Gwyn Griffiths, G3ZIL

14 Bassett Close,
Southampton SO16 7PE, UK;
gwyn@autonomousanalytics.com

Rob Robinett, AI6VN

357 Gravatt Drive,
Berkeley, CA 94705, USA;
rob@robinett.us

Glenn Elmore, N6GN

7319 Silvermoon Lane,
Fort Collins, CO 80525, USA;
n6gn@sonic.net

Estimating LF – HF Band Noise While Acquiring WSPR Spots

Systematic measurements of noise at your receiver can help improve your station performance and provide insights into the ionosphere and propagation.

Noise levels are a topic of long-standing concern to radio amateurs. Over recent years a perceived upward trend in noise levels has added to that concern. This has led to a number of initiatives by national amateur radio societies including in the UK [1], the Netherlands [2], South Africa [3] and Germany [4], and ARRL have recently laid the groundwork for a noise study in a rural environment. Other spectrum users are also affected; for over 30 years the US Government has supported a *Signal-to-Noise Enhancement Program* to measure, identify and mitigate noise and interference. There is an example in the public domain of the methods used and results obtained during a 2010 investigation at a Department of Defense receiver facility at Key West, Florida [5].

Our own motivation to measure noise came from our use of the Weak Signal Propagation Reporter (WSPR) mode. While WSPR reception records include an estimate of the signal to noise ratio (SNR) they do not include an estimate of the noise. Attempts at gleaning reliable information on, for example, the hour-by-hour or day-to-day changes in signal level for a particular path can be hampered by changes in the noise level. This was the challenge communicated to us by Larry Plummer, W6LVP, as part of his study of his propagation path to Danie Mans, ZS3D. It was also our expectation that systematic noise measurements would help us understand and then minimize the noise at our own locations and within our own receiving systems.

We were also keen to make measurements at a site with very low local noise levels that one of us (AI6VN) has access to. The Maritime Radio Historical Society (MRHS)

preserves the ex-RCA coast radio receiving station KPH at Point Reyes, California [6]. We considered that this remote site, within the Point Reyes National Seashore, with proper care, could act as a reference station not subject to local noise or interference. This, indeed, proved to be the case.

Our approach has been to add two noise estimation algorithms to a script, *wsprdaemon*, a robust WSPR decoding and reporting program written by Robert Robinett, AI6VN [7]. Plots of noise from several stations on WSPR bands from 2200 m to 10 m are reported in real time at **wsprdaemon.org** and data are also publicly available from an Influx database via a Grafana web tool [8].

In this article we outline the types of noise we are trying to measure, and those we are not, the time and frequency domain algorithms currently in use, and matters of calibration. We show examples of use, comparing noise and SNR for a ground wave signal, the daily pattern of noise at KPH, patterns of variation in noise at two receivers 1000 km apart, and the change in daily noise patterns with season. We close with a discussion on these results and how we intend to develop the noise-measuring capability within *wsprdaemon*.

Types of Noise We are Trying to Measure

At this stage of our study our noise measurements are intended to include the contributions from the following sources:

• *Galactic noise from the sun or other celestial sources*. For example, at a low noise location above 20 MHz one might expect to see noise variations from galactic sources in synchrony with a sidereal day

(about 23 hours 56 minutes) when the sun is inactive.

• *Natural, atmospheric noise from global lightning (sferics)*. The result of global lightning has been likened to a "drizzle" of broadband noise, very different from the individual, identifiable crashes from local and regional thunderstorms. The local noise level from distant thunderstorms decreases with frequency and shows a diurnal variation with the lowest levels below about 10 MHz during the day, due to absorption in the lower ionosphere, with a significant increase at night.

• *Artificial noise from the myriad of electrical and electronic items of the modern world*. Compact sources are most likely to produce only a very local near-field impact, when their dimensions, or the lengths of wiring, are much smaller than the wavelength(s) of the unwanted emissions. Some equipment with power and or signal leads, without suppression, of a length approaching or exceeding quarter of a wavelength may radiate and have an impact over a wider area.

• *Common mode coupling is an avoidable ingress mechanism for noise*. Where it allows ingress of the types of noise listed above our objective has been to use these measurements as one of several tools to identify and then eliminate, *e.g.* appropriate use of baluns, including tuned baluns, grounding, elimination of current loops and, specifically for the KiwiSDR, using a WiFi connection rather than Ethernet.

Types of Noise We Do Not Want to Measure

Future work may be undertaken to

provide separate estimates of noise level and frequency of occurrence from these sources, but for now, our chosen measurement methods deliberately minimize their contribution to our noise estimates.

• *Noise from local thunderstorms.* While this type of noise is broadband it is intermittent, and hence can be excluded by selecting a time window when not present.

• *Artificial impulsive noise.* Even if this type of noise is repetitive, for example, arcing switch contacts, there will likely be times when it is not present.

• *Interference from radio signals.* There are times when amateur transmissions of CW or data modes such as RTTY or PSK31 occur in the bands used by WSPR. Continuous and sporadic carriers have also been seen. These types of signals can be excluded, even if continuous, by making our measurements within gaps in the frequency spectrum.

Despite the good degree of clarity in these intentions for the types of noise we do and do not seek to measure much remains unknown. We have imperfect understanding of the incoming and local sources of noise and of the coupling mechanisms. There are, therefore, limitations to the measurements described here, but as we show through several examples, they do provide useful insights and ideas for further investigation.

Appropriate Noise Measurement Methods

From the requirements outlined above of what noise we do and do not want to measure, two general approaches emerge. First, measure the minimum true RMS level over one or more short, fixed-width windows that are allowed to occur anywhere within a longer time window. Second, over a specified overall time window measure the spectral components over an appropriate bandwidth in a number of short time windows and take a fraction of the total number of spectral estimates, irrespective of frequency or time, that form the lowest *x*% of the values. We've called these our RMS and FFT methods respectively.

The following sections illustrate these two methods in more detail as we have applied them to estimating noise concurrent with the reception of WSPR signals. For both methods, to minimize programming effort, we have used the cross-platform *Sound eXchange* (*SoX*) software package [9]. Specifically, we use *SoX's* trim and stats options to set the time window and calculate a trough value for the RMS method, and the stat -freq option for the FFT method.

Both methods are applied to the same audio capture file in *wav* format obtained

from a receiver. For the work reported here the audio capture is done by the program *wsprdaemon* that also decodes and reports WSPR spots [7]. Options are provided for analog baseband audio via a sound card, or network audio from an RTL or other SDR using the *Soapy* remote protocol, or as used for the measurements in this paper, via a network connection to a KiwiSDR using the *kiwirecorder.py* utility [10]. A diagram that outlines our noise measurement system and these options is shown in **Figure 1**.

The RMS Method Applied to WSPR

WSPR transmissions start one second past each even minute and end at about 111 seconds. This schedule, in theory, provides a gap of one second before, and eight seconds after a transmission. Of course, in practice a transmission may start early or finish late. We chose to use 0.25 – 0.75 s and 113 – 118 s as the overall pre- and post- transmission time windows, but they can be altered. The *SoX* stats option calculates the lowest RMS level over a 50 ms window within the set interval.

Figure 2 shows histograms of uncalibrated noise level obtained using the RMS method over a set of 405 two-minute WSPR intervals between 7039.94 kHz and 7040.26 kHz on 14 – 15 March 2019 at G3ZIL, Southampton, UK. The three histograms are for the pre- and post-transmission intervals, and for the minimum of the two in each interval. To the right of each histogram outliers are shown as dots, the box being the central distribution, and the quantiles between 25% and 100% are listed. At each quantile the noise level estimate from the post-transmission period is lower than for pre-transmission. This is an

expected result as the post-transmission time window is longer, allowing a greater margin for a mistimed transmission. Taking the minimum of the pre- and post-transmission estimates results in fewer outliers. Therefore, while we will see outlier estimates of noise when using this method, they can easily be removed, if required, by post-processing the data using a median filter.

The FFT Method Applied to WSPR

The *SoX* stat -freq option automatically splits an input audio file into time segments appropriate to a 4096-point FFT, providing 2048 Fourier coefficients for the positive half-space. Given the 12000 Hz sample rate and the two-minute duration of the baseband *wav* file produced by *kiwirecorder.py* this results in 352 blocks of Fourier estimates. Our analysis covers the band from 1338.98 Hz to 1661.27 Hz, which for convenience we refer to as 1340 – 1660 Hz. This is sufficient to cover the WSPR band of 1400 – 1600 Hz and matches the band-pass set in the *kiwirecorder.py* option list. We illustrate our method using a single *wav* file recorded at KPH, Point Reyes, California at a WSPR dial frequency of 7038.6 kHz at 1116 UTC on 19 April 2019.

Figure 3A shows the spectrum from 1340 – 1660 Hz where the power in each 2.9 Hz band in each of the 352 blocks has been summed. From the LF end the SNRs reported by WSPR for the three high peaks were −9, −7 and −7 dB. The three lower, but distinct, peaks at 1431, 1557 and 1574 Hz were also decoded with SNRs of −27, −28 and −23 dB. The signal with an observable peak at 1450 Hz was not decoded. In **Figure 3B** the individual spectrum power estimates are plotted in ascending order (this is not
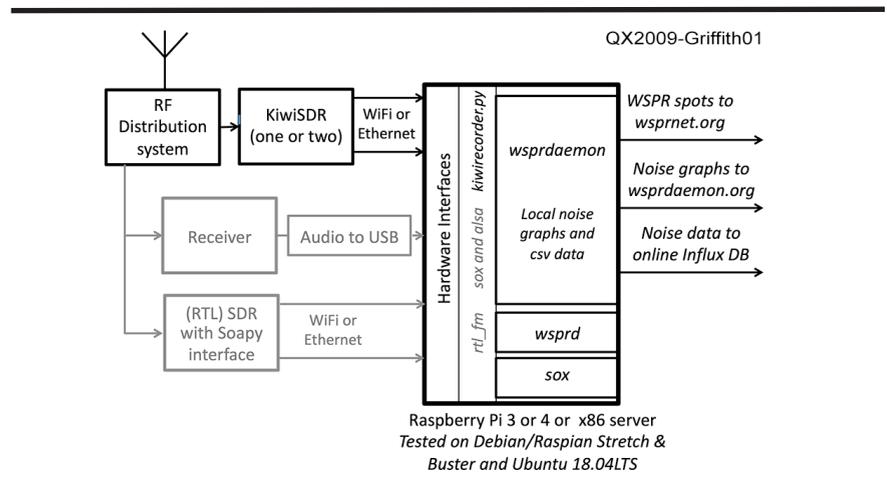
Figure 1 — Block diagram of the noise measurement arrangement. Software modules are in italics. While *wsprdaemon* and its noise measurement can use an audio input or an RTL-SDR via a *Soapy* interface, they are greyed out because they were not used for the examples in this article.

a cumulative graph), and a dashed line is shown at 30% of the entire spectrum. In this case 30% represents 33 frequency bins out of the full set of 111. While our decision to consider the sum of the power spectrum to this 30% point, and to use that bandwidth (97 Hz) as the basis to derive the noise level in 1 Hz, is to some degree arbitrary. There are three reasons for this choice. First, it is well above the region with higher

slope below about 15%. Second, it is well below the start of another upward slope above 60%. Third, it is less than 37% that represents the ratio of our guard band of ±60 Hz either side of the 200 Hz WSPR band. One hypothesis for the rise above 60% is that it represents the sum of the power from WSPR signals that are, individually, below the decode threshold of about −30 dB SNR in a 2.5 kHz bandwidth.

Over a set of 33 two-minute WSPR intervals recorded at KPH at the WSPR dial frequency of 7038.6 kHz on 19 April, the histogram in **Figure 4** illustrates the number of times spectral bands (here in 10 Hz groups) were within the lower 30%. Preferentially, but not exclusively, the algorithm used noise power estimates within the 60 Hz guard bands above or below the 1400 – 1600 Hz WSPR band; 30.4% being
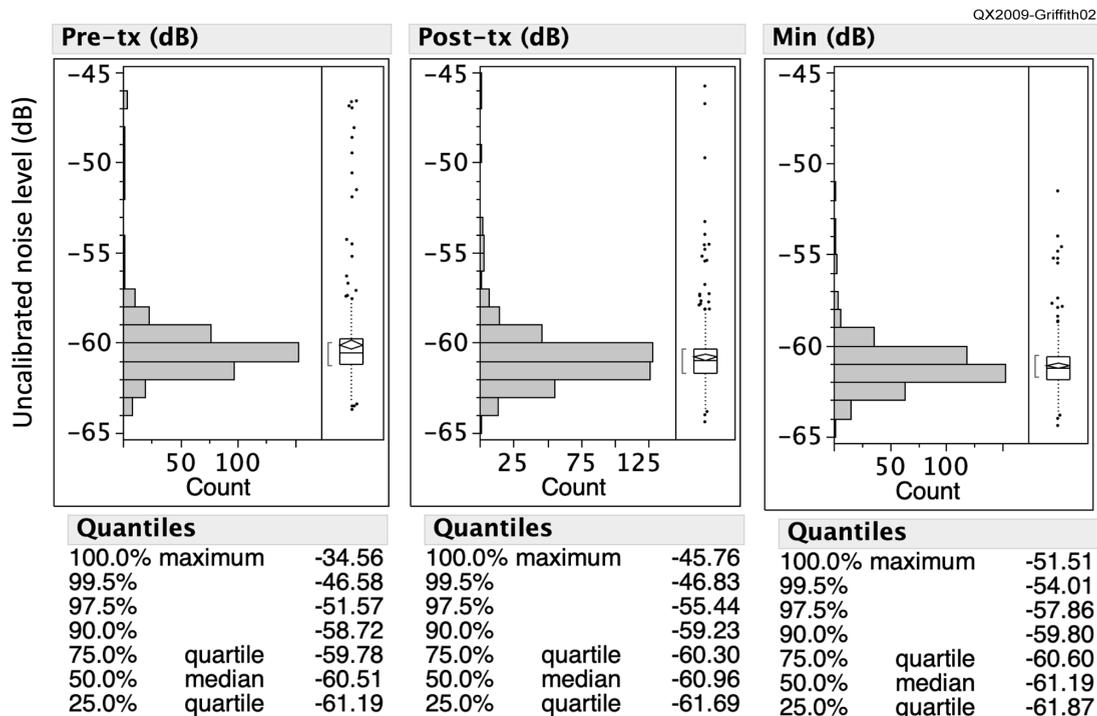
Figure 2 — Histograms of uncalibrated noise level from *SoX* stats RMS 50 ms trough measurements pre- and post-transmission, and the minimum of these minima for 405 two-minute WSPR intervals between 7039.94 kHz and 7040.26kHz on 14 – 15 March 2019 at G3ZIL, Southampton, UK.
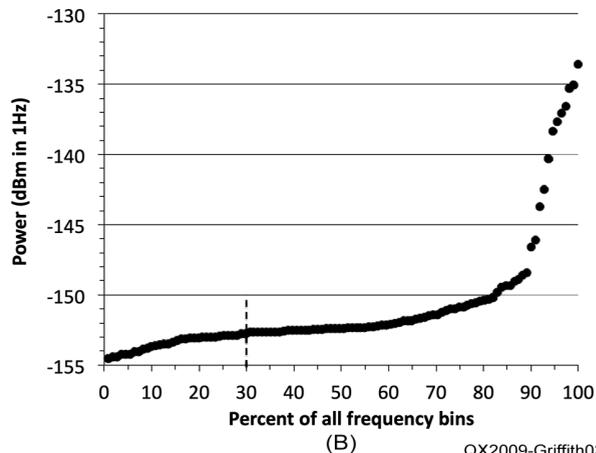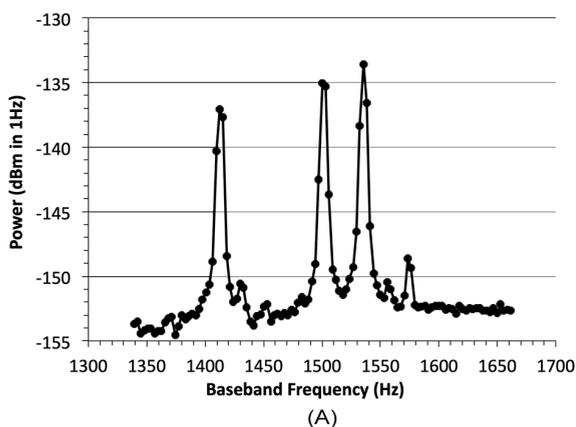
Figure 3A — Baseband power spectrum from a single two-minute *wav* file recorded at KPH, Point Reyes, California with a WSPR dial frequency of 7038.6 kHz at 1116 UTC on 19 April 2019 showing three strong and four weak WSPR signals.

Figure 3B — Power spectrum estimates as in Figure 3A but sorted in ascending order. Our noise estimate is based on the sum of the noise power in the lower 30%.

below 1400 Hz, 26.5% being above 1600 Hz and 43.1% within the WSPR band, representing 0.55, 0.44 and 0.21% per Hz respectively. While the detailed distribution changes from hour to hour we have found that the broad characteristic shape of the histogram remains stable.

This method is similar to, but not exactly the same, as that used within the WSPR decoder. Within the WSPR decoder (*wsprd*) the noise is estimated over ±150 Hz, *i.e.* from 1350 – 1650 Hz, averaging by frequency the amplitude in 512-point FFTs, applying a 7-point rectangular window filter, sorting the resulting coefficients by amplitude and taking the single coefficient at the 30th percentile as the noise level.

Calibration

At this stage of the project, our primary aim is to report noise levels at the antenna socket of the receiver in units of dBm in 1 Hz using a well-understood calibration over the full amplitude and frequency range. The practical details of the calibration methods will depend to some extent on the type of receiver used, in particular whether analogue or SDR, and how signal level measurements can be made. We have described in detail our measurement approach for the KiwiSDR in a technical report available on-line [11]. The following are the key points relevant to the examples described in this article.

Despite the KiwiSDR being a direct sampling receiver the gain is not constant over the range 10 kHz to 30 MHz. In summary, the deviation from a flat response is −2 dB at 10 kHz and a broad peak between 12 MHz and 30 MHz with a maximum of +3 dB at 21 MHz. This response is broadly consistent between units and is not due to the theoretical or measured transfer function of the 7-pole Chebyshev low pass filter. Consequently, the *wsprdaemon* script provides a default amplitude correction for each WSPR band, which the user can adjust.

The *wsprdaemon* noise measurements are made in a commanded 320 Hz bandwidth, but reported in a 1 Hz bandwidth. As the shoulders of the KiwiSDR digital filters intrude (asymmetrically) into the commanded passband the calculated equivalent noise bandwidth is used to convert the measurements in the nominal 320 Hz passband to 1 Hz.

Using a fixed gain rather than AGC and a 16-bit *wav* file as the audio source restricts the dynamic range. Through empirical tests we arrived at a fixed gain setting that results in the receiver self-noise exercising 3 – 5 bits while providing sufficient dynamic range such that clipping typically occurred

only during 1 in 200 two-minute intervals at KPH.

As the KiwiSDR does not use front-end band-pass filters the user must be sure that the sum total of all signals does not cause overload, *wsprdaemon* now logs any instances of overload to help interpret biased noise level estimates. However, non-linearity and imperfect alias rejection can result in a higher noise level within the measurement band due to strong, but below overload, out-of-band signals. We have made measurements of the Noise Power Ratio using both single tone and broadband noise, which suggests that at an input level of −15 dBm (just below overload) the additional noise is about −158 dBm in 1 Hz.

Finally, there are questions of definition and robustness for the noise estimate. While the methods to obtain the RMS and FFT estimators are based on well-defined algorithms, both lack robustness in the presence of some types of noise, which inevitably bias a measurement high. The choice of "best" estimator for the background noise level has been studied extensively within the long-running US Government Signal-to-Noise Enhancement Program [12]. Their preferred estimator is based on a kernel smoothing model fit (KSnf) to the first peak of a histogram of frequency of occurrence against binned median received power per frequency bin. For the example distribution shown in **Figure 4** the difference between our 30% quantile FFT estimator and the KSnf
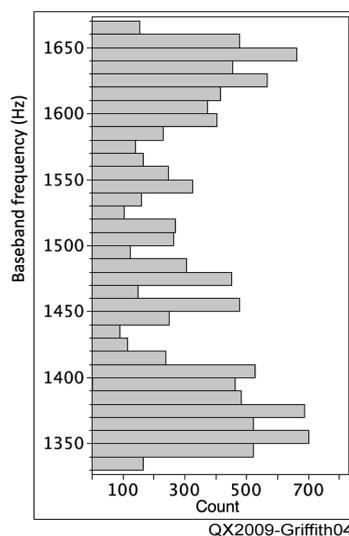
Figure 4 — An example histogram in 10 Hz bands of the frequency bins forming the lowest 30% with lowest power coefficients, comprising data gathered over 33 two-minute intervals at KPH on 19 April 2019 at a WSPR dial frequency of 7038.6 kHz.

method is very small at 0.1 dB (−152.7 dB and −152.6 dB respectively), but undoubtedly the difference will depend on the characteristics of the noise.

Comparison of Our Noise Estimates with the WSPR SNR for a Ground Wave Signal

One approach to validating our noise estimates is by assuming that the received level of a ground wave signal remains constant, and that changes in the estimated SNR of the signal are from changes in noise rather than the signal. We consider this a reasonable assumption if ground conductivity is constant and there is no reception via the ionosphere during an experiment. WSPR transmissions on 40 m from a QRP Labs U3S at N2AJX using an AS-2259 turnstile inverted **V** antenna modified for 80 and 40 m [13] were received 20 km distant on a Beverage antenna at KD2OM with a KiwiSDR and *wsprdaemon* software.

Figure 5A shows scatter plots of the RMS and **Figure 5B** shows FFT noise estimates, and the 1 dB quantized WSPR SNR for each 2-minute interval when a WSPR spot was received. The noise estimates have been adjusted to a bandwidth of 2.5 kHz to match how the WSPR SNR is reported. Of course, the actual WSPR SNR is about 32 dB higher, given that the signal processing bandwidth for each tone is about 1.5 Hz and not 2.5 kHz. Also shown are non-parametric density contours from 10% (outer) to 90% (inner); these contours aid our interpretation of these scatter plots. Also shown is a 1:1 line representing the hypothesis that, for S constant, changes in SNR should have an inverse 1:1 relationship with changes in N. Therefore, we would expect the scatter plots to show a slope parallel to this line.

Visually, the FFT estimator comes closer to the expected form, although with a tendency to show a lower noise than implied by the SNR as the SNR decreases. That is, while the ridge of the non-parametric density contours is below, but close to, the line at an SNR of −3 dB the ridge is successively further below the line as SNR decreases. The RMS estimates are consistently below the line and, from the trend of the density contours, have a shallower slope. From our wider observations of noise on the LF to HF bands using these two methods it is clear that characteristics of the noise do affect each algorithm differently. By characteristics we mean that the noise may not be additive Gaussian white noise. For the examples in **Figure 5A** and **5B** the implication is that the RMS algorithm consistently found 50
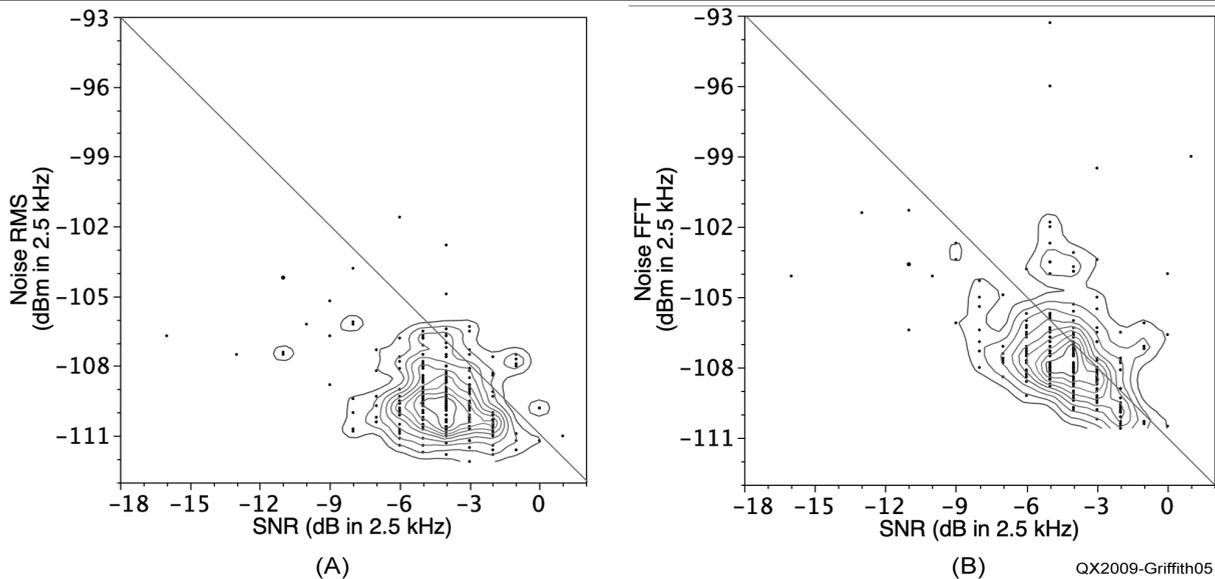
Figure 5 — Scatter plots of the RMS (A) and FFT (B) noise estimates against the 1 dB quantized WSPR SNR for each 2-minute interval. The observations were made between 5 November 2019 0000UTC and 7 November 2019 1156UTC, a total of 235 spots. The sloping line represents the hypothesized inverse 1:1 relationship between SNR and N if S is constant.
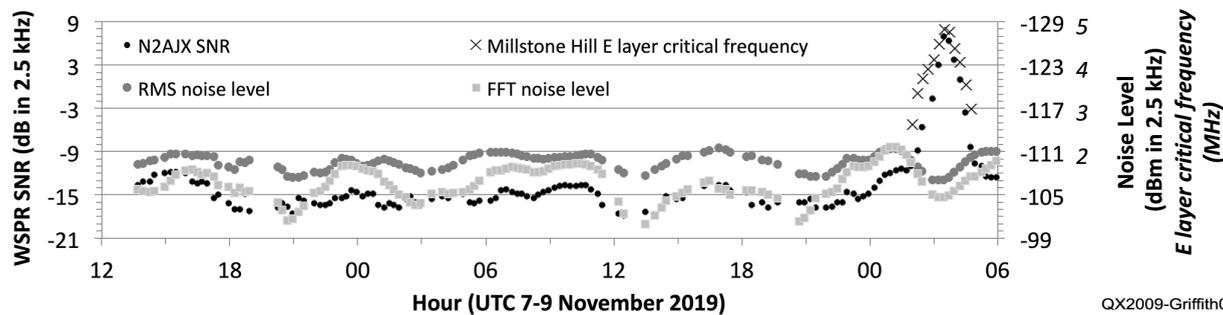
Figure 6 — Time series of the SNR of N2AJX WSPR transmissions received at KD2OM together with the noise level estimated by the RMS and FFT algorithms and the E layer critical frequency for 9 November. The data have been filtered using a running mean over one hour.

ms periods when the noise was substantially less than from the FFT estimator over the lower 30% of frequency bins. Bursts of broadband noise would, as one example, be compatible with this observed difference. Given that the WSPR SNR estimate was derived using a binned FFT approach one would expect the SNR to show a closer relationship to our FFT method than to the RMS when the estimates from the FFT and RMS algorithms differed.

Times of convergence and divergence in the RMS and FFT noise estimates can be seen in the 42-hour time series in **Figure 6**. The data are from N2AJX ground wave signals received at KD2OM with N2AJX using a clone of a five-band CH250B 80 – 10 m vertical antenna and where one-hour moving average filters have been applied to the data. Near 0000 UTC (~1700 local time) on both days the two noise estimates,

RMS in dark grey and FFT in light grey converged, yet at other times they differed by up to 6 dB. As yet, we do not have an explanation for this pattern of convergence and divergence.

Noting that the reverse order secondary y-axis of **Figure 6** there are times of clear visual correlation between the SNR and both the RMS and FFT estimators of noise, for example from 1100 UTC on 8 November to 0100 UTC on 9 November. An analysis of the data between 1314 UTC on 7 November and 0114 UTC on 9 November gave a correlation coefficient of −0.60 for SNR and FFT, and −0.66 for SNR and RMS, both inverse correlations being statistically significant. For comparison, the correlation between the RMS and FFT noise estimators was +0.73. The correlation coefficient squared is an estimate of the proportion of the total variance explained

by the independent. Here the correlation coefficient squared show 36% and 43% of the SNR variance to be explained by the RMS and FFT noise estimators — certainly an important factor, but suggesting the assumption of the variations of SNR only being due to variations in noise was an oversimplification.

Our correlation analysis between SNR and noise ended at 0114 UTC on 9 November because of the onset of ionospheric propagation. An ionosonde record for that time from the Millstone Hill Observatory, Massachusetts, about 500 km east of KD2OM, showed a sporadic E layer at 100 km height beginning to develop at 0130 UTC, persisting until 0515 UTC; the one-hour filtered E layer critical frequency f0Es from Millstone Hill is shown in **Figure 6 [14]**. During this interval of Es propagation N2AJX SNR rose 19 dB to

peak at +7 dB from a baseline of about −12 dB. At the same time the noise increased by about 5 dB from both RMS and FFT algorithms. The implication is that the true signal level rose by 24 dB and not 19 dB. One possible explanation for the observed change in noise, assuming it was dominated by distant, propagated-in noise, is from the large difference in the reception distance annulus at KD2OM between prior to and during the Es event. Characterizing the reception annulus by its lower quartile, median and upper quartile, the change was from 1661 – 5493 – 5922 km between 0000 – 0100 UTC prior to the event compared with 464 – 725 – 1798 km between 0300 – 0400 UTC during the Es peak.

Daily Patterns of Noise on 40 m and 60 m at KPH

Few amateurs are lucky enough to live in areas with such a low level of local noise as at KPH, Point Reyes, California. However, arriving at an RF distribution system and a receiver installation that matched the site's low level of noise was not immediate or straightforward. A Clifton/DX Engineering 23 dB preamplifier was installed after an 8-way splitter, a NooElec Distill AM broadcast band band-stop filter and a 30 MHz low-pass filter. Low-noise 1 MHz switched-mode power supplies are used, with Mini-Circuits T1-1 isolation transformers at the KiwiSDR inputs. At 7 MHz the noise floor of the KiwiSDR is −154 dBm in 1 Hz, equal to −120 dBm in 2.5 kHz, about S-1 (given 50 µV for S-9 and 6 dB per S-unit). **Figure 7** shows that the

minimum received noise level at the antenna socket from the TCI530 log-periodic antenna, via the RF distribution system was about −147 dBm in 1 Hz, 9 dB above the receiver noise floor.

There is a clear daily pattern to the noise. The sharply defined minimum occurred consistently at around 2000 UTC, that is, at noon local time. There was more variability in timing and in level for the peak, but over a ten-day period the highest peak was around 1900 local time followed by, on average, a plateau until 0600 local time, but with significant day-to-day variation including, on some days, a second peak. This observed diurnal noise level is consistent with rising absorption in the D layer from local sunrise, peaking at local noon then decaying by local sunset, with refraction from the merged F layers persisting, but decaying, due to the slow recombination of free electrons and ions.

While the RMS and FFT noise estimates in **Figure 7** are highly correlated (correlation coefficient of 0.87 with 1353 samples) and with a similar span (standard deviation of 6.0 dB for the RMS and 6.5 dB for the FFT) there is a clear offset, with the median of the FFT noise estimates 3.3 dB greater than for the RMS. In addition, in this example, the short-term variability of the RMS estimator is significantly less than for the FFT. However, we have found that offsets and differences in short-term variability are not always present, or if present, have the opposite trend between RMS and FFT than seen in this example [15].

As we proposed that the explanation for the diurnal noise pattern in **Figure 7** was due to the diurnal pattern of propagation we have

to show that the noise estimates are not being biased by higher signal levels between sunset and sunrise. To that end, **Figure 8** shows one 24-hour period 23 – 24 May 2019 at KPH for the 40 m and 60 m EU WSPR bands for the RMS noise estimator with local time. The broad pattern of the noise is the same for the two bands, with the minimum lower on 60 m than on 40 m, consistent with greater D layer absorption at the lower frequency.

WSPR spot counts per hour are shown using the right-hand axis in **Figure 8**. While the profile for the 40 m noise estimate does broadly follow the number of WSPR spots per hour, this is not the case for the 60 m EU segment. Only two transmissions were received during 2000 – 2059 and one between 2100 – 2159. Our conclusion is that the noise estimates were not due to bias from the WSPR signals, despite the visual correlation in **Figure 7** for 40 m; the noise was propagated-in from a multitude of different, distant sources.

Noise Pattern Variations at Two Receivers 1000 km Apart

Having looked at the daily pattern of noise at a single receiver, we now compare the pattern of noise variations at KPH (CM88mc) and KA7OEI in Northern Utah (DN31uo), about 1000 km apart on an approximately east-west path. Both use KiwiSDR receivers and TCI530 antennas, but differ in the splitters, filters and amplifiers between the antenna and the receiver. Hence offsets between noise level estimates at different sites are expected. However, those offsets do not detract from the usefulness of the measurements for studying patterns in the noise.
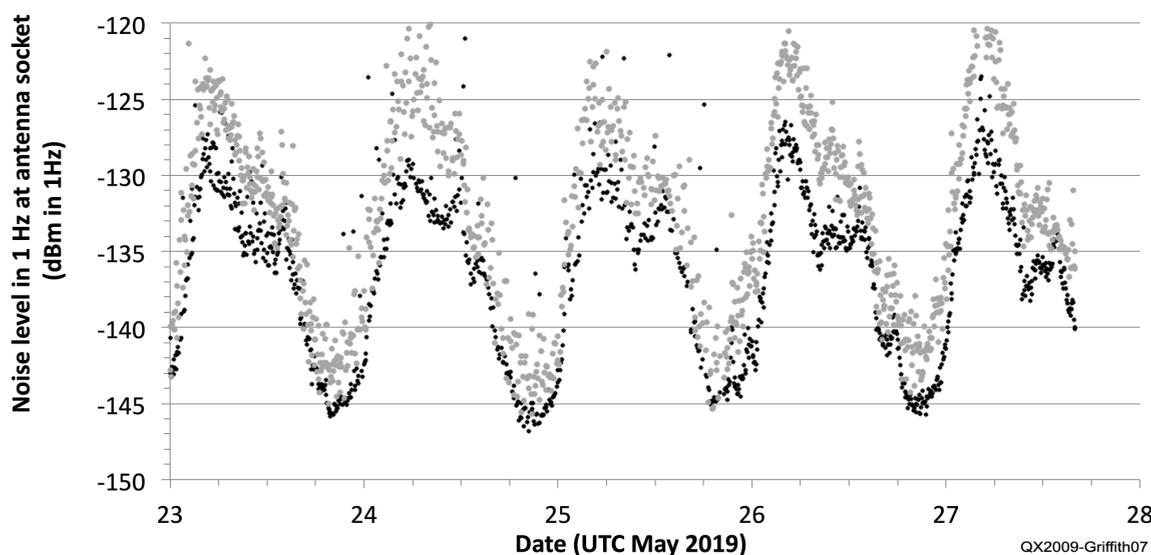
Figure 7 — Time series of the noise level recorded using a KiwiSDR at KPH, Point Reyes, California on 23-27 May 2019 using the RMS algorithm (black) and the FFT algorithm (grey) showing a clear, diurnal pattern, repeatable in the essential features of the time of minimum, and two peaks at the maximum.

An example time series for the FFT noise estimator over four days on the 80 m WSPR band (dial frequency 3.5686 MHz) for KPH and KA7OEI is shown in **Figure 9**. A 10-minute running mean filter has been applied to the 2-minute data. There are several points to note.

• At this time the noise floor at KA7OEI was higher by at least 12 dB than at KPH, seen as a plateau from about 1400 – 0000 UTC each day at KA7OEI without the round-bottomed dip seen at KPH [16], 12 dB being the difference at the KPH minimum after the overall level reported by KA7OEI was adjusted so that the peaks coincided with those at KPH.

• Diurnal variations were dominant; with the late afternoon (local time) rise in noise level occurring consistently earlier at KA7OEI than at KPH. This is consistent with KA7OEI being east and north of KPH, sunset at KA7OEI on 1 September was 39 minutes ahead of KPH [17]. Prior to a cross-correlation analysis at 2-minute intervals spanning ±60 minutes a sixth order Butterworth low-pass filter with a cut-off period of eight hours was applied to both time series. This filter removed potential bias from shorter-period variations that may have different cross-correlation properties. The graph of correlation coefficient against lag / lead in **Figure 10** shows a broad peak centered at a lead of 38 minutes for KA7OEI (R = 0.939 compared with 0.924 at zero), in keeping with the expected time difference.

• Between 0230 and 1100 UTC each day, when D layer absorption was absent, there were variations in the noise, over periods of hours, which showed a similar form at the two receivers. Taking the data on 31 August

and 3 September as examples, after applying cascaded low pass and high pass sixth order Butterworth digital filters to pass variations with periods between two and eight hours, results in the correlation coefficient profiles shown in **Figure 10**. The peaks were at zero lead on 31 August and at two minutes lead on 3 September. The implication being that variations with periods between two and eight hours occurred simultaneously at the two sites, 1000 km apart.

• Between 0600 – 0630 UTC on 2 September the noise level at KA7OEI dropped precipitously by 6 dB. This drop was followed by a train of cycles with a similar peak-to-peak magnitude and a period of the order of one hour. While a 6 dB precipitous drop in noise level was seen at KPH it occurred two hours thirty eight minutes after that at KA7OEI, and it was not followed by an obvious wave train. We currently have no sure explanation for these observations, although we suspect they may be associated with traveling ionospheric disturbances.
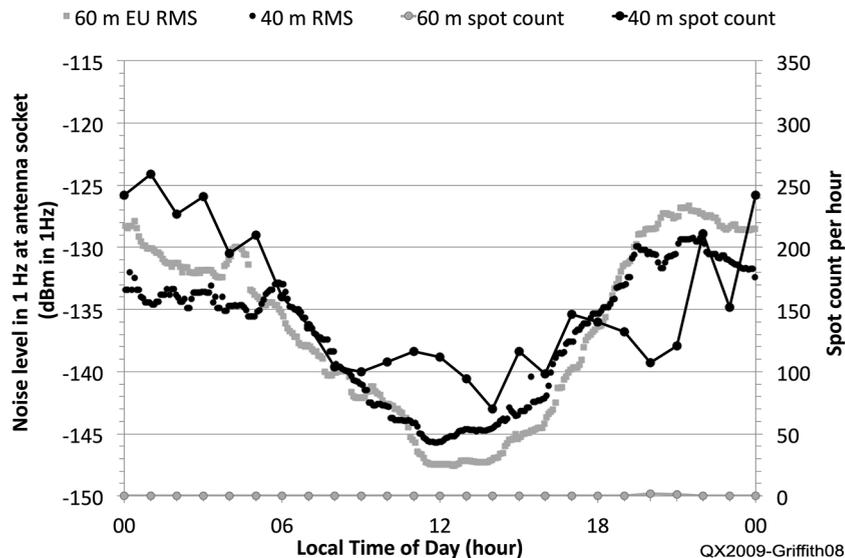
Figure 8 — A 24-hour time series of the noise level at KPH on 23-24 May 2019 using the RMS algorithm with local time of day for the 40 m and 60 m WSPR bands (dial frequencies of 7038.6 and 5364.7 kHz). The number of WSPR spots received in each hour is also shown, while the spot count for 40 m has the same overall pattern as the 40 m noise level only three spots were decoded on 60 m, showing that it is not the WSPR signals biasing the noise measurements.
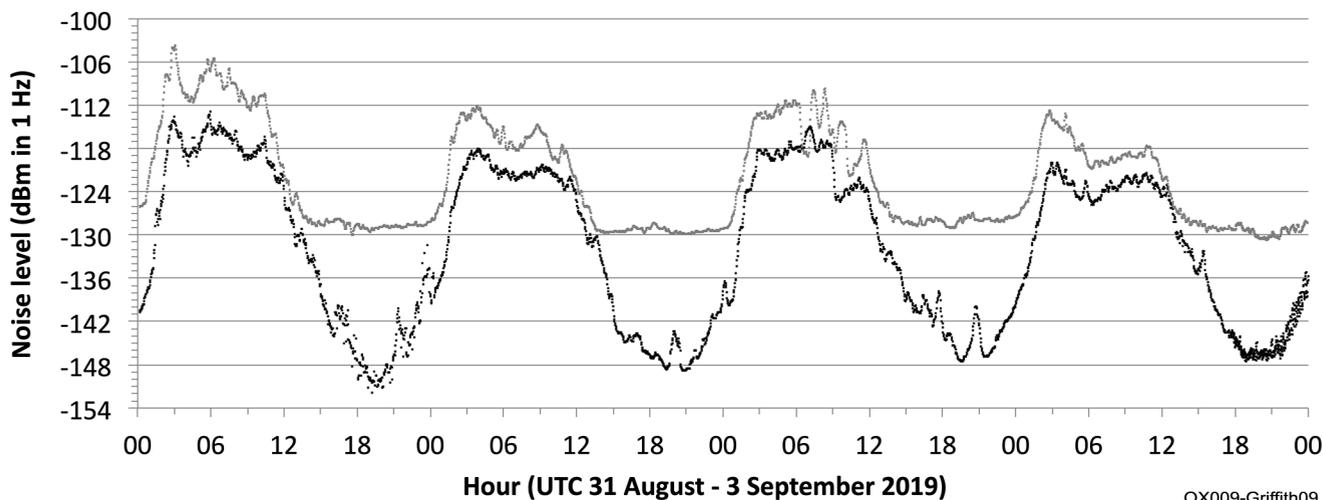
Figure 9 — Time series for the FFT noise level estimator on the 80 m WSPR band for KPH (black) and KA7OEI (grey) at a separation of about 1000 km for 31 August to 3 September 2019.

Changes in the Diurnal Pattern of Noise with Season

Seasonal changes in the diurnal pattern of noise on a single band can be illustrated using a shaded surface graphic with time of day on the *y*-axis, day of the year on the *x*-axis with the shading representing the noise level. **Figure 11** spans 12 June to 7 November for the 40 m WSPR band at OE9GHV, a quiet rural location near Alberschwende, Austria.

The RMS algorithm noise estimates over two minutes were averaged into twenty-minute intervals, hence there are 72 points on the *y*-axis each day. The dark spots or rectangles are times when the KiwiSDR receiver was disconnected from the 80 m vertical loop antenna. These points therefore represent the noise level of the receiver itself. Black lines show the times in UTC of local sunrise and sunset.

During the summer, a prominent noise peak occurred consistently around sunset, with a weaker peak just after sunrise. The noise was at a minimum during daylight hours. However, after around day 264 (21 September, the autumn equinox) there was a gradual change, over about three weeks, from peak noise around sunset to a broader noise plateau starting just after sunrise and persisting until just before sunset. After this gradual change, the minimum noise occurred between sunset and sunrise.

Our working hypothesis for the springtime diurnal pattern at KPH (**Figures 7** and **8**) centered on D layer absorption causing the minimum noise during daylight hours, we suggest that it is also the cause of the observed summer daytime noise minimum. However, we do not have a hypothesis for the change in daily pattern in September, but we do know that there were no changes to the equipment or the local environment during that time. Despite an extensive literature search we have not found published results showing HF band noise in Europe by time of day and by season to provide a comparison or insights.

Discussion

The noise measurement capability within the *wsprdaemon* software is proving useful to a number of KiwiSDRs users in several countries. While the two current estimators, a short-term RMS measurement in the time domain and a simple FFT and threshold approach in the frequency domain, have
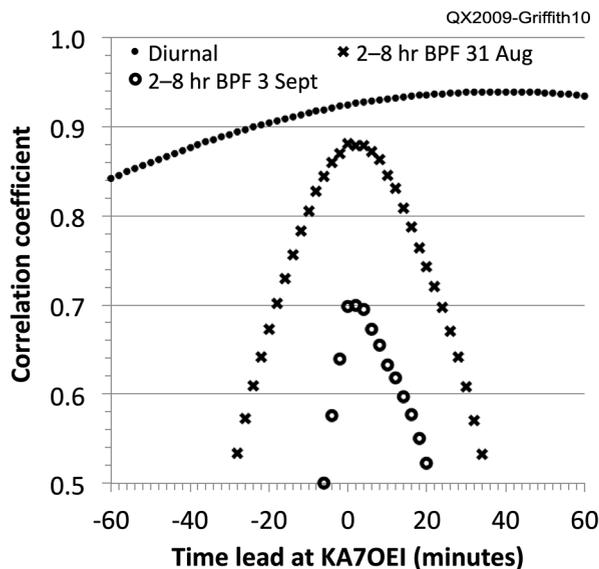
Figure 10 — Cross correlation coefficient between the noise level times series at KPH and KA7OEI shown in Figure 9 against lead time for the diurnal variation (periods over 8 hours) and for variations with periods between two and eight hours.
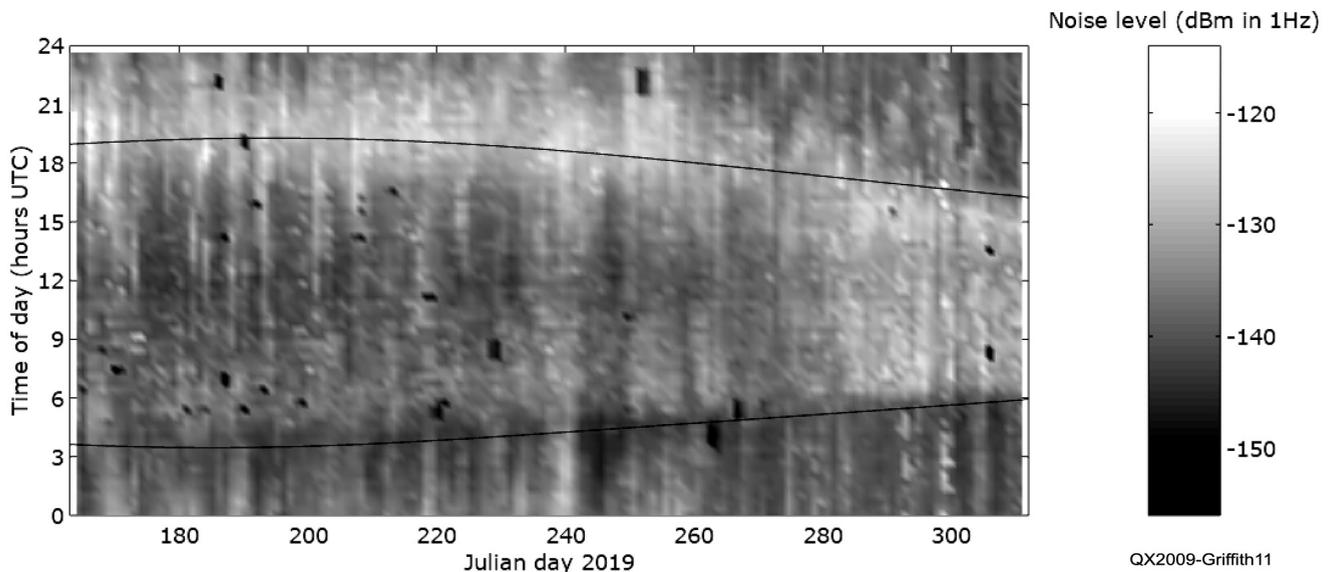
Figure 11 — Seasonal change in the pattern of diurnal noise on 40 m at OE9GHV, a quiet rural hilltop location near Alberschwende, Austria (JN47wk) from 12 June to 7 November 2019 together with lines at sunrise and sunset. After around day 264 (21 September, the autumn equinox) there is a change over about three weeks from peak noise slightly after sunset to a broader noise plateau starting just after sunrise persisting until just before sunset. The receiver is a KiwiSDR and the antenna an 80 m vertical loop via a splitter.

their limitations they do provide the basis for further analysis as shown by the examples in this article. We have started to experiment with adding a Kernel Smoothing model-fitting algorithm to the FFT data, both to evaluate its performance and the additional computational load.

Estimating noise level at the antenna socket of the receiver was the initial goal of this work, and it remains the best-controlled reference point as regards noise level calibration with frequency and receiver settings. However, *wsprdaemon* does provide a simple means in a configuration file for the user to set a band-by-band correction. To date, users that have applied these corrections have done so to account for gain or loss in their distribution system from the antenna. In these cases, the reference point moves to the antenna terminals. In either case, the noise measurement includes noise that may come from common mode coupling into the antenna cabling or from the power supplies of preamplifiers for example. This is very much a practical measurement, of direct relevance to the user of the receiver, of the noise affecting their ability to read or decode signals.

The band-by-band corrections in *wsprdaemon* can also be used, in principle, to refer the noise to free space, that is, as an electric field intensity at the antenna expressed as dBμV/m. This is important when the aim is to characterize the local noise environment external to the distribution system and receiver, and especially when making direct comparisons of absolute noise levels between different locations. However, it is not straightforward to obtain the antenna factor, ground wave gain and equivalent monopole antenna noise factor that are required to determine the field strength as described in ITU-R P.372-13 **[18]**. This is because the relative antenna factor and ground characteristics, as well as polarization and wave arrival angle contributions, may not be well known for a given comparative installation.

Even these initial steps for us as amateurs, of estimating noise at a number of locations in different countries and making the data available to all in real time has attracted the attention of professional ionospheric scientists within the HamSci community **[19]**. They appreciate the capability of radio amateurs to gather and disseminate carefully calibrated data and they are keen to apply their knowledge to interpreting the results — for example, the reasons for the seasonal changes in noise shown in

Figure 11. There is also the puzzle of the staggered precipitous drops in noise level at KPH and KA7OEI shown in **Figure 10**. WO7I in northern Nevada and ND7M in southern Nevada, both with KiwiSDRs, have recently installed *wsprdaemon* and started reporting noise estimates. We look forward to using their data, from north and south of a line joining KPH and KA7OEI to study further the patterns of noise variation and their origin.

Acknowledgement

Gwyn Griffiths, G3ZIL, was first licensed in 1970 as GW3ZIL. He spent his career developing instruments for marine scientists, becoming UK National Oceanography Center's Chief Technologist and Professor of Underwater Systems at Southampton University. Since retiring in 2012 he has returned to amateur radio, concentrating on using WSPR and data analysis to understand HF band noise and minimize its impact on reception.

Rob Robinett, AI6VN, is CEO and Founder of Mystic Video, a Silicon Valley developer of professional TV broadcasting products. After a 40-year hiatus he resumed his interest in amateur radio in 2017 with a particular interest in optimizing RF receiving systems. He has deployed SDR and WSPR receive systems at the historic KPH receive site and a second set at a rural location on Maui.

Glenn Elmore, N6GN, was first licensed in 1961 as WV6STS. He has been active in RF-microwave pursuits as a career and as his radio hobby since then. He has discovered, made advances in the theory and application and published papers on, surface wave transmission lines. He has been actively involved in applying WSPR to VHF/UHF measurements and lately in pursuing the limits of small broadband antenna systems in the LF-HF region as part of the KiwiSDR project.

Notes

[1] See **rsgb.org/main/technical/propaga-tion/noise-floor-study/hf-noise-monitor-ing-campaign/**.

[2] See **hf.r-e-f.org/c4_iaru_r1/16_Vienne/ VIE16_C4_15_VERON_Provisional%20 Results%20of%20Measurement%20 Campaign.pdf**.

[3] See **rfnoise.amsatsa.org.za/about_ us.php**.

[4] See **vienna.iaru-r1.org/wp-content/ uploads/2019/01/VIE19-C7-002-DARC-ENAMS.pdf**.

[5] Available at **www.researchgate.net/pub-lication/283088496_HF_radio_noise_ emissions_measured_at_Key_West_ Florida_March_-_October_2010**.

[6] See **www.radiomarine.org/**.

[7] The source code for the *wsprdaemon* program that includes our noise analysis methods is available online at **github.com/ rrobinett/wsprdaemon**.

[8] Data can be selected, plotted and down-loaded from **grafana.int8.com/**.

[9] Available at **sox.sourceforge.net/**.

[10] Information on the KiwiSDR is available at **kiwisdr.com/** and on **kiwirecorder.py** at **github.com/jks-prv/kiwiclient**.

[11] A detailed technical report on the method development and calibration is available at **www.researchgate.net/ publication/334612025_Estimating_ LF-HF_band_noise_while_acquir-ing_WSPR_spots**.

[12] G. Lott, W. Stark, and M. Bail. "Multi-location long-term HF noise measure-ments and comparison to ITU-R P. 372-8," *Proceedings 10th IET International Conference on Ionospheric Radio Systems and Techniques (IRST 2006)*, 2006 pp. 133–137.

[13] See **arrl-ohio.org/SEC/nvis/ Modified%20AS-2259%20NVIS%20 Antenna.pdf**.

[14] The ionograms are available at **lgdc.uml. edu/common/DIDBDayStationStatisti c?ursiCode=MHJ45&year=2019&mont h=11&day=9**. Further information on the Global Ionospheric Radio Observatory can be found in B. W. Reinisch and I. A. Galkin, "Global ionospheric radio observatory (GIRO)," *Earth, Planets, and Space*, 63, (2011) pp. 377-381, **https://link.springer. com/article/10.5047/eps.2011.03.001**.

[15] Daily plots of FFT and RMS noise esti-mators from 2200 m to 10 m for a number of reporters are available at **wsprdaemon. org** and are likely to show a range of differ-ences between the two noise estimators.

[16] The daytime noise plateau is not always present at KA7OEI, interested readers can explore this and other wsprdaemon noise data sets using the Grafana tool, the KA7OEI 80 m data is at **grafana.int8. com/**.

[17] Use the calculator at **www.esrl.noaa. gov/gmd/grad/solcalc/**.

[18] Available at **www.itu.int/rec/R-REC-P.372-14-201908-I/en**.

[19] See **hamsci.org/**.