narda Safety Test Solutions

White Paper 010

RF Transmitter Localization

There are five major tasks in the field of radio monitoring:

- Transmission detection
- Transmission analysis
- Transmission classification
- Transmission metadata or even content extraction
- Transmitter localization

This paper only covers transmitter localization.

There are several reasons for localizing transmitters. For example, network providers search for harmful interferers, regulators look for unlicensed transmitters, the military is interested in the location of potential enemies, the police searches for jammers, and intelligence agencies are interested in the location of wireless communication devices used by terrorists or equipment used by eavesdroppers.

The two main methods of localizing transmitters are based on the angle of arrival (AoA) and the time difference of arrival (TDoA) at different receiving locations. The ratios of the powers of arrival may also be used to find the location of a transmitter. Hybrid methods that use more than one property for the localization process also exist. This paper deals with the two main methods.

By: Helmut Keller

Narda Safety Test Solutions GmbH
SUBSIDIARY OF L3HARRIS TECHNOLOGIES, INC.
Sandwiesenstr. 7
72793 Pfullingen, Deutschland
Tel.: +49 7121 9732-0
Fax: +49 7121 9732-790
E-mail: info@narda.com
www.narda-sts.com

1 Localization based on AoA

The angles of arrival at a minimum of two different locations must be known for the localization method based on AoA. A device that can determine the AoA of a signal is called a direction finder (DF). The AoA can be visualized by a bearing line that originates at the location of the DF. The localization process is called triangulation. In the minimal configuration, the bearing lines from two DF locations intersect at the location of the transmitter. The locations of the two DF stations and the transmitter thus form a triangle, which gives the process its name. The best localization accuracy is obtained when the bearing lines intersect under an angle of ninety degrees. When more than two bearings are available, each pair of bearings defines a triangle and so a location can be calculated for each pair. The most likely transmitter location is determined statistically in such cases, particularly when the uncertainties in the bearings and direction finder positions are also properly taken into account.

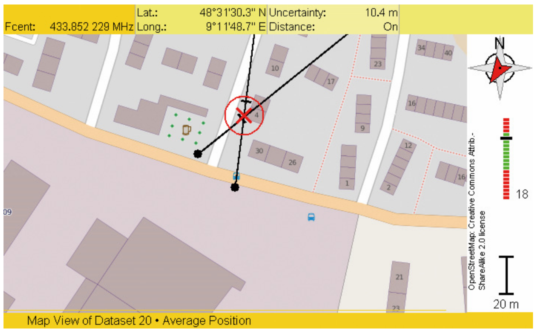

Figure 1. Triangulation using one pair of bearings

1.1 The DF process

Several methods can be used to determine the AoA of a signal. Some of these will be described in this section along with some important aspects of the DF process.

Directional antennas

The first and obvious way to determine the AoA is to use a directional antenna. The power received by a directional antenna depends on the AoA of the signal. The maximum power is received along the main axis. To determine the AoA, the directional antenna is rotated horizontally so that the main axis of the antenna scans the complete azimuth range of 360 degrees, and the received power is plotted versus the azimuth angle. The AoA is the azimuth angle at which the maximum power is received. This horizontal scan can be done manually or automatically.
Two helpful tools make direction finding with directional antennas very convenient with instruments from NARDA
The first tool is an active antenna handle, which can be used with IDA and *SignalShark*. The handle has a built-in electronic compass and therefore always "knows" the direction in which the handle and the attached directional antenna are pointing. The user can find the direction of the maximum received power, for example, by moving the antenna and listening to a tone that has a pitch that is proportional to the received power. When he is certain that the AoA has been found, simply pressing a button on the handle stores the current antenna orientation (azimuth, elevation and polarization) together with the location determined by the GNSS receivers in the IDA or *SignalShark*. This process is called manual bearing.
The second tool supports systematic direction finding by means of a horizontal scan process. The user slowly scans the complete azimuth range with a directional antenna attached to the active antenna handle. Once the scan is finished, the IDA or *SignalShark* automatically calculates the AoA and stores it in a "horizontal scan bearing" data set.

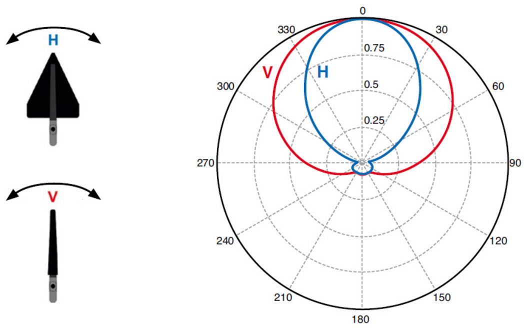

Figure 2. Directional pattern of a logarithmic periodic antenna

Figure 3. Active antenna handle with attached directional antenna

There are also three methods of automatic direction finding with directional antennas.
The first method uses a motorized horizontal scan instead of a manual scan.
The second method uses an array of directional antennas pointing to equally spaced azimuth directions and a rotary switch connects each antenna in turn with a single channel receiver.
These two methods are faster than manual methods, but they all only work well when the transmitted power is constant during the scan.
The third method also uses an array of directional antennas pointing to equally spaced azimuth directions, but each antenna is connected to its own receiver. This method allows reliable bearings to be determined for transmissions with large power variations and even for transmissions of infrequent and short pulses. The receivers only need to be roughly time synchronized for this purpose. The phase synchronization needed for real time antenna array processing, which is discussed later in this paper, is not required.

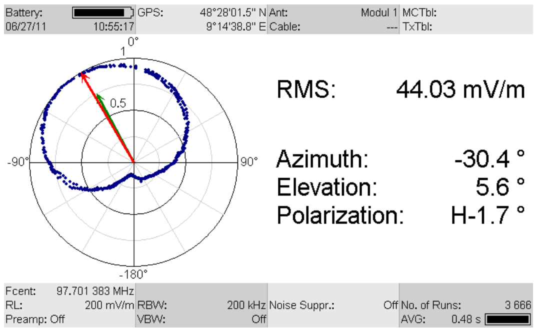

Figure 4. Horizontal scan view of IDA

The Doppler effect

For the DF method based on the Doppler effect, an omnidirectional antenna is moved in a circle periodically. As this periodically changes the distance to the transmitter, the phase of the received signal also changes periodically. The azimuth angle of the omnidirectional antenna position on the circle where the most negative phase shift occurs indicates the AoA of the signal. In practice, a FM demodulator in the connected receiver is often used to detect the phase modulation. Alternatively, a circular antenna array and a rotary switch can be used to emulate the mechanical rotation. This method works well when the transmitted signal is not phase modulated. However, it is not suitable for phase modulated broadband signals.

The real time antenna array processing

Real time antenna array processing requires an antenna array with a dedicated receiver for each antenna element. The receivers must also be phase synchronized, because the DF process requires the determination of the phase difference between the antenna elements. Generally, this phase synchronization requires that the high frequency local oscillators of all the receivers are identical. All the receivers must therefore be fed from the same local oscillator. Such multichannel receivers are rather expensive and are hard to calibrate, but they do offer outstanding DF performance.
One simple, commonly used DF antenna array is a uniform circular array of vertically polarized omnidirectional antennas. Many other array configurations are also used and work well in practice. The number of antenna elements in the array is N.
For a given carrier frequency, the narrow band baseband signal received from a virtual reference element located at the center of the array is a complex scalar $s_0(t)$, which is a function of time. This signal has a variance of $\sigma_{s_0}^2$ and a mean of zero. The output signals of the array build a complex column vector $\mathbf{y}(t)$ of size N. The antenna array itself is

described by its antenna manifold $\mathbf{a}(\theta)$, which is a complex column vector of size N and is a function of the angle of arrival θ. The noise signals of the receivers form a complex column vector $\mathbf{n}(t)$ of size N, and are also functions of time. The noise in the receiver channels is assumed to be statistically independent white Gaussian noise with the same variance σ_n^2 and zero mean in all channels. The following equation is valid if $s_0(t)$ is caused by a single transmitter located in the direction θ_0 and if free space and far field conditions can be assumed:

$$\mathbf{y}(t) = \mathbf{a}(\theta_0) \cdot s_0(t) + \mathbf{n}(t)$$

The M samples of a sequence of time domain signal samples are stored in subsequent columns and build up the row vector $\mathbf{s}_0$ for the signal at the virtual reference element, and the matrices $\mathbf{Y}$ and $\mathbf{N}$ for the received signals and the noise. The following equation is therefore valid for sample sequences:

$$\mathbf{Y} = \mathbf{a}(\theta_0) \cdot \mathbf{s}_0 + \mathbf{N}$$

The antenna array manifold must be known for DF methods based on real time array processing. In practice, the antenna manifold has only been measured for a constrained set of frequencies and directions. The antenna manifold must therefore be interpolated for intermediate frequencies and directions. It is also important to note that the virtual reference element cannot be used during the measurements. The measured manifold is therefore the virtual manifold multiplied by an unknown complex scaling factor.

Beamformer

The beamforming method is based on the idea that a weighted sum of the antenna element output signals can emulate the output signal of a directional antenna pointing in a specific direction. The direction of this virtual directional antenna can be changed by changing the weighting factors. It is therefore possible to scan an azimuth range in different directions and make power measurements for each direction. The direction with maximum power is the estimate for the AoA. This is basically the horizontal scan method already described for directional antennas, the difference being that it uses virtual instead of physical directional antennas, so a finer azimuth grid can be used for the scan. The complete scan can also be executed in real time for each signal sample if desired.
The optimum weighting factors for the antenna element outputs are the complex conjugates of the elements of the antenna array manifold in the desired direction. By normalizing the weightings to the magnitude of the antenna manifold, the output signal $\mathbf{y}_{\mathrm{beam}}(\theta)$ of the beamformer is defined by the following equation where the superscript ^{H} denotes the Hermitian transpose:

$$\mathbf{y}_{\mathrm{beam}}(\theta) = \mathbf{a}^H(\theta)/|\mathbf{a}(\theta)| \cdot \mathbf{Y}$$

The received power $p(\theta)$ of the beamformer averaged over M samples is described by the following equations:

$$p(\theta) = \frac{\mathbf{y}_{\mathrm{beam}}(\theta) \cdot \mathbf{y}_{\mathrm{beam}}^H(\theta)}{M}$$

$$p(\theta) = \frac{|\mathbf{a}^H(\theta)/|\mathbf{a}(\theta)| \cdot (\mathbf{s}_0 \cdot \mathbf{a}(\theta_0) + \mathbf{n})|}{M}$$

$$p(\theta) \cong \sigma_{s_0}^{\mathbf{2}} \cdot |\mathbf{a}^H(\theta)/|\mathbf{a}(\theta)| \cdot \mathbf{a}(\theta_0)|^{\mathbf{2}} + \sigma_n^{\mathbf{2}}$$

The maximum of $p(\theta)$ occurs at $\theta = \theta_0$. It is worth mentioning that for the signal model used so far, the beamformer method is the generalized maximum likelihood estimator of θ_0.

Correlative interferometer

The correlative interferometer method is based on measurement of the current $\mathbf{a}(\theta_0)$ and its correlation with the antenna array manifold $\mathbf{a}(\theta)$. This correlation reaches its maximum when $\theta = \theta_0$.
The current $\mathbf{a}(\theta_0)$ can be measured with the aid of an additional reference antenna element located at the same position as the virtual reference element of the array. Its receiver signal $\mathbf{y}_{\mathrm{ref}}$ is the signal at the reference point plus the noise signal $\mathbf{n}_{\mathrm{ref}}$ in the reference channel of the receiver. This noise signal is also assumed to be statistically independent of the other channels and to be white Gaussian noise with variance σ_n^2 and zero mean.

$$\mathbf{y}_{\mathrm{ref}} = \mathbf{s}_0 + \mathbf{n}_{\mathrm{ref}}$$

The column vector $\mathbf{r}$ contains the covariance values from the N antenna element signals referred to the reference element signal. For the signal model used so far, this gives:

$$\mathbf{r}(\theta_0) = \sigma_{s_0}^2 \cdot \mathbf{a}(\theta_0)$$

An estimate of the covariance vector $\mathbf{r}(\theta_0)$ is given by $\hat{\mathbf{r}}(\theta_0)$:

$$\hat{\mathbf{r}}(\theta_0) = \frac{\mathbf{Y} \cdot \mathbf{y}_{\mathrm{ref}}^H}{M} \cong \sigma_{s_0}^2 \cdot \mathbf{a}(\theta_0)$$

The desired vector is obtained after normalization to the vector magnitudes:

$$\frac{\hat{\mathbf{r}}(\theta_0)}{|\hat{\mathbf{r}}(\theta_0)|} \cong \frac{\mathbf{a}(\theta_0)}{|\mathbf{a}(\theta_0)|}$$

The correlation function $c(\theta)$ to be maximized is therefore given by:

$$c(\theta) = \frac{\mathbf{a}^H(\theta)}{|\mathbf{a}(\theta)|} \cdot \frac{\hat{\mathbf{r}}(\theta_0)}{|\hat{\mathbf{r}}(\theta_0)|} \cong \frac{\mathbf{a}^H(\theta)}{|\mathbf{a}(\theta)|} \cdot \frac{\mathbf{a}(\theta_0)}{|\mathbf{a}(\theta_0)|}$$

A first version of the correlative interferometer maximizes the magnitude of $c(\theta)$:

$$c_{\mathrm{abs}}(\theta) \cong |c(\theta)|$$

Due to the existence of the additional reference element, it can be assumed that the normalized antenna array manifold $\mathbf{a}(\theta)/|\mathbf{a}(\theta)|$ has also been measured by normalized covariance vectors referred to the reference element. There is therefore no phase ambiguity in either the antenna manifold or the covariance vector measured during the DF process. A second version of the correlative interferometer method that requires less computation power also works in this case. It maximizes the real part of $c(\theta)$:

$$c_{\text{real}}(\theta) \cong \Re\big(c(\theta)\big)$$

Note that the best possible correlation value for both methods is unity, due to normalization to the vector magnitudes. A value of unity can only be reached under ideal conditions. The maximum correlation value will be less at low signal to noise ratios, or with multipath propagation. The maximum correlation value multiplied by 100% is therefore used as a DF quality indicator for the bearing view of the *SignalShark*. The *SignalShark* also has a parameter called “min. DF Quality” that can be used to discard bearings with low DF quality.

Comparsion of beamformer and correlative interferometer

The following equation compares the first correlative interferometer method with the beamforming approach:

$$p(\theta) \cong \sigma_{s_0}^2 \cdot |\mathbf{a}(\theta_0)|^2 \cdot c_{\text{abs}}^2(\theta) + \sigma_n^2$$

Clearly, the beamformer method maximizes a scaled and squared $c_{\text{abs}}(\theta)$ plus a constant noise variance, and is therefore identical with the correlative interferometer method. It is worth noting that the additional reference antenna element of the correlative interferometer does not result in better performance. Indeed, if the aforementioned phase ambiguity is accepted it is also possible to estimate $\mathbf{a}(\theta_0)$ without an additional reference antenna element to achieve the same performance.
In practice, different weightings could be used with the beamformer method in order to suppress side lobes. The antenna manifold phases could be used instead of its complex values in the correlative interferometer method. Of course, such modifications can cause differences in the performance of the two methods, but they may not be optimal for the signal model used so far.

Multiple signals

So far, the assumption has been that there is only one signal arriving from a single angle. Now, assuming that there are P signals arriving from P different azimuth angles:

$$\mathbf{y} = \sum_{\mathrm{p}=0}^{\mathrm{P}-1} \mathbf{a}\big(\theta_{\mathrm{p}}\big) \cdot \mathbf{s}_p + \mathbf{n}$$

If the differences between the angles of arrival are much larger than the beam width of the antenna array, the signals can be separated by a beamformer or correlative interferometer. There will be a separate maximum for each signal. If the differences between the angles of arrival are smaller than the beam width of the antenna array, the beamformer or correlative interferometer cannot separate the signals. There will only be a single, biased maximum. In such cases, solution of the DF problem will require the use of one of the so-called “super resolution” algorithms.
Maximum likelihood direction estimation is the optimum method for such problems, but its computational cost is too high for real world applications. More cost-effective computational methods can be used if the P signals are not coherent. One of these methods is the multiple signal classification (MUSIC) algorithm. Unfortunately, multipath signals are in

general coherent and cannot be resolved using such methods. Nevertheless, a brief description of the MUSIC algorithm is given below.

MUSIC algorithm

The MUSIC algorithm is based on the covariance matrix $\mathbf{R}_y$ of the received signals. For non-coherent signals, for the signal model considered so far and the N by N eye matrix $\mathbf{I}$:

$$\mathbf{R}_y = \sum_{p=0}^{P-1} \sigma_{s_p}^2 \cdot \mathbf{a}(\theta_p) \cdot \mathbf{a}^H(\theta_p) + \sigma_n^2 \cdot \mathbf{I}$$

$\hat{\mathbf{R}}_y$ is an estimate of this covariance matrix:

$$\hat{\mathbf{R}}_y = \frac{\mathbf{Y} \cdot \mathbf{Y}^H}{M} \cong \mathbf{R}_y$$

The first step in the MUSIC algorithm calculates the covariance matrix $\hat{\mathbf{R}}_y$ and its eigenvalues and eigenvectors. The eigenvectors with eigenvalues significantly lower than the largest eigenvalues are considered not to belong to signals but to the noise. They encompass the so-called noise space. The remaining eigenvectors cover the so-called signal space.
In the second step, the antenna manifold is correlated with each of the noise eigenvectors, and the squared magnitudes of all correlations are summed. The minima in this sum are assumed to occur at the angles of arrival of the signals, because it can be assumed that the signal space is orthogonal to the noise space.
It is worth noting that the number of signal samples M necessary in order to estimate the covariance matrix is at least three times greater than the number of antenna elements N. In practice, much higher numbers may be necessary in order to separate the signals and estimate their angles of arrival. The maximum number of signals that can be separated is $N-1$.

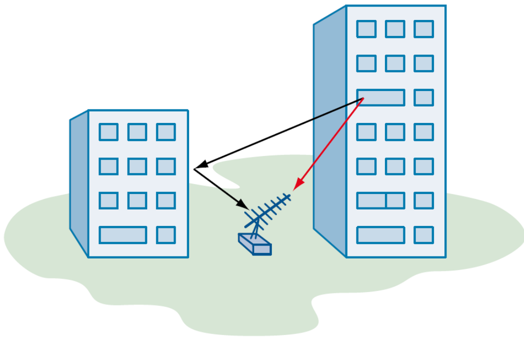

Figure 5. Multipath propagation from a transmitter in the taller building to the receiving antenna on the floor

Real time antenna array processing versus directional antenna array

A main feature of DF methods based on real time antenna array processing is their speed. At high signal to noise ratios, the beamformer or correlative interferometer can estimate the AoA from a single sample of the base band signals of the array, so the DF rate can be as high as the base band signal sample rate. Thus, even extremely short signal transmissions can be detected and their angles of arrival determined. As already mentioned, an array of directional antennas, each with a dedicated receiver, is also fast enough to solve this problem without needing phase synchronization of all the local oscillators in all receivers. Time synchronization accuracy in the order of the base band sampling period is sufficient. Conversely, an antenna array with real time antenna array processing may be easier to construct than an array of directional antennas, and fewer channels may be needed. Real time antenna array processing is also more flexible. It can even use super resolution algorithms if appropriate and desired.

Large aperture antenna arrays

The aperture of an antenna array is the maximum array dimension. DF methods based on real time antenna array processing often use arrays with an aperture in the order of one wavelength. Increasing the ratio of the aperture to the signal wavelength also increases the bearing accuracy and the ability to separate signals from different directions. This is true until phase ambiguities and thus also ambiguities in the estimated AoA occur.

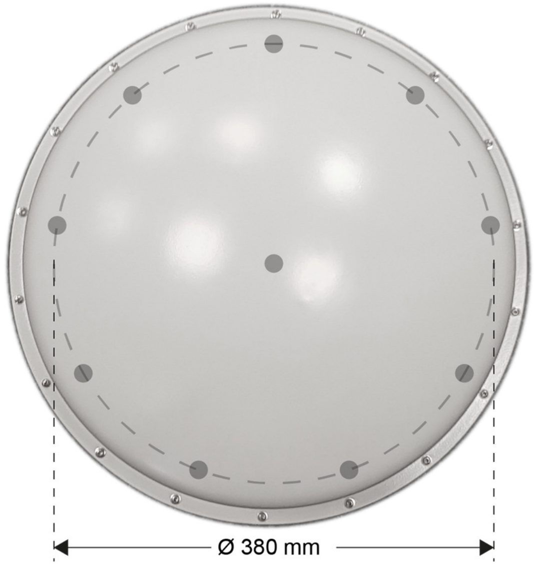

Figure 6. Large aperture uniform circular antenna array with 9 elements and an additional reference element in the center as used in the ADFAs for frequencies from 200 MHz to 2.7 GHz

Multi-path propagation immunity and upper frequency limit

The upper frequency limit of the array is defined by the aforementioned AoA ambiguities and the multipath signal immunity. Increasing the number of antenna elements increases the upper frequency limit if the array aperture is kept constant.

A comparison of uniform circular arrays comprised of even and odd numbers of antenna elements shows that odd-numbered arrays have much higher usable frequency limits. NARDA therefore uses uniform circular arrays with nine elements. Their upper frequency limit is much higher than arrays of eight, ten or even twelve elements, and is reached at a diameter to wavelength ratio of about 3.6. The upper frequency limits of uniform five- and seven-element circular antenna arrays are much lower and are reached at diameter to wavelength ratios of about 1.1 and 2.4, respectively.

NARDA has determined these upper frequency limits by numerical simulation of the multipath propagation immunity of uniform circular arrays. The signal model used in the simulation assumes noise-free receivers and two coherent signals with an arbitrary phase shift, one signal being half the amplitude of the other. The azimuth angle of both signals was selected randomly between zero and 360° and the elevation angle was selected randomly between –15° and +15°. The RMS uncertainty in the estimated angle of arrival of the stronger signal was calculated for 9,216 signal parameter combinations. The RMS uncertainty due to multipath reception below the upper frequency limit was approximately 3.5° multiplied by the ratio of the array diameter to the wavelength. This uncertainty is independent of the number of antenna elements and decreases in proportion to the wavelength of the signals. At a diameter to wavelength ratio of 3.6 the RMS uncertainty due to multipath reception is only 0.97°. Above the upper frequency limit however, the uncertainty can rise dramatically and can easily reach values of more than 10°.

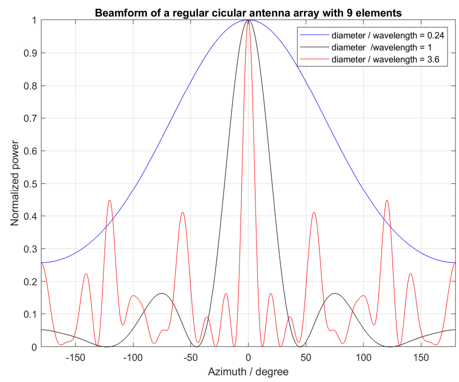

Figure 7. Beamform of a uniform circular antenna array with 9 elements for diameter to wavelength ratios of 0.24, 1, and 3.6

Compact antenna arrays for low frequencies

The size of a large aperture antenna array is proportional to the wavelength of the signals if constant DF accuracy and beam width is desired. For this reason, compact direction finders for low frequencies use antenna arrays that are based on the directional pattern of elementary dipoles rather than on phase differences between the elements.

Watson-Watt

The popular Watson-Watt method uses an array of two orthogonal magnetic dipoles (coils) in the horizontal plane and one vertically polarized electric dipole. The antenna array manifold $\mathbf{a}(\theta)$ for vertically polarized transmitters in the horizontal plane is:

$$\mathbf{a}(\theta) = \begin{bmatrix} \cos(\theta) \\ \sin(\theta) \\ 1 \end{bmatrix}$$

The historic Watson-Watt method used an oscilloscope as the receiver. The first coil was connected to the x-axis, the second coil was connected to the y-axis and the sign of the signal of the electric dipole was used to blank the beam while the electric dipole was receiving negative signals. The visual result on the oscilloscope screen was a bearing line pointing into the direction of the transmitter.
The same antenna array can also be used with a modern beamformer. For the received power $p(\theta)$:

$$p(\theta) \cong \sigma_{s_0}^2 \cdot \frac{|1+\cos(\theta-\theta_0)|^2}{2} + \sigma_n^2$$

This shows that the principle of the Watson-Watt method is defined more by the antenna array than by the processing of the three antenna signals. The multipath propagation immunity of a Watson-Watt array does not depend on the signal frequency. A numerical simulation undertaken by NARDA assumed noise free receivers, vertical polarization for the direct and reflected signal, zero elevation, an amplitude ratio of the direct to the reflected signal of two, and random phase and AoA of the reflected signal. The simulation result for the RMS bearing uncertainty due to the reflected signal was 14.85°.
A uniform circular antenna array reaches the same bearing uncertainty due to reflections at a diameter to wavelength ratio of about 0.24. This means that a circular array with a higher diameter to wavelength ratio is less immune to multipath reception than a Watson-Watt array.
Unfortunately the Watson-Watt method indicates angles of arrival shifted by 90° from the real direction for horizontally polarized transmitters that are elevated above or below the horizontal plane.

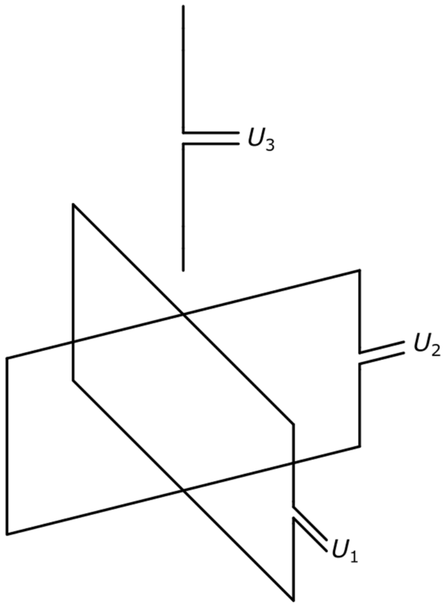

Figure 8. Watson-Watt antenna array

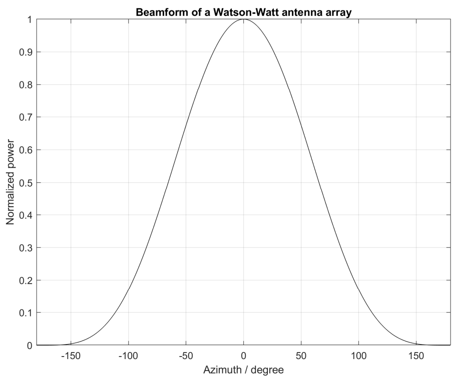

Figure 9. Beamform of a Watson-Watt antenna array

Adcock

An Adcock antenna array can be used instead of the Watson-Watt array to overcome the problem of incorrect bearings for elevated transmitters with horizontal polarization. A pair of vertically polarized dipoles is used instead of each coil. The two dipoles of a pair are arranged on opposite sides of a circle around the center of the array. The only output from each dipole pair is the difference signal. The second dipole pair is arranged orthogonally to the first pair. A dipole in the center can still be used or it can be substituted by a third output that carries the sum signal of all four dipoles on the circle. The antenna array manifold for vertically polarized transmitters in the horizontal plane is in principal the same as for the Watson-Watt array, but the Adcock antenna array is insensitive to horizontally polarized signals and thus avoids giving incorrect bearings for elevated horizontally polarized transmitters. The disadvantages of the

Adcock compared to the Watson-Watt antenna array are that it is larger for low frequencies and that its useful relative frequency range is limited to a few octaves.

Poynting vector

The optimum antenna array for compact direction finders at low frequencies consists of three orthogonal magnetic dipoles and three orthogonal electric dipoles, all located in the same position. The Watson-Watt array is obviously just a subarray of this extended array that allows measurement of the magnitude and direction of the Poynting vector for both azimuth and elevation. This array works well for any combination of azimuth, elevation, and polarization.

Automatic DF with the SignalShark

The automatic DF antennas (ADFA) for the *SignalShark* contain nine-element vertically polarized uniform circular arrays for frequencies above 200 MHz. There is also an additional, central, vertically polarized reference antenna element, and a Watson-Watt array for frequencies below 200 MHz.

Single channel correlative interferometer

The DF method used in conjunction with the circular arrays is the correlative interferometer. However, the *SignalShark* is a single channel receiver. A method of measuring the covariance vector $\hat{\mathbf{r}}$ of the antenna array signals using a single channel receiver is therefore needed.
For a single element $\hat{r}(n)$ of the covariance vector $\hat{\mathbf{r}}$ it can be proven that:

$$\hat{r}(n) = \left|\frac{y_{ref+}y(n)}{2}\right|^2 - \left|\frac{y_{ref-}y(n)}{2}\right|^2 + j\cdot\left(\left|\frac{y_{ref+}y(n)\cdot \mathrm{e}^{\frac{-\mathrm{j}\cdot\pi}{2}}}{2}\right|^2 - \left|\frac{y_{ref+}y(n)\cdot \mathrm{e}^{\frac{\mathrm{j}\cdot\pi}{2}}}{2}\right|^2\right)$$

$$\hat{r}(n) = \hat{r}_0(n) - \hat{r}_1(n) + \mathrm{j}\cdot\left(\hat{r}_2(n) - \hat{r}_3(n)\right)$$

Each element of the covariance vector can therefore be calculated by using the results of four power measurements. For each power measurement, the power of the sum of the signal of the reference element and a phase shifted version of the signal of the selected antenna element must be measured. The phase shifts necessary for the power measurements are: no phase shift at all for the first, signal inversion for the second, a –90° phase shift for the third, and a +90° phase shift for the fourth. $4 \cdot N$ power measurements are necessary for the complete covariance vector. The power measurements do not need to be done concurrently, but can be done sequentially under the following three conditions:

1. The average power of the signal $\mathbf{s}_0$ at the reference element is the same during all power measurements, which are executed for the determination of a single covariance vector.
2. The signal to noise ratio is sufficiently high during all power measurements, which are executed for the determination of a single covariance vector.

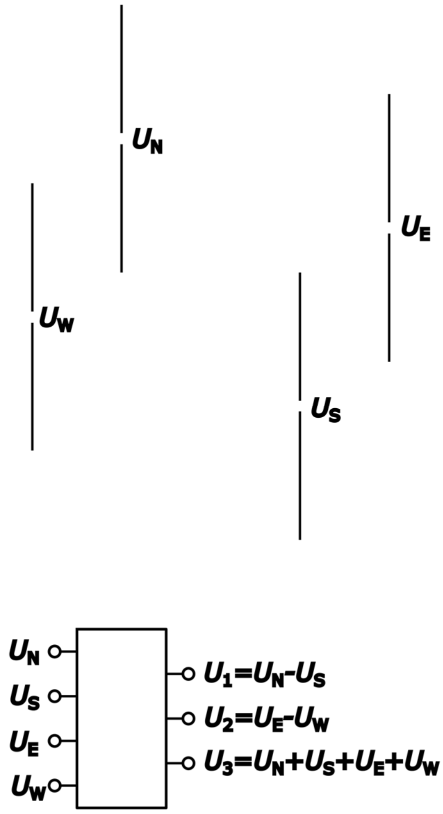

Figure 10. Adcock antenna array

3. The angle of arrival does not change during all power measurements, which are executed for the determination of a single covariance vector.

If these conditions are met, the desired covariance vector can be measured using a single channel receiver, some switches, phase shifters, and an antenna array.

The ADFAs contain the necessary multiplexers, switches and phase shifters. The switch states are synchronized with the power measurements by the *SignalShark*. There is a pause in measurement between each switch state change and the start of a new power measurement to allow the *SignalShark* channel filters to settle. The omnidirectional power and spectrum of the reference antenna element are also measured at the end of each bearing cycle. There are 37 switch states in a complete bearing cycle. A measurement time of less than the necessary settling time cannot be set on the *SignalShark* because such a setting would not be reasonable. This means that the shortest possible bearing cycle time is 74 times the required settling time. The minimum bearing cycle time is also 1.2 ms. This short cycle time can only be achieved for channel bandwidths that are greater than or equal to 3 MHz. For lower channel bandwidths the required settling time is the limiting factor.

NARDA ADFAs are calibrated for the complete azimuth range and for an elevation range of -20° to +40°. The reference element is raised slightly above the horizontal plane of the circular array, which allows positive and negative elevations to be distinguished. The *SignalShark* executes a two-dimensional search for the DF process. It correlates the measured covariance vector with all the covariance vectors determined during the calibration process. The azimuth and elevation value of the point with the maximum correlation gives a rough estimate of the desired AoA pair. Subsequent two-dimensional quadratic interpolation uses the 8 adjacent correlation values as well to give the precise values for the azimuth and elevation of arrival.

Some competing DF systems are calibrated and tested in the horizontal plane only, and are also incapable of determining the elevation of arrival. Measurement of the elevation of arrival is the first obvious advantage of NARDA's extensive calibration and DF process. However, this is not the main advantage, because the accuracy of the azimuth of arrival is about five times better than that of the elevation of arrival. The main advantage is that the accuracy of the azimuth of arrival stays more or less constant when the elevation of arrival changes within the calibrated range. The accuracy specifications of competitors are valid for zero elevation only, and it must be expected that the azimuth accuracy decreases considerably when the elevation differs significantly from zero.

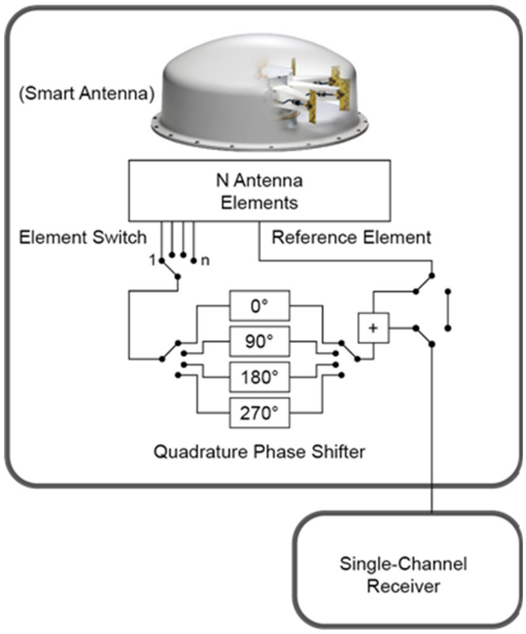

Figure 11. Single channel correlative interferometer

Single channel Watson-Watt

The *SignalShark* uses the Watson-Watt antenna array in the ADFAs for frequencies below 200 MHz. Signal processing is basically the same as for the circular arrays, with some minor changes. The three-element array is treated as a two-element array with an additional reference antenna element. The vertically-polarized omnidirectional antenna is used as the reference antenna element and the two coils are used as the two-element

antenna array. The covariance vector $\hat{\mathbf{r}}$ thus has only two elements. Since the imaginary parts of the covariance vector elements do not contain useful information, they are not calculated and set to zero. This also means that there is no need to measure $\hat{r}_2(n)$ and $\hat{r}_3(n)$. Each element of the covariance vector is therefore measured in two steps only. There are only five steps in a complete bearing cycle that includes measurement of the omnidirectional power and spectrum of the reference signal. A bearing cycle time of 1.2 ms can therefore already be achieved at a channel bandwidth of 400 kHz.
The Watson-Watt antenna array is only calibrated in the horizontal plane because the transmitter elevation cannot be determined by this method. However, the complete Watson-Watt DF system is also tested for the elevation range of –20° to +40°.

Optimum DF parameter settings for the *SignalShark*

Before starting any DF process, the part of the spectrum occupied by the suspicious transmitter must be determined. An initial spectrum measurement is useful for this task. The *SignalShark* uses a channel filter to separate the signal under investigation from other signals during the DF process. The center frequency of the channel filter is set by the Ftune parameter, and its bandwidth is set by the CBW parameter. The CBW denotes the 6-dB bandwidth of a digital Parks-McClellan filter. The passband is flat with a maximum deviation of 0.001 dB, and occupies the middle 84% of the CBW. The stopband attenuation is greater than 85 dB and is reached outside 116% of the CBW. The channel filter must be positioned such that the ratio of the received signal power to the received power of the noise floor or adjacent signals is maximized. A setting very close to the optimum is achieved by setting Ftune to the center frequency Fmid of the occupied spectrum and CBW to the occupied bandwidth OBW. The occupied bandwidth OBW and the center frequency Fmid of the signal under investigation can conveniently be measured with the *SignalShark* using the “occupied bandwidth” spectrum marker function. The OBW definition should be set to 99% of the total power. The optimum Ftune and CBW settings can of course also be estimated by visual inspection of the spectrum.
The attenuator of the *SignalShark* should be set to the lowest value possible for no overload to occur. This is often the 0 dB setting, due to the huge dynamic range of the *SignalShark*. The omnidirectional spectrum can also be observed on the *SignalShark* during the DF process. This feature is useful to detect any RF frontend overload during the DF process. The *SignalShark* automatically detects any overload of the IF digitizer and displays an overload indicator.
The three necessary conditions already mentioned for the single channel approach must be kept in mind for the remaining parameter settings. Clearly, signals that are received at constant and sufficient power and with a constant AoA fulfill all three conditions. Care must be taken to ensure that all three conditions are met if the received power or AoA varies with time.
The measurement time spent on each power measurement in a bearing cycle is an adjustable parameter on the *SignalShark*. Users can change the bearing cycle time indirectly and minimize the negative effects of a

variable signal power or AoA by means of a proper setting of the measurement time.
If the received power varies periodically over time, the measurement time should be set to the period duration or an integer multiple of it. This ensures that the average power received during each measurement time is constant, and the first condition is met exactly.
If the received signal contains a frame structure, the measurement time should be set to the frame length or an integer multiple of it. The average power received during each measurement time will then be more or less constant, and the first condition will be met approximately. All modern mobile communications networks are frame based. The GSM frame length is 4.6154 ms. The frame length for UMTS, LTE, and 5G NR is 10 ms.
If the received signal power varies randomly, the measurement time should be set as high as is feasible when other factors are considered. The average power received during each measurement time will then be more or less constant, and the first condition will be met approximately.
The first and second conditions are very hard to meet if the signal is transmitted in bursts. The bearings should only be taken during the burst period. The *SignalShark* can use omnidirectional power measurements at the end of each bearing cycle to achieve automatic burst detection. The “DF Squelch” parameter determines whether the actual bearing cycle is used for the AoA calculation. The covariance vector of the actual bearing cycle is calculated and used for the AoA calculation only if the omnidirectional power of the previous and the actual bearing cycle are greater than the squelch value. The “min. Stability” parameter exists for the same purpose. The covariance vector of the actual bearing cycle is calculated and used for the AoA calculation only if the magnitude of the level difference between the omnidirectional power of the previous and the actual bearing cycle is less than the specified stability value. For DF of burst transmissions, the DF squelch should be set to a value of at least 10 dB above the noise floor, and the required level stability value should be set to about 1 dB. Using these settings it can be proven that *SignalShark* only uses bearing cycles for DF that are entirely within a burst transmission when the following two conditions are met:

- the minimum burst length is greater than two bearing cycle times plus one measurement time
- the minimum pause length between bursts is greater than one bearing cycle time plus one measurement time

The *SignalShark* automatically discards bearing cycles that contain level transitions or transmission pauses under these conditions.
The desired bearing cycle times for signals with short minimum burst and pause lengths will also be very short. This means that the measurement time will also be very short, so the bearing rate will probably be much higher than necessary. The short measurement time may result in less accurate bearings, because the power fluctuation due to the signal modulation and the additive noise may not be sufficiently reduced. The *SignalShark* can overcome this problem by averaging a number of covariance vectors before the AoA is calculated. This is accomplished using the post averaging time parameter, which can be set to integer multiples of four times the bearing cycle time. A new bearing is calculated

every quarter of the selected post averaging time, using the average of all valid covariance vectors that are not older than the post averaging time. Good bearing accuracy can be achieved even with very short bearing cycle times due to this additional averaging of covariance vectors.
If the direction finder or the transmitter are moving, the third condition can be met if the bearing cycle time is set short enough so that the AoA does not change significantly during a bearing cycle.

DF uncertainty, sensitivity and immunity against multi-path propagation

The DF uncertainty is influenced by theoretical and practical restrictions. The Cramér-Rao bound (CRB) is the lowest theoretical bound for the variance of any AoA estimation algorithm, assuming the signal model used so far for the single transmitter situation. It is assumed that the antenna manifold is known perfectly and therefore that an unbiased estimate can be expected from any reasonable estimation algorithm. The CRB in squared radians is given by:

$$CRB = \frac{1}{2 \cdot M \cdot SNR \cdot |\dot{\mathbf{a}}(\theta)|^2}$$

Note that $\dot{\mathbf{a}}(\theta)$ is the derivative of the antenna manifold with respect to θ in radians and SNR is the ratio of the signal to the noise power at the receiver inputs.
For a uniform circular antenna array with a diameter d and for a signal wavelength λ:

$$CRB = \frac{\lambda^2}{M \cdot SNR \cdot N \cdot \pi^2 \cdot d^2}$$

For the Watson-Watt antenna array:

$$CRB = \frac{1}{2 \cdot M \cdot SNR}$$

Assuming that SNR and M are constant and that there are nine antenna elements in the circular array, a comparison of the two arrays reveals that the Watson-Watt array performs better for diameter to wavelength ratios lower than about 0.15. The circular array performs better for higher ratios, and its accuracy increases with the signal frequency. Note that for a direction finder with a single channel receiver, M is determined by the measurement time of a single power measurement. If the bearing cycle time of a single channel direction finder is kept constant, the break-even for array performance occurs at a diameter to wavelength ratio of about 0.41.
The CRB decreases with the number of samples M, so a longer measurement time will increase the accuracy of the bearings. This is true as long as the impact of other factors that also affect the accuracy remains negligible compared to the impact of the noise of the receiver. For low SNR values of the order of unity, the CRB cannot be reached even by optimum estimation algorithms. Most algorithms get very close to the CRB for SNR values greater than about 4.
The CRB does not describe the accuracy of real-world direction finders for high values of SNR and M. Another aspect dominates the DF uncertainty in this case:

The antenna manifold is not known perfectly in practice, because the accuracy of the calibration process is limited and the antenna manifold may also change with time and due to environmental influences. If only a few devices for one type of DF system are calibrated and an averaged antenna manifold is used for all devices for this type of DF system, the manufacturing tolerances will also impact DF uncertainty. This means that the AoA estimate of a direction finder will normally be biased for a given frequency and AoA. The DF uncertainty specified in direction finder data sheets is usually the RMS value of these biases averaged over the complete azimuth and frequency range. State-of-the-art mobile direction finders are likely to have a typical DF uncertainty of 1° RMS in the frequency range above 200 MHz. ITU-R Recommendation SM.2060-0 defines a measurement procedure for the verification of the DF uncertainty.

The DF sensitivity defines a minimum field strength value for which the RMS value of the DF uncertainty due to the receiver noise remains below a predefined threshold. ITU-R Recommendation SM.2096-0 defines a threshold of 3°, a total integration time of 1 s, and a CBW of 1 kHz for the DF sensitivity specification. The DF sensitivity is in general a function of frequency, and is specified in the direction finder data sheets.

DF sensitivity is influenced primarily by the antenna factors of the antennas and the effective noise floor at the receiver inputs. Note that the sensitivity may be achieved at SNR values smaller than unity. An estimation of the DF sensitivity based on the CRB may therefore result in lower than realistic DF sensitivity values.

In practice, the RF environment also has a huge impact on the DF uncertainty. Immunity against multipath propagation has already been covered by the introduction of various types of antenna arrays. It has been shown that the diameter to wavelength ratio of a circular array determines its immunity and that the immunity of a Watson-Watt antenna array is not dependent on frequency or on any design parameter. It is worth noting that the DF uncertainty due to multipath reception decreases proportionally with the ratio of the amplitude of the reflected signal to the amplitude of the direct signal. ITU-R Recommendation SM.2061-0 defines a measurement procedure for the verification of multipath propagation immunity.

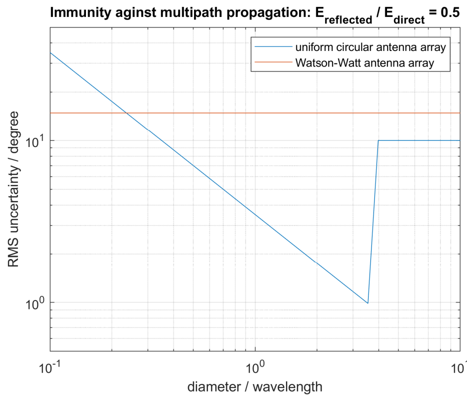

Figure 12. Immunity against multipath propagation for a uniform circular antenna array with nine elements and for a Watson-Watt antenna array

Direction finder position and north reference

The position and orientation of the direction finder must be known for a bearing from it to be meaningful. In practice, the geolocation of the direction finder is determined with a GNSS (global navigation satellite system) receiver located in the antenna array or the receiver. The GNSS position uncertainty is about 15 m RMS and thus precise enough for most applications.

The antenna array orientation is not so easy to determine. All DF antenna arrays have a reference direction, which is also marked on the array construction. Only AoA values relative to this reference direction are used during calibration of the antenna array. The antenna array is normally mounted on a rotary table during calibration. The relative azimuth value of zero is reached when the test transmitter is exactly aligned with the reference direction of the antenna array. The rotary table is zeroed very precisely with the help of optical tools. The table is rotated by the desired angle to set azimuth angles relative to the reference direction. The relative angle uncertainty of the rotary table is normally much better than the bearing uncertainty of the array and thus does not contribute significantly to the calibration uncertainty. When the DF system is used in practice, it indicates the AoA of the suspicious transmitter relative to its reference direction. However, it is desirable to determine the AoA relative to geographic north for localization purposes. There are several ways to achieve this. The *SignalShark* has three “North Reference” settings for bearings: “Ref. Mark Dir.”, “Compass”, and “GNSS velocity”.

Use of the reference mark direction is the most precise method. This is somewhat time-consuming, so it should be employed only if the DF antenna is to be used at a fixed location for at least some hours. The DF antenna reference direction is marked roughly by an arrow and precisely by optical bearing marks. The antenna should be adjusted so that the bearing marks are aligned with a landmark. The azimuth angle of this landmark relative to geographic north and the location of the DF antenna must be determined from the geographic positions of the landmark and the DF antenna. The “Ref Mark Dir.” parameter is set to the azimuth angle of the landmark. The AoA indicated by the *SignalShark* is the sum of the “Ref. Mark. Dir.” and the internal AoA value, which is relative to the antenna reference direction.

The extensive adjustment procedure may make the “Ref. Mark Dir.” method too time-consuming when the DF antenna is only stationary during the DF process but otherwise changes its position frequently. All NARDA DF antennas have a built-in electronic compass with an azimuth uncertainty of typically 1.5° RMS when the DF antenna is located in the undisturbed magnetic field of the earth. This can be used as the north reference if the DF antenna is mounted such that no ferromagnetic materials and no DC currents are present in its vicinity. A good example of this is to mount the antenna on a wooden tripod. It is important that the correct declination of the earth’s magnetic field at the current location is also entered when the “Compass” setting is used for the north reference.

If the DF antenna is mounted on a vehicle, the earth’s magnetic field in the vicinity will normally be heavily disturbed, especially if mounting is by means of a magnetic plate, which is very convenient and often used. If the vehicle is also moving during the DF process, as is often the case, only one option remains for the north reference: “GNSS velocity”. This method assumes that the DF antenna reference direction is the same as the movement direction of the vehicle. The position of the DF antenna on the vehicle must therefore be carefully adjusted so that the antenna reference direction is parallel to the vehicle’s normal direction of travel. Note that the DF antenna bearing marks may also be helpful for this adjustment.

The accuracy of the movement direction measured by the DF antenna GNSS receiver is proportional to the magnitude of the vehicle velocity. The typical uncertainty is about 0.3° RMS at a velocity of 100 km/h. This method therefore works only well when the vehicle is moving faster than about 10 km/h. The GNSS receivers in NARDA DF antennas will try to recall the last valid direction estimate even when the vehicle stops. However, this direction is only a useful estimate of the true direction if the

vehicle did not change direction considerably while moving very slowly. Bearings taken when the vehicle is parked are thus only useful if the vehicle did not change direction abruptly during the period of slow movement before reaching the parked position. If this is not the case, the north reference will need to be changed to “Compass” when the vehicle is parked. The DF antenna must then also be unmounted from the vehicle if low bearing uncertainty is required in the fixed position.
If there is a known reference direction mismatch, it can be corrected in every “North Reference” settings by a parameter called “Azimuth Corr.”. The value of this parameter is just added to the uncorrected azimuth result.
It is worth noting that a perfect north reference is always assumed for the specification of the DF uncertainty. The uncertainty in the north reference is an additional source of uncertainty that must be taken into account when calculating the overall DF accuracy.

1.2 The localization process

The suspicious transmitter can be localized once the bearing lines from two known and significantly different direction finder locations are known and intersect. Localization accuracy increases with the number of direction finder locations and the number of bearings. It is therefore usual to collect bearings from more than two direction finder locations before the localization process is completed.
If only one direction finder is available, it will have to be moved to different locations over time. One way that this can be done is to move to different locations one by one and to take bearings with the direction finder stationary at each location. This procedure is time-consuming, but can be very accurate when the ideal fixed positions are selected. Such positions have a line of sight to the suspicious transmitter and are more or less equally distributed on a close circle around the suspicious transmitter, and there are no reflectors or obstacles close to the direction finder. Since the location of the suspicious transmitter is not known, it is difficult to select the direction finder positions in advance. Selection of the next location will often depend on the previous results.
Another approach is to take the bearings while the direction finder is mounted on a moving vehicle. The advantage of this method is that bearings can be collected from many positions within a short time span. The disadvantage is that it is very likely that most of the bearings will be more or less random because they are taken from positions without line of sight. The localization algorithm that is used must therefore also work reliably even under such conditions. This localization method is often referred to as homing-in.
The third approach requires one direction finder at each location. The advantages of this method are that it is fast, and that it can even localize moving targets. The disadvantages are that it is hard to find the right locations when dealing with arbitrary targets, and that a network connection is required in order to exchange the bearing data.

The SignalShark map view

The *SignalShark* provides a map view that can be used to visualize data on a geographic map, which can be imported in the form of so-called slippy map tiles. NARDA provides a download tool for the slippy map tiles rendered by OpenStreetMap© and other public tile servers. The *SignalShark* also accepts slippy map tiles rendered by any geographic information system tool. Military or other official maps can also be imported into the *SignalShark* in this way. The data to be visualized are overlaid on the geographic map. General data sets are represented by simple pins. Power levels are represented by color coded pins. Bearings are shown as pins with bearing lines. Localization results from different transmitters are shown as crosses and uncertainty ellipses in different colors. The map view is also used to provide a localization heat map and to display the current estimate of the transmitter position during the localization process for a single transmitter.

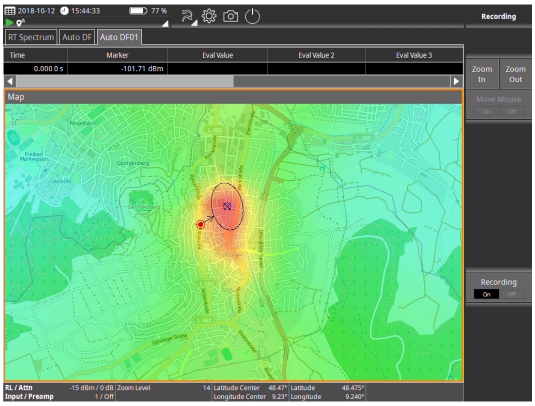

Figure 13. SignalShark's localization algorithm in action

The SignalShark localization algorithm

The *SignalShark* uses an innovative maximum likelihood algorithm for the localization of transmitters. The algorithm assumes a certain probability for line of sight situations. It is assumed that, in these line of sight situations, the probability density function (PDF) of the bearings is Gaussian, with a mean value of the true AoA and an assumed standard deviation. Equal distribution over the complete azimuth range is assumed for non-line of sight situations. In this context it should be noted that a line of sight situation does not necessarily mean that there actually is a line of sight between the direction finder and the transmitter. It is sufficient that the bearings in these situations have a PDF that roughly approximates to the assumed PDF. In other words, a line of sight situation is one where the bearings still correlate with the true AoA. Conversely, a non-line of sight situation is defined by the absence of any useful information in the bearings. The user must enter settings for both parameters of the model. The *SignalShark* user interface uses the terms "LOS Prob." for the assumed probability of line of sight situations, and "Bearing Error" for the assumed standard deviation of the useful bearings.

Under ideal conditions, the "Bearing Error" should be set to the RMS value of the DF uncertainty specified in the DF antenna data sheet, and "LOS Prob." should be set to 100%. More realistic scenarios take the uncertainty in the north reference and the bearing uncertainty due to multipath propagation into account in the "Bearing Error".

Evaluation of homing-in tests conducted by NARDA using the ADFA-1 with transmitter frequencies of around 950 MHz shows that a "Bearing Error" between 5° and 10° and a "LOS Prob." of 50 % describe a good approximation of the true PDF of the bearings taken during homing-in drives in moderate urban environments.

Note that the "Bearing Error" and "LOS Prob." values are not critical for the rate of convergence and the accuracy of the localization algorithm. However, they do have a significant influence on the assumed localization uncertainty, which is displayed as an uncertainty ellipse around the estimated position of the transmitter. If realistic parameters are entered in the model, the ellipse will delineate the area where the transmitter is located with a probability of 95%.

The localization algorithm of the *SignalShark* is based on a matrix that represents a rectangular area on a geographic map. Each element of the matrix represents a discrete position on the map. The element values are related to the probability that the suspicious transmitter is located at the assigned geographic position. The probability values are normalized to the maximum value of the matrix, color coded, and displayed as an overlay on the geographic map. This matrix is therefore called a localization heat map. Red areas denote likely and blue areas unlikely locations of the suspicious transmitter. The entire heat map is recalculated for each new bearing to make a new estimate of the location of the suspicious transmitter.

It is important that the heat map area selected is large enough to ensure that the suspicious transmitter is located within it. The heat map area is highlighted on the geographic map and can be adjusted as desired. By default, the spatial resolution of the heat map is the same as that of the slippy map tiles at the zoom level used when the heat map area is selected. The spatial resolution can also be set to values between 1/64 and 8 times the default resolution. The use of relative spatial resolutions down to 1/4 is recommended if a very large area is selected. The default value of unity is a good choice for medium-sized areas. Values higher than unity make sense if a very high resolution external display is used to make the selection.

It is possible that the transmitter position will be outside the heat map if the area selected is too small. In this case, the localization process must be stopped, a new heat map area selected, and the heat map recalculated for bearings already taken. This can be time-consuming if thousands of bearings have already been taken. It is therefore recommended that a heat map area that is probably oversized should be selected from the outset.

The *SignalShark* localization algorithm can be applied to discrete bearings or to a record of continuous bearings. For the former, the algorithm uses all the bearings stored in a specified folder on the data logger. This is useful when bearings from fixed locations are to be processed. For the latter, the algorithm uses the bearings stored in a bearing record. This method is useful for homing-in drives. When a new discrete bearing is saved or a new bearing occurs during a continuous recording, the algorithm automatically updates the heat map and the localization result.

In most cases, the heat map at the start of a homing-in drive in an urban environment will not show any useful information. This is because the localization algorithm needs line of sight situation bearings from a sufficient number of significantly different locations of the direction finder. It is therefore important to keep moving and scan the complete area where the suspicious transmitter might be located. The useful bearings will intersect at a point close to the position of the suspicious transmitter, while the random bearings will intersect at random positions. It therefore makes sense after driving around the area for some time to drive then into the direction of the current hot spot and then drive around the hot spot area and observe the changes in the heat map. The transmitter has probably been located if the red area on the heat map is reducing in size and its center is not moving significantly.

In most cases, the estimated transmitter location will anyway converge on the true transmitter location as the drive time gets longer. There are situations however where this is not the case:

- Dominant reflectors are present that have more line of sight situations than the transmitter.
- Other transmitters that use the same part of the spectrum as the suspicious transmitter are present close to the search area.

Unfortunately, no reliable criterion can be extracted from the bearing data alone that could be used to decide whether the suspicious transmitter has been localized or not. The user must therefore decide when to end the homing-in drive. This decision can be aided by checking the plausibility of the estimated location of the suspicious transmitter. For example, if the transmitter antenna can be seen then the homing-in process can be ended. If the situation is unclear, it can make sense to pause the drive in the vicinity of the estimated location and inspect this area visually or with handheld DF equipment. If the result of this inspection is negative, the homing-in drive can be resumed in areas that have not yet been visited.

Two additional parameters can also be used to optimize the localization algorithm.

The first is the “Velocity Squelch” parameter, which is important for homing-in drives. It prevents possibly large bearing uncertainties that are due to large uncertainty in the GNSS velocity direction at low speeds or when stopping. It should be set to a value of about 10 km/h if the average driving speed is much higher than 10 km/h. A lower squelch value may be necessary if the traffic situation does not allow average speeds above 30 km/h.

The second parameter, “min. DF Quality”, can be used to discard bearings that have low DF quality values. The DF quality was defined above in the introduction of the correlative interferometer. The DF quality of useful bearings is often greater than 70 %. However, it is better not to discard low DF quality bearings if this means that too few bearings are available.

2 Localization based on TDoA

Localization based on TDoA requires at least three receivers that can provide timestamped I/Q data. The receivers must be synchronized to the same time source. The UTC time distributed by a GNSS is often used for this purpose. The I/Q data from the receivers are transmitted via a network connection to a control computer, where they are correlated to determine the differences in the arrival times of the transmitter signals. The position of the suspicious transmitter is calculated from these TDoA values and the known positions of the receivers using a TDoA localization algorithm. The uncertainty of the TDoA values increases with the reciprocal of the signal bandwidth. Thus localization based on TDoA is not the first choice for narrowband signals.

2.1 *SignalShark* as a TDoA receiver

The accuracy of the timestamps relative to an external time source is an essential property of a TDoA receiver. A pulse per second (PPS) signal is used for the time synchronization process in the *SignalShark*. The PPS signal can come from the GNSS receivers built into the *SignalShark*, or from a NARDA antenna, or possibly via a dedicated SMA connector. An innovative hybrid phase locked loop (PLL) ensures that all the local oscillators, sample rates and timestamps of the *SignalShark* are synchronized to the selected PPS signal. This hybrid PLL ensures that the typical RMS uncertainty of the internal PPS signal of the *SignalShark* relative to the PPS signal of the time reference is only 1.4 ns. The *SignalShark* timestamps have a resolution of 4.9 ns. Therefore an additional RMS uncertainty of 2 ns due to the timestamp quantization must be taken into account. The PPS signals of the GNSS receivers used in the *SignalShark* and in the NARDA antennas have an extended time uncertainty of < 20 ns under clear sky conditions. This means that the extended uncertainty of a time difference measured between two receivers is < 30 ns under clear sky conditions. If this is not precise enough the PPS signals of external time sources can be used for receiver time synchronization.

The I/Q data bandwidth is another important property of a TDoA receiver. The 6-dB bandwidth CBW of the *SignalShark* channel filter ranges from 25 Hz to 40 MHz. At least ten values are available in each decade. The CBW can therefore be tuned precisely to the bandwidth of the suspicious transmitter. Just as with AoA measurements, the recommendation is to set the CBW to approximately the same as the occupied bandwidth of the transmitter. The CBW can also be set to smaller values if required due to a limited network bandwidth. However the optimal TDoA accuracy can't be reached in this case. The I/Q data sampling rate is usually 1.28 times the CBW. The *SignalShark* also has oversampled channel filters that have a sample rate of 2.56 times the CBW. These oversampled filters are recommended for TDoA because they enable more precise TDoA measurements.

If there is a high bandwidth network connection between the control computer and the receivers, the receivers I/Q data can be streamed directly to the central computer. This configuration allows reliable real time localization even of brief and infrequent transmissions. The *SignalShark* can stream I/Q data with sample rates of up to 25.6 MHz via Ethernet.

The available network connection is often too slow for I/Q streaming in practice. In such cases, the I/Q data must be stored by the receiver before being transmitted over the network. In a future firmware release the *SignalShark* will have a sophisticated, built-in I/Q recorder with a capacity of more the 200 million I/Q samples that can be used with sample rates up to 51.2 MHz for this purpose.

The group delay variations in the frontend of the *SignalShark* are compensated in its digital IF equalizer, so they do not affect the precision of time difference of arrival measurements.

The *SignalShark* has been successfully integrated into third party TDoA systems and has proven its high timestamp accuracy. The integration process is quick and simple, as *SignalShark* uses the well-documented

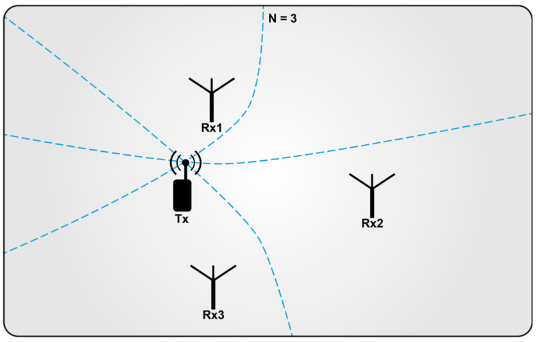

Figure 14. Three receivers localizing a transmitter using TDoA hyperbolas

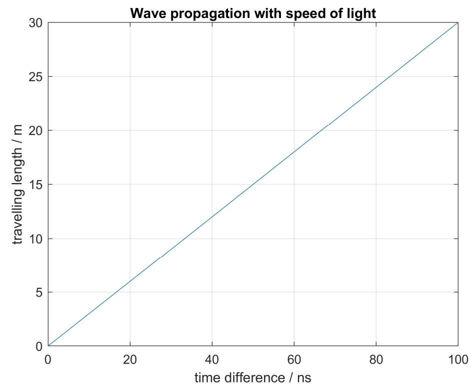

Figure 15. Travelling length versus time difference for speed of light

remote command reference guide for its SCPI commands and the widely-used VITA 49 format for its I/Q data streaming.

2.2 TDoA system

A network that connects the control computer to the receivers is an essential precondition for a TDoA system. If the network is run by a private company or an institution, all the devices in the TDoA system can have a fixed IP address that allows them to communicate directly. However, a fixed IP address is often unavailable if the devices only have access to the public internet, for example via a mobile network or DSL modems. In such situations, the devices usually communicate over a VPN (virtual private network) server. The VPN server can be run by the company or institution that uses the TDoA system. Alternatively, a commercial VPN server can be used. In any case, VPN client software must be installed on the central computer and on the receivers if VPN communication is desired. The operating system of the *SignalShark* is Windows 10, so the preinstalled operating system's VPN client software can be used or other VPN client software can be installed and used instead.

Dedicated TDoA software running on the central computer controls the TDoA receivers. The TDoA software usually displays the spectrums of all the receivers, so that the on-air transmitters can be seen and any suspicious transmissions identified and selected for the localization process. If network capacity allows, the I/Q data of the selected channel are streamed to the control computer. Otherwise, equal time span blocks of I/Q data are transferred periodically to the central computer. Triggered recordings can be transmitted to the central computer if the transmissions are infrequent or short. The central computer correlates the transmitted I/Q data for each possible receiver pair. The position of the maximum of the correlation function indicates the TDoA between the two receivers of each pair.

Once the TDoAs of all possible receiver pairs have been calculated, the suspicious transmitter can be localized. Under ideal conditions, the time differences of arrival of each receiver pair determine a hyperboloid on which the suspicious transmitter must lie. The rotational axis of this hyperboloid runs through the positions the two receivers in a pair. The hyperboloids of different receiver pairs intersect at the position of the suspicious transmitter. The transmitter position can be determined with three receivers if it lies on an assumed plane. This is often assumed to be the plane defined by the three receiver positions. The 3D position of the transmitter can be determined using four or more receivers.

Figure 16 demonstrates the shape of the TDoA hyperbolas of two receivers in the X–Y plane. 21 hyperbolas with equidistant TDoA values are shown. The X and Y axes are normalized to the distance between the receivers. The TDoA values of the hyperbolas are normalized to the propagation time between the two receivers, and range from –1 to +1 in steps of 0.1. Note that the rotational axis of the associated hyperboloids runs through the positions of the two receivers.

Figure 17 shows a localization example with three receivers. The three hyperbolas intersect at the position of the suspicious transmitter.

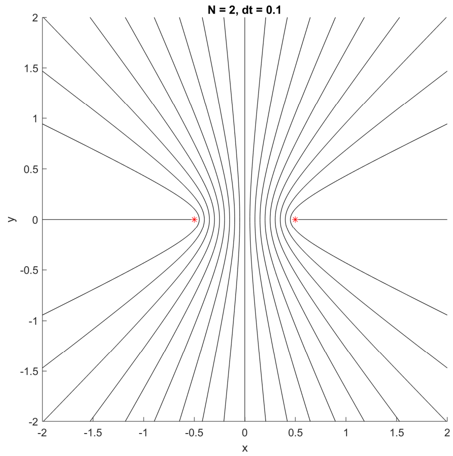

Figure 16. TDoA hyperbolas for a receiver pair

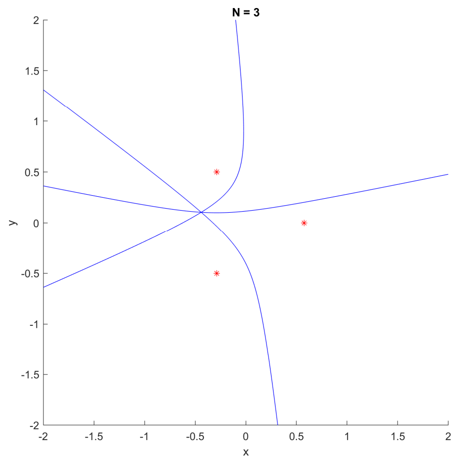

Figure 17. Localization example with three receivers

The geometry of the receiver positions affects the accuracy of the estimated transmitter position. Figures 18 and 19 show the TDoA hyperbolas of receivers that all lie on the X–Y plane. 21 hyperbolas with equidistant TDoA values are shown for each receiver pair. The X and Y axes are normalized to the distance between two adjacent receivers. The TDoA values of the hyperbolas are normalized to the propagation time between two adjacent receivers, and range from –1 to +1 in steps of 0.1.
Figure 18 shows the minimum TDoA system configuration with three receivers. The points where three hyperbolas intersect are the transmitter locations for the corresponding TDoA values. The areas where such points lie close together are areas with a high localization accuracy. The highest accuracy is obviously reached at the center of the receiver positions. The accuracy is significantly reduced in the area outside the triangle delineated by the three receivers. The accuracy also depends on the direction. The accuracy is worst in the directions behind the receivers. Transmitters outside the triangle can still be localized with reasonable accuracy if they lie in the directions between the receivers.
Figure 19 shows a TDoA configuration with four receivers. The points where six hyperbolas intersect are the transmitter locations for the corresponding TDoA values. Here too, it is obvious that the highest accuracy is achieved at the center of the receiver positions. The accuracy is significantly reduced outside the rectangle defined by the four receivers. However, the accuracy is not so dependent on the direction, so in most directions transmitters that are outside the rectangle can still be localized with reasonable accuracy.
A TDoA system works well if all the receivers receive a line of sight component of the transmitter signal and if this component is dominant compared to the reflected signal components. If the reflected signal components at a certain receiver are dominant, the magnitude of the TDoA values of all receiver pairs containing this receiver will be higher than expected. The localization result will therefore be incorrect if those receiver pairs are not excluded from the calculation of the transmitter position. Thus, TDoA systems with many receivers and automatic recognition of those receivers without line of sight should be used in situations where line of sight may not be present and strong reflections are expected.
Many TDoA systems only produce a 2D position result as they assume that the suspicious transmitter lies on a presumed plane. Many TDoA systems even assume that all the receivers lie on the same plane. Even if the elevation of the transmitter is not of interest, it is worth mentioning that the 2D result of such systems may be inaccurate if the assumptions are not fulfilled. True 3D TDoA systems with at least four receivers may therefore improve accuracy if it has to be assumed that the receivers and the suspicious transmitter are located at significantly different heights.
In principle, a TDoA system can also work if the receivers or suspicious transmitters are moving. However, the frequency shift due to the Doppler effect would need to be taken into account in the calculation of the transmitter position.

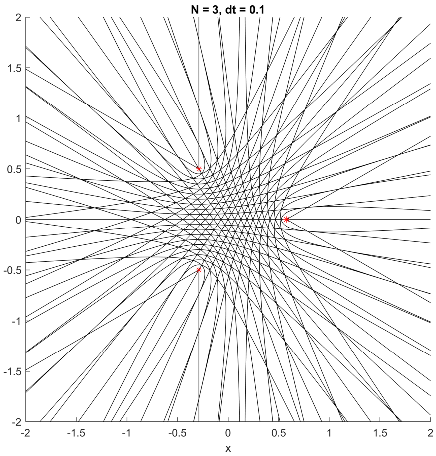

Figure 18. TDoA hyperbolas for three receivers

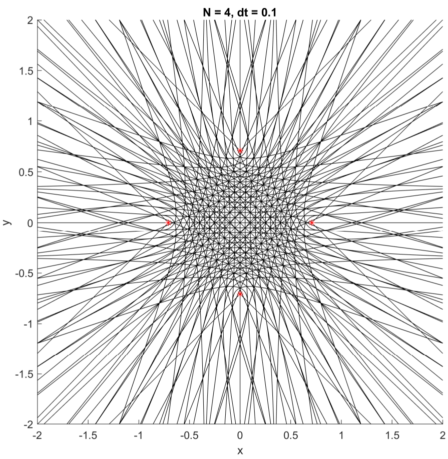

Figure 19. TDoA hyperbolas for four receivers

3 Summary

Several localization methods based on AoA are discussed in depth. In particular, localization methods based on cost effective single channel receivers are described in detail, and the special features of related NARDA equipment are discussed. The outstanding localization algorithm used by the *SignalShark* is described. The different localization methods based on TDoA are also described in depth. The outstanding TDoA properties of the *SignalShark* in particular are discussed in detail.

4 Literature

ITU: “Recommendation ITU-R SM.2060-0, Test procedure for measuring direction finder accuracy”, ITU, Geneva 2014

ITU: “Recommendation ITU-R SM.2061-0, Test procedure for measuring direction finder immunity against multi-path propagation”, ITU, Geneva 2014

ITU: “Recommendation ITU-R SM.2096-0, Test procedure for measuring direction finder sensitivity in the VHF/UHF frequency range”, ITU, Geneva 2016

Tuncer, T.E. and Friedlander, B.: “Classical and Modern Direction-of-Arrival Estimation”, Academic Press, Burlington, 2009

Gemayel, N.E.: “Abschlussbericht: Signalverarbeitungsalgorithmen und Methoden zur Unterstützung des dynamischen Spektrummanagements”* Karlsruhe Institute of Technology, Karlsruhe, 2013

* “Final report: Signal processing algorithms and methods that support dynamic spectrum management”