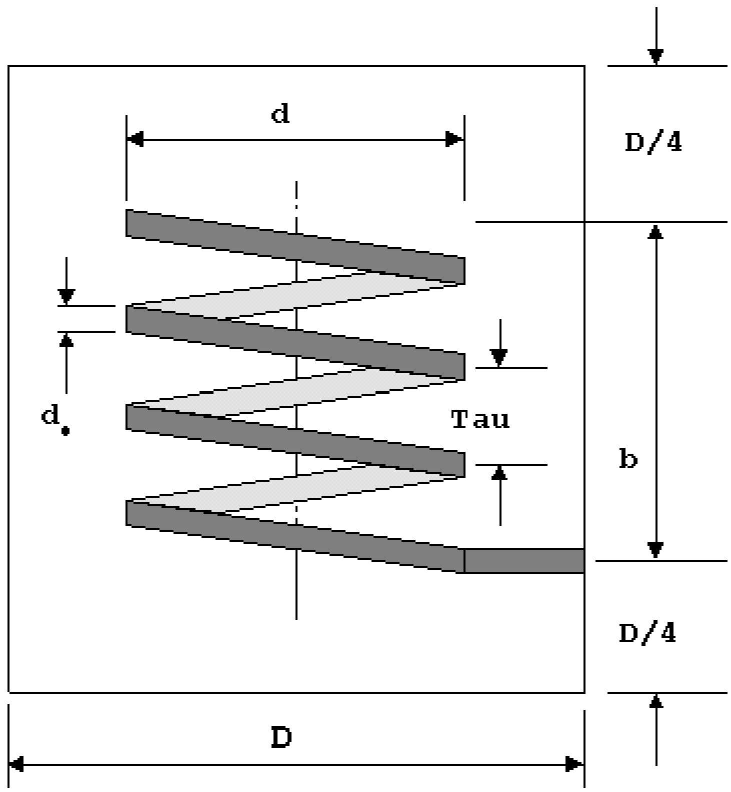
d
D/4
d_o
Tau
b
D/4
D

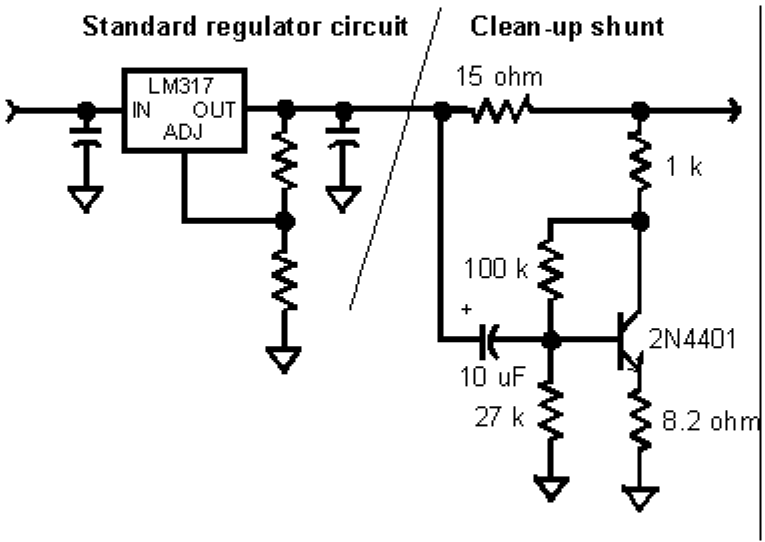
Standard regulator circuit
Clean-up shunt
LM317
IN
OUT
ADJ
15 ohm
1 k
100 k
+
2N4401
10 uF
27 k
8.2 ohm

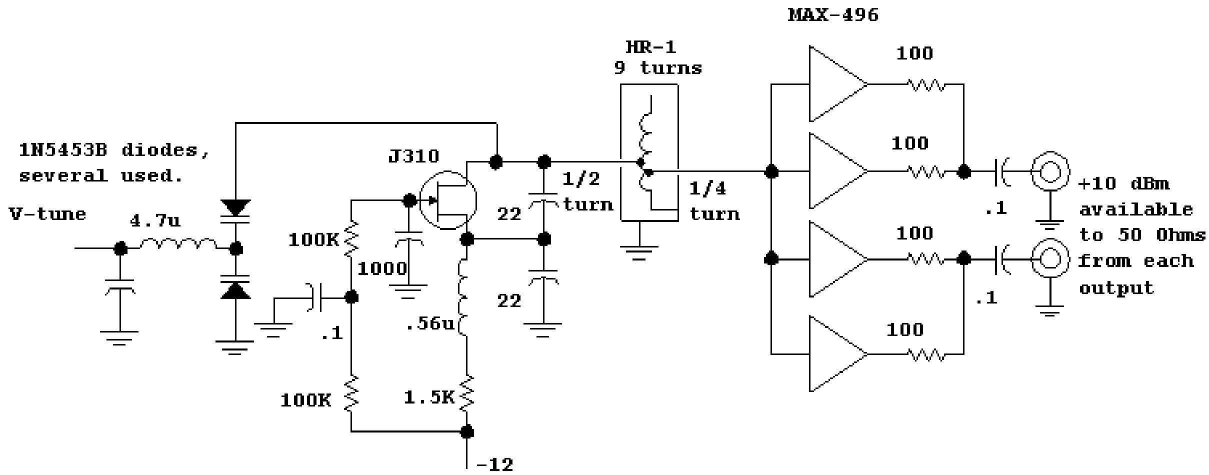

HR-1: Helical Resonator.