


Some Thoughts on E-field Whistler Receiver Design

Please note addendum

My introduction to natural radio listening came on a midsummer's eve this past July. The solder was likely still warm on my McGreevy BBB-4¹ as I listened in on the VLF cacophony accompanying the major magnetic storm that raged overhead. Nearly continuous multiple hop, diffuse whistlers mixed with chorusing and tweeks competed with the constant background crackle of strong sferics. I was so taken with the beautiful sounds that I began an immediate search for alternative receiver designs incorporating more elaborate filtering. Although hum free locations are plentiful here in Northern Vermont, they all require either a short hike or a long drive to access – not conducive to pre-work sunrise listening. After reviewing the various E-field receiver designs found on the web I settled on one in particular that offered more comprehensive filtering, in addition to higher gain, than the BBB-4. My first tests of this new receiver where disappointing, all signals where greatly attenuated (as compared with the BBB-4). Additionally, what seemed to be a RTTY signal was present, stronger than barely audible natural signals save the sferics. After unsuccessfully troubleshooting this receiver, I replaced its higher gain front end with the single FET BBB-4 front end – leaving all other filtering, pre and post, intact. Signal strength was depressed even further! But the RTTY signal was no longer present. Intrigued as to why this, apparently, more robust design would not outperform (or even equal) the BBB-4 I began to do a bit of research on E-field VLF/ELF receivers.



Per the literature^{2,3} an electrically short E-field antenna (in our case vanishingly short) may be modeled as a voltage source signal field strength times the equivalent height of the whip, in series with the reactance given by the isotropic capacitance of the whip, the radiation resistance and the loss resistance at the frequencies of interest. For our case I will consider the frequency range from 500 Hz to 20KHz and a 1-meter long vertical whip. Given these parameters, the radiation resistance is almost nil and the loss resistance so many orders of magnitude below the capacitive reactance that the antenna can be modeled as a capacitor in series with the signal source. The isotropic capacitance of a long thin antenna is approximately 10 pF per meter. Given the above criteria, this results in source impedance that ranges from a low of 800K Ω to a high of 32 M Ω across the frequency range of interest. To recover >90% of the signal available at the base of the whip, the input impedance of our receiver needs to be larger than the source impedance by a factor of 10. To prevent loading the 500 Hz signal present at the base of our 1-meter whip, the receivers input impedance needs to be in excess of 300 M Ω ! Receiver input resistance to ground neatly forms a high pass filter while any input capacitance to earth a straight capacitive voltage divider. Most receiver designs I have seen have a gate resistor no larger than 10 M Ω , many are smaller. More importantly, the input low pass or notch filters often contain 100's of pF in shunting capacitance. You can see where this is leading. A Pspice model of the BBB-4 predicts less than 10% of the available signal appears at the gate of the FET (with a 1-meter whip). In the more elaborate receiver I mentioned above this drops to <1%! In the common source amplifier configuration often encountered, the Miller effect complicates things further by multiplying the capacitance of the FET gate /drain junction roughly by the gain of that stage, further stomping our signal⁴.

To further explore this idea I constructed a “simulated antenna”^{2,3} to aid in bench testing and help give evaluations that are more quantitative and less subjective. This consists of nothing more than a small value capacitor housed within a driven shield to mitigate stray capacitance. I choose 12 pF to approximate a \approx 1-meter whip. Using a 10 mV RMS, 50 Ω signal source, I monitored the output of my BBB-4 under three different conditions – signal injected directly into the gate, signal injected prior to low pass filter (normal input) and signal injected at normal input with “simulated antenna” inline. The results are tabulated below:

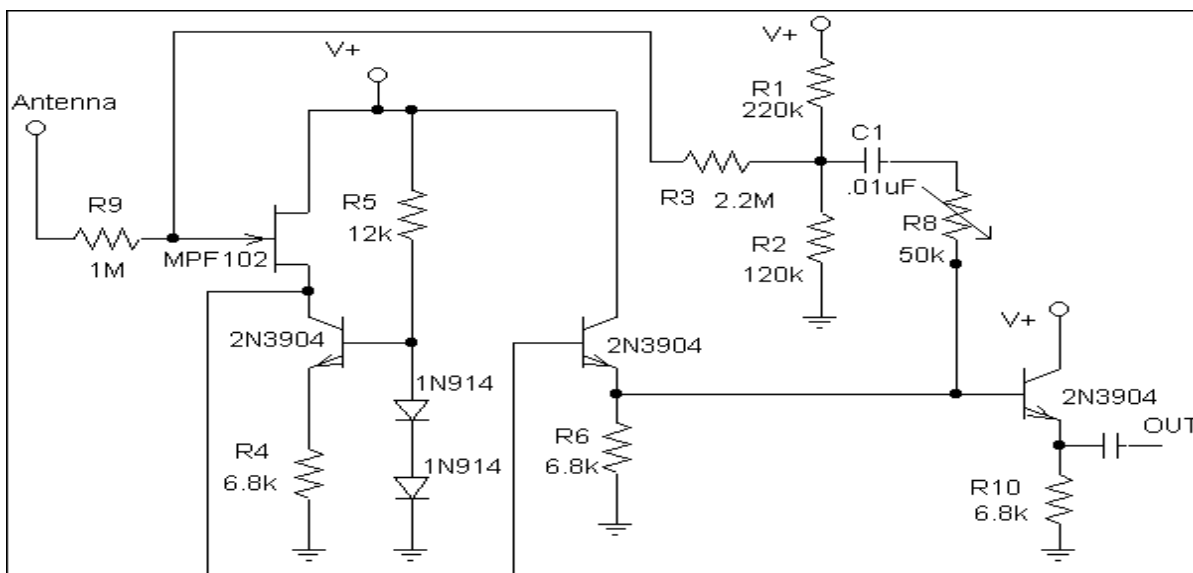
BBB-4 output given 10 mV RMS signal in

Frequency	Direct (mV RMS)	Normal Input (mV RMS)	Simulated Antenna (mV RMS)
50 Hz	5.3	4.7	0
100 Hz	12.6	11.6	0
200 Hz	37.6	33.3	2.5
400 Hz	123	108	9.45
800 Hz	354	306	30.4
1.6 kHz	804	630	70
3.2 kHz	1250	788	109
6.4 kHz	1120	448	86.7
12.8 kHz	624	135	35.5
25.6 kHz	261	28	8.9
51.2 kHz	78.9	4.6	2
102.4 kHz	17.1	1.4	0

At peak frequency response (≈ 3 kHz) my BBB-4 is capable of greater than 40 dB of actual gain. When sourced as it would be in normal use (with a short vertical whip) the effective gain drops to 20 dB at peak. The above is not meant to in any way denigrate Mr. McGreevy's elegant receiver, but rather to demonstrate that E-field natural radio receivers, designed to be used with short whips, must have extremely high input impedance. 300 M Ω would be a good value to target. Almost any form of input filtering will grossly violate this impedance requirement.

This is starting to sound more like the requirements of an electrometer than a radio receiver! A 300 M Ω input impedance at DC would indeed make a crude electrometer (typical electrometer input impedance is in the Teraohms). This could make the receiver sensitive to charges induced on the antenna by the operator as he / she moves around near it. That in turn may drive the input device out of its linear region and, perhaps, into saturation. Seems an impasse, too low an input impedance and signal is lost, too high and it becomes sensitive to quasi-static fields that we don't want to see (hear). Additionally, the Johnson noise of a 300 M Ω gate resistor would not be a pretty thing.

My solution to this conundrum has been to use a boot strapped source follower in front of any filtering or voltage gain stages. This allows all manner of filtering to be driven from the low impedance output of the follower **prior** to any voltage gain stages. The front end's dynamic range can be made much larger than is normally encountered in natural radio receivers. As the follower has no voltage gain, strong out of band energy does not drive the (follower) FET into a non-linear region – no mixing of NAA down into the audio pass band as occurred above. In fact, tight filtering can now be used to remove the majority of the unwanted crud before so much as a dB of voltage gain is used without killing the desired signal. The following circuit is a modified version of a source follower lifted from the pages of Nuts and Volts magazine. Testing it with the above mentioned simulated antenna yields >85% signal recovery, the 15% being lost to strays and junction capacitance I believe.



At DC, the input impedance is $\approx 3.3 \text{ M}\Omega$ - no problems with the “electrometer” effect. As frequency rises, increasing amounts of signal are fed back in phase to the gate resistor via C1 and R8, dramatically increasing the apparent resistance of R3 at these frequencies. The apparent value of R3 will become $R3_{\text{actual}}/(1-\text{gain})$. The 2N3904 active current source gives the follower a gain of better than .99 so the effective value of R3 can appear huge. This gives the effect we are after, an input impedance that is very high at the frequencies of interest but low enough at DC to prevent sensitivity to the motion of nearby charged objects (which is just about everything).

The only shunting capacitance appearing at the input is that of any strays occurring in the wiring to the gate and the internal junction capacitance of the FET itself. As the source “follows” the gate, the capacitance of this junction does not charge and is therefore not an issue. The drain is fixed, so the Miller effect is also not a problem. This holds the input capacitance of the FET to a couple of pF. There are techniques that will further reduce the apparent junction capacitance^{7,8} but I have found the residual capacitance to be more of an aid than a hindrance, as I will discuss later. R9 and the MPF102 should have the bare minimum of lead length needed to reach the input connector. I have taken to mounting mine “flying lead” style on the antenna connector, a technique used in ultra high impedance instrumentation (well, in my Keithley electrometer anyway). Alternatively, a piece of coax could be used but the shield must be driven by the output of the follower – not grounded. In short, whatever technique that works⁵ should be used to keep any stray capacitive coupling with the gate at a minimum – a few pF here and there is a killer when you are looking at the signal source through only 10 – 20 pF. This is, in my opinion, the major reason that E-field receivers are so sensitive to nearby trees and other vertical conductive items. The world is full of objects (including the ground) that are more than willing to be the “other plate” of a capacitive divider. Absorption of the wave front energy by these objects obviously also plays a role, but it is likely minor compared with loading effect.

Without R8 some peaking will occur just before the -3 dB frequency. This may or may not be desirable depending on your application. I included it, as my intentions were to have as flat a frequency response prior to the -3 dB point as possible. Its exact value will vary depending on the size of your whip. You can use a signal generator and a “simulated antenna” sized to approximate your whip to set it up, or the low tech, but very effective, cut and try and listen method. Alternatively, it can be left out altogether; the peaking may be desirable. C1 should be between .001 and .1 μF , this gives a -3 dB point that varies from $\approx 1.5 \text{ kHz}$ down to near 300 Hz. Be aware that using the larger values also brings up any hum nicely. Nothing besides care in avoiding stray capacitive coupling to the input gate is at all critical about this circuit. All parts are available at Radio Shack and it will happily run on anything from 6 – 18V.

I live within 1 mile of 2 powerful AM transmitters and as such was pleased to find that auto-rectification was almost non-existent with this circuit, but it was still detectable way down in there. Adding R9 forms a low pass filter with the residual junction capacitance of the FET, neatly taking care of my BCB interference problems. If strong medium wave signals are not an issue, eliminate R9 as it just provides more opportunity for stray shunting capacitance to sneak in.

I have had much success with this little circuit, both as an impedance converter for the front of my BBB-4 and as a front end for my ever-changing home brew receiver. There are obviously other ways to accomplish the same thing, an op-amp will work fine in the boot strapped follower configuration (but much worse auto-rectification in my experience). Should you choose this route, Burr Brown manufactures low input capacitance (1 pF) op amps. Look for electrometer grade devices such as the OPA129. Alternatively the antenna can just be made larger, increasing its isotropic capacitance and lowering its source impedance. This will inevitably lead to long antennas if the rule of keeping the source impedance to 1/10 of the receiver impedance is held. If the input contains just a few hundred pF of shunting capacitance, the antenna must then have a few thousand pF – that’s BIG.

I want to mention that Helliwell⁶ reports medium latitude whistler field strengths that range from 4 mVmeter⁻¹ to 5 μVmeter^{-1} (unfortunately he doesn’t give the distribution). Although not the local BCB station, neither are these really weak signal strengths. I would propose that the idea that hunting whistlers is weak signal work has grown out of the fact that many receiver designs are disposing of the majority of the energy prior to amplification, giving the impression that the signals we are chasing are very weak. I don’t think this is necessarily the case, we just have to keep from squashing them down to nothing prior to trying to detect them.

In conclusion I want to say that I have no particular expertise in this area and am completely prepared to be wrong on all of this. I would welcome any correspondence on the matter and would particularly like to hear from anyone that tries this idea out.

Scott Fusare N2BJW
sfusare@adelphia.net

¹ Steven McGreevy, *McGreevy BBB-4 Natural Radio Receiver*, <http://www.triax.com/vlfradio/bbb4b.htm>

² Toshimi Okada and Akira Iwai, *Natural VLF Radio Waves*, Research Studies Press 1988

³ Arthur Watt, *VLF Radio Engineering*, Pergamon Press 1967

⁴ Paul Horowitz and Winfield Hill, *The Art of Electronics*, Cambridge University Press 1989

⁵ Keithley Instruments, Inc., *Low Level Measurements 5th Edition*

⁶ Robert A. Helliwell, *Whistlers and Related Ionospheric Phenomena*, Stanford University Press 1965

⁷ Larry K. Baxter, *Capacitive Sensors Offer Numerous Advantages*,

<http://www.planetee.com/planetee/servlet/DisplayDocument?ArticleID=1843>

⁸ U.S. Patent #4390852, *Buffer Amplifier*, <http://www.uspto.gov/patft/index.html>

Addendum 14.05.01

Correspondence received after this paper was published has brought to light a major flaw with the boot strapped follower approach. Dave Ewer has kindly shared the results of his extensive field-testing and careful side by side comparisons of different front end topologies. His work shows the boot strapped follower to have a poor noise floor as compared with the common source front ends he normally uses. After taking more care in my own testing (simultaneous comparisons were not originally done), as well as educating myself better on the issue of low noise design, this is what I have come up with. Yes, the bootstrapped follower has wonderfully high input impedance at the signal frequencies of interest, recovering nearly the entire signal. Unfortunately the gain peaking mentioned in the text also applies to the thermal noise generated by the gate resistor. This results in a noise gain peak near 1 kHz. Making matters worse, the phase shift network used to smooth the signal response has no effect on the noise gain. The final nail in the coffin is my attempt to hold shunting capacitance to an absolute minimum, which, quite unintentionally, raises the corner frequency of the Johnson noise contribution from the gate resistor. All of this results in a noise floor that is an order of magnitude (or more depending on the approach) above that encountered when moderate amounts of shunting capacitance are used with no boot strapping. The paper is otherwise correct as far as I know. My goal of recovering as much signal as possible, while not paying attention to noise, is however a grievous error in the signal to noise department which is always the final arbiter. The circuit presented should perform better if the feedback is removed (R8 and C1), the gate resistor value raised to 22M and some 20 to 30 pF of shunting capacitance added to the input (after R9). Additionally R9 should be only as large as is needed to control interference or, if possible, eliminated entirely. Finally, both emitter follower 2N3904s should have a 220 ohm resistor inserted into the base lead to prevent oscillation.

On another topic, I have had some correspondence in which the observation was made that terrible broadcast band and short-wave interference was experienced when using this circuit. Upon further inquiry it became apparent that the output of the follower was being run into a high gain audio amp with no filtering. I must emphasize that the follower, as presented, is meant only as a front end device – not as a complete receiver. Bandpass filtering for the frequencies of interest must be used prior to any substantial voltage gain. The response of this circuit, especially without R9, extends well into the HF bands and perhaps higher. It will happily pass along powerful AM and short-wave signals that will end up being demodulated somewhere along the line only to completely obliterate the natural radio signals we are after.