

Notes on Improving Station Noise Performance

N6GN

xxxx 2024	0.5
5 Jan 2024	0.4
5 Dec 2021	0.3
30 Nov 2021	0.2
1 st Draft	0.1

One of the great values of the K1JT digital modes and of WSPR in particular is the increased insight into receive station performance they provide. The continuous, world-wide and multi-frequency reports and posts into the databases combined with the [wsprdaemon](#) and [ClickHouse](#) tools provide the ability to monitor and compare results versus changes in an individual receive sites configuration as well as among many other sites worldwide. While at any instant there may be local variations, considering the data over significant time intervals, days to years, allows one to come to a very useful assessment of how well a system is doing.

In comparing contemporaneous WSPR spots among similarly located stations and recognizing that for ionospherically propagated signals one may expect incoming signal strengths to be similar over time among all stations in the same general area, perhaps separated by several hundred km, one can learn a great deal by comparing the signal-to-noise-ratio (SNR) at similar or identical decoders. Differences in reported SNR and spot numbers then become an indicator of differences in the individual noise floors and ultimately to receive system performance.

That the impinging fields from a DX noise source would be similar within a locality is not too hard to reconcile with common thought. What is perhaps more difficult to integrate into one's thinking is that the SNR on any non-directive antenna from LF through HF is approximately the same at a given time. There tends to be a mindset that considers "bigger is better" while in terms of SNR, as opposed to absolute signal level delivered to a receiver, any antenna size whether a full half-wave dipole (or quarter wave monopole) down to very small probes dipoles, loops and whips actually have nearly identical SNR. Large antenna size does not change the antenna SNR due to DX signal and DX (inverse-square propagated) noise levels for antennas located region and having the same polarization.

To understand this, it helps if one recognizes that energy delivered from any antenna is entirely from an equivalent "radiation resistance". This is the real part of the impedance presented by any antenna structure and can be the only source of power coupled to a receiving system since no real power can be coupled from the reactive portion of an antenna's impedance. For antennas of any size, from extremely small up to about one-tenth of a wavelength a good approximation of the antenna's impedance is:

$$\text{Impedance} = R_a + jX_a = 20 \left(\frac{\pi L}{\lambda} \right)^2 - j \frac{120 \lambda}{\pi L} \left(\ln \left(\frac{L}{2a} \right) - 1 \right)$$

where L = conductor length is significantly less than $\lambda/4$, a = conductor radius

In the equation above only the R_a portion of the impedance is useful for delivering power to any receiver. The antenna's aperture and therefore the intercepted power varies only slightly. A short dipole

has about 1.76 dB gain compared to a half-wave dipole having 2.13 dB gain. Beam pattern and ‘capture area’ are very nearly constant.

As the equation describes it, the resistive portion of the impedance varies as the square of antenna dimension, L. Electrically small antennas, those with small maximum dimension, intercept equal powers but produce it from within a comparatively small source resistance, Ra. If this is difficult to accept, it may be useful to consider the situation when a small antenna is transmitting and using the reciprocity theorem to reconcile it.

Electrically Small Transmit Antennas

To help reconcile the above premise that the power received by any electrically small antenna is constant irrespective of the physical size and to overcome a standing common wisdom to the contrary it may be helpful to think about what happens with a transmitting antenna when driven by a perfectly matched source of power. Then by applying the [Lorentz Reciprocity](#) theorem the equivalent situation when that antenna is receiving can be recognized.

The real part of an antenna’s impedance/admittance is the only mechanism through which real power can be coupled. Any complex part, the reactance can not source or sink real power. This “radiation resistance” is not a real resistor that can convert RF power to heat, rather it is the mechanism by which waves in the far field (DX) can be converted to and from current and voltage at a transmitter or receiver. In a very real sense, this is the actual antenna while the part we see and often call the antenna is actually only a matching network. For more detail see the 2012 QEX article [A New Antenna Model](#).

Energy is conserved. If a one watt transmitter were perfectly matched to an electrically small antenna then all the transmitter power would be converted to radiation. There is no other place for it to go since there would be no reflection or conversion to heat or to another form of energy. The previous equation shows that although the value of this resistance changes it affects only the voltage/current ratio and not the fact that a watt coupled to this Ra creates one watt of radiated power. For a small antenna the power is spread out in three dimensional space in a “doughnut” pattern but if one were to collect all the power incident on the surface of a sphere located in the far field, energy is conserved so the original one watt would be recovered.

The pattern of any electrically small antenna is constant so the gain is as well and is about 1.76 dBi, only .5 dB less than that for a half-wave dipole. Completely tied to this is the antenna’s aperture. Aperture or “capture area” is roughly the area over which it “catches signal”. For a short dipole this is

$$\left(\frac{3}{2}\right) \frac{\lambda^2}{4\pi} \approx .12 \text{ wavelength}^2$$

This is an electrical size not a physical size! What one *sees* as the antenna is different from what the antenna actually is and does. For antennas smaller than about a half-wavelength for a dipole or a quarter wave-length for a monopole, it *doesn't matter* what physical size it is, the pattern, aperture, gain, capture area and intercepted power are the same. On transmit any such antenna will create the same (DX) result and if perfectly matched all the available transmitter power will go toward creating identical ERP.

Turning things around and considering the reciprocity theorem, it can be recognized that the situation must be the same when that antenna is used for receiving. If instead of power going away from the source (transmitter) it is arriving due to identical conditions at the distant sphere, except for the direction of power flow, then all the power in the aperture will “collapse” into the radiation resistance IF it is perfectly matched.

Size doesn't fundamentally matter for an electrically small antenna. But it does matter very much in a practical sense because minuscule radiation resistance associated with large reactance becomes extremely high-Q. This can be impossible to match with real materials, at least short of super-conduction. Even if it could be matched well, the operational bandwidth eventually becomes too narrow to be of use for communications.

Perhaps one of the reasons for the perception that SNR changes with antenna size is due to confusion around this R_a and the antenna's impedance measured at an available feed-point. For a half-wave dipole where the antenna is almost entirely resistive, the two values may be similar. A half-wave dipole in free space exhibits about 70 ohms and near resonance there is little to no reactive component. For this situation the two numbers will be about the same. A monopole over an extremely good ground or radial system may also measure near an expected 35 ohms. But as the antenna gets electrically and perhaps physically small, the measured value of typical antennas may still show a moderately high real component of 20-50 ohms, when significant matching such as a series 'loading' coil has been applied to bring the reactive portion to near zero for coupling to a, say, 50 ohm transmitter, while the radiation resistance actually has fallen precipitously. This difference between measurement and theory can tend to mask the fact that for real antennas and even ones over rather good ground (image) planes or radial systems what is actually being measured is the combination of ground resistance and equivalent series resistance of matching network(s) along with the antenna's radiation resistance, not the radiation resistance alone.

Once this misconception is recognized it becomes easier to accept that the actual SNR due to DX signals and propagated noise produced in the radiation resistance does not change. The intercepted power does not change with antenna size either – in the real world it simply gets increasingly difficult and finally impossible to efficiently couple to a minuscule radiation resistance in the presence of high reactance and ground plus matching losses by using available conductors and dielectrics. Necessary Q's for matching networks become so high they are impossible to achieve even with the best of components and for monopoles ground system losses can make efficient matching impossible.

As an example, a 10 m high monopole installed over a typical earth ground/radial system and fed with a suitable series inductor might produce a narrow resonance and reasonably good match to a 50 ohm transmitter on 630 m but most of the transmitter power will be dissipated in ground and matching structure losses. The radiation resistance of the actual monopole is on the order of a few 10's of milliohms at 630m compared to a 50 ohm system impedance. Although the antenna is theoretically capable of sourcing within about .5 dB of the same power that a full half-wave dipole might, in actuality only a very small fraction of that, perhaps 25-30 dB down from the well-matched value, ever gets radiated and arrives at a receiver. Small and 'full size' dipoles have similar capture areas, similar apertures, but it is impossible to efficiently couple to the smaller antennas. This same inefficiency is present for both transmitting and receiving but high levels of propagated noise may make the difference in recovered SNR less noticeable.

Considered as a receive antenna, the misconception that size matters may be further strengthened because of the nature of propagated noise and its change as a function of wavelength. The ITU noise data show an approximate -25 dB/decade negative slope in propagated noise. For a small antenna, which every amateur antenna is at some suitably low wavelength, this slope works in parallel with a positive, +20 dB/decade slope in voltage level due to the constant power sourced by the radiation resistance. Thus the composite interpretation may be something like “everything must be OK, noise is increasing somewhat at longer wavelengths, as expected”. What may not be recognized is that were the radiation resistance efficiently matched that increase would be much greater than is observed as would the recovered signal amplitude.

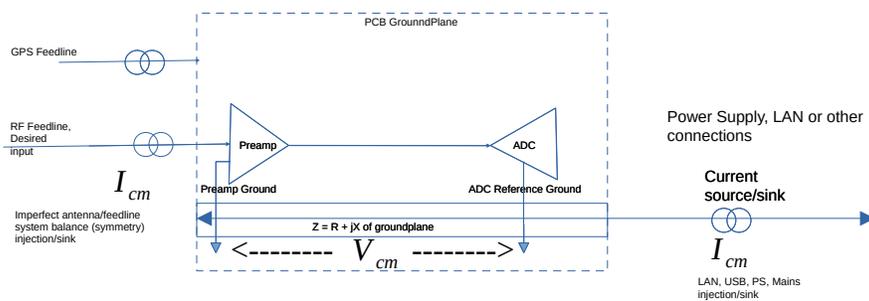
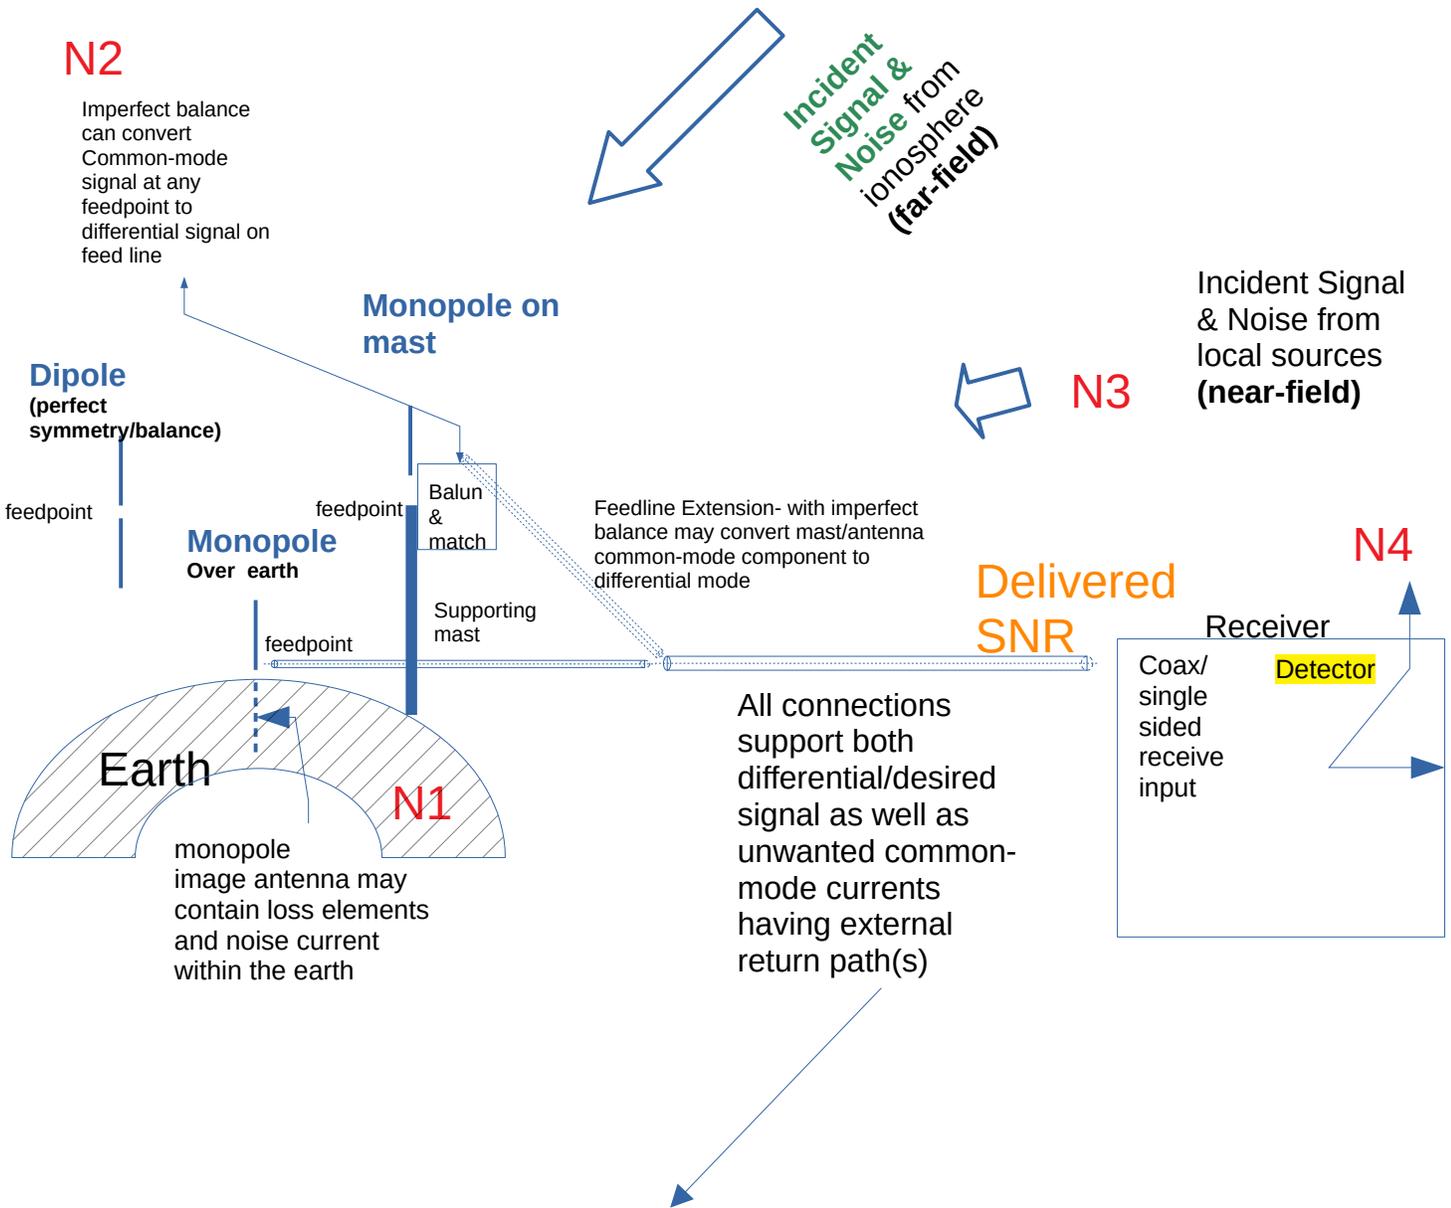
This overall situation can mask the realities of the problem of reproducing the available antenna SNR at a receive system detector. It is this problem which produces widely varying SNR results as reported by the WSPR database as well as by the large differences between “good” amateur stations and more typical ones regardless of mode. The result is commonly explained away with statements such as “It’s all those digital devices we have nowadays and there’s nothing to be done about it.” when in fact very substantial improvements in performance may be possible.

Once the problem is understood as one of reproducing the SNR present within an antenna’s radiation resistance at a receiver detector, whether WSPR, SSB or something else, and once the relative levels involved are appreciated and compared to mechanisms and levels of unwanted sources which can greatly degrade the noise floor an effective course of action may be begun that can greatly improve the situation.

Using the WSPR database to examine total number of spots, unique spots as well as SNRs of spots common with nearby ‘good’ stations, with a little practice one may directly get a moderately accurate estimate of system noise floor and thereby the amount of degradation due to unwanted factors. Only a few dB of change in system noise and SNR can easily make many-ten’s of percent difference in spot count as well as greatly improved DX reception. It’s not uncommon to see WSPR reporters that clearly have 20 dB or more degradation across many or all LF-HF amateur bands. These impairments can and do affect all amateur operations, not just WSPR and digital modes.

A completely different method for assessing station performance is also possible. This is to observe the average noise floor level and character over the course of a day or two. If system noise is predominantly established and limited by propagated noise then the observed diurnal variation in level will be smooth and obvious. 10 or more dB of change on an HF band below the maximum usable frequency (MUF) is common for a system which is working well and is propagation noise limited. If this diurnal variation is not visible or is suppressed, or if there are sudden steps in local noise floor, it is a very good indication that the receive system is being undesirably compromised. Fortunately, there is likely much that can be done to improve the situation once it has been recognized.

Having recognized the problem as one of reproducing the SNR present within an antenna’s radiation resistance SNR at a detector, the situation may be examined by modeling some of the noise impairment coupling mechanisms and sources as below:



$R = \text{milliohms}$
 $X_l = \omega * \text{length} * 32 \text{ nH/inch} \gg \text{milliohms}$
 $V_{cm} = I_{cm} Z_l$ in series with input RF signals

Conversion from Common Mode (CM) current being sourced from either antenna or PS/USB/LAN side **through** the receiver may be converted to differential signal and noise **across** that receivers input. It may require only a few microamperes of CM current to impair the noise floor of a sensitive receiver!

As depicted, an incident DX SNR can be compromised by local ingress by multiple mechanisms, here shown as N1-N4. Even if the entire signal available from an antennas actual radiation resistance is efficiently matched to a receiver, the total noise floor may be raised relative to the desired incoming, far-field noise. The result of this compromise is reduced delivered SNR and poorer system performance relative to what is possible.

Four local noise ingress mechanisms are highlighted:

1. N1

- The earth under an antenna is far from a perfect conductor. Particularly in residential or business environments, significant current may be flowing between end mains users, even though a designated return path to a common ground reference is provided by the utility company.
- The earth itself is lossy. This means that it has an associated noise temperature and radiates. Particularly for horizontally polarized antennas (not depicted) near typical earth this may cause an extra 7-10 dB of loss due to absorption which may be particularly significant when propagated noise levels have been attenuated or are mismatched. This appears as a source of noise and cause of reduced SNR.

2. N2

- Practical baluns and in particular common broadband ferrite baluns, no matter the architecture or core material, may only have a few dB of rejection of balance so degradation from common mode to single ended conversion may still occur. When considered relative to the desired propagated signal and noise levels, this performance may fall very far short of what is required to keep local ingress from compromising system performance.
- When a dipole or other symmetric antenna is used, the connections may be shorted without simultaneously shorting and removing the influence of the common-mode ingress, thus providing a method to assess the degree of compromise. While with a monopole or other unbalanced antenna, shorting the feed-point also shorts and removes the common-mode component so no such assessment can be made. Note that the antenna is 'grounded' by a feed line even if not by a conductive mast. Noise current can flow in this conductor.
- The matching structure(s) may include components commonly thought to be "the antenna". Wire or other conductors may be all be considered part of the matching structure, along with any L or C components that might be placed between the actual source of signal (near the tips for an electrically small antenna) and the feed line. Mismatch loss in these areas serves to attenuate the desired propagated signal and causes the relative influence of all the unwanted sources to increase.

3. N3

- This is an actual local and at longer wavelengths almost always near-field noise coupled source which may commonly but erroneously be considered to have far-field and inverse-square amplitude-vs-distance characteristic. Far field distance is generally considered to be $2D^2/\lambda$ where D is no less than about $\lambda/2$ even for small probes. Because this noise is near-field, slight changes in position may quickly and greatly mitigate it.

Once all other mitigation methods have been exhausted, moving the antenna and, as a last resort, quenching a particular source, may be the only remaining candidates for system improvement. At shorter wavelengths, far-field radiation is possible but it is surprisingly

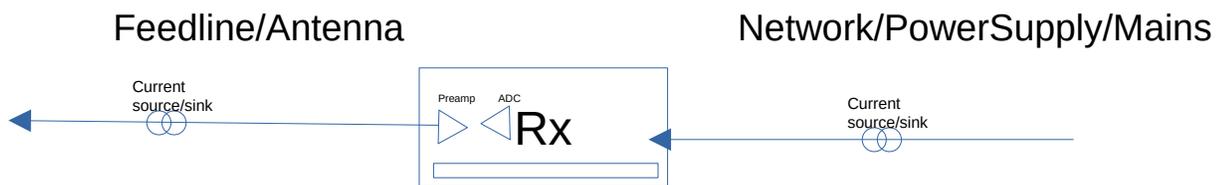
uncommon below VHF. True radiated far-field noise also has much wider-area consequences and can be recognized by distant monitors as well as local ones.

4. N4

- This source is like N3 in that it is local and near-field but it is coupled by conductors rather than radiated. The actual noise source may be more distant but may produce a common mode current that travels *through* the receiver to produce voltage drop across even the small impedance of ground and power planes within a receiver.

These currents may be sourced by LAN and power supply connections and may appear superimposed on the desired input. In this situation, moving the receiver or antenna has little effect. Also it should be recognized that the source impedance for these noise currents may be considerably higher than the [50 ohm] antenna system impedance so common mitigating measures such as ferrite chokes may have a bigger burden and may not produce even good limited results. Breaking or increasing the impedance of conductive paths reduces this current but may be difficult to achieve since power supplies and wired LAN or USB connections prevent this. Using a wireless interface instead of a LAN connection may help considerably.

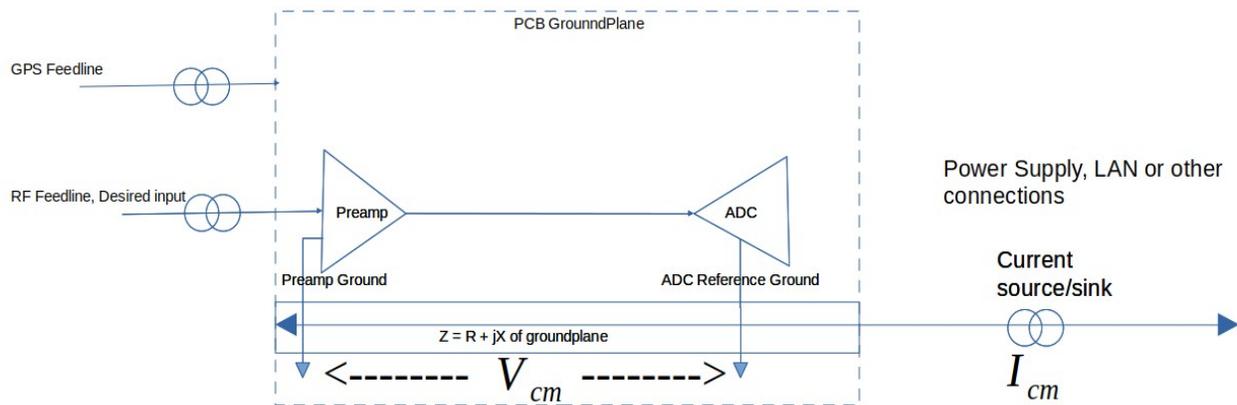
The KiwiSDR provides a practical example of noise type N4. Using the KiwiSDR as an example is not meant to imply that its performance is necessarily worse than other SDRs it's simply one that is easy to measure.



Common mode current sourced (or sunk) on the network or power side travels through the “ground” of the the Rx and out whatever impedance is presented by the ‘end fed antenna’ created by the feedline and intended antenna, *even though there is no physical ground return*. In our universe any line $>\lambda/2$ long has an impedance <377 ohms. Such a line may couple to radiation resistance at its end which causes the impedance presented at the Rx end to swing wildly but it has low values at odd quarter wavelengths (at least). In practice it never shows infinite impedance and can become similar to the current source impedance which may be a “disconnected line”. No physical “ground return” is required.

It's worth considering what happens inside the receiver and the coupling factor between CM current through the receiver from one end and out the other and differential voltage at the input. Of interest is not just the voltage developed across any preamp input but actually, for an SDR, in the voltage from the ADC reference to the input. Single ended input receivers like the Kiwi and most others convert some portion of any common mode current to a voltage, $V_{cm} = I_{cm}Z_{rx}$, that appears across the differential input and produces unwanted output. This conversion is produced by the non-zero impedance in the ground planes and traces of the printed circuit board and separation of the the ADC and input components.

For a KiwiSDR, a 50 ohm source connected to a 50 ohm load with the Kiwi inserted in the path produces unwanted QRN about 80 dB down across much of HF and only ~100 dB at LF.



$$R = \text{milliohms}$$

$$X_l = \omega * \text{length} * 32 \text{ nH / inch} \gg \text{milliohms}$$

$$V_{cm} = I_{cm} Z_l \text{ in series with input RF signals}$$

Common Mode current through the "ground" produces unwanted voltage at the preamp input. It should be noted that the Feedlines need not be 'connected' to anything, e.g. a 5 m long GPS cable appears as approximately a 14 MHz monopole antenna with a feedpoint impedance of ~35 ohms so is an excellent current sink for a higher impedance current being sourced from the power supply or LAN side.

Reduction of CM current through the PCB may best be achieved by shunting CM currents around the board and choking them from going through it by way of isolating flux-coupled transformers (Ant input) and/or small value coupling capacitors (GPS input).

It takes very little current to develop enough V to generate a signal that reaches the noise floor in a WSPR bandwidth of a modern SDR. The common mode impedance across the length of a KiwiSDR PCB is on the order of $Z = R + jX$ where R is ~4 mill-ohms and at 14 MHz the inductance due to the length of the PCB is more than 120 nH. At 14 MHz this results in an impedance of about 10 ohms. A 50 ohm signal generator set to -74 dBm passing its output through the groundlane to a 50 ohm termination can result in an increase in the noise floor and reduced SNR of the desired input. This level is less than a *microampere* of CM current!

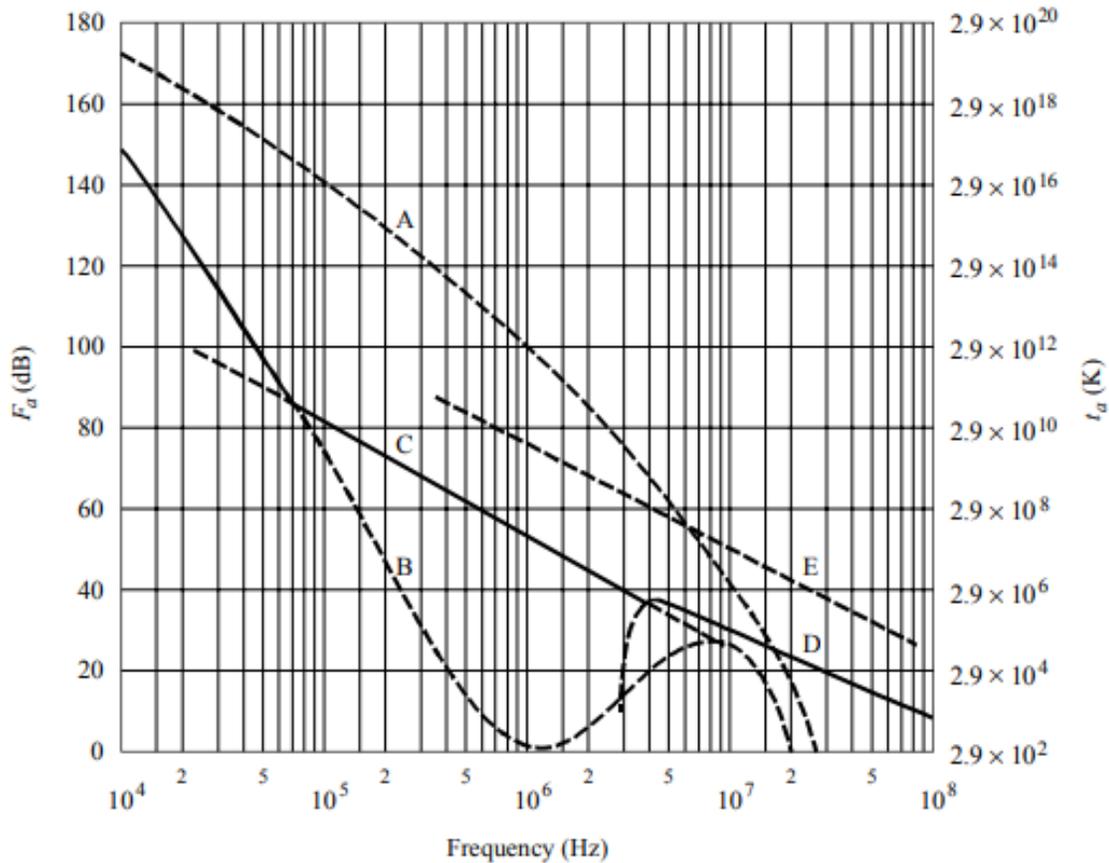
This model and example shows that improving a receive system may be a tractable problem but also that it isn't simple. Methods must be found to identify and reduce or eliminate the contributors such that they no longer cause severe degradation to the SNR available on the intended antenna. Generally it is desirable to reduce them 10 dB or more below the desired noise floor so that propagated noise establishes the noise floor at the receiver detector.

Carefully constructed probe antennas, either dipoles or loops, paired with highly symmetric preamplifiers which have very high common mode rejection ratio (CMRR) may be useful tools in this overall endeavor. If the coupling mechanisms can be identified and made to produce insignificant unwanted impairment then the propagated SNR may be detected and so produce

the best communications possible. Quenching recognized impairment sources should be attempted only as a last resort since it is much more effective to remove the coupling mechanisms.

Rec. ITU-R P.372-8

FIGURE 2
 F_n versus frequency (10^4 to 10^8 Hz)



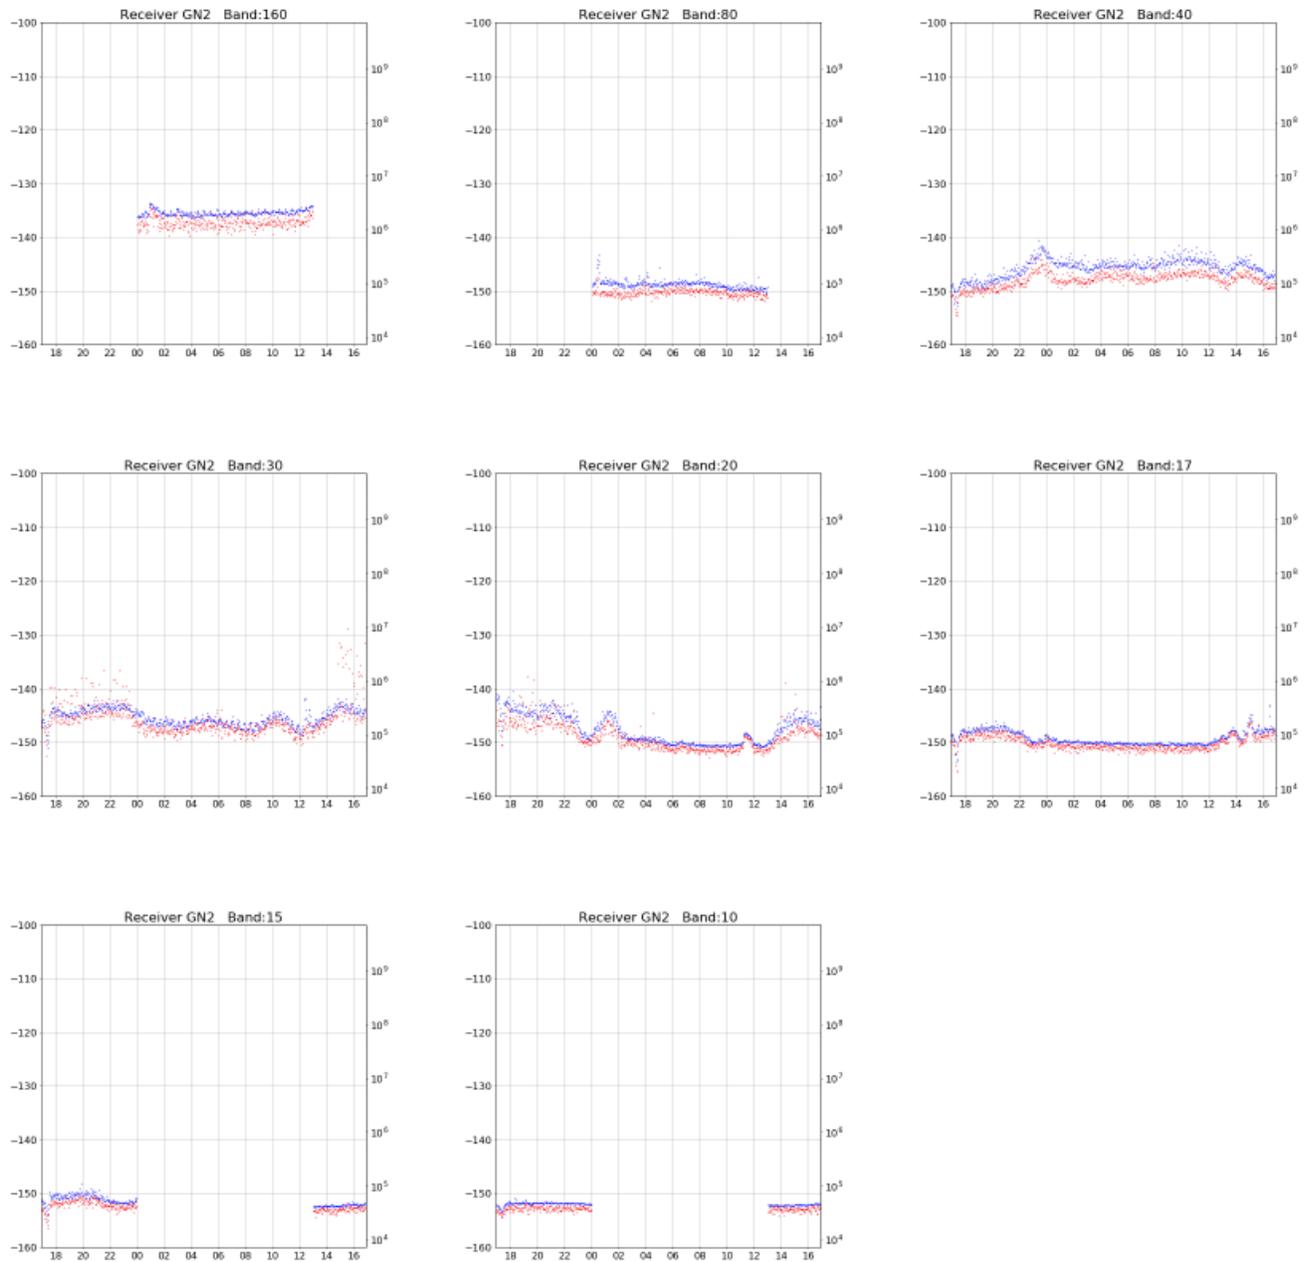
- A: atmospheric noise, value exceeded 0.5% of time
- B: atmospheric noise, value exceeded 99.5% of time
- C: man-made noise, quiet receiving site
- D: galactic noise
- E: median business area man-made noise
- minimum noise level expected

0372-02

Curve C of this ITU plot provides a reasonable guideline for the target noise floor in an antennas radiation resistance.

Methods for Assessing Station Performance

Besides careful calculation of the expected $KTB + F_a$ noise floor due to power in the radiation resistance and comparison of it with measured performance another method to get an approximate idea of station performance is to use the diurnal propagated noise variation as a crude “Y-factor system noise figure meter”. When a system is truly mostly limited by propagated noise, the ITU-R P.372 data, perhaps as averaged by curve C (above), combined with calculated KTB in the radiation resistance and mapped via circuit analysis to level at the detector can provide a guideline for the target level that is to dominate a system but simply watching daily variation can be very helpful too. It should be recognized that this target noise level is reduced by mismatch loss between the radiation resistance and the receiver input so make the target more difficult to achieve.



Example of diurnal noise variation on various amateur bands being used as indication of the margin between the effective system noise floor and the propagated noise. 10 dB is generally enough to reduce the system noise degradation of the SNR to about .5 dB.

A well-performing receive site will show smooth fluctuation in levels, not constant or discontinuous stepped ones as propagated noise varies and as the MUF moves and the example above demonstrates on some but not all of the HF bands. The min/max ratio of these variations can give one an impression of the 'headroom' a system has, how close non-propagated interfering noise is limiting performance.

The wsprdaemon.org site and tools are really helpful for this. At present no meaning can be attached to comparing absolute levels among multiple systems. If collectively we arrive at a good reference, KTB in radiation resistance plus the Fa value from ITU is my present candidate, calculation of this and adding it to the wsprdaemon.sh/python tools would be useful but is probably beyond the skill of many

operators at present. In the meantime only the shape and diurnal characteristics are very useful. It may be useful to match the indicated absolute noise levels of two stations being compared during times when DX propagation is absent due to absorption or too-low MUF.

WSPR itself along with the nice tools provided by <http://wspr.rocks/> and the ClickHouse database can be extremely useful. Along with an MUF map and comparison with local good stations, over time a rather good and representative station assessment may be made, one that does a good job predicting performance against other 'good' stations.

Summary

The goal of noise reduction is to substantially achieve the SNR present within the radiation resistance of a local 'antenna' from DX/ionospherically propagate signals at the receiver detector. Any practical antenna has approximately the same SNR on LF through most of HF when located at the same position and polarization. Achieving reduction then primarily involves analyzing and what those levels should be, mitigating detrimental forces and eventually verifying that the results are confirmatory.

For the most part, it is local mechanisms within control of the user, shown as N's in the drawing, that are impediments to this overall goal. Analyzing them, measuring them and removing the coupling mechanisms rather than their sources, per se, is the desired method. Environments change and only as a very last resort does quenching individual sources come into play. Preventing access is by far the better way.

This can become a quantifiable RF engineering problem with lots of pieces to keep track of and monitor. But the results have shown themselves to be achievable and very worthwhile. The WSPR database and tools help greatly toward this and then, after there is substantial success, provide a way to study propagation, the ionosphere and generally have more fun in amateur radio.