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

Part 1 of 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. Part 2 describes further measurements and analyses after significant KiwiSDR code upgrades which were precipitated by Part 1.

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

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

than the WSPR extension. We present some of these observations below.

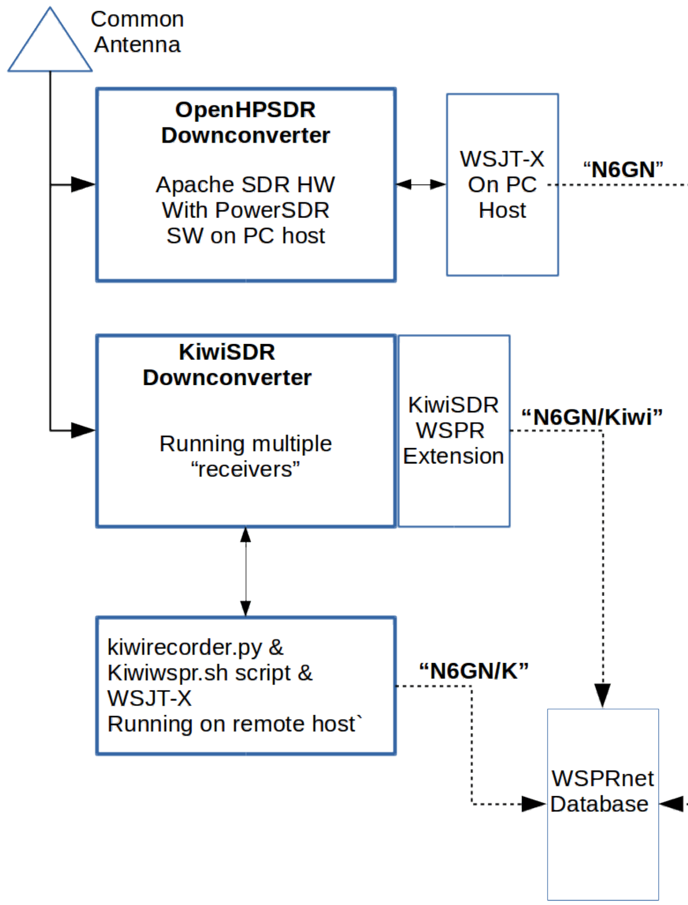

Illustration 1 The three methods of generating spots to the WSPRnet.org database include two paths involving KiwiSDR downconversion and a third using an Apache "Angelia" SDR.

1.1. *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.
4. Scatterplots of SNR against frequency for the Apache and /K paths to show the spurs present on WSPR signals with SNR above +13dB for stations at constant frequency and time series of SNR where stations had a strategy of frequency hopping.

Part 1: Initial Analysis

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.

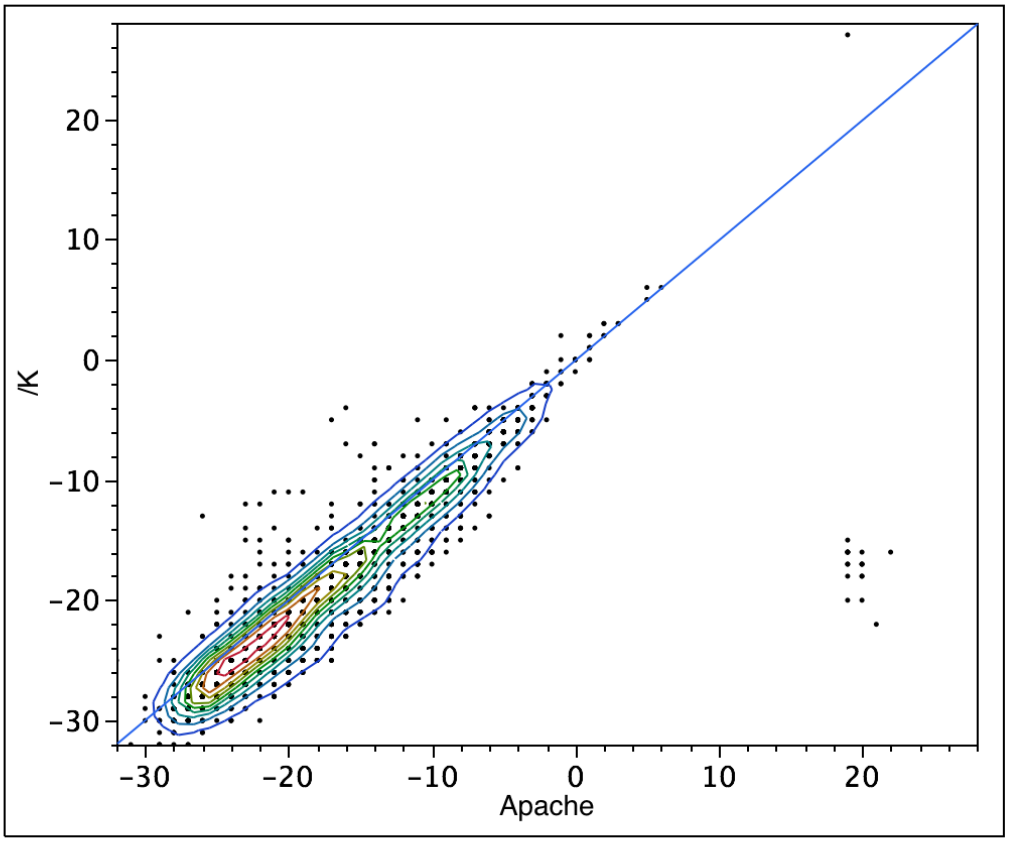

Figure 3.1 Scatter plot of SNR via routes /K and Apache. The outliers are discussed in the text.

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.

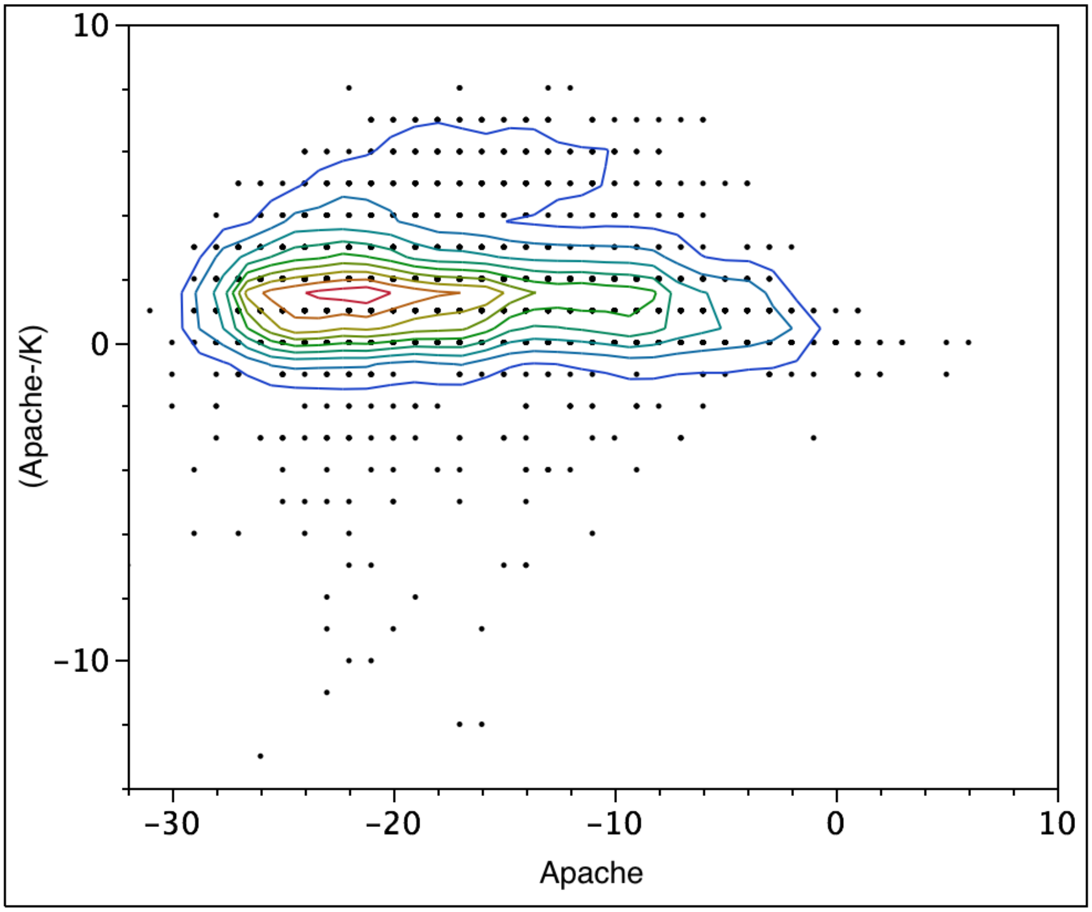

Figure 3.2.Scatter plot of SNR difference(Apache-/K) against SNR Apache.

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.

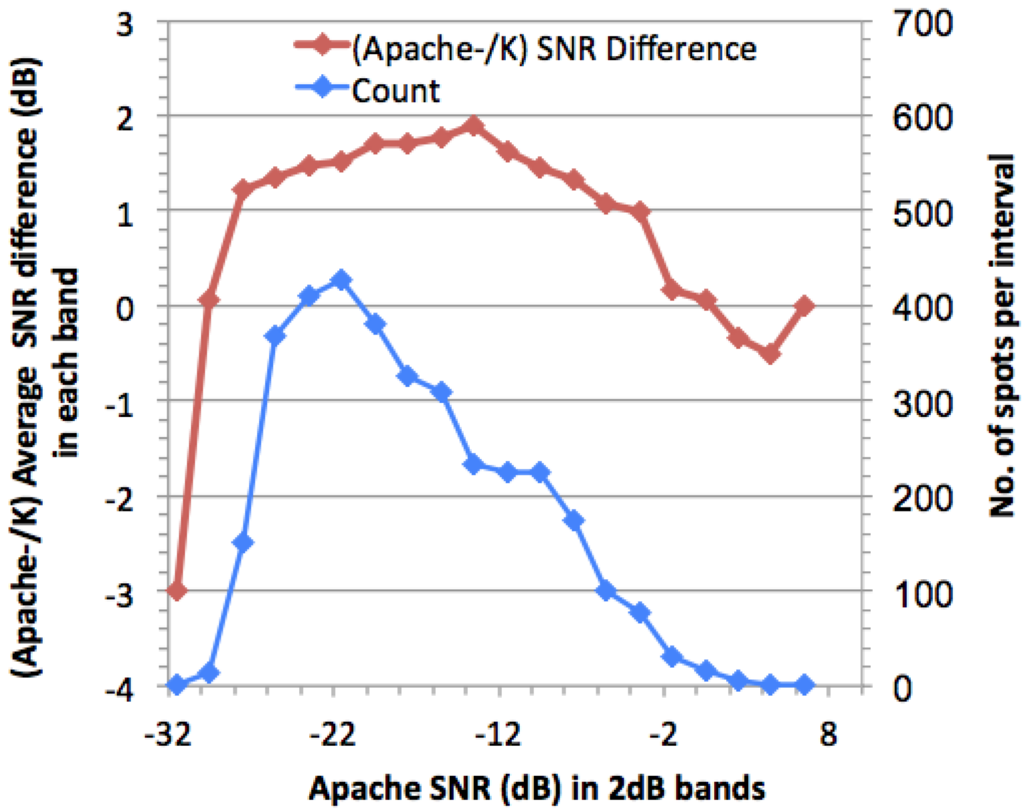

Figure 3.3. Average SNR difference (Apache-/K) in 2dB bands together with the number of spots in each bin. The data for bins centered on -31 and -29dB are most probably anomalous due to proximity to the decoding threshold.

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.

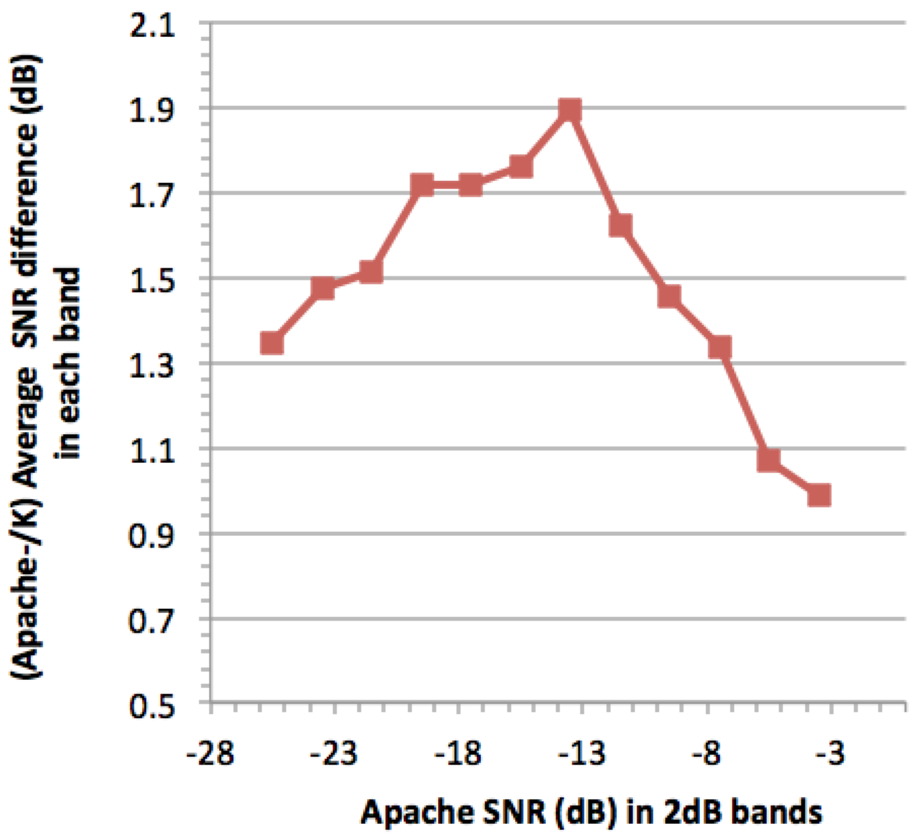

Figure 3.4 Average SNR difference (Apache-/K) in2dB bands where there were >75 spots.

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.

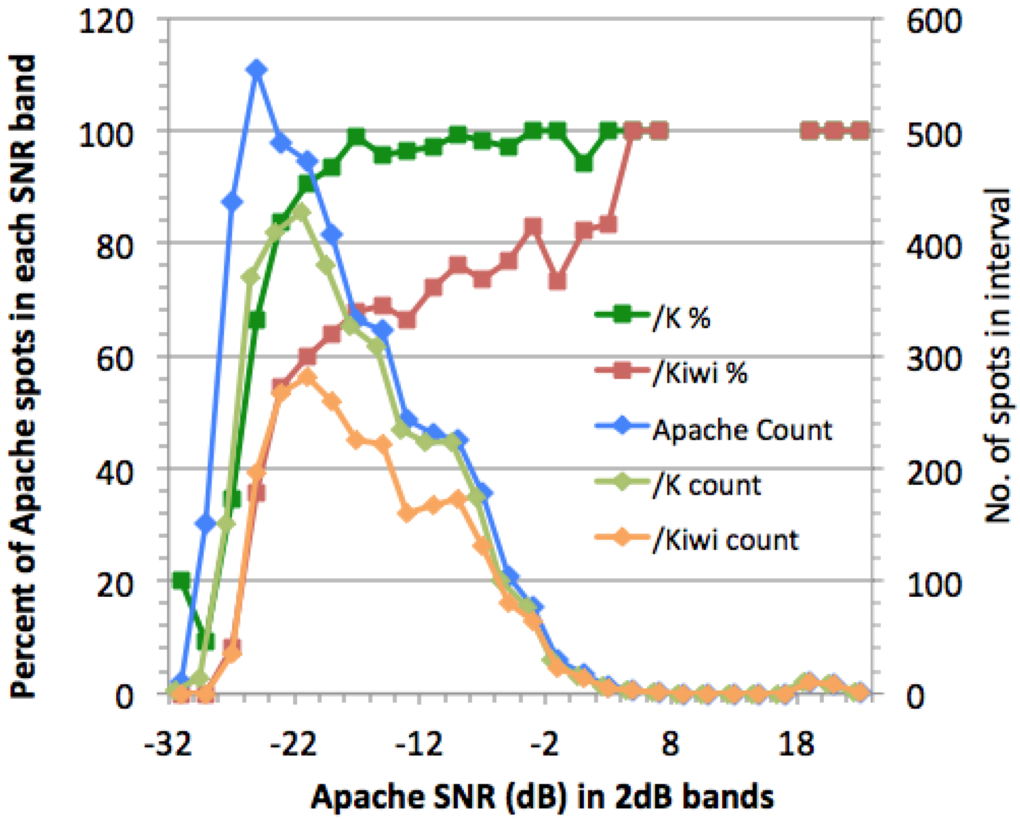

Figure 3.5. Number of spots in 2dB bands for the Apache, /K and /Kiwi, also expressed for the /K and/Kiwi as a percentage of the Apache spots per SNR band.

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.

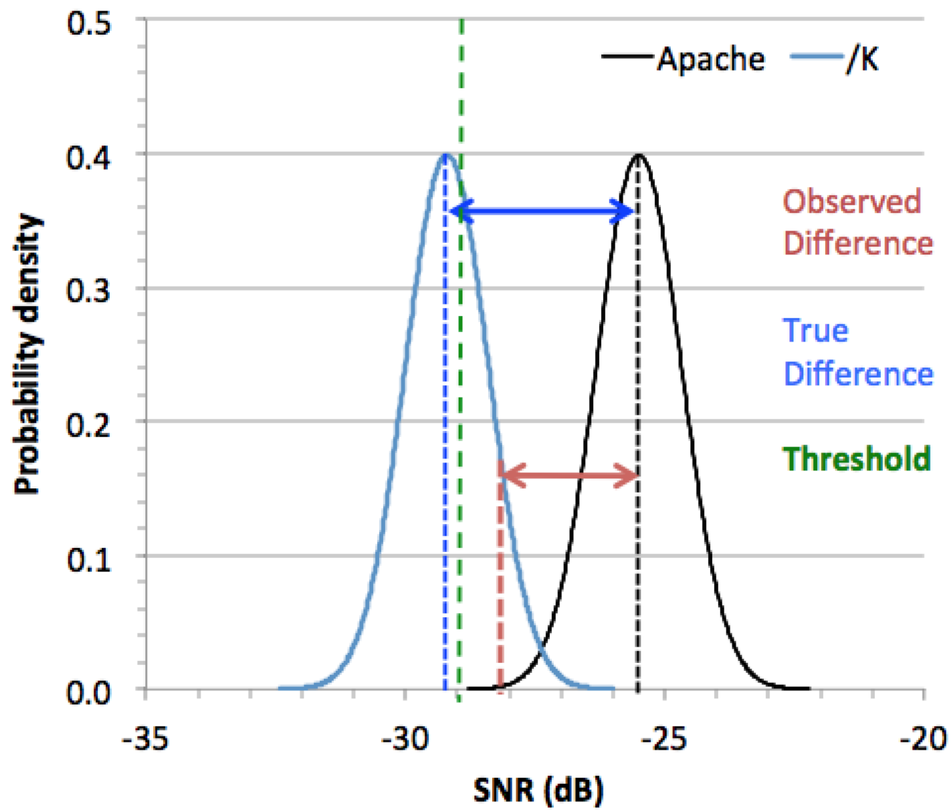

Figure 3.6. Diagram to help explain the impact of the WSPR decoding SNR threshold of about -29 dB on the calculation of SNR difference.

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).

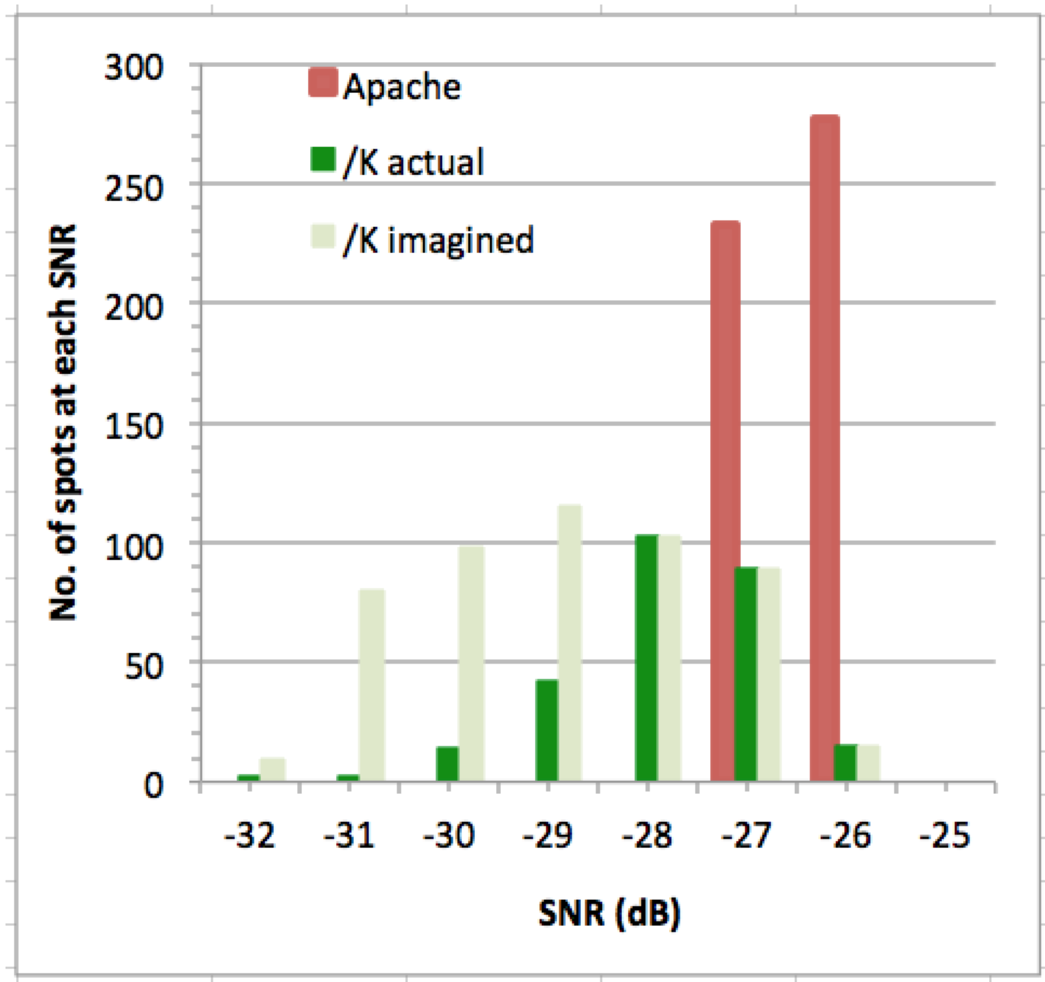

Figure 3.7. Actual Apache and /K distributions with the "imagined" /K distribution below the threshold.

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.

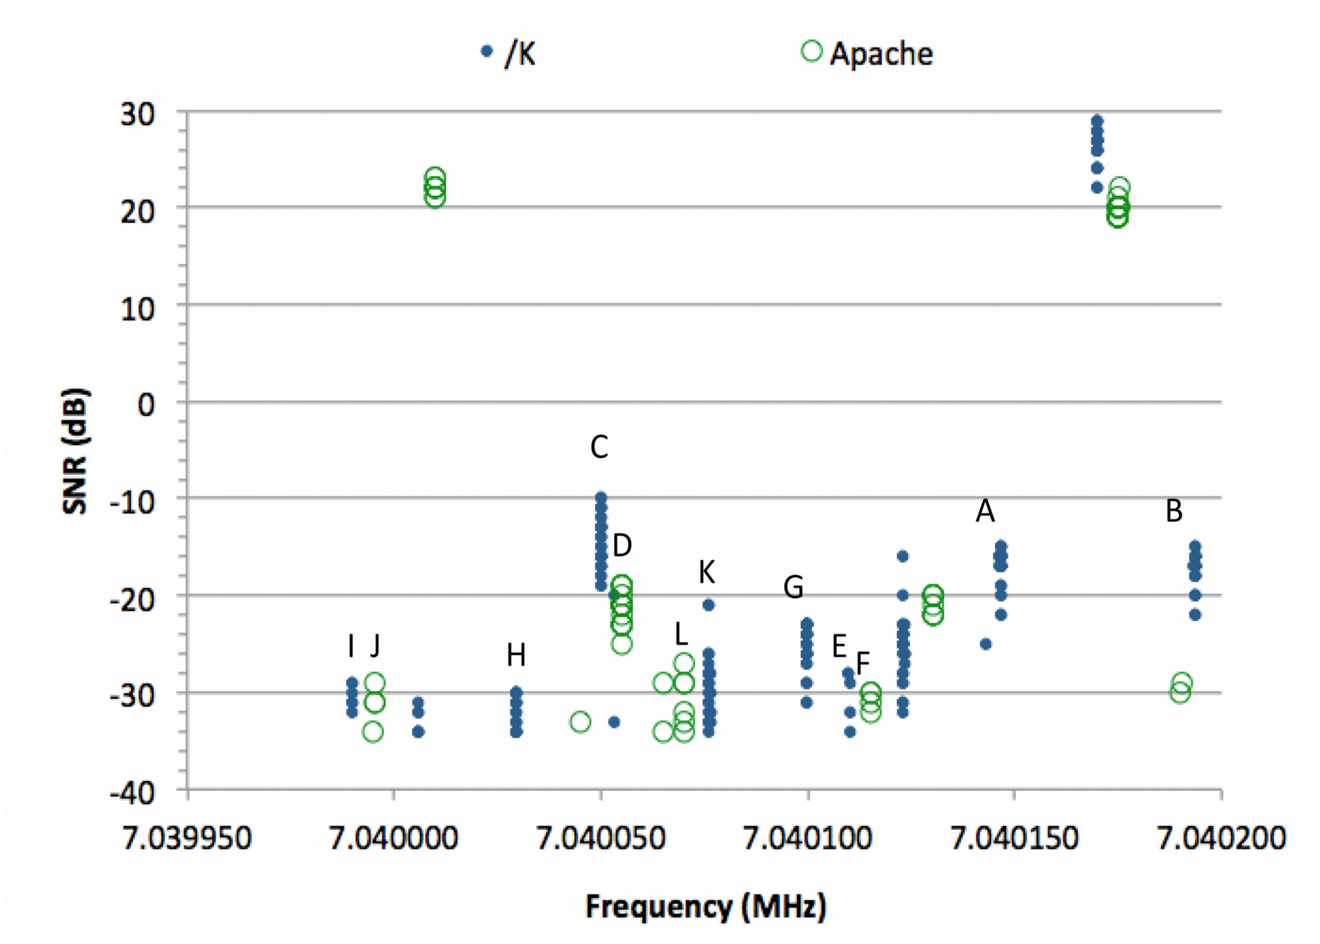

Figure 4.1. Scatter plot of SNR vs frequency for spots from WV0Q at a range of~1.2 km received via the /K route and via the Apache showing multiple spurs, some of which are common to both routes, and others not, as discussed in the text above. Note the Apache frequency has been offset by 5Hz for clarity

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.

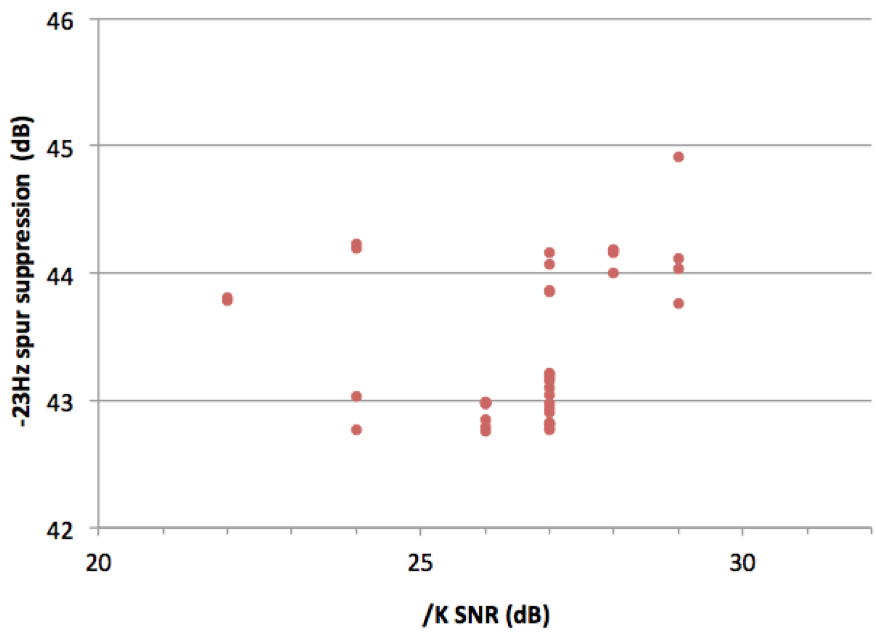

Figure 4.2. Suppression level of the -23Hz spur for changing noise rather than changing signal level then WV0Q as a function of /K SNR. There is no clear dependence, but as we cannot separate changes in S from changes in N we could be at risk of drawing a false conclusion.

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.

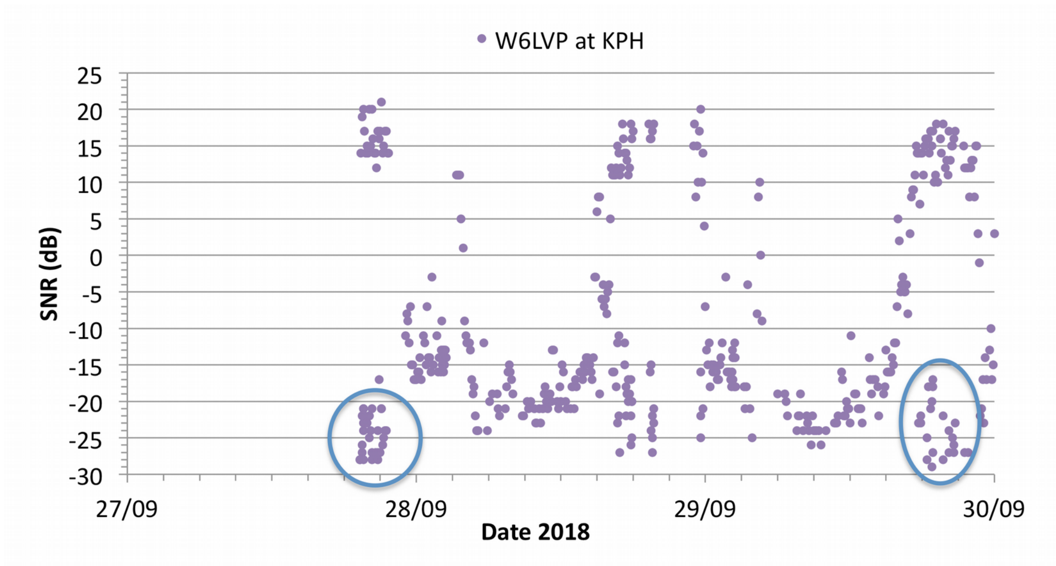

Figure 4.3. Time series of 40m W6LVP SNR at Kph where propagation results in high daytime SNR leading to visible spurs, e.g. within the two ovals

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,

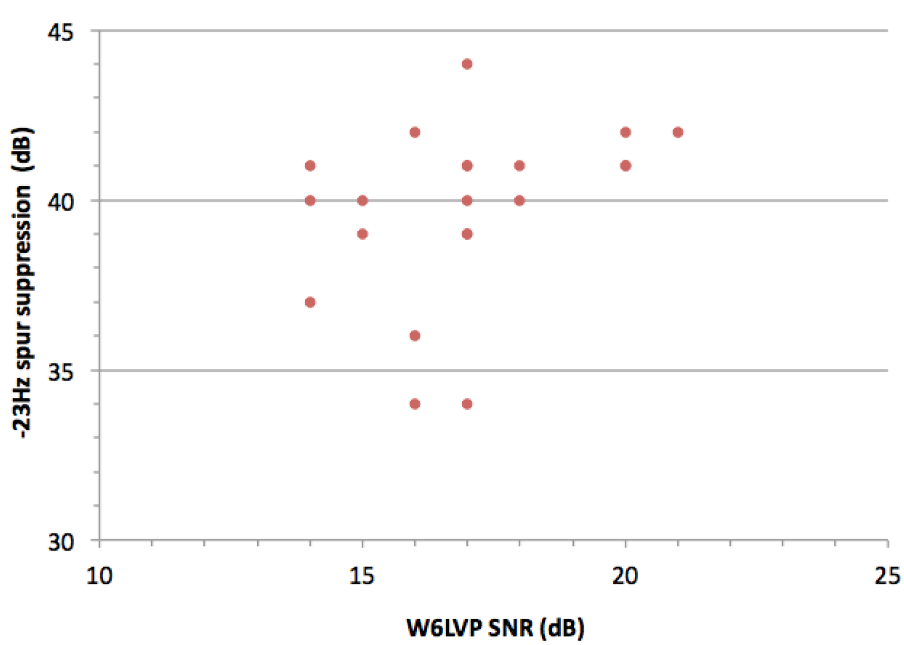

Figure 4.4 Suppression level of the -23Hz spur forW6LVP at KPH on 40m as a function of SNR.

.

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.

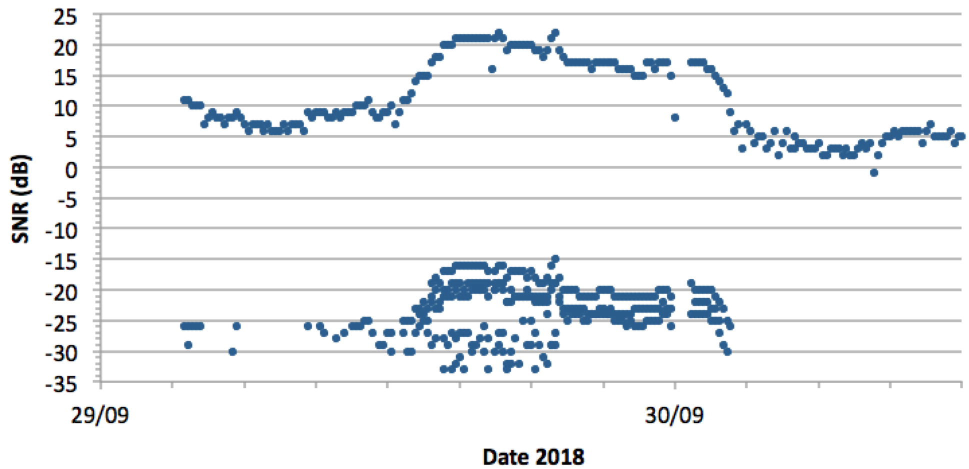

Figure 4.5. Time series of 630m KJ6MKI SNR at KPH where propagation results in high daytime SNR leading to spurs being visible – all the points with an SNR below -15dB.

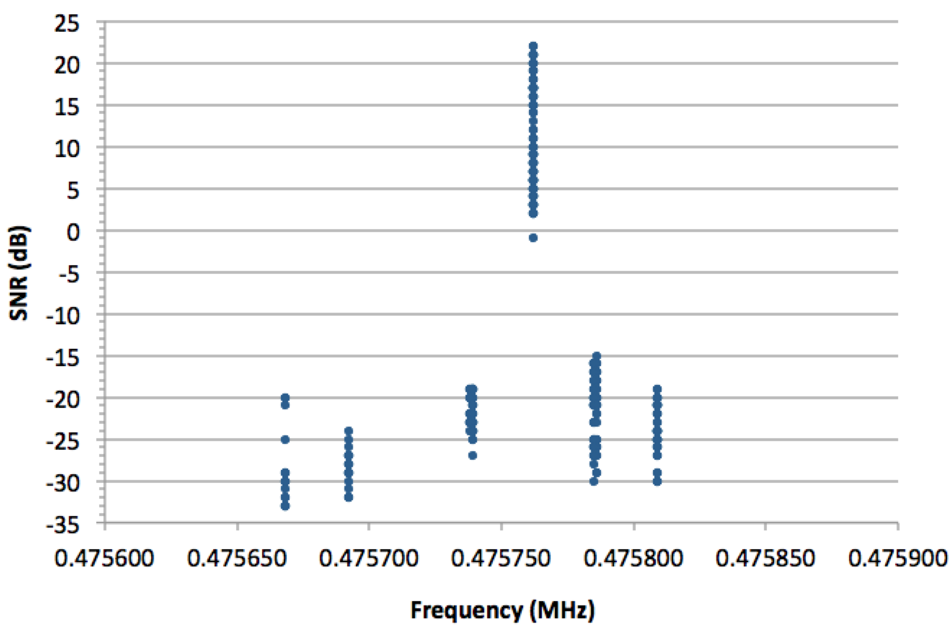

Figure 4.6. Scatter plot of SNR vs frequency for spots from KJ6MKI on 630m received at KPH showing spurs.

.

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.

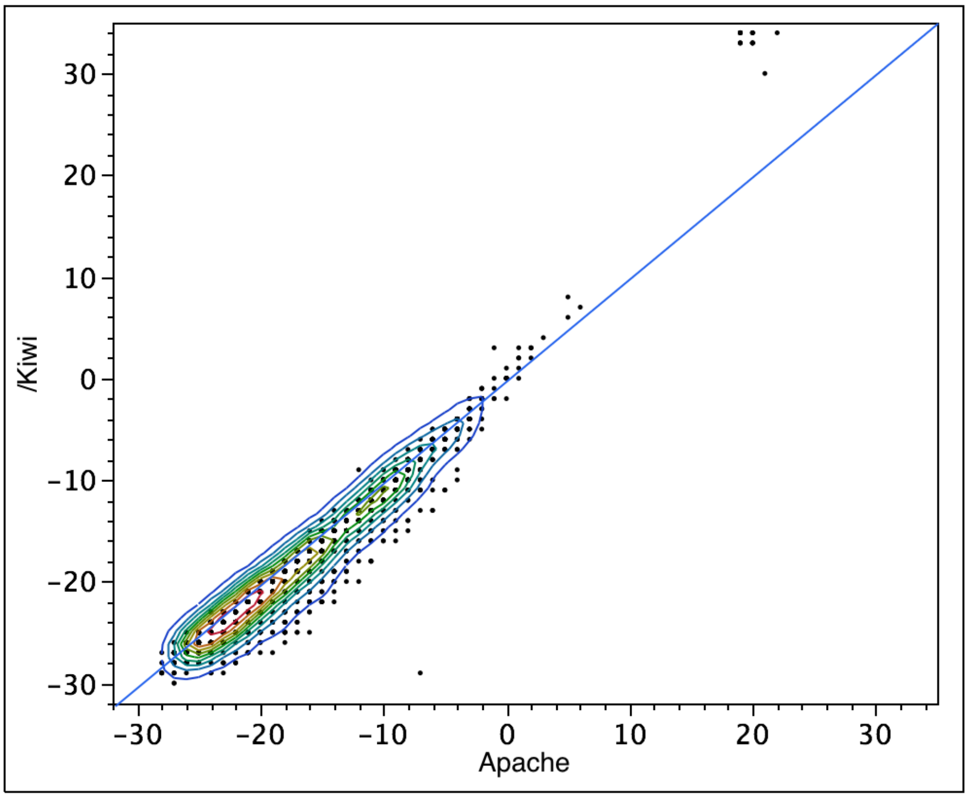

Figure 5.1. Scatter plot of SNR via routes /Kiwi and Apache. The outliers are discussed in the text.

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

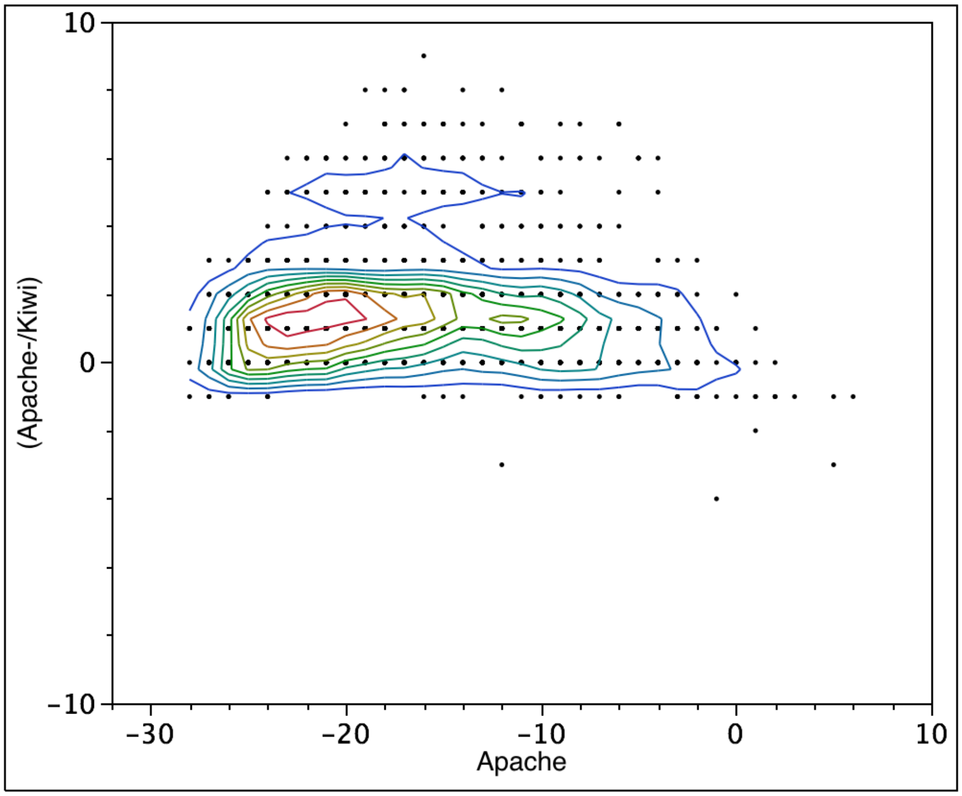

Figure 5.2. Scatter plot of SNR difference(Apache-/Kiwi) against SNR Apache.

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.

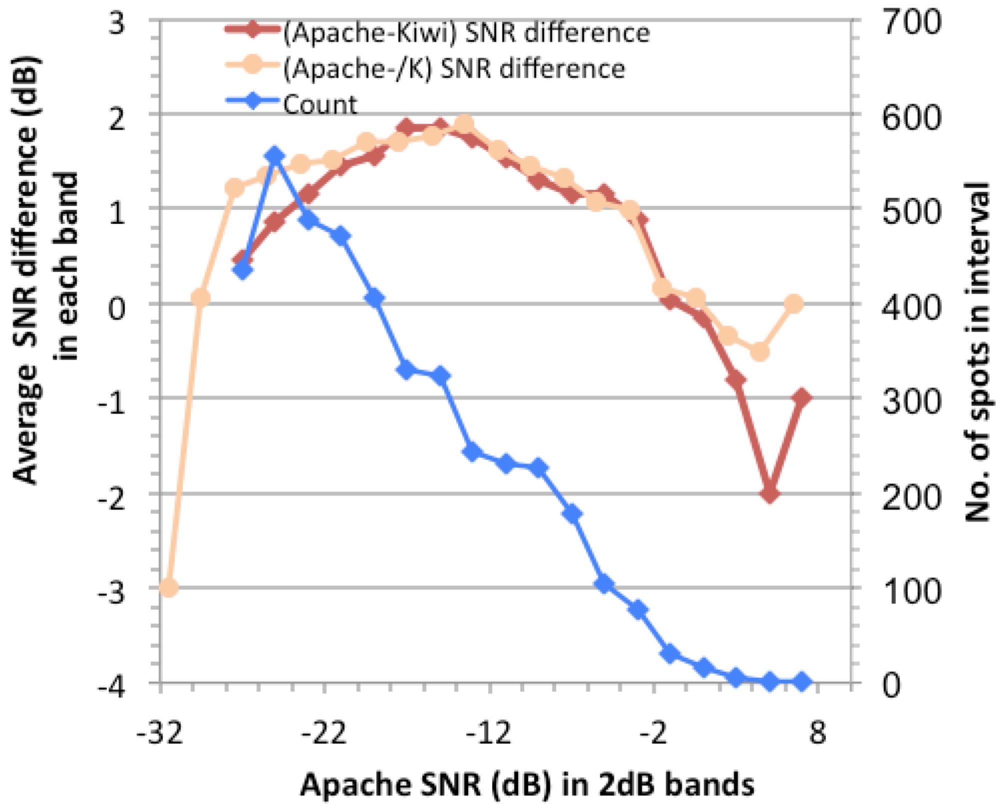

Figure 5.3. Average SNR difference (Apache-/Kiwi) in 2 dB bins, together with the number of spots per bin. For comparison the SNR difference for (Apache-/K) is also shown.

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.

Part 2: Post-upgrade Analysis

6. SNR Comparisons Apache SDR (Apache) vs. KiwiSDR (/K)

Since the time Part 1 was written, KiwiSDR software was upgraded to address issues raised. The result appears to have been very significant reduction or elimination of the 23 Hz spurious signals as well as degraded SNR in the KiwiSDR as compared to the Apache SDR. A large amount of data was again taken with both SDRs connected to the same source antenna. This new data has been examined to compare performance since the software upgrades, KiwiSDR v1.242 and beyond.
Figure 6.1 combines a histogram (left), an outlier box plot, and a normal quantile plot (right , in effect a cumulative probability plot). 10560 out of 11261 spots (93.8%) showed an SNR difference of -1, 0 or +1dB. This is excellent; the systematic SNR difference described in section 3 has been eliminated by the software change that cleared the +/-23Hz spurs problem.

However, the dashed red line in Figure 6.1, which shows the expected range of values for a normal distribution, is far from the data for SNR differences more negative than -3dB. Those spots (the black lines at each SNR difference value) lie outside the Lilliefors confidence bounds (dotted red lines). Therefore, there must be a mechanism at play that is responsible for the greater-than-statistical variation. Note that the same is not true for the SNR differences above +3dB – they are close to the red lines. So it looks as if we have two distributions, one, which we will call the Normal Core, comprises the data within the Lilliefors bounds in Figure 6.1, the second comprises those spots more negative than -3dB that we will call the Extended Tail. Next we will set out characteristics of the spot data and Extended Tail with the aim of finding its root cause.

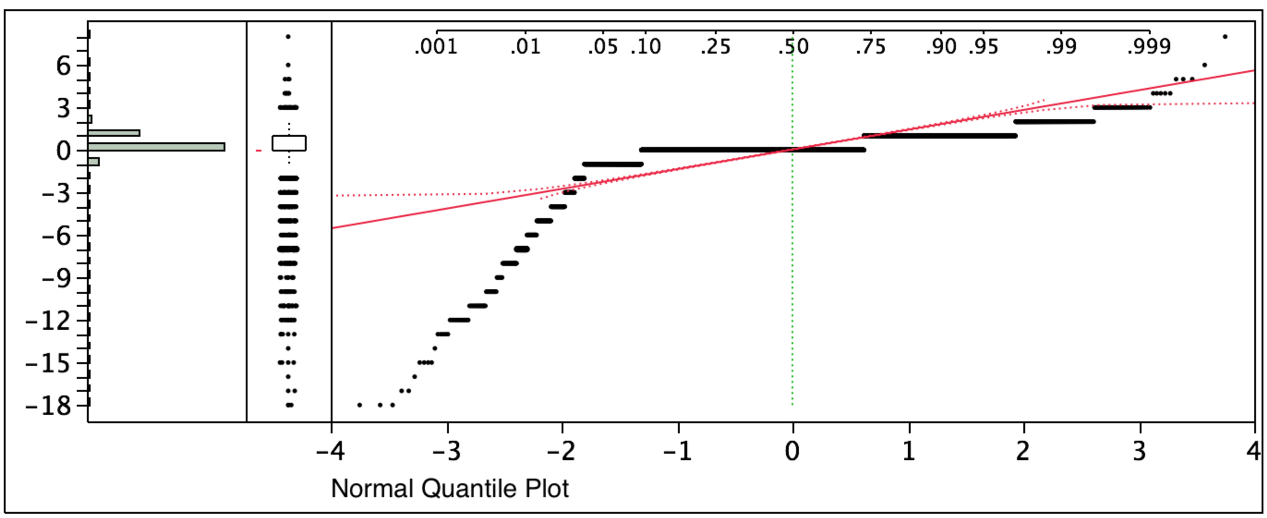

Figure 6.1 Left: Histogram of (Apache-/K) SNR differences for 11261 spots. Center: Outlier box plot. Right: Normal Quantile plot with the red line showing the distribution expected if Normal (Gaussian), with the red dotted lines showing the Lilliefors confidence bounds for the Normal distribution. We term the outliers with an SNR difference more negative than -3dB as the "Extended Tail"

The scatter plot in Figure 6.2 for the entire data set shows a least squares best fit slope in red of 0.961, very close to the 1:1 line in blue. This linear fit captures 96.5% of the variance. But what is also clear is that the spots forming the Extended Tail lie between an Apache SNR of -30 to -10dB.

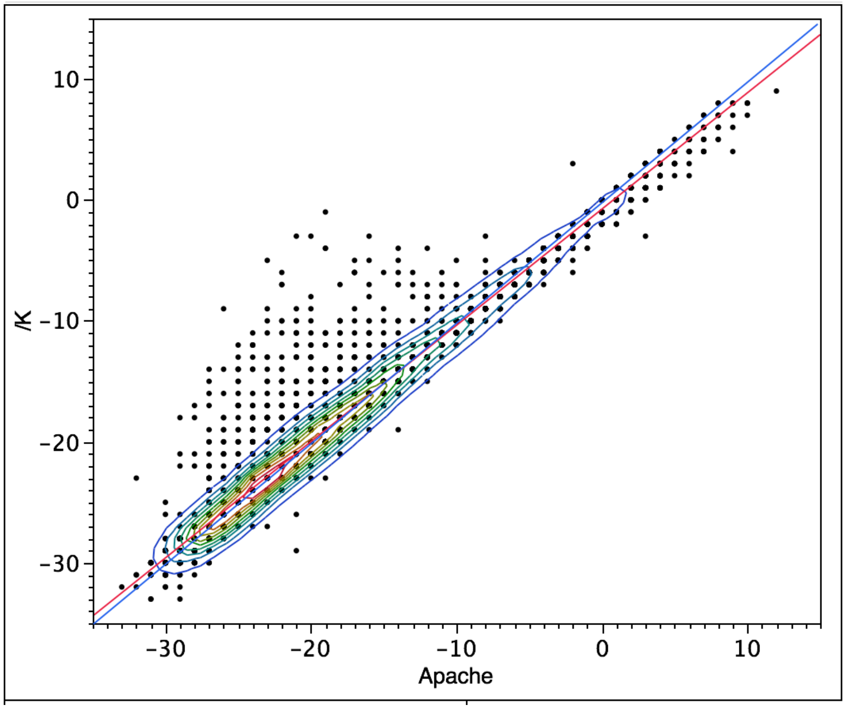

Figure 6.2 Scatterplot of SNR from the /K against that from the Apache. This is a different view of the Extended Tail. But note the excellent agreement for the vast majority of spots with the closely spaced non-parametric density contours around the best fit line (red) and the 1:1 line (blue).

The fact that the extended tail is not symmetric results in a negative bias for the All Spots plot of average (Apache-/K) SNR difference in 2dB bands shown in Figure 6.3 between -30 and -10dB Apache SNR. When the spots from the Extended Tail are removed, the Normal Core curve shows the average difference to be less than 0.5dB between -32dB and -6dB, above which the Apache has an increasing advantage. The 1dB difference for the single point at -33dB Apache SNR is not shown.

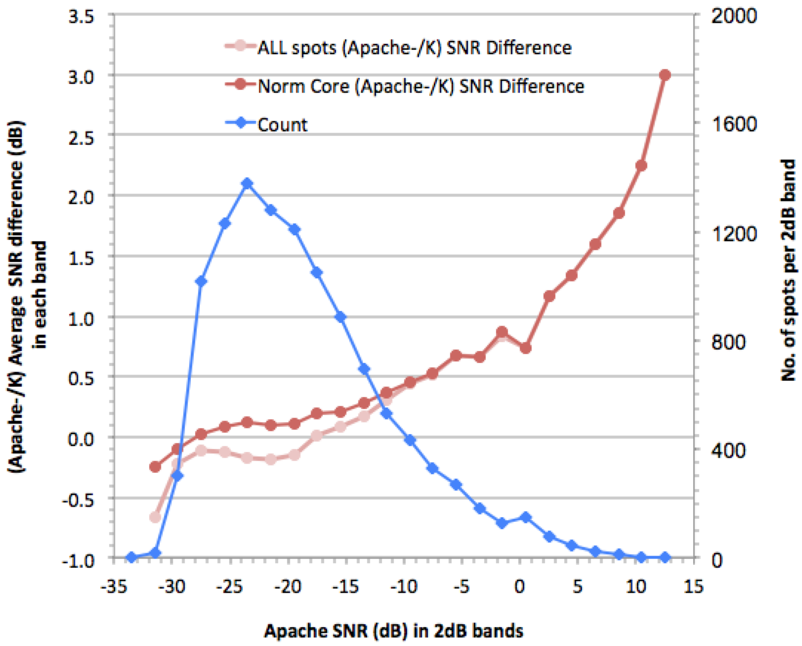

Figure 6.3 (Apache-/K) SNR difference averaged in 3dB bands plotted against Apache SNR together with the number of spots in each 3dB SNR band. In salmon, the entire data set, in red, the Extended Tail outliers have been removed.

6.1 Possible causes of the large SNR differences in the Extended Tail

1. **Spurs:** Hypothesis: the spots comparison for those in the Extended Tail is for a (120Hz or 60Hz) spur on the Apache (which would be lower SNR) and the correct frequency on the /K. **Dismissed**. There are no spurs in this entire data set; the maximum frequency differences are +4 and -1Hz.

2. **AGC:** A very strong signal in a 2-minute interval could reduce the gain of the Apache thereby reducing the SNR of weaker stations. **Dismissed**.
   A) The Apache AGC threshold was set to be higher than any likely WSPR signal in the band.
   B) There were no exceptionally strong WSPR signals, the highest SNR was +8dB.
   C) Examination of the spots showed that those with high SNR difference were not within 2-minute intervals with particularly strong signals.

3. **Very close proximity to a strong signal**: One of our previous informal studies had shown that close proximity (<6Hz) to a strong signal (in our test this was >0dB SNR) reduced the probability that a weaker spot would be successfully decoded. Reducing the SNR is a less drastic effect than not decoding at all. As well as the effect itself there would need to be a receiver-dependent difference in the effect.
   Partially proven – we show that proximity is a cause for the (Apache-/K) SNR differences in the Extended Tail, but that the adjacent signal need not be strong.
   Figure 6.4 shows a scatterplot of the SNR difference (Apache-/K) against the frequency difference between the spot with the SNR difference plotted and the closest spot, whichever side it happens to be. Care has been taken to only look at the frequency differences within each 2-minute interval. We see the large majority of spots that form the Normal Core spread across the frequency difference axis. This would be expected if there was no relationship between SNR difference and proximity to an adjacent signal. Intuitively, the tapers towards the extreme ends of delta frequency arise from those 2-minute intervals with very few spots where one would expect the separation to be larger than when the band was crowded.

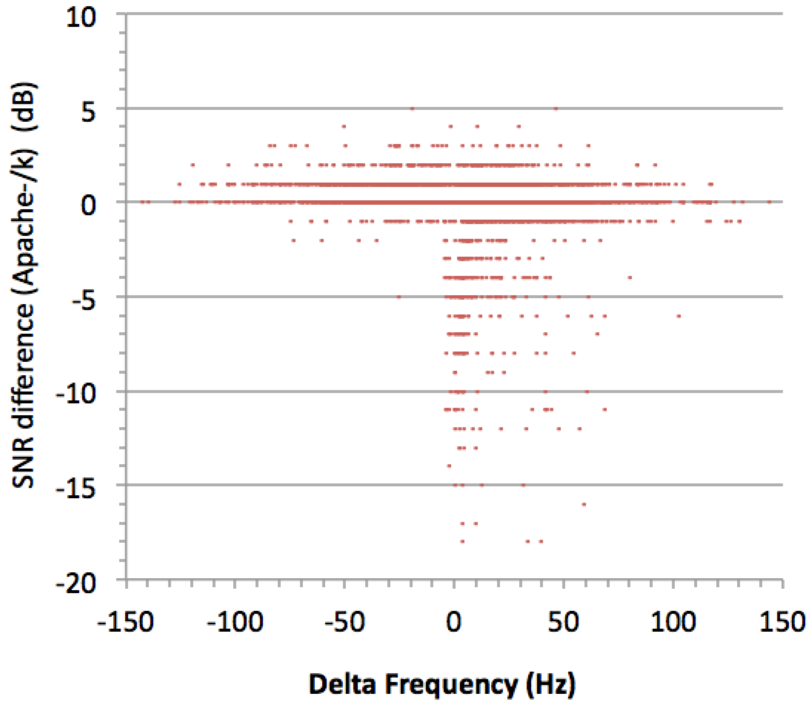

Figure 6.4 Scatterplot of (Apache-/K) SNR difference against the difference in frequency to the closest adjacent signal.

What this plot shows is a very definite “proximity-in-frequency” effect for those strongly negative (Apache-/K) SNR difference outliers that we have called the Extended Tail. There is an asymmetric “funnel-like” pattern where there is a sudden onset at about -3Hz for all values of SNR difference, a cluster of spots at each SNR difference near 0Hz, an extended tail with fewer spots to over +50Hz, and for greater SNR differences the width of the range of frequency difference reduces within the main part of the funnel.

Having confirmed a clear proximity-in-frequency effect the missing factor in Figure 6.4 is the absolute SNR of the adjacent signals or, alternatively, the SNR difference between the spot under consideration and the close-in-frequency adjacent spot. For this we need a 3D plot. Figure 6.5 shows (Apache-/K) SNR difference and delta frequency (as in Figure 6.4) together with the SNR difference to the closest spot (which can be either side). Importantly, in this section the SNR difference to closest spot is from the /K. It's easier to interpret Figure 6.5 while spinning the cube orientation, but even in this view it is clear there is a central column that tapers with increasing SNR difference. This column is centered pretty close to zero SNR difference to the closest spot.

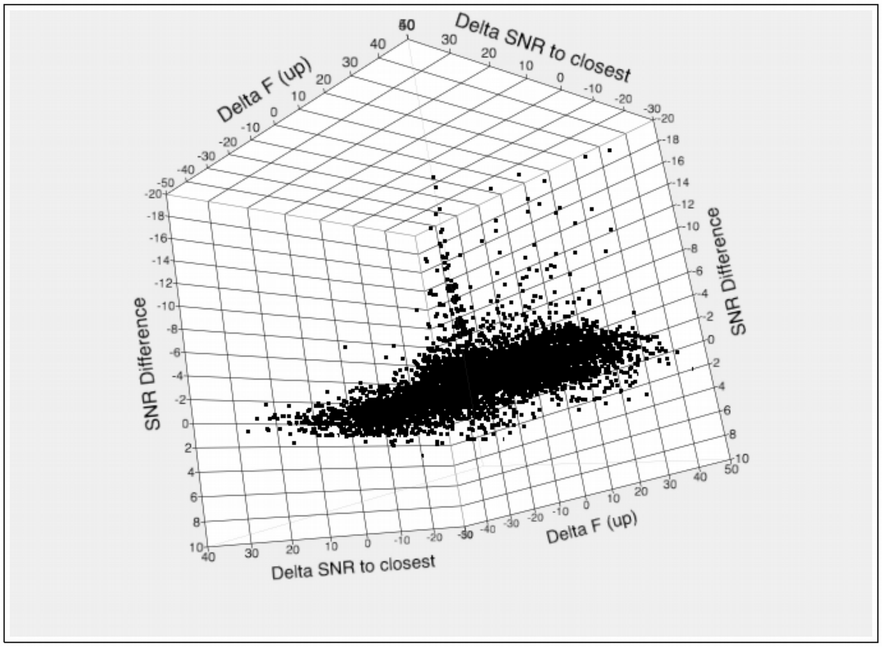

Figure 6.5 3D representation of (Apache-/K) SNR difference, difference in frequency to the closest adjacent signal and difference in SNR for the /K between the spot under consideration and the close-in adjacent spot. The column with high (Apache-/K) SNR difference is clear.

We can see the central column in more detail in the plots of Figure 6.6a and 6.6b that show average (Apache-/K) SNR difference (here, only where each difference is more negative than -2dB, so just within the Extended Tail) in bins of (a) frequency difference in 3Hz bands and (b) SNR difference to nearest signal in 3dB bands.
While the SNR difference changes are not dramatic, the number of spots in each band do show very clear peaks. In frequency, the width is just over 6Hz – probably not a coincidence that this is the bandwidth of a WSPR signal - and the two peak values are at +1.5 and +4.5Hz. In delta SNR to the closest signal, the single peak is at +1.5dB.

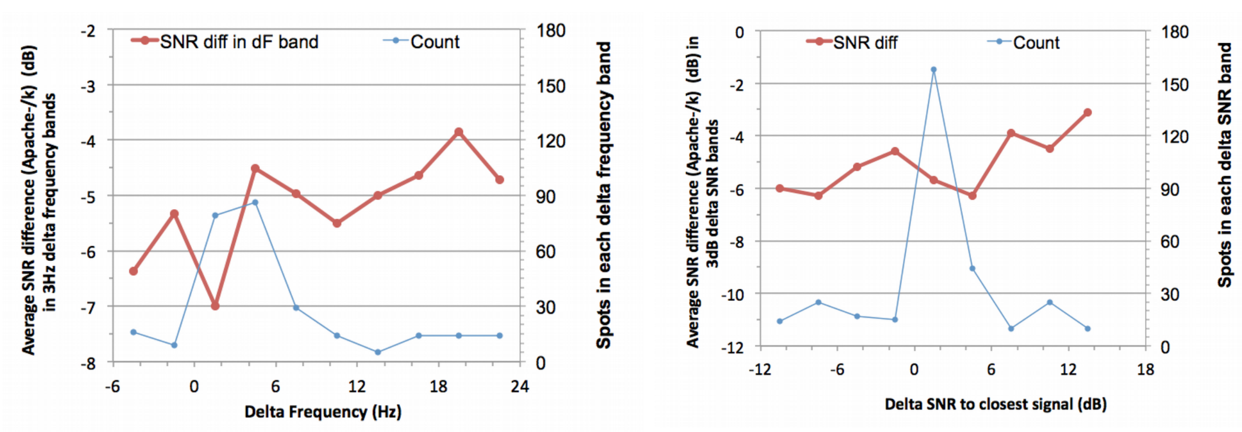

Figure 6.6 Average (Apache-/K) SNR difference in 3dB bands against (Left) frequency difference and (Right) SNR difference for the /K to the nearest spot.

4. **The problem is within the /K route. Proven**. Having shown above that proximity in frequency is a factor, we can look at the spots database and look for examples of large SNR difference (Apache-/K) for

closely proximate signals and then look at the SNR values for the Apache and the /K for the previous and the next transmission by the affected sender where those do not also occur in close frequency proximity to another sender. The Table below, taken from spots on 24 October 2018, gives sufficient examples to show beyond doubt that the problem is within the /K route.

Time	Callsign	Apache SNR	/K SNR	(Apache-/K) SNR diff	Delta F (Hz)	SNR at Delta F	/K SNR minus SNR at Delta F
01:08:00	W6EXT	-19	-19	0	4	-19	
01:18:00		-20	-9	-11	3	-7	-2
01:30:00		-22	-22	0	27	-8	
05:54:00	K3EA	-25	-22	-3	3	-22	
06:04:00		-27	-16	-11	2	-14	-2
06:16:00		-27	-27	0	31	-21	
09:16:00	W2RCL	-26	-26	0	18	-25	
09:18:00		-26	-14	-12	4	-11	-3
09:22:00		-29	-30	1	14	-18	
11:50:00	WA4DT	-16	-16	0	2	-17	
11:58:00		-19	-4	-15	2	-3	-1
12:02:00		-21	-22	1	23	-23	

From this table we can see the following:
a.) Apache SNR either shows small changes or shows a clear trend over the three adjacent time periods for each station.
b.) The /K shows a jump in the middle period – the period where there is an (Apache-/K) SNR difference that falls within what we have called the Extended Tail. It is clear that it is the /K that has the problem.

5. **For completeness, the problem is not with the WSPR noise measurement.** As always with SNR we have to be sure the problem is not the noise measurement. The WSPR noise measurement is used to calculate the SNR of all spots within the 2-minute period. The fact that the /K SNR is affected only when two stations are close in frequency means that it is an "S" problem rather than an "N" problem.
6. **In the /K there is "Top-up" of the S of the weaker of two closely adjacent signals by the stronger one. Proven**. In the Table above, for the middle spot pair for each station we see that the /K SNR is "Topped-up" towards the SNR of the closely adjacent station. For example, for W6EXT at 01:18 the /K SNR should be about -20dB (the Apache SNR), but it is "Topped-up" to -9dB due to being 3Hz away from the adjacent station with an SNR of -7dB. W6EXT's reported SNR has been "Topped-up" to 2dB less than the adjacent station. From the final column, in these examples, the "top-up" is to within 1-3dB of the stronger adjacent station, and from Figure 6.6b we see that the peak is at 1.5dB.
As the S measurement is essentially that of amplitude, and not affected by the coding of the two signals, it is conceptually quite possible for there to be leakage in the measurement of S from the stronger signal to the closely adjacent weaker one where the frequency difference is of the order of the WSPR bandwidth or less.

7. Are spurs still present following the upgrade to v1.242 on 23 October 2018?

The fix by John Seamons of KiwiSDR was made available on 23 October 2018 and the immediate test at N6GN with a local WSPR source showed no spurs or images down to almost 60dB down from the carrier, as in the Table below where the spots at 1838 were before the upgrade and the spot at 1848 after.

Timestamp	Call	MHz	SNR
2018-10-23 18:48	N6GN	7.040100	+29
2018-10-23 18:38	N6GN	7.040053	-25
2018-10-23 18:38	N6GN	7.040170	-25
2018-10-23 18:38	N6GN	7.040124	-19
2018-10-23 18:38	N6GN	7.040100	+28
2018-10-23 18:38	N6GN	7.040076	-21
2018-10-23 18:38	N6GN	7.040006	-33

We have confirmed this early result at N6GN by repeating the analysis for W6LVP on 40m and KJ6MKI on 630m at KPH, both with signals at times well over the +13dB SNR threshold for observing spurs at +/-23Hz, and with clear spurs shown in Figures 4.3 and 4.5 respectively. After the fix the equivalent plots are shown in Figures 7.1 and 7.2. Both are entirely clear of the spurs that were present at about 44dB down from the main.

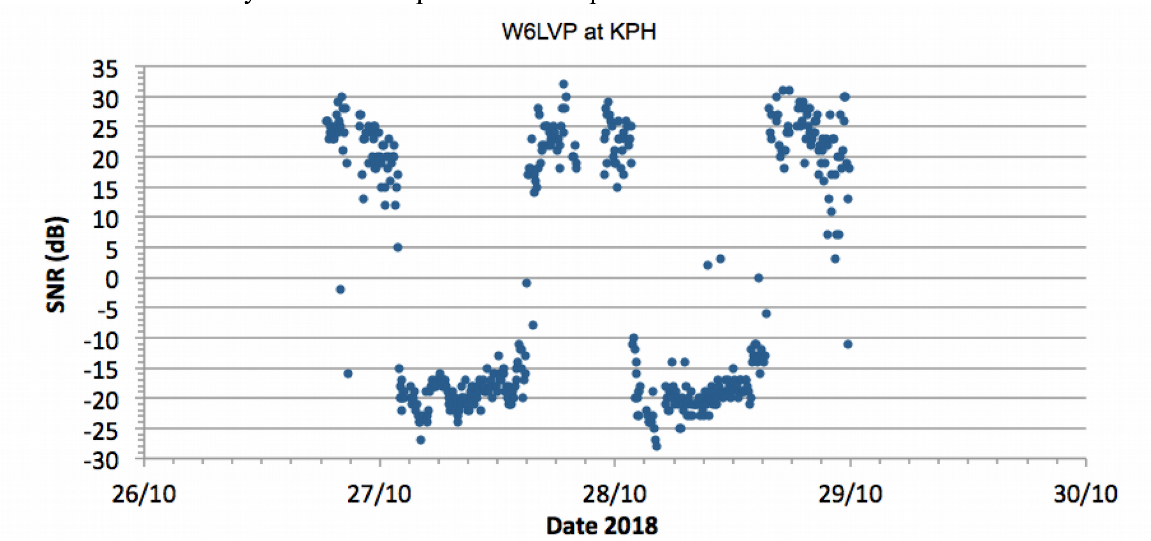

Figure 7.1 SNR of W6LVP on 40m at KPH. Compare with Figure 4.3; here is no evidence of spurs in this plot.

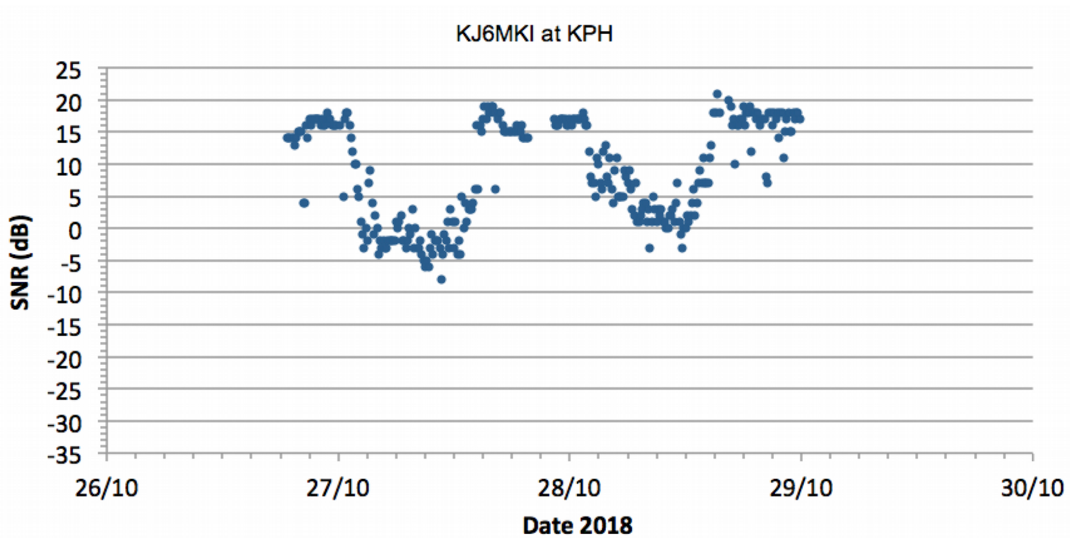

Figure 7.2 SNR of KJ6MKI on 630m at KPH. Compare with Figure 4.35; here is no evidence of spurs in this plot.

7.1 Asymmetric spurs within 25Hz of the main WSPR signal

While not detracting from our conclusion above that the symmetrical spurs at 23Hz, occurring when the main signal SNR was above +13dB, have been removed following the changes in KiwiSDR v1.242 and above we have observed the presence of asymmetric spurs, i.e. those that only occur on one side of the main at any one time. At present, we think there are two types:

a) Very close-in to the main, typically within 5Hz. While the rate of occurrence may vary, we have found such spurs in the spots reported by one station using a receiver other than the KiwiSDR. However, we do consider it possible that the "SNR top-up" within the KiwiSDR, described in section 6, may contribute to a higher incidence of these very close-in spurs. Our observations are documented in section 7.1.1.

b) Close-in spurs, typically between 10 and 23Hz of the main. Our working hypothesis is that these spurs are not generated within the KiwiSDR but are due to aircraft scatter (ACS). Our observations are documented in section 7.1.2. ACS is an engrossing topic in its own right; we have only done enough analysis to show to our satisfaction that ACS is much more likely to be the cause of these spurs than a problem within the KiwiSDR.

7.1.1 Spurs within 5Hz

K6SRO drew our attention to the presence of very close-in spurs on some of his 80m WSPR signals received by the KiwiSDR at KPH, for example:

```
Date & Time      Frequency (MHz)     SNR (dB
06/11/18 21:28     3.570096           -10
06/11/18 21:28     3.570100            30
```

Our understanding from the man page for the wsprd decoder (within WSJT-X or as a standalone program) is that it treats decodes from the same callsign and within 1Hz as the same signal. The duplicate checking in wsprd can be turned off with a -v flag[3]. However, as shown by the example above, same callsign, same time-slot decodes can happen at larger frequency offsets than 1Hz.

To give us some idea of how frequently such decodes happen we have examined 932 spots from K6SRO received at KPH on 80m, leading to the following observations / questions:

1. Frequency of occurrence: 22 out of 932 (2.4%) spots were very close-in spurs at 5Hz or less from the main signal. The incidence was: 5Hz: 4 (0.43%), 4Hz: 14 (1.5%); 3Hz 4 (0.43%). There were also two close-in spurs at 7Hz and 21Hz (0.2%).
2. Are they symmetric around the main signal? No. All are one-sided.
3. Is there a dominant side? Yes. In 22 out of the 24 cases the spur was to the low frequency side of the main.
4. SNR difference between the main and spurs: Average SNR difference between main and spur was 35dB, ranging from 4 to 40dB, lower quartile 26dB, median 35dB, upper quartile 37dB.
5. Is the presence of a spur related to the SNR of the main signal? Probably. The lower quartile SNR of the main when a close-in spur was present was +24dB, but there were spurs when the main SNR was -7, +12 and +16dB. That is, they are more prevalent at KPH at a main signal SNR at or above 24dB with a 5.6% (17 out of 304) chance of occurrence, but with 0.5% (3 out of 628) when the SNR was below 20dB.
6. Do other receiver stations show these very close-in spurs? There is very limited evidence, we have seen that the prevalence as KPH is higher at SNR above +24dB and there are very few sender-receiver combinations with the two attributes of very high SNR and several hundred spots to provide a rigorous comparison for this few-in-a hundred effect. We have however found 3 out of 566 very close-in spurs for K6SRO at W7OWO (RSP receiver), but 0 out of 173 at KP4MD (Flex 1500). For the ground-wave pair

3 See man page at https://www.mankier.com/1/wsprd

DL4XJ and DF5FH there were 0 out of 225 despite a median SNR of +15dB and a maximum SNR of +24dB (Red Pitaya).

Based on the information available to us these very close-in spurs do not seem to be deterministic. This is in contrast with the spurs at +/-23Hz that we have documented that were always present above a main SNR of +13dB. We have not been able to determine what other factors may contribute to the occurrence of the very close-in spurs.

7.1.2 Spurs between 5 and 25Hz

Rob, AI6VN, observed what was initially thought to be a sporadic, SNR-dependent, frequency offset for WSPR signals from K6PZB, located some 38km from KPH. A time series plot showing SNR and the frequency difference in Hz from K6PZB's GPS-locked frequency of 14.097090MHz is shown in Figure 7.3.

We note the following:

1. The large majority (78%) of spots are reported at 14.097090MHz (with the 1Hz resolution data).
2. There are two periods when the KiwiSDR appears to have lost its GPS-aiding and we see frequency errors from the variations in frequency of the KiwiSDR crystal oscillator.
3. There are some sporadic very close-in spurs within +/-5Hz.
4. There are spots either side to -20 / +24Hz.
5. While 96% (1072 out of 1116) of the spots have an SNR of -14dB or above, the SNR of 44 spots, apparently sporadic, was significantly lower.

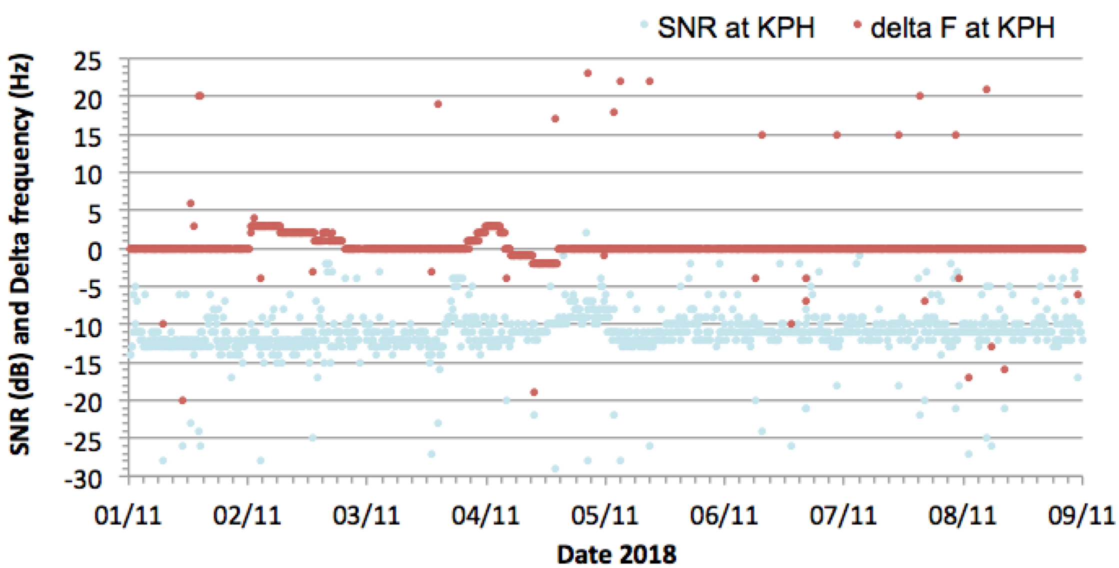

Figure 7.3 Time series of K6PZB SNR at KPH (blue) with the offset from his GPS locked frequency of 14.097090MHz showing periods when the KiwiSDR lost its GPS-aiding but also showing the sporadic spurs discussed in this section.

The coincidence of frequency offset and low SNR, first made visually from text records, is confirmed in the scatterplot of Figure 7.4, from which we note:

1. Ideally, the SNR spots in green would form a vertical pillar at 0Hz, but the loss of GPS-aiding results in a broader distribution out to -2/+3Hz.
2. The SNR outliers below -16dB do lie beyond -3Hz and +6Hz, confirming the initial observation that there appears to be an SNR dependent frequency offset. However, closer inspection showed that all of these points were duplicates – there was a decode within -2/+3Hz and at an SNR of -14dB or greater

associated with each outlier, hence these are spurs.

3. Figure 7.4 also shows the reported drift rate in Hz/min for each spot. For the spurs beyond -3Hz and +6Hz, with two exceptions the spots at +6Hz and -23dB SNR and at -20Hz and -26dB SNR[4], they have a drift rate of -1 to -4Hz/min. There are no spots with positive drift rates.
4. For the spurs with with a positive frequency offset and a non-zero drift rate there is a clear trend towards lower SNR with greater frequency offset. This trend is not obvious for negative frequency offsets.
5. For the spurs with with a positive frequency offset and a non-zero drift rate there is a clear gap, the first spur in this cluster is at +15Hz.

From the Table below we observe that the spur for K6PZB was present even though he was not the strongest signal within this interval. However, he was the only nearby ground wave signal.

Figure 7.4 Scatterplot of K6PZB SNR at KPH against frequency offset (green), together with drift rate (red). Many spots overlie each other in the central core between -2 and +3Hz (1072 spots in all).

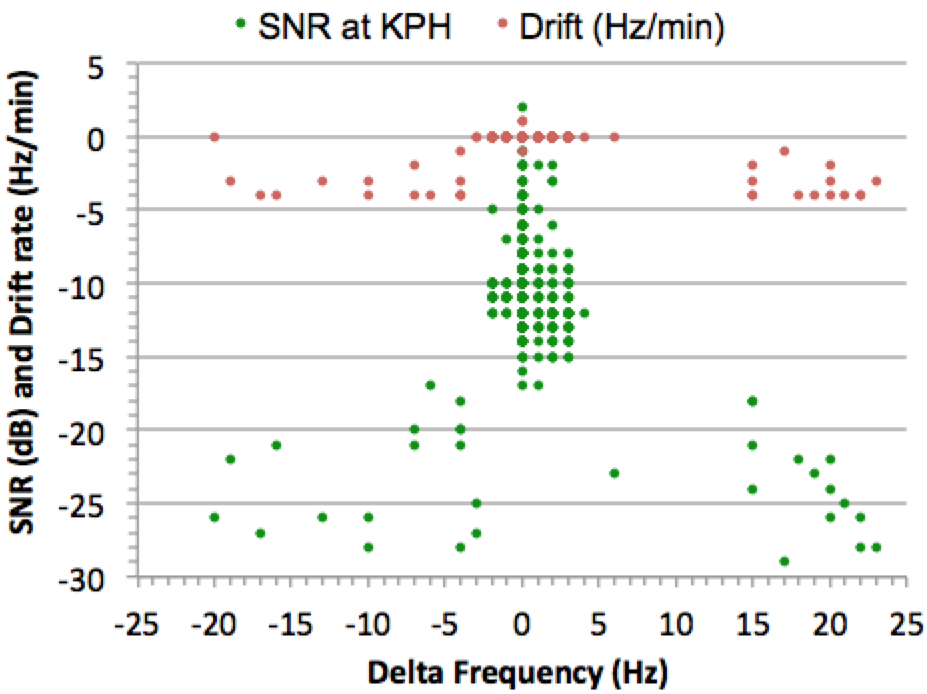

2018-11-03:14:10,	KV4PD,	14.097101,	-8,	0,	3437
2018-11-03:14:10,	K6PZB,	14.097109,	-23,	-4,	38
2018-11-03:14:10,	AE5HO,	14.097127,	-13,	0,	2403

Based on all these observations our hypothesis is that these spurs beyond -3 and +6Hz were due to aircraft scatter (ACS) and not due to an issue with the KiwiSDR. To test this hypothesis we have examined the following:

1. **Are there flight-paths near K6PZB and KPH?:** Air traffic patterns were studied using the flightrader24.com website. There is a well-travelled flight path to the east of the K6PZB-KPH pair, generally with flight northbound to Alaska and a great circle route from California and Mexico to the Far East. Southbound flights or flights from the Far East to California tend to follow a track either over KPH or to the west. The local route from Santa Rosa to San Francisco lies between K6PZB and KPH. An example track of a China-bound Airbus A380 aircraft from Los Angeles (LAX) is shown in Figure 7.5.

4 It is possible that the spot at -20Hz and -26dB SNR with a drift rate of 0Hz/min has had its true drift rate rounded down to zero.

Conclusion: There is sufficient air traffic to make the hypothesis tenable.

2. **Is the range of frequency offsets observed possible given the ground speed of commercial aircraft?** For the case where the aircraft is very much greater distance than the separation of the transmitter and receiver, the Doppler shift in Hz is simply twice the speed of the aircraft divided by the speed of light multiplied by the frequency, which for a ground speed of 530 knots and 14MHz is 25Hz. However, where the distance to the aircraft is not much greater than the separation, as in Figure 7.5, the frequency shift is given by the bi-static Doppler equation:
where lambda is the wavelength, *d/dt* is the time derivative, and the two ranges are between the transmitter and aircraft and receiver and aircraft. The Table in Figure 7.5 shows the Doppler shift at the stated times along the track of this aircraft, at 500 knots, starting at time zero over Vallejo in the south east of the map. The Doppler shift approaches its asymptotic value with the aircraft at a distance well over ten times the separation of the receiver and transmitter.
Conclusion: The range of frequency offsets observed are consistent with aircraft speeds of up to 530knots, with the variations arising from the shape of the curve of Doppler shift with range as in the example Table in Figure 7.5

3. **Is the rate of change of Doppler shift sufficiently low, at a short enough range, so that there is adequate SNR?** The upper limit for the rate of change of frequency for successful WSPR decode is often given as +/-4Hz/minute. Inspection of the Table in Figure 7.5 suggests that while near closest approach the rate of change of Doppler shift exceeds -4Hz/min, at a total round trip range in excess of about 125km the rate of change drops to -4Hz/min and is down to -1Hz/min at over 200km total round trip range.
Note that there is no need for a course change or speed change by the aircraft to produce a non-zero drift rate, a change in bearing between the aircraft and the baseline between transmitter and receiver is sufficient. **Conclusion: While acknowledging that the rate of change of Doppler shift is dependent on aircraft speed, the cluster of spots with positive frequency offsets in Figure 7.4 does suggest we do see spots with sufficient SNR at drift rates (range) of less than -4Hz/min.**

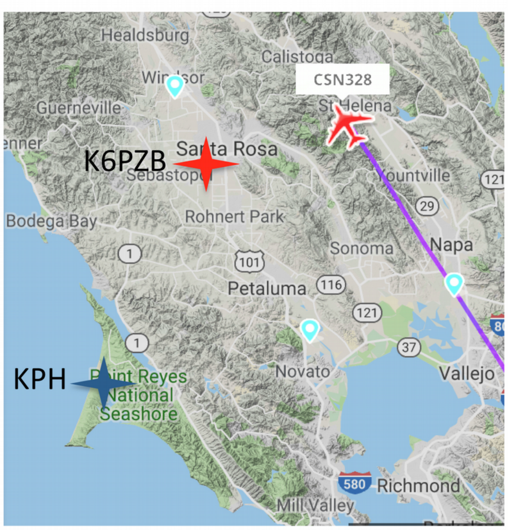

Time min	Total range km	d/dt(tot range) m/s	Doppler shift Hz
0	117	246	11.6
1	102	137	6.5
2	94	-4	-0.2
3	94	-171	-8.1
4	104	-351	-16.5
6	146	-446	-20.9
8	200	-477	-22.4
10	257	-491	-23.1
14	376	-503	-23.6
20	558	-509	-23.9

Figure 7.5 (Left) Map and track from flightradar24.com showing the track of a China-bound A380 jet. Southbound flights to San Francisco and Los Angeles follow a track just offshore of Point Reyes. The approximate positions of K6PZB and KPH are shown. (Right) Table of range, range derivative and Doppler shift for 20 minutes on from the aircraft being over Vallejo in the south-east of the map, 20 minutes takes the aircraft well north of the area shown on this map.

4. **K6PZB transmits at a typical interval of 10 minutes. We have not observed consecutive spots with**

frequency offsets. Are these two facts compatible with ACS? Figure 7.6 shows the frequency shift, rate of change of frequency and range for an aircraft traveling at 500 knots along the track of Figure 7.5. We have shaded that part of the track where the rate of change of frequency exceeds 4Hz/minute (the dotted horizontal line). The horizontal lines with arrows each span 10 minutes. The extreme left and right arrow heads are at a total round-trip range of over 450km where it would not be unreasonable that the SNR would be too low to decode. The center arrowed line suggests that two consecutive decodes should be possible if total round trip ranges of 160km and 190km would produce adequate SNR. Looking at the frequency of occurrence of drift rates we have 16 at 4Hz/min, 7 at 3Hz/min, 3 at 2Hz/min and 3 at 1Hz/min. So a decode is most likely just at the boundary of the shaded box at 4Hz/min – that is, at the minimum possible range at which the drift rate is low enough to decode. This pushes the next transmission time to a greater range, making a successful decode less likely due to lower SNR. **Conclusion: For aircraft traveling at 500 knots we would expect, at most, two consecutive spots to be decoded, and the lack of consecutive spots with frequency offsets in this set of data is not at variance with the ACS hypothesis.**

5. **We observed negative and positive frequency offsets, but both have negative drift rates. Is this observation consistent with ACS?** See Figure 7.6, reading from the left, the aircraft is approaching, with its maximum positive Doppler shift, which goes through zero and then becomes maximum negative. But the slope is always negative; the magnitude of the slope changes but not its sign. This is true whatever the direction of approach; what will change with direction of approach is the shape of the slope with time curve. **Conclusion: This observation is entirely consistent with the ACS hypothesis, indeed, it could be considered the clincher.**

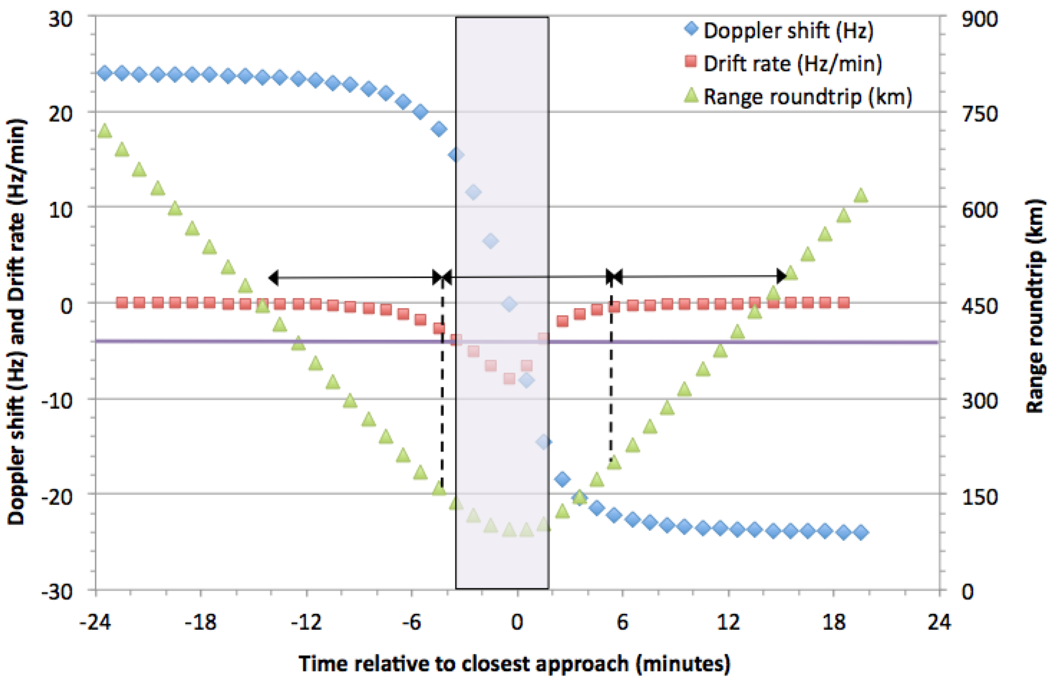

Figure 7.6 Time series of Doppler shift, drift rate and round trip range for an aircraft traveling at 500 knots along the track of Figure 7.5. The shaded area has a drift rate of over -4Hz/min and so no decodes are possible. The horizontal lines with arrows span 10 minutes, and illustrate that two consecutive spots should be possible, even at 10 minute transmission interval, but that for 10 minutes earlier or later the aircraft would be at too great a distance to provide an adequate SNR.

While these points have convinced us that these spurs are from ACS, a great deal of work could be done to be absolutely sure, including use of CW or 2 or 4 minute interval WSPR, simultaneous monitoring of flight radar websites noting aircraft speed and positions and possibly two-frequency observations (e.g. 20m and 15m).

Other mechanisms exist that induce Doppler shift, e.g. rising and falling of layers of the ionosphere, but so many of the characteristics we have observed fit with these simple simulations of ACS that we have not delved into other mechanisms.

Summary

From our measurements and interpretation it appears that the KiwiSDR did have some additional room for improvement. We found evidence that in comparison to other SDRs the delivered SNR was not yet as good as it might be. Spurious sidebands were generated within the KiwiSDR which resulted in false decodes on WSPR. Extremely quick response by John Seamons and Christoph to the original data presented in Part 1 resulted in elimination of the spurious spots as well as a large reduction in the SNR differences which had been observed between the Apache SDR and KiwiSDR platforms.

Further measurements have revealed a remaining difference wherein spots of two closely spaced WSPR stations on the KiwiSDR can result in the SNR of the weaker of the two being spuriously assigned an SNR higher than it should be, typically to within 1.5dB of the stronger of the two. . These findings are presented in Part 2 of this article.

Appendix 1
WV0Q Detail

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.

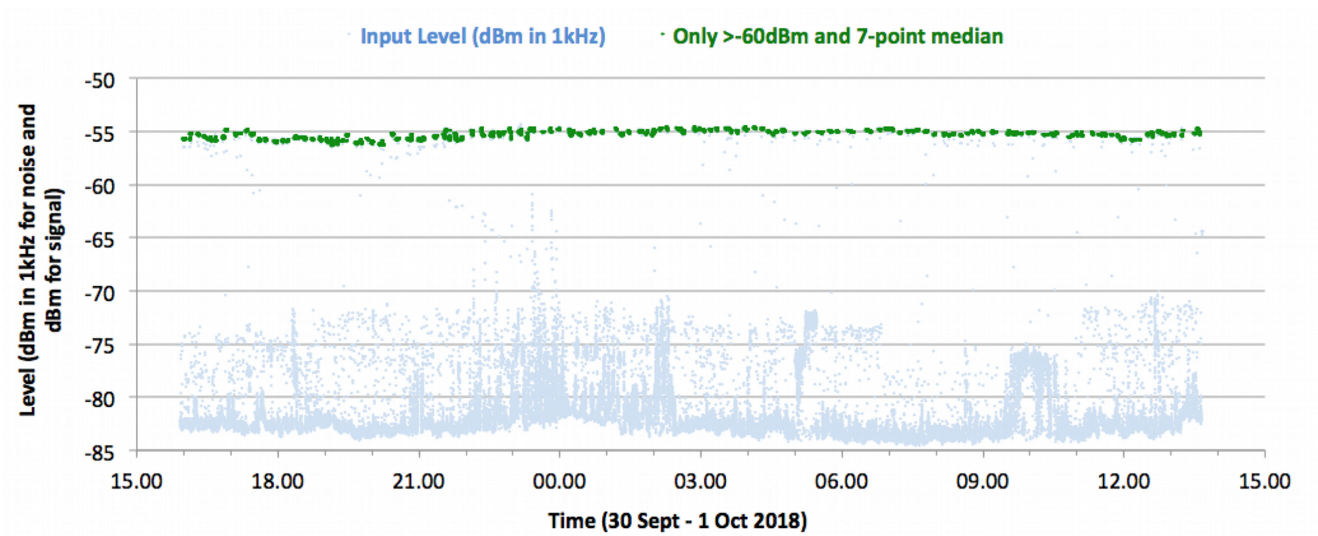

Figure 6.1. Time series of raw input level measurements at 1.25s intervals (light blue) together with the level ofWV0Q's main transmission after a 7-point median filter (green)

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.

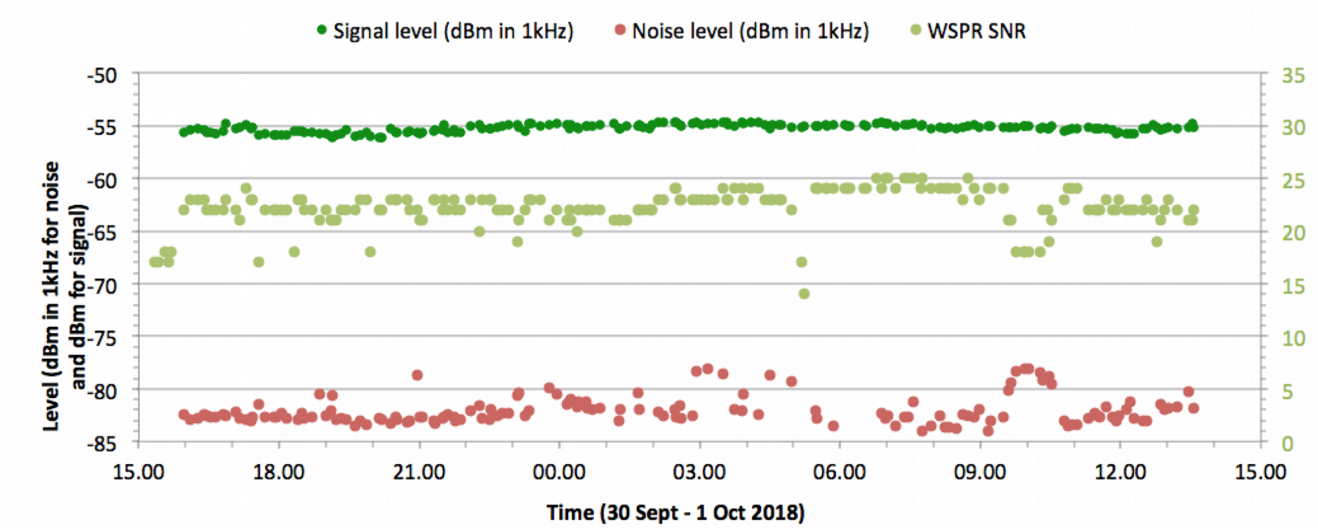

Figure 6.2. Time series of the level of WV0Q's transmission (dark green) together with WSPR-reported SNR in light green and our best estimate of noise and interference in the 8 seconds immediately after the end of each transmission. Gaps in the latter are likely due to interference within the 1kHz noise 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.

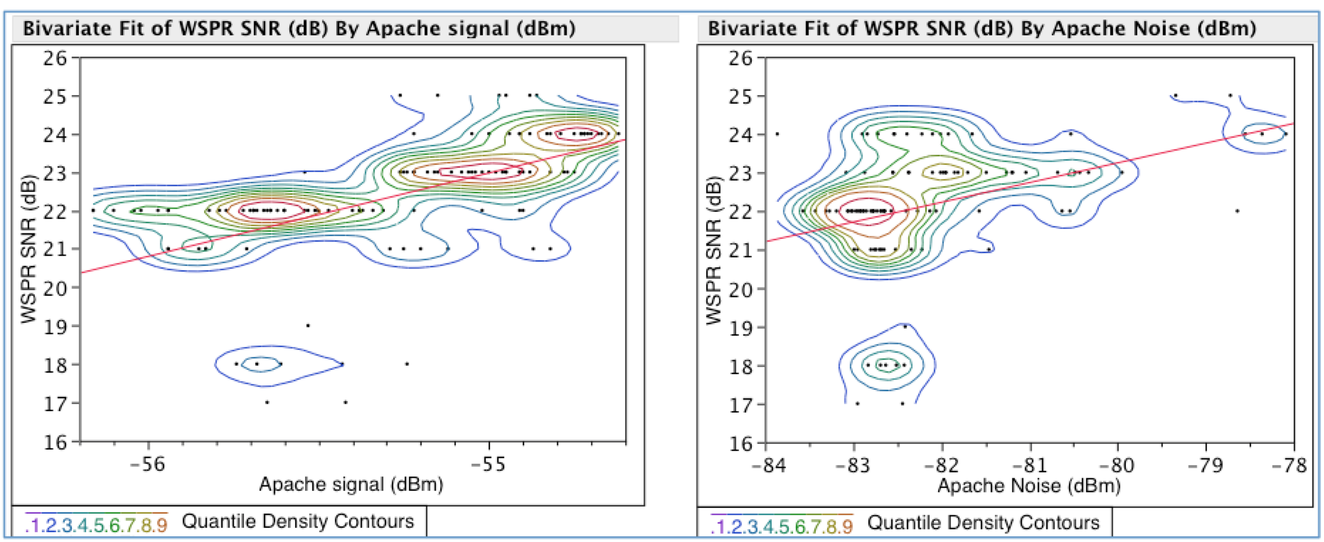

Figure 6.3. Scatter plots of WSPR SNR against Apache signal level (left) and noise in a 1kHz bandwidth 8 seconds after the end of each WSPR transmission.

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.