TDOA Localization: From Theory to the Field

Overview

—

—

Version 1.0
10 October 2023

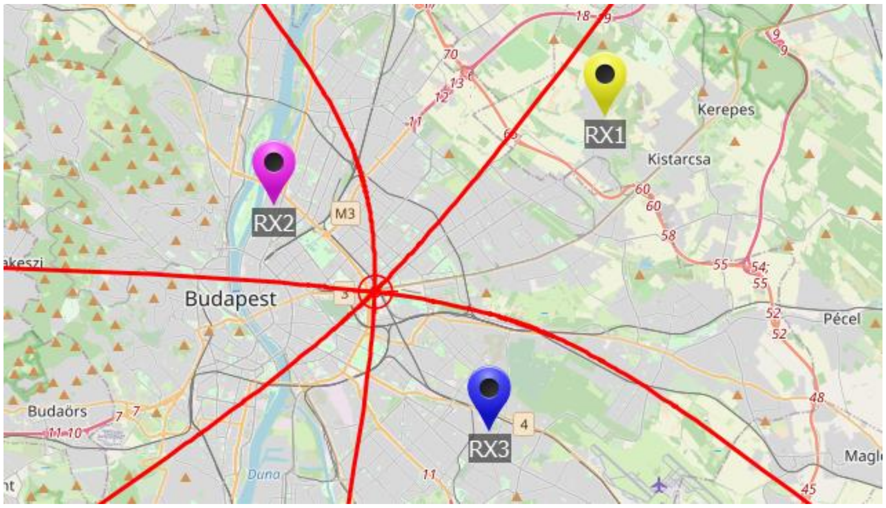

Localization is essential in spectrum monitoring applications for helping operators pinpoint the source of transmissions they are interested in.

Time-difference of arrival (TDOA) localization has emerged as an attractive solution to passively find the source of signals of interest. In short, **TDOA localization uses three or more receivers to find the position of a target based on the difference between the time it takes the signal to arrive at each receiver**, hence the name "TDOA."

TDOA localization has numerous advantages. First, it can be used with commercial off-the-shelf (COTS) or application-specific antennas, e.g., wideband monitoring antennas, directional antennas, or antennas with a focus on special frequencies. Owing to the correlation gain, TDOA can be applied on signals close to the noise floor or even below it. Also, TDOA is more resilient to multipath than alternative solutions like direction-finding (DF) localization.

In this article, we will delve into the technical details behind TDOA localization, which will help us understand how to benefit from its full potential. We will also discuss the quality of location estimates obtained with TDOA. Finally, we will compare TDOA and DF localization, which will help us select the most suitable method for our application.

The Principles of TDOA Localization

The conventional TDOA localization needs at least three receivers to locate a target in a 2D plane. The receivers need to be time-synchronized, usually through GNSS.

One common way to estimate the TDOA of the signal arriving at two receivers is based on the cross-correlation between the received signals. When the signal arrives at each of the three receivers, it is streamed to a central processing station. There, the signals from each two receivers are cross-correlated in order to find the TDOA.

The cross-correlation function (CCF) of two signals f and g is obtained by multiplying one of the signals (for instance, f) by the complex conjugate of the other signal (g), integrating over the total duration of the two signals, and then repeating this process for a set of time lags τ.

We will explain the computation of the CCF using as an example a signal that arrives at two receivers, which we will call RX1 and RX2. The signal travels through two different paths, so the waveforms that arrive at RX1 and RX2 will not be identical. Let's denote by f and g the signals arriving at RX1 and RX2, respectively. We consider that the distances between the signal source and each receiver are different, so the signal will arrive at different times at RX1 and RX2.

Figure 1 illustrates this example: the first plot shows the signal f that arrives at RX1 at time τ_{RX1}. Similarly, the second plot shows the signal g that arrives at RX2 at time τ_{RX2}. The signals f and g have similar shapes, but g has a higher noise level.

Intuitively, cross-correlating two signals is equivalent to shifting one of them until we find the time lag at which they are the most similar. The measure of similarity is given by the integral: because we are summing up the product of the two functions at all time values, the sum will be the largest when all the peaks (both negative and positive) are perfectly aligned. At this time lag, the CCF will have the maximum value.

In Figure 1, the third and the fourth plots show signal g shifted with the delays $\Delta\tau_1$ and $\Delta\tau_2$, respectively. In this example, $\Delta\tau_1 = \tau_{RX1} - \tau_{RX2}$; in other words, $\Delta\tau_1$ is the delay applied to g such that the useful parts of the signals f and g are aligned. The last plot shows the resulting CCF. The maximum of the CCF is found for the delay $\Delta\tau_1$. The time shift $\Delta\tau_1$ is also the TDOA between RX1 and RX2, since it is equal to the difference between the arrival times of the signal at the two receivers.

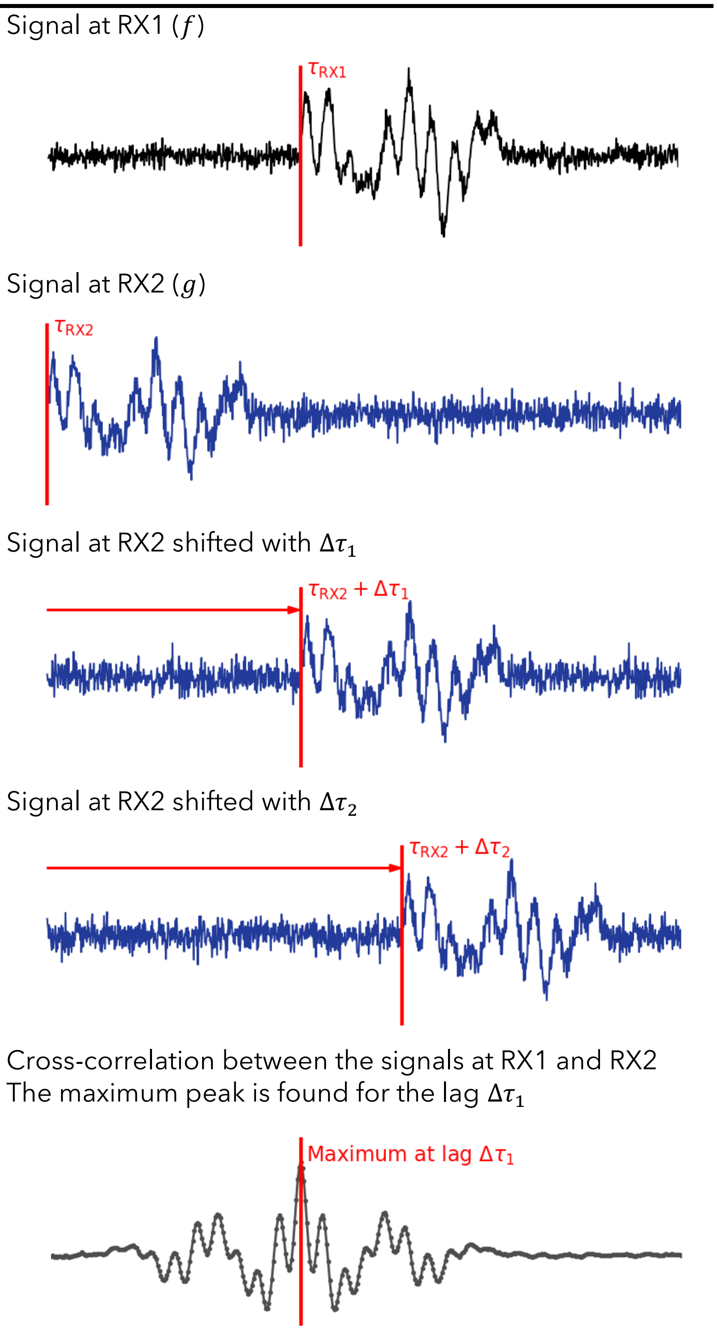
Signal at RX1 (f)

Signal at RX2 (g)

Signal at RX2 shifted with $\Delta\tau_1$

Signal at RX2 shifted with $\Delta\tau_2$

Cross-correlation between the signals at RX1 and RX2
The maximum peak is found for the lag $\Delta\tau_1$

Figure 1. The receivers RX1 and RX2 receive a signal arriving through different paths at different times. The received signals are cross-correlated in order to find the time lag at which they are perfectly aligned. This time lag represents the TDOA of the signal between RX1 and RX2. In this example, the perfect alignment is obtained for the lag $\Delta\tau_1$.

> A signal of interest arrives at receivers placed at different locations. By cross-correlating the versions of the signal that arrive at two receivers, we obtain the difference in their arrival times, also known as the TDOA.

So far, we have explained how to acquire the TDOAs. What might not be obvious is how to turn them into locations, which is what we are interested in.

First, the *time* difference can be easily converted to a *distance* difference by multiplying it by the speed of light.

Then, one distance difference gives us a set of possible locations for the transmitter. Let's take as example Figure 2 and suppose that we have computed the *distance* difference between the paths traveled by the signal to receivers RX1 and RX2. This value determines a specific hyperbola, which has the *foci* at the two receivers. Using the *sign of the difference*, we limit the search space to only one branch of the hyperbola, with the focus at RX1 (the solid red line). The transmitter of interest can be anywhere on this line.

By computing at least three TDOAs (for instance, between the pairs RX1–RX2, RX1–RX3, and RX2–RX3), we can find the target of interest at the intersection of all hyperbolas, as shown in Figure 3.

In the **Decodio Localizer** software, TDOA localization can be applied on both continuous transmissions and burst signals.

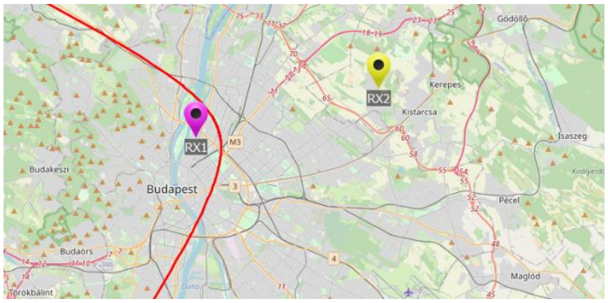

Figure 2. The TDOA is transformed into a *distance* difference of arrival at two receivers. We consider that the signal arrives at the two receivers (RX1 and RX2). The signed distance difference determines a branch of a hyperbola (the solid red line) with the focus at RX1. The source can be anywhere on this line.

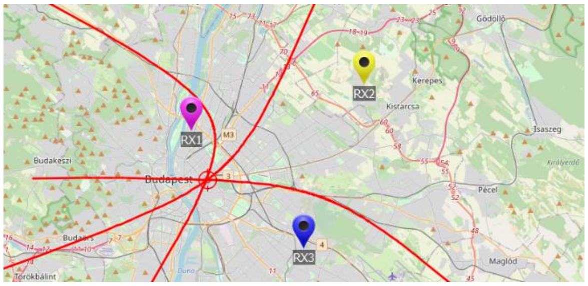

Figure 3. By computing more TDOAs, the target can be localized at the intersection of three hyperbolas.

The Quality of TDOA Localization

A frequently-asked question that arises at this point is: How good is TDOA localization? First, we need to define what "good" means. In other words, we need to define our quality metrics. Although we usually think of the *accuracy* of location estimates as being the most important quality factor, we will see that their *precision* can be just as (if not more) critical.

How to define the quality?

The difference between accuracy and precision is often not well understood. However, it is essential in order to interpret our results.

A localization solution typically returns a set of location estimates for a target signal. Because of the noise present in the signal, the location estimates will also be noisy. Therefore, we can improve the results if we look at the *average* location over a time period. This will mitigate part of the measurement noise. In Figure 4, we consider that the average location is the center of the circle which wraps the location estimates.

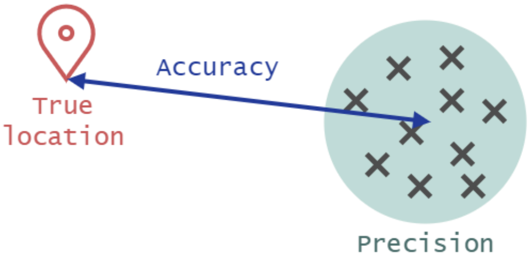

Figure 4. Accuracy vs. precision in localization.

The distance between the mean of all computed locations and the true location of the signal source is the *accuracy* of the location estimate.

The spread of the estimated locations around their mean gives us the *precision* of the location estimate.

We can define the accuracy and precision in terms of the location error, which we denote by $\boldsymbol{e}$. We define the location error as the Euclidean distance between an estimated location and the true location:

$$e_i = \|\boldsymbol{p}_t - \boldsymbol{p}_i\| = \sqrt{(x_t - x_i)^2 + (y_t - y_i)^2},$$

where $\boldsymbol{p}_t$ is a vector that contains the 2D Cartesian coordinates x_t, y_t of the true location. Similarly, the vector $\boldsymbol{p}_i(x_i, y_i)$ contains the 2D coordinates of the i^{th} location estimate.

The accuracy is then the mean location error for a set of N location measurements:

$$\text{accuracy} = \sum_{i=1}^{N} e_i.$$

The precision is most commonly defined as the standard deviation of the location errors. Since the standard deviation is sensitive to outliers, other measures (e.g., the interquartile range) can be chosen instead depending on the application.

Accuracy and precision in practice

We might be tempted to always desire the highest possible accuracy from a localization solution (e.g., on the order of decimeters or meters). However, this might not be necessary in many scenarios. Relaxing the requirements can bring many benefits such as a lower cost of the setup or an easier operation.

For example, if we want to localize a target inside a city, then an accuracy of several hundred meters can be enough to pinpoint the target to several buildings. However, if the target needs to be localized at a particular street number, the localization method should yield an accuracy of at least several meters.

The precision should not be ignored either. Even if the distance between the center of the cluster and the true location (i.e., the mean distance error) is small, we can have a large spread of location estimates around the mean location. This means that our precision is low and we cannot locate a target with high certainty. In practice, this could mean that a larger area should be searched for the source of interest, which might require more time and resources.

What errors to expect from TDOA?

While it is hard to give a general answer that is applicable to all localization scenarios, we will highlight the most important factors that influence the quality of TDOA localization: the bandwidth, the time synchronization between the receivers, the multipath profile, and the geometry of the receivers.

To illustrate the influence of these factors, we will use data from a measurement campaign to give an idea about the accuracy we can expect from **Decodio Localizer**'s TDOA solution.

The setup consisted of three receivers spread out around Budapest, shown previously in Figure 3. We used Narda SignalShark receivers. In this case, we localize TETRA signal sources, but the TDOA can also be applied on other signals.

While we do not know the exact true location of the signal source, in most cases we can pinpoint it with high confidence to nearby landmarks (e.g., the train station), as this is where TETRA stations are likely to be found. Please note that this example is for illustration purposes and, in other applications, the achievable accuracy will highly depend on the particular setup.

Bandwidth. The signal bandwidth is critical for the precision of TDOA localization. This is because the auto-correlation function (ACF) of a signal is the Fourier transform of its power spectral density. Therefore, as Figure 5 shows, the wider the signal bandwidth is, the sharper the main lobe of the ACF becomes and the more precise the timing estimate (and hence the TDOA) will be.

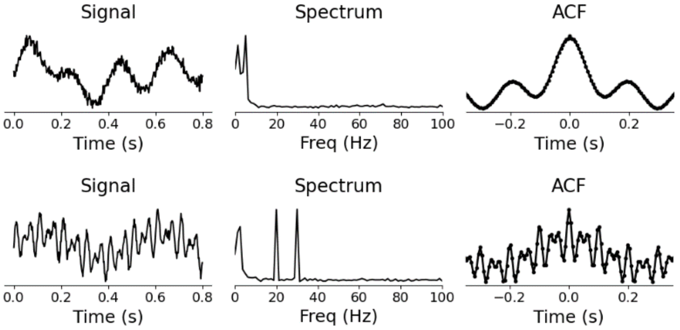

Figure 5. Signal bandwidth vs. shape of the autocorrelation function (ACF) with narrow band signals (top) vs. wider band signals (bottom). Note how the width of the main lobe in the ACF gets narrower with a wider band.

In the following, we will use a TETRA signal from the measurement campaign to show the impact of bandwidth on the precision of TDOA localization.

Narrowband

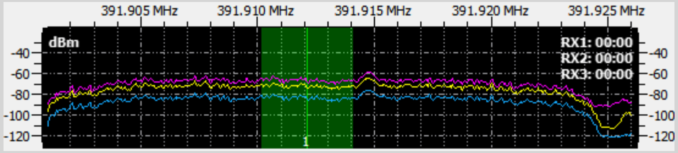

We want to localize a TETRA signal with a bandwidth of 25 kHz. However, to simulate the impact of a signal with a low bandwidth, we select only a portion of 4 kHz from the signal (the green stream from the figure).

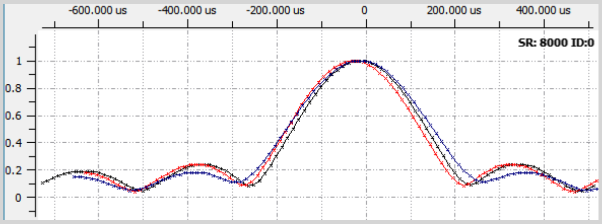

The main lobe of the resulting CCF has a width of approximately 450 μs.

We obtain a localization precision of approx. 188 m. The total span of the heatmap is 1.2 km.

Wideband

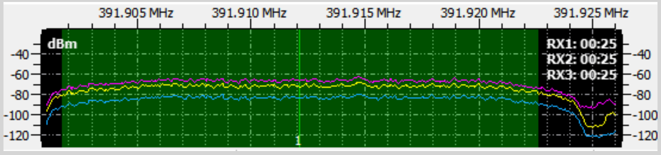

In this example, we have the same signal as before, only that this time we use almost the entire bandwidth of the signal to compute the TDOA (the green area now spans the width of the signal bandwidth).

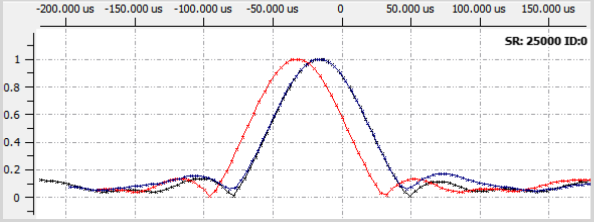

The width of the main lobe in the CCF is now approx. 125 μs, so more than 2x sharper than in the previous narrowband example.

This leads to a precision of 67 m. The span of the heatmap is now approx. 500 m. The precision is more than twice as high for WB signals than for NB signals. In practice, this means that we need to search for the source of interest over a smaller area.

Receiver synchronization. The receivers need to be synchronized using GNSS, which is affected by jitter errors and can decrease the localization precision. The receivers must be placed outdoors, where they can have a clear sky view, in order to minimize GNSS errors. If the receivers are not properly synchronized to account for factors such as cable length and processing time, systematic biases of up to 100 ns (equivalent to 30 m) can affect the TDOAs.

Multipath and non-line of sight (NLOS). Multipath propagation is the phenomenon through which signals arrive at a receiver through multiple paths. NLOS propagation occurs when the signal traveling through the direct path between the transmitter and the receiver is obstructed by an object.

Multipath propagation can lead to multiple peaks in the ACF. In this case, it is harder to select the correct peak especially if the direct path is (partially) obstructed. As a result, the location estimate can be biased.

Both multipath and NLOS are inevitable in crowded urban environments and can introduce errors of even hundreds of nanoseconds in the TDOAs. Even in remote areas, trees, hills, or mountains can lead to multipath and NLOS propagation. This decreases the accuracy of TDOA localization. For this reason, it is advisable to place the receivers on high places (e.g., tall buildings, hills) to maximize the chances of the receiver being in line of sight (LOS) with the transmitter.

Geometry. The placement of the receivers with respect to the target also has an important influence on the precision. Usually, the localization accuracy and precision are the highest when the target is inside the convex hull defined by the receivers. However, there can be cases in which the signal source is outside this area, as shown in Figures 5 and 6. In this example, the estimated locations are spread over an area of 2 km due to the poor geometry. Note that the receivers must not be placed in a line, as this leads to ambiguous locations.

SNR. While the SNR affects the precision of location estimates, its effect becomes evident only for wideband signals. In fact, TDOA can be applied on signals close to the noise floor or even below it thanks to the correlation gain [1]. When computing the cross-correlation, signals that are uncorrelated (such as the noise) are attenuated.

On the other hand, the signal of interest will correlate with itself and will be preserved. Some signals, such as those using spread spectrum modulation, can have a signal level even below the noise floor and still be recovered thanks to the good auto-correlation properties of the spreading codes.

> Multipath propagation is the main factor that affects the localization accuracy. The precision is mainly limited by the signal bandwidth.

To summarize, the main factors that affect the localization accuracy are the receiver synchronization and the propagation conditions. Since a perfect receiver synchronization is hard to achieve and the system will likely operate in multipath environments, we should expect a localization accuracy of tens to hundreds of meters.

The signal bandwidth is the main factor driving the precision, since PMR signals are typically narrowband. Only for wideband signals do factors such as the GNSS jitter and the SNR start to become more important.

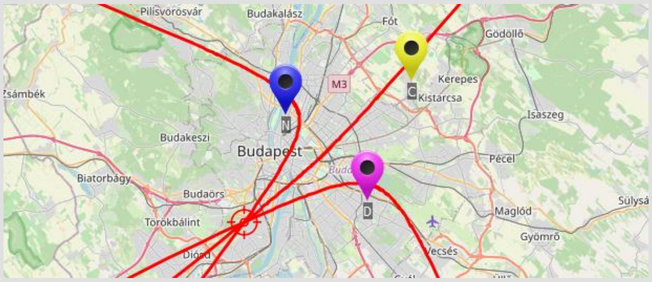

Figure 6. The source of the signal is located outside the convex hull delimited by the receivers.

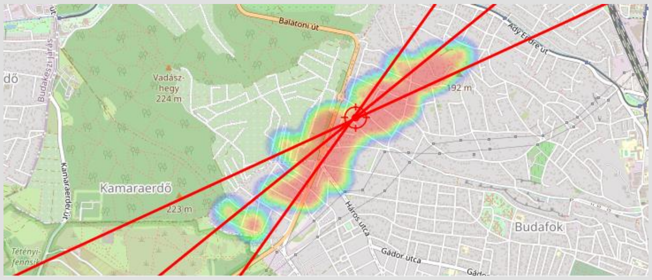

Figure 7. Close-up of Figure 6. The precision is lower than if the source were inside the area delimited by the receivers. The point cloud spans an area of approx. 1.3 km.

TDOA vs. DF Localization

Having understood the principles of TDOA localization, we can now compare it with its most popular alternative, DF localization. This will help us make an informed decision about which method is best suited for a particular application. Table 1 summarizes the main differences between the two methods.

Perhaps the most crucial difference between the two techniques is the type of antenna they require. In order to perform DF localization, we need complex, specialized DF antennas. Most DF antennas have a frequency range limited to 8 GHz [2, 3, 4] and are vertically polarized [5]. This limits the type of signals we can localize using DF.

In contrast, TDOA can be used with wideband monitoring antennas, directional antennas, or antennas with a focus on special frequencies. This could mean, for instance, that we can use TDOA localization even with signals from the K band intended for satellite communications. Satellite communications are becoming crucial in military operations, as they can provide internet connectivity and communication capabilities even in remote areas or where ground communications structures are damaged [6]. Therefore, being able to localize satellite terminals based on their uplink transmissions can prove to be an important tactical advantage.

On the other hand, it is worth noting that TDOA localization systems need at least three receivers (as opposed to minimum two for DF) with a stable network connection between them in order to transfer the data needed to perform the cross-correlation.

Table 1. Comparison between TDOA and DF localization.

	TDOA Localization	**DF Localization**
SNR	Works even in low SNR	Needs good SNR
Antennas	Works with COTS antennas	Needs bulky, complex DF antennas
Frequency	Limited to the frequency range of regular antennas	Limited to the frequency range of DF antennas (typ. up to 8 GHz)
Number of receivers	Needs min. 3 RXs	Needs min. 2 static RX or 1 moving RX (e.g., in a drive test)
Geometry of receivers	Best if source within the area defined by the RXs; RXs cannot be placed in a line	Flexible
Signal bandwidth	Works best with WB signals	Preferably NB signals

Conclusion

In conclusion, TDOA localization can be an important ally for passive source localization. It is especially useful in cases in which we need to monitor possibly weak signals with lean, wideband, COTS antennas. The accuracy and precision of TDOA localization can easily reach tens to hundreds of meters in regular applications. Therefore, TDOA localization can mitigate some of the disadvantages of DF or even complement it in hybrid localization method.

References

[1] *Reducing the Noise Floor and Improving the SNR with Cross-Correlation Techniques*, by Jithesh Srinivas, Zurich Instruments. Link: https://www.zhinst.com/ch/en/blogs/how-reduce-noise-floor-and-improve-snr-employing-cross-correlation-techniques
[2] Direction-finding antennas from Rohde and Schwarz. The maximum frequency covered is 8500 MHz, by the R&S®ADD597 antenna. https://www.rohde-schwarz.com/products/aerospace-defense-security/direction-finding-antennas_334198.html
[3] Alaris DF antennas cover frequencies up to 8 GHz. https://alarisantennas.com/by-application/direction-finding/
[4] CommsAudit DF antennas currently have a maximum range of up to 6 GHz. https://www.commsaudit.com/product/antennas/
[5] "The world's first VHF-UHF direction finding antennas for all polarizations," Rohde & Schwarz. Link: https://scdn.rohde-schwarz.com/ur/pws/dl_downloads/dl_common_library/dl_news_from_rs/194/n194_DDF0xAE-DDF195_e.pdf
[6] *Starlink satellite services in Ukraine*, Wikipedia. Link: https://en.wikipedia.org/wiki/Starlink_satellite_services_in_Ukraine