Jan M. M. Simons, PAØSIM
Prins Mauritsstraat 14, 5923 AZ Venio, The Netherlands; pa0sim@veron.nl

Finding Signals in the Noise Using Two Antennas

Use noise cancelling and noise reduction techniques to extract signals from noise.

Man-made noise, and hence, the noise floor on the HF bands in urban areas has significantly increased over the years[1]. Especially in the evening, local man-made noise (QRM) levels can be very high. When living in such an urban area and having limited space for antennas, noise cancelling can be used to attenuate a single identifiable man-made noise source. The remaining noise floor is the sum of multiple man-made noise sources and natural noise. Noise reduction can be used to mitigate this noise floor. The presented approach uses the signals of two small orthogonal antennas to find the signals in the noise. It incorporates one mouse-click cancelling and measuring propagation polarization behavior over frequency. It is a new approach to the one published by Simons[2].

Two Orthogonal Antennas

Two orthogonal antennas, like two identical small active receiving loops (Figure 1), provide two different but co-located antenna signals. Noise arrives from all directions with varying power and polarization. The resulting signal is random in time. The phase difference and amplitude ratio between the antenna signals are also random in time. Even two man-made noise sources with about equal signal strength can result in random behavior. This makes the noise floor random and renders a noise canceller less effective. Of course a single strong man-made noise source can be present. Then noise cancelling is the first line of attack.

Received sky-wave signals like SSB and CW however will only drift slowly in time in direction of arrival and in polarization. The signals of both antennas will be highly correlated and the phase difference and amplitude ratios will be almost constant over a period of time. When receiving circular polarized signals the phase difference is ±90° and the amplitudes are equal all over time.

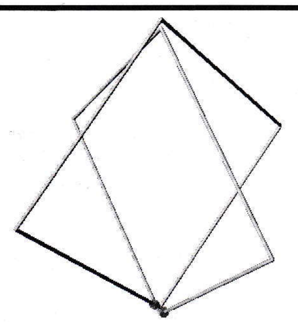

Figure 1 — Two orthogonal small active receiving loop antennas, each 1.25 m by 1.25 m.

Summary of the New Noise Reduction Approach

The processing is performed in the frequency domain at audio frequencies and is based on a three-dimensional space for estimating whether it is a signal or noise. The amplitude ratio and phase difference can be calculated for each frequency in the audio spectrum of the two antenna signals *A* and *B* that arrive from the two orthogonal antennas, and used as a pointer to a location in a two-dimensional (2D) space. The 2D spaces of all frequencies together make a single three-dimensional (3D) space. The amplitude ratio and phase difference, representing the polarization, and frequency are the coordinates of this 3D space.

When just noise is present the locations are uncorrelated and randomly spread over time. The locations of a signal with noise however will be concentrated in that space

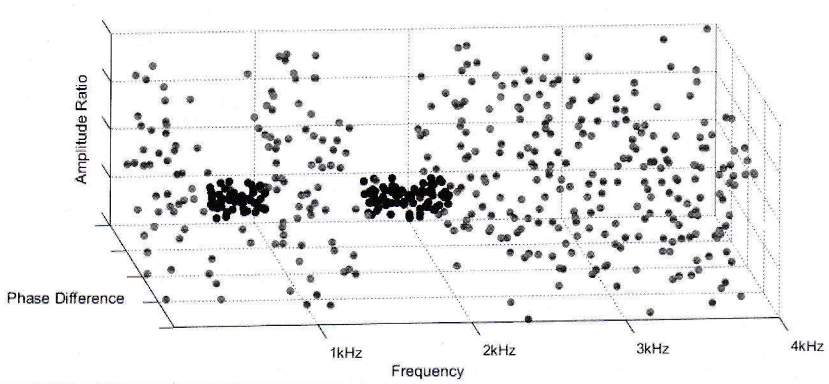

Figure 2 — Possible locations of a SSB signal power with noise over a short period of time in the 3D space.

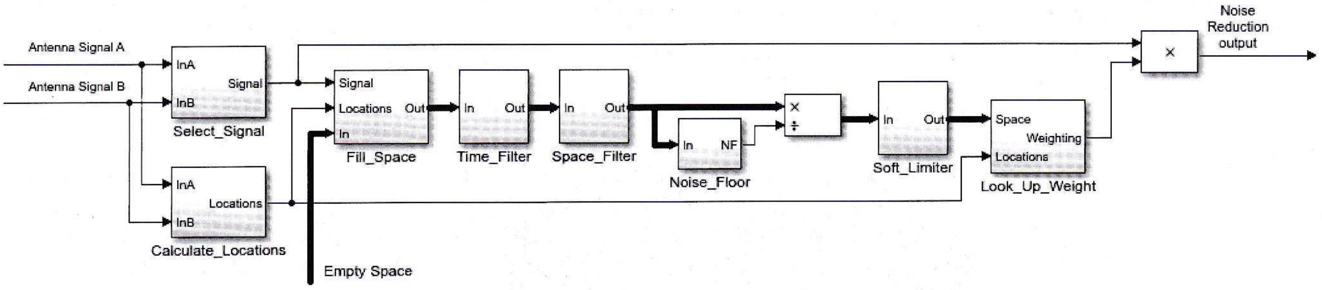

Figure 3 — Noise reduction processing using a 3D space.

over time. They not only will be concentrated in the 2D space of a single frequency, but also over a certain frequency span. Because of modulation the frequency components in a signal cover a bandwidth. In Figure 2 signal components of an SSB signal are present around 600 Hz and 1.7 kHz. They are concentrated on each frequency and over a limited frequency span. The higher power level of the signal components is indicated by a darker dot.

The thick lines in Figure 3 indicate the 3D spaces in the processing. An empty 3D space is filled on the corresponding locations with the power of the frequency components of the selected signal. The noise reduction is based on the assumption that the concentration of power in the 3D space over time represents an estimate of the location of the signal in that space. The concentration of power over time is calculated by performing in sequence a low-pass filtering in time and a 3D low pass filtering in space. The result is a low pass filtered space.

The noise floor is a sample of the low-pass filtered space when only noise is present. The power level in the noise floor is a measure for the expected signal levels. It is needed to normalize the low-pass filtered space. By scaling it to the maximum power in the noise floor the space can be weighted using a *sigmoid* function. The sigmoid function soft limits the space. The soft-limited concentration of power in the 3D space over time now not only represents an estimate of the location of the signal in space, but also whether it is a signal or noise.

The selected antenna signal is then weighted (filtered) by using the actual location as a pointer to the weighting factor in this soft-limited low-pass filtered 3D space.

Processing in Frequency Domain

The audio output signals from two phase coherent receivers are processed with 8 kHz sample frequency. The processing is done in the frequency domain using a 512 sample long FFT/IFFT and 8 times overlap with a 64 ms long Hann window. Each frequency bin is 15.625 Hz. With the Hann window this results in a 22.5 Hz bandwidth (-3 dB). Equivalent noise bandwidth (ENBW) is about 23.5 Hz. In order to have enough measurements in time of the spectra (125 per second) at least 8 times overlap is needed. The Hann window is applied also after the IFFT, because rather extreme filtering is possible[3].

Both antenna signals *A* and *B* are real signals in time. The FFT provides for each antenna 512 frequency bins with complex numbers. The output of the FFT however is conjugate symmetric and only 256 of the 512 bins need to be processed. The bin containing the dc component is skipped.

Spatial Information

Each pair of signal *A* and signal *B* bins can be plotted in a 2D space. The location in that space is set by the phase difference *P* (*x*-axis) and the amplitude ratio *R* (*y*-axis),

$$P = \text{angle}(A) - \text{angle}(B);\ [0° < P < 360°]$$

It is not practical to use the amplitude ratio directly. If the magnitude of antenna *A* is $|A|$ and of antenna *B* is $|B|$ the ratio is set by,

$$R = \frac{|A| - |B|}{|A| + |B|};\ [-1 < R < 1]$$

This allows one of the signals to be zero. Noise will have random locations in that space. In Figure 4A the locations in a single bin over 4 seconds are plotted. Each dot is one of 4×125 measurements. The distribution in space depends on the statistics of the noise in a bin. The *randn* function of Matlab[4] is used to generate normally-distributed pseudorandom numbers.

If a circularly polarized carrier wave signal is added with a 3 dB signal-to-noise ratio (SNR) on the antenna signals, the locations in space (Figure 4B) will become concentrated at the signal location. The relative phase of the circular polarized signal components is -90° and their amplitude ratio is 1. With an SNR of 3 dB in a single bin the concentration in space is already significant.

Concentration in space is the number of measurements per unit of area over time. However not only the location in space, but also the power is available in the measurements. Noise will have random locations and random power (Figure 5A). A carrier wave signal with a 3 dB SNR (Figure 5B) will express itself by a higher concentration and by a higher power. The sum of the noise and the carrier wave signal is concentrated around the location of the carrier wave signal. A higher signal to noise ratio results in a greater concentration of power in space.

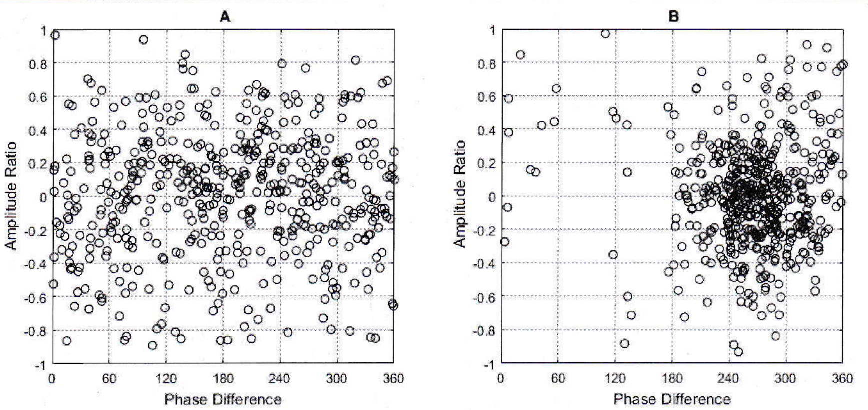

Figure 4 - The concentration in space in a single bin (23 Hz bandwidth) of just noise on both antennas (A), and the same noise with a circular polarized signal with 3 dB SNR (B). Each dot is one of 4x125 measurements over 4 seconds.

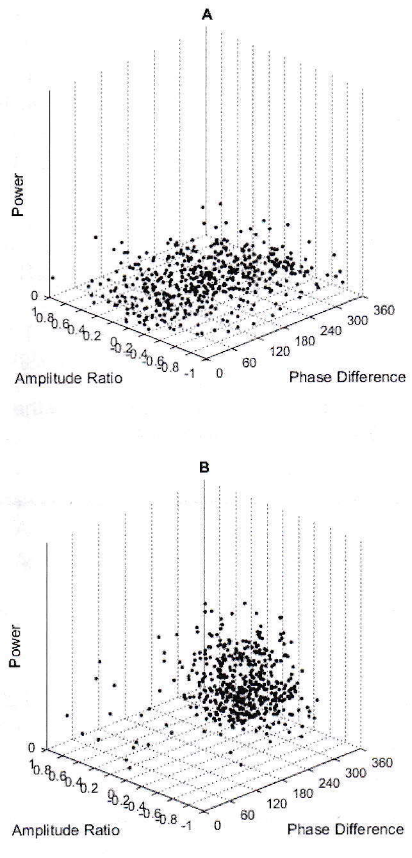

Figure 5 — The concentration of power in a single bin (23 Hz bandwidth) of only noise (A) on both antennas and the same noise with a circularly polarized signal with 3 dB SNR (B). Each dot is one of 4×125 measurements over 4 seconds.

Extracting the Spatial Information

The spatial information we need is in the concentration of power in space over time. It is derived by calculating per unit area the sum of the power over time. Calculating over time is a low-pass filter operation of the power in space over the measurements. Calculating per unit area the sum of the power in space is a 2D low-pass filtering operation. Two filters are needed: low-pass filtering in time and low- pass filtering in space.

The Spaces

First the locations in space of the measurements have to be quantized to reduce the size of the space for processing. The processing is performed for all 256 frequency bins. The space of each frequency bin is called a subspace. The size of the subspace is set to 16 by 16 forming 256 locations. A further reduction is implemented for SSB by combining the measurements of 4 consecutive subspaces into a single subspace. This reduces the number of subspaces from 256 to 64. In this way 64 (16×16 size) subspaces have to be processed. CW needs only about a quarter of the 4 kHz spectrum and a single frequency bin is used in a subspace.

Figure 6 shows the 64 (16×16 size) subspaces. Four subspaces are sliced out of the full 3D space. The white square indicates a possible location of a signals power in space on a single frequency in subspace 45. Instead of processing it as a 3D space, all subspaces are combined in a single 2D space. The single 2D space consists of 64 subspaces in a row. Total size is 16×1024. Thus for SSB the row consists of 64 frequency ranges of 4 bins and for CW of 1 bin.

The 16×1024 space is clearly not a convenient display presentation. The total space is reorganized for displaying as a 2D image consisting of 8×8 subspaces (Figure 7).

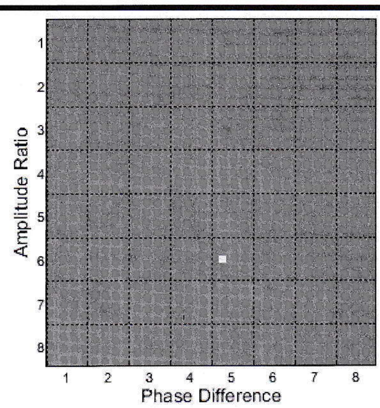

Figure 7 — All 64 subspaces are combined for displaying in a 128×128 image. The white square indicates a possible location of a signal power in subspace 45. Top-left lowest frequency, bottom-right highest frequency.

Low Pass Filtering in Time

The filtering in time of the 16×1024 2D space is implemented as a *RC*-like filter with an attack and a decay time constant. Each location in that space is filtered over time. In Figure 8 the default time constants are 25 ms for attack and 38 ms for decay. Each dot is a measurement in 8 ms steps.

In Figure 9 the low-pass filter is applied, otherwise there would be only 4 measurements in the plot. Noise is spread over the whole subspace (Figure 9A). The signal concentrates the measurements in subspace (Figure 9B). Because 4 bins are combined in the subspace for SSB, the noise from the other 3 bins is also present. The locations in these bins are also affected by the carrier wave signal as a result of

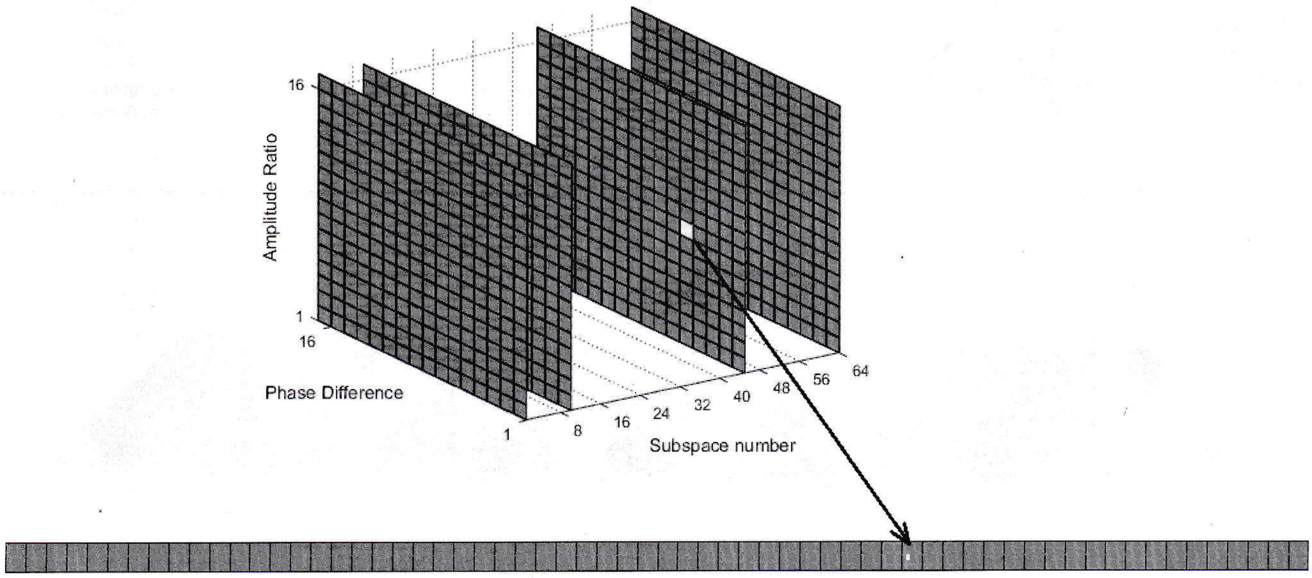

Figure 6 — The 64 (16×16 size) subspaces are combined in a single 16×1024 2D space for processing.

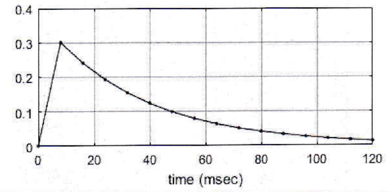

Figure 8 — The *RC*-like low-pass time filter default response over 120 ms on a single measurements with amplitude of 1.

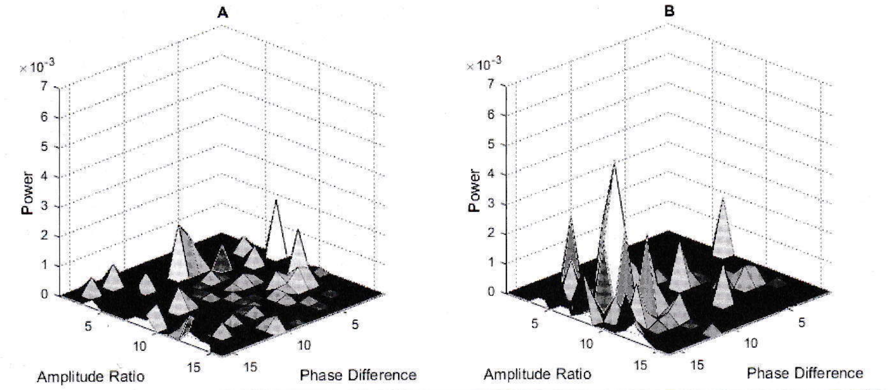

Figure 9 — The effect of the low-pass filtering in time on a SSB subspace with (A) only the noise, and (B) the same noise with a circularly polarized signal at 3 dB SNR.

spectral leakage. For CW the default decay is scaled four times longer. The number of measurements in a subspace is four times lower and a longer decay increases the number of effective measurements.

Low Pass Filtering in Subspace

The low-pass filtering in space is performed by filtering the spectrum of the 2D space. A 2D FFT calculates the spectrum of the 2D space. The spectrum is multiplied with a 2D low-pass filter (LPF). The result is transformed back with a 2D IFFT (Figure 10).

The coefficients for the LPF can be calculated from the required impulse response in a subspace. If the required response is known, the coefficients are obtained by the inverse-IFFT (or FFT) on that response. The FFT of the response is the product of the spectrum of the input space with a single location filled with an amplitude of 1 and of the LPF. So the LPF coefficients can be calculated by dividing the FFT of the response and the FFT of the impulse in the input space.

As a response a Tukey (tapered cosine) window is used in the phase and in the ratio direction in space, because it is flexible in setting window shapes. The total 2D low-pass filter is the product of the LPF for the phase direction and for the ratio direction.

A signal will concentrate the power at the signal's location (Figure 11). This increases the level at that location at the cost of the other locations in the subspace. This is a positive effect, because it enhances the signal's location.

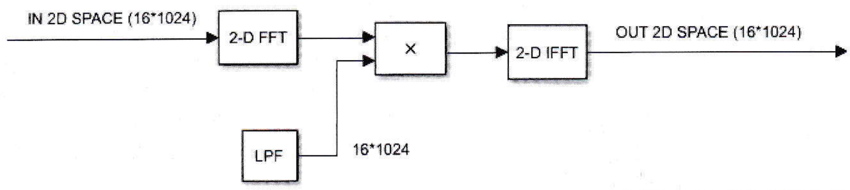

Figure 10 — Block diagram of the low-pass filtering in space.

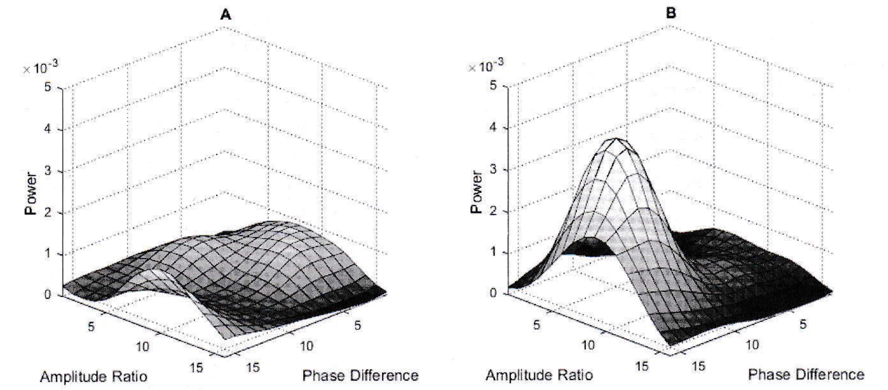

Figure 11 — The effect of combining a time filter and subspace low-pass filter in a SSB subspace with (A) only noise and, (B) with a circular polarized signal at 3 dB SNR.

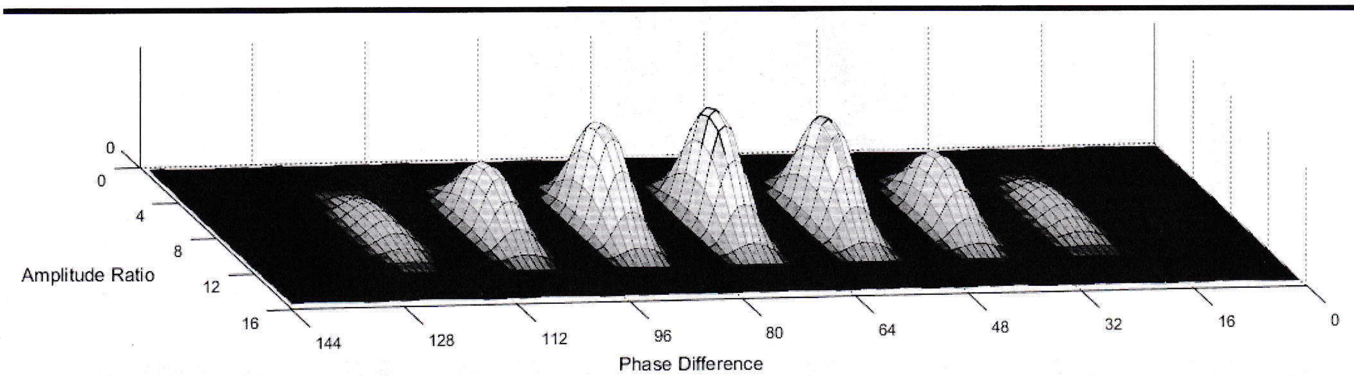

Figure 12 — The response of the 3D filtering in subspace and over frequency displayed in the 2D space. Nine subspaces are sliced to show the response of the total LPF calculating the concentration of power in a volume in 3D space.

Low Pass Filtering over Frequency

Until now the filtering was limited to each single subspace. It is expected that the signal components will overlap in consecutive subspaces. So there is also a concentration of power in subspaces over frequency. The LPF can be extended in the same way to perform a low-pass filtering of locations over subspaces. Figure 12 shows the response using a Tukey window over 9 subspaces.

The combination of 4 bins in a single subspace for SSB acts in-principle as a low-pass filter over frequency. Just like the quantization of the locations in subspace to 16×16 acts as a low-pass filter in the subspace. The span of the 4 bins is about 4.5×15.625 Hz = 70 Hz.

Spatial Filtering

A signal is spatial filtered by weighting the signal components in the frequency bins according the location in a weighted space. The weighted space is based on an estimate of the signal locations — an estimate of the noise as a reference and a weighting function. The locations in the weighted space contain a measure of whether it is a signal or noise. The locations of the signals in the bins are used as pointers to the weighted space.

Estimating the Location of the Signal

The current time-filtered and space-filtered subspaces contain the estimate of the location in space of the present signal. It is a result of calculating the concentration of power in the 3D space over time.

Estimating the Noise Floor

The sum of the power on each location over all subspaces is called bulk space. When it is sampled, if only noise is present, it represents the noise floor of the receiving system and the received noise. It is expected that the level and the statistics of the noise is stationary over time. The noise floor is a simple representation or estimate of this statistics.

The noise floor is shaped by the statistics of the locations in space and the corresponding power: (*concentration* × *power*). The power depends on the selected antenna signal. The locations in space are set by the noise (amplitude) statistics of both antennas. The result is that at the borders of the ratio R on the y-axis the measurements are less frequent. This frequency dominates in the noise floor level, because it sets the concentration over time. At the borders of the y-axis the noise floor level will be lower.

In the example (Figure 13) instead of using the sum of the antenna A and B signals, the strongest signal according location in space is selected for the power. The noise floor has to be scaled to the actual receivers (audio) bandwidth to get the noise floor in a subspace. Only subspaces within the receiver's bandwidth contain noise. If the bandwidth is smaller the sum of subspaces is lower.

In a practical situation the SSB/CW signal can be present making it difficult to sample the noise floor at the right moment. The sum over all locations in bulk space represents the total power. It is expected that this total power is minimal when only the noise is present. The noise floor is detected by monitoring this power over a period of time and sampling the bulk space when the power has its minimum. Because there will be variations on the noise floor, an averaged minimum is needed. A moving average over 8 noise floor estimations is calculated and sampled.

Weighting by Soft Limiting

Weighting uses information of the noise floor and the current time filtered and space filtered subspaces (2D space) containing the location estimation of the signal. The maximum power in the noise floor is an indication of the expected signal levels. The 2D space is scaled (normalized) to this maximum. This allows an alternative sigmoid function (Figure 14) to be used as a soft limiter. It calculates the weighting gain in space based on this scaled space. The gain is called soft-limited gain (Figure 15).

The soft-limited 2D space locations contain a measure for whether it is a noise (Figure 16A) or signal (Figure 16B). The locations of the signals in the frequency bins are used as pointers to the soft-limited 2D space. The resulting soft-limited gain is used for weighting the signal.

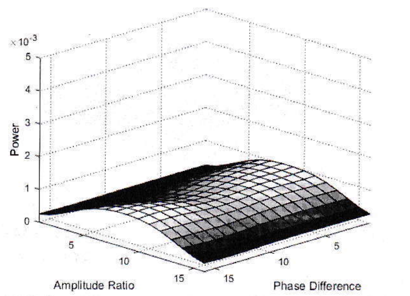

Figure 13 — The scaled sampled noise floor as a result of combining time filtering and space filtering of only noise.

Selecting Antenna Space and Signals

Antenna Space

The location in subspace is the result of the selected antenna signal space. Up until now the antenna signals provided by the linear polarized antennas A and B are used. Any linear polarized signal can be seen as the sum of two circular polarized signals. The Left and Right hand circular polarized antenna signals L and R can be calculated by summing both signals with ±90° phase shift. In the frequency domain the phase shift can be calculated much easier by multiplying the complex numbers of the spectrum of one antenna with +/-1i. In this way the linear polarized antennas A and B provide also the signals of a left hand and of a right hand circular polarized antenna.

Signals

The location in subspace is always set by the amplitude of the two signals of the selected antenna space. The power used in the subspaces however is set by the selected signal. For the best estimate of the location of the signal in space, the power of the selected signal needs to have the best SNR. Even if the location of the signal in space is found, the sum of the signal and the noise will be passed by the noise reduction. It is only filtered by the noise reduction; the noise is

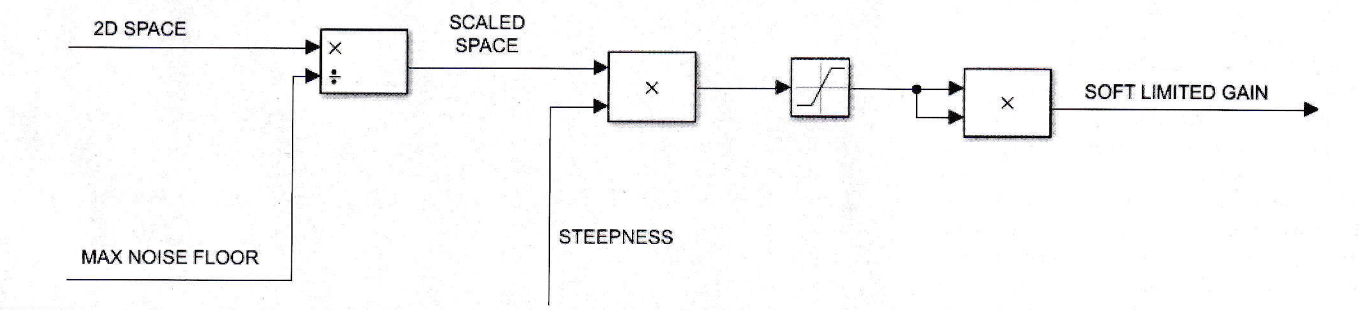

Figure 14 — Block diagram of the alternative *sigmoid* function used as a soft limiter. The scaled space is multiplied by the steepness limited between 0 and 1, and squared.

not subtracted. A few dB or more in SNR can be gained by selecting the signal with the best SNR, especially when using noise cancelling. In that way we gain from having two antennas.

Five signals can be selected:

- Noise cancelled
- Diversity (Mono)
- Diversity (Stereo)
- *A*/*L* (antenna *A* or *L*)
- *B*/*R* (antenna *B* or *R*).

Any combination of the two antenna signals results in a location in space that is noise cancelled.

The polarization diversity signal is a selection between *A* and *B* (or *L* and *R*) signals based on the location in subspace. The strongest antenna signal is selected and used for estimating the location. With Mono the noise reduction (soft-limiter gain) acts on this signal and with Stereo on both antenna signals (enabling Stereo Diversity Noise Reduction). In principle, selection of one of the antenna signals is not needed, but it skips the need for setting the noise cancelling.

Implementation and GUI

The processing, including the GUI of Figure 17, was developed and prototyped in Matlab/Simulink. At the right, the spaces are displayed as an image. The subspaces of a SSB signal are displayed showing frequency selective locations in the subspaces (white dots).

All noise reduction settings are located on the bottom row. The AGC is at the end of the processing. The required space and signal can be selected, the detection of the noise floor can be enabled, the noise floor itself and a dc level can be used as an optional threshold. The detection Noise Floor switch enables finding and updating the noise floor. While it is switched ON, a new noise floor is tracked until switched OFF. The settings of the Time Filter, Space Filtering and the Soft Limiter are at the bottom right.

One Mouse-Click Noise Cancelling

In the GUI bulk space, the sum of all subspaces, can be selected for displaying. In Figure 18 bulk space shows the location of an identifiable man-made noise source in space. It is resized to 128×128 using bilinear interpolation for displaying and for a higher location resolution. The maximum in the display (white cloud) indicates the location. The location can be selected by clicking on it in the displayed image. The acquired coordinates in the image can then be recalculated to the settings needed for the noise canceller. The antenna with the strongest man-made noise is known from the location and it's signal is scaled down before subtracting. So only a single mouse-click is needed to set the noise canceller. The location is valid only if it is not already noise-cancelled by accident. Clicking on the opposite (or another) location in space will reveal the correct location and will show an increase in noise level on the S-meters.

If the man-made noise is present only on one of the antennas *A* or *B*, the location will be spread over the *x*-axis, because the phase difference is in the noise. The maximum is then spread also and is less visible. Switching to the *L*/*R* space will reveal the location of the maximum at ratio $R = 0$ (amplitude *L* equals amplitude *R*) and at a phase difference of 0° or 180°, because it is a linearly polarized signal with a fixed orientation. Instead of noise cancelling in the *A*/*B* space, the *B* or *A* signal can then be selected for processing.

Local man-made noise can also be circular polarized, especially in the near field. The signal at antenna *L* or *R* can be zero. In that case the man-made noise will become

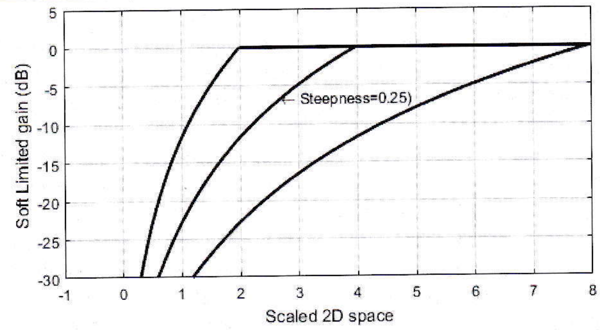

Figure 15 — The alternative *sigmoid* function in decibels. The soft-limited gain as function of the scaled 2D space with steepness: 0.5, 0.25 and 0.125.

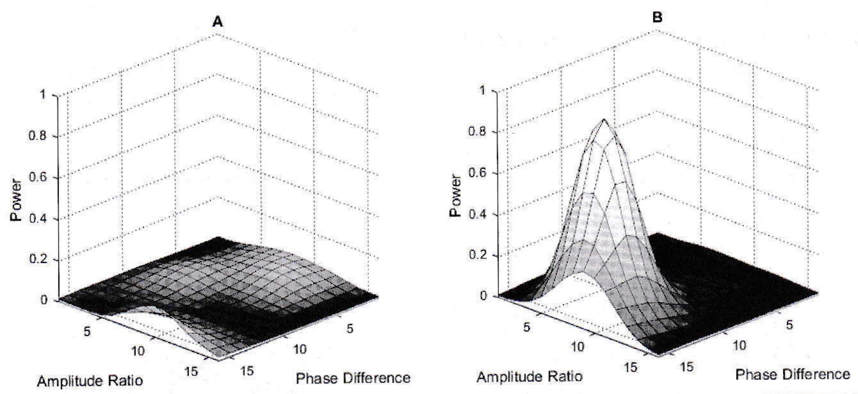

Figure 16 — The soft-limited (weighted) subspace with (A) only noise and, with (B) a signal at 3 dB SNR. Vertical scale is the soft-limited gain.

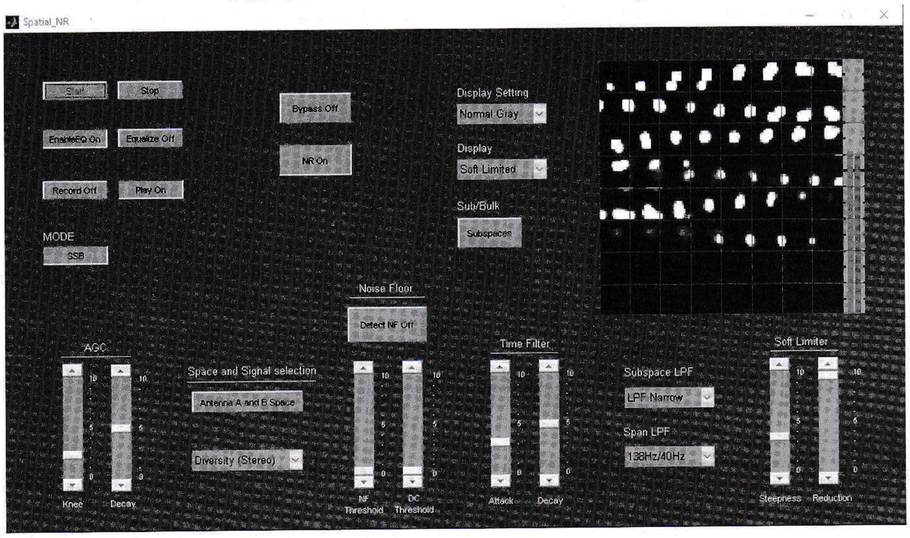

Figure 17 — The GUI with subspaces displaying a frequency selective fading SSB signal.

best visible in the A/B space at ratio $R = 0$ with a ±90° phase difference. Instead of noise cancelling in the L/R space, the R or L signal can then be selected if polarization of the signal allows.

The noise canceller setting result in one space can be converted to the setting needed in the other space. So noise cancelling can be done in the space showing the man-made noise location best and so giving the best noise cancelling setting. This makes it also possible to switch between spaces without needing to update the noise cancelling setting.

In practice, the noise on both antennas will not be perfectly uncorrelated and the noise can show a preference in location. In an urban area this can be caused by a dominant man-made noise source, by poor balancing of the loop antennas or by coupling with surrounding objects or other antennas. This location will be visible in bulk space and the noise canceller can be set to this location for best noise cancelling.

Noise cancelling in effect maximizes the polarization mismatch and/or the radiation pattern for maximum attenuation of the noise source.

Propagation

Knowledge of propagation is needed to understand the signal location[5,6] in space. All shortwave signals travelling through ionized media [in the presence of Earth's magnetic field, — *Ed.*] propagate by two modes, the ordinary (O) and extraordinary (X) mode[7,8]. The signals of both modes are opposite-elliptically polarized. NVIS signals on 80 m in the Netherlands are nearly circularly polarized during the day, because the ordinary wave predominates. Selecting the left-hand circular signal increases the SNR up to 3 dB. When selecting the right-hand circular signal the SNR can decrease more than 20 dB. Also, on the other bands like 40 m, the signals can be elliptical polarized. When both modes are present at about equal amplitude the resulting polarization becomes linear and the orientation will rotate (drift) over time. Predominantly linearly polarized signals are rarely stable in linearity and orientation. The propagation attenuation and delay of the X and O mode signals drift independently in time. If the signal on one antenna drifts below the noise the phase difference will become noisy. The location of a drifting elliptically polarized signal traces a nice circle in subspace over time.

Multipath reception can cause frequency selective fading (FSF). It is not only expressed in the frequency selective amplitude, but also in the frequency selective polarization. As a result, the locations in space will be frequency selective. Because the subspaces are filtered independently, the noise reduction can very well cope with that. It limits however the use of the frequency span filtering. A lower

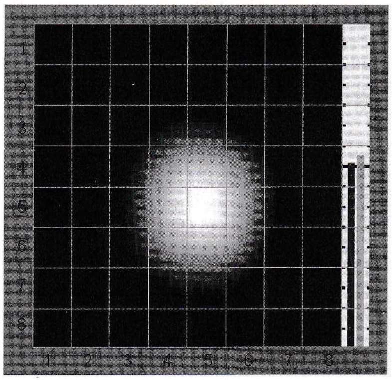

Figure 18 — Bulk space showing the white cloud representing a single 10 dB man-made noise over noise source (scaled to maximum).

span bandwidth can be selected when FSF is present. If noise cancelling has to be used the signal can fade on the noise cancelled location. Diversity is the preferred signal selection when frequency selective fading is present and there is no dominant man-made noise source.

When the polarization of the signal is mainly circular or elliptical, the A/B space is preferred. The locations are mainly on the x-axis and the phase difference is a well defined ±90°.

If the signals are linearly polarized and rotating, the location can be at the edges of A/B space and the phase difference can become random over the entire 360°. The L/R space is then preferred, because the linear signals are on the x-axis and the phase difference is better defined.

Both the spaces A/B and L/R use the phase difference. If one of the antenna signals drops in amplitude because of drifting polarization, the phase difference becomes noisy and even can be set by the noise if that signal becomes zero. This affects the concentration in the subspaces. It is however compensated by the 3 dB higher signal level of the other antenna signal.

Phase Coherent Receivers

Both receivers must be as identical as possible. Most critical is the phase behavior. Common names are phase coherent, phase locked and phase synchronous. However, they don't necessarily indicate the same phase. It is essential that both receivers use the same local oscillator signals for equal frequency and equal phase. It is not practical when the phase difference is not fixed and jumps to other values when tuning over the band. As a result every change in frequency needs measuring the changed phase difference again or resetting the noise cancelling again.

Less critical are the phase/gain frequency responses of both receivers. As long as the phase/gain frequency response over the audio band is fixed it can be equalized in the processing. A non-flat frequency response can be flattened in the processing. The pass band must be flat, because a single noise floor reference is used for all frequency bins. HF filtering affects the phase difference also when tuning over the band. Fixed differences like gain differences and phase offsets can be calibrated out and so are less critical. The AGC must be performed after the noise cancelling and noise reduction.

There are professional analog receivers designed for diversity reception (such as Collins R390A) or direction finding (such as Telefunken Telegon) that can be used. However they are rare or expensive. Some ham transceivers support diversity reception, but the phase difference is not fixed. Another option is to make two analog receivers (such as two Elecraft K2 receivers) phase-fixed by sharing the same local oscillators.

SDR receivers are a modern and practical implementation of potentially identical phase-fixed receivers. During the development of the noise reduction, I switched from using two analog K2 receivers to using the ANAN200D.[9] The ANAN200D receivers however — as practically all SDR receivers for ham radio with multiple spectral capture units — are not phase fixed. The phase difference is not known and can jump to other values when tuning. That is the main reason receiver equalization in the processing is still needed. Equalization uses an external noise source connected to both receiver inputs. *PowerSDR* software runs on a second separate computer so all processing power of the first processing computer stays available for the noise reduction. The analog line-out of the ANAN200D is fed via a sound card to the processing computer. Running it all on one computer might be possible, but the phase difference in the *PowerSDR* audio interfacing is not controlled.

Test Signals and Screen Recordings

Most actual band signals vary too much in SNR for testing purposes. The question is at what minimum SNR level are the signals still readable. Because of fading actual band signals are usually not constantly at that controlled level.

One set of test signals can be made artificially in a separate *Simulink* model by combining actual band man-made noise and actual band signals with a very high SNR. Signal parts with about equal signal strength are selected out. The content of the conversation is not relevant. In that way the overall SNR can be controlled by adding the signal to the noise with a tunable level. This is valid, because actual signals are also the sum of band noise and signal. In the *Simulink* model the rms levels of the noise and of the signal can be measured so the SNR level can be set to known values. A time constant of 125 ms is used in the rms measurements. The peak rms value of the signal represents the peak value of the transmitter's power and so of the signal strength.

A second set of test and demonstration signals are actual band signals without modifications. Artificial noise and CW are also used as a test signal. Audio noise reduction examples and screen recordings are presented[10] in audio samples and screen recordings.

Discussions

A single strong man-made noise source has to be resolved for best noise reduction results. It behaves like a signal and has a preferred location in space. This affects the noise reduction, because the location of a signal is pulled to the location of this noise source depending on the signal's amplitude. For the same reason, diversity is not effective when a strong man-made noise source is present. Diversity is preferred if no single dominant man-made noise source is present. A signal cannot then be cancelled or attenuated by noise cancelling.

Conclusions

With the one mouse-click noise cancelling even the weakest identifiable man-made noise source or any noise preference can be cancelled fast and easily. The presented stereo diversity noise reduction is very effective by using the signals of two orthogonal antennas. It is based on the concentration of signal power over time in a 3D space. Only minimal information of the SSB and CW signal is needed. The noise reduction has one problem. Except for the very weak signals you can't tell anything about the signal strength just by listening. You must switch off the noise reduction for that.

The filtering in the subspaces can be regarded as polarization filtering, not including direction of arrival. In that way the noise reduction uses filtering over time, frequency and polarization. At the same time it can benefit from selecting the best space and best signal for the actual band conditions. It enables stereo diversity noise reduction.

Selecting the *best space* and *best signal* are the most relevant settings. In an urban area the noise cancelled signal is the best selection most of the time, unless the location of the man-made noise source coincides with the signal location in space.

Displaying the polarization behavior of the signals over frequency gives feedback for a better understanding of the propagation trends over time. The presence of elliptical polarization or frequency selective fading becomes visible. It also helps in selecting the best space and *best signal* for the band conditions.

Displaying both the location of the signals and of the man-made noise enables you to see if the location of the man-made noise and the location of the signal coincides.

Jan Simons, PAØSIM, has been a licensed radio amateur since 1975. He works as a mixed analog/digital electronic engineer for Oce/Canon, a printer and copier company in The Netherlands. Developing mixed signal ASICs for piezo inkjet print heads and methods for using the piezo as a sensor for detecting underperforming nozzles is a part of his job. As a radio amateur, besides the technical part of the hobby, he especially enjoys working DX and CW on shortwave.

Notes

[1]Additional information and noise reduction examples of SSB and of CW signals; **www.pa0sim.nl/QEX.htm**.

[2]J. M. M. Simons, PAØSIM, "Effective Directivity for Shortwave Reception by DSP," *QEX*, July/Aug. 2006, pp. 37-45.

[3]**www.katjaas.nl/FFTwindow/FFTwindow&filtering.html**.

[4]**www.mathworks.com**.

[5]B. A. Witvliet, PE5B, "Near vertical incidence skywave: Interaction of antenna and propagation mechanism," PhD dissertation, University of Twente, Enschede, The Netherlands, 2015.

[6]M. C. Walden, "High-Frequency Near Vertical Incidence Skywave Propagation," *QEX*, Jul./Aug. 2017, pp. 21-34.

[7]**www.pa0sim.nl/XOpropagation.htm**.

[8]Eric Nichols, KL7AJ, "Gimme an X, Gimme an O," *QST* Dec. 2010.

[9]**https://apache-labs.com/**.

[10]Audio samples and screen recordings: **www.pa0sim.nl/QEX.htm**.