*38th International Electric Vehicle Symposium and Exhibition (EVS38)*
*Göteborg, Sweden, June 15-18, 2025*

# Optimization of Outdoor Positioning for Autonomous Charging

Hannes Reindl, Simon Wellenzohn, Gregor Eckhard
Easelink GmbH, Austria, hannes.reindl@easelink.com

**Executive Summary**

Ultra-wideband (UWB) technology is highly promising for outdoor vehicle positioning, offering centimeter-level precision due to its wide frequency range and resilience against interferences. This paper explores the application of UWB as a positioning and guidance system for electric vehicles (EVs) with Matrix Charging® – a technology for fully automated charging of EVs. Through extensive field and lab tests of the NXP NCJ29D5 UWB chip, key error sources such as multipath effects, line-of-sight loss, and environmental interference were identified and mitigated using a Numerical Maximum Likelihood Estimator (NMLE) algorithm. The system achieved an average accuracy of 20 mm, underscoring the role of UWB in the advancement of secure, precise, and efficient EV charging and future smart mobility infrastructure.

*Keywords: Electric Vehicles, Retrofitting EVs, AC & DC Charging technology, Smart charging, Human-Machine/computer interactions*

## 1 Introduction

Easelink has developed an automatic conductive charging system, called Matrix Charging®, which automatically connects an EV to the grid for charging. This paper explains how vehicle and infrastructure components are aligned by supporting the driver with guidance through Ultra-Wide-Band (UWB) technology. This positioning system is one of the main interface points between the user and the Matrix Charging® technology, therefore representing a crucial part of the Matrix Charging® system. It needs to work to a high degree of reliability from not only an objective/number-based view but also a subjective/user-experience view to convey the feeling of a solid, well-engineered product.

### 1.1 Guidance for the Matrix Charging® System

The Matrix Charging® system charging solution consists of two main components and a mechanism to physically connect them. The first component is the Matrix Charging Connector (MCC) that sits on the underbody of a car and is connected to the EVs On-Board-Charger (OBC). The second is the Matrix Charging Pad (MCP), which is placed on the ground and connected to the grid. Figure 1 depicts both components in a lab setup (*left)* and in a real-world setup (*right*).

Connecting the EV to the grid, i.e., the MCC to the MCP, is done by lowering the Core Unit from the MCC and landing inside the Matrix of the MCP. This means the user has to guide the EV so that the MCC is directly on top of the MCP. While this might be feasible from an outside perspective, a person sitting inside the car can not see through the car's underbody and has to guess how much further the driver has to go. Figure 2 shows the driver's perspective sitting inside an EV, underlining why a positioning system is required.

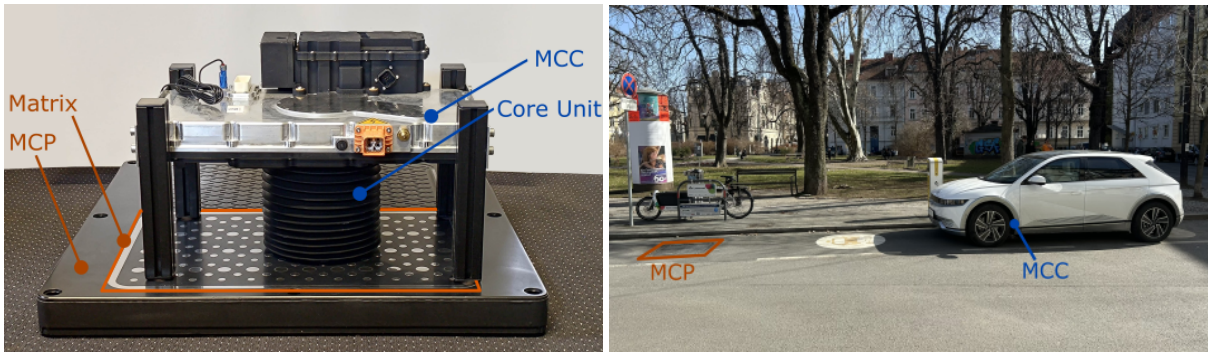


Figure 1: Matrix Charging Connector (MCC) and Matrix Charging Pad (MCP) in lab setup (*left*) and in a real-world scenario (*right*).


Figure 2: View from inside the car of a person trying to place the MCC right on top of the MCP.

Matrix Charging® is currently deployed in the eTaxi Austria project of the cities of Vienna and Graz, the largest pilot project in the world for automated charging. The project aims to demonstrate barrier-free, automated charging points at taxi stands as a key step towards the electrification of urban taxi fleets. In the scope of the eTaxi Austria project, more than 40 EVs have been equipped with Matrix Charging®.

For these retrofitted vehicles, some sort of positioning system was required to precisely guide the drivers onto the MCP while sitting inside the vehicle. Figure 2 already presents Easelink's solution in the bottom left corner - a small display showing the position of the Core Unit relative to the Matrix in real time (see Fig. 3 and 1).

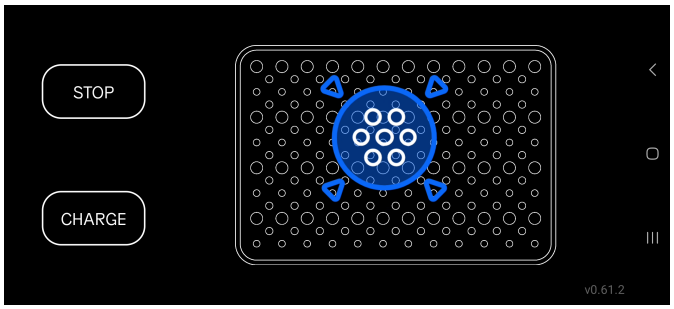


Figure 3: User Interface during a ranging process. The blue highlighted circle shows the current position of the Core Unit relative to the MCP Matrix.

Due to the retrofit application of the Matrix Charging® technology to the target vehicle, an additional display was installed. In the future, for series integration, the system will be fully integrated into the vehicle's original display.

## 2 Using electromagnetic waves to determine distances

Easelink uses the NXP NCJ29D5 UWB chip [1] to solve the positioning task. Unlike traditional wireless technologies, UWB operates by transmitting short, low-power electromagnetic waves over a broad frequency range (typically 3.1 to 10.6 GHz). This enables UWB to measure the Time of Flight (ToF) of signals between devices as depicted in Fig. 4.

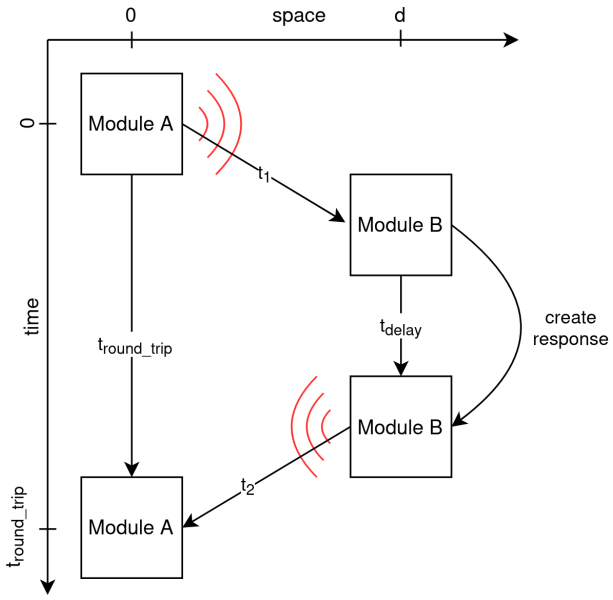


Figure 4: Two UWB modules communicating with each other via electromagnetic waves to determine the distance $d$ between them.

Module A wants to know the distance to Module B and therefore emits an electromagnetic wave that takes time $t_1$ until it reaches Module B. Right after Module B receives this wave, it also starts a timer and computes a response that takes time $t_{\text{delay}}$. Next, Module B sends a wave back to Module A, arriving after time $t_2$. When Module A receives this wave, it stops its timer to determine the round-trip time $t_{\text{round\_trip}}$ and the distance $d$ between the modules can be calculated using equation 1

$$d = \frac{\left(t_{\text{round\_trip}} - t_{\text{delay}}\right) \cdot c}{2} = \frac{(t_1 + t_2) \cdot c}{2} \tag{1}$$

, where $c$ is the speed of light. By subtracting the computation time $t_{\text{delay}}$ from the round-trip-time $t_{\text{round\_trip}}$, only the time-of-flight periods remain. The division by 2 is necessary since $t_1 + t_2$ contains the two-way path and distances are one-way. Even though the computation time $t_{\text{delay}}$ is only in the order of a few nanoseconds and one might consider neglecting it, an error of one nanosecond magnified by the speed of light already results in a distance error of 300 mm.

The NXP NCJ29D5 UWB chip [1] is specifically designed for automotive applications. Operating in the 6-8.5 GHz frequency band, it offers centimeter-level precision required for high-accuracy real-time localization systems. Real-time is key to providing good feedback to the user. Too much delay will cause the displayed position to significantly lag behind the real position, causing the user to overshoot the target and making it uncomfortable to use. The high accuracy aspect is also key since the Core Unit has to land inside the Matrix (see Fig. 1) to establish a proper connection and the positioning algorithm can only be as accurate as the underlying distance measurements allow for.

One of the main benefits of UWB is its ability to measure distance accurately even through plastic, which was not possible with our initial system using ultrasonic waves. This allows sealing the MCC and MCP, simplifying their construction with respect to IP rating. UWB also allows detection of multipath components due to their ultra short time domain pulse [2], increasing its robustness in multipath environments such as the underbody of a car where the Matrix Charging® system operates.

# 3 Physical Setup

Figure 5 shows the location of the UWB modules inside the Matrix Charging® system. The MCP has one module in each corner labeled A1 through A4 and since the entire MCP is fixed in its location, those modules are called anchors. The MCC has two UWB modules labeled T1 and T2, which are moving since the MCC is mounted inside an EV. These modules are usually called tags.

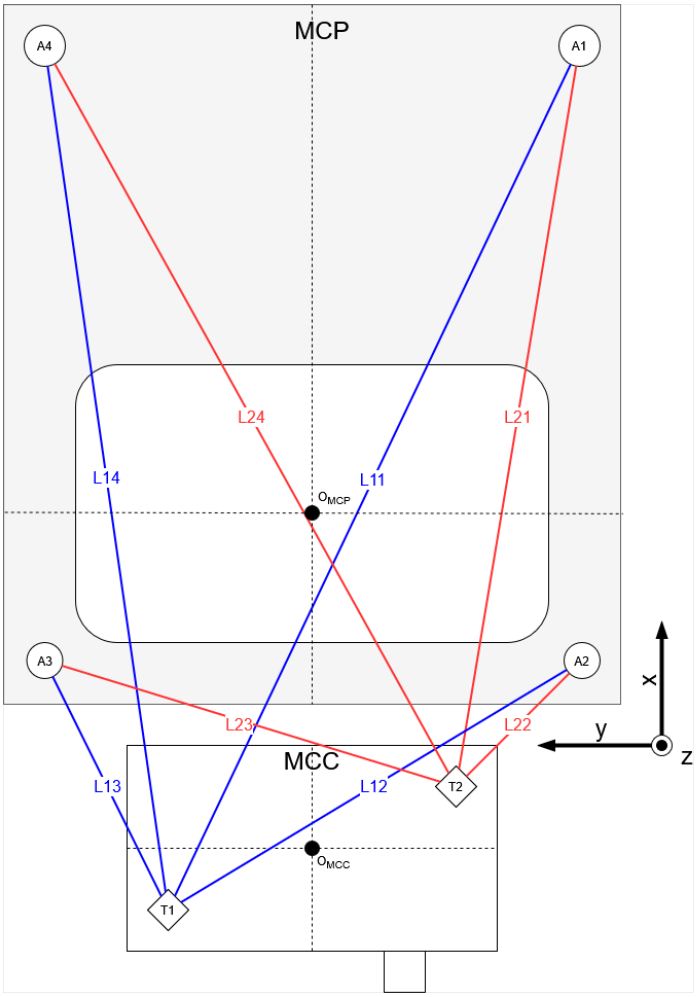


Figure 5: Top-down view of the physical setup consisting of two UWB modules inside the MCC and four modules inside the MCP

The goal of the positioning system is to find the current position of the MCC origin $O_{MCC}$ relative to the MCP origin $O_{MCP}$ . Note that the MCP origin $O_{MCP}$ is placed in the center of the Matrix and the MCC origin $O_{MCC}$ is in the middle of the Core Unit. Using the time-of-flight measurements, as discussed in the previous section, the UWB modules can determine the distances L11 through L24. Although three distances are technically sufficient to estimate a position using triangulation (just like in GPS), an overdetermined system is highly beneficial with respect to accuracy and robustness. Electromagnetic waves are highly susceptible to metallic objects, causing the measured distances to differ greatly from their true value. Since the MCC is mounted underneath a car, there are more than enough metallic objects nearby to interfere with electromagnetic waves, causing inaccuracies. Using this overdetermined system still allows solving the estimation problem even if selected distances are nonsense due to interferences. If all eight distances are useful, the accuracy of the estimated position is greatly improved.

The setup depicted in Fig. 6 shows the used validation setup consisting of a computer-controllable gantry, an MCC and an MCP. The actuators of the gantry allow the MCC to be moved with millimeter accuracy anywhere within the test bench while returning its current position. Using a little math, the true distances L11 through L24 can be calculated and easily compared against the measured distances returned from the UWB modules, as done in the next section.

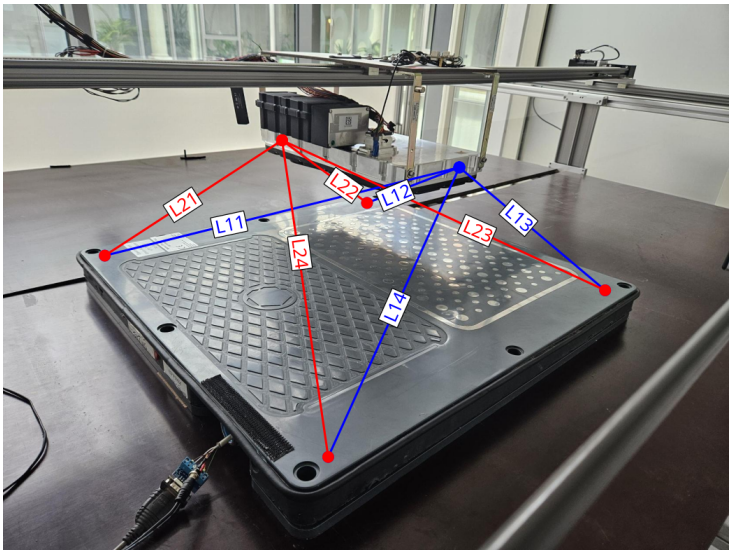


Figure 6: Validation setup with moveable MCC using a gantry.

### 3.1 Real world distance measurements

The probability distributions of the distances measured by the UWB modules are depicted in Fig. 7.

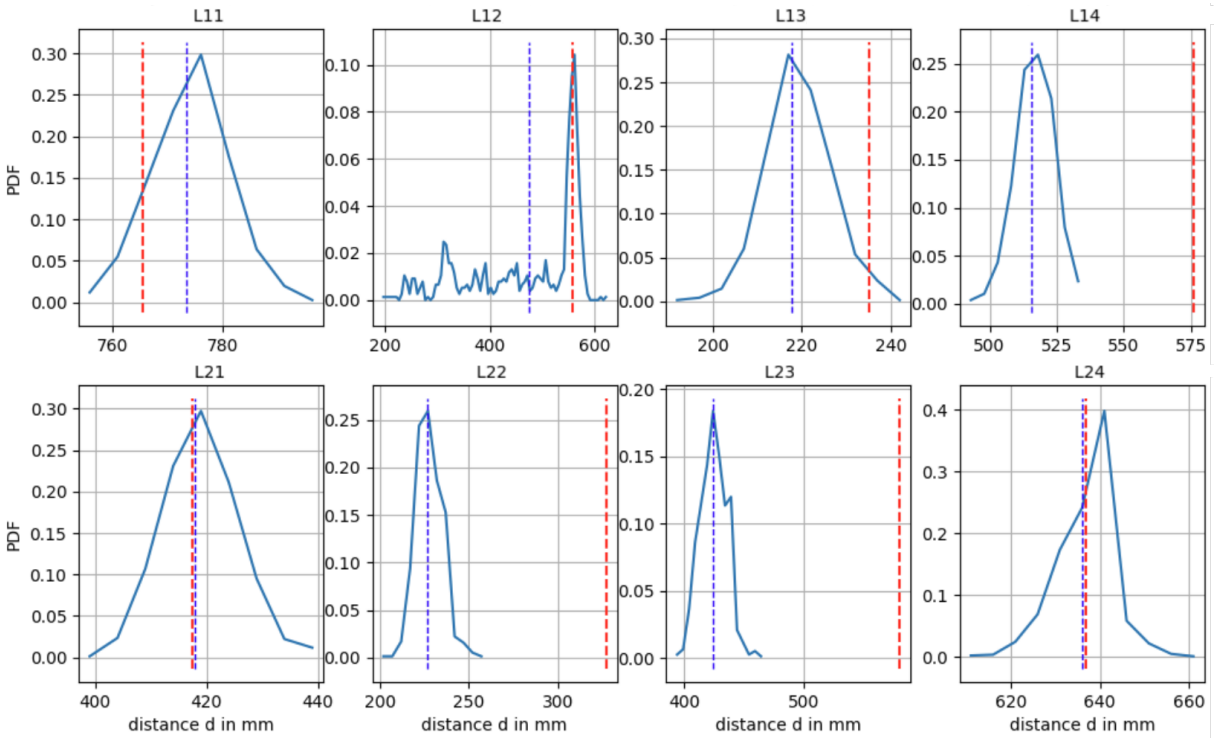


Figure 7: Distance probability distribution of all eight distances. The blue dashed line marks the distribution mean and the red dashed line marks the true distance value.

These distributions were created by moving the gantry to a certain position and measuring the distances several hundred times to gain some statistical confidence. The blue dashed line indicates the mean value of the distribution, while the red dashed line indicates the true distance. Each distance shows some variance during the recorded period, which was expected. However, there is also a significant bias between the true and measured distances.

Distance L14 in Fig. 7 returns an average distance of roughly 515 mm, while the true distance is about 575 mm (a bias of 60 mm). The distance L12 spans all the way from 200 mm up to 600 mm, most likely caused by interferences or maybe a blocked Line-of-Sight component. Blocked Line-of-Sight refers to the case where an object is located between two UWB modules. The modules cannot ”see” each other and the electromagnetic wave has to take an indirect path to the other module by reflecting, falsifying the distance measurement. Note that the results plotted in Fig. 7 change significantly if the MCC is moved slightly left or right by a few millimeters due to the short wavelength.

# 4 Multilateration

Multilateration refers to the calculation of spatial coordinates of unknown points from their distances to known points [3]. In the case of Matrix Charging, the known points are the anchors inside the MCP and the unknown points are the tags inside the MCC (see Fig. 5). Several different algorithms were implemented and tested, but the algorithm with the best result was the Numerical Maximum Likelihood Estimator (NMLE). Before explaining the NMLE algorithm, the following section gives some context on why a more elaborate algorithm is required for robust and accurate position estimation.

## 4.1 The naive approach for solving the multilateration problem (LERPINV)

In 2D, the multilateration problem can be solved by intersecting circles (see Fig. 8 (*left*)). To keep the figures readable, the second tag inside the MCC will be elided i.e., only four distances are displayed. The radius of the circle is the same as the distance measured between the anchor-tag pair. The position of the tag can be found by calculating the point where all four circles intersect.

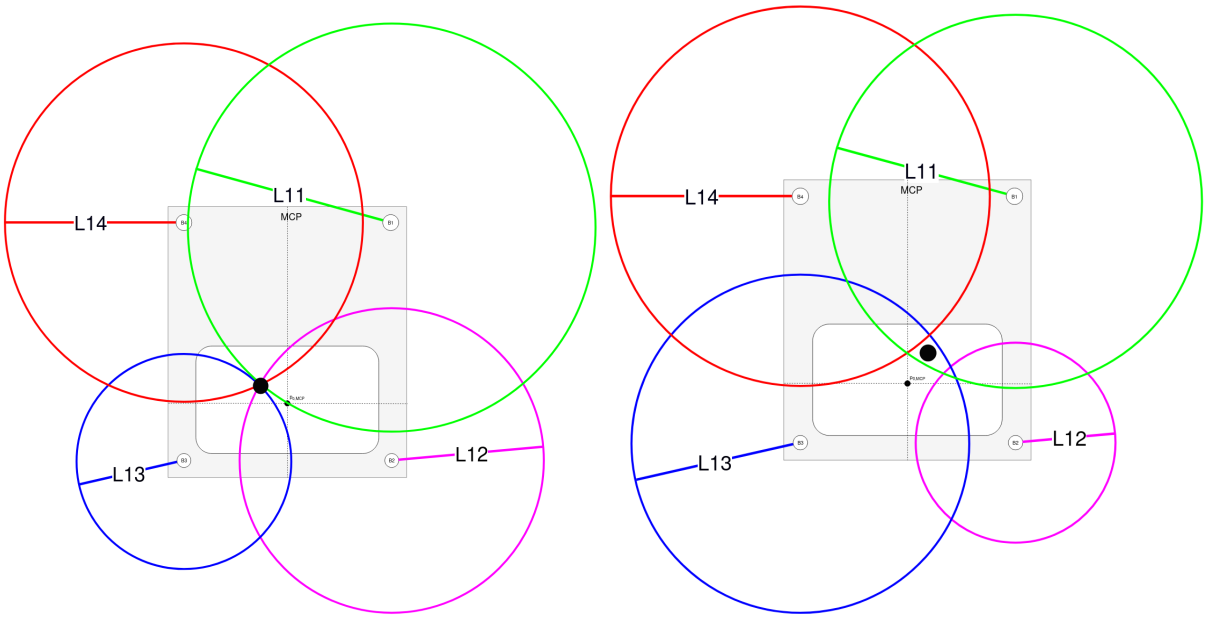


Figure 8: Estimating a position by intersecting circles using ideal distances (*left*) and real-world distances (*right*).

The benefit of this approach is that it allows for a simple algebraic solution by reducing the problem to linear equations and solving a matrix system using the pseudo-inverse [4]. The disadvantage is its high sensitivity to noise in the measurements (see Fig. 8 (*right*)), leading to poor overall performance if the distances are not ideal. As discussed in the previous section, the distances include some inaccuracies (see Fig. 7). This algorithm was labeled LERPINV and its inferiority to the NMLE algorithm will be displayed shortly in section 5.

## 4.2 Numerical Maximum Likelihood Estimator (NMLE)

The Numerical Maximum Likelihood Estimator (NMLE) frames the multilateration problem in a probabilistic manner. While the approach of section 4.1 modeled the distance as a deterministic circle, the NMLE models them as a random variable and creates a probability distribution for each distance. Figure 9 shows an exemplary probability distribution/likelihood.
These likelihoods are constructed by rotating a Gaussian distribution along a circle in Fig. 8, extending it to the third dimension. Note that any distribution could be used, but as depicted in Figure 7, a Gaussian

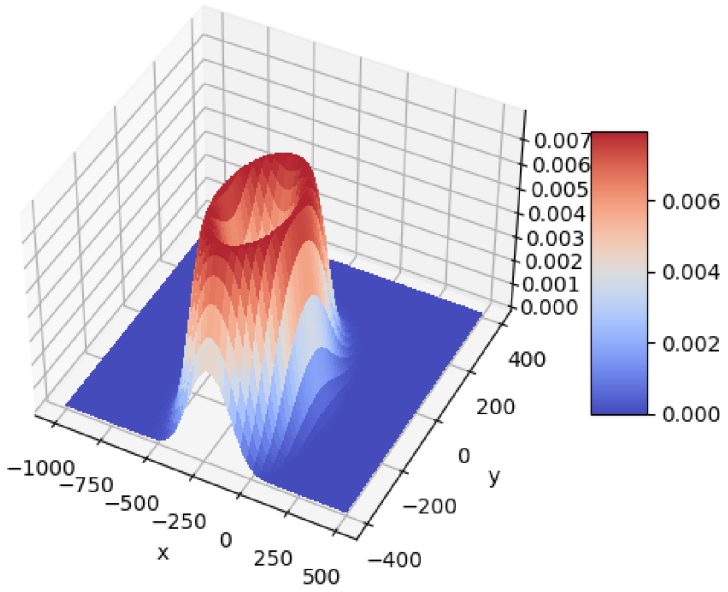


Figure 9: Exemplary likelihood function for one of the eight distances over x- and y-coordinates.

distribution fits the data best. By describing the distances this way, the inherent randomness of the measurements is taken into account and the real world is modeled better. Constructing these likelihoods from Fig. 9 for all eight distances and multiplying them together results in the joint likelihood shown in Fig. 10.

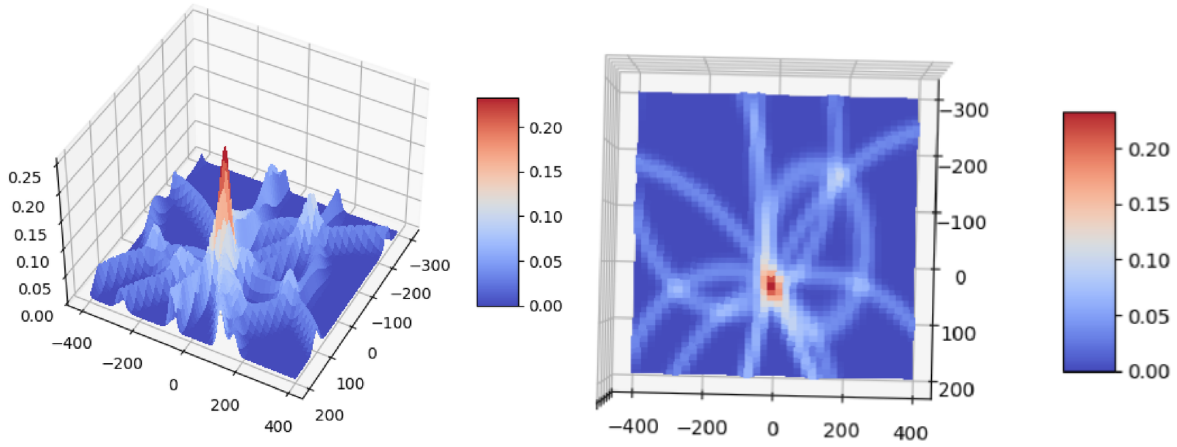


Figure 10: Joint likelihood which, when maximized, returns the tag position. Both figures show the exact same data but viewed from different angles. Note that the distribution was slightly modified for demonstrativeness purposes.

The location where most of them overlap creates the highest peak and the tag position can be estimated by finding the argument that maximizes this joint likelihood. Expressed as a formula, the position can be estimated with

$$\hat{\boldsymbol{p}}_{\mathrm{ML}} = \underset{\boldsymbol{p}}{\operatorname{argmax}} \prod_{i=0}^{8} P(\tilde{d}_i|\boldsymbol{p}) \tag{2}$$

, where $\hat{\boldsymbol{p}}_{\mathrm{ML}}$ is the estimated position and $P(\tilde{d}_i|\boldsymbol{p})$ the likelihood of a single distance $\tilde{d}_i$. The vector $\boldsymbol{p}$ is the argument to be maximized and holds the coordinates $\boldsymbol{p} = [x, y]$

Besides being a more robust framework, NMLE also has the advantage that the estimated position is a random variable itself and all of this can be put into a Bayesian framework. This means that prior knowledge can be incorporated into the algorithm and the posterior probability can be maximized.

$$\hat{\boldsymbol{p}}_{\text{Bayes}} = \underset{\boldsymbol{p}}{\operatorname{argmax}} \overbrace{\underbrace{\left(\prod_{i=0}^{8} P(\tilde{d}_i|\boldsymbol{p})\right)}_{\text{joint likelihood}} \cdot \underbrace{P(\boldsymbol{p})}_{\text{prior}}}^{\text{posterior}} \tag{3}$$

Prior information is obtained through the continuity of the problem. If the car was at position $\boldsymbol{p} = [x_0, y_0]$ a few milliseconds ago, it is probably near that position in the next iteration. With this information, a prior probability distribution can be created that has its peak at $[x_0, y_0]$, as can be seen in Fig. 11.

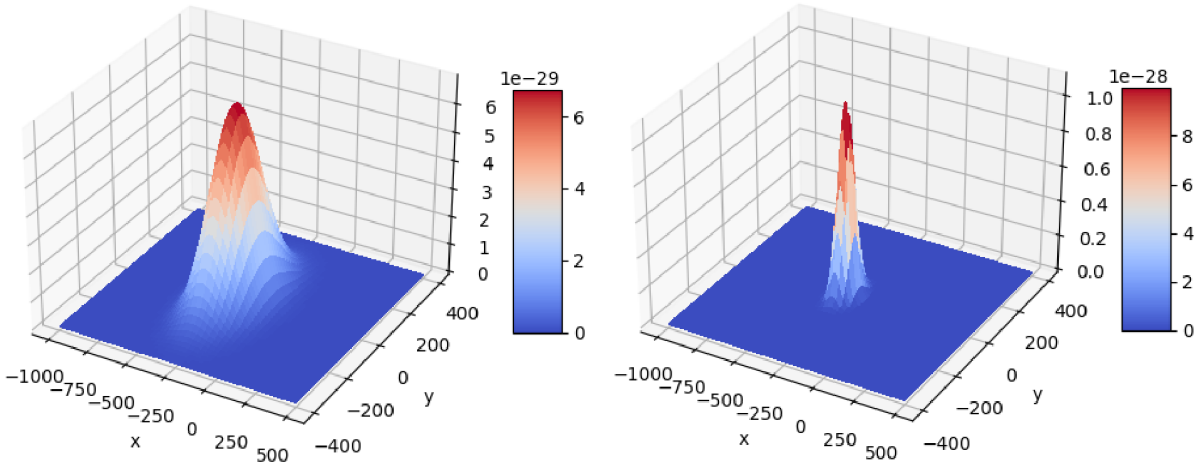


Figure 11: Two constructed probability distributions modeling prior knowledge with a higher (*left*) and lower variance (*right*).

The variance of the prior distribution is adaptively changed by using the information from the previous iterations. This adaptive process will not be explained in this paper, but the general idea is to calculate the average movement of the last few estimated positions and if it is close to zero, the variance of the prior should decrease.

The variance of the measurements causes the estimated position to jiggle, even if the EV is standing still, but by incorporating prior knowledge, the estimated position converges towards a single point. This greatly enhances the user experience with the positioning system, as the displayed position should not jiggle if the driver does not move the car. Additionally, the prior distribution is shaped so that it trusts changes along the y-axis less since movement purely along the y-axis would mean that the car moves sideways, which they are usually not capable of.

Further problem-specific knowledge can be incorporated by including the roughly known height difference between MCC and MCP. This reduces the problem from 3D to 2D, shrinking the problem domain and leading to a more robust algorithm.

Another big improvement is a clever transformation that allows the use of all eight distances at once to estimate the position. The original setup (Fig. 5) consists of two tags inside the MCC, meaning that the multilateration problem has to be solved two separate times, once for T1 and once for T2. This reduces the accuracy of the estimated position as only four distances can be used at once instead of all eight. But by shifting both tags to the MCC origin using a simple coordinate transformation, a virtual tag is created that can use all eight distances at once. From the perspective of the MCC, this appears like there are two virtual MCPs, as displayed in Fig. 12.

This transformation has the downside that it implicitly assumes that the MCC and MCP coordinate systems are aligned, meaning the x-axis of the MCC also points in the same direction as the MCP x-axis and there is no angular displacement between them. While this might seem restrictive, it does not cause any problems in the currently targeted use cases, where no big angular shifts between the different coordinate systems are possible.

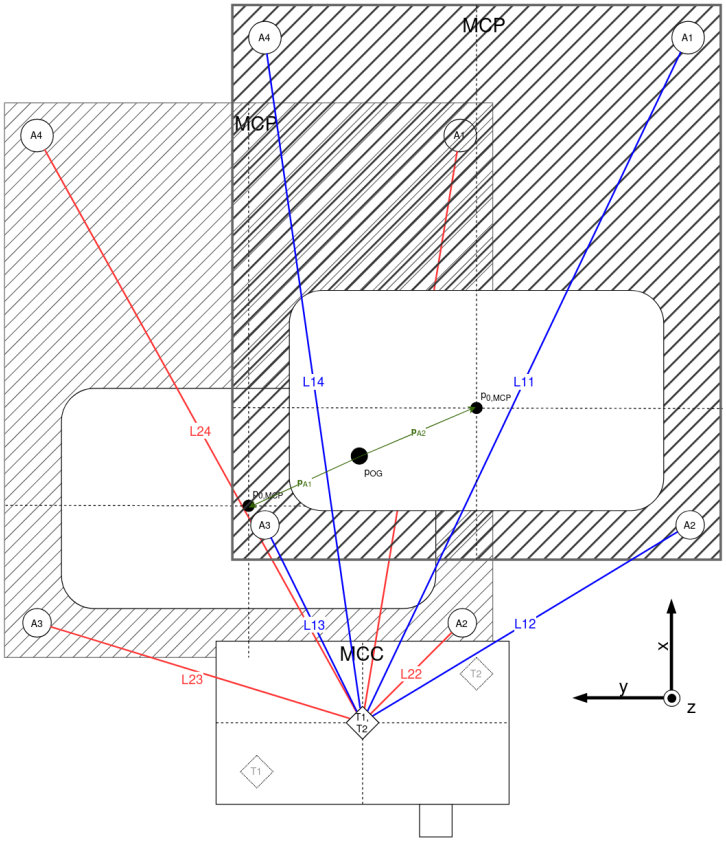


Figure 12: Transformed setup where both tags were shifted into the MCC origin, allowing all eight distances to be used at once to estimate the position.

# 5 Results

In Fig. 13, the gantry was moved to a certain position $[x_0, y_0]$ and hundreds of distance measurements were recorded. These distances were then fed into three different algorithms and the estimated positions were plotted in the left subfigure. The right subfigure displays the Empirical Cumulative Density Function (ECDF), allowing statements such as: ”99% of all estimated positions using NMLE no prior are within an error of 50 mm.”

The LERPINV algorithm (introduced in section 4.1) performs the worst with respect to the Root Mean Squared Error (RMSE) and also shows some outliers in the top left corner. These outliers are caused by highly erroneous distance measurements, which LERPINV is susceptible to.

The RMSE of the NMLE algorithm without using prior knowledge (NMLE no prior) is much better, having only an average estimation error of 33 mm and showing no outliers. The NMLE algorithm outperforms LERPINV because its framework is more robust to inaccurate measurements. The estimation position still deviates from the true position because of the bias within the distance measurements, which no algorithm can fix. Although the NMLE no prior version shows less variance in the estimated position compared to LERPINV, both of them are no match for the NMLE with prior version.

The NMLE with prior uses the Bayesian approach to incorporate prior knowledge gained from the previously estimated position, as explained in section 4.2. Note that the RMSE is basically the same between the two NMLE variants, meaning the average estimation error is still the same. The Bayesian framework, first and foremost, reduced the variance.

While Fig. 13 is easy to understand, it is not a proper evaluation of the NMLE algorithm since moving the MCC slightly to the left or right will show an entirely different result. Therefore, a heat map (Fig. 14) was created, which depicts the RMSE over a reasonable domain.
The blue areas in Fig. 14 mark positions with low estimation error and the red areas mark positions with a high estimation error. While it would be ideal for the entire domain to be deep blue, interferences and nearby metallic objects do not allow for that.

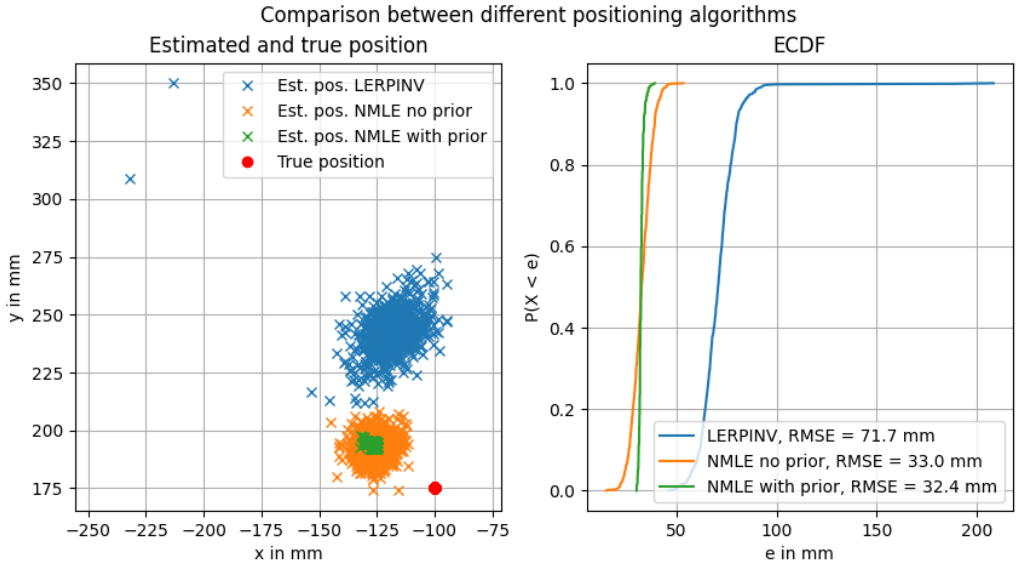


Figure 13: Comparison between LERPINV and NMLE using static data.

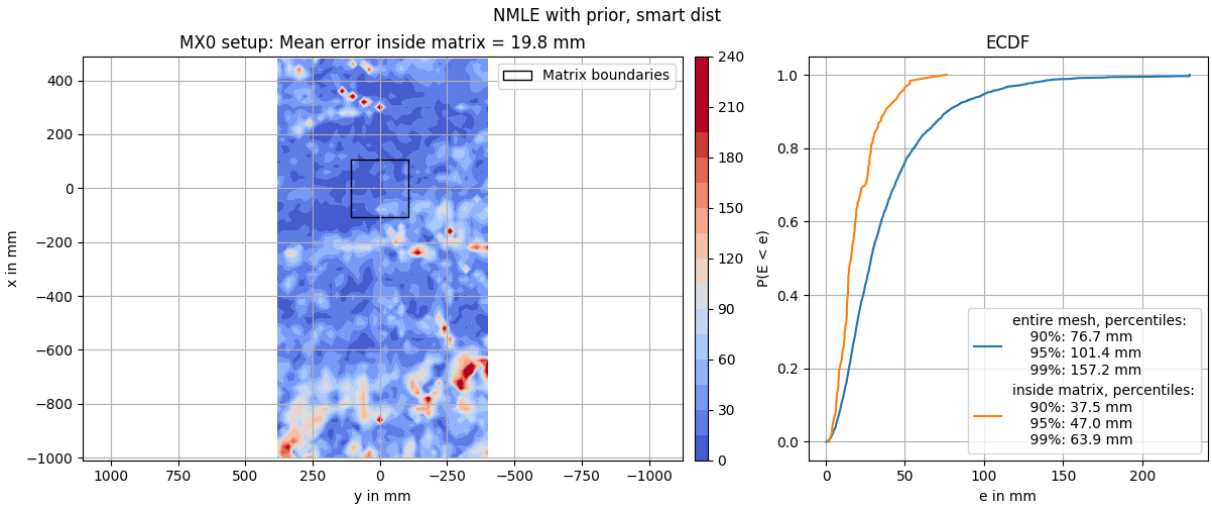


Figure 14: Heatmap and ECDF of the estimation error.

For the final benchmark, the domain was divided into two parts: Inside and outside the Matrix boundaries. The required accuracy is higher inside the Matrix since a faulty estimated position could cause the Core Unit to land outside the Matrix, requiring the user to reposition the car and leading to frustration. NMLE achieves an average estimation error of 19.8 mm inside the Matrix and ensures the error is below 47 mm in 95% of all cases, showcasing a solid algorithm.

# 6 Outlook and Conclusion

The results of the developed positioning system represent a highly promising interim result. During the course of the eTaxi Austria project, more than 40 taxi drivers tested the system over 12 months during their standard taxi operation and the feedback received was very positive. However, as can be seen in Figure 14, there are still selected areas where position estimation does not perform as expected and needs to be further improved.

In addition, we developed a very detailed understanding of which areas of the system need to be further investigated to enhance performance. On top of that, we started to collaborate closely with chip manufacturer NXP to get the most out of their UWB technology.

As a next step, we focus on the following areas for improvement: Using the newest generation of NXP UWB chips, optimization through antenna selection and placement within the Matrix Charging® system, and improvement of the positioning algorithm by sensor fusion with an additional accelerometer or other sensors. By the time of the writing of this paper, the improvement potential of those optimization levels can already be confirmed, which gives us great confidence in our further activities to commercialize UWB technology for our targeted use-case.

## References


[1] "Trimension NCJ29D5 for automotive applications," https://www.nxp.com/products/NCJ29D5, accessed 2025-03-11.

[2] "The ins and outs of uwb technology," https://www.nxp.com/design/design-center/training/TIP-THE-INS-AND-OUTS-OF-UWB-TECHNOLOGY, 2020, accessed 2025-03-11.

[3] A. Norrdine, "An algebraic solution to the multilateration problem," in *2012 International Conference on Indoor Positioning and Indoor Navigation, 13-15th November 2012*, 04 2015.

[4] "Indoor positioning using time of flight with respect to wifi access points," https://web.archive.org/web/20220925104244/https://people.csail.mit.edu/bkph/ftmrtt_location, accessed 2025-03-11.


## Presenter Biography

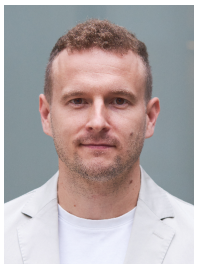

Gregor Eckhard is CTO at Easelink and has been working in the automotive industry for over 12 years in a variety of roles. Gregor holds a Master's degree in Automotive Engineering and Business Economics followed by an MBA in Business Administration from Austria, Argentina and Spain. As a management consultant at McKinsey and Company, he implemented technology and strategy projects with automotive manufacturers in Europe, the USA and Asia. He then took over the management of the automotive start-up MultiMaterial-Welding AG, leading it to a successful investment deal with the Bossard Group and SKion. Gregor is still a member of the supervisory board of MultiMaterial- Welding AG. At Easelink, Gregor leads the Easelink team together with Hermann Stockinger, Founder and Managing Director, to set the charging standard of tomorrow's electric vehicles.