WHITE PAPER

PLATH SIGNAL PRODUCTS

INTERCEPTION CONFIDENCE · DIRECTION FINDING ASSURANCE

High Resolution Methods for Direction Finding

1 The Need for High Resolution Direction Finding

One of the most important functions of an ESM (Electronic Support Measures) or a SIGINT (Signal Intelligence) system is the geolocation of the intercepted emissions. Knowing the location of an emitting target is essential for several purposes. First and foremost, it indicates the location and disposition of forces. Second by clustering different types of emissions in given areas it can also give an indication of the composition of forces. Since the function of direction finding is an essential element in ESM or SIGINT it is the geolocation of emissions that is of particular interest. Two or more (mostly three) line of bearings assumed to be emitted from the same target and measured almost at the same time will intersect at more or less the origin of the emission.

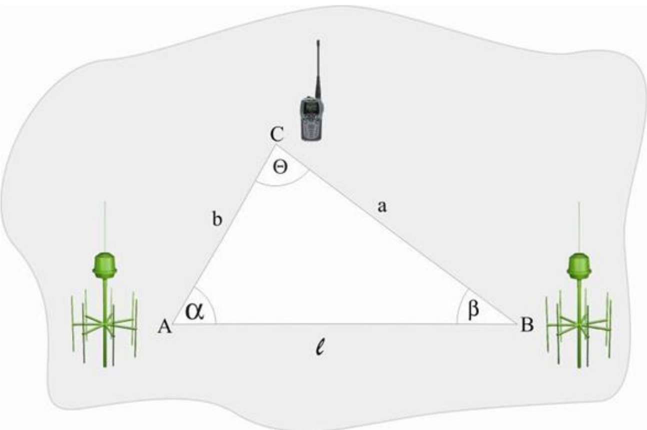

Figure 1: Geolocation of an emitter: the first DF at site A and the second DF at site B determine the bearings of an unknown emitter at site C

Another possibility is to measure the time of arrival of the signal at different sensors. The difference of the interception time – often referred to as the time difference of arrival and abbreviated TDOA – directly relates to the range difference from the emitter to the different sensors, which in turn is taken to compute the position of the emitter.

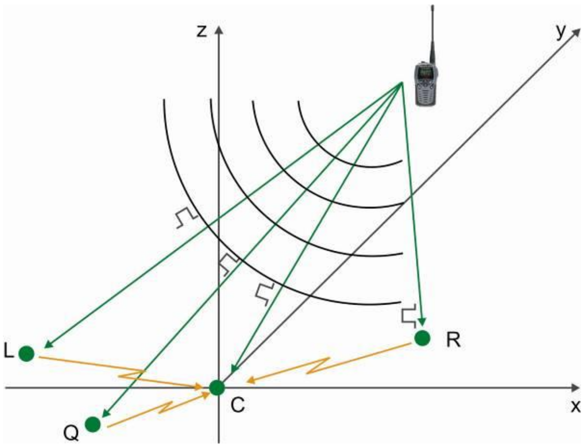

Figure 2: Principle of operation for TDOA

Although the geolocation of the emitter is an essential element in ESM or SIGINT we will first and foremost focus on the underlying direction-finding methods as these are dominating prerequisites for an exact location. This paper discusses especially high-resolution direction-finding methods in the context of their typical use-cases.

2 Use Cases for High Resolution Direction Finding

To determine bearings to the signal emitters we can rely on principles like Watson-Watt antenna array or correlative interferometry. From those bearings (using triangulation) we may determine the geo-location of the emitter. The geo-location of the emitter is important information, especially in the establishment of fundamental knowledge, which feeds information to military decision makers on every level. Additionally, geo-location of the emitters in tactical scenarios brings unquestionable benefits to tactical operation. But how reliable is the geo-location computation if more than one signal occupies one frequency?

2.1 Strategic decisions based on HF frequency spectrum monitoring

Strategic systems are fed with information from their subsystems, where one such can be a HF COMINT system. The nature of HF radio waves propagation allows for global communication, but also brings challenges: As HF signals cover large distances - even around the world - there is a possibility, that multiple signals might be received on the same channel. Also due to HF frequency phenomena we might receive groundwave and skywave of the same signal that leads to signal fading. Both aspects in combination means we are dealing with two cases of overlapped signals: correlated signals and non-correlated signals. Non-correlated signals can be separated and geo-located by means of Super Resolution algorithms, while these algorithms tend to fail for correlated signals. In this case holistic statistical methods are the weapon of choice for exact geo-location.

Figure 3: Typical antenna array for strategic HF interception and direction-finding system

2.2 Tactical decisions based on V/UHF frequency spectrum monitoring

In tactical operations V/UHF (i.e. frequencies above 30 MHz) communication plays an important role. But also, V/UHF signals overlapping on a single channel may occur. In typical land-based operation, the chance that we may observe more than one transmission on the same frequency is very small. Electromagnetic waves in V/UHF do not propagate on a big range since they are limited by the quasi-optical horizon. The optical horizon takes the earth's curvature into account by limiting the propagation to the optical horizon an observer would see at his current eye level – or in our case the current antenna level. The term quasi-optical pays tribute to the fact that electromagnetic wavelengths in the frequency range of V/UHF are orders of magnitudes longer than optical wavelengths. Therefore, theses waves tend to travel somewhat further than the optical horizon, almost 4/3 of the optical range to be (more) precise. However, electromagnetic waves in the V/UHF domain share another similarity to optical waves: They are also subject to attenuation due to different obstacles such as concrete buildings. At sea unintended overlapping of V/UHF signals may be observed, but only in a very rare case, when particular atmosphere conditions are causing radio waves super-refraction. This can unexpectedly enlarge the propagation distance of the radio wave. In the case of airborne COMINT systems the problematic with overlapping transmissions on the same frequency is noticeable. There line of sight to the emitters on ground is unobscured and due to the height of the antenna the range to quasi-optical horizon is increased enormously.

Figure 4: The system MoGeFA was introduced into the Bundeswehr for tactical V/UHF frequency spectrum monitoring and direction-finding

Hence, the problematic with overlapped signals in V/UHF frequency spectrum monitoring can be observed and limit reliability of sound decision making especially in airborne operations.

It is also important to mention that Super Resolution is going in tandem with beamforming algorithms-SR/BF. This methodology allows us first to distinguish the exact signal of interest, even if it is weaker and co-existing on the same channel other signals. Then the selected signal of interest undergoes the beamformer algorithm. Thanks to beamforming we can put the weight in our computation on a desired direction thereby improving the reception possibilities by even several orders of magnitude as well as extrapolating the elevation component of that particular signal.

3 High Resolution Direction Findings Methods

In relation to the different use cases - detailed in the section above - also different methods are being used. These generally comprise the use of spinning antennas, amplitude comparison, phase interferometry, array-based high-resolution and finally time difference of arrival methods. Whilst a mechanical or electronical rotating antenna, e.g. a three-channel direction finder with a nine-element directional antenna, may not cover the entire azimuth simultaneously and hence transient signals might escape. In contrast, systems that use amplitude comparison methods employ several - often up to sixteen or more - fixed antennas that either point in different directions as directional antennas or are set up in a special arrangement relative to each other. When a signal reaches the antenna array, the two antennas with the strongest or second strongest antenna signal are determined. The ratio of these two antenna-signals as well as the underlying antenna pattern is used to determine the angle-of-arrival of the electromagnetic waves.

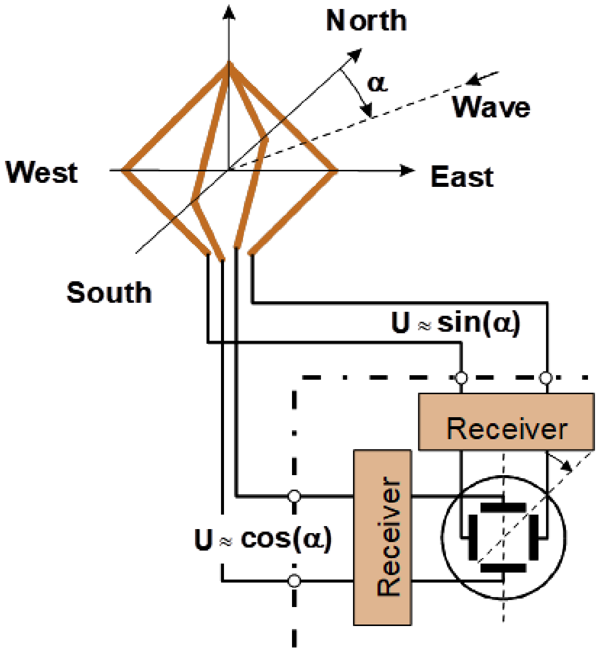

Figure 5: A well proven amplitude comparison is the Watson-Watt direction finding method

This approach is much more robust while detecting transient signals. If the phases of the antenna signals are compared instead of the amplitudes, another method can be described, namely that of phase interferometry for determining the angle of incidence of electromagnetic waves. This method uses the relative phase between the individual antenna signals or the sensor outputs to determine the direction. However, there is a disadvantage in the ambiguities that can arise when the antenna base spacing is greater than half the wavelength.

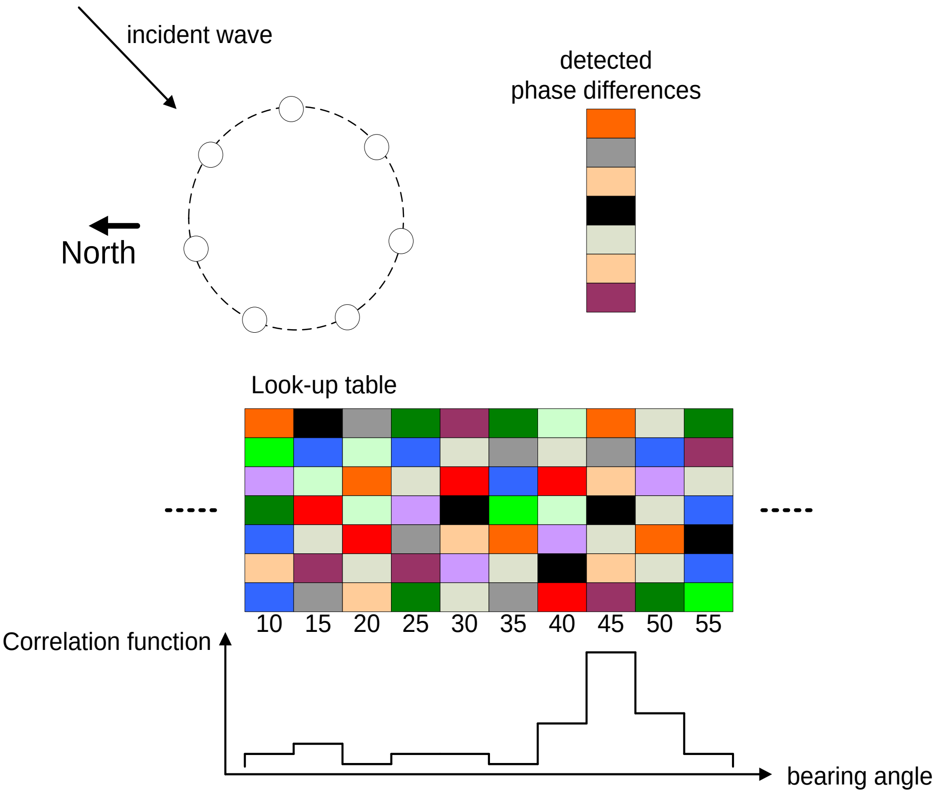

Figure 6: Phase interferometry is widely implemented into so-called Correlative Interferometry Procedure, whereby pre-calibrated values are correlated to actual phase-differences

As with the amplitude comparison method, phase interferometry does not require moving parts and also detects transient signals unless switching between antennas is done electronically. This becomes necessary, for example, when the number of sensor channels does not match the number of antennas. Furthermore, the amplitude ratio or phase relationship on which these two methods are based is no longer valid if the signal paths to the two sensors are completely different. This is especially the case when there are multiple, overlapping signals or interferences as mentioned above. Another possibility is based on measuring the time difference of arrival (TDOA) of the received signals. For example, the method is based on the property that a cross-correlation between a signal and a delayed version of the signal is maximized if the delay used in the cross-correlation is equal to the TDOA. Cross-correlation is a measure of similarity of two series as a function of the displacement of one relative to the other. It is also known as a sliding dot product because it can be imageable as two functions sliding over one another whilst for every position each overlap in the functional domain is multiplied and the results of the multiplications are then summarized. This now yields a new function that depends on the position of the two functions relative to one another. This particular function has a peak once the similarity of the overlap of the two function is maximised. However, this property also does not hold when there is interference and multiple signals of interest. Since the peak of the cross-correlation function at the TDOA occurs when only a single signal is present, and since the width of the peak is approximately equal to the coherence time of the signal, it follows that the peaks due to multiple signals merge into one peak when the distance between adjacent TDOAs of the signals is less than the coherence times of the signals.

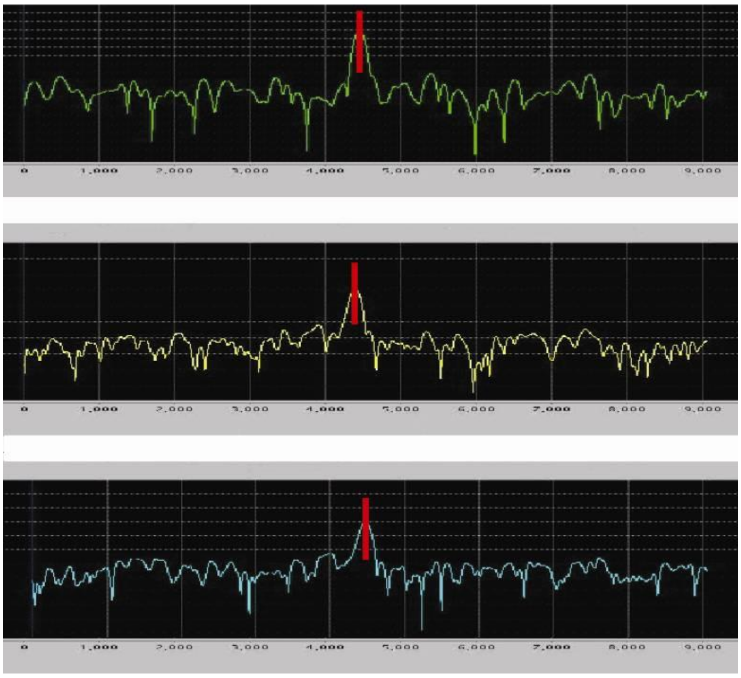

Figure 7: Example for correlation results of signals intercepted at different sites yielding a sharp peak

Thus, all methods for determining the direction of incidence of electromagnetic waves have in common that they become inaccurate or strongly distorted as soon as several signals spectrally overlap or interfere with each other. It follows equally that under these circumstances methods for High Resolution Direction Finding must be used.

3.1 Implementation of Super Resolution

The algorithms developed to determine the direction of multiple, overlapping signals are called Super Resolution (direction finding). This name is derived from the ability to resolve two signals, which arrive at a close azimuth; in principle the difficulty to separate two (or more) signals increases the more they overlap in azimuth. Furthermore, for these signals it is important whether there is a statistical distortion of the calculated azimuths. This distortion results in an erroneous calculation of the azimuths - regardless of how well they are resolved. Since the measurements of the angles of incidence are statistical calculations, this distortion is a distortion of the mean value of the probability density function respectively. In probability theory, a probability density function (PDF) is a function whose value at any given sample in the sample space can be interpreted as providing a relative likelihood that the value of the random variable would be close to that sample. In more simple words the PDF shows where most values accumulate and how likely characteristic values are in certain areas. For most PDFs the accumulation would correspond to the mean value and the possible (or probable) range of values to standard deviation. Both of the previously mentioned aspects are often traded off against each other, since an increase in one aspect can have a negative effect on the other. In addition, the variability plays an important role as well. The variability is the standard deviation of the PDF and which represents the range of azimuths over which the computations are expected to vary in the presence of noise. Most Super Resolution methods require a search in the spaces spanned by the steering vectors. That is, the signals are assumed to come from a particular direction, and the steering vectors are computed using just that assumption. It is further assumed that the correct answer is the bearing that gives the largest spectrum. The smallest spectrum is attributed to the underlying noise. Thus, a series of linear equations are solved for the incremental azimuths around the antenna array. The two methods discussed here for calculating the azimuth angles of signals incident on an antenna array are MUSIC and ESPRIT.

3.2 MUSIC

The MUSIC algorithm (multiple signal classification) is an estimation algorithm in digital signal processing. This algorithm renders it possible to determine the frequency and the direction of reception of a mixture of several signals with interference. Today, it is one of the most frequently used algorithms for determining the angle of an incoming signal. It thus belongs to the procedures of Super Resolution (direction finding), which is based on the eigenvalue decomposition of the signal correlation matrix. In this technique, the eigenvalues and eigenvectors of the signal correlation matrix are determined. Eigenvalues characterise essential properties of matrices, such as whether the corresponding system of linear equations is uniquely solvable or not. In this context, it is assumed that the largest eigenvalues are associated with signal vectors and the smallest eigenvalues are associated with noise vectors. Successive samples of antenna data are collected. Each of these samples corresponds to a single frame. Since RF signals are generally stochastic in nature, due to random noise and other reasons such as signal fading and modulation effects, the samples will generally vary from frame to frame - or from observation to observation.

The MUSIC method for spectral estimation is sensitive to fully coherent signals because the covariance matrix is then singular- which in other words means, if the signals are having the same phase and frequency, trying to solve the MUSIC matrix (which is then singular) can lead to incorrect answers (no solution or infinite number of solutions). In other words, the underlying system of equations of the signal correlation matrix cannot be solved unambiguously since the correlation is too explicit. The performance of the algorithm then depends on the degree of coherence in the signals. An advantage of the MUSIC method is that it can work with a non-uniform antenna array and no prior knowledge of the number of signals is required.

There are several methods for calculating the angular spectrum of an antenna and a particular algorithm, and each method has individual characteristics and limitations. These methods differ in the assumptions made about the signals. For example, it may be assumed that the two signals are not correlated (correlated signals usually mean that one signal is a multipath reflection of the other).

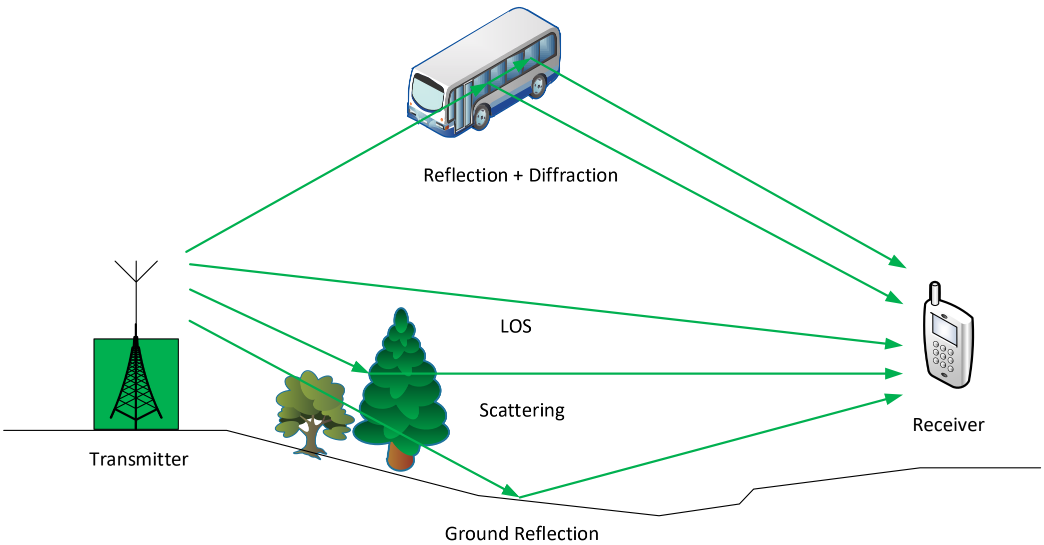

Figure 8: Signals that are reflected or diffracted concurrently to the line-of-sight (LOS) or direct propagated signal are strongly correlated

There are also assumptions about the configuration of the antenna array. For example, one technique might require a circular array and another a linear array. It should be noted that the vertical angle of arrival can be calculated in the same way with an appropriately oriented antenna, even though we are talking here about measuring the azimuthal angle of arrival. This is particularly useful in calculating the locations of shortwave emitters arriving at a non-zero elevation angle that could be reflected by the ionosphere.

3.3 ESPRIT

The ESPRIT algorithm uses the signal and noise space decomposition based on the correlation signal model as introduced for MUSIC to provide expectation-true estimates of the signal frequencies. However, while MUSIC works on the noise space, ESPRIT uses the eigenvalues of the rotation matrix of two time-shifted received signals. These eigenvalues are calculated by means of a similarity transform of this matrix, which is estimated from the signal space. This is to overcome a major shortcoming of the MUSIC algorithm, namely the need to generate, store and periodically check calibration data or to know analytical expression for the array multiplicity. With this need comes, on the one hand, the need for large data storage and, on the other hand, a continuous determination of the actual array manifold, which may vary over time due to disturbances in sensor locations, weather, nearby reflecting and absorbing bodies, and so on. The Estimation of Signal Parameters via Rotational Error (ESPRIT) method avoids these requirements by imposing a certain structure on the array geometry and then using this structure to its advantage. In particular, the array is assumed to consist of two identical subarrays, one of which is displaced by a known distance relative to the other. If the direction of the shift is unknown, the resulting direction estimates will all differ from the true values by the same unknown amount. A simplified interpretation of ESPRIT is that it is a generalised interferometer that can pick up multiple signals by using more than two sensors.

3.4 Holistic Statistical Methods

Unlike the previously mentioned methods that work a single channel or a single signal, Holistic Statistical Methods tend to take work on a larger proportion of the electromagnetic spectrum. Since the approach is holistic their advantage is based on the fact that also correlated overlapping signals can be separated. This refers in particular to the HF frequency domain where also backward reception is a phenomenon. This phenomenon occurs when signal is intercepted by an (mostly directional) antenna from the main lobe and the same signal travels around the globe and is then also intercepted by the back lobe. Also, unlike the MUSIC or ESPRIT these algorithms comprise a collection of harmonised algorithms that need to be processed in a coordinated sequence. Usually the first step in this sequence is the removal of noise. This is done by the determination of the noise floor for each frequency segment (frequency bin) in order to define the average value where no signal can be seen. This can be understood as signal normalisation for an entire spectrum by adjusting the detected noise to a previously defined noise model. Within the second step follows an elimination of broadband distortions. Spikes, lightning strokes and other broadband distortions are characterized by frames where almost all frequency bins show a significant signal above the noise floor. These frames are excluded from further analysis. After these pre-processing steps the actual signal processing follows, which in the presented case is a cluster search. This is no longer done in a spectrum, but - due to the fact that conventional direction findings methods are used – it takes the direction-finding histogram into account.

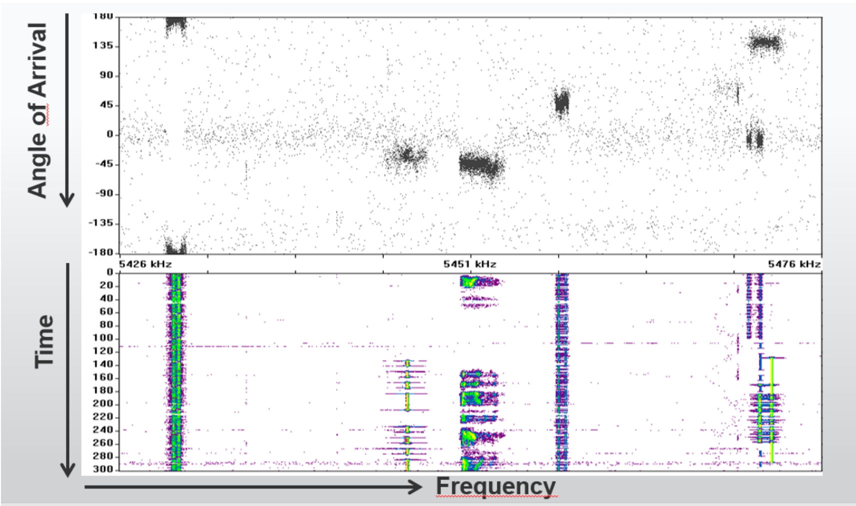

Figure 9: Result Data of a Wideband Direction Finder represented as a direction-finding histogram and a concurrent waterfall display

Within this direction-finding histogram certain peaks are identified as the nucleus for a clustering process by which distortions that may harm the direction-finding process are eliminated. Once these clusters have been compiled the clusters are continuously analysed according to time of occurrence, centre frequency, bandwidth, bearing and all deviations from these given parameters. Finally, one decisive advantage that is gained by the Holistic Statistical Methods, is presented by a burst analysis that in turn leads to a detection and classification of frequency agile or frequency hopping signals. Signals up to a certain length are analysed according to dwell time, burst length and bearing to find and characterize as frequency hopping signals.

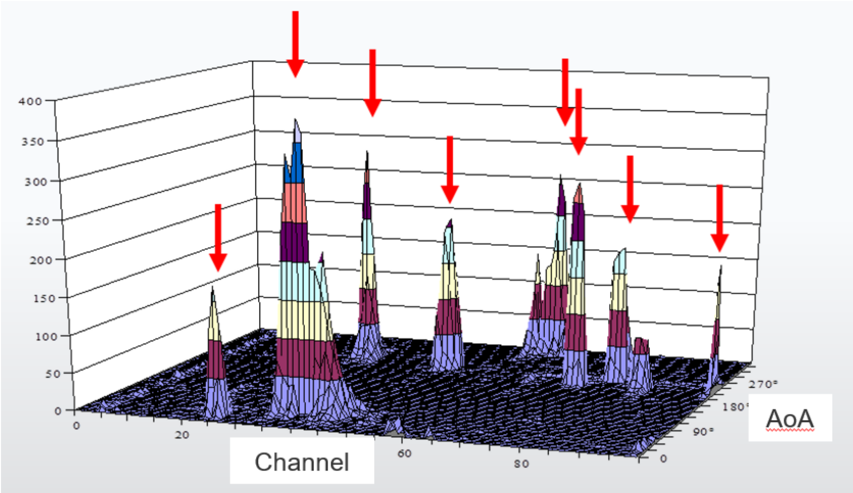

Figure 10: Cluster analysis of a set of distinct signals derived from the direction-finding histogram

Whilst the main objectives are the separation co-channel signals the topic, which will be discussed in the following, is the performance of DDA 5 in combination with the DFP 51xx series direction finders with respect to separation of co-channel signals in the HF domain. Although Super Resolution algorithms have the best ability to resolve two signals, which arrive at a close azimuth, SR is very expensive and requires large antenna fields and – according to the best of our current knowledge – are not yet ready for broadband automation.

The approach that is discussed in this paper is based on compact antenna configurations at a diameter of approximately 32 m at most (or even less) and a software-based solution with a resulting performance sufficient even for most of the use-cases.

Fundamentals

As shown in the Figure 5, the Watson-Watt DF system uses two receivers or two channels of a receiver connected either to an antenna array or, as in this example, to a cross-loop antenna. This antenna consists of four antenna elements arranged in pairs, orthogonal to each other, usually in a NORTH-SOUTH or WEST-EAST configuration. With the help of this antenna setup, a differential signal is generated per antenna pair, which is also referred to as AB or CD. Thus, each incident electromagnetic wave is split into electrical signals with a sine and a cosine component. By applying these signals as a voltage to the x and y deflection plates of a cathode ray tube, a Lissajous figure in the form of an ellipsis is created on the display of the tube. The axis through the focal points then corresponds to the angle of incidence. However, due to the use of only two receivers or only two channels, there is an ambiguity in the direction of bearing; for each direction of incidence can also be rotated by 180°.This ambiguity can be avoided by using another antenna (H) placed in the centre of the antenna array. This adds another channel that feeds the differential signal HA to the system.

Modern direction finders no longer display the signal components on a cathode ray tube, but process the signals digitally. Modern broadband direction finders do not use narrow, limited channels, but rather entire spectra are used for direction finding. In the case of the direction finders of the DFP51xx family for every measurement, i.e. for every time/frequency bin, there is exactly one value for the azimuth. As a general rule the azimuth for a given time/frequency corresponds to the strongest signal at the receiving antenna. Now, in order to separate two overlapping signals, each signal needs to dominate at least occasionally.

Simplex Communication/Time-Division

Signals will be separated reliably by the DDA5 in case there is a pause or interrupt of at least 720ms (standard configuration) independent of the angular separation. By this the DDA5 will reliably separate ALE 2G and HFDL as well as Morse keying.

Overlapping Signals

In case of truly overlapping signals the chance for successful separation is a complex function of the signal-to-noise ratios, the widths of the azimuthal distributions, and – of course – the azimuthal distance. Figure 11 shows examples (real-world) of azimuthal distributions, that are noise-free, i.e. all azimuths belong to one signal.

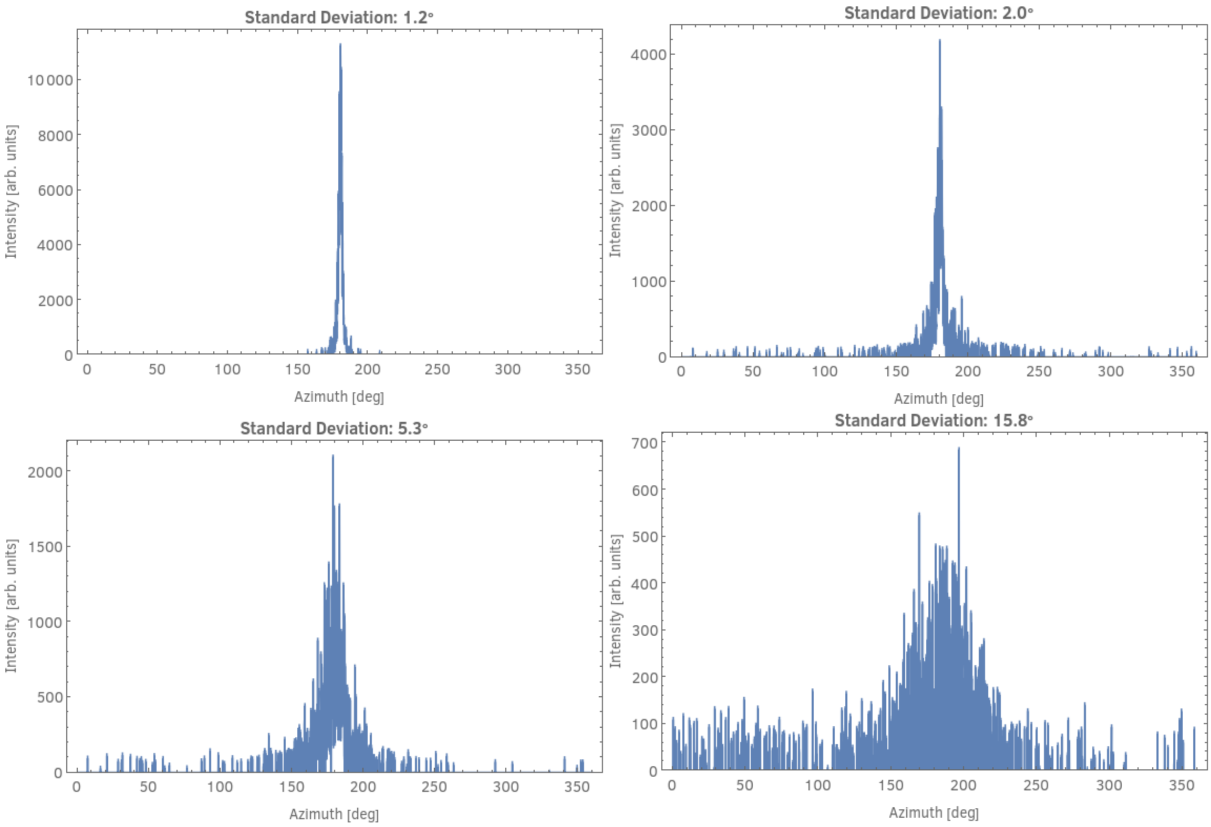

Figure 11: Azimuthal distributions of real-world signals

Apparently, each signal in a real environment is not limited to one exact angle-of-arrival. More over each signal adheres to a certain distribution. The width of the distributions – i.e. the uncertainty of the measurements – is caused by imperfect propagation conditions, when it comes to a propagation through and especially reflected by the ionosphere. This is due to the fact that the ionosphere is an amorphous, location-dependent, time-varying, and complex multilayer cloud of free electrons rather than a coherent mirror. Also, with increasing elevation of the incoming wave front the azimuth becomes less defined. Although at a first glance the distribution appears to be a normal distribution, the distributions follow a more a Lévy-type shape. The Lévy distribution is a continuous probability distribution for a non-negative random variable. Its exceptionally long and dominant tails make separation hard for close distances and concurrently it is evident that the azimuthal distance alone is insufficient to characterize the ability to separate co-channel signals. As a rule of thumb the azimuthal distance needs to be at least 3.8 times the azimuthal uncertainty (standard deviation of the distribution of azimuths) in order to be able to separate corresponding signals. Example of two combined (real-world) azimuthal distributions are given in Figure 12.

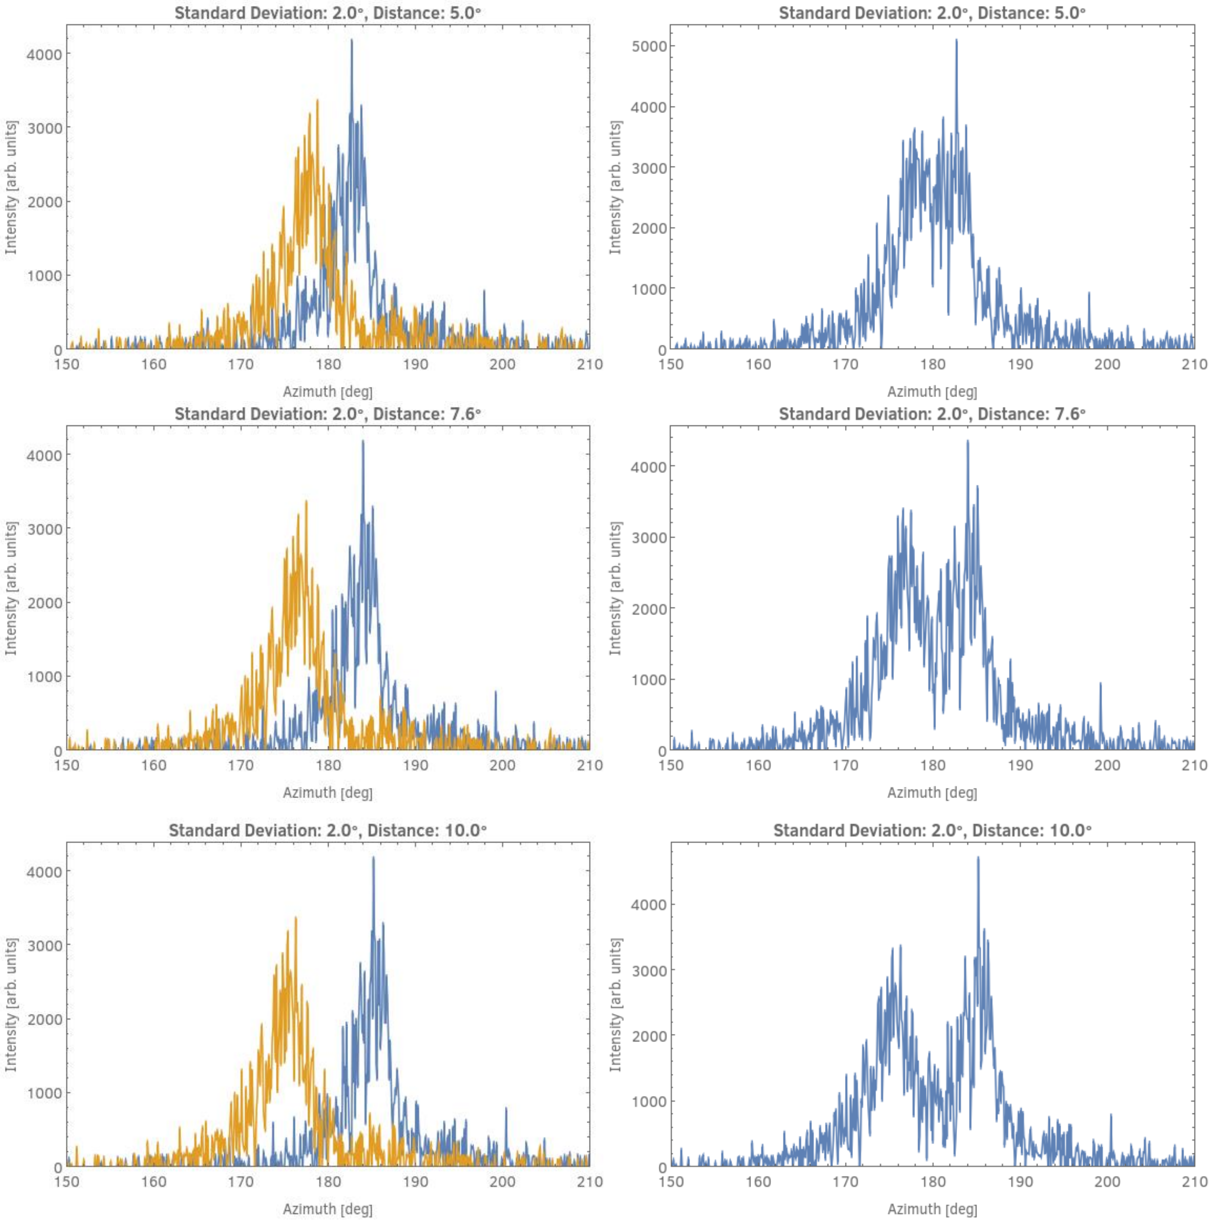

Figure 12: Example of two combined azimuthal distributions

Also, from the previous Figure 12 it is evident that without a-priori knowledge, i.e. number of given signals, types of their distributions, it is impossible to separate signals with azimuthal distances significantly below the mentioned threshold.

The distribution (real-world) of the standard deviations of measured azimuths is given in the following Figure 13. It can be seen that the distribution has its maximum at 2.1°. As a consequence, the minimum required azimuthal separation for signals of that kind would be 8.0°.

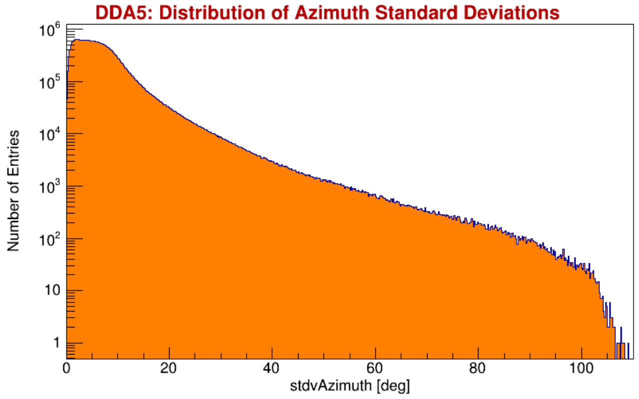

Figure 13: Distribution of azimuth standard deviations

The corresponding quantile function is shown in Figure 14. Half of the azimuth measurements (50% quantile) have standard deviations of less than 5.7°. By this, the minimum required azimuthal separation for signals of that kind would be 21.7°.

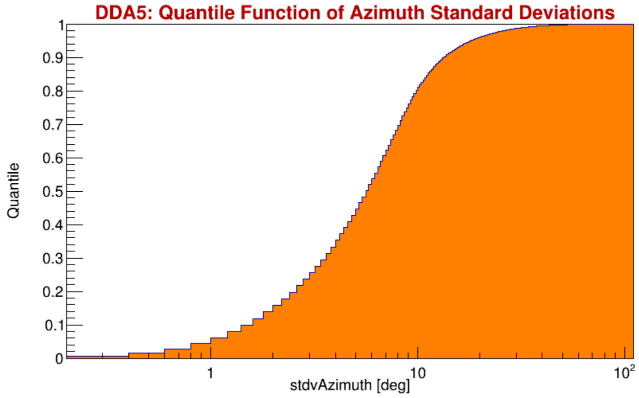

Figure 14: Quantile function of azimuth standard deviations

Real-world Examples

Now, in order to demonstrate the performance of our product line the results for three challenging real-world scenarios are presented in the following sections.

Typical Scenario

Figure 15 shows a scenario of two overlapping signals. Bandwidth/duration: 12 kHz/8s. The azimuthal distance measures 14.3°. From left to right:

1) raw azimuthal distribution (contrast enhanced), azimuth (ordinate) vs. frequency (abscissa)
2) raw spectrogram, time (ordinate) vs. frequency (abscissa)
3) segmentation result #1: FSK-2, 3.1° σ, 47.4 dB SNR
4) segmentation result #2: left tone of FSK-2, 1.8° σ, 37.0 dB SNR
5) segmentation result #3: right tone of FSK-2, 1.9° σ, 39.0 dB SNR

It can be seen from the results that a good separation of the two overlapping signals is possible. The fainter FSK-2 signal is segmented into the two tones. This does not impose problems, neither for triggering production nor for correlating to results of other direction finders in order to geo-locate the signal.

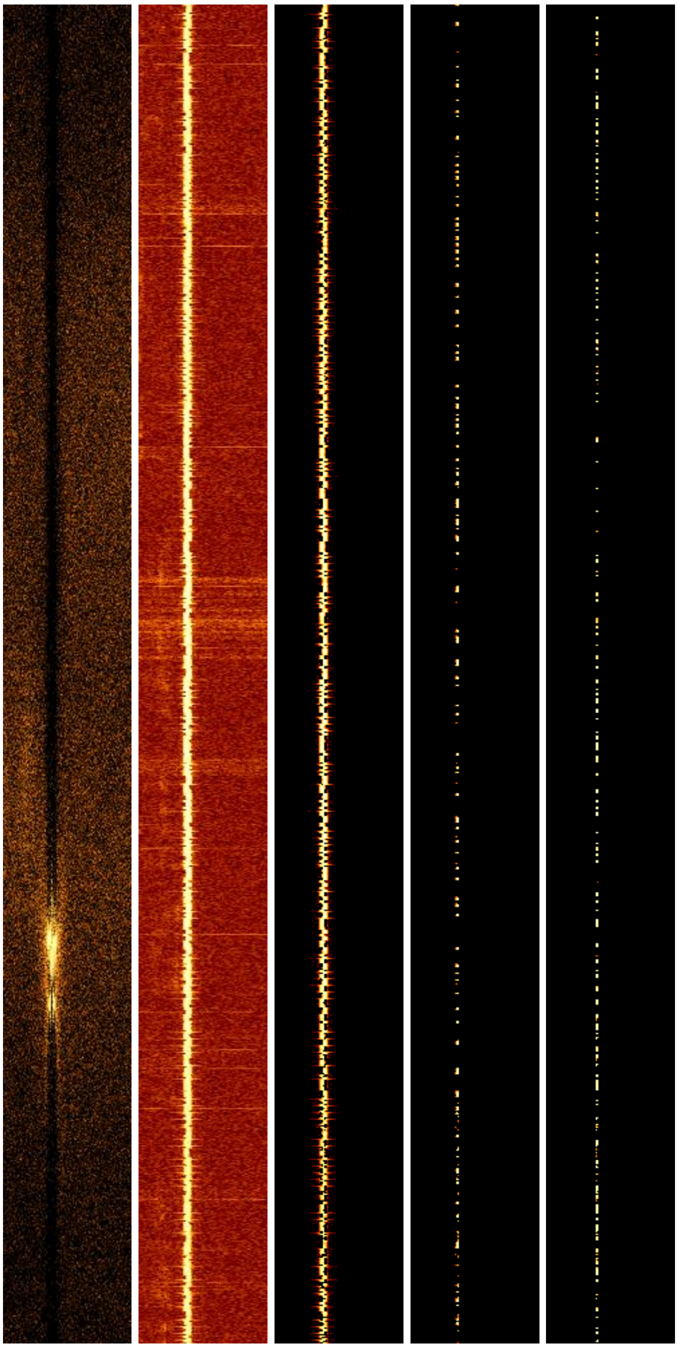

Figure 15: Scenario of two overlapping signals

Complex Scenario

Figure 16 presents a scenario of four overlapping signals. Bandwidth/duration: 18 kHz/8s. The smallest of the azimuthal distances measures 37.3°. From left to right:

1) raw azimuthal distribution (contrast enhanced), azimuth (ordinate) vs. frequency (abscissa)
2) raw spectrogram, time (ordinate) vs. frequency (abscissa)
3) segmentation result #1: STANAG-4285(?), 13.1° σ, 14.2 dB SNR
4) segmentation result #2: Link-11, 3.9° σ, 25.3 dB SNR
   apparently another (faint) Link-11 present at the same azimuth (appears two/three times)
5) segmentation result #3: Link-11 (LSB), 8.4° σ, 21.8 dB SNR
6) segmentation result #4: Carrier(?), 18.8° σ, 16.2 dB SNR

Here a perfect separation of all four overlapping signals can be demonstrated.

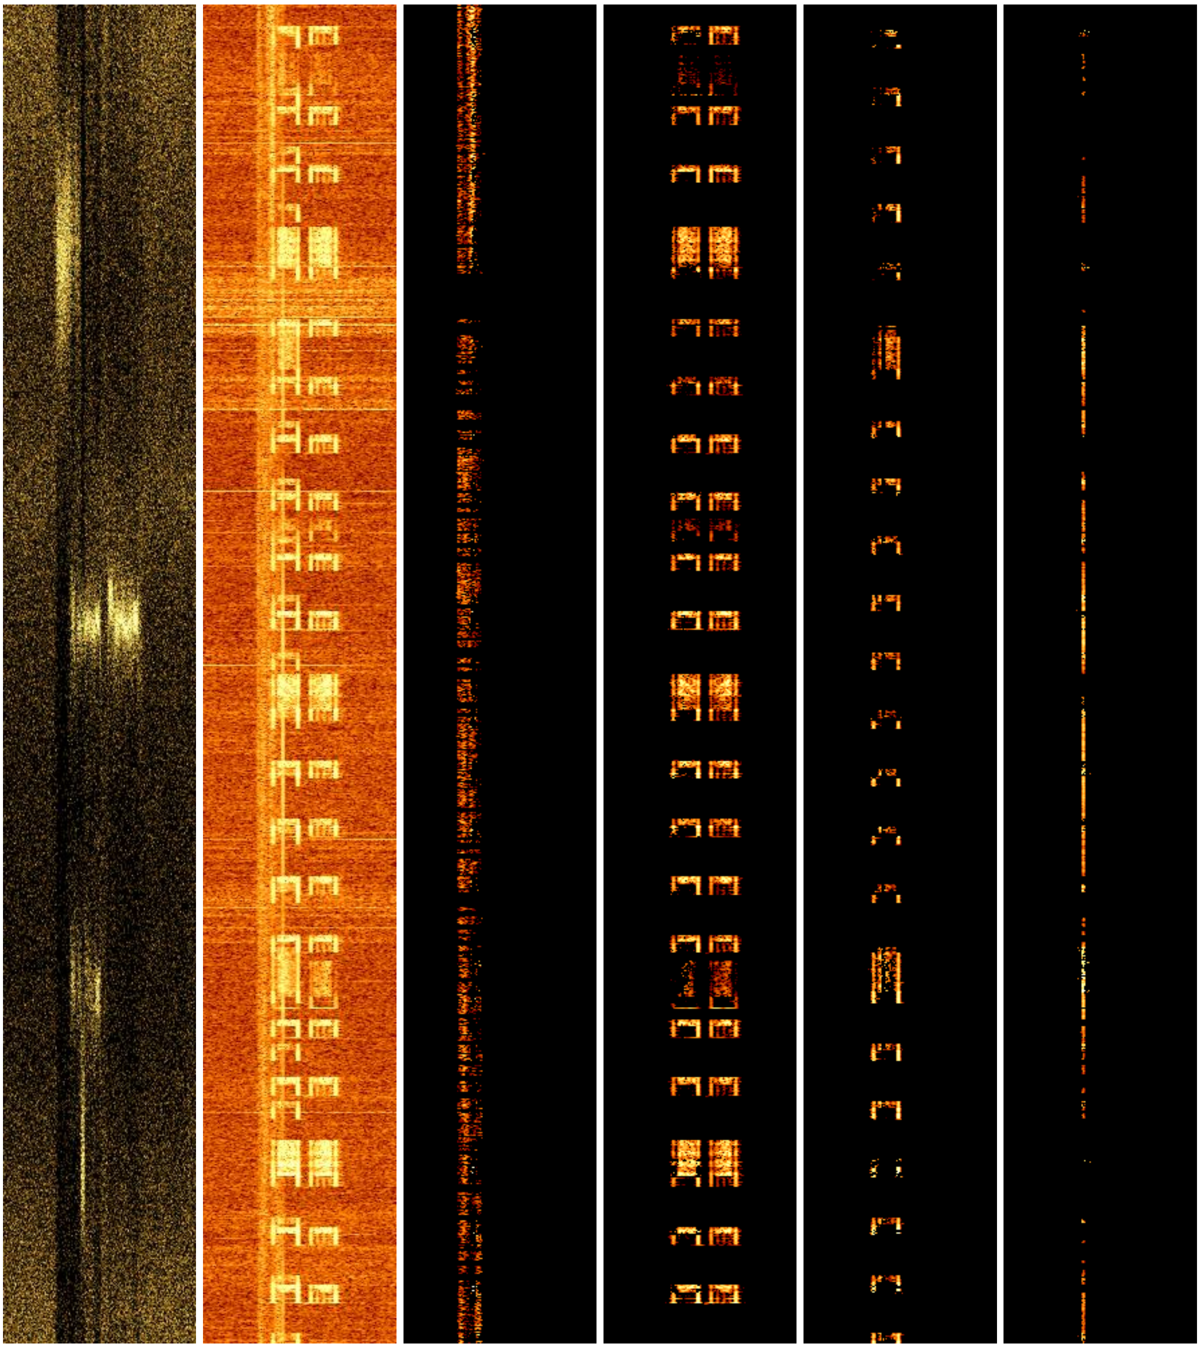

Figure 16: Scenario of four overlapping signals

Scenario near Threshold

Figure 17 presents a scenario of (at least) two overlapping signals. Bandwidth/duration: 18 kHz/8s. The azimuthal distance measures just 7.5°. From left to right:

1) raw azimuthal distribution (contrast enhanced), azimuth (ordinate) vs. frequency (abscissa)
2) raw spectrogram, time (ordinate) vs. frequency (abscissa)
3) segmentation result #1: Link-11, 0.8° σ, 31.8 dB SNR
   apparently another (faint) Link-11 present at the in-between azimuth (appears two times)
4) segmentation result #2: Link-11 (USB), 1.9° σ, 32.6 dB SNR
   apparently parts of the faint Link-11 (appears two times)
5) segmentation result #3: right Carrier of Link-11, 1.4° σ, 29.3 dB SNR
6) segmentation result #4: right Carrier of Link-11, 2.9° σ, 25.9 dB SNR
   separated from #3 due to large time gap

As a conclusion it is demonstrated that nearly perfect separation of the two overlapping signals is possible. Just the right carrier of the fainter signal is segmented as an independent signal. However, this aspect does not impose problems, neither for triggering production nor for correlating to results of other direction finders in order to geo-locate the signal.

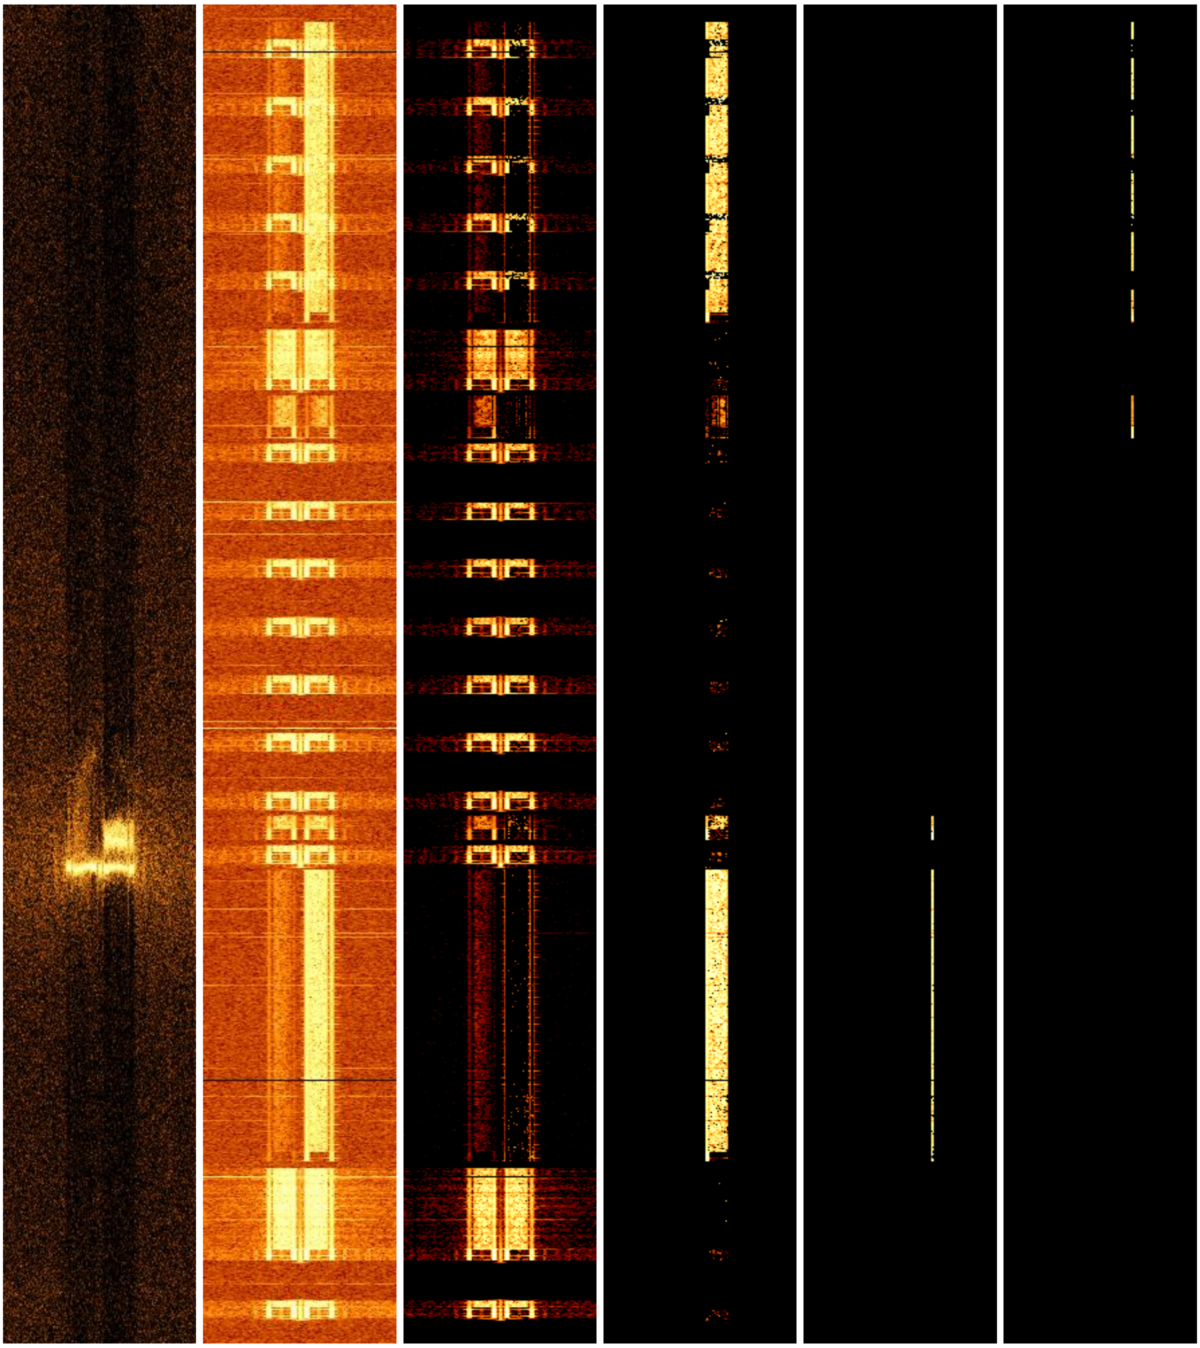

Figure 17: Scenario of (at least) two overlapping signals

4 Summary

The question which algorithm or which high-resolution direction-finding method is superior cannot be answered in an easy manner. First and foremost, the aspect, which needs to be addressed first, is the application in connection with the frequency range. While in the HF frequency range a multiple occupancy of individual frequencies is quite possible, in the V/UHF frequency range Super Resolution might actually only be applicable for airborne applications. Since for all other applications it cannot be determined a-priori whether correlated or non-correlated signals are involved, combinations of the previously discussed procedures could be conceivable in principle. Finally, it must be noted that up to now, with the Super Resolution methods, one computing unit must be provided for each channel that is to be processed. With a small number of channels, the effort is still quite manageable, but it can increase quickly. Therefore, it is not common solution to perform Super Resolution broadband using broadband direction finders. And yet there is no need to relinquish operations with broadband direction finders to secure that possible overlapping signals will be distinguished. Real-world examples above have proven that - thanks to DDA5 - separation of overlapping signals can be achieved on a satisfactory level.