Some Observations While Using the KiwiSDR to Spot WSPR Stations

Gwyn Griffiths, G3ZIL gwyn@autonomousanalytics.com , Glenn Elmore, N6GN
n6gn@sonic.net, Rob Robinett, AI6VN, rob@robinett.us

Abstract

This article is intended to document some of our recent experiences using a KiwiSDR on WSPR and observed degradations of those spots when compared with other receiving and decoding techniques. Three techniques are examined: complete decode of WSPR stations by way of the KiwiSDR WSPR extension, use of the KiwiSDR to downconvert to an audio file which is decoded by a remote host, and for contrast a non-KiwiSDR path using an Apache Angelia SDR board. These comparisons show degradations in the two KiwiSDR paths both in terms of SNR of spotted stations and in number of stations spotted. Additionally these investigations have identified spurious signals associated with the downconversion process within the KiwiSDR.

1. Background

During recent months, several WSPR stations around the world have begun using KiwiSDRs as a means of measuring transmitter, receiver, antenna and propagation by spotting into the WSPRnet.org data base. The KiwiSDR has shown itself to be a very attractive candidate for these sorts of studies. With its good performance, built-in Web support, spectrum analyzer capability, low cost and with a large number of stations worldwide the data produced has shown itself to very useful for many purposes. Many of us have become quite enthusiastic supporters of this platform for both communications and for measurement purposes. KiwiSDRs at AI6VN/KH6, KPH the Maritime Radio Historical Site at Point Reyes, California, N6GN and WA2ZKD have produced excellent results compared to the global WSPR reporters, particularly so when the KiwiSDR was used with a broadband antenna and when multiple ‘receivers’ within a single KiwiSDR were employed for spotting WSPR stations.

But several stations doing this have noticed anomalies. As part of these uses for the KiwiSDR we have noticed some issues that if better understood could possibly lead to improvements in this already very useful platform. Towards that end, we have made some studies using three approaches, two involving the KiwiSDR and one using a separate SDR path for comparison. The two KiwiSDR paths each use the KiwiSDR for downconversion to audio. One uses the available WSPR extension in the KiwiSDR, and the second path uses a KiwiSDR ‘receiver’ to generate a .wav file resulting from down conversion to audio of a particular WSPR band’s 2 minute segment. This second KiwiSDR path uses AI6VN’s Kiwiwspr.sh bash script[1] to remotely access a KiwiSDR via the Kiwirecorder.py script running on a remote computer host and communicating over a network to further decode WSPR spots and post the results to WSPRnet.org. The three paths spot into the WSPRnet.org database as

- N6GN — Apache Angelia SDR/WSJT-X v 1.9.1
- N6GN/k — KiwiSDR downconvert, remote WSJT-X 1.9.1 decode
- N6GN/Kiwi — KiwiSDR downconvert & KiwiSDR WSPR extension decode and spot

An illustration of these three paths are shown in Illustration 1.

Two anomalies we have seen are that WSPR spots involving the KiwiSDR downconversion path have somewhat lower SNR than those from the Apache/Angelia path. We’ve further noticed that there were more spots from the Angelia SDR KiwiSDR than from the KiwiSDR-downconvert-remote decode which had more than the WSPR extension. We present some of these observations below.

1 See http://valentfx.com/vanilla/discussion/1331/8-channel-kiwi-wspr-decoding-script-kiwiwspr-sh-using-raspberry-pi-or-other-external-server

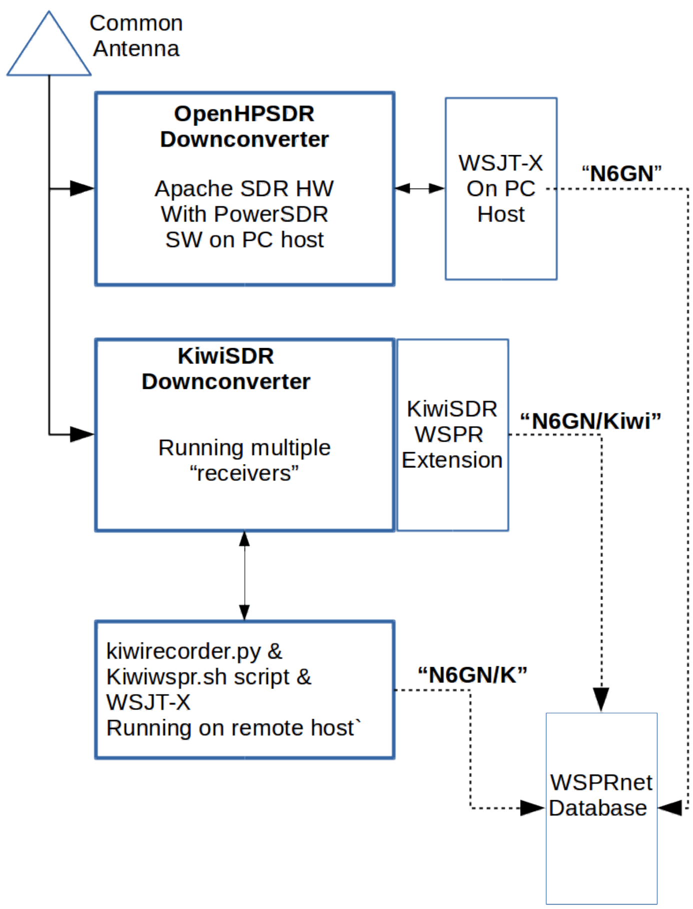

Illustration 1 The three methods of generating spots to the WSPRnet.org database include two paths involving KiwiSDR downconversion and a third using an Apache "Angelia" SDR.

2. ***Apache and /K comparison at multiple SNR***

Glenn, N6GN, observed that the KiwiSDR WSPR extension often does not produce quite as large SNR from its WSPR spots (route N6GN/Kiwi in Illustration 1). This shortcoming does not seem to be due only to a difference of decoder in the WSPR extension which is evidently built from an earlier WSPR decoder. It shows up even when a KiwiSDR is used only to provide a wav file that is subsequently decoded in a current-version WSJT-X running elsewhere (route N6GN/K in Illustration 1). WSPR spots from the Apache SDR path often produce greater values for SNR, even using identical versions of WSJT-X for the decode. It appears that the downconverted audio file produced by KiwiSDR has lower SNR than that produced by the Apache SDR.

In order to better understand the situation and with the hope of working toward a root cause that might allow KiwiSDR improvements, all three paths were run simultaneously from the same antenna. Signal levels were verified to be well under top-of-ADC for each SDR and it was also verified that noise coming from the antenna was effectively establishing the noise floor for both SDRs. If the downconversion processes were similarly good then similar audio files were expected to result.

2. Methods

To better capture the situation, several plotting methods were used:

1. Simple scatter plots of SNR 'A' vs SNR 'B' for each coincident in time spot from 'A' and 'B' augmented with non-parametric density contours. Given the large number of points and the 1 dB quantization, the contours give a more correct visual impression of the details of the relationship.
2. Scatter plots of the SNR difference against the SNR of one of the routes augmented with non-parametric density contours. These plots can sometimes pick out detail not so obvious in the 'A' vs. 'B' plots.
3. Calculation of the average SNR in each 2dB bin from -32 to -30dB upwards, with the number of spots in each bin. This provides a quantitative assessment of any trends seen in the scatter plots.

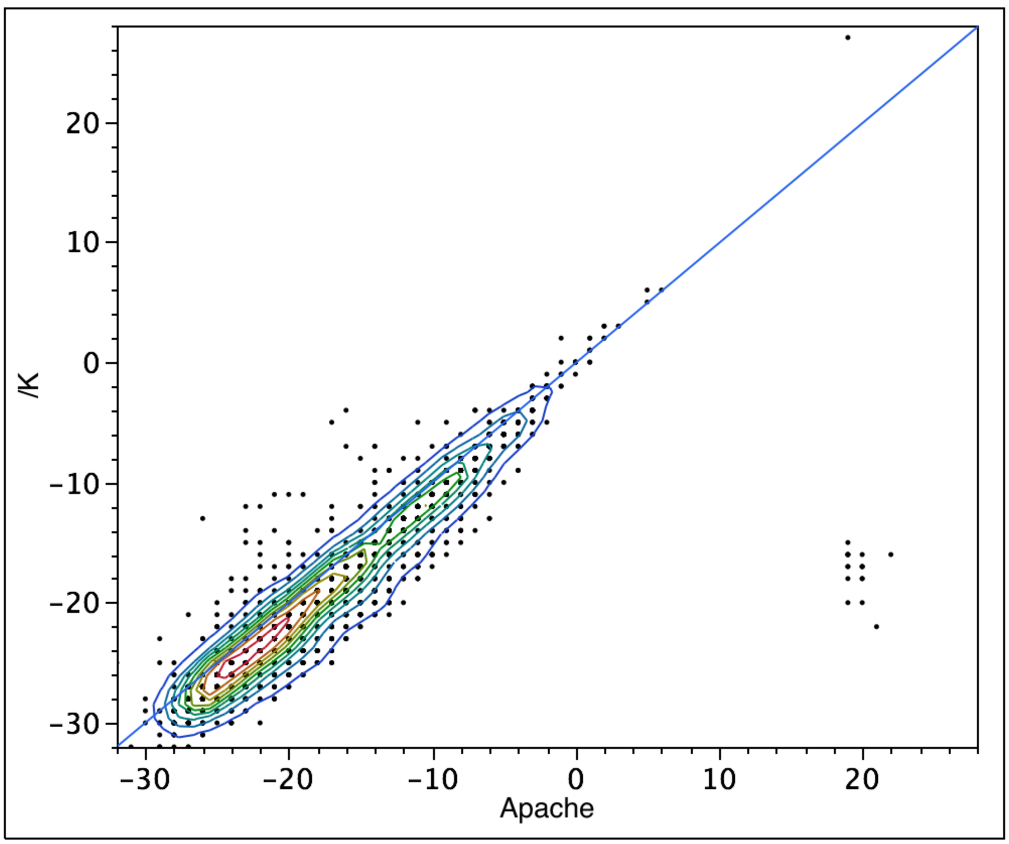

Figure 3.1 Scatter plot of SNR via routes /K and Apache. The outliers are discussed in the text.

3. Results: Apache vs. /K

Scatter plots:

The simple scatter plot, Figure 3.1, with non-parametric density contours at 10% intervals and a 1:1 line included, shows that the large majority of SNR from the /K spots to be lower than the Apache. But, at Apache

SNR of about -4dB and above the difference becomes smaller, to the point that the /K route reports *higher* SNR above +2dB.

Outliers: A) There is a cluster around an Apache SNR of +19 to +22dB where the /K SNR is -15 to -22dB. All of these are from WV0Q. These are thought to be spurious sidebands for this very strong signal. B) There is a single outlier at a /K SNR of +27dB. This is also from WV0Q. C) There are outliers (beyond the 5% - purple - line above the 1:1 line in Figure 2 between an Apache SNR of -30 to -10 dB.

Figure 3.2 draws out that these outliers where /K SNR is greater than Apache SNR form a diffuse cloud rather than the tighter-packed points where Apache SNR is greater than the /K. There is not yet an obvious cause for these.

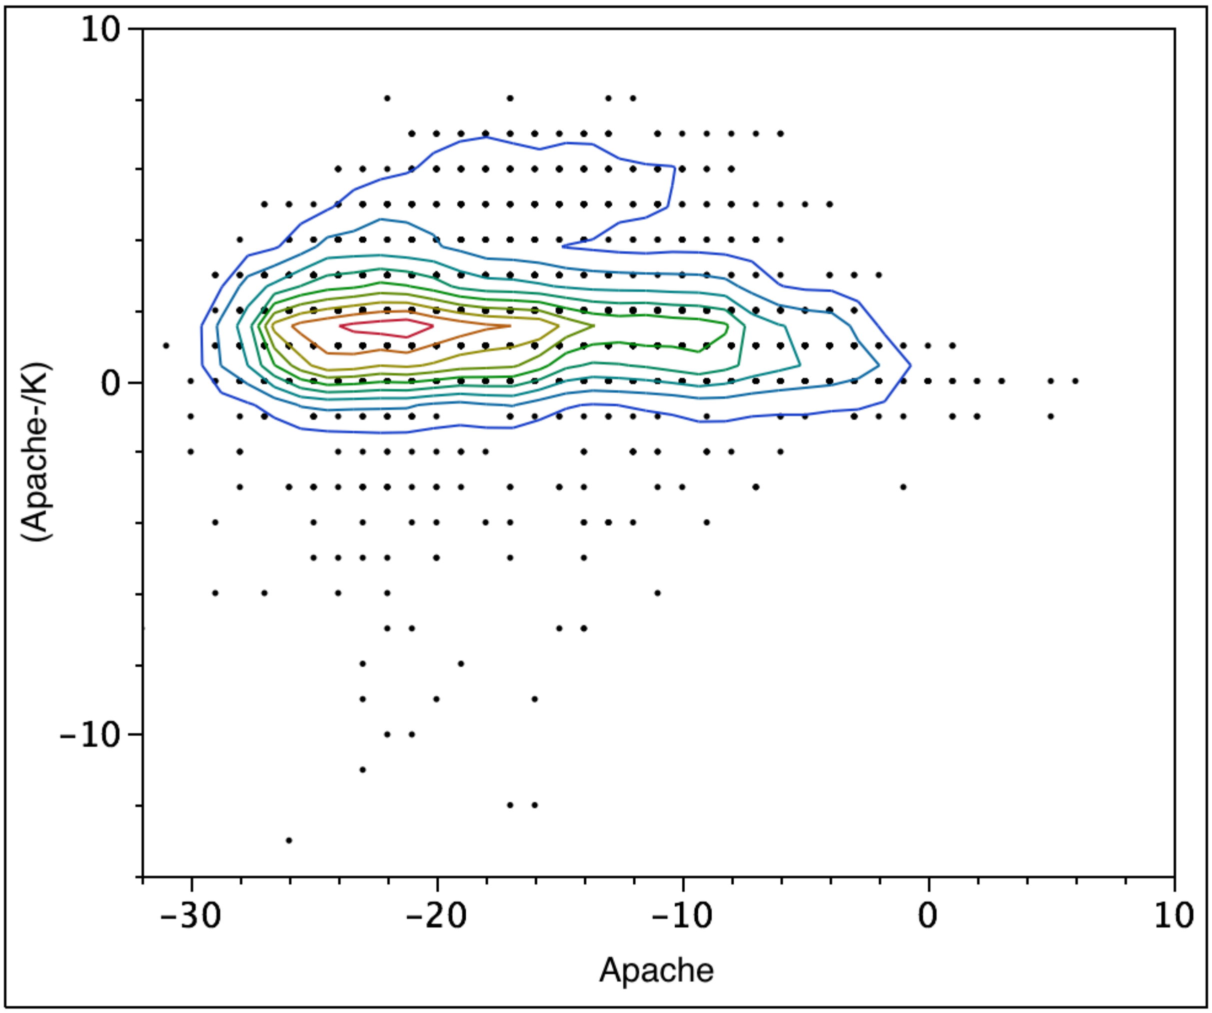

Figure 3.2.Scatter plot of SNR difference(Apache-/K) against SNR Apache.

The 'noses' of the contours towards the right of Figure 3.2 show the trend observed in Figure 3.1 for the SNR difference to decrease to zero as the SNR increases.

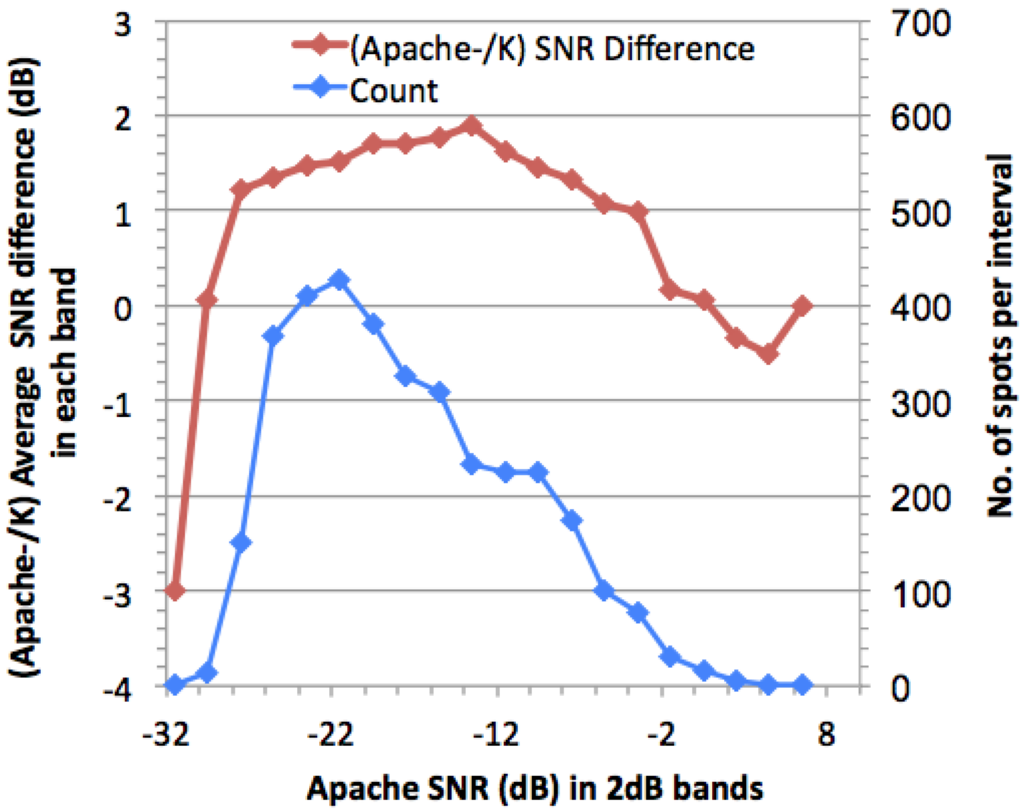

Figure 3.3. Average SNR difference (Apache-/K) in 2dB bands together with the number of spots in each bin. The data for bins centered on -31 and -29dB are most probably anomalous due to proximity to the decoding threshold.

Difference in SNR bands:

Both of these plots are constrained to a resolution of 1dB. By looking at the SNR difference within bands of SNR, we choose 2dB, a higher resolution view is obtained. Figure 3.3 shows the average SNR difference (Apache-/K) together with the number of spots in each 2dB band from -32 to -30dB upwards. First, a caution: the averages within the -32 to -30dB band (two spots) and the -30 to -28dB band (14 spots) may be anomalous. This is because we are at or very near the WSPR decoding threshold, as explained below; this is not an issue for the few points at high SNR.

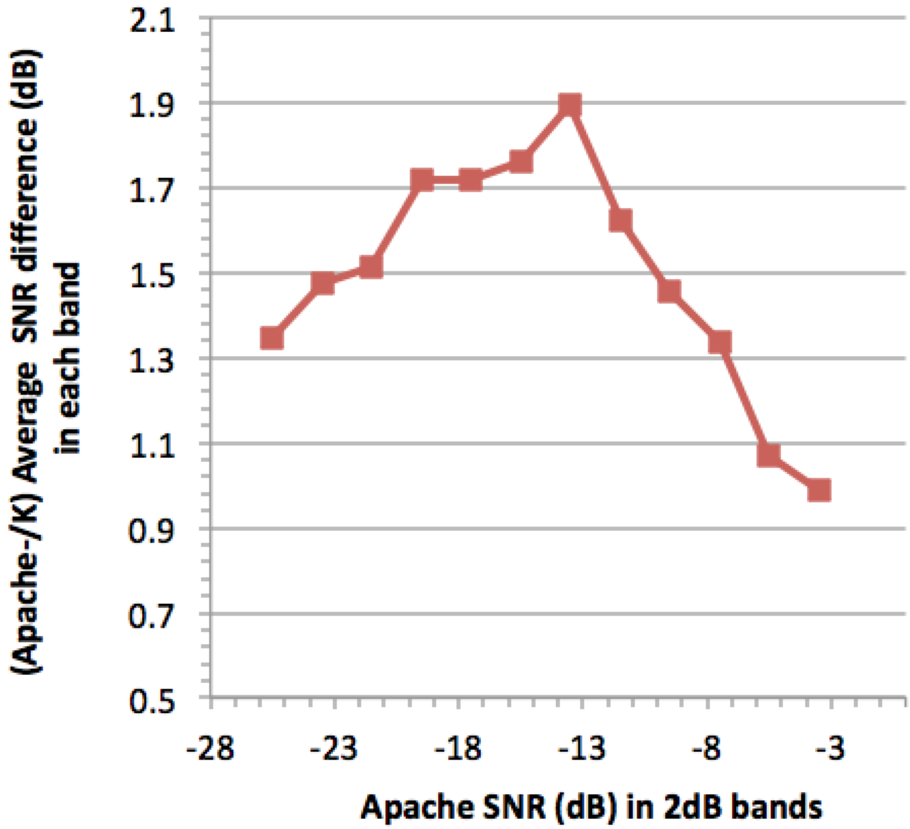

Figure 3.4 Average SNR difference (Apache-/K) in2dB bands where there were >75 spots.

Looking at the average SNR in the bands with more than 75 spots, that is, between bands centered on -27 to -3dB, there is a trend towards increasing difference from -27 to -13 dB after which there is an apparent change in slope to a decreasing difference to -3 dB, Figure 3.4. The visual impression is of two linear slopes rather than a continuous curve. However, this impression is very dependent on the difference at the inflexion at -13dB.

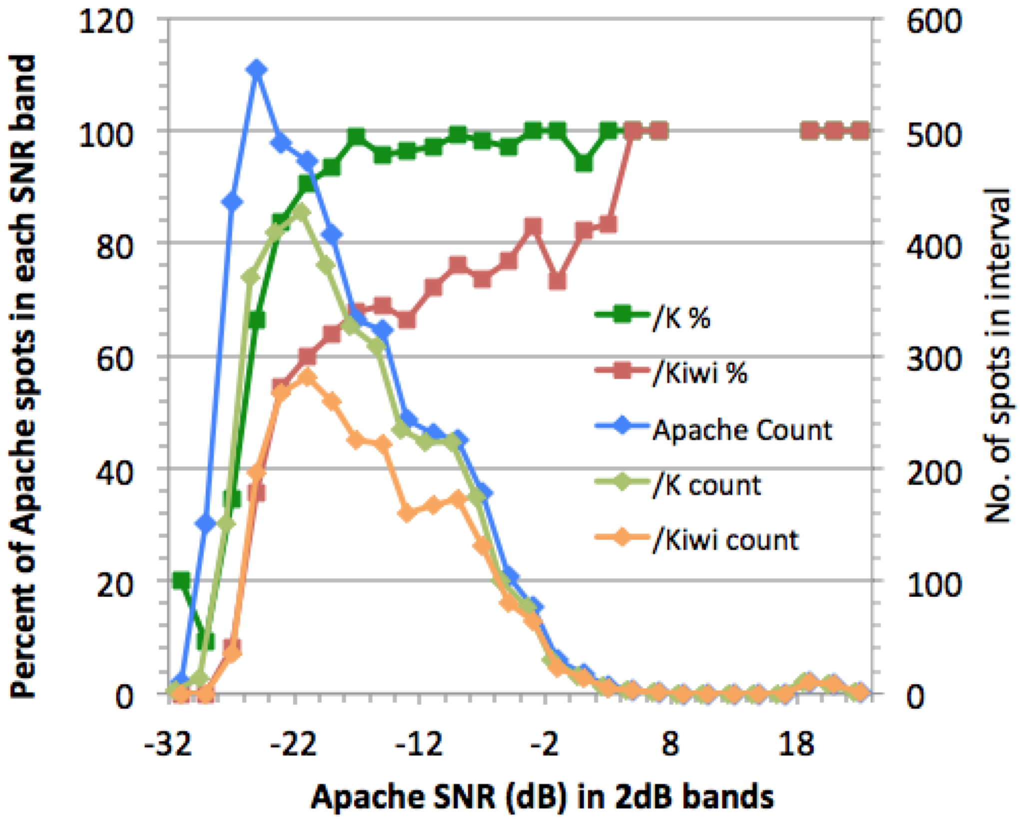

Figure 3.5. Number of spots in 2dB bands for the Apache, /K and /Kiwi, also expressed for the /K and/Kiwi as a percentage of the Apache spots per SNR band.

Percentage of Apache spots decoded:

Figure 3.5 shows the number of spots in each 2dB SNR band from the Apache and the /K and /Kiwi routes, and expressed as percent of the Apache. Above an SNR of -20 to -18 dB the /K route decodes over 93% of those decoded by the Apache. However, the percentage decoded drops rapidly below -22 to -24 dB. Given the peak of the distribution of number of spots received for the Apache is at -24 to -26dB this causes a decrease in the number of spots decoded for /K.

At first sight it is curious that this decrease in percentage of spots decoded occurs as the mean difference between the Apache and /K SNR apparently decreases, Figure 3.4. The apparent magnitude of that difference, at about 1.5 dB in this region, does not seem likely to be the cause of the low and decreasing number of spots decoded.

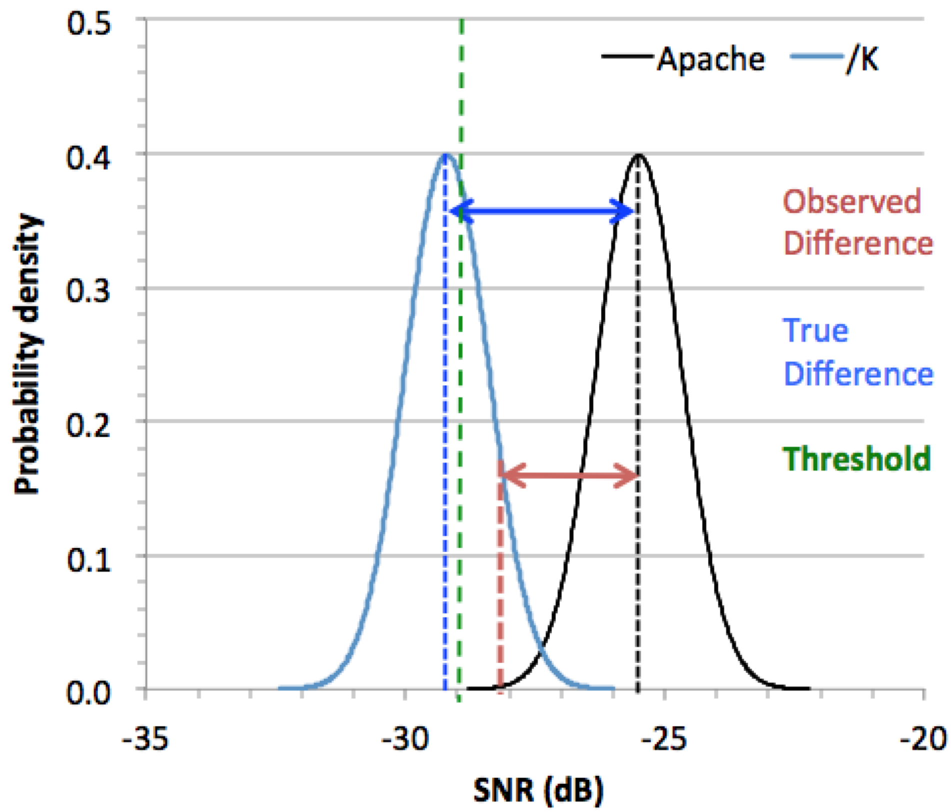

Figure 3.6. Diagram to help explain the impact of the WSPR decoding SNR threshold of about -29 dB on the calculation of SNR difference.

Effect of being close to the WSPR threshold:

The puzzle of the reducing percentage of spots from /K as SNR reduces and the apparent decrease in SNR difference can both be explained as a consequence of the /K SNR approaching the WSPR threshold. Figure 3.6 shows this as a diagram using Gaussian distributions. The Apache SNR is high enough to not be affected by the threshold, but the /K is affected. If, say, only spots above the threshold are decoded then less than 50% will be seen, but the average of those that *are* seen (red) will be higher than the average of the true (but unseen) /K distribution (blue).

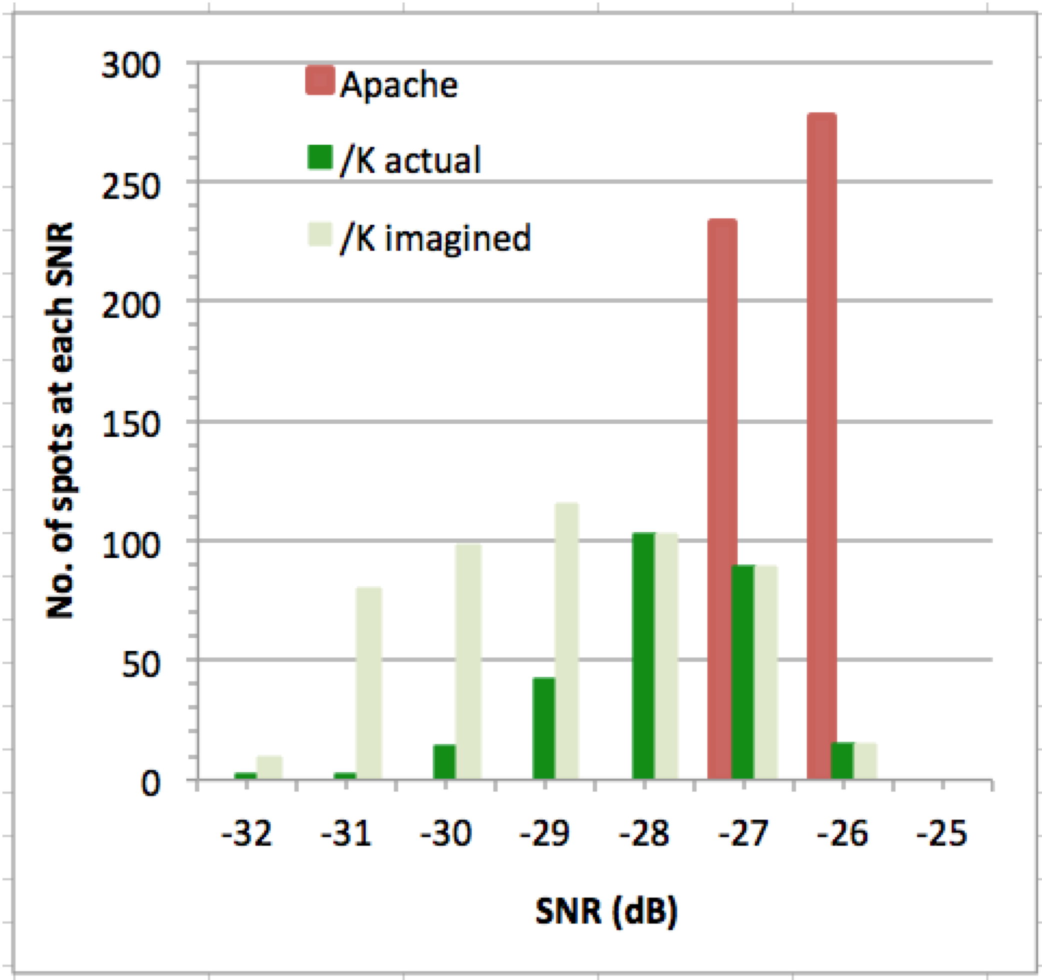

Figure 3.7. Actual Apache and /K distributions with the "imagined" /K distribution below the threshold.

Figure 3.7 shows the full Apache data set for -27 and -26 dB SNR and in dark green the observed distribution for /K, comprising 52% at a mean difference of 1.4 dB. In pale green is the "imagined" true /K distribution with 100% of the spots; at -29dB and below fewer spots are actually decoded than should be present; here the real SNR difference between the Apache and the imagined /K distribution is 2.5 dB.

4. Apache and /K comparison at high SNR - the case of WV0Q

WV0Q, located about 1.2km from N6GN, an occasional sender on 40m WSPR, provides a high SNR point of comparison for the three routes, not only for the difference in the mean SNR at the correct transmission frequency but also as a test of inter-modulation or other non-linear mechanism(s) that can give rise to numerous spurious decodes. We suspect that some of these may be related to re-sampling imperfections within the KiwiSDR.

The mean SNR difference at the correct transmission frequency (Apache - /K) is -6.33dB, from 28 spots on the Apache and 40 spots via the /K route, that is the SNR is **higher** from /K. With a standard error of 0.36dB for the difference, the -6.33dB is statistically highly significant.

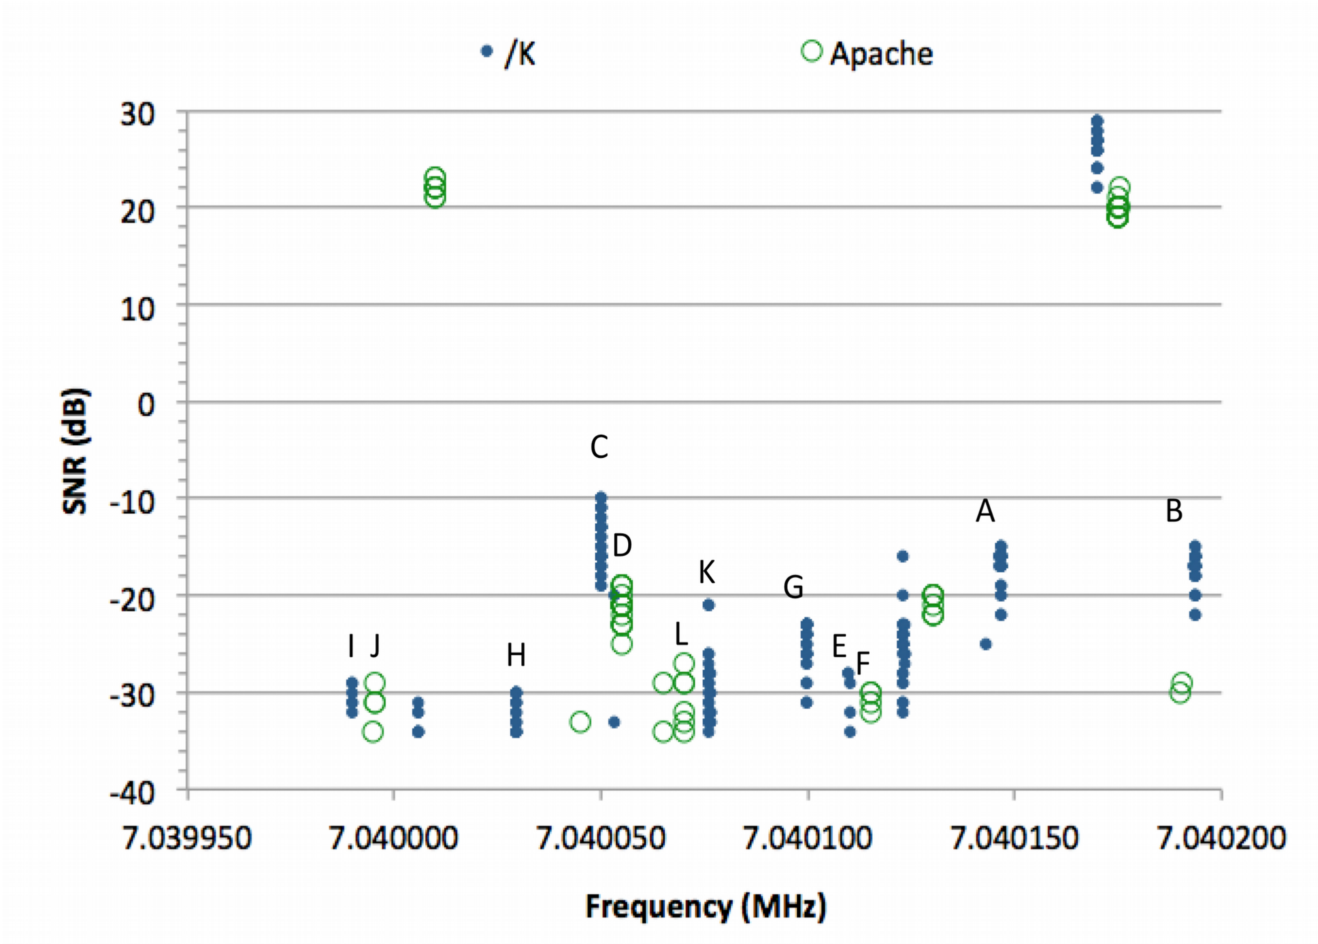

Figure 4.1. Scatter plot of SNR vs frequency for spots from WV0Q at a range of~1.2 km received via the /K route and via the Apache showing multiple spurs, some of which are common to both routes, and others not, as discussed in the text above. Note the Apache frequency has been offset by 5Hz for clarity

Figure 4.1 shows a scatter plot of the SNR against the decoded frequency for the /K and the Apache routes. For clarity the Apache frequencies have been shifted high by 5 Hz. There are a number of observations that can be made on this plot, noting that given a transmission frequency of 7.040170 MHz most of the higher-frequency spurs lie outside the WSPR band and therefore are not decoded so we cannot test for symmetry:

- Spurs **only** present via /K - consequently these are likely to be receiver-generated spurs. Spurs 'A' and 'B' are symmetrical and 23.5Hz either side of the transmission frequency. On average these spurs are 43.5dB down on the level at the correct frequency. Spur 'G' is 70.24Hz below the transmission frequency and 51dB down. Spur 'H' is 141Hz below and about 60dB down.
- Spurs present on **both** the Apache and /K - consequently likely to come from the transmitter.

  Spur 'C' for the /K is matched by spur 'D' for the Apache. For /K the frequency offset is 120.04Hz, undoubtedly due to full-wave rectified hum derived from the 110V AC supply modulating WV0Q's transmission. The average level for spur 'C' from /K is 41dB down.

  Spur 'E' for the /K is matched by spur 'F' for the Apache. For /K the frequency offset is 59.94Hz, the fundamental AC supply frequency. The average level for spur 'E' from /K is 57dB down.

  Spur 'I' for the /K is 175Hz below the transmission frequency, matched by spur 'J' for the Apache.
- Spurs where there appears to be a frequency offset between the Apache and /K - this is 'K' for the /K,

where there appears to be Apache spots but on the LF side by 6Hz whereas a HF shift of 5Hz was applied.

The root cause of the spurs at +/-23Hz is currently unknown. Qualitative examination of high SNR spots from KiwiSDRs at KPH has shown these +/-23Hz spurs occur with W6LVP on 40m and KJ6MKI on 630m.

A key question, in two parts, is whether these +/-23Hz spurs are always present, regardless of SNR (it's just that we only see them when the SNR exceeds the WSPR threshold by the spur suppression). The second part of the question is whether the suppression is constant at about 43dB or does it depend on SNR? The complicating factor is the use of SNR as a proxy for absolute signal level; if the spur suppression does depend on signal level but the changes we see in SNR reflect changing noise rather than changing signal level then we could be at risk of drawing a false conclusion.

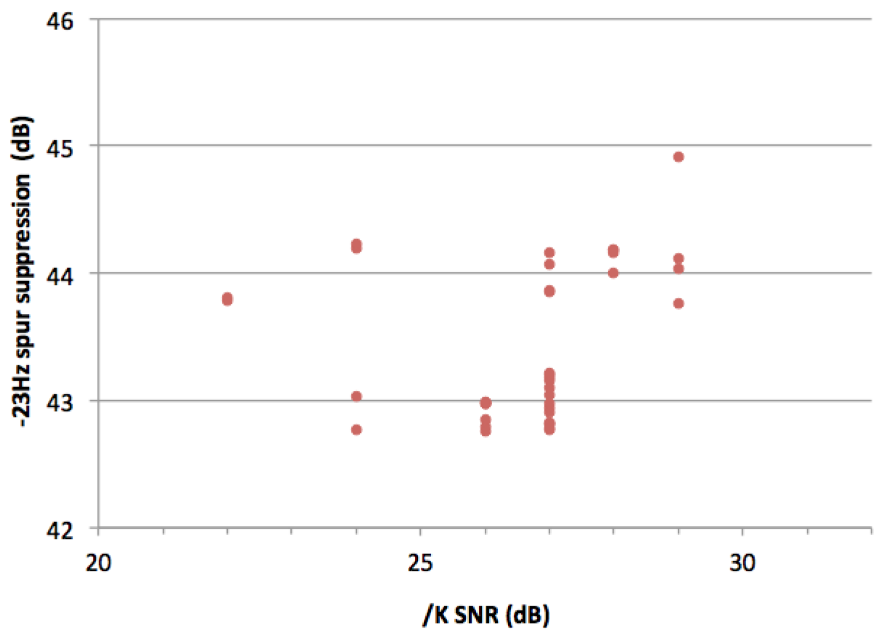

Figure 4.2. Suppression level of the -23Hz spur for changing noise rather than changing signal level then WV0Q as a function of /K SNR. There is no clear dependence, but as we cannot separate changes in S from changes in N we could be at risk of drawing a false conclusion.

This is probably the case with WV0Q at N6GN. The ground wave signal level from a distance of 1.2km could reasonably be expected to be constant. From the data set, the -23Hz spur suppression level against SNR, for 40 spots, is shown in Figure 4.2. Because of the 1dB quantization for both axes many spots would lie on top of each other, so a random +/-0.25dB jitter has been added to the suppression level. There is no obvious trend here, but we do not know whether the changes in SNR are due to S or to N. What we do know is that there was a spur present with every WV0Q spot.

4.1 W6LVP at KPH on 40m

This sub-section addresses the +/-23Hz spur question using data from the KiwiSDR at KPH on W6LVP spots on 40m gathered from the wsprnet.org database as an aid perhaps to understanding the results at N6GN. W6LVP is about 545 km from KPH and received spots exhibit a wide range of SNR via ionospheric propagation. We can make the assumption that the range of SNR from W6LVP greatly exceeds the variation in noise level at KPH, given its low noise location.

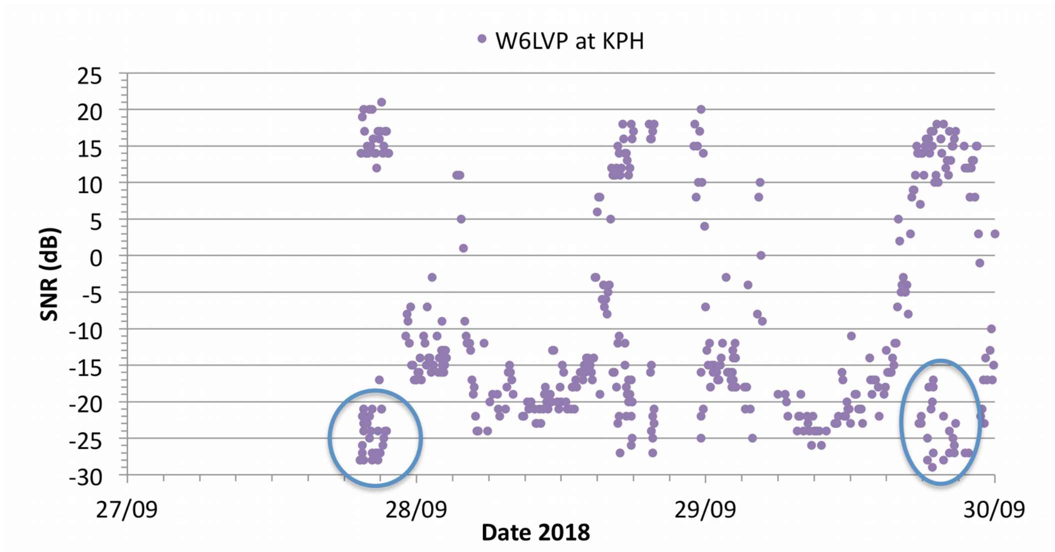

Figure 4.3. Time series of 40m W6LVP SNR at Kph where propagation results in high daytime SNR leading to visible spurs, e.g. within the two ovals

Unfortunately, W6LVP's frequency hops around the WSPR band and so the equivalent plot to Figure 4.1 is not informative. However, a time series, Figure 4.3, is helpful in that we can see the periods when spurs are present. As for WV0Q, there are many spur frequencies so this analysis has been limited to -23Hz spurs,

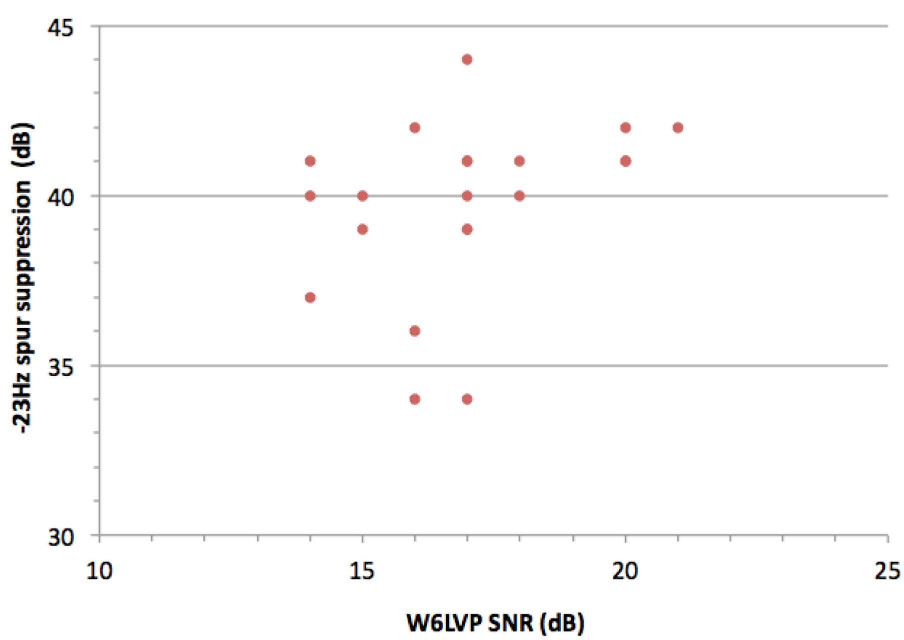

Figure 4.4 Suppression level of the -23Hz spur forW6LVP at KPH on 40m as a function of SNR.

.

Figure 4.4. A linear fit gave a coefficient of determination (R^2) of 0.12, explaining only 12% of the variance, so there is no clear trend of suppression level with SNR.

What can be said is that we see no -23Hz spurs when the main spot SNR is +13dB or below. With a mean spur suppression in this case of 40dB, present at and above +14dB SNR, the absence of spurs in the database is NOT because there are no +/- 23Hz spurs, but because the SNR of the spurs given the main spot SNR is below the WSPR threshold. This conclusion is supported by the observations of the 630m spots of KJ6MKI at KPH, Figures 4.5 and 4.6.

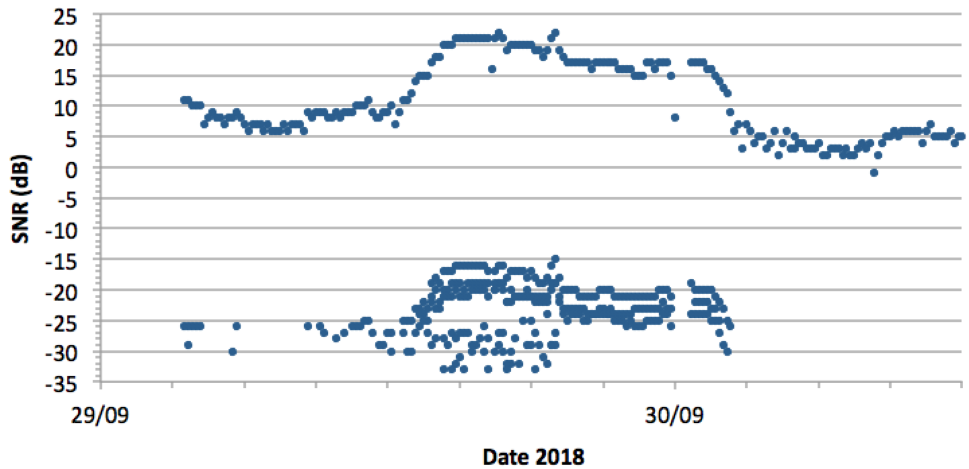

Figure 4.5. Time series of 630m KJ6MKI SNR at KPH where propagation results in high daytime SNR leading to spurs being visible – all the points with an SNR below -15dB.

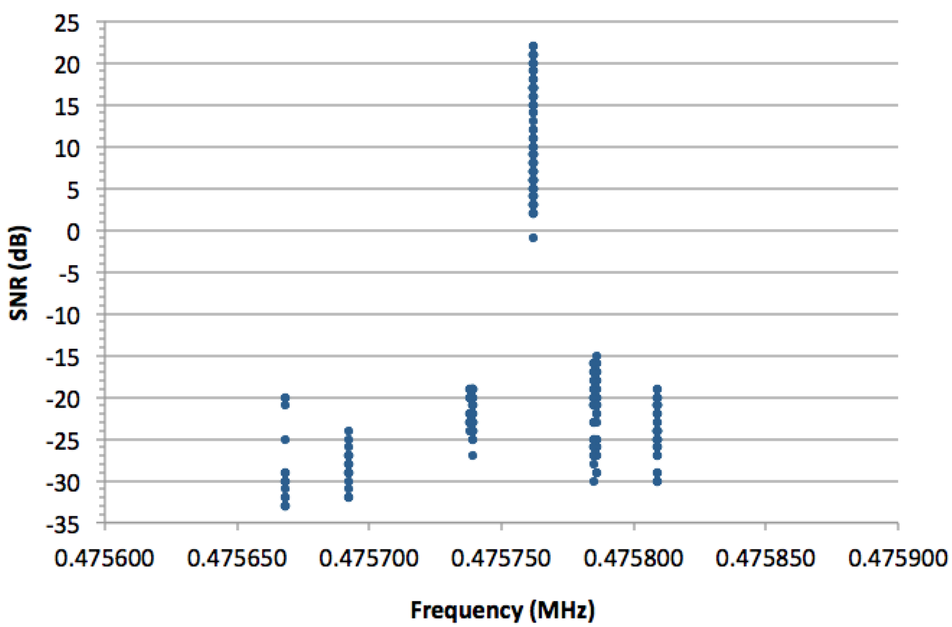

Figure 4.6. Scatter plot of SNR vs frequency for spots from KJ6MKI on 630m received at KPH showing spurs.

.

5. Results: Apache vs. /Kiwi

The KiwiSDR/Beaglebone is known to run out of time to process received spots within a 2-minute WSPR window if there are many spots received. In the data set used for this analysis there were 2768 /Kiwi spots to 4310 from the Apache, a ratio of 64.2%. However, this "out of time" argument may not be as simple as it first appears, in that the decoding algorithm in the WSPR extension to the KiwiSDR decodes a high percentage of spots with higher SNR, Figure 3.5. Our understanding is that the KiwiSDR WSPR extension uses a single-pass decoder, that is, there is is no second pass after coherent removal of the signals decoded in the first pass[2]. The two-pass decoder from K9AN has been included in wsjt-x releases since version 1.6.

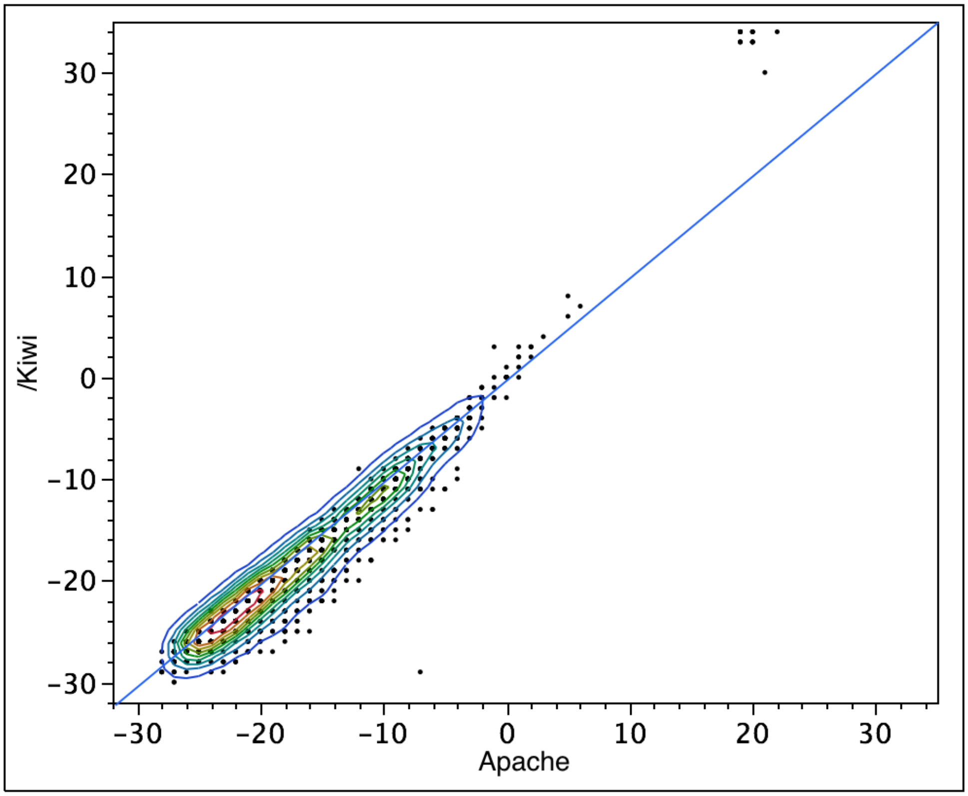

Figure 5.1. Scatter plot of SNR via routes /Kiwi and Apache. The outliers are discussed in the text.

As for SNR performance, the essential story is the same as for the /K route, with the Kiwi showing lower SNR for the vast majority of spots, but the difference reducing to zero around an SNR of 0dB and the Kiwi WSPR extension showing higher SNR for a cluster of WV0Q spots at an Apache SNR of around +20dB, Figure 5.1. Note that in this figure the /Kiwi outliers are quite different from those via the /K route. Except for one spot the wsprnet.org database listed WV0Q spurs, hence a low SNR (outliers A in section 3). For /Kiwi there are 19 WV0Q spots at an Apache SNR of +18-24dB where the average /Kiwi SNR is 13.8dB higher.

2 See https://valentfx.com/vanilla/discussion/comment/3982/#Comment_3982

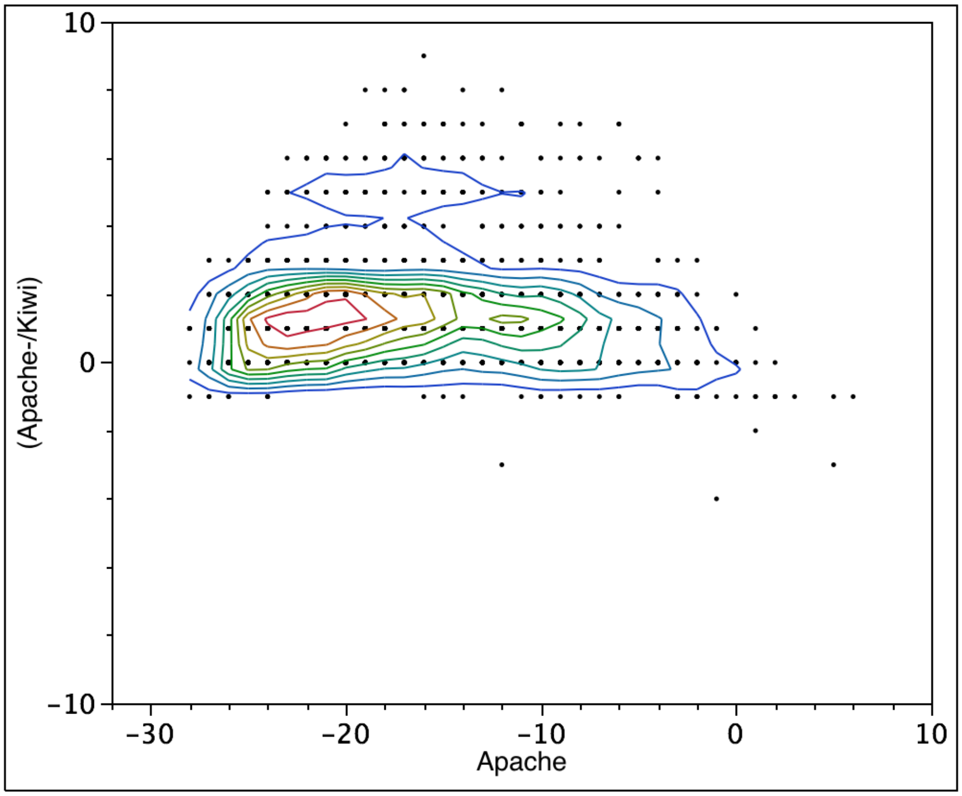

Figure 5.2. Scatter plot of SNR difference(Apache-/Kiwi) against SNR Apache.

Figure 5.2 shows that the diffuse cloud of spots in Figure 3.2 for the /K route where the /K SNR was higher than the Apache is missing; there is a very sharp cut-off at an SNR difference of -1dB.

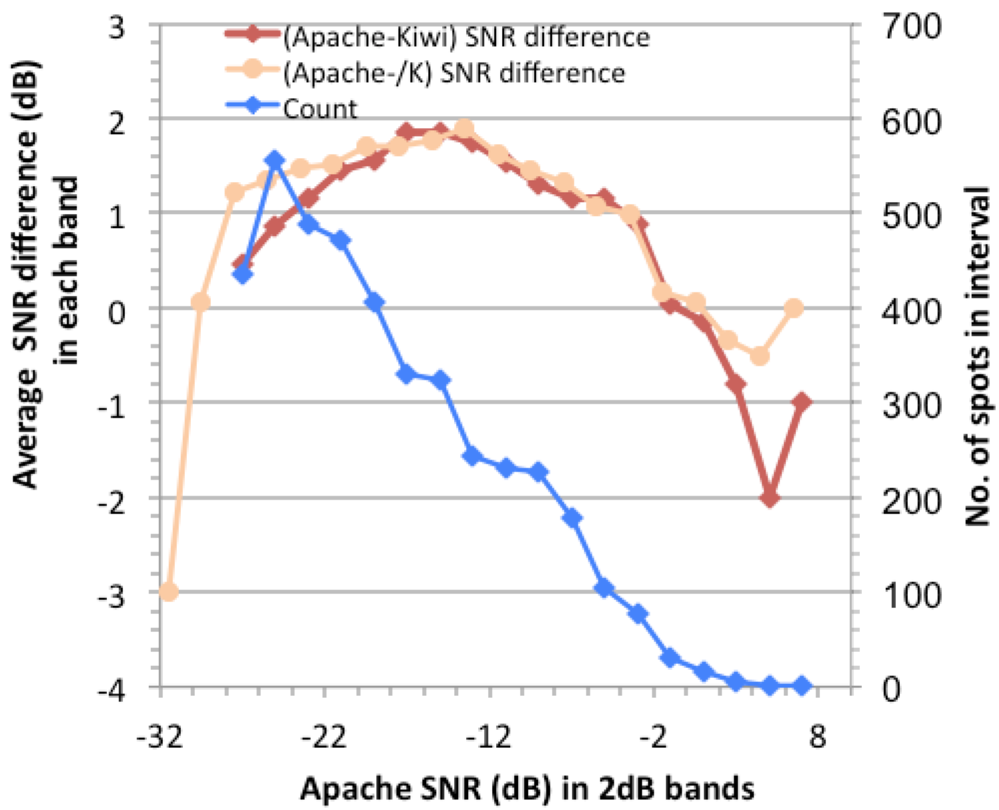

Figure 5.3. Average SNR difference (Apache-/Kiwi) in 2 dB bins, together with the number of spots per bin. For comparison the SNR difference for (Apache-/K) is also shown.

The shape of the curve of SNR difference against Apache SNR is almost identical for the /Kiwi and the /K, Figure 5.3, for where there are over 50 spots per SNR bin; the one part where there is noticeable deviation is at Apache SNR below -23dB where the /Kiwi route shows a smaller SNR difference.

The conclusion therefore is that the SNR difference and the shape of the SNR difference with SNR are set within the KiwiSDR and its WSPR extension, confirming the observations of N6GN. Additionally, although these data are specific to 40m, the general disparity between KiwiSDR-derived spots and Apache SDR spots holds across different amateur bands from 630m to 20m. It appears that the the degradation of delivered SNR is not correlated with input frequency, it has the characteristics of a post-downconversion effect on all SNRs at any LF-HF input, thus it almost appears as an audio phenomenon. Whether there is a correlation with total signal power within the downconverted spectrum remains to be determined.

6. WV0Q - Comparison of constancy of signal level and SNR

In our analysis of WSPR spots received at N6GN from WV0Q in section 4 we could not separate whether changes in SNR over the 1.2km ground wave path were due to changes in signal level or changes in noise. Consequently, N6GN and WV0Q collaborated to perform an experiment to identify the contributions of variations in signal level and noise to variations in SNR. The 'S' meter readings of the Apache receiver (Angelina) were recorded every 50ms and 20 measurements averaged per recorded data point, giving points every 1.25s. N6GN's approximate start time was corrected by -31.5 seconds using the start of WV0Q's WSPR transmissions. These measurements are in a 1kHz bandwidth centered on the WSPR band center.

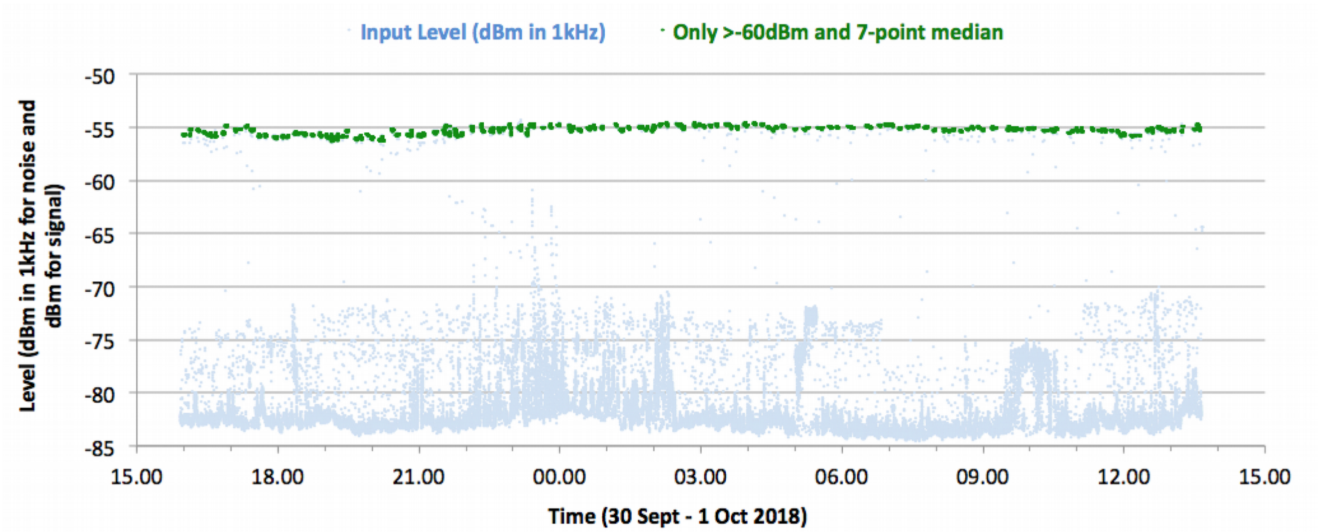

Figure 6.1. Time series of raw input level measurements at 1.25s intervals (light blue) together with the level ofWV0Q's main transmission after a 7-point median filter (green)

Figure 6.1 shows the (about 63,000) raw input level measurements as faint blue dots and, as dark green points above -60dBm after a 7-point median filter, the signal level from WV0Q. The "noise" level includes WSPR and other transmissions (e.g. CW, RTTY) within the 1kHz measurement bandwidth as well as any local interference. We obtain our best estimate for the appropriate noise level to compare with the WSPR SNR by taking the median of the noise measurements in the gap between 112 and 119 seconds after the start of each two-minute interval. It should be noted that the effective antenna factor, that of the short dipole and the following preamplifier is not precisely known at this time, thus absolute levels of both signal and noise power are likely skewed. Even if they were not, the effects of real earth ground on efficiency and antenna pattern at both ends are not known.

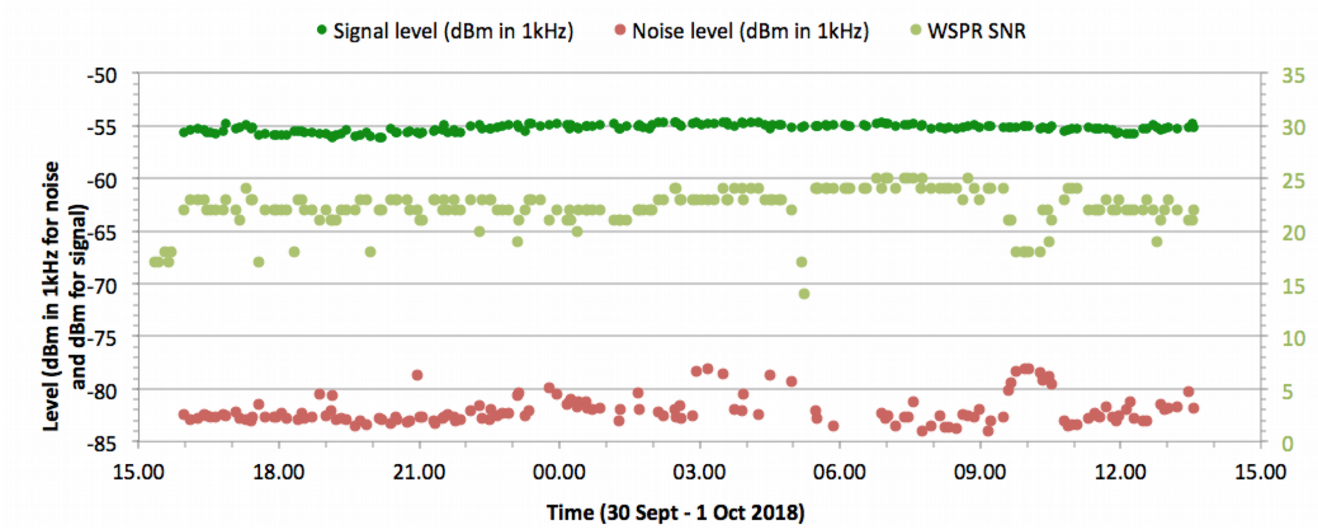

Figure 6.2. Time series of the level of WV0Q's transmission (dark green) together with WSPR-reported SNR in light green and our best estimate of noise and interference in the 8 seconds immediately after the end of each transmission. Gaps in the latter are likely due to interference within the 1kHz noise measurement bandwidth but outside the 200Hz WSPR bandwidth leading to those noise measurements being removed by the filter.

Figure 6.2 shows this best estimate noise level (red) at the end of each WV0Q transmission. The gaps are due to interfering signals being present in the 1kHz band. WSPR SNR is in light green, and the signal level as measured by the Apache in dark green. By visual inspection, and taking the standard deviations:

- The SNR variation (1.84dB) is much greater than the variation of measured signal level (0.36dB).
- The SNR variation is greater than the variation in noise level (1.34dB), but only by about 40%. The SNR variation also includes the real variation in signal level and the quantization noise of the WSPR SNR measurement.

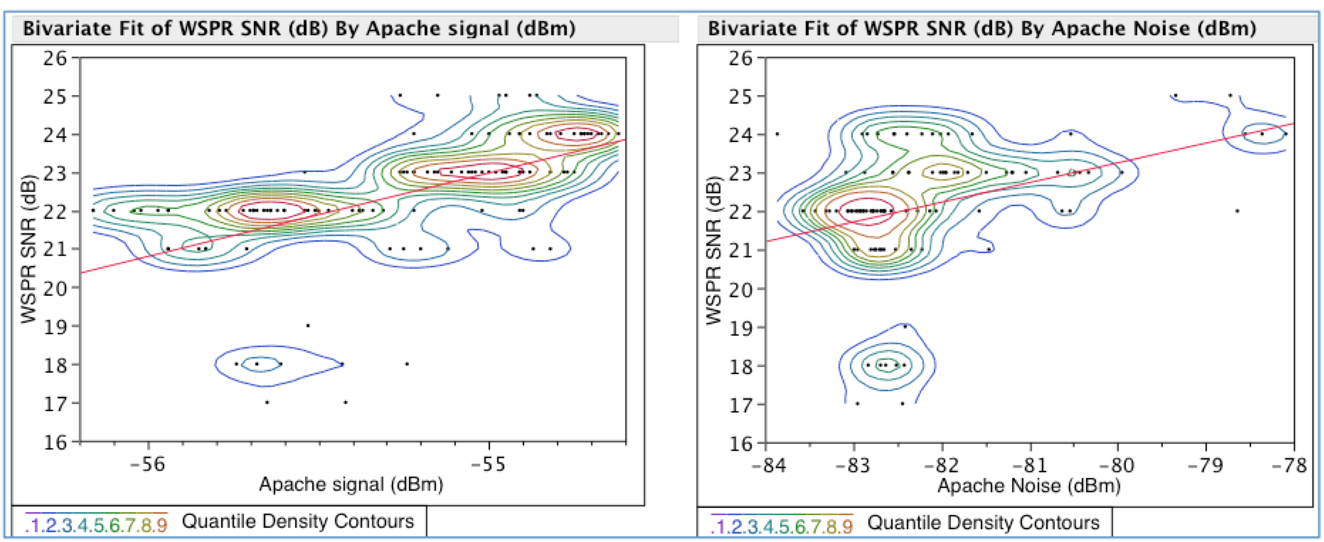

Figure 6.3. Scatter plots of WSPR SNR against Apache signal level (left) and noise in a 1kHz bandwidth 8 seconds after the end of each WSPR transmission.

Apart from the period around 1000UTC in Figure 6.2 there appears to be little visual correlation between the variations in WSPR SNR and the noise plus interference level in the following 8 seconds. The scatter plot of WSPR SNR with Apache signal level during WV0Q transmissions shows the expected form, Figure 6.3 (left) although the slope is not 1 but just over 2, which does raise a question. Given the WSPR SNR is quantized at 1dB the horizontal scatter in Apache signal level is expected, while there is a monotonic increase in the peak of the Apache signal level as SNR increases from 22 to 24dB, for a 1:1 relationship there should be a greater increase for the Apache (or less of an increase for the WSPR SNR). The linear fit explains 30% of the variance, leaving 70% unaccounted for, and our assumption is that this is due to variations in the noise level.

For the scatter plot of WSPR SNR against our best estimate of noise plus interference, Figure 6.3 (right) the result is unexpected in that the least squares fit suggests WSPR SNR increases as the noise level increases, and only 13% of the variance is explained. More thinking is needed!

It is believed that this record of the S meter of the Angelia (Apache) is a reasonably accurate record of absolute signal strength at the receiver's input. Of course this number is offset by the antenna factor of the active antenna and by the in situ characteristics of ground, foliage etc. between N6GN and WV0Q. Very roughly the antenna factor is probably at least several dB positive since there is about 6 dB of gain in the preamplifiers, perhaps as much in the CAT5 driver/receivers that provide the transmission line and who-knows-what in the ground effect at each station and between them. Both of antenna patterns are de-steered by imperfect grounds but there is also a half-hemisphere gain counteracting that. At a separation of 1.2 km, in free space one could expect a path loss of about 37+20log(.8) + 20log(7) = 52 dB. If both antennas were isotropic and co-polarized one would expect +20-53= –33 dBm. Measured signal levels are much lower than that, with values at least 20 dB and perhaps 40 dB lower if antenna factors and grounds are taken into account. Ground shielding and beam de-steering are probably very significant. Even though foliage loss no doubt persists down at 40m (compare with measurements performed at 10m in an article by N6GN published in QEX magazine) these are probably not nearly so large. The AGC on the Angelia was set high enough that it should not have been bumped by WV0Q's signal. The difference of AGC actions between the Angelia and KiwiSDR, which had unknown AGC level setting, could have had an effect but it would seem it seems unlikely that his accounts for the differences observed.

For now, we conclude that over the period of these measurements the WV0Q signal level at N6GN was essentially constant with a mean reported level of -55.2dBm with a standard error of the mean of 0.03dB and that the changes in WSPR-reported SNR from the mean of 22.2dB were mostly due to changes in noise and interference level. However, we have not been able to calculate a robust quantitative relationship between the SNR variations and noise plus interference due to the shortcomings of the noise measurement.

Summary

From our measurements and interpretation it appears that the KiwiSDR may have some additional room for improvement. We have found evidence that in comparison to other SDRs the delivered SNR may not yet be as good as it might be. Additionally, it appears that spurious sidebands are generated within the KiwiSDR which can result in false decodes on WSPR. We hope that this documentation may help identify and perhaps even correct these characteristics and result in an already fine measurement and communications SDR becoming even better.